

A Folk Theorem for Repeated Elections with Adverse Selection

John Duggan
Dept. of Political Science
and Dept. of Economics

February 14, 2013

Abstract

I establish a folk theorem for a model of repeated elections with adverse selection: when citizens are sufficiently patient, arbitrary policy paths through arbitrarily large regions of the policy space can be supported by a refinement of perfect Bayesian equilibrium. Politicians are policy-motivated (so office benefits cannot be used to incentivize policy choices), the policy space is one-dimensional (limiting the dimensionality of the set of utility imputations), and politicians' preferences are private information (so punishments cannot be targeted to a specific type). The equilibrium construction relies critically on differentiability and strict concavity of citizens' utility functions.

1 Introduction

Dynamic models of elections have increasingly become an important focus of political economy and formal political science. An advantage of these models over static models is that they allow, in principle, the possibility of policy dynamics. They enable analysis of the interplay between dynamic incentives and information on the part of voters and politicians, and they permit the analysis of essentially dynamic issues such as incumbency advantage, the consequences of term limits, and the effects of re-election incentives on policy choices. And to the extent that equilibrium policies resemble the ideal policy of the median voter, these models can provide non-cooperative underpinnings for the classical median voter theorem. This paper considers a framework for repeated elections, analyzed in Duggan (2000), in which politicians' preferences are private information, and voters update their prior beliefs about an office holder's type on the basis of her past policy choices; thus, the model is one of pure adverse selection. As a consequence, an office holder's policy choice reflects her short-term preference to implement desirable policy for the current period and long-term considerations of the impact on her chances for re-election; and each voter compares a relatively familiar incumbent with the prospect of an untried challenger.

The subject of analysis has naturally been on equilibria that have a stationary structure: the policy choice of a politician is a constant function of her type, and voters condition only on an office holder's current policy choice. This structure often implies that equilibria have a partitional form, with centrist politicians choosing their ideal policies and being re-elected, extremist politicians choosing their ideal policies and being removed from office, and moderate politicians pooling at a policy close to the median in order to gain re-election. The analysis of these equilibria reveals the possibility of incumbency advantage, as voter risk aversion creates some scope for a centrist politician to gain re-election by choosing a policy close (but not equal) to the median, and moderate politicians compromise in order to gain re-election. Banks and Duggan (2008) show that when citizens become arbitrarily patient, politicians almost always compromise by choosing policies arbitrarily close to the median ideal policy, delivering a version of the median voter theorem based on a model of dynamic elections with private information. Adding a term limit, an office holder simply chooses her ideal policy in her last term of office, and the partition in general depends on the incumbent's tenure in office. See Duggan (2000) and Banks and Duggan (2008) for analysis of the model without term limits, and see Bernhardt et al. (2004) for the model with term limits. Further applications include the analysis of parties (Bernhardt et al. (2009)), valence (Bernhardt et al. (2011)), and taxation (Camara (2012)).

Although the focus on stationary equilibria is a natural starting point, and can be viewed as an identifying assumption to generate testable implications of the model, it is important from a theoretical perspective to understand the restrictions on behavior implicit in this refinement. To this end, I formulate a more permissive equilibrium concept, called "perfect Bayesian electoral equilibrium," that allows politicians and voters to condition on past policy and election outcomes. I establish a folk theorem for the electoral model: when politicians and voters are sufficiently patient, arbitrary paths through an arbitrarily large region of the policy space can be supported as outcomes of perfect Bayesian electoral equilibria. This result does not follow from extant results on the folk theorem in the context of repeated games. For one reason, the game is not a repeated game, as voters alternate moves with politicians, and the identity of the office holder may change over time. More importantly, a politician's type is private information, so voters and future politicians cannot condition their choices on the current office holder's type; thus, the equilibrium concept must specify beliefs as well as strategies, and the appropriate concept should be consistent with the spirit of perfect Bayesian equilibrium. In fact, among other things, the refinement I propose excludes the possibility of targeting individual voters for their votes, and it imposes a stage-game dominance condition on voting strategies, so a voter votes for the incumbent if the expected discounted payoff from re-electing the incumbent is at least equal to that from electing an untried challenger.

Several more comments on the strength of this result are in order.

- I assume politicians are purely policy-motivated. This assumption means that a politician cannot be punished simply by removing her from office, and it significantly complicates the problem of incentivizing desired policy

choices.

- I assume the policy space is the unit interval. This assumption precludes the use of punishments involving arbitrarily bad policy choices, and it restricts the dimensionality of the space of payoff imputations and precludes the possibility of targeting punishments to particular office holder types—even if the politicians’ types were observed.
- Reinforcing the latter remark, because voters cannot observe the type of an office holder, punishment strategies in principle cannot be type-specific, so they must “thread the needle” by responding to deviations in a way that avoids rewarding one type while punishing another.
- Nevertheless, the folk theorem establishes that arbitrary policy paths can be supported in a strong sense: in each period, all types of office holder choose the same pre-specified policy.
- I do not exploit impatience of politicians by using an initial finite number of periods to elicit an office holder’s type; instead, I construct equilibria that generate arbitrary histories from the beginning of the game.

The equilibrium construction is quite delicate. For a policy path that takes a constant value near one extreme of the policy space, the construction must induce an office holder from the other side of the ideological spectrum to choose that policy (which is arbitrarily close to the worst possible) over an infinite horizon. And paths that are not constant are supported by punishments that “splice” these extremal equilibria together and alternate between the two extremes with appropriate frequency. The construction relies critically on differentiability of utility functions to address the former issue, and on strict concavity of utilities to address the second, in contrast to other folk theorem results that do not require such fine structure.

The conclusion of the analysis is that the restriction to stationary equilibria in existing work on models of repeated elections with adverse selection is substantial. Without some limitation on the extent of conditioning on past policy choices and electoral outcomes (or an a priori bound on the rate of discount), almost any policy path is consistent with equilibrium, and the model loses all predictive power. This is reminiscent of the cycling results of Plott (1967), McKelvey (1976, 1979), and Schofield (1978, 1983) in the multidimensional model with an empty majority core, but the current analysis holds in one dimension, and it is based on a non-cooperative equilibrium analysis that explicitly addresses the strategic incentives of politicians and voters. Closest to the current paper is Duggan and Fey (2006), who analyze repeated Downsian elections, in which two office-motivated parties compete in an infinite sequence of elections. In each period, the two parties simultaneously announce policy platforms, and the winning party is bound to its promise. That paper considers the general multidimensional model, and Theorem 1 of the paper establishes that when voters are relatively patient (discount factors greater than one half), every policy path can be generated by a subgame perfect equilibrium. Whereas

the latter result holds generally, Duggan and Fey's (2006) Theorem 2 applies when the majority core is nonempty and is perhaps most comparable to the folk theorem of this paper: it establishes that when parties are relatively patient (and voters' discount factors are positive), every policy path is generated by a subgame perfect equilibrium.

Of course, folk theorems have been proved for abstract dynamic games. Fudenberg and Maskin (1986) state the folk theorem for subgame perfect equilibria of discounted repeated games, and this is extended to repeated games with imperfect public monitoring by Fudenberg et al. (1994). Because voters and office holders alternate moves, however, the repeated elections model is not the simple repetition of a fixed stage game. Wen (2002) proves a folk theorem in repeated "sequential games," which assumes that in each period, players play an extensive form game such that players are partitioned into groups, groups move sequentially, players within a group choose simultaneously, and their feasible action sets are independent of others' choices. The timing of moves in this framework has some similarity to the electoral model, with an office holder first choosing policy and then voters simultaneously casting ballots, but it differs in the respect that the electoral model assumes a countably infinite set of candidates, and the identity of the office holder can conceivably change every period. Dutta (1995) establishes a folk theorem for stochastic games, which generalize sequential games, and Horner et al. (2011) and Fudenberg and Yamamoto (2011) extend the result to games with imperfect public monitoring, but again this work assumes a finite number of players. In addition, the above papers assume finite action sets, whereas the electoral model allows office holders to choose from a convex policy space. More importantly, the above papers assume that payoff functions are complete information, precluding the possibility of adverse selection and the difficult problems it entails.

As perfect Bayesian equilibria refine the standard concept of perfect Bayesian equilibrium, it stands that voter beliefs are updated via Bayes rule along the path of play, and the construction I describe in fact delivers more: even off the path of play from the initial history, voter beliefs about the type of a challenger are given by a fixed, common prior, and voters apply Bayes rule separately along the personal path of play of every office holder. And along the path of play from the initial history, all politician types choose the same policy (which may vary across periods), the first office holder is re-elected in every period, and voters do not revise their beliefs. But I do exploit some flexibility in specifying beliefs off the path of play: following an out of equilibrium policy choice, I specify that voters believe that with probability one the politician is an extreme type. An additional refinement that further restricts beliefs might preclude the construction provided here and entail some limits on policies chosen in equilibrium. I leave open the question of the effect of stronger refinements on policies supported by equilibrium strategies and beliefs.

The model analyzed in this paper is one of pure adverse selection. An alternative form of incomplete information is moral hazard, where voters do not observe the action of an office holder but only some noisy signal of the action; see Ferejohn (1986) for a model of pure moral hazard and Banks and Sundaram

(1993, 1998) for models that combine adverse selection and moral hazard. In such a model, punishments are necessarily conditioned on the noisy signal and equilibria take a much different form, and the possibilities for a folk theorem result would appear limited. The characterization of all perfect Bayesian equilibria (or bounds on equilibria) in moral hazard models is one that deserves attention in future research but is outside the scope of the current paper.

2 Electoral Model

Let $X = [0, 1]$ denote the one-dimensional policy space, let N denote the set of voters, and let M denote the countably infinite set of potential candidates. It may be that $M \subseteq N$, in which case N must be infinite, but this assumption is not needed for the analysis; in general, the electorate N may be finite, countably infinite, or a continuum. The set $M \cup N$ of citizens is partitioned into a finite set T of types, denoted $\tau_0, \tau_1, \dots, \tau_n$. Let $q = (q_1, \dots, q_n)$ denote the proportions of types within the set of voters. Note that the policy space is assumed to be compact, which precludes the possibility of equilibria using punishment strategies that involve policy choices at arbitrarily large (and increasing) distances from the center of the space.

The policy preferences of a type τ citizen are represented by a utility function $u_\tau: X \rightarrow \mathfrak{R}$ that is continuous and strictly concave with unique maximizer x_τ , the ideal policy of the type τ citizen. Assume without loss of generality that ideal policies are ordered by type indices, i.e., $x_{\tau_0} \leq x_{\tau_1} \leq \dots \leq x_{\tau_n}$. For simplicity, assume that the extreme types are located at the extreme points of the policy space and are uniquely defined by their ideal policies, i.e., $0 = x_{\tau_0} < x_{\tau_1}$ and $x_{\tau_{n-1}} < x_{\tau_n} = 1$. Assume that the median voter type is not extremist, i.e., $\max\{q_0, q_n\} < \frac{1}{2}$. Finally, assume that the utility functions of extreme types are differentiable at their ideal policies, i.e., u_{τ_0} is differentiable at $x = 0$ and u_{τ_n} is differentiable at $x = 1$, and that the first order condition holds at these ideal policies, i.e., $u'_{\tau_0}(0) = u'_{\tau_n}(1) = 0$.

Elections proceed as follows. Denumerating politicians as $M = \{c_1, c_2, \dots\}$, politician c_1 holds a public office and selects a policy $x_1 \in X$ in the first period. In each period $t = 2, 3, \dots$, a member of M is chosen to fill the office as follows. In period t , a politician i_t is the incumbent office holder, and politician c_t runs as challenger in a majority rule election. Voters simultaneously select ballots from $\{i_t, c_t\}$, and if the proportion of voters who vote for i_t is at least one half, then the incumbent wins the election, selects a policy, and becomes the incumbent in period $t + 1$, i.e., $i_{t+1} = i_t$. Otherwise, the challenger wins the election, selects a policy, and becomes the incumbent in period $t + 1$, i.e., $i_{t+1} = c_t$.

I assume that a politician's type is private information and is not observed by voters, whereas past policy choices are observable. Thus, the environment is one of pure adverse selection, and in equilibrium voters must condition their beliefs about an office holder's type based on her past performance in office. Assume that voters' prior beliefs are that politician types are independently and identically distributed according to the distribution $p = (p_0, \dots, p_n)$. Thus,

in the period t election, when voters compare the incumbent i_t to the challenger c_t , they condition their beliefs about i_t on past policy choices, and they believe c_t is type τ_k with probability p_k . Assume that the extreme politician types are possible, i.e., $\min\{p_0, p_n\} > 0$. Note that it may be that $p \neq q$, so the distribution of politician types may differ from the distribution of types in the electorate; and in particular, it may be that $p_0 + p_n = 1$, so that policies are always chosen by extreme liberal or conservative politicians.

Let $\mathbf{x} = (x_1, x_2, \dots) \in [0, 1]^\infty$ denote a path of policies over time. Assume that citizens' preferences over policy paths are given by the discounted sum of utility,

$$\sum_{t=1}^{\infty} \delta^{t-1} u_\tau(x_t),$$

where $\delta \in [0, 1)$ is a common discount factor. Furthermore, assume these preferences extend to distributions over policy paths via expected utility. Note that politicians are purely policy-motivated and do not receive a benefit from holding office, consistent with the citizen-candidate approach to elections (Osborne and Slivinski (1996), Besley and Coate (1997)). The problem of supporting policy paths in equilibrium with patient players would be much simpler if politicians' preferences include a fixed, positive office benefit term, so the current paper addresses the more difficult case of pure policy motivation.

A *complete public history of length t* describes the publicly observed events in the first t periods, namely, the vote tallies from previous elections and the policies selected by office holders. A *partial public history of length t* is a complete history of length $t - 1$ together with the policy choice x_t of the current office holder. The *initial history*, denoted \emptyset , describes the electoral game at the beginning of period 1. A *policy strategy* for politician c_t is a mapping that assigns a policy choice to every complete history of length $t' \geq t + 1$ such that c_t is the office holder in period t' , i.e., $c_t = i_{t+1} = \dots = i_{t'}$. A *voting strategy* for a voter assigns a ballot, i_t or c_t , to every partial public history of length t .¹ A strategy profile, which specifies a policy strategy for every politician and a voting strategy for every voter, determines a distribution over finite (complete and partial) public histories, and therefore over infinite policy paths, beginning at the initial history. Because this is a dynamic game of incomplete information, a description of strategically relevant details must include not only the strategies used by all citizens but also a *belief assessment* that specifies the beliefs of all citizens at each finite (complete or partial) history. Given such a belief assessment and an arbitrary public history, a distribution over longer public histories, and therefore over paths of policies in future periods, is determined. Note that I consider only pure strategy equilibria in this paper.

The most permissive equilibrium concept consistent with dynamic incentives of citizens is that of *perfect Bayesian equilibrium*. This is a specification of pol-

¹In case N is infinite, we generally must add the assumption of measurability, so that the proportion of voter types casting a given ballot can be calculated. Since I will later restrict attention to type-symmetric strategies, this technical issue is moot.

icy strategies, voting strategies, and a belief assessment for all citizens such that (i) for every finite complete public history, the policy choice of the office holder maximizes her expected discounted payoff, (ii) for every finite partial public history, the ballot of each voter maximizes her expected discounted payoff, and (iii) for every finite (complete or partial) public history on the path of play from the initial node, citizens' beliefs are determined by Bayes rule given the policy strategies used by office holders and the private information (if any) held by the citizen.

In fact, the proof of the folk theorem employs an equilibrium construction that satisfies a number of additional refinements that preclude equilibrium behavior that appears especially untenable.

- (a) *outcome measurability*: politicians and voters condition their actions only on the history of previous policy choices and election winners.
- (b) *Bayesian independence*: at a finite partial public history, voter beliefs about an untried challenger are given by the prior, and voters update their prior beliefs about an office holder's type conditioning only on observed policy choices of that politician.
- (c) *Bayesian consistency*: at a finite partial public history that is off an office holder's personal path of play, and given voters' beliefs at that history, voters update these beliefs on the basis of future policy choices using Bayes rule when possible.
- (d) *common beliefs*: beliefs about an office holder's type off her personal path of play are the same for all voters.
- (e) *type symmetry*: for every partial history, all voters of the same type cast the same ballot.
- (f) *weak dominance*: given a partial history at which the expected discounted payoff of a voter from re-electing the incumbent is not equal to that from electing an untried challenger, the voter votes for the candidate offering the higher payoff.
- (g) *deferential voting*: given a partial history at which the expected discounted payoff of a voter from re-electing the incumbent equals that from electing an untried challenger, the voter votes for the incumbent.

I refer to a perfect Bayesian equilibrium satisfying the above conditions as a *perfect Bayesian electoral equilibrium*.

Some discussion of these restrictions may be warranted. I refer to the policy choice and election winner in period t as the "electoral outcome" in the period, and condition (a) requires that politicians and voters condition only on past electoral outcomes. This means, in particular, that a single voter (or voter type) cannot be targeted for punishments for voting the wrong way at some partial history. Conditions (b) and (c) impose further restrictions on the specification of voter beliefs off the equilibrium path from the initial history. Condition (b) means that even at a history such that a deviation has occurred (either an out of

equilibrium policy choice or election winner), voters' beliefs about a challenger are always given by the prior. I term the policy choices that an office holder has taken while in office the politician's "personal history," and if these choices are consistent with the politician's strategy for some type with positive prior probability in every period along the personal history, then it is on the politician's "personal path of play." Condition (b) furthermore implies that once a politician is elected, voters condition their beliefs only on the politician's personal history when it is on her path of play. In particular, if an office holder chooses a policy consistent with voter beliefs, and if voters' strategies determine that the incumbent is removed from office, but the incumbent is instead re-elected (contrary to the specification of voting strategies), then voters still update their beliefs about the office holder's type according to Bayes rule. Condition (c) means that at a personal history that is off an office holder's personal path of play, voter beliefs are thereafter updated via Bayes rule (if possible) from that point on. In particular, once voter beliefs are specified at a partial history that is off an office holder's personal path of play, if the office holder is re-elected and then chooses policy that is consistent with those beliefs, then voter beliefs are updated via Bayes rule. Condition (d) is self-explanatory, and it reinforces condition (e), which relies on the fact that in equilibrium, calculations of expected payoffs are the same for voters of the same type. Condition (f) is a standard restriction in voting games, requiring that each voter vote as though she were pivotal in the election. When the electorate is finite, this can be formulated as a stage-game dominance refinement, and it simply precludes implausible voting behavior due to no one voter being pivotal. When the electorate is a continuum, each voter has mass zero and can never affect the election, and (f) is a standard condition imposed to generate plausible voting behavior. Finally, condition (g) requires that all voters vote for the incumbent when indifferent, a condition that typically arises from the common stationarity refinement used in the literature.

3 Main Theorem

In this section, I establish a folk theorem for perfect Bayesian electoral equilibria. Because the proof is rather involved, I give a detailed and less formal discussion of the approach following the statement of the theorem, deferring the formal proof until the next section. Let $\mathbf{x} = (x_1, x_2, \dots) \in [0, 1]^\infty$ denote a path of policies over time, and for future use, given policy $a \in X$, let $\mathbf{x}^a = (a, a, \dots)$ denote the path that is constant at a . Given discount factor $\delta < 1$, let $P(\delta)$ consist of every path that is supportable by a perfect Bayesian electoral equilibrium, i.e., \mathbf{x} belongs to $P(\delta)$ if and only if there is a perfect Bayesian electoral equilibrium such that along the path of play from the initial history, in every period t , all types of office holder choose policy x_t .

Theorem: *For all $a, b \in X$, with $0 < a < b < 1$, there exists $\bar{\delta} < 1$ such that for all $\delta \in (\bar{\delta}, 1)$, we have $[a, b]^\infty \subseteq P(\delta)$.*

The proof consists of four steps. Steps 1 and 2 exploit differentiability to

construct equilibria that support policy paths that are constant at a policy a near zero and constant at a policy b near one; the steps are symmetric, so I focus on the former. In Step 3, I define a “splicing” procedure that allows me to construct equilibria that generate policy paths that alternate between a and b with essentially arbitrary frequency. Then Step 4 exploits strict concavity and uses these alternating equilibria to punish deviations from arbitrary policy paths.

In Step 1, the path \mathbf{x}^a is supported using a punishment strategy that replaces a deviating incumbent with an untried challenger and then generates a policy path that depends on the newly elected office holder’s type: if she is type $\tau \neq \tau_0$, it generates a string of m choices of $x = 0$ followed by m choices of $x = x_{\tau_1}$; and if she is type τ_0 , it generates a string of $2m$ choices of $x = 0$. Here, m is a carefully chosen natural number that will be appropriately large, and as such policy choices beyond $2m$ periods have a second-order effect. The problem is to specify a and m so that office holders of all types are incentivized to choose a along the path of play, rather than trigger the punishments. Punishing a deviation by moving to a centrally located policy would deter type τ_0 politicians from deviating, but it creates opportunities for manipulation by type $\tau \neq \tau_0$ politicians, who could prefer the punishment to remaining at a . Punishments that involve policies $x = 0$ must be used, but in a nuanced way that continues to deter type τ_0 politicians.

To gain some geometric insight into the problem, consider the politicians’ payoffs from a string of length m of policy x followed by a string of length m of policy y , namely,

$$\begin{aligned} & (1 - \delta)[u_\tau(x) + \dots + \delta^{m-1}u_\tau(x) + \delta^m u_\tau(y) + \dots + \delta^{2m-1}u_\tau(y)] \\ &= (1 - \delta^m)u_\tau(x) + \delta^m(1 - \delta^m)u_\tau(y). \end{aligned}$$

In Figure 1, I depict indifference curves over sequences

$$\underbrace{(x, x, \dots, x)}_{m \text{ periods}}, \underbrace{(y, y, \dots, y)}_{m \text{ periods}},$$

with the horizontal axis representing x and the vertical representing y , for a type τ_0 citizen and a type $\tau \neq \tau_0$ citizen. Note that the indifference curve of the type τ_0 citizen is flat at $(0, x_{\tau_1})$, reflecting the fact that $u'_{\tau_0}(0) = 0$, and the indifference curve of the type $\tau \neq \tau_0$ has finite slope at $(0, x_{\tau_1})$, reflecting the fact that $u'_\tau(0) > 0$.

Given the above description of punishments, any deviation by an office holder is followed by a lottery over strings

$$\underbrace{(0, 0, \dots, 0)}_{m \text{ periods}}, \underbrace{(x_{\tau_1}, x_{\tau_1}, \dots, x_{\tau_1})}_{m \text{ periods}} \quad \text{and} \quad \underbrace{(0, 0, \dots, 0)}_{m \text{ periods}}, \underbrace{(0, 0, \dots, 0)}_{m \text{ periods}}$$

with probability p_0 on the latter and $1 - p_0$ on the former. These strings are indicated in Figure 1, and the indifference curves corresponding to the expected payoff from this punishment lottery are indicated by the heavy level sets in the




Figure 1: Indifference curves over sequences




Figure 2: Lens of incentives

figure for type τ_0 and type $\tau \neq \tau_0$ politicians. To reconcile the incentives of office holders, the policy a must belong to the lens formed by the indifference curves of the type τ_0 and all $\tau \neq \tau_0$ politicians. In Figure 2, we see that for the given setting of m , there is no policy a belonging to the lens of incentives that is also within ϵ of zero.

Note that the marginal rate of substitution between policy in the first m periods and in the next m periods is

$$\frac{(1 - \delta^m)u'_\tau(0)}{\delta^m(1 - \delta^m)u'_\tau(x_{\tau_1})} = \frac{u'_\tau(0)}{\delta^m u'_\tau(x_{\tau_1})}.$$

Choosing m higher, so δ^m becomes smaller, citizens put greater weight on the first m periods than the next m , making the indifference curves in Figure 1 steeper. The bulk of Step 1 consists of showing that when a is small, we can




Figure 3: Choosing a and δ^m

choose δ^m in a critical range such that (a, a) is preferred to the punishment lottery by all citizen types; geometrically, the lens of incentives intersects the 45 degree line sufficiently close to the origin, as depicted in Figure 3. Once this is established, we can perform the construction for discount factors close to one: when δ is close to one, we can choose m so that δ^m belongs to the critical range. This entails that the punishment lottery is worse than the constant policy a for $2m$ periods for all types, and in addition we then specify that δ is close enough to one that the potential gain from any one-shot deviation is more than offset by the cost of entering the punishment phase.

The punishment strategy used in Step 1 is in fact somewhat more complex than the simple description above, because we must specify policies beyond the $2m$ periods following a deviation, and because we must address the possibility of deviations from the punishment strategies. Of particular importance, we must ensure that voters have incentives to re-elect the incumbent if, $m + 1$ periods following a deviation, she chooses $x = 0$. This choice reveals that the politician is type τ_0 , and voters then update that the office holder will continue to choose $x = 0$ in the next $m - 1$ periods if re-elected. The alternative is to elect a challenger and restart the punishment phase, but by construction of the punishment strategies, the challenger, if elected, will also choose $x = 0$ for m periods, regardless of her type. A disadvantage of electing a challenger to a type $\tau \neq \tau_0$ voter is that with probability p_0 , the challenger is type τ_0 and will choose $x = 0$ for an additional m periods. The advantage of electing a challenger to such a voter is that with probability $1 - p_0$, the politician may choose preferable policies $m + 1$ periods hence. Figure 4 depicts the distribution over policy paths following a deviation. We specify that after the first $2m$ periods following a deviation, a type $\tau \neq \tau_0$ office holder choose $x = x_{\tau_1}$, while a type τ_0 office holder alternately choose $x = x_{\tau_1}$ for k periods and $x = 0$ for one period. When k is large, there-




Figure 4: Punishment paths

fore, the benefit of electing a challenger to a type $\tau \neq \tau_0$ voter is outweighed by the cost, and a majority of voters indeed prefer to re-elect the incumbent.

Step 3 provides a procedure for splicing equilibria that support the extreme policies from Steps 1 and 2. For example, consider two equilibria supporting \mathbf{x}^a and \mathbf{x}^b . To construct an equilibrium supporting the path (a, b, a, b, \dots) , I specify, roughly, that the equilibrium construction from Step 1 be applied to histories consisting of odd periods, and that the construction from Step 2 be applied to histories consisting of even periods. Thus, a deviation in an odd period t is punished in future odd periods; and moreover, in future even periods, we simulate punishments as though an analogous deviation had occurred in the previous even period $t - 1$. If an office holder chooses $x \neq a$ in period three, for example, then play in future even periods proceeds as though she had chosen $x \neq b$ in period two, so the incumbent is removed from office, and we commence with the punishments in Figure 4 in all future even and odd periods. This means that the newly elected office holder chooses $x = 0$ for the next m odd periods and $x = 1$ for the next m even periods. After that, play depends on the politician's type: if she is type τ_0 , for example, then she chooses $x = x_{\tau_{n-1}}$ in all subsequent even periods; and she chooses $x = 0$ in an additional m odd periods, then alternately chooses $x = x_{\tau_1}$ for k odd periods and $x = 0$ for one odd period. Because punishments in odd periods now alternate with even periods, payoffs in odd periods are now effectively discounted by δ^2 (and similarly for payoffs in even periods), so the construction requires a discount factor greater than the cutoffs from Steps 1 and 2; if \mathbf{x}^a is supported for discount factors above δ' and \mathbf{x}^b is supported for discount factors about δ'' , then the splicing procedure requires a discount factor above $\sqrt{\max\{\delta', \delta''\}}$. Further details related to the transmission of information across odd and even periods arise as well. When an office holder chooses $x = x_{\tau_1}$ after m odd punishment periods, for example, she reveals that she is not type τ_0 , and the updated probability that she is type τ_n increases; this means that she is more likely to choose $x = 1$ for an additional m even periods, and we must confirm that a majority of voters still have an incentive to re-elect the incumbent. These details are straightforward and are addressed in the proof.

With the splicing procedure defined, it can be used recursively to generate policy paths that alternate between a and b at nearly arbitrary frequencies. For example, we can splice equilibria supporting (a, a, a, a, \dots) and (a, b, a, b, \dots) to

produce a perfect Bayesian electoral equilibrium that supports policy choice a with frequency .75.

$$\begin{array}{cccccccc} a & a & a & a & a & a & a & a \dots \\ a & b & a & b & a & b & a & b \dots \\ \hline a & a & a & b & a & a & a & b \dots \end{array}$$

And we can splice equilibria supporting a with frequency .5 and .75 to produce an equilibrium supporting a with frequency .625.

$$\begin{array}{cccccccccccccccc} a & b & a & b & a & b & a & b & a & b & a & b & a & b & a & b \dots \\ a & a & a & b & a & a & a & b & a & a & a & b & a & a & a & b \dots \\ \hline a & a & a & b & a & a & b & b & a & a & a & b & a & a & b & b \dots \end{array}$$

Thus, using this logic, we can apply the splicing procedure a finite number of times to approximate an arbitrary frequency to within any given level of precision. Each application imposes a higher cutoff for the discount factor, but because only a finite number of iterations are required for a given level of precision, we can take the maximum required for each application of the procedure.

Step 4 uses the alternating equilibria from Step 3 to punish deviations from arbitrary paths of policies. Choose policies $a, b \in X$ as in the theorem. Then use Steps 1 and 2 to choose $a' \in (0, \frac{a}{2})$ and $b' \in (\frac{b+1}{2}, 1)$ such that for discount factors sufficiently close to one, there exist perfect Bayesian electoral equilibria supporting $\mathbf{x}^{a'}$ and $\mathbf{x}^{b'}$. Consider any path $\mathbf{x} \in [a, b]^\infty$ in the interval $[a, b]$. Note that for a type τ citizen, the discounted utility from the path starting in period $t + 1$, i.e., $(x_{t+1}, x_{t+2}, \dots)$, can be written as the integral

$$E_{\lambda_{t+1}}[u_\tau(x)] = \int u_\tau(x) \lambda_{t+1}(dx) = (1 - \delta) \sum_{t'=t+1}^{\infty} \delta^{t'-t-1} u_\tau(x_{t'})$$

with respect to a probability measure λ_{t+1} on $[a, b]$, where I normalize payoffs by $1 - \delta$. In the space of utility imputations, depicted in Figure 5 for the case of $n = 1$, λ_{t+1} induces a probability measure on the thick portion of the utility frontier, and the vector of expected utilities lies in the shaded region below the frontier. Moreover, each x in the support of λ_{t+1} can be written as a strict convex combination of a' and b' with weight $\alpha(x)$ on a' and $1 - \alpha(x)$ on b' . By strict concavity, the convex combination

$$\alpha(x)u_\tau(a') + (1 - \alpha(x))u_\tau(b')$$

is less than $u_\tau(x)$ by an increment $\eta > 0$ that is uniform across $x \in [a, b]$. Thus,

$$\begin{aligned} E_{\lambda_{t+1}}[u_\tau(x)] &\geq \int [\alpha(x)u_\tau(a') + (1 - \alpha(x))u_\tau(b')] d\lambda_{t+1} \\ &= E_{\lambda_{t+1}}[\alpha(x)]u_\tau(a') + (1 - E_{\lambda_{t+1}}[\alpha(x)])u_\tau(b') + \eta. \end{aligned}$$

We can use the splicing procedure to approximate the weights $E_{\lambda_{t+1}}[\alpha(x)]$ and $1 - E_{\lambda_{t+1}}[\alpha(x)]$ by $\frac{k}{2^\ell}$ and $\frac{2^\ell - k}{2^\ell}$ so that the corresponding convex combination is uniformly $\frac{\eta}{2}$ less than λ_{t+1} . Choosing δ sufficiently close to one,




Figure 5: Extreme punishments

the maximum gain (normalized by $1 - \delta$) from a one-shot deviation in period t is less than $\frac{\delta\eta}{2(1-\delta)}$, and it follows that we can specify punishment strategies to deter any deviation in period t . Finally, the above discussion has fixed a period $t + 1$ and the corresponding probability measure λ_{t+1} , but the set of probability measures with support in $[a, b]$ is compact in the weak* topology, and the approximation $\frac{k}{2^\ell}$ and $\frac{2^\ell - k}{2^\ell}$ holds for an open set around λ_{t+1} , so the open covering generated by choice of k and ℓ must have a finite subcover. Thus, we can choose δ sufficiently close to one such that for all t , Step 3 can be applied to specify an effective punishment strategy to deter a deviation in period t , delivering the required perfect Bayesian electoral equilibrium.

4 Proof of Theorem

The proof proceeds in four steps.

Step 1: For all $\epsilon > 0$, there exist $a \in (0, \epsilon)$ and $\hat{\delta} \in (0, 1)$ such that for all $\delta \in (\hat{\delta}, 1)$, we have $\mathbf{x}^a \in P(\delta)$.

Fix $\epsilon > 0$. I begin by choosing parameters $a, \delta, \beta \in (0, 1)$ with $a < \epsilon$ and a natural number $k \geq 1$ to satisfy three inequalities that are critical for the equilibrium construction, namely (5)–(7), below. Following these choices, I will specify $\hat{\delta}$ to satisfy three further inequalities, namely (12)–(14). First, choose k sufficiently large that for all $\tau \neq \tau_0$, we have

$$2(1 - p_0) \left[u_\tau(x_{\tau_1}) - \frac{ku_\tau(x_{\tau_1}) + u_\tau(0)}{k + 1} \right]$$

$$< p_0 \left[\frac{k u_\tau(x_{\tau_1}) + u_\tau(0)}{k+1} - u_\tau(0) \right]. \quad (1)$$

For each τ , define

$$W_\tau = \frac{1}{1-\delta^{k+1}} [(1-\delta^k)u_\tau(x_{\tau_1}) + \delta^k(1-\delta)u_\tau(0)],$$

which implicitly depends on k and δ . This is the discounted utility (normalized by $1-\delta$) of alternating sequences of $x = x_{\tau_1}$ for k periods and $x = 0$ for one period. Note that by l'Hospital's rule, we have for all τ ,

$$\lim_{\delta \uparrow 1} W_\tau = \frac{k u_\tau(x_{\tau_1}) + u_\tau(0)}{k+1}, \quad (2)$$

and in addition that

$$\lim_{\beta \downarrow 0} \frac{(1-p_0)\beta(1-\beta)}{(1-p_0)\beta + p_0\beta^2} = 1. \quad (3)$$

Using (1)–(3), there exists $\bar{\beta} \in (0, 1)$ such that for all $\beta \in (0, \bar{\beta})$, we have

$$\frac{(1-p_0)\beta(1-\beta)}{1+p_0(\beta-\beta^2)-\beta} > \frac{1}{2} \cdot \frac{(1-p_0)\beta + p_0\beta^2}{1+p_0(\beta-\beta^2)-\beta}, \quad (4)$$

and such that there exists $\underline{\delta} \in (0, 1)$ such that for all $\delta \in (\underline{\delta}, 1)$ and all $\beta \in (0, \bar{\beta})$, we have for all $\tau \neq \tau_0$,

$$\begin{aligned} & (1-p_0)[\beta(1-\beta)(u_\tau(x_{\tau_1}) - W_\tau) + \beta(u_\tau(x_{\tau_1}) - W_\tau)] \\ & < p_0\beta(1-\beta)(W_\tau - u_\tau(0)). \end{aligned} \quad (5)$$

Inequality (4) is used in the selection of a , below, and (5) is used to give voters proper incentives in the equilibrium construction.

Given $\delta \in (\underline{\delta}, 1)$ and $\beta \in (0, \bar{\beta})$, the next inequality is used to address incentives of type τ_0 office holders along the path of play:

$$\begin{aligned} & (1-\delta)(u_{\tau_0}(0) - u_{\tau_0}(a)) \\ & \leq \delta \left[p_0[(1-\beta^2)(u_{\tau_0}(a) - u_{\tau_0}(0)) + \beta^2(u_{\tau_0}(a) - W_{\tau_0})] \right. \\ & \quad \left. + (1-p_0)[(1-\beta)(u_{\tau_0}(a) - u_{\tau_0}(0)) + \beta(u_{\tau_0}(a) - u_{\tau_0}(x_{\tau_1}))] \right]. \end{aligned} \quad (6)$$

The following inequality addresses incentives for type $\tau \neq \tau_0$ office holders along the path of play: for all $\tau \neq \tau_0$,

$$\begin{aligned} & (1-\delta)\Delta \\ & \leq \delta \left[p_0[(1-\beta^2)(u_\tau(a) - u_\tau(0)) + \beta^2(u_\tau(a) - W_\tau)] \right. \\ & \quad \left. + (1-p_0)[(1-\beta)(u_\tau(a) - u_\tau(0)) + \beta(u_\tau(a) - u_\tau(x_{\tau_1}))] \right], \end{aligned} \quad (7)$$

where $\Delta = \max_{\tau, x, y} [u_\tau(x) - u_\tau(y)]$. To see that inequalities (6) and (7) are compatible for some $\delta \in (\underline{\delta}, 1)$ and some $\beta \in (0, \bar{\beta})$ and for sufficiently small $a \in (0, \epsilon)$, define

$$\begin{aligned} A_\tau &= u_\tau(a) - u_\tau(0) \\ A_\tau^* &= u_\tau(x_{\tau_1}) - u_\tau(a), \end{aligned}$$

where the latter two quantities implicitly depend on a . It will be convenient to define $\underline{A} = \min_{\tau \neq \tau_0} A_\tau$ and $A^* = \max_{\tau \neq \tau_0} A_\tau^*$. Of course, $\underline{A} \rightarrow 0$ and $A^* \rightarrow \max_{\tau \neq \tau_0} [u_\tau(x_{\tau_1}) - u_\tau(0)] > 0$ as $a \rightarrow 0$. Assuming without loss of generality that $a < x_{\tau_1}$, the right-hand side of (6) is positive if

$$\frac{(1-p_0)\beta}{1+p_0(\beta-\beta^2)-\beta} > \frac{u_{\tau_0}(0) - u_{\tau_0}(a)}{u_{\tau_0}(a) - u_{\tau_0}(x_{\tau_1})}, \quad (8)$$

and using $u_\tau(x_{\tau_1}) \geq W_\tau$, the right-hand side of (7) is positive if for all $\tau \neq \tau_0$,

$$\frac{u_\tau(a) - u_\tau(0)}{u_\tau(x_{\tau_1}) - u_\tau(a)} > \frac{(1-p_0)\beta + p_0\beta^2}{1+p_0(\beta-\beta^2)-\beta}. \quad (9)$$

Note that

$$\begin{aligned} & \lim_{a \downarrow 0} \left(\frac{u_{\tau_0}(0) - u_{\tau_0}(a)}{u_{\tau_0}(a) - u_{\tau_0}(x_{\tau_1})} \right) \left(\frac{A^*}{\underline{A}} \right) \\ & \leq \lim_{a \downarrow 0} \left(\frac{\frac{u_{\tau_0}(0) - u_{\tau_0}(a)}{a}}{u_{\tau_0}(a) - u_{\tau_0}(x_{\tau_1})} \right) \left(\frac{\max_{\tau \neq \tau_0} u_\tau(x_{\tau_1}) - u_\tau(a)}{\min_{\tau \neq \tau_0} \frac{u_\tau(a) - u_\tau(0)}{a}} \right) \\ & = \frac{u'_{\tau_0}(0)}{u_{\tau_0}(0) - u_{\tau_0}(x_{\tau_1})} \left(\frac{\max_{\tau \neq \tau_0} u_\tau(x_{\tau_1}) - u_\tau(0)}{\min_{\tau \neq \tau_0} u'_\tau(0)} \right) \\ & = 0. \end{aligned}$$

Therefore, we can choose $a \in (0, \epsilon)$ and $\beta \in (0, \bar{\beta})$ such that

$$\frac{(1-p_0)\beta}{1+p_0(\beta-\beta^2)-\beta} > \frac{u_{\tau_0}(0) - u_{\tau_0}(a)}{u_{\tau_0}(a) - u_{\tau_0}(x_{\tau_1})} > \frac{1}{2} \cdot \frac{(1-p_0)\beta}{1+p_0(\beta-\beta^2)-\beta}, \quad (10)$$

automatically fulfilling (8), and such that

$$\frac{u_{\tau_0}(0) - u_{\tau_0}(a)}{u_{\tau_0}(a) - u_{\tau_0}(x_{\tau_1})} < \frac{1}{4} \cdot \frac{u_\tau(a) - u_\tau(0)}{u_\tau(x_{\tau_1}) - u_\tau(a)}. \quad (11)$$

Combining (10) and (11) and using (4), we have

$$\begin{aligned} \frac{u_\tau(a) - u_\tau(0)}{u_\tau(x_{\tau_1}) - u_\tau(a)} &> 4 \left(\frac{u_{\tau_0}(0) - u_{\tau_0}(a)}{u_{\tau_0}(a) - u_{\tau_0}(x_{\tau_1})} \right) \\ &> 2 \left(\frac{(1-p_0)\beta(1-\beta)}{1+p_0(\beta-\beta^2)-\beta} \right) \\ &> \frac{(1-p_0)\beta + p_0\beta^2}{1+p_0(\beta-\beta^2)-\beta}, \end{aligned}$$

fulfilling (9), as required. It follows that for these choices of a and β , the right-hand sides of (6) and (7) are positive, and we can therefore choose $\delta \in (\underline{\delta}, 1)$ sufficiently close to one to satisfy (6) and (7).

Fixing a as above, note that there is some latitude in the choice of β , as (8) and (9) involve strict inequalities. Next, I exploit this flexibility to choose $\hat{\delta}$ close to one as in the statement of Step 1. Let $\beta_1, \beta_2 \in (0, \bar{\beta})$ be such that $\beta_1 < \beta_2$ and for all $\beta \in [\beta_1, \beta_2]$, inequalities (8) and (9) hold. Then choose $\hat{\delta} \in (\underline{\delta}, 1)$ sufficiently high that for all $\delta \in (\hat{\delta}, 1)$ and all $\beta \in [\beta_1, \beta_2]$, inequalities (6) and (7) hold, and so that we have for all $\tau \neq \tau_0$,

$$(1 - \delta)\Delta < \delta p_0[\beta(1 - \beta)(u_\tau(x_{\tau_1}) - u_\tau(0)) + \beta^2(u_\tau(x_{\tau_1}) - W_\tau)] \quad (12)$$

$$(1 - \delta)\Delta < \delta p_0(u_\tau(x_{\tau_1}) - W_\tau) \quad (13)$$

$$(1 - \delta)\Delta < \delta(1 - p_0)(W_{\tau_0} - u_{\tau_0}(x_{\tau_1})), \quad (14)$$

and so that $\hat{\delta}\beta_2 > \beta_1$. Inequalities (12) and (13) are used to address incentives of type $\tau \neq \tau_0$ office holders, and (14) to address type τ_0 office holders, in the punishment phases. Using the inequality $\hat{\delta}\beta_2 > \beta_1$, I claim that for each $\delta > \hat{\delta}$, there exists a natural number m such that $\delta^m \in (\beta_1, \beta_2)$. Indeed, we have $\delta > \beta_1$ by assumption, and $\delta^m < \beta$ for sufficiently high m . The claim holds with $m = 1$ if $\delta < \beta_2$. Otherwise, $\delta \geq \beta_2$, and we can let $m - 1$ be the smallest natural number such that $\delta^{m-1} \geq \beta_2$. Then we have $\delta^m < \beta_2$ and $\delta^m = \delta\delta^{m-1} \geq \delta\beta_2 > \beta_1$, i.e., $\delta^m \in (\beta_1, \beta_2)$, as claimed.

Next, given discount factor $\delta > \hat{\delta}$, I specify strategies and beliefs that support the constant path $\mathbf{x}^a = (a, a, \dots)$. Note that by the preceding argument, we can choose m so that (6) and (7) are satisfied by specifying $\beta = \delta^m \in (\beta_1, \beta_2)$. In the construction, I assume that voter beliefs about an untried challenger are given by the prior p , and that beliefs along an office holder's personal path of play are derived from Bayesian updating, and that off the politician's personal path of play, voters believe that with probability one the incumbent is type τ_0 . Strategies and beliefs are specified by assigning each history to one of five phases below, three of which are parameterized by a natural number j , and one of those is partitioned into three sub-phases, depending on voter beliefs. Phases 2–4 are used to punish deviations along the path of play, which takes place in Phase 1, and Phase 5 is used to punish deviations in the punishment phases.

Phase 1 All types of office holder choose $x = a$. If the incumbent chooses $x = a$, then all voters' beliefs are given by the prior p , and all voters vote for the incumbent; and if the incumbent chooses $x \neq a$, then all voters believe that with probability one the incumbent is type τ_0 , and all type $\tau \geq \tau_1$ voters vote for the challenger, and type τ_0 voters vote for the incumbent.

Phase 2(j) All types of office holder choose $x = 0$. If the incumbent chooses $x = 0$, then all voters' beliefs are given by the prior p , and all type $\tau \geq \tau_1$ voters vote for the incumbent and type τ_0 voters vote for the challenger; and if the incumbent chooses $x \neq 0$, then all voters believe that with

probability one the incumbent is type τ_0 , and all type $\tau \geq \tau_1$ voters vote for the challenger, and type τ_0 voters vote for the incumbent.

Phase 3(j) The type τ_0 office holder chooses $x = 0$, and all type $\tau \neq \tau_0$ office holders choose $x = x_{\tau_1}$. If the incumbent chooses $x = 0$, then all voters believe that with probability one the incumbent is type τ_0 , and all type $\tau \geq \tau_1$ voters vote for the incumbent, and type τ_0 voters vote for the challenger; if $x = x_{\tau_1}$, then all voters' beliefs are given by \hat{p} , the updated prior conditional on the politician's type being not equal to τ_0 ; and all type $\tau \geq \tau_1$ voters vote for the incumbent, and type τ_0 voters vote for the challenger; and otherwise, if $x \notin \{0, x_{\tau_1}\}$, then all voters believe that with probability one the incumbent is type τ_0 , and all type $\tau \geq \tau_1$ voters vote for the challenger, and type τ_0 voters vote for the incumbent.

Phase 4.0(j) Type τ_0 office holders choose $x = x_{\tau_1}$ if $j \in \{1, \dots, k\}$ and $x = 0$ if $j = k + 1$, and type $\tau \neq \tau_0$ office holders choose $x = x_{\tau_1}$ for all $j = 1, \dots, k + 1$. All voters believe that with probability one the incumbent is type τ_0 . If $j \in \{1, \dots, k\}$ and the incumbent chooses $x = x_{\tau_1}$, or if $j = k + 1$ and $x \in \{0, x_{\tau_1}\}$, then all type $\tau \geq \tau_1$ voters vote for the incumbent, and type τ_0 voters vote for the challenger; and otherwise, type $\tau \geq \tau_1$ voters vote for the challenger, and type τ_0 voters vote for the incumbent.

Phase 4.1 Type $\tau \neq \tau_0$ office holders choose $x = x_{\tau_1}$ and type τ_0 office holders choose $x = 0$. If the incumbent chooses $x = x_{\tau_1}$, then all voters' beliefs are given by \hat{p} , the updated prior conditional on the politician's type being not equal to τ_0 , and all type $\tau \geq \tau_1$ voters vote for the incumbent and type τ_0 voters vote for the challenger; and otherwise, if $x \neq x_{\tau_1}$, then all voters believe that with probability one the incumbent is type τ_0 , and all type $\tau \geq \tau_1$ voters vote for the challenger, and type τ_0 voters vote for the incumbent.

Phase 4.2 Type $\tau \neq \tau_0$ office holders choose $x = x_{\tau_1}$ and type τ_0 office holders choose $x = 0$. If the incumbent chooses $x = x_{\tau_1}$, then all voters' beliefs are given by \hat{p} , the updated prior conditional on the politician's type being not equal to τ_0 , and all type $\tau \geq \tau_1$ voters vote for the incumbent and type τ_0 voters vote for the challenger; and if the incumbent chooses $x = 0$, then all voters believe that with probability one the politician is type τ_0 , and all type $\tau \geq \tau_1$ voters vote for the incumbent and type τ_0 voters vote for the challenger; and if the incumbent chooses $x \notin \{0, x_{\tau_1}\}$, then voters believe that with probability one the politician is type τ_0 , all type $\tau \geq \tau_1$ voters vote for the challenger, and type τ_0 voters vote for the incumbent.

Phase 5 All types of office holder choose their ideal policies. All voters believe that with probability one the incumbent is type τ_0 , and all type $\tau \geq \tau_1$ voters vote for the challenger, and all type τ_0 voters vote for the incumbent.

The specification is completed by assigning the initial history \emptyset to Phase 1, and by defining the rule for transitioning between the phases.

- In Phase 1, if the office holder chooses $x = a$, then remain in Phase 1; if the office holder chooses $x \neq a$ and is re-elected, then move to Phase 5; and if the office holder chooses $x \neq a$ and is not re-elected, then move to Phase 2(1).
- In Phase 2(j), if the office holder chooses $x = 0$ and is re-elected, or if $x \neq 0$ and the incumbent is not re-elected, then move to Phase 2($j + 1$) if $j < m$ and to Phase 3(1) if $j = m$; if the office holder chooses $x = 0$ and is not re-elected, then move to Phase 2(1); and if the office holder chooses $x \neq 0$ and is re-elected, then move to Phase 5.
- In Phase 3(j), if the office holder chooses $x = 0$ and is re-elected, then move to Phase 3($j + 1$) if $j < m$ and to Phase 4.0(1) if $j = m$; if the office holder chooses $x = x_{\tau_1}$ and is re-elected, then move to Phase 3($j + 1$) if $j < m$ and to Phase 4.1 if $j = m$; if the office holder chooses $x \notin \{0, x_{\tau_1}\}$ and is re-elected, then move to Phase 5; if the office holder chooses $x \in \{0, x_{\tau_1}\}$ and is not re-elected, then move to Phase 2(1); and if the office holder chooses $x \notin \{0, x_{\tau_1}\}$ and is not re-elected, then move to Phase 4.2.
- In Phase 4.0(j), first suppose the incumbent is re-elected. If $j \in \{1, \dots, k\}$ and the incumbent chooses $x = x_{\tau_1}$, then move to Phase 4.0($j + 1$); if $j \in \{1, \dots, k\}$ and the incumbent chooses $x \neq x_{\tau_1}$, then move to Phase 5; if $j = k + 1$ and $x \in \{0, x_{\tau_1}\}$, then move to Phase 4.0(1); and if $j = k + 1$ and $x \notin \{0, x_{\tau_1}\}$, then move to Phase 5. Supposing the incumbent is not re-elected, if $j \in \{1, \dots, k\}$ and the incumbent chooses $x = x_{\tau_1}$, or if $j = k + 1$ and $x = \{0, x_{\tau_1}\}$, then move to Phase 2(1); and otherwise, move to Phase 4.2.
- In Phase 4.1, if the office holder chooses $x = x_{\tau_1}$ and is re-elected, then remain in Phase 4.1; if the office holder chooses $x \neq x_{\tau_1}$ and is re-elected, then move to Phase 5; if the office holder chooses $x = x_{\tau_1}$ and is not re-elected, then move to Phase 2(1); and if the office holder chooses $x \neq x_{\tau_1}$ and is not re-elected, then move to Phase 4.2.
- In Phase 4.2, if the office holder chooses $x = 0$ and is re-elected, then move to Phase 4.0(1); if the office holder chooses $x = x_{\tau_1}$ and is re-elected, then move to Phase 4.1; if the office holder chooses $x \notin \{0, x_{\tau_1}\}$ and is re-elected, then move to Phase 5; if the office holder chooses $x \in \{0, x_{\tau_1}\}$ and is not re-elected, then move to Phase 2(1); and if the office holder chooses $x \notin \{0, x_{\tau_1}\}$ and is not re-elected, then remain in Phase 4.2.
- In Phase 5, if the office holder is re-elected, then remain in Phase 5; and otherwise, if the incumbent is not re-elected, then move to Phase 2(1).

This assignment of histories to phases entails a specification of strategies and beliefs. In order to verify that they form a perfect Bayesian equilibrium, I will focus on the most critical best response conditions, addressing the more straightforward conditions with less formal arguments. I do not consider the votes of type τ_0 voters, as they do not affect electoral outcomes.

It is useful to describe the paths of play at the beginning of Phases 1–5. Of course, the path of play beginning in Phase 1 is simply policy choice a and re-election, ad infinitum. Beginning in Phase 2(1), policy $x = 0$ is chosen for m periods and the incumbent is re-elected. Beginning in Phase 3(1), if the office holder is type τ_0 , then she chooses $x = 0$ and is re-elected for m periods; and otherwise, if the office holder is type $\tau \neq \tau_0$, she chooses x_{τ_1} and is re-elected for m periods. We then move to Phase 4.0 or 4.1, depending on whether the office holder is type τ_0 or $\tau \neq \tau_0$. In the latter case, the office holder chooses $x = x_{\tau_1}$ and is re-elected ad infinitum; and in the former case, the office holder alternately chooses $x = x_{\tau_1}$ for k periods and $x = 0$ for one period, and is re-elected ad infinitum.

Next, I calculate the (normalized) continuation value of each citizen type at the beginning of Phase 2(1) after electing an untried challenger:

$$V_{\tau} = p_0[(1 - \delta^{2m})u_{\tau}(0) + \delta^{2m}W_{\tau}] + (1 - p_0)[(1 - \delta^m)u_{\tau}(0) + \delta^m u_{\tau}(x_{\tau_1})].$$

This reflects the fact that with probability p_0 the office holder is type τ_0 and chooses $x = 0$ for $2m$ periods and then alternately chooses $x = x_{\tau_1}$ for k periods and $x = 0$ for one period; and with probability $1 - p_0$ the office holder is type $\tau \neq \tau_0$ and chooses $x = 0$ for m periods and chooses $x = x_{\tau_1}$ thereafter. I employ the one-shot deviation principle to confirm that no office holders of any type can gain by deviating, and that all voter types (in particular, types $\tau \geq \tau_1$ in Phases 2 and 3) vote as though pivotal.

Phase 1: First, consider the choice of a type τ_0 office holder. Following the above strategy, the politician chooses $x = a$ and is continually re-elected. Following any deviation, the politician is removed from office, so the best one-shot deviation is to choose $x = 0$ and to receive V_{τ_0} in expectation after that. The one-shot deviation is unprofitable if

$$u_{\tau_0}(a) \geq (1 - \delta)u_{\tau_0}(0) + \delta V_{\tau_0},$$

or equivalently,

$$\begin{aligned} & (1 - \delta)(u_{\tau_0}(0) - u_{\tau_0}(a)) \\ & \leq \delta \left[p_0[(1 - \delta^{2m})(u_{\tau_0}(a) - u_{\tau_0}(0)) + \delta^{2m}(u_{\tau_0}(a) - W_{\tau_0})] \right. \\ & \quad \left. + (1 - p_0)[(1 - \delta^m)(u_{\tau_0}(a) - u_{\tau_0}(0)) + \delta^m(u_{\tau_0}(a) - u_{\tau_0}(x_{\tau_1}))] \right], \end{aligned}$$

which after substituting $\beta = \delta^m$ is equivalent to (6). Next, consider a type $\tau \neq \tau_0$ office holder. The best one-shot deviation is to choose $x = x_{\tau}$ followed by V_{τ} . The one-shot deviation is unprofitable if

$$u_{\tau}(a) \geq (1 - \delta)u_{\tau}(x_{\tau}) + \delta V_{\tau},$$

which is implied by

$$(1 - \delta)\Delta \leq \delta \left[p_0[(1 - \delta^{2m})(u_{\tau}(a) - u_{\tau}(0)) + \delta^{2m}(u_{\tau}(a) - W_{\tau})] \right]$$

$$+(1-p_0)[(1-\delta^m)(u_\tau(a)-u_\tau(0))+\delta^m(u_\tau(a)-u_\tau(x_{\tau_1}))].$$

which after substituting $\beta = \delta^m$ is equivalent to (7). Now consider a voter of any type. In case the incumbent chooses $x = a$, if the incumbent is re-elected, then she will continue to choose a and be re-elected. If a challenger is elected, then we remain in Phase 1, and the newly elected office holder continues to choose a . Thus, all voters are indifferent between the incumbent and a challenger, and voting for the incumbent is consistent with being pivotal. In case the incumbent chooses $x \neq a$, then all voters believe that with probability one the politician is type τ_0 . If the incumbent is re-elected, then we move to Phase 5, and voters expect that the office holder will choose $x = x_{\tau_0} = 0$, and then an untried challenger is elected, and we move to Phase 2(1). If the challenger is elected, then we move directly to Phase 2(1). Thus, the expected discounted payoff from electing a challenger exceeds that from re-electing the incumbent for a type $\tau \geq \tau_1$ voter if

$$V_\tau \geq (1-\delta)u_\tau(0) + \delta V_\tau,$$

which indeed holds.

Phase 2(j): First, consider the choice of a type τ_0 office holder. Following the above strategy, the politician chooses $x = 0$ and is re-elected for $m - j$ periods, after which we move to Phase 3(1) and she chooses $x = 0$ for m more periods, after which we move to Phase 4.0(1) and she alternately chooses $x = x_{\tau_1}$ for k periods and $x = 0$ for one period. Following any deviation, the politician is removed from office, we continue to Phase 2($j + 1$) if $j < m$ or to Phase 3(1) if $j = m$, and clearly the office holder receives a higher payoff by choosing her ideal policy and avoiding the election of a type $\tau \neq \tau_0$ challenger. Next, consider the choice of a type $\tau \neq \tau_0$ office holder. Following the above strategy, the politician chooses $x = 0$ and is re-elected for $m - j$ periods, and then chooses $x = x_{\tau_1}$ and is continually re-elected. The best one-shot deviation is to choose $x = x_\tau$, after which the politician is removed from office, and we move to Phase 2($j + 1$) or 3(1) with a new office holder. The incentives to deviate are maximal when $j = 1$, as then the cost of deviating is realized with probability p_0 after $m - 1$ periods. The one-shot deviation is unprofitable if

$$\begin{aligned} & (1-\delta)u_\tau(0) + \delta \left[(1-\delta^{m-1})u_\tau(0) + \delta^{m-1}u_\tau(x_{\tau_1}) \right] \\ & \geq (1-\delta)u_\tau(x_\tau) + \delta \left[p_0[(1-\delta^{2m-1})u_\tau(0) + \delta^{2m-1}W_\tau] \right. \\ & \quad \left. + (1-p_0)[(1-\delta^{m-1})u_\tau(0) + \delta^{m-1}u_\tau(x_{\tau_1})] \right], \end{aligned}$$

which is implied by

$$(1-\delta)\Delta \leq \delta p_0 \left[\delta^{m-1}(1-\delta^m)(u_\tau(x_{\tau_1}) - u_\tau(0)) + \delta^{2m-1}(u_\tau(x_{\tau_1}) - W_\tau) \right],$$

which after substituting $\beta = \delta^m$ is implied by (12).

Now consider a type $\tau \geq \tau_1$ voter. In case the incumbent chooses $x = 0$, then the voters' beliefs about the incumbent's type are given by the prior p . If the incumbent is re-elected, then she will choose $x = 0$ for $m - 1$ more periods, and we move to Phase 3(1). If the challenger is elected, then we move back to Phase 2(1). Thus, the expected discounted payoff to the type $\tau \geq \tau_1$ voter from re-electing the incumbent is greater than from electing a challenger. In case the incumbent chooses $x \neq 0$, then the voters believe that with probability one the politician is type τ_0 . If the incumbent is re-elected, then we move to Phase 5, and she chooses $x = 0$, is removed from office, and we move to Phase 2(1). If the challenger is elected, then we remain in Phase 2 for $m - 1$ more periods, and we move to Phase 3(1). Thus, the expected discounted payoff of electing the challenger is greater for the type $\tau \neq \tau_0$ voter than that of re-electing the incumbent.

Phase 3(j): First, consider the choice of a type τ_0 office holder. Following the above strategy, the politician chooses $x = 0$ and is re-elected for $m - j$ more periods, after which we move to Phase 4.0(1) and she alternately chooses $x = x_{\tau_1}$ for k periods and $x = 0$ for one period. Following any deviation $x' \notin \{0, x_{\tau_1}\}$, the politician is removed from office, and we move to Phase 4.2, where a challenger is elected and chooses policy. If the newly elected office holder is type τ_0 , then she chooses $x = 0$, and we move to Phase 4.0(1); and if she is type $\tau \neq \tau_0$, then she chooses $x = x_{\tau_1}$, and we move to Phase 4.1. The gains from the one-shot deviation are highest when $j = m$, and the deviation is unprofitable in this case if

$$W_{\tau_0} > p_0[(1 - \delta)u_{\tau_0}(0) + \delta W_{\tau_0}] + (1 - p_0)u_{\tau_0}(x_{\tau_1}),$$

which is implied by

$$(1 - \delta)u_{\tau_0}(x_{\tau_1}) + \delta W_{\tau_0} > (1 - \delta)u_{\tau_0}(0) + \delta[p_0 W_{\tau_0} + (1 - p_0)u_{\tau_0}(x_{\tau_1})],$$

which is implied by (14). Following the deviation $x' = x_{\tau_1}$, the office holder is re-elected, and we move to Phase 3($j + 1$) if $j < m$ and to Phase 4.1 if $j = m$. In the latter case, the office holder then chooses $x = 0$, is removed from office, and we move to Phase 4.2, where a type $\tau \neq \tau_0$ challenger always leads to Phase 4.1, and a type τ_0 leads to Phase 4.0(1). The office holder's expected discounted payoff is greater following the initial strategy and choosing her ideal policy $x = 0$ before entering Phase 4.0(1). In the former case, after moving to Phase 3($j + 1$), the office holder continues to choose $x = 0$ for $m - j$ periods before moving to Phase 4.0(1), and she does not gain from the deviation.

Next, consider a type $\tau \neq \tau_0$ office holder. Following the above strategy, the politician chooses $x = x_{\tau_1}$ and is re-elected for $m - j$ more periods, after which we move to Phase 4.1, and she continues to choose $x = x_{\tau_1}$ and be re-elected. The best one-shot deviation is to choose $x = x_{\tau}$, after which the politician is removed from office, and we move to Phase 4.2. The one-shot deviation is unprofitable if

$$u_{\tau}(x_{\tau_1}) > (1 - \delta)u_{\tau}(x_{\tau}) + \delta \left[p_0[(1 - \delta)u_{\tau}(0) + \delta W_{\tau}] + (1 - p_0)u_{\tau}(x_{\tau_1}) \right],$$

which is implied by

$$u_\tau(x_{\tau_1}) > (1 - \delta)u_\tau(x_\tau) + \delta \left[p_0 W_\tau + (1 - p_0)u_\tau(x_{\tau_1}) \right],$$

which is implied by (13).

Now consider a type $\tau \geq \tau_1$ voter. In case the office holder chooses $x = x_{\tau_1}$, then the voters' posterior puts probability zero on type τ_0 . If the incumbent is re-elected, then we move either to Phase 3($j + 1$) or to Phase 4.1, and the politician continues to choose $x = x_{\tau_1}$ and be re-elected. If the incumbent is removed from office, then we move to Phase 2(1), leading to a lower expected discounted payoff.² In case the office holder chooses $x = 0$, then the voters' believe that with probability one the incumbent is type τ_0 . If the incumbent is re-elected, then we move either to Phase 3($j + 1$) or to Phase 4.0(1). If the incumbent is removed from office, then we move to Phase 2(1). Thus, the expected discounted payoff from re-electing the incumbent exceeds that from electing the challenger for a type $\tau \neq \tau_0$ voter if

$$(1 - \delta^{m-j})u_\tau(0) + \delta^{m-j}W_\tau > V_\tau. \quad (15)$$

This inequality is most restrictive when $j = 1$, in which case it is implied by (5) after substituting $\beta = \delta^m$.³ In case the office holder chooses $x \neq x_{\tau_1}$, then all voters believe that with probability one the incumbent is type τ_0 . If the incumbent is re-elected, then we move to Phase 5, she chooses $x = 0$, is removed from office, and we move to Phase 2(1). If the challenger is elected, then we move to Phase 4.2, and if the newly elected office holder is type τ_0 , then she chooses $x = 0$, and we move to Phase 4.0(1); and if she is type $\tau \neq \tau_0$, then she chooses $x = x_{\tau_1}$, and we move to Phase 4.1. Since re-electing the incumbent leads to policy $x = 0$ in the subsequent period, whereas the election of a challenger leads to $x \in \{0, x_{\tau_1}\}$, it suffices to compare the stream of payoffs beginning two periods hence. The expected discounted payoff from electing the challenger exceeds that from re-electing the incumbent for a type $\tau \geq \tau_1$ voter if

$$p_0 W_\tau + (1 - p_0)u_\tau(x_{\tau_1}) > (1 - \delta^m)u_\tau(0) + \delta^m \left[p_0 [(1 - \delta^m)u_\tau(0) + \delta^m W_\tau] + (1 - p_0)[u_\tau(x_{\tau_1})] \right],$$

where we use the fact that in Phase 2(1), all types of office holder choose $x = 0$ for the first m periods. The above inequality holds if

$$W_\tau > (1 - \delta)u_\tau(0) + \delta \left[p_0 [(1 - \delta^m)u_\tau(0) + \delta^m W_\tau] \right]$$

²For this argument, it is not crucial that voter beliefs are given by the posterior \hat{p} conditional on the office holder's type belonging to $\{\tau_1, \dots, \tau_n\}$. What is required is merely that beliefs put probability zero on τ_0 .

³Note that (5) is actually sufficient for the stronger conclusion with $j = 0$, so that the voter's payoff would decrease even if the incumbent were removed immediately prior to choosing $x = 0$ in Phase 3(1) and we moved instead directly to Phase 2(1).

$$+(1 - p_0)u_\tau(x_{\tau_1}) \Big],$$

which is implied by setting $j = 1$ and deleting the first $m - 1$ realizations of $x = 0$ from both sides of (15).

Phase 4.0(j): First, consider the choice of a type τ_0 office holder. Following the above strategy, the politician alternately chooses $x = x_{\tau_1}$ for k periods and $x = 0$ for one period and is continually re-elected. The gains from deviating are highest when $j = 1$, and the best one-shot deviation is to choose $x' = 0$, after which the incumbent is replaced by a challenger and we move to Phase 4.2. The one-shot deviation is unprofitable if

$$W_{\tau_0} > (1 - \delta)u_{\tau_0}(0) + \delta[p_0W_{\tau_0} + (1 - p_0)u_{\tau_0}(x_{\tau_1})],$$

which is implied by

$$(1 - \delta)u_{\tau_0}(x_{\tau_1}) + \delta W_{\tau_0} > (1 - \delta)u_{\tau_0}(0) + \delta[p_0W_{\tau_0} + (1 - p_0)u_{\tau_0}(x_{\tau_1})],$$

which is implied by (14). Next, consider a type $\tau \neq \tau_0$ office holder. Following the above strategy, the politician chooses $x = x_{\tau_1}$ and is re-elected in every period. The gains from deviating are the same as in Phase 3, and it follows from (13) and the above arguments that one-shot deviations are unprofitable.

Now consider a type $\tau \geq \tau_1$ voter. All voters believe that with probability one the incumbent is type τ_0 . In case $j \in \{1, \dots, k\}$ and the incumbent chooses $x = x_{\tau_1}$, or in case $j = k + 1$ and $x \in \{0, x_{\tau_1}\}$, if the incumbent is re-elected, then we remain in Phase 4.0, and the office holder alternately chooses $x = x_{\tau_1}$ for k periods and $x = 0$ for one period. If the incumbent is removed from office, then we move to Phase 2(1). The expected discounted payoff from re-electing the incumbent is greater than in Phase 3(1), whereas the expected discounted payoff from electing the challenger is the same as in Phase 3(1), so (5) and the arguments above imply that voting for the incumbent is consistent with the voter being pivotal. In the complementary case, if the incumbent is re-elected, then we move to Phase 5, the office holder chooses $x = 0$ and is removed from office, and we move to Phase 2(1). If the challenger is elected, then we move to Phase 4.2, where the new office holder chooses $x = 0$ and we move to Phase 4.0(1) if she is type τ_0 , and she chooses $x = x_{\tau_1}$ ad infinitum otherwise. The expected discounted payoff from electing the challenger exceeds that from re-electing the incumbent, because we have

$$\begin{aligned} & p_0[(1 - \delta)u_\tau(0) + \delta W_\tau] + (1 - p_0)u_\tau(x_{\tau_1}) \\ & > (1 - \delta)u_\tau(0) + \delta W_\tau \\ & > p_0[(1 - \delta^{2m})u_\tau(0) + \delta^{2m}W_\tau] \\ & \quad + (1 - p_0)[(1 - \delta^m)u_\tau(0) + \delta^m u_\tau(x_{\tau_1})], \end{aligned}$$

where the second inequality above follows by setting $j = m - 1$ in (15).

Phase 4.1: First, consider the choice of a type τ_0 office holder. Following the above strategy, the politician chooses $x = 0$, is removed from office, and

we move to Phase 4.2. Deviating to $x' \notin \{0, x_{\tau_1}\}$ again leads to election of a challenger and Phase 4.2, so this cannot be profitable. If the politician deviates to $x' = x_{\tau_1}$, then she is re-elected and we remain in Phase 4.1, she then chooses $x = 0$, is removed from office, and we move to Phase 4.2. The office holder's expected discounted payoff is higher following the initial strategy and obtaining her ideal policy immediately before moving to Phase 4.2. Next, consider a type $\tau \neq \tau_0$ office holder. Following the above strategy, the politician chooses $x = x_{\tau_1}$ and is re-elected in every period. The gains from deviating are the same as in Phase 3, and it follows from (13) and the above arguments that one-shot deviations are unprofitable.

Now consider a type $\tau \geq \tau_0$ voter. In case the incumbent chooses $x = x_{\tau_1}$, then all voters' beliefs are given by \hat{p} . If the incumbent is re-elected, then she continues to choose $x = x_{\tau_1}$ and be re-elected. If the incumbent is removed from office, then we move to Phase 2(1). The expected discounted payoff from re-electing the incumbent exceeds that from electing a challenger if $u_\tau(x_{\tau_1}) > V_\tau$, which is implied by (15). In case the incumbent chooses $x \neq x_{\tau_1}$, then all voters believe that with probability one the incumbent is type τ_0 . If the incumbent is re-elected, then we move to Phase 5, she chooses $x = 0$ and is removed from office, and we move to Phase 4.2. If the challenger is elected, then we move to Phase 4.2. Thus, the expected discounted payoff from electing the challenger exceeds that from re-electing the incumbent.

Phase 4.2: First, consider the choice of a type τ_0 office holder. Following the above strategy, the politician chooses $x = 0$, is re-elected, and we move to Phase 4.0(1), where she alternately chooses $x = x_{\tau_1}$ for k periods and $x = 0$ for one period. If the politician deviates to $x = x_{\tau_1}$, then she is re-elected, we move to Phase 4.1, she then chooses $x = 0$ and is removed from office, and we move to Phase 4.2. The one-shot deviation is unprofitable if

$$(1 - \delta)u_{\tau_0}(0) + \delta W_{\tau_0} > (1 - \delta)u_{\tau_0}(x_{\tau_1}) + (1 - \delta)\delta u_{\tau_0}(0) + \delta^2 \left[p_0[(1 - \delta)u_{\tau_0}(0) + \delta W_{\tau_0}] + (1 - p_0)u_{\tau_0}(x_{\tau_1}) \right],$$

which is implied by

$$(1 - \delta)u_{\tau_0}(0) + (1 - \delta)\delta u_{\tau_0}(x_{\tau_1}) + \delta^2 W_{\tau_0} > (1 - \delta)u_{\tau_0}(x_{\tau_1}) + (1 - \delta)\delta u_{\tau_0}(0) + \delta^2 \left[p_0[(1 - \delta)u_{\tau_0}(0) + \delta W_{\tau_0}] + (1 - p_0)u_{\tau_0}(x_{\tau_1}) \right],$$

which is implied by

$$W_{\tau_0} > p_0[(1 - \delta)u_{\tau_0}(0) + \delta W_{\tau_0}] + (1 - p_0)u_{\tau_0}(x_{\tau_1}),$$

and finally, following arguments in Phase 3, this is implied by implied by (14).

Phase 5: Since the incumbent is always removed from office, it is optimal for office holder's of all types to choose their ideal policies. Now consider a type $\tau \geq \tau_1$ voter. All voters believe that with probability one the incumbent is type τ_0 .

If the incumbent is re-elected, then we remain in Phase 5, the politician chooses $x = 0$, is removed from office, and we move to Phase 2(1). If the challenger is elected, then we move directly to Phase 2(1). Thus, the expected discounted payoff from electing the challenger exceeds that from re-electing the incumbent.

To conclude, for all $\epsilon > 0$, we can choose $a \in (0, \epsilon)$, natural number k , and $\hat{\delta} \in (0, 1)$ such that for all $\delta \in (\hat{\delta}, 1)$, there exists m such that (5)–(7) and (12)–(14) are satisfied by a , δ , and $\beta = \delta^m$. Given these parameters, I have specified strategies and beliefs such that the policy a is chosen along the path of play by all types of office holder, voter beliefs satisfy Bayes rule along each office holder's personal path of play, and no politician has a profitable one-shot deviation, and such that all voter types vote as though pivotal. This completes Step 1.

An argument symmetric to that for Step 1 yields the next step of the proof.

Step 2: *For all $\epsilon > 0$, there exist $b \in (1 - \epsilon, 1)$ and $\hat{\delta} \in (0, 1)$ such that for all $\delta \in (\hat{\delta}, 1)$, we have $\mathbf{x}^b \in P(\delta)$.*

Given the argument of Step 1, I do not provide the (redundant) details of Step 2. But it will be important to note that a and b can be chosen close enough to zero and one, respectively, to allow the same choice of k and m in both steps. To be more explicit, choose k such that for all $\tau \neq \tau_0$, (1) holds, and such that for all $\tau \neq \tau_n$, the corresponding inequality for Step 2 holds. Choose $\bar{\beta}$ such that for all $\beta \in (0, \bar{\beta})$, (4) and the corresponding inequality for Step 2 hold, and such that there exists $\underline{\delta} \in (0, 1)$ such that for all $\delta \in (\underline{\delta}, 1)$ and all $\beta \in (0, \bar{\beta})$, we have (5) for all $\tau \neq \tau_0$ and the corresponding inequality for Step 2 for all $\tau \neq \tau_n$. Furthermore, define

$$\begin{aligned} B_\tau &= u_\tau(b) - u_\tau(1) \\ B_\tau^* &= u_\tau(x_{\tau_{n-1}}) - u_\tau(b) \\ \underline{B} &= \min_{\tau \neq \tau_n} B_\tau \\ B^* &= \max_{\tau \neq \tau_n} B_\tau^*, \end{aligned}$$

and note that by the arguments in Step 1, we have

$$\lim_{b \uparrow 0} \left(\frac{u_{\tau_n}(1) - u_{\tau_n}(b)}{u_{\tau_n}(b) - u_{\tau_n}(x_{\tau_{n-1}})} \right) \left(\frac{B^*}{\underline{B}} \right) = 0.$$

Thus, we can choose $a \in (0, \epsilon)$, $b \in (1 - \epsilon, 1)$, and $\beta \in (0, \bar{\beta})$ such that (10) and (11) hold, along with the corresponding inequalities for Step 2. Then the arguments in Step 1 allow us to specify $\hat{\delta} \in (\underline{\delta}, 1)$ so that for all $\delta \in (\hat{\delta}, 1)$, there exists m such that (12)–(14) hold, along with the corresponding inequalities for Step 2. This permits the equilibrium construction from Step 1 to be applied in Step 2 with the same choice of k and m .

Next, I define a procedure for splicing the equilibria from Steps 1 and 2 to produce a path of policies that alternates between a and b , and I use the

construction recursively to generate paths that cycle between the alternatives to approximate any empirical frequency.

Step 3: For all $\epsilon > 0$, there exist $a \in (0, \epsilon)$ and $b \in (1 - \epsilon, 1)$ such that for all natural numbers ℓ and all k with $k = 0, 1, \dots, 2^\ell$, there exists $\tilde{\delta} \in (0, 1)$ such that for all $\delta \in (\tilde{\delta}, 1)$, there exists $\mathbf{x} \in P(\delta)$ such that in consecutive runs of 2^ℓ periods, the path \mathbf{x} visits a for k periods and b for $2^\ell - k$ periods, i.e., for all $m = 0, 1, 2, \dots$,

$$\begin{aligned} \#\{t \mid m2^\ell + 1 \leq t \leq (m+1)2^\ell \text{ and } x_t = a\} &= k \\ \#\{t \mid m2^\ell + 1 \leq t \leq (m+1)2^\ell \text{ and } x_t = b\} &= 2^\ell - k. \end{aligned}$$

The proof is by induction. Fix $\epsilon > 0$. By Step 1, there exist $a \in (0, \epsilon)$ and $\delta' < 1$ such that for all $\delta \in (\delta', 1)$, we have $\mathbf{x}^a \in P(\delta)$. Similarly, by Step 2, there exist $b \in (1 - \epsilon, 1)$ and $\delta'' < 1$ such that for all $\delta \in (\delta'', 1)$, we have $\mathbf{x}^b \in P(\delta)$. For the basis step, consider $\ell = 1$. For $k = 0$ and $k = 2$, the desired result follows from Steps 2 and 1, respectively. Consider $k = 1$, so either $\mathbf{x} = (a, b, a, b, \dots)$ or $\mathbf{x} = (b, a, b, a, \dots)$. I focus on the former case without loss of generality. Set $\tilde{\delta} = \sqrt{\max\{\delta', \delta''\}}$, choose $\delta \in (\delta', 1)$, and define strategies and beliefs by “splicing” the specifications in Steps 1 and 2. That is, we specify actions and beliefs according to the equilibrium from Step 1 in odd periods and using the equilibrium from Step 2 in even periods.

To be more precise, we begin in Phase 1 in both odd and even periods, so that the path of play from the initial history is the sequence (a, b, a, b, \dots) . Following any deviation in an odd period t , for example, we specify that the transition between phases in subsequent odd periods follow the protocol from Step 1 (applied to odd periods). And we specify that the phase in subsequent even periods follow the same transition, so that the phase in each odd period $t' > t$ and the next even period $t' + 1$ are synchronized, with the exception that when entering Phase 4.0 in an odd period, we enter Phase 4.1 in the next even period, and when entering Phase 4.1 in an odd period, we enter Phase 4.0 in the next even period. Policy and voting choices are given by the strategies from Steps 1 and 2, although now we must take care in specifying voter beliefs when we enter Phase 5 and when information is revealed in Phases 3(1) and 4.2. First, note that Phase 5 is entered only off a politician’s personal path of play and only when she is re-elected, so that voter beliefs may be specified arbitrarily. Let $\bar{x} \in \{0, 1\}$ be the policy that is weakly preferred among $\{0, 1\}$ by a majority of voters, and let $\underline{x} \in \{0, 1\}$ be the remaining policy, so that

$$\sum \{q_j \mid u_{\tau_j}(\bar{x}) \geq u_{\tau_j}(\underline{x})\} \geq \frac{1}{2}.$$

Thus, by concavity, the policy \underline{x} is the worst possible policy for a majority of voters. Let $\underline{\tau} \in \{\tau_0, \tau_n\}$ be the type such that $x_{\underline{\tau}} = \underline{x}$. In Phase 5, whether in an odd or even period, we specify that voters believe with probability one that the office holder is type $\underline{\tau}$. This ensures that a majority of voters prefer to replace the incumbent with a challenger, as dictated in Steps 1 and 2.

Second, after entering Phase 3(1) in an odd period t , if the office holder chooses $x = 0$, then voters believe that with probability one the office holder is type τ_0 , and if she is re-elected, then we move to Phase 3(1) in the following even period $t+1$. As noted in footnote 2, the information revealed by this choice does not affect the incentives of type $\tau \neq \tau_n$ voters to re-elect the incumbent prior to period $t+1$. If the office holder chooses $x = x_{\tau_1}$, then voters beliefs are given by the posterior \hat{p} conditional on the office holder's type belonging to $\{\tau_1, \dots, \tau_n\}$, and if she is re-elected, then we move to Phase 3(1) in the following even period. As noted in footnote 3, this revelation of information does not affect the incentives voters to re-elect the incumbent. If the office holder chooses $x \notin \{0, x_{\tau_1}\}$, then voters believe she is type $\underline{\tau}$, and if she is re-elected, then we move to Phase 5; thus, the office holder will be removed from office, as required. Supposing the incumbent is re-elected in period t , voter beliefs are updated in period $t+1$ and phases transition as follows. If voter beliefs are that the incumbent is type τ_0 and she chooses $x = x_{\tau_{n-1}}$, then beliefs are unchanged, and if she is re-elected, then we move to Phase 3(2) in the following odd period; again footnote 3 implies that the incumbent will be re-elected. If voter beliefs are that the incumbent is type τ_0 and she chooses $x \neq x_{\tau_{n-1}}$, then voters believe she is type $\underline{\tau}$, and if she is re-elected, then we move to Phase 5, and she is removed from office. If voter beliefs are \hat{p} and she chooses $x = x_{\tau_{n-1}}$, then voter beliefs are given by the posterior conditional on the office holder's type belonging to $\{\tau_1, \dots, \tau_{n-1}\}$, and if she is re-elected, then we move to Phase 3(2) in the next odd period $t+2$. If voter beliefs are \hat{p} and she chooses $x = 1$, then voters believe that with probability one the office holder is type τ_n , and if she is re-elected, then we move to Phase 3(2); again footnote 2 implies that the incumbent will be re-elected. And if the office holder chooses $x \notin \{x_{\tau_{n-1}}, 1\}$, then voters believe she is type $\underline{\tau}$, and we move to Phase 5. Supposing the incumbent is re-elected in $t+1$, voter beliefs are again given by Bayes rule, with deviations punished by Phase 5.

Third, after entering Phase 4.2 in an odd period t , if the office holder chooses $x = 0$, then voters believe that with probability one the office holder is type τ_0 , and if she is re-elected, then we move to Phase 4.1 in the following even period period. If the office holder chooses $x = x_{\tau_1}$, then voter beliefs are given by \hat{p} , and if she is re-elected, then we move to Phase 4.2 in the following even period. Supposing the incumbent is re-elected in period t , voter beliefs are updated in period $t+1$ and phases transition as follows. In Phase 4.1, voters believe that the office holder is type τ_0 . If she chooses $x = x_{\tau_{n-1}}$, then beliefs are unchanged, and if she is re-elected, then we move to Phase 4.0(1) in the next period. If she chooses $x \neq x_{\tau_{n-1}}$, then voters believe she is type $\underline{\tau}$, and if she is re-elected, then we move to Phase 5. In Phase 4.2, voters' beliefs about the office holder's type are given by \hat{p} . If the office holder chooses $x = 1$, then voters believe that with probability one the office holder is type τ_n , and if she is re-elected, then we move to Phase 4.1. If the office holder chooses $x = x_{\tau_{n-1}}$, then beliefs are given by the posterior conditional on the politician's type belonging to $\{\tau_1, \dots, \tau_{n-1}\}$, and if she is re-elected, then we move to Phase 4.1. And if the office holder chooses $x \notin \{x_{\tau_{n-1}}, 1\}$, then voters believe she is type $\underline{\tau}$, and if

she is re-elected, then we move to Phase 5. Because removal of the incumbent following her specified choice of policy leads to Phase 2(1) in the next odd and even periods, a majority of voters prefer to re-elect the incumbent, as dictated in Steps 1 and 2. Supposing the incumbent is re-elected in $t + 1$, voter beliefs are again updated via Bayes rule.

In effect, this splicing of strategies punishes a deviation in an odd period t by following the punishments from Step 1 in all future odd periods, and by following Step 2 in all future even periods, simulating punishments as though an analogous deviation had occurred in the previous even period $t - 1$. A similar construction is used following a deviation in an even period. This establishes a perfect Bayesian electoral equilibrium with path of play (a, b, a, b, \dots) and provides a splicing procedure that can be iterated generally.

For the induction step, assume the claim is true for $\ell - 1$, choose any $k = 0, 1, \dots, 2^\ell$, and write $k = k' + k''$, where $k', k'' \leq 2^{\ell-1}$. By the induction hypothesis, there exists $\delta' \in (0, 1)$ such that for all $\delta \in (\delta', 1)$, there exists $\mathbf{y} \in P(\delta)$ such that for all $m = 0, 1, 2, \dots$,

$$\begin{aligned} \#\{t \mid m2^{\ell-1} + 1 \leq t \leq (m+1)2^{\ell-1} \text{ and } y_t = a\} &= k' \\ \#\{t \mid m2^{\ell-1} + 1 \leq t \leq (m+1)2^{\ell-1} \text{ and } y_t = b\} &= 2^\ell - k'. \end{aligned}$$

And there exists $\delta'' \in (0, 1)$ such that for all $\delta \in (\delta'', 1)$, there exists $\mathbf{z} \in P(\delta)$ such that for all $m = 0, 1, 2, \dots$,

$$\begin{aligned} \#\{t \mid m2^{\ell-1} + 1 \leq t \leq (m+1)2^{\ell-1} \text{ and } z_t = a\} &= k'' \\ \#\{t \mid m2^{\ell-1} + 1 \leq t \leq (m+1)2^{\ell-1} \text{ and } z_t = b\} &= 2^\ell - k''. \end{aligned}$$

We then set $\tilde{\delta} = \sqrt{\max\{\delta', \delta''\}}$, and for every $\delta > \tilde{\delta}$, we splice the equilibria supporting paths \mathbf{x} and \mathbf{y} according to the above procedure. This yields a perfect Bayesian electoral equilibrium that generates the path $\mathbf{x} = (x_1, x_2, \dots) = (y_1, z_1, y_2, z_2, \dots)$ from the initial history. Formally, we define \mathbf{x} by

$$x_t = \begin{cases} y_{\frac{t+1}{2}} & \text{if } t \text{ is odd,} \\ z_{\frac{t}{2}} & \text{if } t \text{ is even.} \end{cases}$$

Then for all $m = 0, 1, 2, \dots$, we have

$$\begin{aligned} &\#\{t \mid m2^\ell + 1 \leq t \leq (m+1)2^\ell \text{ and } x_t = a\} \\ &= \#\{t \mid m2^{\ell-1} + 1 \leq t \leq (m+1)2^{\ell-1}, t \text{ odd, and } y_t = a\} \\ &\quad + \#\{t \mid m2^{\ell-1} + 1 \leq t \leq (m+1)2^{\ell-1}, t \text{ even, and } z_t = a\} \\ &= k' + k'' \\ &= k, \end{aligned}$$

where the first equality follows by noting that when t (odd) ranges from $m2^\ell + 1$ to $(m+1)2^\ell - 1$, the ratio $\frac{t+1}{2}$ ranges from $m2^{\ell-1} + 1$ to $(m+1)2^{\ell-1}$, and similarly, when t (even) ranges from $m2^\ell + 2$ to $(m+1)2^\ell$, the ratio $\frac{t}{2}$ ranges

from $m2^{\ell-1} + 1$ to $(m + 1)2^{\ell-1}$. This establishes the desired perfect Bayesian equilibrium.

The final step of the proof uses the equilibria from Step 3 to punish deviations from arbitrary paths through a region of the policy space that becomes arbitrarily large as citizens become patient.

Step 4: For all $a, b \in X$, with $0 < a < b < 1$, there exists $\bar{\delta} \in (0, 1)$ such that for all $\delta \in (\bar{\delta}, 1)$, we have $[a, b]^\infty \subseteq P(\delta)$.

Consider $a, b \in X$ with $0 < a < b < 1$. Fix $\epsilon = \frac{1}{2} \min\{a, 1 - b\}$, and let $a' \in (0, \epsilon)$ and $b' \in (1 - \epsilon, 1)$ be as in Step 3. For each $x \in [a, b]$, we can write $x = \alpha(x)a' + (1 - \alpha(x))b'$, where

$$\alpha(x) = \frac{b' - x}{b' - a'},$$

and then we have

$$0 < \frac{b' - b}{b' - a'} \leq \alpha(x) \leq \frac{b' - a}{b' - a'} < 1.$$

Since each u_τ is strictly concave, it follows that for all $x \in [a, b]$, we have $u_\tau(x) > \alpha(x)u_\tau(a') + (1 - \alpha(x))u_\tau(b')$. More generally, let Λ denote the set of probability measures λ with support contained in $[a, b]$, and note that for all $\lambda \in \Lambda$, we have

$$\begin{aligned} \int u_\tau(x)\lambda(dx) &> \int [\alpha(x)u_\tau(a') + (1 - \alpha(x))u_\tau(b')]\lambda(dx) \\ &= \alpha(E[\lambda])u_\tau(a') + (1 - \alpha(E[\lambda]))u_\tau(b'). \end{aligned}$$

Therefore, since u_τ is continuous and Λ is compact in the weak* topology, we have

$$\min_{\lambda \in \Lambda} \left[\int u_\tau(x)\lambda(dx) - \alpha(E[\lambda])u_\tau(a') - (1 - \alpha(E[\lambda]))u_\tau(b') \right] \equiv \eta > 0.$$

For all $\lambda \in \Lambda$, we can approximate $\alpha(E[\lambda])$ by $\frac{k}{2^\ell}$ through choice of the natural numbers ℓ and $k = 0, 1, \dots, 2^\ell$. Thus, there exist k and ℓ such that

$$\int u_\tau(x)\lambda(dx) > \left(\frac{k}{2^\ell}\right)u_\tau(a') + \left(1 - \frac{k}{2^\ell}\right)u_\tau(b') + \frac{\eta}{2}. \quad (16)$$

Define $\Lambda^{k,\ell}$ as the set of probability measures λ such that (16) holds. It follows that $\{\Lambda^{k,\ell}\}_{k,\ell}$ is an open covering of Λ , and since the latter set is compact, it possesses a finite subcover. Let F be a finite set of (k, ℓ) pairs, with ℓ a natural number and $k \in \{0, 1, 2, \dots, 2^\ell\}$, such that $\Lambda \subseteq \bigcup_{(k,\ell) \in F} \Lambda^{k,\ell}$. By repeated application of Step 3, we can choose $\tilde{\delta} \in (0, 1)$ such that for all $\delta \in (\tilde{\delta}, 1)$ and all $(k, \ell) \in F$, there exists $\mathbf{x}^{k,\ell} \in P(\delta)$ such that for all $m = 0, 1, 2, \dots$,

$$\begin{aligned} \#\{t \mid m2^\ell + 1 \leq t \leq (m + 1)2^\ell \text{ and } x_t = a'\} &= k \\ \#\{t \mid m2^\ell + 1 \leq t \leq (m + 1)2^\ell \text{ and } x_t = b'\} &= 2^\ell - k. \end{aligned}$$

Given any pair (k, ℓ) , discount factor $\delta < 1$, and type τ , the discounted sum of utility from $\mathbf{x}^{k, \ell}$, denoted $U_\tau^{k, \ell}$, satisfies

$$U_\tau^{k, \ell} = \sum_{t=1}^{2^\ell} \delta^{t-1} [I_{a'}(x_t^{k, \ell})u_\tau(a') + I_{b'}(x_t^{k, \ell})u_\tau(b')] + \delta^{2^\ell} U_\tau^{k, \ell},$$

and therefore, after normalizing by $(1 - \delta)$, we have

$$(1 - \delta)U_\tau^{k, \ell} = \frac{1 - \delta}{1 - \delta^{2^\ell}} \sum_{t=1}^{2^\ell} \delta^{t-1} [I_{a'}(x_t^{k, \ell})u_\tau(a') + I_{b'}(x_t^{k, \ell})u_\tau(b')],$$

where $I_{a'}$ and $I_{b'}$ are the indicator functions for a' and b' . Taking the limit as δ goes to one and applying l'Hospital's rule, we have

$$\begin{aligned} \lim_{\delta \uparrow 1} (1 - \delta)U_\tau^{k, \ell} &= \frac{1}{2^\ell} \sum_{t=1}^{2^\ell} [I_{a'}(x_t^{k, \ell})u_\tau(a') + I_{b'}(x_t^{k, \ell})u_\tau(b')] \\ &= \left(\frac{k}{2^\ell}\right) u_\tau(a') + \left(1 - \frac{k}{2^\ell}\right) u_\tau(b'). \end{aligned}$$

Now let Λ^δ consist of all probability measures $\lambda \in \Lambda$ such that for some $(k, \ell) \in F$, we have

$$\int u_\tau(x)\lambda(dx) > (1 - \delta)U_\tau^{k, \ell} + \frac{\eta}{4}. \quad (17)$$

Since $\Lambda \subseteq \bigcup_{(k, \ell) \in F} \Lambda^{k, \ell}$ and $(1 - \delta)U_\tau^{k, \ell} \rightarrow \left(\frac{k}{2^\ell}\right) u_\tau(a') + \left(1 - \frac{k}{2^\ell}\right) u_\tau(b')$, it follows from (16) that for each $\lambda \in \Lambda$, we have $\lambda \in \Lambda^\delta$ for δ sufficiently close to one. That is, $\{\Lambda^\delta \mid \delta \in (0, 1)\}$ is an open covering of Λ . Again, since Λ is compact, this collection possesses a finite subcover, indexed by a finite set $D \subseteq (0, 1)$. Therefore, setting $\delta' = \max D < 1$, it follows that for all $\delta \in (\delta', 1)$ and all $\lambda \in \Lambda$, there exists $(k, \ell) \in F$ such that (17) holds. Furthermore, choose $\bar{\delta} \in (\delta', 1)$ such that

$$\Delta \leq \left(\frac{\bar{\delta}}{1 - \bar{\delta}}\right) \left(\frac{\eta}{4}\right). \quad (18)$$

Finally, consider any $\delta \in (\bar{\delta}, 1)$ and any $\mathbf{x} \in [a, b]^\infty$. I specify a perfect Bayesian equilibrium that generates \mathbf{x} from the initial history. Note that for each t , we can write the discounted utility of a type τ citizen from $(x_t, x_{t+1}, x_{t+2}, \dots)$ as the integral with respect to a probability measure λ_t , i.e.,

$$\sum_{t'=t}^{\infty} \delta^{t'-t} u_\tau(x_{t'}) = \frac{1}{1 - \delta} \int u_\tau(x)\lambda(dx), \quad (19)$$

where λ_t is defined by

$$\lambda_t(Y) = (1 - \delta) \sum_{t': x_{t'} \in Y} \delta^{t'-t}$$

for all measurable $Y \subseteq [0, 1]$. In period t of Phase 0, all types of office holder choose x_t , and for all policy choices, voters' beliefs are given by the prior p , and all voter types vote to re-elect the incumbent. In Phase t , for $t = 1, 2, \dots$, let $(k, \ell) \in F$ satisfy (17) with $\lambda = \lambda_{t+1}$, and use Step 4 to specify strategies and beliefs so that they form a perfect Bayesian equilibrium that generates the path $\mathbf{x}^{k, \ell}$ beginning in period $t + 1$, i.e., the path beginning in period $t + 1$ is $(x_1^{k, \ell}, x_2^{k, \ell}, \dots)$. The initial history is labelled Phase 0, and in period t of Phase 0, if it is not the case that the office holder chooses x_t and is re-elected, then transition to the punishment Phase $t + 1$ in the next period $t + 1$.

To verify that this specification constitutes a perfect Bayesian equilibrium, note that the discounted payoff to a type τ citizen from a one-shot deviation in period t of Phase 0 is bounded above by $\max_{x \in X} u_\tau(x) + \delta U_\tau^{k, \ell}$. Furthermore,

$$\begin{aligned} & \max_{x \in X} u_\tau(x) + \delta U_\tau^{k, \ell} \\ & \leq \min_{x \in X} u_\tau(x) + \left(\frac{\delta}{1 - \delta} \right) \left(\frac{\eta}{4} \right) + \delta U_\tau^{k, \ell} \\ & = \min_{x \in X} u_\tau(x) + \frac{\delta}{1 - \delta} \left[\frac{\eta}{4} + (1 - \delta) U_\tau^{k, \ell} \right] \\ & < \min_{x \in X} u_\tau(x) + \frac{\delta}{1 - \delta} \int u_\tau(x) \lambda_{t+1}(dx) \\ & \leq \frac{1}{1 - \delta} \sum_{t'=t}^{\infty} \delta^{t'-t} u_\tau(x_{t'}), \end{aligned}$$

where the inequalities follow from (18), (17), and (19), respectively. Therefore, one-shot deviations are unprofitable, and we conclude that there exists a perfect Bayesian equilibrium that generates the path \mathbf{x} from the initial history.

5 Conclusion

When citizens are sufficiently patient, arbitrary policy paths through arbitrarily large regions of the policy space can be generated as the result of equilibrium play. Interestingly, Duggan and Fey (2006) conclude as a result of their analysis of repeated Downsian elections that the background assumptions of the Downsian model must be re-examined.

As in the electoral accountability approach, alternatives [to the Downsian model] may involve policy motivations for candidates, dropping the commitment assumption (as in the literature on citizen-candidates), allowing for imperfect information about voter preferences (as in the literature on probabilistic voting), or some combination of these directions. (p.56)

The current paper has combined two of the three alternatives listed above, with the finding that the indeterminacy of equilibrium policy paths continues to hold—even when politician types are private information. Because voters only

remove the incumbent when they have a strict preference for a challenger, and because they are indifferent only when all types of office holder are expected to choose the same policies, the folk theorem result would persist if the model were modified so that in each period, the distribution of ideal policies in the electorate were subject to small idiosyncratic shocks. (This could be modeled by adding noise to the distribution q each period, or by perturbing the ideal policy x_τ for each voter type.) Thus, it appears that the indeterminacy is fundamental.

As mentioned in the Introduction, the addition of moral hazard to the model would significantly alter the equilibrium analysis and limit the scope for punishment of politicians, so this modeling approach may yet entail substantial restrictions on equilibrium behavior. In lieu of this, or another departure from the pure adverse selection framework, the application of dynamic electoral models will rely on equilibrium refinements (e.g., the common restriction to stationary equilibria) beyond the concept of perfect Bayesian electoral equilibrium considered in this paper.

References

- Banks, J., Duggan, J., 2008. A dynamic model of democratic elections in multi-dimensional policy spaces. *Quarterly Journal of Political Science* 3, 269–299.
- Banks, J., Sundaram, R., 1993. Adverse selection and moral hazard in a repeated election model. In: William Barnett et al. (Ed.), *Political Economy: Institutions, Information, Competition, and Representation*. Cambridge University Press.
- Banks, J., Sundaram, R., 1998. Optimal retention in agency problems. *Journal of Economic Theory* 82, 293–323.
- Bernhardt, D., Camara, O., Squintani, F., 2011. Competence and ideology. *Review of Economic Studies* 78 (487-522).
- Bernhardt, D., Campuzano, L., Squintani, F., Camara, O., 2009. On the benefits of party competition. *Games and Economic Behavior* 66, 685–707.
- Bernhardt, D., Dubey, S., Hughson, E., 2004. Term limits and pork barrel politics. *Journal of Public Economics* 88, 2383–2422.
- Camara, O., 2012. Economic policies of heterogeneous politicians.
- Duggan, J., 2000. Repeated elections with asymmetric information. *Economics and Politics* 12, 109–136.
- Duggan, J., Fey, M., 2006. Repeated downsian electoral competition. *International Journal of Game Theory* 35, 39–69.
- Dutta, P., 1995. A folk theorem for stochastic games. *Journal of Economic Theory* 66, 1–32.

- Ferejohn, J., 1986. Incumbent performance and electoral control. *Public Choice* 50, 5–25.
- Fudenberg, D., Levine, D., Maskin, E., 1994. The folk theorem with imperfect public information. *Econometrica* 62, 997–1040.
- Fudenberg, D., Maskin, E., 1986. The folk theorem for repeated games with discounting or with incomplete information. *Econometrica* 54, 533–554.
- Fudenberg, D., Yamamoto, Y., 2011. The folk theorem for irreducible stochastic games with imperfect public monitoring. *Journal of Economic Theory* 146, 1664–1683.
- Horner, J., Sugaya, T., Takahashi, S., Vieille, N., 2011. Recursive methods in discounted stochastic games: An algorithm for $\delta \rightarrow 1$ and a folk theorem. *Econometrica* 79, 1277–1318.
- McKelvey, R., 1976. Intransitivities in multidimensional voting models and some implications for agenda control. *Journal of Economic Theory* 12, 472–482.
- McKelvey, R., 1979. General conditions for global intransitivities in formal voting models. *Econometrica* 47, 1085–1112.
- Plott, C., 1967. A notion of equilibrium and its possibility under majority rule. *American Economic Review* 57, 787–806.
- Schofield, N., 1978. Instability of simple dynamic games. *Review of Economic Studies* 45, 575–594.
- Schofield, N., 1983. Generic instability of majority rule. *Review of Economic Studies* 50, 695–705.