Political Science Research and Methods Vol 2, No. 2, 213–242 October 2014
© The European Political Science Association, 2014
doi:10.1017/psrm.2014.4

A Folk Theorem for Repeated Elections with Adverse Selection*

JOHN DUGGAN

This article establishes a folk theorem for a model of repeated elections with adverse selection: when citizens (voters and politicians) 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. An extension of the arguments allows policy paths to depend on the office holder's type, subject to incentive compatibility constraints.

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 article 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. In equilibrium, therefore, 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

* John Duggan is Professor Political Science Economics, W. Allen Wallis Institute of Political Economy, University of Rochester, Rochester, NY 14627 (dugg@ur.rochester.edu). This article was presented at the Second Warwick Political Economy Conference in Venice, 2013. I am grateful for discussions with Paulo Barelli and Adam Meirowitz and to two anonymous referees for extremely helpful comments.

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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 (voters and politicians) 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 public goods with privately observed budget constraints (Meirowitz (2007)), parties (Bernhardt et al. (2009)), valence (Bernhardt, Camara and Squintani (2011)) and taxation (Camara (2012)).[1]

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,' which 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. In fact, the theorem extends to support any incentive-compatible assignment of policy paths to types. The results do 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

[1] The current article contains the one-dimensional version of the Banks and Duggan (2008) model as a special case; Duggan (2000) and later applications assume a continuous distribution of types, whereas the current article assumes a finite set of types to simplify the analysis.

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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 holders 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 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 politician 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—and necessary for generating empirically meaningful predictions—for without some limitation on the extent of conditioning on histories (or an *a priori* bound on the rate of discount), almost any policy path is consistent with equilibrium. 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 article 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 article considers the general multidimensional model, and Theorem 1 of the article 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 article: 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, Levine and Maskin (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

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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 articles assume finite action sets, whereas the electoral model allows office holders to choose from a convex policy space. More importantly, the above articles assume that payoff functions are complete information, precluding the possibility of adverse selection and the difficult problems it entails.[2]

As perfect Bayesian electoral 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.[3] I leave open the question of the effect of stronger refinements on policies supported by equilibrium strategies and beliefs.

The model analyzed in this article 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 article.

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

[2] Chapter 5 of Aumann and Maschler (1995) considers infinitely repeated, two-player, finite-action games in which Player 1 has two possible types. Their Theorem 6.1 gives a partial characterization of equilibrium payoffs under the assumptions that: the informed player can make a cheap talk announcement prior to the game, that the players can use correlated strategies in stage games, and that deviations are evaluated according to the overtaking criterion.

[3] Traditional refinements are not easily applied in the current framework. Because the electoral game assumes infinite action sets for politicians, consistency of beliefs obtained through perturbed games—and therefore sequential equilibrium—are not easily defined. Because the game is multiplayer and infinite horizon, belief refinements from signaling games are also not easily adapted. The equilibrium construction in Step 1 of the Appendix satisfies the intuitive criterion in spirit, as only type τ_1 office holders receive their maximum payoff in any continuation game, so that deviations can be addressed by specifying that voters believe the politician is type τ_0. As well, D1 appears to have limited bite, because voters have only two actions (vote for the incumbent or the challenger) following the observation of an office holder's policy choice.

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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 (voters and politicians) is partitioned into a finite set T of types, denoted $\tau_0, \tau_1, \ldots, \tau_n$.[4] Let $q = (q_0, q_1, \ldots, q_n)$ 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 $u_\tau : X \to \Re$ 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; that is, $x_{\tau_0} \leq x_{\tau_1} \leq \ldots \leq x_{\tau_n}$.[5] 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; that is, $0 = x_{\tau_0} < x_{\tau_1}$ and $x_{\tau_{n-1}} < x_{\tau_n} = 1$. Assume that the median voter type is not extremist; that is, $\max\{q_0, q_n\} < \frac{1}{2}$. Since utilities are continuous and concave, one-sided derivatives exist, and we let $u'_\tau(0)$ and $u'_\tau(1)$ denote, respectively, the derivatives from the right and left at the extreme policies. Finally, assume that the first-order condition holds at the ideal policies of the extreme types; that is $u'_{\tau_0}(0) = u'_{\tau_n}(1) = 0$, and note that the derivatives for all other types are non-zero; that is, $u'_\tau(0) > 0$ for all $\tau \neq \tau_0$, and $u'_\tau(1) < 0$ for all $\tau \neq \tau_n$.

Elections proceed as follows. Enumerating politicians as $M = \{c_1, c_2, \ldots\}$, politician c_1 holds a public office and selects a policy $x_1 \in X$ in the first period.[6] Each period $t = 2, 3, \ldots$ begins with an incumbent office holder, denoted i_t, and this politician selects a policy. Voters observe the policy choice, and an election is held between the incumbent and challenger c_t in a majority rule election. Voters simultaneously select ballots from $\{i_t, c_t\}$, and if the proportion of voters that votes for i_t is at least one-half, then the incumbent wins the election and becomes the incumbent in period $t+1$, so that $i_{t+1} = i_t$. Otherwise, the challenger wins the election and becomes the incumbent in period $t+1$, so that $i_{t+1} = c_t$.

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 $p = (p_0, \ldots, p_n)$. 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; that

[4] The specification of the pool M of potential candidates is merely to simplify the assumption of independent types, later, and is without loss of generality. See Banks and Duggan (2008) for details on drawing challengers from the electorate at large in a way that satisfies a finite-dimensional independence condition.

[5] The assumption of a finite set of types simplifies the analysis substantially, as two 'most extreme' types are easily identified, and, for example, I can easily separate type τ_0 from type $\tau \neq \tau_0$ politicians by having the latter choose policy x_{τ_1} and the former choose x_{τ_0}. The argument with a continuum of types would require me to separate types into two groups using a suitably chosen policy x instead of x_{τ_1}, with the two groups defined by an indifferent cutoff type. I conjecture that the equilibrium construction can be extended to the continuum model at the cost of a more complicated and lengthy analysis.

[6] If an initial election were modeled explicitly, then we would specify that the voters' beliefs about the incumbent and challenger were both equal to the prior, and so there is no need to model this election explicitly.

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

is, $\min\{p_0, p_n\} > 0$. 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. For future use, let $\hat{p}$ denote the voters' updated prior conditional on the politician's type being not equal to τ_0; that is, for $\tau \neq \tau_0$, we have $\hat{p}_\tau = p_\tau / \sum_{\tau' \neq \tau_0} p_{\tau'}$.

Let $\mathbf{x} = (x_1, x_2, \ldots) \in [0,1]^\infty$ 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 article 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 Ø, 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; that is, $c_t = i_{t+1} = \cdots = i_{t'}$. A *voting strategy* for a voter assigns a ballot, i_t or c_t, to every partial public history of length t.[7] 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 asymmetric 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 article.

The most permissive equilibrium concept consistent with dynamic incentives of citizens is that of *perfect Bayesian equilibrium*. This is a specification of policy strategies, voting strategies, and a belief assessment for all citizens such that (1) for every finite complete public history, the policy choice of the office holder maximizes her expected discounted payoff, (2) for every finite partial public history, the ballot of each voter maximizes her expected discounted payoff, and (3) 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.

[7] In 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.

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

the same for voters of the same type. Condition (f) is a standard restriction in voting games, requiring that each voter votes 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.[8]

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, relegating the formal proof to the Appendix. Let $\mathbf{x} = (x_1, x_2, \ldots) \in [0,1]^\infty$ denote a path of policies over time, and for future use, given policy $a \in X$, let $\mathbf{x}^a = (a, a, \ldots)$ 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. That is, $\mathbf{x}$ 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 $a, b \in \mathrm{X}$ with $0 < a < b < 1$, there exists $\overline{\delta} < 1$ such that for all $\delta \in (\overline{\delta}, 1)$, we have $[a,b]^\infty \subseteq P(\delta)$.

The proof of the theorem 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.

[8] The equilibrium construction does not rely on this restriction in an essential way. On the path of play, all types of office holder choose the same policy, and the equilibrium could be modified so that an office holder is replaced with positive probability (or even probability one). Off the equilibrium path, a majority of voters has a common strict preference for one candidate or the other.

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

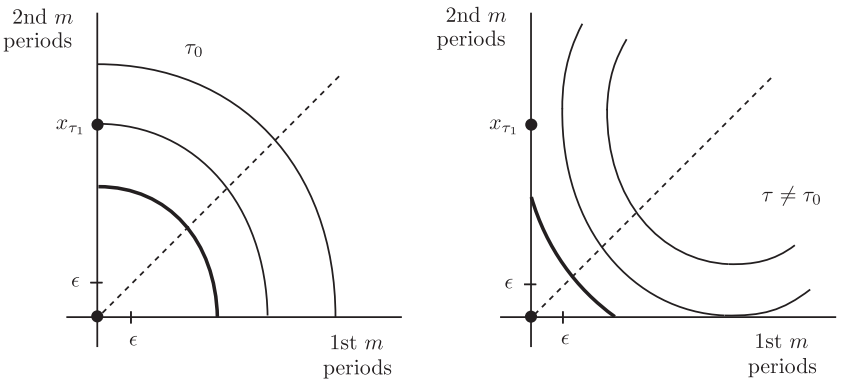


Fig. 1. Indifference curves over sequences

To gain some geometric insight into the problem, consider the politicians' payoffs (normalized by $1-\delta$) from a string of length m of policy x followed by a string of length m of policy y, namely,

$$(1-\delta)[u_\tau(x) + \cdots + \delta^{m-1}u_\tau(x) + \delta^m u_\tau(y) + \cdots + \delta^{2m-1}u_\tau(y)]$$
$$= (1-\delta^m)u_\tau(x) + \delta^m(1-\delta^m)u_\tau(y).$$

In Figure 1, I depict indifference curves over sequences

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

with the horizontal axis representing x and the vertical axis 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 a negative 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

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

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 for type τ_0 and type $\tau \neq \tau_0$ politicians. To reconcile the incentives of office holders, choosing policy a must be weakly preferable to the punishment lottery for all types, and so the point (a, 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

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

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

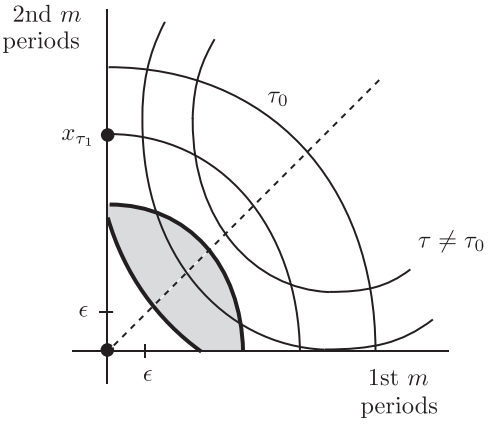


Fig. 2. Lens of incentives

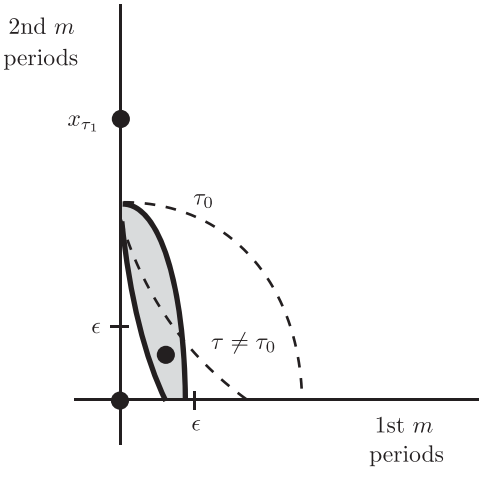


Fig. 3. Choosing a and δ^m

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 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 strategies used in Step 1 are in fact somewhat more complex than the simple description above, because (1) we must specify policies beyond the $2m$ periods following a deviation, and because (2) we must address the possibility of deviations from

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

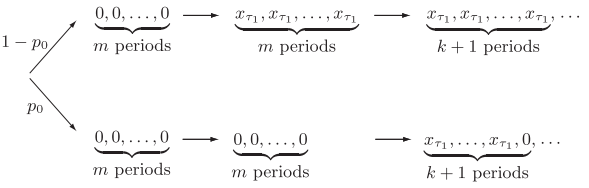


Fig. 4. Punishment paths

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 beyond the first m. The advantage of electing a challenger to such a voter is that with probability $1-p_0$, the politician will choose preferable policies after the first m terms of office. 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 chooses $x=x_{\tau_1}$, while a type τ_0 office holder alternately chooses $x=x_{\tau_1}$ for k periods and $x=0$ for one period. When k is large, therefore, 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 prefers 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 $(a,b,a,b,\ldots)$, 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 3, for example, then play in future even periods proceeds as though she had chosen $x \neq b$ in Period 2, 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

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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 has 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 $(a,a,a,a,\ldots)$ and $(a,b,a,b,\ldots)$ to produce a perfect Bayesian electoral equilibrium that supports policy choice a with frequency 0.75.

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

And we can splice equilibria supporting a with frequency 0.5 and 0.75 to produce an equilibrium supporting a with frequency 0.625.

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

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; given $\alpha \in [0,1]$, for example, we can construct an equilibrium such that in sufficiently long runs of periods, the policy outcome is a with frequency $\hat{\alpha}$ close to α. Each application imposes a higher cutoff for the discount factor, but because only a finite number of iterations is 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 $b' \in \left(\frac{b+1}{2}, 1\right)$ 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 $t+1$, namely $(x_{t+1}, x_{t+2}, \ldots)$, 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

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

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

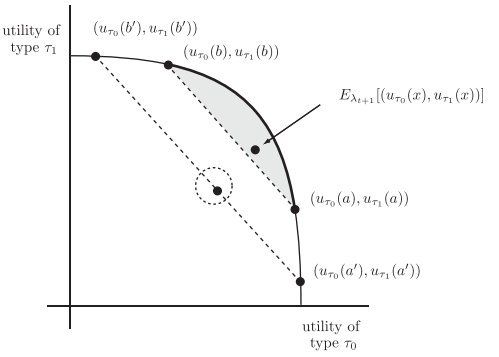


Fig. 5. Extreme punishments

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

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

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 $\hat{\alpha}$ and $1 - \hat{\alpha}$, so that the corresponding convex combination is uniformly $\frac{\eta}{2}$ less than $E_{\lambda_{t+1}}[u_\tau(x)]$, and we use the corresponding strategies from Step 3 to punish any deviation in period t.

For these punishments to deter all one-shot deviations, we must choose δ sufficiently close to one. Specifically, the punishment is effective if for all τ,

$$(1 - \delta) \max_x u_\tau(x) + \delta \left(E_{\lambda_{t+1}}[u_\tau(x)] - \frac{\eta}{2} \right) \leq (1 - \delta) \min_y u_\tau(y) + \delta E_{\lambda_{t+1}}[u_\tau(x)],$$

and letting $\Delta = \max_{\tau,x,y}[u_\tau(x) - u_\tau(y)]$ denote the maximum gain from a one-shot deviation, the above inequality holds if $\Delta \leq \frac{\delta\eta}{2(1-\delta)}$, which holds for δ close to one. Thus, 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 $\hat{\alpha}$ and $1 - \hat{\alpha}$ holds for an open set around λ_{t+1}, so the open covering generated by choice of $\hat{\alpha}$ as we vary λ 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.

TYPE-DEPENDENT PATHS

In this section, I give an informal argument for a corollary of the above theorem that extends the result to allow for policy paths to depend on the office holder's type. Let $\xi : T \to [0, 1]^\infty$ be a *path assignment* that associates to each type τ a path $\xi(\tau) = (\xi_1(\tau), \xi_2(\tau), \ldots)$ of policies. Say the path assignment ξ is *incentive compatible* if

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

the path associated with any type is no worse for that type than any other assigned path; that is, for all τ and τ',

$$\sum_{t=1}^{\infty} \delta^{t-1} u_\tau(\xi_t(\tau)) \geq \sum_{t=1}^{\infty} \delta^{t-1} u_\tau(\xi_t(\tau')).$$

Of course, if ξ is constant, so that the policy specified is independent of the office holder's type (the case considered in the theorem), then it is incentive compatible. A path assignment ξ is *supportable* given $\delta < 1$ if there is a perfect Bayesian equilibrium such that along the path of play from the initial history, in every period t, the type τ office holder chooses policy $\xi_t(\tau)$. Clearly, incentive compatibility is necessary for supportability; the following corollary establishes that when citizens are patient, it is essentially sufficient as well.

COROLLARY: For all $a, b \in X$ with $0 < a < b < 1$, there exists $\overline{\delta} < 1$ such that for all $\delta \in (\overline{\delta}, 1)$ and all path assignments ξ with $\xi(T) \subseteq [a, b]^\infty$, if ξ is incentive compatible, then it is supportable given δ.

To establish the corollary, let ξ be an incentive-compatible path assignment such that for all τ, $\xi(\tau) \in [a, b]^\infty$. To support this assignment, we specify that after a public history of length $t-1$, such that a type τ politician has held office and chosen 'correct' policies $(x_1, \ldots, x_{t-1}) = (\xi_1(\tau), \ldots, \xi_{t-1}(\tau))$, the office holder chooses $x_t = \xi_t(\tau)$ and is re-elected. Punishments are as follows. If past policy choices $x_1, \ldots, x_{t-1}$ of an office holder are consistent with ξ, so that the set

$$T(x_1, \ldots, x_{t-1}|\xi) = \{\tau \in T | (\xi_1(\tau), \ldots, \xi_{t-1}(\tau)) = (x_1, \ldots, x_{t-1})\}$$

is nonempty, but x_t is not consistent with these choices, then select any $\tilde{\tau} \in T(x_1, \ldots, x_{t-1}|\xi)$, and use the punishments constructed in the proof of the theorem for a deviation from $\xi_t(\tilde{\tau})$. If x_t is consistent with ξ but the office holder is not re-elected, then—given that voter beliefs may not place probability one on the office holder's true type (that is, $T(x_1, \ldots, x_{t-1}, x_t|\xi)$ may not be singleton)—we must specify punishments in a more nuanced way. Let $p(x_1, \ldots, x_t|\xi)$ represent the voters' posterior beliefs about the office holder's type, so that if the incumbent is re-elected, then a type τ voter's expected discounted utility is

$$\sum_{\tau'} p_{\tau'}(x_1, \ldots, x_t|\xi) \sum_{t'=t+1}^{\infty} \delta^{t'-t-1} u_\tau(\xi_{t'}(\tau')),$$

which is the expected utility with respect to a particular lottery $\lambda(x_1, \ldots, x_t|\xi)$. By the arguments of Step 4, we may specify that if the incumbent is not re-elected, then punishment strategies alternate between a' and b' at a frequency that yields a lower discounted utility for all voter types.

Choosing $\overline{\delta}$ sufficiently close to one to offset the one-period gains from any deviation, the above punishment strategies support the path assignment ξ. In particular, consider the optimization problem of a type τ office holder in period t, assuming that she has chosen the correct policies $(x_1, \ldots, x_{t-1}) = (\xi_1(\tau), \ldots, \xi_{t-1}(\tau))$ in all previous periods. The strategies above specify that she choose $x_t = \xi_t(\tau)$ in period t and continue to follow the path $\xi_{t+1}(\tau), \xi_{t+2}(\tau), \ldots$ thereafter. I claim that the politician cannot gain from deviating to the policies assigned to some other type $\tau' \in T(x_1, \ldots, x_{t-1}|\xi)$. Indeed, consider any

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

$\tau' \in T(x_1, \ldots, x_{t-1}|\xi)$, and let $\xi(\tau) = (x_1, x_2, \ldots)$ and $\xi(\tau') = (x'_1, x'_2, \ldots)$ denote the policies assigned to types τ and τ'. Note that

$$\sum_{t'=t}^{\infty} \delta^{t'-t} u_\tau(x_{t'}) \geq \sum_{t'=t}^{\infty} \delta^{t'-t} u_\tau(x'_{t'})$$

holds, in light of the fact that $(x'_1, \ldots, x'_{t-1}) = (x_1, \ldots, x_{t-1})$, if and only if

$$\sum_{t'=1}^{\infty} \delta^{t'-t} u_\tau(x_{t'}) \geq \sum_{t'=1}^{\infty} \delta^{t'-t} u_\tau(x'_{t'}),$$

which holds by incentive compatibility, establishing the claim. Moreover, it cannot be profitable to follow the policies assigned to $\tau' \in T(x_1, \ldots, x_{t-1}|\xi)$ and deviate at a later point t', as the punishment for deviating from $\xi_{t'}(\tau')$ is worse than following $\xi(\tau')$ for all types. Recall that if the office holder deviates to a policy x_t' that is inconsistent with ξ, then the strategies above invoke punishments for deviating from the policies of an arbitrary type $\tilde{\tau} \in T(x_1, \ldots, x_{t-1}|\xi)$. Since it is not profitable to follow the policies assigned to $\tilde{\tau}$, and since the punishment for deviating from $\xi_t(\tilde{\tau})$ is worse than following $\xi(\tilde{\tau})$ for all types, it cannot be profitable to deviate to policies that are inconsistent with ξ. We conclude that the office holder's expected discounted utility is maximized by following the path assignment.

CONCLUSION

This article shows that 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. Thus in the absence of an *a priori* bound on discount factors, some form of stationarity that limits the extent of conditioning on histories is needed to generate empirically meaningful predictions. 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 beyond the concept of perfect Bayesian electoral equilibrium considered in this article.

REFERENCES

Aumann, R., and M. Maschler. 1995. *Repeated Games with Incomplete Information*. Cambridge, MA: MIT Press.

Banks, J., and J. Duggan. 2008. 'A Dynamic Model of Democratic Elections in Multidimensional Policy Spaces'. *Quarterly Journal of Political Science* 3:269–99.

Banks, J., and R. Sundaram. 1993. 'Adverse Selection and Moral Hazard in a Repeated Election Model'. In *Political Economy: Institutions, Information, Competition, and Representation*, edited by William Barnett et al. Cambridge University Press.

——. 1998. 'Optimal Retention in Agency Problems'. *Journal of Economic Theory* 82:293–323.

Bernhardt, D., O. Camara, and F. Squintani. 2011. 'Competence and Ideology'. *Review of Economic Studies* 78:487–522.

Bernhardt, D., L. Campuzano, F. Squintani, and O. Camara. 2009. 'On the Benefits of Party Competition'. *Games and Economic Behavior* 66:685–707.

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

Bernhardt, D., S. Dubey, and E. Hughson. 2004. 'Term Limits and Pork Barrel Politics'. *Journal of Public Economics* 88:2383–422.
Besley, T., and S. Coate. 1997. 'An Economic Model of Representative Democracy'. *Quarterly Journal of Economics* 112:85–114.
Camaro, O. 2012. 'Economic Policies of Heterogeneous Politicians'. Working paper, Marshall School of Business, University of Southern California, available at http://www-bcf.usc.edu/~ocamara/taxes.pdf, accessed 6 May 2014.
Duggan, J. 2000. 'Repeated Elections with Asymmetric Information'. *Economics and Politics* 12:109–36.
Duggan, J., and M. Fey. 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., D. Levine, and E. Maskin. 1994. 'The Folk Theorem with Imperfect Public Information'. *Econometrica* 62:997–1040.
Fudenberg, D., and E. Maskin. 1986. 'The Folk Theorem for Repeated Games with Discounting or with Incomplete Information'. *Econometrica* 54:533–54.
Fudenberg, D., and Y. Yamamoto. 2011. 'The Folk Theorem for Irreducible Stochastic Games with Imperfect Public Monitoring'. *Journal of Economic Theory* 146:1664–83.
Horner, J., T. Sugaya, S. Takahashi, and N. Vieille. 2011. 'Recursive Methods in Discounted Stochastic Games: An Algorithm for $\delta \to 1$ and a Folk Theorem'. *Econometrica* 79:1277–318.
McKelvey, R. 1976. 'Intransitivities in Multidimensional Voting Models and some Implications for Agenda Control'. *Journal of Economic Theory* 12:472–82.
——. 1979. 'General Conditions for Global Intransitivities in Formal Voting Models'. *Econometrica* 47:1085–112.
Meirowitz, A. 2007. 'Probabilistic Voting and Accountability in Elections with Uncertain Policy Constraints'. *Journal of Public Economic Theory* 9:41–68.
Osborne, M., and A. Slivinski. 1996. 'A Model of Competition with Citizen Candidates'. *Quarterly Journal of Economics* 111:65–96.
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* 45:575–94.
——. 1983. 'Generic Instability of Majority Rule. *Review of Economic Studies* 50:695–705.
Wen, Q. 2002. 'A Folk Theorem for Repeated Sequential Games'. *Review of Economic Studies* 69:493–512.

APPENDIX

Proof of Theorem

The proof proceeds in four steps.

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

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

$$\begin{aligned}(1-p_0)\left[u_\tau(x_{\tau_1}) - \tfrac{k u_\tau(x_{\tau_1}) + u_\tau(0)}{k+1}\right] \\ < p_0\left[\tfrac{k u_\tau(x_{\tau_1}) + u_\tau(0)}{k+1} - u_\tau(0)\right].\end{aligned} \tag{1}$$

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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'Hôpital's rule, we have for all τ,

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

and in addition that

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

Using Inequalities 1–3, there exists $\overline{\beta} \in (0,1)$ such that two conditions hold: (1) for all $\beta \in (0, \overline{\beta})$, we have

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

and (2) there exists $\underline{\delta} \in (0,1)$ such that for all $\delta \in (\underline{\delta}, 1)$ and all $\beta \in (0, \overline{\beta})$, we have for all $\tau \neq \tau_0$,

$$(1-p_0)(u_\tau(x_{\tau_1}) - W_\tau) < p_0(1-\beta)(W_\tau - u_\tau(0)). \tag{5}$$

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

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

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

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

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

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 $\beta \in (0, \overline{\beta})$ and for sufficiently small $a \in (0, \varepsilon)$, define

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

where these 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, $\underline{A} \to 0$ and $A^* \to \max_{\tau \neq \tau_0}[u_\tau(x_{\tau_1}) - u_\tau(0)] > 0$ as $a \to 0$. Assuming without loss of generality that $a < x_{\tau_1}$, the right-hand side of Inequality 6 is positive if

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

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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

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

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 $a \in (0, \varepsilon)$ and $\beta \in (0, \overline{\beta})$ such that

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

automatically fulfilling Inequality 8, and such that for all $\tau \neq \tau_0$,

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

Combining Inequalities 10 and 11 and using Inequality 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 Inequality 9, as required. It follows that for these choices of a and β, the right-hand sides of Inequality 6 and 7 are positive, and we can therefore choose $\delta \in (\underline{\delta}, 1)$ sufficiently close to one to satisfy Inequalities 6 and 7.

Fixing a as above, note that there is some latitude in the choice of β, as Inequalities 8 and 9 involve strict inequalities. Next, I exploit this flexibility to choose $\hat{\delta}$ close to one and fulfill the statement of Step 1. To do so, let $\beta_1, \beta_2 \in (0, \overline{\beta})$ 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 in addition to Inequality 5, so that we have for all $\tau \neq \tau_0$,

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

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

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

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 Inequality 14 is used to address type τ_0 office holders, in the punishment phases.

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

Next, I determine a number m of periods to be used in punishment phases, as a function of $\delta \in (\hat{\delta}, 1)$. Using the inequality $\hat{\delta}\beta_2 > \beta_1$, I claim that for each $\delta > \hat{\delta}$, there exists 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 $\delta^m = \delta\delta^{m-1} \geq \delta\beta_2 > \beta_1$, so that $\delta^m \in (\beta_1, \beta_2)$, as claimed. The remainder of the construction takes such m as given.

Now, given discount factor $\delta > \hat{\delta}$ and m as above, I specify strategies and beliefs that support the constant path $\mathbf{x}^a = (a, a, \ldots)$. Note that by the preceding argument, we can choose m so that Inequalities 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 $j \in \{1, \ldots, k\}$, they choose $x = 0$ if $j = k+1$, and type $\tau \neq \tau_0$ office holders choose $x = x_{\tau_1}$ for all $j = 1, \ldots, k+1$. All voters believe that with probability one the incumbent is type τ_0. If $j \in \{1, \ldots, k\}$ 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.

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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 Ø 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 $j \in \{1, \ldots, k\}$ and the incumbent chooses $x = x_{\tau_1}$, then move to Phase 4.0(j + 1); if $j \in \{1, \ldots, k\}$ 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 $j \in \{1, \ldots, k\}$ 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 $\tau \neq \tau_0$, then 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*. In Phase 5, all types of office holder choose their ideal policies. When this phase is reached, voters believe the office holder is type τ_0, the politician is removed from office, and we move to Phase 2 with the challenger assuming office.

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

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

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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,

$$(1-\delta)(u_{\tau_0}(0) - u_{\tau_0}(a))$$

$$\leq \delta[p_0[(1-\delta^{2m})(u_{\tau_0}(a) - u_{\tau_0}(0)) + \delta^{2m}(u_{\tau_0}(a) - W_{\tau_0})]$$

$$+(1-p_0)[(1-\delta^{m})(u_{\tau_0}(a) - u_{\tau_0}(0)) + \delta^{m}(u_{\tau_0}(a) - u_{\tau_0}(x_{\tau_1}))]],$$

which after substituting $\beta = \delta^m$ is equivalent to Inequality 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

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

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

which after substituting $\beta = \delta^m$ is equivalent to Inequality 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.

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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

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

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

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

which is implied by

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

which after substituting $\beta = \delta^m$ is implied by Inequality 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

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

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 Inequality 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

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

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

which is implied by

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

which is implied by Inequality 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) with a challenger, leading to a lower expected discounted payoff.[9] 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

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

This inequality is most restrictive when $j = 1$ and actually holds with $j = 0$. Indeed, substituting $\beta = \delta^m$, the inequality then becomes

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

which is equivalent to Inequality 5 after manipulation.[10] 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

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

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

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

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

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

[9] For 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 $\{\tau_1, \ldots, \tau_n\}$. What is required is merely that beliefs put probability zero on τ_0.

[10] Since this argument shows that Inequality 5 is sufficient for the more restrictive inequality with $j = 0$, 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).

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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 Inequality 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 Inequality 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 $j \in \{1, \ldots, k\}$ 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 Inequality 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

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

where the second inequality above follows by setting $j = m-1$ in Inequality 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 best one-shot deviation is as in Phase 3, and it follows from Inequality 13 and the above arguments that one-shot deviations are unprofitable.

Now consider a type $\tau \geq \tau_1$ 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 Inequality 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 2(1). If the challenger is elected, then we move to Phase 4.2. By the arguments for Phase 3, 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

$$\begin{aligned}
(1-\delta)u_{\tau_0}(0) + \delta W_{\tau_0} > (1-\delta)u_{\tau_0}(x_{\tau_1}) + (1-\delta)\delta u_{\tau_0}(0) \\
+ \delta^2[p_0[(1-\delta)u_{\tau_0}(0) + \delta W_{\tau_0}] + (1-p_0)u_{\tau_0}(x_{\tau_1})]
\end{aligned}$$

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

Using Inequality 14 and the arguments of Phase 3, the above inequality is implied by

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

Since $W_{\tau_0} > u_{\tau_0}(x_{\tau_1})$, this in turn holds if $u_{\tau_0}(0) > u_{\tau_0}(x_{\tau_1})$, which holds by assumption. If the politician deviates to $x' \notin \{0, x_{\tau_1}\}$, then she is removed from office and we remain in Phase 4.2, which is not profitable. 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 best one-shot deviation is as in Phase 3, and it follows from Inequality 13 and the above arguments that one-shot deviations are unprofitable.

Now consider a type $\tau \geq \tau_1$ voter. In case the office holder 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), and again Inequality 15 implies that the voter's expected discounted payoff from re-electing the incumbent exceeds that from electing a challenger. In case the incumbent chooses $x = 0$, 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 4.0(1), and if the incumbent is not re-elected, then we move to Phase 2(1). Again, Inequality 15 implies that the voter's expected discounted payoff from re-electing the incumbent exceeds that from electing a challenger. In case the office holder chooses $x \notin \{0, x_{\tau_1}\}$, then all voters believe 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 remain in Phase 4.2. By the arguments for Phase 3, the expected discounted payoff from electing the challenger exceeds that from re-electing the incumbent.

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 $\varepsilon > 0$, we can choose $a \in (0, \varepsilon)$, natural number k, and $\hat{\delta} \in (0,1)$ such that for all $\delta \in (\hat{\delta}, 1)$, there exists m such that Inequalities 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 $\varepsilon > 0$, there exist $b \in (1-\varepsilon, 1)$ 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 $\tau \neq \tau_0$, Inequality 1 holds, and such that for all $\tau \neq \tau_n$, the corresponding inequality for Step 2 holds. Choose $\overline{\beta}$ such that for all $\beta \in (0, \overline{\beta})$, Inequality 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 $\beta \in (0, \overline{\beta})$, we have Inequality 5 for all $\tau \neq \tau_0$ and the corresponding inequality for Step 2 for all $\tau \neq \tau_n$. Furthermore, define

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

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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

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

Thus we can choose $a\in(0,\varepsilon)$, $b\in(1-\varepsilon,1)$ and $\beta\in(0,\overline{\beta})$ such that Inequalities 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 Inequalities 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 $\varepsilon>0$, there exist $a\in(0,\varepsilon)$ and $b\in(1-\varepsilon,1)$ such that for all natural numbers ℓ and all k with $k=0,1,\ldots,2^\ell$, 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; that is, for all $m=0,1,2,\ldots$,

$$\#\{t|m2^\ell+1\le t\le(m+1)2^\ell \text{ and } x_t=a\}=k$$
$$\#\{t|m2^\ell+1\le t\le(m+1)2^\ell \text{ and } x_t=b\}=2^\ell-k.$$

The proof is by induction. Fix $\varepsilon>0$. By Step 1, there exist $a\in(0,\varepsilon)$ 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 $b\in(1-\varepsilon,1)$ 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 $k=1$, and set $\mathbf{x}=(b,a,b,a,\ldots)$. I focus on the former case without loss of generality. Set $\tilde{\delta}=\sqrt{\max\{\delta',\delta''\}}$, choose $\delta\in(\tilde{\delta},1)$, 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 $(a,b,a,b,\ldots)$, and we specify strategies from Step 1 in odd periods and from Step 2 in even periods. Following any deviation in an odd period t, for example, the transition between phases in subsequent odd periods follows the protocol from Step 1 (applied to odd periods), and we specify that the phase transitions in the next even period $t+1$ as though an analogous deviation had occurred in period $t-1$; from that point on, the transition between phases in even periods follows the protocol from Step 2 (applied to even periods). So, if along the equilibrium path the office holder deviates to $x'\neq a$ in an odd period t, then the politician is removed from office, and we move to Phase 2(1) for the next m even periods (with actions determined by Step 2) and for the next m odd periods (with actions determined by Step 1). After that span of $2m$ periods has elapsed, we move to Phase 3 for the next m even periods and for the next m odd periods. We then enter Phase 4.0 or Phase 4.1, depending on policy choices in odd and even periods. Subsequent deviations are treated the same way. Suppose in an odd period t', we are in Phase 3, and the office holder chooses $x'\notin\{0,x_{\tau_1}\}$ and is not re-elected. Then we transition to Phase 4.2 in the next even period $t'+1$ and the next odd period $t'+2$. If the newly elected office holder is type τ_n, for example, then she chooses $x=x_{\tau_n}$ in period $t'+1$, we transition to Phase 4.0 in all subsequent even periods, she chooses $x=x_{\tau_1}$ in period $t'+2$, and we transition to Phase 4.1 in all subsequent odd periods. On the other hand, if the new office holder is type τ with $\tau_1\le\tau\le\tau_{n-1}$, then she chooses $x_{\tau_{n-1}}$ in period $t'+1$ and x_{τ_1} in period $t'+2$, and we transition to Phase 4.1 in all subsequent periods.

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. To simplify exposition, we focus on the case in which the transition first occurs in an odd period, so that information is revealed to voters in the next even period. 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 $\overline{x}\in\{0,1\}$ be the policy

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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

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

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 prefers 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 10, 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 $\{\tau_1, \ldots, \tau_n\}$, and if she is re-elected, then we move to Phase 3(1) in the following even period. As noted in footnote 9, this revelation of information does not affect the incentives of 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 $x = x_{\tau_{n-1}}$ in period $t+1$, then beliefs are unchanged, and if she is re-elected, then we move to Phase 3(2) in the next odd period $t+2$; again, footnote 10 implies that the incumbent will be re-elected. If voter beliefs are that the incumbent is type τ_0 and she chooses $x = 1$, 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 $\{\tau_1, \ldots, \tau_{n-1}\}$, and if she is re-elected, then we move to Phase 3(2) in the next odd period $t+2$; and if voter beliefs are $\hat{p}$ and the incumbent chooses $x = 1$, then voters believe she is type τ_n, and if she is re-elected, then we move to Phase 3(2). Again, footnote 9 implies that the incumbent will be re-elected following either choice. And if the office holder chooses $x \notin \{x_{\tau_{n-1}}, 1\}$, then voters believe she is type $\underline{\tau}$, and the politician is removed from office. Supposing the incumbent is re-elected in $t+1$, voter beliefs are again given by Bayes' rule, as applicable.

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 4.2 in the following even period period $t+1$. 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 also move to Phase 4.2 in the following even period $t+1$. Again, in either case, type $\tau \geq \tau_1$ voters will avoid transitioning to Phase 2(1) by re-electing the incumbent in period t. Then 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 $x = x_{\tau_{n-1}}$ in period $t+1$, then beliefs are unchanged, and when she is re-elected, we move to Phase 4.0 in all subsequent odd periods and to Phase 4.1 in all subsequent even periods, until a further deviation occurs. If voter beliefs are that the incumbent is type τ_0 and she chooses $x = 1$, 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 $\{\tau_1, \ldots, \tau_{n-1}\}$, and when she is re-elected, we move to Phase 4.1 in all subsequent odd periods and to Phase 4.0 in all subsequent even periods, until a further deviation occurs. If voter beliefs are $\hat{p}$ and she chooses $x = 1$, then voters believe that she is type τ_n, and when she is re-elected, we move to Phase 4.1 in all subsequent odd periods and to Phase 4.0 in all subsequent even periods. And if the office holder chooses $x \notin \{x_{\tau_{n-1}}, 1\}$, then voters believe she is type $\underline{\tau}$, and she is removed from office. Supposing the incumbent is re-elected in $t+1$, voter beliefs are again updated via Bayes' rule, as applicable.

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 $(a, b, a, b, \ldots)$ and provides a splicing procedure that can be iterated generally.

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

For the induction step, assume the claim is true for $\ell-1$, choose any $k=0,1,\ldots,2^{\ell}$, 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 $m=0,1,2,\ldots$,

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

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 $m=0,1,2,\ldots$,

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

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 $\mathbf{x}=(x_1,x_2,\ldots)=(y_1,z_1,y_2,z_2,\ldots)$ 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 $m=0,1,2,\ldots$, we have

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

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 $\overline{\delta}\in(0,1)$ such that for all $\delta\in(\overline{\delta},1)$, we have $[a,b]^{\infty}\subseteq P(\delta)$.

Consider $a,b\in X$ with $0<a<b<1$. Fix $\varepsilon=\frac{1}{2}\min\{a,1-b\}$, and let $a'\in(0,\varepsilon)$ and $b'\in(1-\varepsilon,1)$ 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

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

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

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 $k = 0, 1, \ldots, 2^\ell$. Thus, there exist k and ℓ such that

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

Define $\Lambda^{k,\ell}$ as the set of probability measures $\lambda \in \Lambda$ such that Inequality 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 $k \in \{0, 1, 2, \ldots, 2^\ell\}$, 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, …,

$$\#\{t | m2^\ell + 1 \le t \le (m+1)2^\ell \text{ and } x_t = a'\} = k$$
$$\#\{t | m2^\ell + 1 \le t \le (m+1)2^\ell \text{ and } x_t = b'\} = 2^\ell - k.$$

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'Hôpital's rule, we have

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

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

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

Since $\Lambda \subseteq \bigcup_{(k,\ell)\in F}\Lambda^{k,\ell}$ and $(1-\delta)U_\tau^{k,\ell} \to \left(\frac{k}{2^\ell}\right)u_\tau(a') + \left(1-\frac{k}{2^\ell}\right)u_\tau(b')$, it follows from Inequality 16 that for each $\lambda \in \Lambda$, we have $\lambda \in \Lambda^\delta$ for δ sufficiently close to one. That is, $\{\Lambda^\delta | \delta \in (0,1)\}$ 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 Inequality 17 holds. Furthermore, choose $\overline{\delta} \in (\delta', 1)$ such that

$$\Delta \le \left(\frac{\overline{\delta}}{1-\overline{\delta}}\right)\left(\frac{\eta}{4}\right). \tag{18}$$

Finally, consider any $\delta \in (\overline{\delta}, 1)$ and any $\mathbf{x} \in [a,b]^\infty$. I specify a perfect Bayesian equilibrium that generates $\mathbf{x}$ from the initial history by again defining actions and beliefs in each of a number of phases and by specifying a rule for transitioning between phases. Now, however, the procedure is simpler. Phase 0 will represent the path of play from the initial history, while a countable number of phases, t=1, 2, …, are used as punishments. Specifically, Phase t is used to punish deviations from

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press

Phase 0 in period t, while further deviations from phase t are handled implicitly in Step 3. Note that for each t, we can write the discounted utility of a type τ citizen from $(x_t, x_{t+1}, x_{t+2}, \ldots)$ as the integral with respect to a probability measure $\lambda_t \in \Lambda$; that is,

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

where λ_t is defined by

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

for all measurable $Y \subseteq [0,1]$. It follows that there exists $(k,\ell) \in F$ such that Inequality 17 holds with $\lambda = \lambda_t$. To define the equilibrium, we specify that in period t of Phase 0, all types of office holder choose x_t, and voters' beliefs along the path of play are given by the prior p, and all voter types vote to re-elect the incumbent. If the office holder chooses $x' \neq x_t$, then voter beliefs are again given by the prior. In Phase $t = 1, 2, \ldots$, let $(k,\ell) \in F$ satisfy Inequality 17 with $\lambda = \lambda_{t+1}$, and use Step 3 to specify strategies and beliefs so that they form a perfect Bayesian electoral equilibrium that generates the path $\mathbf{x}^{k,\ell}$ beginning in period $t+1$; that is, the path beginning in period $t+1$ is $(x_1^{k,\ell}, x_2^{k,\ell}, \ldots)$. The initial history is labeled 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 $t+1$ occurs in the next period $t+1$.

To verify that this specification constitutes a perfect Bayesian electoral 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,

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

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.

https://doi.org/10.1017/psrm.2014.4 Published online by Cambridge University Press