Dynamic Political Agency with Adverse Selection and Moral Hazard*

Marcus Berliant[†] John Duggan[‡]

June 7, 2008

Abstract

We define a general model of electoral accountability in which an electorate chooses between an incumbent and a challenger in an infinite sequence of elections. We assume that politician types are private information and that the actions of an officeholder are unobserved by the electorate, so that adverse selection and moral hazard are both present. We establish existence of a perfect Bayesian equilibrium in pure strategies such that voters and politicians condition on the past only through the voters' beliefs regarding the type of the current officeholder, i.e., the equilibrium is "belief-stationary."

JEL numbers: D72, C73 Keywords: electoral accountability, principal-agent, moral hazard, adverse selection, democratic politics, political cycles, career concerns

1 Introduction

A distinguishing feature of democratic politics is the principal-agent relationship between the electorate and an elected politician. This relationship is typified by multiple,

*The authors retain responsibility for any errors.

†Department of Economics, Washington University, Campus Box 1208, 1 Brookings Drive, St. Louis, MO 63130-4899 USA. Phone: (314) 935-8486. Fax: (314) 935-4156. E-mail: berliant@artsci.wustl.edu.

‡W. Allen Wallis Institute of Political Economy, Department of Political Science, and Department of Economics, University of Rochester, Rochester, NY 14627 USA. Phone: (585) 273-4999. Fax: (585) 271-3900. E-mail: dugg@UR.rochester.edu.

heterogeneous principals (voters), and by ongoing interaction over time. The interaction between electorate and officeholder is complicated by the presence of moral hazard and adverse selection. Indeed, adverse selection is present in that voters cannot directly observe characteristics of political candidates such as competence or policy preference. It is evident that voters cannot directly observe decisions made by politicians in smoke-filled rooms, but even explicitly codified policies and legislation are observable to most voters only at prohibitive cost; rather, voters observe the repercussion of political decisions through policy outcomes, insofar as they affect their quality of life, that serve as a noisy signal of the choices of their representatives. The final feature of democratic politics we mention is that, owing to the size of the electorate and conflicts of interest among voters, it is implausible that voters could design and commit to a socially optimal incentive contract with a politician. These characteristics are not unique to democratic polities, but are also descriptive of corporate governance, where shareholders are engaged in a similar relationship with their board of directors. Even in autocratic regimes, a dictator must typically maintain the support of a political elite (wealthy interests or military leaders), who play a role similar to the electorate with their delegate, now a dictator instead of an elected officeholder.

In this paper, we conduct a theoretical analysis of democratic politics in light of this principal-agent relationship. We set forth a general model of electoral accountability, in which voters participate in an infinite sequence of elections between incumbents and challengers whose types are unobserved by voters; in each period, the elected officeholder chooses an unobserved action, after which a publicly observed outcome is realized; voters update their beliefs regarding the officeholder's type, and we move to the next period when another election is held. Because of commitment problems, which we view as inherent in democratic politics, we do not take a second-best approach to analyzing the model, where we would characterize a socially optimal contract (subject to incentive compatibility constraints) specifying the officeholder's compensation and re-election probabilities as a function of observed outcomes. Rather, we analyze the model as game between voters and politicians in which voters rationally anticipate the strategic choices of politicians and factor in the strategies of politicians when updating beliefs via Bayes rule; an officeholder, on the other hand, anticipates the reaction of the electorate to realized policy outcomes and chooses her action optimally given those expectations. We stipulate that this game may admit a plethora of complex, history-dependent, perfect Bayesian equilibria, but we view as implausible the idea that a large electorate could coordinate on such equilibria. Instead, we focus our analysis on the simplest form of equilibrium in this model, which is stationary in the sense that the decisions of voters and politicians depend only on common beliefs regarding the current incumbent's type. We refer to such equilibria as "belief-stationary."

Examples: For instance, the application to existence of pure strategy equilibrium when there are many issues, where equilibrium strategies vary across time.

non-existence example

A fundamental issue for the analysis of the model, and the principle focus of this paper, is the existence of belief-stationary equilibria. The issue is similar to that of existence of stationary Markov perfect equilibria in stochastic games, where the voters' beliefs about the officeholder's type correspond to a state in a stochastic game, but our model poses a difficult challenge not present in that literature. There, the transition probability to next period's state depends on the current state and the actions of players, which are endogenous, but the transition probability *function* is fixed. That is, the distribution of next period's state depends on the current state and actions in a way that is exogenously given in the model. In our framework, however, the transition from the voters' prior in the current period to a posterior (i.e., next period's prior) is governed by Bayes rule, which requires knowledge of the officeholder's strategy. That is, the distribution of next period's prior depends on today's prior and the officeholder's action in a way that is endogenous to the model. As in the literature on stochastic games, our approach to existence involves the introduction of an amount (perhaps arbitrarily small) of noise in the model, ensuring a minimal amount of smoothness in voters' and politicians' continuation values and preventing them from becoming arbitrarily "jagged." This noise takes the form of preference shocks, drawn independently across time, to the actors in the model and an exogenous signal about the incumbent officeholder's type that is independent of her actions and observed by all voters. Though the challenges we confront are slightly different (owing to the presence of incomplete information), this approach is adapted from Duggan and Kalandrakis (2007), who establish existence and regularity properties in a dynamic model of legislative bargaining with an endogenous policy status quo.

Marcus, you had an extra clause here. Too much information??

We provide a general existence result for belief-stationary equilibria. We impose only compactness conditions on the sets of feasible officeholder actions and policy outcomes, we impose only continuity properties on state utilities of voters and politicians, and we assume a general voting rule that simply specifies a collection of decisive coalitions of voters. We add a finite set of exogenous political states that evolve according to a controlled Markov process, allowing the officeholder's choice of action to influence next period's state. The state enters the stage utilities of the actors and can capture underlying economic variables, and it determines the collection of decisive coalitions in any period, thereby capturing demographics that can shift over time. The voters' prior beliefs about the challenger's type also depend on the current state, so that the state can incorporate variation across time in the pool from which challengers are drawn, whether due to economics or partisan forces. As well, political states can serve as memory that allows us to keep track of the passage of time, so that we can impose any finite term limit on officeholders, or we can divide the term of office into any finite number of sub-periods and allow voters to update their beliefs in each sub-period leading up to an eventual election.

We can let p depend on x too.

uhc

The timing of the model is as follows. In each period, we begin with an incumbent politician, a challenger, voter beliefs about the candidates, and a political state. A preference shock is drawn for each voter type, summarizing voters' perceptions of observable candidate qualities that are independent of policy outcomes. These qualities may derive from party preferences of the voters (partisan voters may give extra weight to a candidate by virtue of her party affiliation), preferences for personality, opinions regarding the candidate's appearance wearing hunting gear, or in a tank, etc. In other applications, the preference shock to the principal may simply capture costs of replacement. Similarly, a preference shock is drawn for each politician type, perturbing the politicians' preferences over actions taken while in office. These shocks are assumed to either be unobserved by the voter (or, equivalently, realized after the election). Voters cast ballots for the incumbent or challenger, each voter calculating the expected discounted utility from electing each candidate and voting for the candidate offering the higher payoff. The winner of the election, now the officeholder for the current period, chooses an unobserved action. Prior to the realization of the outcome of the officeholder's action, voters receive an exogenous signal about the officeholder's type. This could take the form of a news story about the politician's conduct in a prior office, or her performance in college, or aspects of her personal life that are outside the model. Voters update beliefs on the basis of that signal. Then the outcome of the officeholder's action is realized, and voters further revise their beliefs, which become the prior in the following period. A political state and challenger are drawn for the next period, and the above protocol is repeated. We prove existence of a perfect Bayesian equilibrium that is in pure strategies and is stationary with respect to the voters' beliefs about the incumbent's type.

The addition of noise, in the form of preference shocks and exogenous signals, to the electoral accountability model is a departure from earlier work. As discussed above, it provides us with some control over the voters' and politicians' continuation values, but it also implies that actors in the model almost always have unique best responses, giving us an existence result that avoids mixed strategies and increasing the scope for numerical computation of equilibria in the model. The use of noise to overcome existence and computational problems is not uncommon in applied literatures such as industrial organization (Aguirregabiria and Mira (2007) and Doraszelski and Satterthwaite (2007)) and intergenerational altruism (Bernheim and Ray (1989) and Nowak (2006)). In our setting, we assume that noise is continuously distributed, but the distributions of preference shocks and exogenous signals may be arbitrarily close to degenerate. We view these shocks as a technical device ensuring at least a minimal level of tractability, yet one that admits a compelling interpretation in the framework of the model; in any case, because preference shocks and exogenous information can be taken to be arbitrarily small, we feel that they do not compromise the applicability of the model.

The remainder of the paper is organized as follows. Section 2 contains a relatively thorough review of the literature on electoral accountability, as well as pointing out connections to related literatures on political budget cycles and career concerns. Section 3 sets forth the framework of the paper, defining the model and the concept of belief-stationary equilibrium. Section 4 contains the statement and proof of the existence theorem. Section 5 concludes.

2 Literature Review

This paper contributes to the literature on principal-agent theory, broadly speaking, and in particular to the literature on electoral accountability in political economy. With the exception of Barro (1973), models of the latter sort assume some form of incomplete information: either the motivations of politicians are known but their actions in office are unobserved (moral hazard), or their actions in office are observed but their motivations are not (adverse selection), or both.[1] Models combining moral hazard and adverse selection present the especially difficult challenge that the beliefs of voters evolve in a non-trivial way, and voters and politicians can conceivably condition their choices on beliefs in complex ways. This conditioning on beliefs poses difficult

[1] Ferejohn (1986) and Austen-Smith and Banks (1989) consider electoral accountability models of pure moral hazard. Our contribution is closer to models with an element of adverse selection.

challenges for the existence of belief-stationary equilibria, and existence of equilibrium in these models has only been obtained by exploiting special structure or by relaxing the restriction of belief-stationarity. The model with pure adverse selection is relatively tractable. Existence of stationary equilibrium still relies on the solution of a difficult fixed point problem, but in stationary versions of that model (i.e., there are no exogenous political states), one can establish the existence of equilibria in which the voters revise their beliefs regarding the officeholder only after her first policy decision, after which updating leaves the voters' beliefs fixed.

In relation to models with both adverse selection and moral hazard, ours is closest to work by Banks and Sundaram (1993, 1998). In both papers, these authors consider a one-dimensional model in which a representative voter has preferences over the actions of officeholders; in contrast, higher actions are costly to politicians, with types ordered according to their cost. A politician's payoff when out of office is zero, independent of policy outcomes. Thus, an action is best interpreted as an effort level, and the voter has a well-defined ranking of politician types. Banks and Sundaram (1993) consider a version of the model with no term limits and prove existence of an equilibrium in which voters and politicians use trigger strategies characterized by a cutoff, say $\overline{x}$, in the space of policy outcomes.[2] In this equilibrium, if the outcome realized following an officeholder's action ever drops below the level $\overline{x}$, the voter's strategy specifies that she remove that officeholder in any future election (so the officeholder is removed immediately), and the officeholder's strategy is that she shirk (choosing the lowest possible action) thereafter. The authors identify a cutoff (actually, a range of cutoffs) that forms a "simple" equilibrium. These equilibria are not belief-stationary, however, and in fact in equilibrium it may be that the voter replaces an incumbent even though she believes the officeholder to be the best possible type with probability arbitrarily close to one. Banks and Sundaram (1998) prove existence of a belief-stationary equilibrium in a version of the model with a two-period term limit that is otherwise identical to the earlier model. Their proof exploits the monotonic structure of the model and the fact that, with a two-period term limit, conditioning on voter beliefs is trivial: the voter's prior over an officeholder's type in her first term is fixed, but an officeholder always chooses her ideal action in her second term, regardless of the voter's beliefs. Thus, the voter updates her beliefs about the incumbent based on the outcome realized in the first term, but no decisions are conditioned on those updated beliefs.

[2] **?** consider a two-period version of this model, capturing some spirit of the model of Banks and Sundaram (1998).

Our model departs from that of Banks and Sundaram in that we drop the monotonic structure they assume. In our model, the sets of actions and outcomes may be multidimensional, and a type may reflect a cost of effort, but it can encode arbitrary information about the officeholder, including her policy preferences, her career prior to running for office (e.g., whether she held office as a Senator or state governor), or the geographic region where she lives. We do not restrict stage utilities beyond continuity, and therefore we may interpret actions as legislation or policy directives and outcomes as belonging to an ideological space (e.g., measuring the conservatism of the Supreme Court) or describing economic variables (e.g., employment or price levels in different sectors of the economy). We allow for a heterogeneous electorate with an arbitrary voting rule, relaxing the assumption of a representative voter. Finally, we view politicians on a par with voters, and we therefore assume that stage utility accrues to politicians even while out of office. With the addition of a small amount of noise to the model, we prove existence of equilibria in pure strategies that are stationary with respect to beliefs with or without a finite term limit.

Among work on models of pure adverse selection, the closest to our paper is Banks and Duggan (2008a,b), who consider a general multidimensional adverse selection model assuming an abstract type space, allowing for a heterogeneous electorate, and allowing the utilities of out-of-office politicians to depend on policy outcomes. The equilibria characterized by Banks and Duggan are described by a "win set" containing policies that are sufficient to ensure reelection of the officeholder; any politician whose ideal point lies in the win set simply chooses that policy in every period she holds office; in general, some types of politician whose ideal points are outside the win set will "compromise" by choosing the most advantageous policy in the win set, thereby winning re-election; and some types with ideal points outside the win set will "shirk" by simply choosing their ideal points, forgoing reelection. In comparison, our model drops the convexity assumptions maintained by Banks and Duggan and adds an abstract political state along with the element of moral hazard to the model. The characterization of equilibria in our general model is obviously difficult, but more...

A number of papers analyze more specialized adverse selection models. Reed (1994) considers a two-period model of effort provision. Duggan (2000) assumes an infinite horizon and a one-dimensional policy space. Bernhardt et al. (2004) add finite term limits to the model. Kang (2005) adds costly signaling in campaigns.[3] Bernhardt et al. (2005) add partisan challenger selection. Casamatta and De Paoli (2007) consider

[3] Le Borgne and Lockwood (2001) study a two-period model of costly signaling by challengers.

the possibility of inefficient public good provision in the model.[4] Meirowitz (2007) considers a model in which officeholders have private information about the feasible set of policies. The structure assumed in much of this work can be embedded in our framework. We can of course specialize the model to one dimension, which may be ideological (so that stage utilities are single-peaked) or an effort level (so voters' preferences are increasing and politicians are decreasing). The presence of the political state allows us to capture any finite term limit or, if desired, to truncate the model after two periods. It also allows us to capture the possibility that the challenger belongs to the opposition party, so that her type is drawn from a different distribution than the incumbent's type. We also allow the set of actions feasible for a politician to depend on her type, so our model incorporates private information about the set of policies feasible in a given period. We cannot incorporate costly signaling by challengers, as in Kang (2005), since our politicians do not take actions prior taking office. In the model of Casamatta and De Paoli (2007), a politician's type is the cost of a public good (high or low), which is fixed across time and known to politicians but unknown to the voter; thus, politician types are perfectly correlated, a situation we do not capture. Kalandrakis (2006) takes up a different class of adverse selection model, in which two long-lived parties compete repeatedly in a sequence of elections, the party ideal points subject to a random (correlated) shock after each period. In that model, the voter's beliefs about the challenger's (the out-party's) type evolve endogenously, in addition to the incumbent's (the party in power), something we do not capture.[5]

Marcus, I'm assuming we put the electoral outcome in the p function. This is very simple: $p^{\mathfrak{I}}$ v. $p^{\mathfrak{C}}$.

But we could capture it!

Fearon (1999)

Our paper is also related to the literature on political budget cycles with rational expectations and to the principal-agent literature on career concerns. Rogoff (1990) considers a finite-horizon model in which politicians take hidden actions and whose competence levels are private information.[6] Each election period is broken into two sub-periods in which an officeholder chooses a public good level and a level of investment. The officeholder knows her own competence level but the voter does not, creating the scope for manipulation of voter beliefs in the sub-period preceding an election. Our model is distinguished from Rogoff's in two important ways. First, in Rogoff's model, the voter observes the actions of the current officeholder before the election is held, and those actions are observed without noise. Since Rogoff focuses on separating equilibrium, updating of beliefs is simple: in equilibrium, the voter is able to infer

[4] Coate and Morris (1995) take up this issue in a two-period version of the model.

[5] Aragonès et al. (2007) consider a model of repeated elections with symmetric information and analyze history-dependent equilibria with some of the flavor of the electoral accountability models.

[6] See also **?**.

the officeholder's type exactly, and voter beliefs following out-of-equilibrium actions can be constructed to support the equilibrium. In contrast, we assume that the link between politician actions and observable outcomes is stochastic, so that there are no out-of-equilibrium realizations, and updating is pinned down by Bayes rule. Second, the officeholder's competence level (either high or low) is subject to a random shock in each sub-period, and the officeholder's competence in one period is revealed to the voter at the beginning of the next. Thus, the voter's beliefs evolve in a trivial way: regardless of the voter's posterior at the end of period t, her prior at the beginning of period $t + 1$ is pinned down by the assumptions of the model. Our model allows for the possibility of sub-periods between elections, and it therefore invites an application to political budget cycles in which officeholder types are never revealed to voters and beliefs follow a complex dynamic.

We need to assume no out of equilibrium outcomes!

Following the work of Holmström (1999) on career concerns, Ashworth (2005) and Martinez (2008a,b) have explored models of moral hazard with symmetric learning about the officeholder's type. The assumption of symmetric information is meant to capture the possibility that a worker, or officeholder, may not know with certainty her suitability for the job at hand. Clearly, this assumption abstracts from the possibility of private information, such as a worker's cost of effort. This abstraction is perhaps of greater concern in models of dynamic elections, however, for then we would like a politician's type to capture particulars of a politician's policy preferences. Our model does not permit learning on the part of the officeholder, as our focus is on adverse selection, but our techniques apply equally well to the career concerns model with adverse selection. We capture as a special case the situation in which a single voter (an employer) decides whether to retain the officeholder (a worker) in a given job or to hire a replacement. This application of our model, as currently configured, is limited to a single worker, but our techniques would extend to a model in which an employer were assigning workers to a finite number of jobs, providing an existence result for belief-stationary equilibria in a general model of career concerns with adverse selection.

3 The Model

Framework. We posit a countably infinite pool M of potential political agents and a set N of voters, or principals, a finite set Ω of exogenous states, and we consider a countable sequence of elections between incumbents and challengers in periods $t =$

Let voter utility depend on officeholder's type: $u_\tau(x, \tau', \omega)$. Let p depend on τ and voting outcome. Give politicians a

$1, 2, \ldots$. Both politicians and voters are characterized by types τ, which lie in the finite set T (note that this set is the same for both kinds of agents). We assume that the state $\omega \in \Omega$ in any period is observed by all players, and that the distribution of types in the electorate in state ω is known to all players. In the analysis of the model, the relevant aspect of the distribution of voter types is the corresponding collection of decisive coalitions of types, i.e., coalitions of types the support of which is sufficient to win an election in state ω. We choose a parsimonious model of the electorate by simply specifying a collection $\mathfrak{D}(\omega)$ of subsets $S \subseteq T$ that can determine the outcome of the election in state ω. Politician types are private information. In any period, a single politician $m \in M$ is active and, if her type is τ, chooses an action a from a non-empty, state-dependent set $A_\tau(\omega) \subseteq \Re^d$ of actions. An outcome x in a set $X \subseteq \Re^e$ is stochastically determined (as described below) and observed by all players. The stage utility function of a type τ officeholder is $u_\tau \colon X \times A_\tau \times \Omega \to \Re$, where $A_\tau = \bigcup_\omega A_\tau(\omega)$ is the set of conceivable actions for the politician. Thus, the officeholder's payoff in a period in which the state is ω, the officeholder takes action a, and the outcome is x is $u_\tau(x, a, \omega)$. The stage utility function of a type τ voter, or to an out-of-office politician, is $u_\tau \colon X \times \Omega \to \Re$. Payoffs in a state ω for a type τ player are discounted by $\delta_\tau(\omega) \in [0, 1)$.

Can contain empty set.

In order to ensure tractability of the model, we modify this basic setup by adding two elements of noise. First, we incorporate preference shocks (θ, ϵ) in the model, where $\theta = (\theta_\tau)_\tau \in (\Re^d)^T$ perturbs the stage utilities of officeholders and $\epsilon = (\epsilon_\tau)_\tau \in \Re^T$ perturbs the stage utilities of the voters. Second, we assume that voters receive a noisy signal, ψ, separate from the observed policy outcome, about the current officeholder's type that is unrelated to her action choice. We incorporate this noise into the model with the following timing.

1) At the beginning of any period, the type τ and preference shock θ_τ of each politician are given and privately known to the politicians.[7] A state ω and preference shocks ϵ are given and publicly observed by all players, and prior beliefs b of the voters regarding the officeholder's type are common knowledge.[8]

Given our setup, the incumbent might not want to run for election...

[7] To be precise, we assume that τ is private information of an individual politician m, while θ_τ is an aggregate shock that is the private information of all type τ politicians. The analysis goes through, at a higher notational cost, if we assume shocks θ_m that are politician-specific.

[8] Because the shock ϵ_τ is an aggregate shock that affects the payoffs of all type τ voters, we maintain the interpretation that these shocks are publicly observed. Our results are unaffected if ϵ_τ is realized prior to the officeholder's action and is the private information of type τ voters.

2) An incumbent politician, i, is matched with a challenger, c, randomly drawn from the pool M of politicians, and the two candidates compete in an election. The voters' prior beliefs about the type of the challenger are given by $\bar{b}(\omega) \in B$.

We can let the distribution of the challenger's type depend on the current officeholder's type...

3) Voters simultaneously cast ballots in $C = \{\mathfrak{I}, \mathfrak{C}\}$, where $\mathfrak{I}$ denotes a vote for the incumbent and $\mathfrak{C}$ denotes a vote for the challenger. Assuming that voters of the same type vote the same way, we can specify the winner of the election as the incumbent if the set of types voting for her is decisive, i.e., $\{n \in N \mid \ n \text{ votes } \mathfrak{I}\ \} \in \mathfrak{D}(\omega)$, and as the challenger otherwise. To keep the analysis relatively simple, we assume that the loser (e.g., Al Gore, John Kerry) is not re-drawn as a challenger in the future. Let $k \in \{i, c\}$ denote the winner of the election and officeholder for the current period.

4) The officeholder, if her type is τ, takes an action $a \in A_\tau(\omega)$ that is unobserved by the voters.

5) Before a policy outcome is realized, we assume that voters observe an exogenous signal ψ about the officeholder's type.[9] This could, for example, take the form of a news story about the politician's conduct in a prior office, or a news story about the politician's personality that is independent of her actions while in office. We suppose that voters update beliefs to $\tilde{b}$, where $\tilde{b}$ is drawn from a density $g_\tau(\cdot|b, \omega)$.

6) A policy outcome x is then realized with density $f_\tau(\cdot|a, \omega)$, and voters update their beliefs from $\tilde{b}$ to b' using Bayes' rule.

7) Each voter of type τ' receives

$$u_{\tau'}(x, \omega) + \chi_i(k)\epsilon_{\tau'},$$

and each politician, m, if her type is τ', then receives stage utility

$$\chi_k(m)(u_{\tau'}(x, a, \omega) + \theta_{\tau'} \cdot a) + (1 - \chi_k(m))(u_{\tau'}(x, \omega) + \chi_i(k)\epsilon_{\tau'}).$$

Here, χ_i is an indicator function that takes a value of one if office holder k was the incumbent from the previous period, zero otherwise, and χ_k is an indicator function that takes a value of one if politician m is the current officeholder, zero otherwise. Thus, stage utility accrues to out-of-office politicians in the same way that it does for voters. The preference shock θ_τ serves to perturb the politician's preferences over actions, while the shock ϵ_τ summarizes the type τ voters' perceptions of candidate qualities that are independent of policy outcomes.

independence of θ and ϵ...

[9]For consistency, we assume that politicians also observe this signal. This is without consequence for the analysis of the model.

8) Finally, a new state ω' is drawn from the distribution $p(\cdot|a,\omega)$, and shocks (θ',ϵ') are drawn from the joint density $h(\cdot|\omega)$. (We write $h(\theta|\omega)$ and $h(\epsilon|\omega)$ for the marginal densities on the politicians' and voters' shocks. We denote the corresponding distribution functions by $H_\theta(\cdot \mid \omega)$ and $H_\epsilon(\cdot \mid \omega)$.) The model transitions to the next period with state ω', beliefs b', shocks (θ',ϵ'), and play resumes as described above.

Assumptions. We impose no structure on the set N of voters outside the collection $\mathfrak{D}(\omega)$ of decisive coalitions. Regarding the latter, we assume only that if one coalition of types S is decisive, and if we add types to that coalition, then the larger coalition is also decisive: for all ω and S, if $S \in \mathfrak{D}(\omega)$ and $S \subseteq S' \subseteq T$, then $S' \in \mathfrak{D}(\omega)$. This allows us to capture weighted majority rule, where the weight of a type indicates the representation of that type within the voting population and can depend on the current state. Formally, we can let $\omega = (\omega_\tau) \in \Re^T_+$, where ω_τ is the representation type τ, and we can then define

$$\mathfrak{D}(\omega) \quad = \quad \left\{ S \subseteq T \mid \sum_{\tau \in S} \omega_\tau > \frac{1}{2} \right\}.$$

We admit the possibility that $\mathfrak{D}(\omega)$ contains the empty set, in which case the incumbent is automatically elected, allowing us to capture any finite number of sub-periods between elections. Of course, our framework admits a single principle, or representative voter, as a special case. We can account, as well, for non-democratic systems in which the current dictator must maintain the support of a winning coalition $S \in \mathfrak{D}(\omega)$ of members of the political elite in order to stay in power. Here, the difference in office and out-of-office payoffs, $u_\tau(a,x,\omega) - u_\tau(x,\omega)$, can incorporate punishment of dictators who are removed from office.

The model allows the possibility that $M \subseteq N$, so that politicians comprise a subset of the electorate, but our equilibrium concept (defined below) implicitly assumes the probability that a voter is selected as challenger is equal to zero. It is nevertheless valid to imbed the set M of politicians within the set N of voters if, for example, N has a continuum of elements and we conceive of challenger selection as a random draw from the electorate. (In this case, we can let M be an arbitrary realization of a countably infinite number of draws from N according to a uniform distribution.) Alternatively, we can maintain that $M \cap N = \emptyset$, so that technically, the set of potential officeholders (who decide policy) is disjoint from the set of voters (who decide elections). Both assumptions capture the idea that the set of politicians is small relative to the electorate.

We assume the set $A_\tau(\omega)$ of feasible actions is compact for all τ and ω, admitting the case of a finite set of actions as a special case, and we assume X is compact. We assume that stage payoffs of politicians, $u_\tau(x, a, \omega)$, and voters, $u_\tau(x, \omega)$, are jointly continuous in their arguments, capturing the case of quadratic stage payoffs, i.e., $u_\tau(x, \omega) = u_\tau(x, a, \omega) = -||x - \hat{x}_\tau(\omega)||^2$, where $\hat{x}_\tau(\omega) \in \Re^e$ is the state-dependent ideal point of a type τ voter or politician. The linear form of the politicians' preference shock, $u_\tau(x, a, \omega) + \theta_\tau \cdot a$, is chosen for simplicity and because, for the special case of quadratic utilities, θ_τ has the convenient interpretation as a perturbation of the politicians' ideal point.[10] Note that when A_τ is finite, the linear shock can be viewed as simply the sum of a noise term $\theta_{\tau,a}$ to a politician's payoff from taking action a, as in $u_\tau(x, a, \omega) + \theta_{\tau,a}$, unifying our treatment of voters and politicians.[11]

In the above, we have black-boxed the signal ψ observed by voters and simply assumed noise on voter beliefs distributed according to the density $g_\tau(\cdot|b, \omega)$. Moreover, we will assume in the sequel that $g_\tau(\tilde{b}|b, \omega)$ is jointly continuous in $(\tilde{b}, b, \omega)$. To give a more detailed account, we put signals in the space $\Re^T$, and we assume that the distribution of ψ depends on the type of the officeholder but not on the officeholder's actions. Thus, we interpret ψ as media coverage that focuses on the politician's type, e.g., a story about the politician's college transcript or actions in earlier office, and not on actions or policy outcomes. Assume that ψ is distributed according to the density $\tilde{h}_\tau(\cdot|\omega)$ when the officeholder is type τ and the current state is ω. Then following the realization of ψ, beliefs are updated according to the mapping $\tilde{\beta}\colon B \times \Omega \times \Re^T \to B$ defined by

$$\tilde{\beta}_\tau(b, \omega, \psi) = \frac{b_\tau \tilde{h}_\tau(\psi|\omega)}{\sum_{\tau'} b_{\tau'} \tilde{h}_{\tau'}(\psi|\omega)}.$$

Here, $\tilde{b} = \tilde{\beta}(b, \omega, \psi)$ gives the voters' posterior on the officeholder's type when the priors are b, the current state is ω, and the signal is ψ. It is important for our purposes that for all b and all ω, the random variable $\tilde{\beta}_\tau(b, \omega, \cdot)$ has a continuous density $g_\tau(\cdot|b, \omega)$.

MARCUS: We would like to claim that for almost all $\{\tilde{h}_\tau\}$, the revised beliefs are indeed distributed according to a continuous density. This is related to our earlier lemma, though simpler in some respects. I haven't thought more about it...

[10] To see this, note that $u_\tau(x, a, \omega) + \theta_\tau \cdot a = -(x - (\hat{x}_\tau(\omega) + \frac{1}{2}\theta_\tau)) \cdot (x - (\hat{x}_\tau(\omega) + \frac{1}{2}\theta_\tau)) + \theta_\tau \hat{x}_\tau(\omega) + \frac{1}{4}\theta_\tau \cdot \theta_\tau$, which is equal to the quadratic utility from x with ideal point $\hat{x}_\tau(\omega) + \frac{1}{2}\theta_\tau$ plus a constant term.

[11] The linear form is merely a convenient method for perturbing stage utilities of politicians in the general case, allowing A_τ infinite. See Mas-Colell's (1985) Theorem I.3.1 for a more general condition that suffices for our purposes.

We assume that the density $f_\tau(x|a,\omega)$ is jointly continuous in (x,a,ω) for all τ. For each x, $\tilde{b}$, and ω, define the mapping $\phi(\cdot|x,\tilde{b},\omega)\colon \bigcup_\tau A_\tau \times B \to \Re^T$ by

$$\phi_\tau(a,b|x,\tilde{b},\omega) \;=\; f_\tau(x|a_\tau,\omega)g_\tau(\tilde{b}|b).$$

Then we assume that each $\phi_\tau(\cdot|x,\tilde{b},\omega)$ is Lipschitz continuous with uniform Lipschitz constant $L_\tau(x,\tilde{b},\omega)$, i.e., for all (a,b) and (a',b'), we have

$$||\phi_\omega(a,b|x,\tilde{b},\omega) - \phi_\omega(a',b'|x,\tilde{b},\omega)|| \;\leq\; L_\tau(x,\tilde{b},\omega)||(a,b)-(a',b')||.$$

We further assume that these Lipschitz bounds can be taken to be uniform over x, $\tilde{b}$, and ω, i.e., $L_1 \equiv \sup_{\tau,x,\tilde{b},\omega} L_\tau(x,\tilde{b},\omega) < \infty$. This condition holds, for example, if f_τ and g_τ are continuously differentiable. We assume that politician and voter payoffs are bounded in expectation: there exists L_2 such that for all $\tau \in T$, all $\omega \in \Omega$, all $a \in A_\tau$, and all $x \in X$, we have

$$\max\left\{\left|\int_\theta (u_\tau(x,a,\omega)+\theta_\tau\cdot a)h(\theta|\omega)d\theta\right|, \left|\int_\epsilon (u_\tau(x)+\epsilon_\tau)h(\epsilon|\omega)d\epsilon\right|\right\} \;\leq\; L_2.$$

This holds as long as θ and ϵ have finite first moments for each state ω. We assume that draws of (θ,ϵ) and $\tilde{b}$ are independent across time. Finally, we assume that each $p_\tau(\omega'|\cdot,\omega)$ is Lipschitz with bound $L_3(\omega,\omega')$ uniform in a, and we define $L_3 = \max_{\omega,\omega'} L_3(\omega,\omega')$, which is well-defined since Ω is finite. This is satisfied if $p_\tau(\omega'|a,\omega)$ is continuously differentiable in a. independence of θ and ϵ

Strategies. A voter history of length t is a sequence

$$\mathfrak{h}_v = ((m_1,\omega_1,\epsilon_1,\psi_1,x_1),\ldots,(m_{t-1},\omega_{t-1},\epsilon_{t-1},\psi_{t-1},x_{t-1}),(m_t,\omega_t,\epsilon_t,\psi_t,x_t)),$$

where m, ω, ϵ, ψ, and x indicate, respectively, the name of the politician holding office in a given period, the current period's state, the shock to voter preferences for that period, an exogenous signal about the type of the officeholder, and the policy outcome in the current period.[12] A history for a politician m is a sequence

$$\begin{aligned}\mathfrak{h}_m \;=\; &(\tau,(m_1,\theta_{1,\tau},\omega_1,\epsilon_1,a_1,\psi_1,x_1),\ldots,(m_{t-1},\theta_{t-1,\tau},\omega_{t-1},\epsilon_{t-1},a_{t-1},\psi_{t-1},x_{t-1}),\\ &(m_t,\theta_{t,\tau},\omega_t,\epsilon_t)),\end{aligned}$$

[12] To economize our setup, we do not define histories for individual voters, and since we assume there is no aggregate uncertainty regarding the distribution of voter types, we do not take a stand on whether a particular voter's type is private information. Rather, in the sequel, we incorporate the dependence of a voter's ballot on her type by assuming that voters' strategies are anonymous with respect to type, and we allow the voting strategy of type τ voters to differ from the strategy of type τ' voters.

where τ is the politician's type, θ_τ is the privately observed shock to type τ politicians' preferences over actions, and a represents the politician's actions each period (we can designate an action $\overline{a}$ as a default action and set $a_t = \overline{a}$ in any period t during which the politician did not hold office, i.e., $m_t \neq m$).[13] Given a politician history $\mathfrak{h}$, let $\tau(\mathfrak{h})$ denote the politician's type in that history. Let $\mathfrak{H}_v$ denote the set of histories for type τ voters, and let $\mathfrak{H}_m$ be the set of histories for politician m.

A strategy for a type τ voter is a mapping $\rho_\tau \colon \mathfrak{H}_v \to C$ that specifies the voter's ballot as a Borel measurable function of history. A strategy for politician m is a mapping $\alpha_m \colon \mathfrak{H}_m \to \bigcup_\tau A_\tau$ such that $\alpha_m(\mathfrak{h}) \in A_{\tau(\mathfrak{h})}$ for all histories $\mathfrak{h} \in \mathfrak{H}_m$, specifying a feasible action as a Borel measurable function of histories for m. Note that to simplify the analysis of the electorate, we impose the restriction that voters of like types cast like ballots. Let $\sigma = (\alpha, \rho)$ denote a strategy profile for the voter and politicians. Due to our assumptions, in particular about noise experienced by agents, there will be no mixing except on a probability zero set of preference shocks, so we can restrict attention to pure strategies. Since politicians are symmetric in this model, we henceforth impose the restriction that politicians condition actions on types in the same way. Formally, consider any two politicians $m, m' \in M$, and let $\mathfrak{h}_m$ be a history for m. Now let $\mathfrak{h}_{m'}$ be any history for m' such that $\tau(\mathfrak{h}_m) = \tau(\mathfrak{h}_{m'})$ and in which the two politicians are interchanged in the sequence of office holders, i.e., for all t, $m_t = m$ if and only if $m'_t = m'$. Then we simplify the analysis by requiring that $\alpha_m(\mathfrak{h}_m) = \alpha_{m'}(\mathfrak{h}_{m'})$.

We focus on a particularly simple class of perfect Bayesian equilibria. We say σ is *belief-stationary* if α and ρ are such that for all t, the choices of politicians and voters depend on previous periods only through the voters' beliefs regarding the current officeholder's type. Formally, we suppose that voters update regarding the officeholder's type based on policy outcomes and states via a belief mapping $\tilde{\beta} \colon X \times B \times \Omega \times \Omega \to B$ as follows. Suppose, given voter history $\mathfrak{h}_v$ of length t, that the current officeholder, m, has held office for the previous ℓ periods. Let $b_{t-\ell}(\mathfrak{h}) = \overline{b}(\omega_{t-\ell})$ be the voter's prior on m's type in state $\omega_{t-\ell}$ at the beginning of period $t - \ell$ (when m is the challenger), and let

$$b_{t'}(\mathfrak{h}) = \tilde{\beta}(x_{t'-1}, b_{t'-1}(\mathfrak{h}), \omega_{t'})$$

for $t' = t - \ell + 1, \ldots, t$ be the voter's updated beliefs on m's type. Thus, $b_t(\mathfrak{h})$ is the voters' prior in period t. Then we require that the voter's strategies are measurable with respect to the voter's type, the state, her preference shock, and the prior for the

[13] We must define histories and strategies for individual politicians, as a politician's payoff can depend on whether she (and not just another politician of the same type) holds office in a period.

current period regarding the incumbent's type: for all types τ and all voter histories $\mathfrak{h}_v, \mathfrak{h}'_v$ of length t with $\omega_t = \omega'_t$, $\epsilon_\tau = \epsilon_{\tau'}$, and $b_t(\mathfrak{h}_v) = b_t(\mathfrak{h}'_v)$, we have $\rho_\tau(\mathfrak{h}_v) = \rho_\tau(\mathfrak{h}'_v)$. Furthermore, we require that the politicians' strategies are measurable with respect to the politician's type, her preference shock, the state, and the voters' updated beliefs: for all m and all histories $\mathfrak{h}_m, \mathfrak{h}'_m$ of length t where m holds office in period t and such that $\tau(\mathfrak{h}_m) = \tau(\mathfrak{h}'_m)$, $\theta_{t,\tau} = \theta'_{t,\tau}$, $\omega_t = \omega'_t$, and $b_t(\mathfrak{h}_m) = b_t(\mathfrak{h}'_m)$, we have $\alpha_m(\mathfrak{h}_m) = \alpha_m(\mathfrak{h}'_m)$. In this case, we say that σ is *stationary with respect to* $\tilde{\beta}$.

Given a belief-stationary profile, we view the voters' strategies simply as mappings $\rho_\tau : B \times \Re \times \Omega \to C$, where $\rho_\tau(b, \epsilon_\tau, \omega)$ is the vote of a type τ voter when her priors about the incumbent's type are b, her preference shock is ϵ_τ, and the current state is ω. Thus, we collapse history dependence into dependence upon beliefs, where the belief mappings are implicitly assumed. Similarly, we view the politician's strategies as a function of the voters' beliefs, the politicians' preference shocks, and the current state. In light of our symmetry assumption on politicians, we model politician strategies as mappings $\alpha_\tau : B \times \Re^d \times \Omega \to A_\tau$, where $\alpha_\tau(b, \theta_\tau, \omega)$ is the action taken by type τ politicians when the voters' priors about the officeholder's type are b, the politician's preference shock is θ_τ, and the current state is ω.

A belief stationary equilibrium σ, which we define shortly, must be stationary with respect to a sensible belief mapping for the voter. Thus, we define the Bayesian belief mapping by

$$\beta^\sigma_\tau(x, b, \omega, \omega') = \frac{b_\tau \int_\theta f_\tau(x|\alpha_\tau(b, \theta_\tau, \omega), \omega) p(\omega'|\alpha_\tau(b, \theta_\tau, \omega), \omega) h(\theta|\omega) d\theta}{\sum_{\tau'} b_{\tau'} \int_\theta f_{\tau'}(x|\alpha_{\tau'}(b, \theta_{\tau'}, \omega), \omega) p(\omega'|\alpha_{\tau'}(b, \theta_{\tau'}, \omega), \omega) h(\theta|\omega) d\theta} \quad ,$$

which yields the probability, from the perspective of the electorate, that the office-holder is type τ conditional on beginning a period in state ω with beliefs b, observing policy outcome x, and then beginning the next period in state ω'. Let $\beta^\sigma(x, b, \omega) \equiv (\beta^\sigma_\tau(x, b, \omega))_\tau \in B$ represent the voters' posterior beliefs about the incumbent's type given priors b and conditional on observing x in state ω.

Equilibrium. In order to formulate the continuation values of voters and politicians in the model, we must define the probability that a single type of voter votes to re-elect the incumbent and the probability that the incumbent wins. First, given a belief-stationary profile σ, define the probability that a type τ voter votes for the incumbent when beliefs are b, the preference shock is ϵ_τ, and the state is ω as $\varpi^\sigma_\tau(b, \epsilon_\tau, \omega) \equiv \chi_{\mathfrak{I}}(\rho_\tau(b, \epsilon_\tau, \omega))$, where $\chi_{\mathfrak{I}}(z)$ is an indicator function taking a value of one if $z = \mathfrak{I}$, zero otherwise. Then the probability that the incumbent is re-elected, conditional on b, ϵ, and ω, is given by the mapping $\varpi^\sigma\colon B \times \Re^T \times \Omega \to \{0,1\}$ and defined by

$$\varpi^\sigma(b,\epsilon,\omega) = \sum_{S\in\mathfrak{D}(\omega)} \left(\prod_{\tau\in S} \varpi^\sigma_\tau(b,\epsilon_\tau,\omega)\right)\left(\prod_{\tau\notin S}(1-\varpi^\sigma_\tau(b,\epsilon_\tau,\omega))\right).$$

Finally, the ex ante probability that the incumbent wins, integrating over the voters' shocks ϵ, is given by the mapping $\pi^\sigma_\tau\colon B\times\Omega\to[0,1]$ and defined by

$$\pi^\sigma(b,\omega) = \int_\epsilon \varpi^\sigma(b,\epsilon,\omega)h(\epsilon|\omega)d\epsilon.$$

We now define, for any belief-stationary profile σ, the voters' expected discounted utility, excluding (for notational reasons) the current period's preference shock, from electing a politician when the voters' beliefs are b and the state is ω. Define the mapping $V^\sigma_\tau\colon B\times\Omega\to\Re$ by

Our definition of cont values assumes θ and ϵ are independent.

$$\begin{aligned}
V^\sigma_\tau(b,\omega) \;=\; & \sum_{\tau'} b_{\tau'}\int_\theta\int_{\tilde b}\int_x \Bigg[u_\tau(x,\omega)+\delta_\tau(\omega)\sum_{\omega'}p(\omega'|\alpha_{\tau'}(b,\theta_{\tau'},\omega),\omega)\\
&\cdot\int_\epsilon\big[\varpi^\sigma(\beta^\sigma(x,\tilde b,\omega,\omega'),\epsilon,\omega)(V^\sigma_\tau(\beta^\sigma(x,\tilde b,\omega,\omega'),\omega')+\epsilon_\tau)\\
&+(1-\varpi^\sigma(\beta^\sigma(x,\tilde b,\omega,\omega'),\epsilon,\omega))V^\sigma_\tau(\bar b(\omega'),\omega')\big]h(\epsilon|\omega')d\epsilon\Bigg]\\
&\cdot f_{\tau'}(x|\alpha_{\tau'}(b,\theta_{\tau'},\omega),\omega)g_{\tau'}(\tilde b|b,\omega)h(\theta|\omega)dxd\tilde b d\theta.
\end{aligned}$$

Given beliefs b regarding the incumbent's type, current state ω, and shock ϵ_τ, a type τ voter's expected discounted payoff from re-electing the incumbent is then $V^\sigma_\tau(b,\omega)+\epsilon_\tau$, and the payoff from electing the challenger is $V^\sigma_\tau(\bar b(\omega),\omega)$.

We next define the officeholder's expected discounted utility, excluding (for notational reasons) the current period's preference shock, from choosing action a when the voters' beliefs are b and the state is ω. Define the mapping $W^\sigma_\tau\colon A_\tau\times B\times\Omega\to\Re$ by

$$\begin{aligned}
&W^\sigma_\tau(a,b,\omega)\\
&= \int_{\tilde b}\int_x\Bigg[u_\tau(x,a,\omega)+\delta_\tau(\omega)\sum_{\omega'}p(\omega'|a,\omega)\int_\theta\big[\pi^\sigma(\beta^\sigma(x,\tilde b,\omega,\omega'),\omega)\\
&\quad\cdot(W^\sigma_\tau(\alpha_\tau(\beta^\sigma(x,\tilde b,\omega,\omega'),\theta_\tau,\omega'),\beta^\sigma(x,\tilde b,\omega,\omega'),\omega')+\alpha_\tau(\beta^\sigma(x,\tilde b,\omega,\omega'),\theta_\tau,\omega')\cdot\theta_\tau)\\
&\quad+(1-\pi^\sigma(\beta^\sigma(x,\tilde b,\omega,\omega'),\omega'))V^\sigma_\tau(\bar b(\omega'),\omega')\big]h(\theta|\omega')d\theta\Bigg]\\
&\quad\cdot f(x|a,\omega)g_\tau(\tilde b|b,\omega)dxd\tilde b
\end{aligned}$$

Given action choice a, beliefs b regarding the officeholder's type, current state ω, and shock θ_τ, a type τ officeholder's payoff is then $W^\sigma_\tau(a, b, \omega) + a \cdot \theta_\tau$.

We define a strategy profile σ to be a *belief-stationary equilibrium* if it is stationary with respect to β^σ and satisfies equilibrium restrictions on voting strategies and politicians' actions. Our equilibrium condition on voting strategies is that a voter cast her ballot with the incumbent if and only if the expected payoff from re-electing the incumbent weakly exceeds the expected payoff from the challenger.

Optimal voting. For all τ, all b, all ϵ_τ, and all ω,

$$\rho_\tau(b, \epsilon_\tau, \omega) \quad = \quad \begin{cases} \mathfrak{I} & \text{if } \epsilon_\tau \geq V^\sigma_\tau(\bar{b}(\omega), \omega) - V^\sigma_\tau(b, \omega) \\ \mathfrak{C} & \text{else.} \end{cases}$$

Thus, in equilibrium, the probability that a type τ voter votes for the incumbent, given beliefs b and state ω, is

$$\pi^\sigma_\tau(b, \omega) = 1 - H_{\epsilon_\tau}\left(V^\sigma_\tau(\bar{b}(\omega), \omega) - V^\sigma_\tau(b, \omega)\right),$$

where H_{ϵ_τ} is the marginal distribution on ϵ_τ. Our equilibrium condition on politicians' strategies is that policy choices maximize the expected discounted utility of the politicians.

Optimal policy choice. For all τ, all b, all θ_τ, and all ω, the action $\alpha_\tau(b, \theta_\tau, \omega)$ solves

$$\max_{a \in A_\tau(\omega)} W^\sigma_\tau(a, b, \omega) + a \cdot \theta_\tau.$$

Thus, the preference shock θ_τ serves to perturb the objective function of the officeholder, and it follows from Mas-Colell's (1985) Theorem I.3.1 that the above maximization problem has a unique solution for almost all θ_τ, allowing us to focus on pure strategies.

We close this section by noting that the regularity properties of continuation values as defined above are limited. In particular, even if the politicians' strategies, α_τ, are continuous, it is difficult to place a Lipschitz bound on $V^\sigma_\tau(\cdot, \omega)$ that is uniform across σ, or more generally to prove equicontinuity of $\{V^\sigma_\tau(\cdot, \omega) \mid \sigma \text{ is belief-stationary}\}$, because b enters the righthand side of the above formula through α_τ, which is endogenous. Thus, placing a Lipschitz bound on $V^\sigma_\tau(\cdot, \omega)$ would rely on establishing regularity properties of α_τ itself. We do not take that approach. In the proof of the existence theorem, we reformulate continuation values in a less intuitive, but technically more efficacious, way.

4 Equilibrium Existence

Theorem. There exists a belief-stationary equilibrium.

The remainder of this section consists of the proof of the theorem. As is standard, we establish existence by means of a fixed point argument in a space of continuation value functions of voters and politicians. We face the standard problems of compactness and continuity. By specifying continuation values appropriately, we are able to exploit the noise in our model, through our uniform Lipschitz assumption, to obtain compactness in a space of continuous mappings with a relatively strong topology (the sup norm). Once that is done, we again exploit the noise in our model to establish that best responses for voters and politicians are almost surely unique, and this in turn, with the sup norm on the domain of continuation values, enables us to establish continuity of the fixed point mapping. We denote the continuation values in the domain of our mapping by v and w to distinguish them from the more intuitive quantities defined in the previous section, and for notational purposes, we often view mappings as vector valued, e.g., $v = (v_\tau)_\tau$, where a τ subscript selects one coordinate in this vector.

Letting $\mathfrak{A} = \{(\tau, a) \in T \times \Re^d \mid a \in A_\tau\}$ denote the graph of the feasible action correspondence, we define the fixed point mapping Ψ on mappings $v \colon \mathfrak{A} \times B \times \Omega \to \Re^T$ and $w \colon \mathfrak{A} \times B \times \Omega \to \Re$, where we embed (v, w) in a compact subspace $\mathfrak{X}$ of $C^1(\mathfrak{A} \times B \times \Omega, \Re^T \times \Re)$.[14] Here, we interpret $v_\tau(\tau', a, b, \omega)$ as the expected discounted utility of a type τ voter, assuming the state in the previous period was ω, the voters' prior on the officeholder at the beginning of the previous period was b, and the officeholder in the previous period (today's incumbent) chose action a in the previous period and is type τ'. Similarly, $w(\tau, a, b, \omega)$ is the expected discounted utility of a type τ incumbent who in chose action a in the previous period when the state was ω and the voters' prior was b. This is not a natural accounting method, but it serves an important technical function: these quantities are sufficient to determine the continuation values V and W from the previous period, yet the arguments of these function enter only into

[14] As a technicality, we imbed the finite sets Ω and T in $\Re$ in an arbitrary way, so that $||(\tau, a, b, \omega) - (\overline{\tau}, \overline{a}, \overline{b}, \overline{\omega})||$ sufficiently small implies $\omega = \overline{\omega}$ and $\tau = \overline{\tau}$.

exogenously specified density functions.

To define $\mathfrak{X}$, let $L_4 = \int_B \int_X dxdb$ be the product of the Lebesgue measure of X and the Lebesgue measure of the simplex B, and let $\overline{\delta} = \max_{\tau,\omega} \delta_\tau(\omega) < 1$, where the inequality follows from finiteness of T and Ω. Then $\mathfrak{X}$ consists of pairs (v, w) of mappings such that (1) for all τ,

$$\max\Big\{\sup_{\tau',a,b,\omega} |v_\tau(\tau', a, b, \omega)|, \sup_{\tau',a,b,\omega} |w(\tau', a, b, \omega)|\Big\} \quad \leq \quad \frac{L_2}{1-\overline{\delta}},$$

and (2) for all τ, all τ', and all ω, $v_\tau(\tau', a, b, \omega)$ and $w(\tau', a, b, \omega)$ are Lipschitz in (a, b) with uniform bound

$$\frac{L_2 L_3 |\Omega| + L_1 L_2 L_4}{1-\overline{\delta}}.$$

That is, for all (a, b) and $(\overline{a}, \overline{b})$,

$$\begin{aligned} &\max\big\{|v_\tau(\tau', a, b, \omega) - v_\tau(\tau', \overline{a}, \overline{b}, \omega)|, |w(\tau', a, b, \omega) - w(\tau', \overline{a}, \overline{b}, \omega)|\big\} \\ \leq \quad &\left[\frac{L_2 L_3 |\Omega| + L_1 L_2 L_4}{1-\overline{\delta}}\right] ||(a, b) - (\overline{a}, \overline{b})||. \end{aligned}$$

We claim that $\mathfrak{X}$ is compact. Obviously, it is bounded, i.e., there is a uniform bound on $||(v(\tau', a, b, \omega), w(\tau', a, b, \omega)||$ for $(v, w) \in \mathfrak{X}$, and closed, i.e., if a sequence $\{(v^\ell, w^\ell)\}$ in $\mathfrak{X}$ converges uniformly to (v, w), then $(v, w) \in \mathfrak{X}$. Furthermore, we claim that $\mathfrak{X}$ is equicontinuous. Consider any $\zeta > 0$, and choose $\eta > 0$ small enough that $||(\tau, a, b, \omega) - (\overline{\tau}, \overline{a}, \overline{b}, \overline{\omega})|| < \eta$ implies $\tau = \overline{\tau}$ and $\omega = \overline{\omega}$. Moreover, choose η small enough that

$$\eta\sqrt{|T| + 1}\left[\frac{L_2 L_3 |\Omega| + L_1 L_2 L_4}{1-\overline{\delta}}\right] \quad \leq \quad \zeta.$$

Then given any $(v, w) \in \mathfrak{X}$ and any (τ, a, b, ω) and $(\overline{\tau}, \overline{a}, \overline{b}, \overline{\omega})$ with $||(\tau, a, b, \omega) - (\overline{\tau}, \overline{a}, \overline{b}, \overline{\omega}))|| \leq \eta$, we have

$$\begin{aligned} &||(v(\tau, a, b, \omega), w(\tau, a, b, \omega)) - (v(\overline{\tau}, \overline{a}, \overline{b}, \overline{\omega})), w(\overline{\tau}, \overline{a}, \overline{b}, \overline{\omega})))|| \\ \leq \quad &\sqrt{|T| + 1}\left[\frac{L_2 L_3 |\Omega| + L_1 L_2 L_4}{1-\overline{\delta}}\right] ||(a, b) - (\overline{a}, \overline{b})|| \\ \leq \quad &\sqrt{|T| + 1}\left[\frac{L_2 L_3 |\Omega| + L_1 L_2 L_4}{1-\overline{\delta}}\right] \eta \\ \leq \quad &\zeta, \end{aligned}$$

where the first equality uses $\tau = \overline{\tau}$, $\omega = \overline{\omega}$, and property (2) in the definition of $\mathfrak{X}$. Therefore, $\mathfrak{X}$ is equicontinuous, and we conclude that it is compact, as claimed.

To define the mapping Ψ, take $(v,w) \in \mathfrak{X}$ as given, and write the value of the mapping at (v,w) as $(\hat{v},\hat{w}) = \Psi(v,w)$. Intuitively, $(\hat{v},\hat{w})$ are the continuation values induced by optimal behavior of voters and politicians, assuming, hypothetically, that (v,w) accurately reflect their future payoffs. We define these best response continuation value functions in a number of steps, at each step constructing a mapping used in the specification of $(\hat{v},\hat{w})$. We use hats to distinguish these "intermediate" objects. At each step, we use (v,w) in our constructions, and we express this dependence through the superscript v,w. Although we view the intermediate mappings we define as depending parametrically on v and w, we will see that joint continuity, in (v,w) as well as in the arguments of the intermediate mappings, is critical for the fixed point argument.

We first construct the continuation value for an officeholder incorporating stage utility from the current period but omitting the current period's preference shock, paralleling the definition of W in the previous section. Define the mappings $\hat{W}^{v,w}_\tau \colon A_\tau \times B \times \Omega \to \Re$ by

$$\hat{W}^{v,w}_\tau(a,b,\omega) \quad = \quad \int_x [u_\tau(x,a,\omega) + \delta_\tau(\omega) w(a,b,\omega,\tau)] f_\tau(x|a,\omega)dx \tag{1}$$

for all τ. Under the interpretation of w given above, $\hat{W}^{v,w}_\tau(a,b,\omega)$ then corresponds to a type τ officeholder's expected discounted utility from choosing action a, not including the current preference shock, when the current state is ω and the voters' beliefs regarding her type are b. Note that $\hat{W}^{v,w}_\tau(a,b,\omega)$ is jointly continuous in (a,b,ω,v,w). Indeed, consider a sequence $\{(a^\ell,b^\ell,\omega^\ell,v^\ell,w^\ell)\}$ converging to some (a,b,ω,v,w). Fix arbitrary x, and note that $u_\tau(x,a^\ell,\omega^\ell) \to u_\tau(x,a,\omega)$ by continuity of u_τ, $f_\tau(x|a^\ell,\omega^\ell) \to f_\tau(x|a,\omega)$ by continuity of f_τ, and because w^ℓ is converges uniformly to w, we have $\delta_\tau(\omega^\ell)w^\ell(a^\ell,b^\ell,\omega^\ell,\tau) \to w(a,b,\omega,\tau)$. Thus, the integrand on the righthand side of (1) converges pointwise to $[u_\tau(x,a,\omega) + \delta_\tau(\omega)w(a,b,\omega,\tau)]f_\tau(x|a,\omega)$. Furthermore, by continuity of u_τ, w, and f, and using compactness of $A \times B \times \Omega$ and X, it follows that the sequence of integrands is uniformly bounded by an integrable function. Therefore, Lebesgue's dominated convergence theorem implies that $\hat{W}^{v^\ell,w^\ell}_\tau(a^\ell,b^\ell,\omega^\ell) \to \hat{W}^{v,w}_\tau(a,b,\omega)$, as claimed. Compactness of X needed here?

Given a realization θ_τ of the officeholder's preference shock, she solves

$$\max_{a \in A_\tau(\omega)} \hat{W}_\tau(a,b,\omega) + \theta_\tau \cdot a.$$

Since $A_\tau(\omega)$ is compact, it follows that the maximization problem in (**??**) admits at least one solution for each $(b,\theta_\tau,\omega,v,w)$. Then Aliprantis and Border's (1999) Theorem 17.18 yields a mapping $\hat{\alpha}^{v,w} \colon B \times \Re^d \times \Omega \to \prod_\tau A_\tau$ such that for each τ, $\hat{\alpha}^{v,w}_\tau$

selects solutions to (**??**) and is jointly measurable in $(b, \theta, \omega, v, w)$. Furthermore, Mas-Colell's (1985) Theorem I.3.1 implies that for all (b, ω, v, w), there is a measure zero set $\Theta^{v,w}_\tau(b, \omega) \subseteq \Re^d$ such that for all $\theta_\tau \notin \Theta^{v,w}_\tau(b, \omega)$, the action $\hat{\alpha}^{v,w}_\tau(b, \theta_\tau, \omega)$ is the unique solution to (**??**). As a consequence, the theorem of the maximum implies that for all (b, ω, v, w) and all $\theta_\tau \notin \Theta_\tau(b, \omega)$, the selection $\hat{\alpha}^{v,w}_\tau(b, \theta_\tau, \omega)$ is continuous at $(b, \theta_\tau, \omega, v, w)$.

Having solved for the (essentially) unique best response actions of officeholders, we can now pin down the updating of the voters' prior via Bayes rule. Define $\hat{\beta}^{v,w}: X \times B \times \Omega \times \Omega \to B$ by

$$\hat{\beta}^{v,w}_\tau(x, b, \omega, \omega') = \frac{b_\tau \int_\theta f_\tau(x|\hat{\alpha}^{v,w}_\tau(b, \theta_\tau, \omega), \omega) p_\tau(\omega'|\hat{\alpha}^{v,w}_\tau(b, \theta_\tau, \omega), \omega) h(\theta|\omega) d\theta}{\sum_{\tau'} b_{\tau'} \int_\theta f_{\tau'}(x|\hat{\alpha}^{v,w}_{\tau'}(b, \theta_{\tau'}, \omega), \omega) p_{\tau'}(\omega'|\hat{\alpha}^{v,w}_{\tau'}(b, \theta_{\tau'}, \omega), \omega) h(\theta|\omega) d\theta}.$$

Here, $\hat{\beta}^{v,w}_\tau(x, b, \omega, \omega')$ is the voters' posterior regarding the type of the incumbent assuming that in the previous period, the voters' prior was b, the state was ω, the outcome realized was x, and the state in the current period is ω'. We claim that $\hat{\beta}$ is jointly continuous in $(x, b, \omega, \omega', v, w)$. To see this, consider a sequence $\{(x^\ell, b^\ell, \omega^\ell, \omega'^\ell, v^\ell, w^\ell)\}$ converging to $(x, b, \omega, \omega', v, w)$. We focus on continuity of the numerator, as continuity of the denominator follows by similar arguments. Consider any θ with $\theta_\tau \notin \Theta^{v,w}_\tau(b, \omega)$, and note that by our above arguments, we have $\hat{\alpha}^{v^\ell,w^\ell}_\tau(b^\ell, \theta_\tau, \omega^\ell) \to \hat{\alpha}^{v,w}_\tau(b, \theta_\tau, \omega)$. Therefore,

$$\begin{aligned} &f_\tau(x|\hat{\alpha}^{v^\ell,w^\ell}_\tau(b^\ell, \theta_\tau, \omega^\ell), \omega^\ell) p_\tau(\omega'^{,\ell}|\hat{\alpha}^{v^\ell,w^\ell}_\tau(b^\ell, \theta_\tau, \omega^\ell), \omega^\ell) h(\theta|\omega^\ell) \\ &\quad \to f_\tau(x|\hat{\alpha}^{v,w}_\tau(b, \theta_\tau, \omega), \omega) p_\tau(\omega'|\hat{\alpha}^{v,w}_\tau(b, \theta_\tau, \omega), \omega) h(\theta|\omega) \end{aligned}$$

for almost all θ. By Lebesgue's dominated convergence theorem, the numerator converges, as claimed.

We next construct continuation values for the voters incorporating stage utility from the current period but omitting the current period's preference shock, paralleling the definition of V in the previous section. Define the mappings $\hat{V}^{v,w}_\tau : B \times \Omega \to \Re$ by

$$\begin{aligned} \hat{V}^{v,w}_\tau(b, \omega) &= \sum_{\tau'} b_{\tau'} \int_\theta \int_x [u_\tau(x, \omega) + \delta_\tau(\omega) v_\tau(\hat{\alpha}^{v,w}_{\tau'}(b, \theta_{\tau'}, \omega), b, \omega, \tau')] \\ &\quad \cdot f_{\tau'}(x|\hat{\alpha}^{v,w}_{\tau'}(b, \theta_{\tau'}, \omega), \omega) h(\theta|\omega) dx d\theta \end{aligned}$$

for all τ. In words, $\hat{V}_\tau(b, \omega)$ is the type τ voter's expected discounted payoff, not including the current preference shock, from electing a candidate whose type is distributed according to b in state ω. Note that $\hat{V}^{v,w}_\tau(b, \omega)$ is jointly continuous in (b, ω, v, w). This follow from arguments similar to those above, exploiting joint continuity of u_τ, f_τ, and continuity and uniform convergence of v_τ, and we omit the argument.

We can now construct best response voting strategies. Define the mappings $\hat{\rho}_\tau^{v,w}\colon B \times \Re^T \times \Omega \to C$ by

$$\hat{\rho}_\tau^{v,w}(b, \epsilon_\tau, \omega) = \begin{cases} \mathfrak{I} & \text{if } \epsilon_\tau \geq \hat{V}_\tau^{v,w}(\bar{b}(\omega), \omega) - \hat{V}_\tau^{v,w}(b, \omega), \\ \mathfrak{C} & \text{else} \end{cases}$$

for all τ. Define the probability that a type τ voter votes for the incumbent when beliefs are b, the preference shock is ϵ_τ, and the state is ω as $\hat{\varpi}_\tau^{v,w}(b, \epsilon_\tau, \omega) \equiv \chi_{\mathfrak{I}}(\hat{\rho}_\tau^{v,w}(b, \epsilon_\tau, \omega))$, where $\chi_{\mathfrak{I}}(z)$ is an indicator function taking a value of one if $z = \mathfrak{I}$, zero otherwise. Then the probability that the incumbent is re-elected, conditional on b, ϵ, and ω, is given by the mapping $\hat{\varpi}^{v,w}\colon B \times \Re^T \times \Omega \to \{0, 1\}$ and defined by

$$\hat{\varpi}^{v,w}(b, \epsilon, \omega) = \sum_{S \in \mathfrak{D}(\omega)} \left(\prod_{\tau \in S} \hat{\varpi}_\tau^{v,w}(b, \epsilon_\tau, \omega) \right) \left(\prod_{\tau \notin S} (1 - \hat{\varpi}_\tau^{v,w}(b, \epsilon_\tau, \omega)) \right).$$

Note that for fixed (b, ω, v, w), the set of preference shocks ϵ_τ that are exactly equal to $\hat{V}_\tau^{v,w}(\bar{b}, \omega) - \hat{V}_\tau^{v,w}(b, \omega)$ has probability zero. Thus, by joint continuity of $\hat{V}_\tau$, if $(b^\ell, \omega^\ell, v^\ell, w^\ell) \to (b, \omega, v, w)$, then $\hat{\varpi}^{v^\ell, w^\ell}(b^\ell, \epsilon, \omega^\ell) \to \hat{\varpi}^{v,w}(b, \epsilon, \omega)$ almost surely in ϵ. The ex ante probability that the incumbent wins, integrating over the voters' shocks ϵ, is given by the mapping $\pi^{v,w}\colon B \times \Omega \to [0, 1]$ and defined by

$$\hat{\pi}^{v,w}(b, \omega) = \int_\epsilon \hat{\varpi}^{v,w}(b, \epsilon, \omega) h(\epsilon|\omega) d\epsilon.$$

Note that because H_{ϵ_τ} has a density, the dominated convergence theorem implies that the mapping $\hat{\pi}$ is jointly continuous in (b, ω, v, w).

Finally, we can specify the best response continuation values $(\hat{v}, \hat{w}) = \Psi(v, w)$. For simplicity, we decompose each coordinate $\hat{v}_\tau$ as $\hat{v}_\tau = \hat{v}_\tau^{\mathfrak{I}} + \hat{v}_\tau^{\mathfrak{C}}$, where $\hat{v}_\tau^{\mathfrak{I}}\colon \mathfrak{A} \times B \times \Omega \to \Re$ is defined by

$$\begin{aligned} \hat{v}_\tau^{\mathfrak{I}}(\tau', a, b, \omega) \quad = \quad & \int_{\tilde{b}} \int_x \sum_{\omega'} p_{\tau'}(\omega'|a, \omega) \int_{(\theta, \epsilon)} \hat{\varpi}^{v,w}(\hat{\beta}^{v,w}(x, \tilde{b}, \omega, \omega'), \epsilon, \omega') \\ & \cdot \left[\int_{x'} (u_\tau(x', \omega') + \epsilon_\tau) f_{\tau'}(x'|\hat{\alpha}_{\tau'}^{v,w}(\hat{\beta}^{v,w}(x, \tilde{b}, \omega, \omega'), \theta_{\tau'}, \omega'), \omega') dx' \right. \\ & \left. + \delta_\tau(\omega') v_\tau(\tau', \hat{\alpha}_{\tau'}^{v,w}(\hat{\beta}^{v,w}(x, \tilde{b}, \omega, \omega'), \theta_{\tau'}, \omega'), \hat{\beta}^{v,w}(x, \tilde{b}, \omega, \omega'), \omega') \right] \\ & \cdot h(\theta, \epsilon|\omega') f_{\tau'}(x|a, \omega) g_{\tau'}(\tilde{b}|b, \omega) d(\theta, \epsilon) dx d\tilde{b}. \end{aligned}$$

We read this formula as follows. Given that the incumbent is type τ', the previous period's prior and state were b and ω, and she chose action a, we first need to "reach back" to determine the current period's prior regarding the incumbent's type. Thus, we integrate over $\tilde{b}$, reflecting exogenous information received by the voters, and over outcomes x. We integrate over the current state ω', and the voters' updated beliefs are $\hat{\beta}^{v,w}(x,\tilde{b},\omega')$. The $\hat{v}^{\mathfrak{I}}$ term collects realizations of (θ,ϵ) for which the incumbent is re-elected. Given such a realization, the officeholder chooses action $\hat{\alpha}^{v;w}_{\tau'}(\hat{\beta}^{v,w}(x,b',\omega,\omega'),\theta_{\tau'},\omega')$, an outcome x' is realized, and the voter's stage utility in the current period is $u_\tau(x',\omega')+\epsilon_\tau$. Completing the recursion, we then move to the next period, discounted by $\delta_{\tau'}(\omega')$, when the voter's expected discounted payoff is given by v_τ using the officeholder's action and the priors on the officeholder for the current period.

We define the mapping $\hat{v}^{\mathfrak{C}}_\tau\colon \mathfrak{A}\times B\times\Omega\rightarrow\Re$ by

$$\begin{aligned}
\hat{v}^{\mathfrak{C}}_\tau(\tau',a,b,\omega) \quad = \quad & \int_{\tilde{b}}\int_x\sum_{\omega'} p_{\tau'}(\omega'|a,\omega)\int_{(\theta,\epsilon)}(1-\hat{\varpi}^{v,w}(\hat{\beta}(x,b',\omega,\omega'),\epsilon,\omega')) \\
& \cdot\sum_{\tau''}\bar{b}_{\tau''}(\omega')\Bigg[\int_{x'}u_\tau(x',\omega')f_{\tau''}(x'|\hat{\alpha}^{v;w}_{\tau''}(\bar{b}(\omega'),\theta_{\tau''},\omega'),\omega')dx' \\
& +\delta_\tau(\omega')v_\tau(\tau'',\hat{\alpha}^{v;w}_{\tau''}(\bar{b}(\omega'),\theta_{\tau''},\omega'),\bar{b}(\omega''))\Bigg] \\
& \cdot h(\theta,\epsilon|\omega')f_{\tau'}(x|a,\omega)g_{\tau'}(\tilde{b}|b,\omega)d(\theta,\epsilon)dxd\tilde{b}.
\end{aligned}$$

We read this similarly, though now we collect (θ,ϵ) realizations for which the challenger is elected, and we integrate over the challenger's type τ'' and preference shock $\theta_{\tau''}$. The officeholder, formerly the challenger, chooses $\hat{\alpha}^{v;w}_{\tau''}(\bar{b}(\omega'),\theta_{\tau''},\omega')$, the outcome x' is realized, and we move to the next period with continuation value v_τ.

We decompose $\hat{w}$ as $\hat{w}=\hat{w}^{\mathfrak{I}}+\hat{w}^{\mathfrak{C}}$, where $\hat{w}^{\mathfrak{I}}\colon\mathfrak{A}\times B\times\Omega\rightarrow\Re$ is defined by

$$\begin{aligned}
\hat{w}^{\mathfrak{I}}(\tau,a,b,\omega) \quad = \quad & \int_{\tilde{b}}\int_x\sum_{\omega'}p_\tau(\omega'|a,\omega)\hat{\pi}^{v,w}(\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\omega') \\
& \cdot\int_\theta\Bigg[\int_{x'}u_\tau(x',\hat{\alpha}^{v,w}_\tau(\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\theta_\tau,\omega')) \\
& \cdot f(x'|\hat{\alpha}^{v,w}_\tau(\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\theta_\tau,\omega'),\omega')dx' \\
& +\hat{\alpha}^{v,w}_\tau(\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\theta_\tau,\omega')\cdot\theta_\tau+\delta_\tau(\omega') \\
& \cdot w(\tau,\hat{\alpha}^{v,w}_\tau(\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\theta_\tau,\omega'),\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\omega')\Bigg] \\
& \cdot h(\theta|\omega')f_\tau(x|a,\omega)g_\tau(\tilde{b}|b,\omega)d\theta dxd\tilde{b}.
\end{aligned}$$

We read this as follows. We reach back to determine the current period's prior regarding the incumbent's type, integrating over $\tilde{b}$ and the outcome x in the previous period. We integrate over the current state ω', and the voters' updated beliefs are $\hat{\beta}^{v,w}(x,\tilde{b},\omega')$. The $\hat{w}^{\mathfrak{I}}$ term accounts for the case in which the incumbent is re-elected, as reflected in the weight $\hat{\pi}^{v,w}(\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\omega')$. In the current period, the officeholder's preference shock θ_τ is realized, she takes action $\hat{\alpha}^{v,w}_\tau(\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'))$, an outcome x' is realized, and stage utility accrues to the officeholder for the current period. Completing the recursion, we then move to the next period, discount, and insert the continuation value w using the officeholder's action and the priors on the officeholder for the current period.

Define the mapping $\hat{w}^{\mathfrak{C}} \colon \mathfrak{A} \times B \times \Omega \to \Re$ by

$$\begin{aligned} \hat{w}^{\mathfrak{C}}(\tau,a,b,\omega) \quad = \quad & \int_{\tilde{b}}\int_x \sum_{\omega'} p_\tau(\omega'|a,\omega)(1-\hat{\pi}^{v,w}(\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\omega')) \\ & \cdot \sum_{\tau'} \overline{b}_{\tau'}(\omega') \int_\theta \Bigg[\int_{x'} u_\tau(x',\omega') f_{\tau'}(x'|\hat{\alpha}^{v,w}_{\tau'}(\overline{b}(\omega'),\theta_{\tau'},\omega'),\omega')dx' \\ & + \delta_\tau(\omega')v_\tau(\tau',\hat{\alpha}^{v,w}_{\tau'}(\overline{b}(\omega'),\theta_{\tau'},\omega'),\overline{b}(\omega'')\Bigg] \\ & \cdot h(\theta|\omega')f_\tau(x|a,\omega)g_\tau(\tilde{b}|b,\omega)d\theta dx d\tilde{b}. \end{aligned}$$

We read this as before, though now we account for the case in which the incumbent loses the election and the challenger is elected. We reach back to the previous period to obtain the voters' prior beliefs regarding the incumbent's type, but only to calculate the weight $1-\hat{\pi}^{v,w}(\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\omega')$. We integrate over the challenger's type and preference shock, an outcome is realized for the current period, and stage utility accrues to the out-of-office politician as it would any other type τ voter. We then discount and move to the next period, where the expected discounted utility is given by v_τ.

We claim that Ψ maps $\mathfrak{X}$ into $\mathfrak{X}$. In showing property (1) in the definition of the domain, we focus on the inequality $\sup_{\tau',a,b,\omega} |\hat{v}_\tau(\tau',a,b,\omega)| \leq \frac{L_2}{1-\overline{\delta}}$, as a similar argument establishes the bound on w. Consider any $(v,w) \in \mathfrak{X}$, and let $(\hat{v},\hat{w}) =$

$\Psi(v,w)$. To simplify algebra, we define

$$\begin{aligned}
U(x,\tilde{b},\omega') \quad = \quad & \int_{(\theta,\epsilon)} \Big\{ \hat{\varpi}^{v,w}(\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\epsilon,\omega') \\
& \cdot \Big[\int_{x'} (u_\tau(x',\omega')+\epsilon_\tau) f_{\tau'}(x'|\hat{\alpha}^{v,w}_{\tau'}(\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\theta_{\tau'},\omega'),\omega')dx' \\
& +\delta_\tau(\omega')v_\tau(\tau',\hat{\alpha}^{v,w}_{\tau'}(\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\theta_{\tau'},\omega'),\hat{\beta}^{v,w}(x,\tilde{b},\omega,\omega'),\omega')\Big] \\
& +(1-\hat{\varpi}^{v,w}(\hat{\beta}(x,b',\omega,\omega'),\epsilon,\omega')) \\
& \cdot \sum_{\tau''} \overline{b}_{\tau''}(\omega') \Big[\int_{x'} u_\tau(x',\omega') f_{\tau''}(x'|\hat{\alpha}^{v,w}_{\tau''}(\overline{b}(\omega'),\theta_{\tau''},\omega'),\omega')dx' \\
& +\delta_\tau(\omega')v_\tau(\tau'',\hat{\alpha}^{v,w}_{\tau''}(\overline{b}(\omega'),\theta_{\tau''},\omega'),\overline{b}(\omega'')\Big]\Big\} h(\theta,\epsilon|\omega')d(\theta,\epsilon)
\end{aligned}$$

for each x, $\tilde{b}$, and ω. Then

$$\hat{v}_\tau(\tau',a,b,\omega) \quad = \quad \int_{\tilde{b}}\int_x \sum_{\omega'} p_{\tau'}(\omega'|a,\omega)U(x,\tilde{b},\omega')f_{\tau'}(x|a,\omega)g_{\tau'}(\tilde{b}|b,\omega)dxd\tilde{b},$$

which implies

$$\sup_{\tau',a,b,\omega} |\hat{v}_\tau(\tau',a,b,\omega)| \quad \leq \quad \sup_{x,\tilde{b},\omega'} |U(x,\tilde{b},\omega')| \leq L_2 + \overline{\delta} \sup_{\tau',a,b,\omega} |v_\tau(\tau',a,b,\omega)|,$$

which yields the claimed inequality. Note further that $\sup_{x,\tilde{b},\omega'} |U(x,\tilde{b},\omega')| \leq \frac{L_2}{1-\overline{\delta}}$.

Now fix ω and τ', consider any τ and any (a,b) and $(\overline{a},\overline{b})$. We argue that

$$\begin{aligned}
& \max\left\{|\hat{v}_\tau(a,b,\omega,\tau')-\hat{v}_\tau(\overline{a},\overline{b},\omega,\tau')|,|\hat{w}(a,b,\omega,\tau')-\hat{w}(\overline{a},\overline{b},\omega,\tau')|\right\} \\
& \leq \quad \frac{L_2L_3|\Omega|+L_1L_2L_4}{1-\overline{\delta}}||(a,b)-(\overline{a},\overline{b})||.
\end{aligned}$$

We focus on $|\hat{v}_\tau(\tau',a,b,\omega)-\hat{v}_\tau(\tau',\overline{a},\overline{b},\omega)|$, as a similar argument holds for $|\hat{w}(\tau',a,b,\omega)-$

$\hat{w}(\tau', \overline{a}, \overline{b}, \omega)|$. Note that

$$\begin{aligned}
& |\hat{v}_\tau(\tau', a, b, \omega) - \hat{v}_\tau(\tau', \overline{a}, \overline{b}, \omega)| \\
\leq \quad & \left| \int_{\tilde{b}} \int_x \sum_{\omega'} p_{\tau'}(\omega'|a, \omega) U(x, \tilde{b}, \omega') f_{\tau'}(x|a, \omega) g_{\tau'}(\tilde{b}|b, \omega) dx d\tilde{b} \right. \\
& \left. - \int_{\tilde{b}} \int_x \sum_{\omega'} p_{\tau'}(\omega'|\overline{a}, \omega) U(x, \tilde{b}, \omega') f_{\tau'}(x|\overline{a}, \omega) g_{\tau'}(\tilde{b}|\overline{b}, \omega) dx d\tilde{b} \right| \\
\leq \quad & \sum_{\omega'} \left| (p_{\tau'}(\omega'|a, \omega) - p_{\tau'}(\omega'|\overline{a}, \omega)) \int_{\tilde{b}} \int_x U(x, \tilde{b}, \omega') f_\tau(x|a, \omega) g_{\tau'}(\tilde{b}|b, \omega) dx d\tilde{b} \right. \\
& \left. + p_{\tau'}(\omega'|\overline{a}, \omega) \int_{\tilde{b}} \int_x U(x, \tilde{b}, \omega') [f_{\tau'}(x|a, \omega) g_{\tau'}(\tilde{b}|b, \omega) - f_{\tau'}(x|\overline{a}, \omega) g_{\tau'}(\tilde{b}|\overline{b}, \omega)] dx d\tilde{b} \right| \\
\leq \quad & \sum_{\omega'} \left[L_3 ||a - \overline{a}|| \frac{L_2}{1 - \overline{\delta}} + p_{\tau'}(\omega'|\overline{a}, \omega) \int_{\tilde{b}} \int_x \frac{L_2}{1 - \overline{\delta}} L_1 ||(a, b) - (\overline{a}, \overline{b})|| dx d\tilde{b} \right] \\
\leq \quad & \left[\frac{L_2 L_3 |\Omega|}{1 - \overline{\delta}} + \frac{L_1 L_2 L_3}{1 - \overline{\delta}} \right] ||(a, b) - (\overline{a}, \overline{b})||,
\end{aligned}$$

as desired. Omitting the analogous argument for $\hat{w}$, we conclude that $\Psi(v, w) = (\hat{v}, \hat{w}) \in \mathfrak{X}$, as desired.

The key, and final, step in the proof is to establish continuity of the mapping Ψ. Using compactness of $\mathfrak{X}$, consider a sequence $\{(v^\ell, w^\ell)\}$ in $\mathfrak{X}$ such that $(v^\ell, w^\ell) \to (v, w)$. Let $(\hat{v}^\ell, \hat{w}^\ell) = \Psi(v^\ell, w^\ell)$ for all ℓ, and let $(\hat{v}, \hat{w}) = \Psi(v, w)$. We must show that $(\hat{v}^\ell, \hat{w}^\ell) \to (\hat{v}, \hat{w})$. For each ℓ, there exist $\hat{\alpha}^\ell$, $\hat{\beta}^\ell$, $\hat{V}^\ell$, $\hat{\rho}^\ell$, $\hat{\varpi}^\ell$, and $\hat{\pi}^\ell$ (suppressing v^ℓ and w^ℓ in the superscript) such that $\hat{v}^\ell = \hat{v}^{\mathfrak{I},\ell} + \hat{v}^{\mathfrak{E},\ell}$, where

$$\begin{aligned}
\hat{v}^{\mathfrak{I},\ell}_\tau(\tau', a, b, \omega) \quad = \quad & \int_{\tilde{b}} \int_x \sum_{\omega'} p_{\tau'}(\omega'|a, \omega) \int_{(\theta,\epsilon)} \hat{\varpi}^\ell(\hat{\beta}^\ell(x, \tilde{b}, \omega, \omega'), \epsilon, \omega') \\
& \cdot \left[\int_{x'} (u_\tau(x', \omega') + \epsilon_\tau) f_{\tau'}(x'|\hat{\alpha}^\ell_{\tau'}(\hat{\beta}^\ell(x, \tilde{b}, \omega, \omega'), \theta_{\tau'}, \omega'), \omega') dx' \right. \\
& \left. + \delta_\tau(\omega') v^\ell_\tau(\tau', \hat{\alpha}^\ell_{\tau'}(\hat{\beta}^\ell(x, \tilde{b}, \omega, \omega'), \theta_{\tau'}, \omega'), \hat{\beta}^\ell(x, b', \omega, \omega'), \omega') \right] \\
& \cdot h(\theta, \epsilon|\omega') f_{\tau'}(x|a, \omega) g_{\tau'}(\tilde{b}|b, \omega) d(\theta, \epsilon) dx d\tilde{b} \\
\hat{v}^{\mathfrak{E},\ell}_\tau(\tau', a, b, \omega) \quad = \quad & \int_{\tilde{b}} \int_x \sum_{\omega'} p_{\tau'}(\omega'|a, \omega) \int_{(\theta,\epsilon)} (1 - \hat{\varpi}^\ell(\hat{\beta}(x, b', \omega, \omega'), \epsilon, \omega')) \\
& \cdot \sum_{\tau''} \overline{b}_{\tau''}(\omega') \left[\int_{x'} u_\tau(x, \omega') f_{\tau''}(x'|\hat{\alpha}^\ell_{\tau''}(\overline{b}(\omega'), \theta_{\tau''}, \omega'), \omega') dx' \right. \\
& \left. + \delta_\tau(\omega') v^\ell_\tau(\tau'', \hat{\alpha}^\ell_{\tau''}(\overline{b}(\omega'), \theta_{\tau''}, \omega'), \overline{b}(\omega'') \right] \\
& \cdot h(\theta, \epsilon|\omega') f_{\tau'}(x|a, \omega) g_{\tau'}(\tilde{b}|b, \omega) d(\theta, \epsilon) dx d\tilde{b},
\end{aligned}$$

and $\hat{w}^\ell = \hat{w}^{\mathfrak{I},\ell} + \hat{w}^{\mathfrak{C},\ell}$, where

$$\begin{aligned}
\hat{w}^{\mathfrak{I},\ell}(\tau,a,b,\omega) &= \int_{\tilde{b}}\int_x \sum_{\omega'} p_\tau(\omega'|a,\omega)\hat{\pi}^\ell(\hat{\beta}^\ell(x,\tilde{b},\omega,\omega'),\omega') \\
&\cdot \int_\theta \Big[\int_{x'} u_\tau(x',\hat{\alpha}^\ell_\tau(\hat{\beta}^\ell(x,\tilde{b},\omega,\omega'),\theta_\tau,\omega')) \\
&\cdot f(x'|\hat{\alpha}^\ell_\tau(\hat{\beta}^\ell(x,\tilde{b},\omega,\omega'),\theta_\tau,\omega'),\omega')dx' \\
&+ \hat{\alpha}^\ell_\tau(\hat{\beta}^\ell(x,\tilde{b},\omega,\omega'),\theta_\tau,\omega')\cdot\theta_\tau + \delta_\tau(\omega') \\
&\cdot w^\ell(\tau,\hat{\alpha}^\ell_\tau(\hat{\beta}^\ell(x,\tilde{b},\omega,\omega'),\theta_\tau,\omega'),\hat{\beta}^\ell(x,\tilde{b},\omega,\omega'),\omega')\Big] \\
&\cdot h(\theta|\omega')f_\tau(x|a,\omega)g_\tau(\tilde{b}|b,\omega)d\theta dx d\tilde{b} \\
\hat{w}^{\mathfrak{C},\ell}(\tau,a,b,\omega) &= \int_{\tilde{b}}\int_x \sum_{\omega'} p_\tau(\omega'|a,\omega)(1-\hat{\pi}^\ell(\hat{\beta}^\ell(x,\tilde{b},\omega,\omega'),\omega')) \\
&\cdot \sum_{\tau'} \bar{b}_{\tau'}(\omega')\Big[\int_\theta\int_{x'} u_\tau(x',\omega')f_{\tau'}(x'|\hat{\alpha}^\ell_{\tau'}(\bar{b}(\omega'),\theta_{\tau'},\omega'),\omega')dx' \\
&+ \delta_\tau(\omega')v^\ell_\tau(\tau',\hat{\alpha}^\ell_{\tau'}(\bar{b}(\omega'),\theta_{\tau'},\omega'),\bar{b}(\omega''))\Big] \\
&\cdot h(\theta|\omega')f_\tau(x|a,\omega)g_\tau(\tilde{b}|b,\omega)d\theta dx d\tilde{b}.
\end{aligned}$$

Let $\hat{\alpha}$, $\hat{\beta}$, $\hat{V}$, $\hat{\rho}$, $\hat{\varpi}$, and $\hat{\pi}$ correspond to $(\hat{v},\hat{w}) = \Psi(v,w)$ in the same way.

Now consider an arbitrary sequence $\{(\tau^\ell,a^\ell,b^\ell,\omega^\ell)\}$ in $\mathfrak{A}\times B\times\Omega$ converging to any (τ',a,b,ω). With compactness of $\mathfrak{A}\times B\times\Omega$, continuity of Ψ follows if we show

1. $\hat{v}^{\mathfrak{I},\ell}(\tau^\ell,a^\ell,b^\ell,\omega^\ell) \to \hat{v}^{\mathfrak{I}}(\tau',a,b,\omega)$
2. $\hat{v}^{\mathfrak{C},\ell}(\tau^\ell,a^\ell,b^\ell,\omega^\ell) \to \hat{v}^{\mathfrak{C}}(\tau',a,b,\omega)$
3. $\hat{w}^{\mathfrak{I},\ell}(\tau^\ell,a^\ell,b^\ell,\omega^\ell) \to \hat{w}^{\mathfrak{I}}(\tau',a,b,\omega)$
4. $\hat{w}^{\mathfrak{C},\ell}(\tau^\ell,a^\ell,b^\ell,\omega^\ell) \to \hat{w}^{\mathfrak{C}}(\tau',a,b,\omega)$.

We focus on the first convergence claim, as the others follow by similar arguments. We have shown that for each b and ω, there is a measure zero set $\Theta^{v,w}_{\tau'}(b,w)$ of preference shocks for type τ' politicians such that for all $\theta_{\tau'} \notin \Theta^{v,w}_{\tau'}(b,w)$, the mapping $\hat{\alpha}_{\tau'}$ is jointly continuous at $(b,\theta_{\tau'},\omega,v,w)$. Define $Z_{\tau'} \subseteq B\times X\times\Omega\times(\Re^d)^T\times\Re^T$ by

$$Z_{\tau'} \;=\; \bigcup_{(\tilde{b},x,\omega')}\Big[\{(\tilde{b},x,\omega'^d)^T \mid \theta_{\tau'}\in\Theta^{v,w}_{\tau'}(\hat{\beta}(x,\tilde{b},\omega,\omega'),\omega)\}\times\Re^T\Big],$$

which has measure zero by Fubini's theorem. Define $Z = \bigcup_{\tau'} Z_{\tau'}$, also measure zero. Now consider any $(\tilde{b},x,\omega',\theta,\epsilon)\notin Z$. Above continuity arguments establish that $\hat{\beta}^\ell(x,\tilde{b},\omega,\omega') \to \hat{\beta}(x,\tilde{b},\omega,\omega')$, and then it follows that $\hat{\alpha}^\ell_{\tau'}(\hat{\beta}^\ell(x,\tilde{b},\omega,\omega'),\theta_{\tau'},\omega') \to$

$\hat{\alpha}_{\tau'}(\hat{\beta}(x,\tilde{b},\omega,\omega'),\theta_{\tau'},\omega')$. Then uniform convergence of v_τ^ℓ to v_τ implies

$$\begin{aligned} & v_\tau^\ell(\tau^\ell,\hat{\alpha}_{\tau'}^\ell(\hat{\beta}^\ell(x,\tilde{b},\omega,\omega'),\theta_{\tau'},\omega'),\hat{\beta}^\ell(x,\tilde{b},\omega,\omega'),\omega') \\ & \quad \to \quad v_\tau(\tau',\hat{\alpha}_{\tau'}(\hat{\beta}(x,\tilde{b},\omega,\omega'),\theta_{\tau'},\omega'),\hat{\beta}(x,\tilde{b},\omega,\omega'),\omega') \end{aligned}$$

for all τ. By above arguments, we have $\hat{\varpi}^\ell(\hat{\beta}^\ell(x,\tilde{b},\omega,\omega'),\epsilon,\omega') \to \hat{\varpi}(\hat{\beta}(x,\tilde{b},\omega,\omega'),\epsilon,\omega')$ almost surely in ϵ. And by our continuity assumptions, $f(x|a^\ell,\omega^\ell) \to f(x|a,\omega)$ and $g(\tilde{b}|b,\omega^\ell) \to g(\tilde{b}|b,\omega)$. Thus, the integrand on the righthand side of $\hat{v}_\tau^{\mathfrak{I},\ell}(a^\ell,b^\ell)$ converges pointwise almost everywhere to

$$\begin{aligned} & \hat{\varpi}(\hat{\beta}(x,\tilde{b},\omega,\omega'),\epsilon,\omega')\bigg[\int_{x'}(u_\tau(x',\omega')+\epsilon_\tau)f_{\tau'}(x'|\hat{\alpha}_{\tau'}(\hat{\beta}(x,\tilde{b},\omega,\omega'),\theta_{\tau'},\omega'),\omega')dx' \\ & \quad +\delta_\tau(\omega')v_\tau(\tau',\hat{\alpha}_{\tau'}^\ell(\hat{\beta}(x,\tilde{b},\omega,\omega'),\theta_{\tau'},\omega'),\hat{\beta}(x,\tilde{b},\omega,\omega'),\omega')\bigg] \\ & \quad \cdot h(\theta,\epsilon|\omega')f_{\tau'}(x|a,\omega)g_{\tau'}(\tilde{b}|b,\omega) \end{aligned}$$

Thus, by Lebesgue's dominated convergence theorem, we have $\hat{v}_\tau^{\mathfrak{I},\ell}(\tau^\ell,a^\ell,b^\ell,\omega^\ell) \to \hat{v}_\tau^{\mathfrak{I}}(\tau',a,b,\omega)$ for each τ, as required.

5 Conclusions

Further work: What do equilibria look like? Multiple incumbents play a stage game to determine stage action.

References

Aguirregabiria, V. and P. Mira (2007) "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 75: 1–53.

Aragonès, E., T. Palfrey, and A. Postlewaite (2007) "Political Reputations and Campaign Promises," *Journal of the European Economic Association*, 5: 846–884.

Ashworth, S. (2005) "Repuational Dynamics and Political Careers," *Journal of Law, Economics, and Organization*, 21: 441–466.

Austen-Smith, D. and J. Banks (1989) "Electoral Accountability and Incumbency," in P. Ordeshook, ed., "Models of Strategic Choice in Politics," University of Michigan Press.

Banks, J. and J. Duggan (2008a) "A Dynamic Model of Democratic Elections in Multidimensional Policy Spaces," *W. Allen Wallis Institute Working Paper Series*, 53.

Banks, J. and J. Duggan (2008b) "A Dynamic Model of Democratic Elections in Multidimensional Policy Spaces," *Quarterly Journal of Political Science*, forthcoming.

Banks, J. and R. Sundaram (1993) "Adverse Selection and Moral Hazard in a Repeated Elections Model," in W. Barnett, M. Hinich, and N. Schofield, eds., "Political Economy: Institutions, Competition, and Representation," Cambridge, Cambridge University Press.

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.

Bernhardt, D., L. Campuzano, and F. Squintani (2005) "On Benefits of Party Competition," *Unpublished manuscript.*

Bernhardt, D., E. Hughson, and S. Dubey (2004) "Term Limits and Pork Barrel Politics," *Journal of Public Economics*, 88: 2383–2422.

Bernheim, D. and D. Ray (1989) "Markov Perfect Equilibria in Altruistic Growth Economies with Production Uncertainty," *Journal of Economic Theory*, 47: 195–202.

Casamatta, G. and C. De Paoli (2007) "Inefficient Public Good Provision in a Repeated Elections Model," *Journal of Public Economic Theory*, 9: 1103–1126.

Coate, S. and S. Morris (1995) "On the Form of Transfers to Special Interests," *Journal of Political Economy*, 103: 1210–1235.

Doraszelski, U. and M. Satterthwaite (2007) "Computable Markov-perfect Industry Dynamics: Existence, Purification, and Multiplicity," *Unpublished manuscript.*

Duggan, J. (2000) "Repeated Elections with Asymmetric Information," *Economics and Politics*, 12: 109–136.

Duggan, J. and T. Kalandrakis (2007) "Dynamic Legislative Bargaining," *W. Allen Wallis Institute Working Paper Series*, 45.

Ferejohn, J. (1986) "Incumbent Performance and Electoral Control," *Public Choice*, 50: 5–25.

Holmström, B. (1999) "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66: 169–182.

Kalandrakis, T. (2006) "A Reputational Theory of Two Party Competition," *W. Allen Wallis Institute Working Paper Series*, 41.

Kang, I. (2005) "Strategic Challengers and the Incumbency Advantage," *Unpublished manuscript.*

Le Borgne, E. and B. Lockwood (2001) "Candidate Entry, Screening, and the Political Budget Cycle," *Warwick Economic Research Papers*, 582.

Martinez, L. (2008a) "Repuation, Career Concerns, and Job Assignments," *Unpublished manuscript.*

Martinez, L. (2008b) "A Theory of Political Cycles," *Unpublished manuscript.*

Meirowitz, A. (2007) "Probabilistic Voting and Accountability in Elections with Uncertain Policy Constraints," *Journal of Public Economic Theory*, 9: 41–68.

Nowak, A. (2006) "On Perfect Equilibria in Stochastic Models of Growth with Intergenerational Altruism," *Economic Theory*, 28: 73–83.

Reed, R. (1994) "A Retrospective Voting Model with Heterogeneous Preferences," *Economics and Politics*, 6(39–58).

Rogoff, K. (1990) "Equilibrium Political Budget Cycles," *American Economic Review*, 80: 21–37.