Rochester Center for
Economic Research

Adverse Selection and Moral Hazard in a Repeated Elections Model

Banks, Jeffrey S. and Rangarajan K. Sundaram

Working Paper No. 283
June 1991

University of
Rochester

Adverse Selection and Moral Hazard in a Repeated Elections Model

Jeffrey S. Banks and Rangarajan K. Sundaram

Rochester Center for Economic Research
Working Paper No. 283

June 1991

Adverse Selection and Moral Hazard in a Repeated Elections Model*

Jeffrey S. Banks

and

Rangarajan K. Sundaram

University of Rochester

and

The Rochester Center for Economic Research
Working Paper No. 283

June 1991

* Prepared for presentation at the Conference on Political Economy, Washington University at St. Louis, May 22–25, 1991. Financial support from the National Science Foundation and Sloan Foundation to the first author is gratefully acknowledged.

1. Introduction

In this paper we consider the problem faced by a (median) voter in an electorate who must, in each period of an infinite horizon, select a candidate for the performance of some task, where we refer to the candidate selected in period t as the period t incumbent. Rewards accrue to the voter as a consequence of her choice, where we can think of these rewards as government–controlled benefits secured by the incumbent, and the voter's objective is to maximize the discounted sum of rewards over the infinite horizon. However, a factor determining the distribution of rewards in any given period – namely, the incumbent's action that period – is unobservable to the voter. Higher actions are preferable by the voter, yet are associated with higher costs for the incumbent. Moreover, the voter is uncertain about a parameter describing these costs to a candidate, where this is labeled a candidate's "type". While the voter does not care about a candidate's type *per se*, learning this information will assist the voter in ascertaining the desirability of a candidate, since different types will take different actions. There are an infinite number of candidates available to the voter (so, in particular, at every point in time there is at least one untried candidate), candidates are all *ex ante* identical and are infinitely lived, and there are no restrictions on how often a candidate can be in office. Finally, each candidate attempts to maximize his own discounted sum of payoffs over the infinite horizon through his action choices while in office.

The voter's choice problem thus includes aspects of ***moral hazard***, in that she must provide incentives for a candidate to take costly actions, as well as ***adverse selection***, in that the voter would like to only choose those candidates who take the higher actions. Thus, to a considerable degree, the moving parts of this model are the same as those that arise frequently in the Principal–Agent literature and in other contracting models of economic theory. On the other hand, there are substantial differences as well, the foremost being the presumed inability of the voter and the candidates to sign contracts determining payments to a candidate as a function of the rewards the latter generates, thereby ruling out equilibria of the form that are standard in economic models. Rather, the voter must rely on the only policy instrument at her disposal – the control of re–election rules – to provide candidate incentives that mitigate the voter's moral hazard and adverse selection problems. Indeed, since re–election rules offer the elected candidates the ***only*** incentives to take desired actions, repetition of the voter–candidate relationship is the *sine qua non* for providing such incentives, in contrast to Principal–Agent models.

Most previous research on repeated elections in the presence of informational asymmetries have studied the voter's decision problem from either the moral hazard or the adverse selection perspective; examples of the former include Barro (1973), Ferejohn (1986), and Austen–Smith and Banks (1989), and of the latter Rogoff (1990), Reed (1991), and Banks and Sundaram (1990).[1] For example, in Ferejohn (1986) the voter knows the preferences of the candidates with certainty, but only observes the action choice by the incumbent with noise. The incumbent observes this noise term prior to taking his action, and the voter selects a re–election rule to provide incentives for incumbents to take costly actions. Conversely, Rogoff's (1990) model of political budget cycles is more in the adverse selection vein, where candidate "types" denote competency at delivering public goods to the voters.[2] In contrast to our model, the actions taken by candidates while in office are observable by the voter, and hence act as a signal of candidate competency. In addition, a candidate's competency varies over time, so in particular competency is uncorrelated across electoral cycles. Thus the issue of the voter learning about candidates through

---

[1]Other models of repeated elections include Alesina (1988) and Alesina and Spear (1988) on credible policy pronouncements, Ledyard (1989) on the transmission of information between candidates regarding voter preferences, and McKelvey and Reizman (1990) on the observance of seniority in legislatures.

[2]See also Rogoff and Sibert (1988).

their performance, which forms a crux of the current model, is avoided.

We analyze the interaction among the voter and the candidates described above as a (stochastic) game of incomplete information, and characterize a particular class of sequential equilibria. In these, the voter employs a simple ***retrospective voting rule*** (Fiorina, 1981): retain the current incumbent as long as rewards remain above a certain level. Faced with this re–election rule, incumbents adopt time–invariant actions as functions of their true type, where lower cost types take higher actions, and are consequently re–elected more often. Thus, as an incumbent's tenure increases the voter is more confident she has selected a (relatively) hard–working type, since the voter's belief about the incumbent is placing greater weight on higher types after every re–election of the incumbent. This behavior also implies that, from the voter's perspective, an incumbent's probability of re–election is a strictly increasing function of his tenure in office.

2. The Model

An individual, whom we refer to as the ***voter***, has a task to be performed in time periods $t = 1,2,...$ In each period the voter selects a single candidate for the performance of this task, where we let $N = \{1,2,...\}$ denote the (infinite) set of available candidates. The chosen candidate, whom we refer to as the period t ***incumbent***, selects an action $a \in A \equiv [\underline{a},\overline{a}] \subset \mathbb{R}$, where this action is unobserved by the voter, and (stochastically) determines the voter's reward for that period. Specifically, the voter's per–period reward is a realization from a continuous density $f(.;a)$, where for any $a \in A$, $\text{supp}.f(.;a) := \{r|f(r;a) > 0\}$ denotes the support of the density $f(.;a)$, $\bar{r}(a) = \int rf(r;a)dr$ denotes the expected reward, and $F(.;a)$ is the associated distribution. We make the following assumptions on $f(.;.)$:

A1 For all $a \in A$, $\text{supp}.f(.;a) = R \subseteq \mathbb{R}$;

A2 For all $a_1, a_2 \in A$, if $a_1 > a_2$ then $\bar{r}(a_1) > \bar{r}(a_2)$; and

A3 There exists $\hat{r} \in \text{int.}R$ such that for all $r \in R$, $r \geq \hat{r}$, $f(r;a_1) > f(r;a_2)$ if $a_1 > a_2$, where $a_1, a_2 \in A$.

Assumptions A1 and A2 are self-explanatory. A3 requires that the family of reward densities parametrized by the actions $a \in A$ are "stacked" beyond some point. Such a condition is met if, for example, the realized reward is given by $(a+\epsilon)$, where ϵ has a unimodal density (say, normal) on the real line.

The per-period payoff for the voter is simply equal to her realized reward, whereas the incumbent's payoff is a function of the action chosen as well as his "type" $\omega \in \{\omega_1,...,\omega_n\} \equiv \Omega$. Let $u(a,\omega)$ denote this payoff, where for all $\omega \in \Omega$, $u(\underline{a},\omega) > 0$, and $u(.)$ is continuously differentiable in a, with $u_1 \equiv \partial u/\partial a < 0$. Thus taking higher actions, which are preferred by the voter, are more costly for an incumbent. We further add Inada-type conditions to guarantee interior action choices by incumbents: $\forall\ \omega \in \Omega$, $u_1(a,\omega) \to 0$ as $a \to \underline{a}$, and $u_1(a,\omega) \to -\infty$ as $a \to \bar{a}$. We assume without loss of generality that for all $a \in A$, $u(a,\omega_1) < ... < u(a,\omega_n)$, so for any fixed action higher types receive higher per-period rewards; we also assume $u_1(a,\omega)$ is non-decreasing in ω, i.e. the marginal disutility from taking any action is lower for higher types.[3] All non-chosen candidates receive a per-period payoff of zero, regardless of type. The voter (resp. each candidate) discounts future payoffs by a factor $\delta \in [0,1)$ (resp. $\rho \in [0,1)$).

Candidate types are drawn independently from Ω according to $b^o = (b^o_1,...,b^o_n)$, where $b^o_j > 0$, $j = 1,...,n$. Each candidate knows his own type, but does not know any other candidate's type, and the voter does not know any candidate's type. Let

[3]For example, we could have $u(a,\omega) = \omega - c(a)$, where $\partial c/\partial a > 0$; these are the incumbent preferences in Ferejohn (1986) and Austen-Smith and Banks (1989).

$P(\Omega)$ denote the set of probability measures on Ω, and set $\Pi(\Omega) = \times_{i=1}^{\infty} P(\Omega)$. Thus the voter's prior belief on candidate types is given by $b_{\infty}^{o} \equiv (b^{o}(1),\ldots,b^{o}(i),\ldots) \in \Pi(\Omega)$, while candidate i's prior belief differs from the voter's only in that $b^{o}(i)$ has a 1 on his true type and 0 on all other types.

A ***history of length t***, denoted h^t, is a specification of all public events through period t, namely, the candidates chosen each period and the rewards realized. Let H^t denote the set of all possible histories of length t, and set $H^0 = \phi$. For a generic candidate i, this history is augmented by the actions taken by i in the periods (if any) where i was the incumbent. We will refer to the set of i's augmented histories by H_i^t, with common element h_i^t. Of course, if i has not figured in the history h^t, then $h_i^t = h^t$.

A ***strategy σ for the voter*** is a sequence of functions $\sigma = \{\sigma^t\}_{t=1}^{\infty}$, where for each t $\sigma^t: H^{t-1} \to N$ is a measurable map defining the candidate to be chosen in period t as a function of the history through period t–1. Let Σ denote the set of all possible strategies for the voter. A ***strategy γ_i for candidate i*** is a sequence of functions $\gamma_i = \{\gamma_i^t\}_{t=1}^{\infty}$, where for each t, $\gamma_i^t: H_i^{t-1} \times \Omega \to A$ is a measurable map specifying the action to be taken if i is the period t incumbent, as a function of the history h_i^{t-1} through period t–1 as well as candidate i's type. Thus i's choice of action in any period can in principle depend on i's previous actions while in office. However since candidate i's actions are not observable by anybody but i, and do not affect i's payoffs in subsequent periods, we can without loss of generality restrict attention to candidate strategies which condition only on the voter's selections and the subsequent rewards. Thus for all t, $h^{t-1} \in H^{t-1}$, let $\gamma_i^t(h^{t-1},\omega)$ denote the action taken by candidate i if i is the period t incumbent, conditional on the (public) history h^{t-1} and i's type. The set of all such strategies for a candidate is evidently the same as that for any other candidate; denote this common set by Γ. A

strategy profile is then a list of strategies, one for the voter and one each for the candidates, and a generic strategy profile will be denoted (σ,γ), where $\gamma = (\gamma_i)_{i\in N}$.

A ***belief system for the voter*** as a sequence of measurable maps $\mu = \{\mu^t\}_{t=1}^{\infty}$, where for each t, $\mu^t:H^{t-1} \to \Pi(\Omega)$. Thus, $\mu^t(h^{t-1})$ gives the voter's beliefs about the candidates' types after observing a history h^{t-1}. A ***belief system for candidate i*** is similarly a sequence of maps $\varphi_i = \{\varphi_i^t\}$ describing i's beliefs about candidate types; of course, for all t, $\varphi_i^t(.)$ will assign probability 1 to i's true type. A ***belief profile*** is then (μ,φ), where $\varphi = \{\varphi_i\}_{i\in N}$.

Given a strategy profile (σ,γ) and a belief $\mu^t(h^{t-1})$ at history h^{t-1}, the voter can compute her expected payoff conditional on being at h^{t-1}; denote this by $W(\sigma,\gamma,\mu;h^{t-1})$. Similarly for candidate i, we have $C_i(\sigma,\gamma,\varphi_i;h^{t-1})$. A ***sequential equilibrium*** (Kreps and Wilson, 1982) consists of a strategy profile (σ,γ), and a belief profile (μ,φ), such that 1) the strategies are sequentially rational given the beliefs, and 2) the beliefs are consistent with the strategies. Sequential rationality requires, for all t, h^{t-1}, i) $W(\sigma,\gamma,\mu;h^{t-1}) \geq W(\sigma',\gamma,\mu;h^{t-1})\ \forall\ \sigma' \in \Sigma$, and ii) $\forall\ i \in N$, $C_i(\sigma,\gamma,\varphi_i;h^{t-1}) \geq C_i(\sigma,\gamma_i',\gamma_{-i},\varphi_i;h^{t-1})\ \forall\ \gamma_i' \in \Gamma$.

Consistency requires at a minimum that for all histories h^{t-1} that are "reached" by the strategy profile (σ,γ), $\mu^t(h^{t-1})$ and $\{\varphi_i^t(h^{t-1})\}$ be derived via Bayes Rule from the strategies (σ,γ) and and the prior b_{∞}^{o}.[4] So suppose h^{t-1} is reached by (σ,γ), and let candidate i = $\sigma^t(h^{t-1})$, i.e. candidate i is the period t incumbent. Given a current belief about candidate i, $b^t(i) = (b_1^t(i),\ldots,b_n^t(i))$, an observed reward r, and a conjectured action rule $a(\omega) \equiv \gamma_i^t(h^{t-1},\omega)$ (stochastically) generating r, in equilibrium the voter (as well as the non-incumbent candidates) updates her belief about candidate i in a Bayesian fashion:

[4]Indeed, this is a requirement of the Nash equilibrium concept itself.

$$b_j^{t+1}(i) = \beta_j(b^t(i),r;a(\omega)) \equiv \frac{b_j^t(i)f(r;a(\omega_j))}{\sum_{k=1}^{n} b_k^t(i)f(r;a(\omega_k))} .$$

On the other hand, by the independence assumption on candidate types, the voter learns only about the incumbent's type; therefore for all non–incumbents $m \in N$, $b^{t+1}(m) = b^t(m)$.

With respect to out–of–equilibrium beliefs, i.e. beliefs at histories that are not reached by the profile (σ,γ), consistency requires such beliefs to be the limit of beliefs formed from completely mixed strategies by the players, where these completely mixed strategies converge to the candidate equilibrium strategies (σ,γ). However, since the reward densities have the same support regardless of a candidate's strategy or his true type (from A1), such histories occur only when the voter selects the "wrong" (according to σ) candidate. Therefore it is sufficient to only assume the ***voter*** "trembles" away from σ. But then out–of–equilibrium beliefs are completely specified by the candidates' equilibrium strategies, since in the above description, even if $i \neq \sigma^t(h^{t-1})$, i's strategy dictates what i would have done if selected, and therefore the "consistent" belief upon observing the subsequent reward is precisely that derived via Bayes' Rule if an fact i were equal to $\sigma^t(h^{t-1})$, i.e. if i were ***supposed*** to have been selected. Therefore, out–of–equilibrium beliefs are derived from (σ,γ), in particular from γ, in exactly the same fashion as equilibrium path beliefs. Thus, in what follows we suppress the belief profile in the characterization of sequential equilibria, since they follow immediately from the description of the strategy profile.

3. Simple equilibria

Clearly, the repeated nature of the elections and the ability of the voter to potentially "learn" about the true types of some or all candidates through the realized rewards foreshadow a possibly large and complex set of sequential equilibria for the above game. This being so, in what follows we focus attention on a particularly manageable class of equilibria, having the following characteristics: i) candidate i's strategy is only a function of his "personal" history with the voter, that is, the rewards i has generated and the voter's response of either retaining or replacing i as the incumbent; ii) all candidates adopt identical strategies (as functions of their personal histories and types), and the voter's strategy treats all candidates symmetrically; and iii) the voter adopts a ***no recall*** strategy, in which previously selected and discarded candidates are never again chosen. One can make a selection argument for ***only*** examining such equilibria as well, based on the structure of the interaction between the voter and the candidates. For instance, condition i) asserts that the only relevant information for a candidate is his own relationship with the voter; condition ii) seems natural given the symmetry of the game from the candidate's perspective and from the independence of candidate types, and similarly condition iii) seems natural given the infinite set of candidates, the independence of candidate types, and the (assumed) symmetry of the candidates' strategies: at every point in time, the voter has the ability to "start over" with an untried candidate, and if "starting over" is preferred to the period t incumbent at the beginning of period t+1, then it should be still preferred in any subsequent period.

We can describe such behavior in a more compact notation than that given in Section 2. Let $(r,I)^t \equiv ((r_1,I_1),\dots,(r_t,I_t)) \in [\mathbb{R}\times\{0,1\}]^t$ denote a t–period personal history of rewards and subsequent decisions to retain ($I = 1$) or replace ($I = 0$) the incumbent. Let $\alpha_i = \{\alpha_i^t\}$ be a measurable sequence of functions, where $\forall\ \omega \in \Omega$

$\alpha_i^0(\omega) \in A$, and for $t \geq 1$, $\alpha_i^t:[\mathbb{R}\times\{0,1\}]^{t-1}\times\Omega \to A$. The strategy $\gamma_i^{\alpha_i}(\omega)$ for candidate i is defined as follows: if i is selected as the period t incumbent after the public history h^{t-1}, then given h^{t-1} i can compute his personal history $(r,I)^{\tau}$, where $\tau \leq t-1$. If i is of type $\omega \in \Omega$, he takes action $\alpha_i^{\tau+1}((r,I)^{\tau},\omega)$. Thus, $\gamma_i^{\alpha_i}$ only depends on the personal history of rewards and responses generated by i, and not the personal histories generated by other candidates. The symmetry requirement on the candidates then is that $\alpha_i = \alpha_m = \alpha$ for all $i,m \in N$, and some α.

We can likewise describe the no recall condition for the voter's strategy in terms of personal histories: for each $i \in N$, there is a candidate–specific re–election rule $\nu_i = \{\nu_i^t\}$, where for all $t \geq 1$, $\nu_i^{t+1}:[\mathbb{R}\times\{0,1\}]^{t-1}\times\mathbb{R} \to \{0,1\}$. The interpretation is that $\nu_i^{t+1}((r,I)^{t-1},r)$ denotes the voter's decision to retain candidate or replace candidate i with a previously untried candidate as a function of i's personal history prior to the last period $(r,I)^{t-1}$, the voter having selected i in the previous period, and i having generated a reward $r \in \mathbb{R}$. If the voter treats the candidates symmetrically, we have $\nu_i = \nu_m \equiv \nu$, $\forall$ $i,m \in N$, some ν. The voter's strategy σ^{ν} is then defined as follows: given a public history h^{t-1}, suppose candidate i is the period t–1 incumbent, where (according to h^{t-1}) i has in addition been selected in τ previous periods. The voter retains i for period t if i's personal history is such that $\nu^{\tau+2}((r,I)^{\tau},r) = 1$, and replaces i with an untried (according to h^{t-1}) candidate otherwise. Hence σ^{ν} (in equilibrium) never recalls a previously selected and discarded candidate, and decides whether to retain or replace an incumbent based solely on the incumbent's personal history.

So we can characterize such equilibria by the pair (α,ν), where $\alpha(.)$ is the (common) candidate strategy describing the actions taken as a function of the candidate's type and personal history; and $\nu(.)$ is the voter strategy specifying when the voter replaces an incumbent with an untried candidate as a function of the incumbent's personal history.[5] Within this class of equilibria there exist some of a quite spartan form, namely, ones where (along the equilibrium path) the voter's replacement rule is only a function of the ***last*** reward generated by the incumbent, and is the same rule regardless of how long the incumbent has been in office. This will give a candidate of type ω an incentive to adopt the ***same*** action in every period in office regardless of personal history (again, along the equilibrium path). These we label ***simple equilibria***.

One type of simple equilibrium is the following: all candidates of all types adopt the lowest action, $\underline{a}$, in every period in office, and the voter always replaces the incumbent with an untried candidate. Thus, if the voter is going to replace the incumbent with probability one regardless of the realized reward, an incumbent has no incentive to take any but the lowest cost action; and if all incumbents take the lowest cost action the voter might as well simply throw out all incumbents. Note that such behavior would also constitute an equilibrium if the game had but a ***finite*** time horizon, since in the last period the incumbent will certainly choose $a = \underline{a}$ regardless of type or history, so the voter is indifferent over all candidates; if she selects from the untried candidates, then the incumbent in the penultimate period will certainly choose $a = \underline{a}$; etc.[6] We can think of such behavior, therefore, as being analogous to "one–shot" Nash behavior in repeated games. On the other hand, the model also generates simple equilibria of a more interesting nature. These we characterize in the following result.

[5]Note that we are not restricting the strategies available to the players in any way; we are merely focusing on a class of equilibria with a particular structure.

[6]However, this will not be the ***only*** equilibrium in the finite horizon game: due to the voter's indifference over candidates in the final period of such a game, she can make her selection a non–trivial function of realized rewards; cf. Austen–Smith and Banks (1989).

Proposition: There exist sequential equilibria with the following structure:

$$\nu^{t+1}((r,I)^{t-1},r_t) = \begin{cases} 1 & \text{if } r_\tau \geq r^* \ \forall\ \tau \leq t \quad \& \quad I_\tau = 1 \ \forall\ \tau \leq t-1 \\ 0 & \text{else} \end{cases},$$

$$\alpha^t((r,I)^{t-1},\omega) = \begin{cases} a^*(\omega) & \text{if } r_\tau \geq r^* \quad \& \quad I_\tau = 1 \quad \forall\ \tau \leq t-1 \\ \underline{a} & \text{else} \end{cases},$$

where $a^*(\omega)$ is strictly increasing in ω and depends on ρ, but where $a^*(.)$ and r^* do not depend on δ.

Proof. We begin by showing that along the equilibrium path a candidate's choice of best response to $\nu(.)$ may be modeled as a two-state dynamic programming problem. So fix a candidate $i \in N$ and a type $\omega \in \Omega$. Define $G = \bigcup_{t=1}^{\infty} R^t$, with generic element g, let $|g|$ denote the length of g, and let e_n denote the n-vector each of whose elements is unity. Define

$G_1 = \{g \in G \mid g \geq r^*.e_{|g|}\}$, and

$G_2 = G \setminus G_1$.

In this notation, the voter's strategy may be written $\nu: G \to \{0,1\}$, where $\nu(g) = 1$ if and only if $g \in G_1$.

Let $S = \{0,1\}$, where $s = 0$ signifies the event that i's history to date satisfies $g \in G_2$, and $s = 1$ signifies either $g \in G_1$ or $g = \phi$. If $s = 1$, the probability i will continue in state $s = 1$ after taking action $a \in A$ is evidently just the probability the action will generate a reward of at least r^*, which is $1 - F(r^*,a)$. On the other hand, $s = 0$ is an absorbing state, no matter what action is taken.

Consequently, we may define transition probabilities $Q(.|.,.)$ from $S \times A$ to S by

$$Q(1|1,a) = 1 - F(r^*,a) = 1 - Q(0|1,a), \text{ and}$$

$$Q(0|0,a) = 1 \text{ for all } a \in A.$$

Next, recall that i receives a time–invariant payoff while in office, and a zero payoff otherwise. This may be represented by $r: S \times A \to \mathbb{R}$, where

$$r(1,a) = u(a,\omega), \text{ and}$$

$$r(0,a) = 0 \text{ for all } a \in A.$$

The tuple $\{S,A,Q,r\}$ now represents a standard stationary dynamic programming problem. All the conditions of Maitra (1968) are seen to be met (in particular, A is compact), and the solution to this problem may be obtained via the Bellman equation

$$V(s) = \max_{a \in A} \{r(s,a) + \rho \int V(s')Q(ds'|s,a)\}. \tag{1}$$

Substituting for $r(.)$ and $Q(.)$, we finally obtain

$$V(1) = \max_{a \in A} \{u(a,\omega) + \rho[1-F(r^*,a)]V(1)\}, \text{ and} \tag{2}$$

$$V(0) = 0.$$

Evidently, if $s = 1$, the constant action which maximizes (2) is optimal for candidate i, while any action is optimal at $s = 0$; denote the former by $a^*(\omega)$. We can therefore suppress the dependence of V on the state, and highlight the dependence of V on the parameter ω, by rewriting (2) as

$$V(\omega) = \max_{a \in A} \{u(a,\omega) + \rho[1-F(r^*,a)V(\omega)\} . \quad (3)$$

By (3) the solution $a^*(\omega)$ solves

$$u_1(a^*,\omega) - \rho F_2(r^*,a^*)V(\omega) = 0 . \quad (4)$$

Also, $V(\omega) > 0 \ \forall \ \omega$, and is strictly increasing in ω, since (by $u(a,\omega)$ increasing in ω) higher types can emulate the behavior of lower types and receive a strictly higher payoff in every period in office. To see $a^*(\omega)$ is increasing in ω, let $a_j = a^*(\omega_j)$ for $j = 1,\dots,n$, and let $k > j$ (so $\omega_k > \omega_j$). Clearly we must have $a_k \neq a_j$, since if (4) holds at a_j for ω_j, then, by V(.) increasing in ω and $u_1(a,\omega)$ nondecreasing in ω, (4) does not hold at a_j for ω_k. For all $a \in A$ let $p(a) = 1 - F(r^*,a)$, and note that p(.) is strictly increasing on A and $p(a) > 0$ for all $a \in A$. By incentive compatibility, we have

$$u(a_k,\omega_k) + p(a_k)V(\omega_k) \geq u(a_j,\omega_k) + p(a_j)V(\omega_k) \quad (5)$$

$$u(a_j,\omega_j) + p(a_j)V(\omega_j) \geq u(a_k,\omega_j) + p(a_k)V(\omega_j) . \quad (6)$$

Subtracting the RHS of (6) from the LHS of (5), and the LHS of (6) from the RHS of (5), and rearranging terms, we get

$$[p(a_k) - p(a_j)][V(\omega_k) - V(\omega_j)] \geq$$
$$[u(a_j,\omega_k) - u(a_k,\omega_k)] - [u(a_j,\omega_j) - u(a_k,\omega_j)]. \quad (7)$$

If $a_j > a_k$, then both terms on the RHS of (7) are negative; yet by $u_1(a,\omega)$

nondecreasing in ω the first of these is (in absolute value terms) less than or equal to the second. Therefore, if $a_j > a_k$, we would have

$$[p(a_k) - p(a_j)][V(\omega_k) - V(\omega_j)] \geq 0, \tag{8}$$

which contradicts V(.) increasing in ω and p(.) increasing in a. Therefore $a_j > a_k$ is ruled out, and since we've shown $a_k \neq a_j$, we have $a_k > a_j$, thus proving $a^*(\omega)$ strictly increasing.

Out of equilibrium, i.e. if the voter has retained candidate i after observing a reward $r < r^*$ or if the voter returns to i after previously replacing him, it is clear $\alpha(.)$ is optimal since, according to $\nu(.)$, the voter will (with probability 1) replace i in the next period, and never select i again; thus choosing the lowest cost action, namely $a = \underline{a}$, is optimal. A similar reasoning holds if the voter ever recalls candidate i when i has (sometime in the past) generated a reward $r < r^*$. Therefore $\alpha(.)$ is a best response to $\nu(.)$ for any possible history.

Alternatively, if the voter has ever observed a reward of $r < r^*$ from candidate i, or has previously replaced i, then the voter will never select i again regardless of her beliefs, since according to $\alpha(.)$ i would now take action $a = \underline{a}$ if selected regardless of type, implying any untried candidate is strictly better for the voter. Thus, it remains to be shown that the voter prefers to retain an incumbent who continually generates rewards greater than r^* independent of the value of $\delta \in [0,1)$.

For notational ease, let $a_i = a^*(\omega_i)$, $i = 1,\ldots,n$. Since $a_1 < \ldots < a_n$, by A3 we know that there exists $\hat{r} \in$ int.R such that for all $r \geq \hat{r}$, $f(r;a_1) < \ldots < f(r;a_n)$. Pick r^* so that $r^* \geq \hat{r}$. Given the candidates' common strategy $\alpha(.)$, we pose the voter's problem of identifying a best response to $\alpha(.)$ as a stationary dynamic programming problem. Associate with candidate i a state $(b(i),s(i)) \in P(\Omega)\times\{0,1\}$, where as before the indicator "s" describes whether all previous rewards by a candidate have been above r* (s = 1), or if some previous reward was below r* (s = 0), and let the set of actions available to the voter be $N = \{1,\ldots,i,\ldots\}$. We can simplify notation by noting that $\alpha(.)$ is played by ***all*** candidates, and is a function only of ***personal*** history; therefore the voter's decision problem is isomorphic to one with a single candidate, and two actions available to the voter, $A = \{0,1\}$, where a = 1 denotes retaining the candidate as the incumbent, and a = 0 denotes starting the process over at the "initial" state $(\pi,1)$, where $\pi = b^{o}$.

So let $Z = P(\Omega)\times\{0,1\}$ denote the state space, and for $(b,s) \in Z$ and $r \in R$, if a = 1 define the new state by $(b,s)(r) = (\beta(b,r;a^*),\mu(s,r))$, where $\beta(.)$ denotes Bayes' Rule, and

$$\mu(s,r) = \begin{cases} 1 & \text{if } s = 1 \text{ and } r \geq r^* \\ 0 & \text{else} \end{cases} .$$

If a = 0, on the other hand, the new state is $(\pi,1)$. In state $(b,s) \in Z$ the expected reward to the voter from the action a = 1 is given by

$$\lambda(b,s,1) = \begin{cases} \bar{r}(\underline{a}) & \text{if } s = 0 \\ E(b) & \text{if } s = 1 \end{cases} ,$$

where for any $b \in P(\Omega)$, $E(b) = \sum_{j=1}^{n} b_j.\bar{r}(a_j)$; while if a = 0, $\lambda(b,s,0) = E(\pi)$.

Let V(.) be the unique fixed–point of the appropriately–defined contraction which satisfies

$$V(b,s) = \max\ \{V(\pi,1),\ LV(b,s)\}\ , \tag{9}$$

where

$$LV(b,s) = \lambda(b,s,1) + \delta\int V[\beta(b,r),\mu(s,r)]f^{b}(r)dr ,$$

and where for all $b \in P(\Omega)$, $f^{b}(r) = \sum_{j=1}^{n} b_j.f(r;a_j)$. It is easily seen that if $s = 0$, then for any $b \in P(\Omega)$ we have $V(b,s) = V(\pi,1) > LV(b,s)$. The above may therefore be written in a manner suppressing the dependence on the indicator s as

$$V(b) = \max\{V(\pi),\ E(b) + \delta[V(\pi)\int_{r<r^*} f^{b}(r)dr + \int_{r\geq r^*} V(\beta(b,r))f^{b}(r)dr]\} \quad (10)$$

Given $b,b' \in P(\Omega)$, say that b ***strongly (stochastically) dominates*** b' if $b_j/b_k \geq b'_j/b'_k$ for $j > k$, i.e. if b places relatively greater weight on higher types. It is not too difficult to see that strong stochastic dominance implies stochastic dominance; therefore, if b strongly dominates b', then for any non-negative numbers $x_1 < \ldots < x_n$, $\sum_{j=1}^{n} b_j.x_j > \sum_{j=1}^{n} b'_j.x_j$. For example, if b strongly dominates b', then $E(b) > E(b')$, since $\bar{r}(a_1) < \ldots < \bar{r}(a_n)$. Further, for all rewards $r \geq \hat{r}$, $f(r,a_j) > f(r,a_k)$ if $j > k$, so $f^{b}(r) > f^{b'}(r)$ for all $r \geq r^* \geq \hat{r}$. Finally, note that as long as the incumbent's rewards remain above r^*, the voter's current belief will strongly dominate the initial belief π, since after one such observation we have

$$\frac{b_j}{b_k} = \frac{\beta_j(\pi,r,a^*)}{\beta_k(\pi,r,a^*)} = \frac{\pi_j f(r,a_j)}{\pi_k f(r,a_k)} > \frac{\pi_j}{\pi_k}; \quad (11)$$

after two such observations, we have

$$\frac{\beta_j(b,r,a^*)}{\beta_k(b,r,a^*)} = \frac{b_j f(r,a_j)}{b_k f(r,a_k)} > \frac{b_j}{b_k} > \frac{\pi_j}{\pi_k}, \quad (12)$$

etc. Therefore, if we can show that $V(b) > V(\pi)$ for all beliefs b which strongly dominate π, then the best response for the voter to a continual stream of rewards above r^* is to retain the incumbent, thereby proving the optimality of $\nu(.)$.

So consider the difference $V(b) - V(\pi)$. By (10), we know that $V(\pi)$ can be written as

$$V(\pi) = E(\pi) + \delta[\int_D V(\pi)f^{\pi}(r)dr + \int_{D'} V(\beta(\pi,r))f^{\pi}(r)dr] \tag{13}$$

where D,D′ partition R, and where $D' \subset [r^*,\infty)$. Now since V(.) is the value of following an ***optimal*** strategy, we know that for any $b \in P(\Omega)$,

$$V(b) \geq E(b) + \delta[\int_D V(\pi)f^{b}(r)dr + \int_{D'} V(\beta(b,r))f^{b}(r)dr] \tag{14}$$

where the sets D,D′ are the same as those in (13). The RHS of (14) is the payoff associated with retaining the current incumbent, replacing him in the subsequent period if $r \in D$ and then proceeding according to the optimal strategy, and retaining him if $r \in D'$ and then proceeding according to the optimal strategy. Thus,

$$V(b) - V(\pi) \geq \{E(b) - E(\pi)\} +$$

$$\delta\{V(\pi)\int_D [f^{b}(r)-f^{\pi}(r)]dr + \int_{D'} V(\beta(b,r))f^{b}(r)dr - \int_{D'} V(\beta(\pi,r))f^{\pi}(r)dr\}. \tag{15}$$

We can rewrite the second term in braces as

$$V(\pi)\int_{D}[f^{b}(r)-f^{\pi}(r)]dr\ +$$

$$\int_{D'}V(\beta(\pi,r))[f^{b}(r)-f^{\pi}(r)]dr\ +\ \int_{D'}[V(\beta(b,r))-V(\beta(\pi,r))]f^{b}(r)dr.$$

As noted above, if b dominates π, then for all $r \in D'$, $f^{b}(r) > f^{\pi}(r)$. Further, since for all $b' \in P(\Omega)$, $V(b') \geq V(\pi)$ by (10), the above term is greater than or equal to

$$V(\pi)\int_{D}[f^{b}(r)-f^{\pi}(r)]dr\ +\ \int_{D'}V(\pi)[f^{b}(r)-f^{\pi}(r)]dr\ +\ \int_{D'}[V(\beta(b,r))-V(\beta(\pi,r))]f^{b}(r)dr$$

$$= V(\pi)[\int_{D\cup D'}f^{b}(r)dr\ -\ \int_{D\cup D'}f^{\pi}(r)dr]\ +\ \int_{D'}[V(\beta(b,r))-V(\beta(\pi,r))]f^{b}(r)dr$$

$$= \int_{D'}[V(\beta(b,r))\ -\ V(\beta(\pi,r))]f^{b}(r)dr. \qquad (16)$$

Therefore,

$$V(b)\ -\ V(\pi)\ \geq\ \{E(b)\ -\ E(\pi)\}\ +\ \delta\{\int_{D'}[V(\beta(b,r))\ -\ V(\beta(\pi,r))]f^{b}(r)dr\}, \qquad (17)$$

where the first term in braces is strictly positive, and the second term is of undetermined sign.

Now if b dominates π then, as noted above, for all $r \in D'$ $\beta(b,r)$ will dominate $\beta(\pi,r)$. Let $p = \beta(b,r)$, and $q = \beta(\pi,r)$ for some $r \in D'$, so p dominates q. Then by definition of D',

$$V(q)\ =\ E(q)\ +\ \delta[\int_{M}V(\pi)f^{q}(r')dr'\ +\ \int_{M'}V(\beta(q,r')f^{q}(r')dr', \qquad (18)$$

where M′ can depend on r, and $M' \subset [r^*,\infty)$. Further,

$$V(p) \geq E(p) + \delta[\int_M V(\pi)f^p(r')dr' + \int_{M'} V(\beta(p,r')f^p(r')dr', \tag{19}$$

where as before the RHS of (19) gives the payoffs from (potentially) departing from the stationary optimal strategy by retaining the current incumbent, and replacing him in the subsequent period if and only if $r \in M$. Now, using precisely the same arguments as above, we get

$$V(p) - V(q) \geq \{E(p) - E(q)\} + \delta\{\int_{M'} [V(\beta(p,r'))-V(\beta(q,r'))]f^p(r')dr'. \tag{20}$$

Placing (20) into (17), then, we get

$$V(b) - V(\pi) \geq \{E(b) - E(\pi)\} + \delta\{\int_{D'} [E(\beta(b,r))-E(\beta(\pi,r))]f^b(r)dr\} +$$

$$\delta^2\{\int_{D'}[\int_{M'}[V(\beta(\beta(b,r),r')) - V(\beta(\beta(\pi,r),r'))]f^{\beta(b,r)}(r')dr']f^b(r)dr\}, \tag{21}$$

where now the first and second terms in braces are strictly positive, and the third is of undetermined sign. But then for all $r \in D'$, $r' \in M'$, we have $\beta(\beta(b,r),r')$ dominating $\beta(\beta(\pi,r),r')$, so we can apply this argument a ***third*** time, thereby generating three terms on the RHS which are strictly positive, and a fourth of undetermined sign but which is multiplied by δ^3. Continuing this logic, we see that $V(b) - V(\pi)$ will be greater than or equal to an infinite sum of strictly positive terms, plus a term on the order of $\delta^\infty.K$, where K is bounded since $\delta < 1$. But then $\delta^\infty.K = 0$, again since $\delta < 1$, thus proving $V(b) > V(\pi)$. □

In the above equilibria the voter has an incentive to play $\nu(.)$, in part through the incumbent playing a "trigger" strategy, and in part through the voter "learning" about the incumbent's type. The trigger strategy aspect comes about by requiring the voter to replace an incumbent whenever a relatively low reward is witnessed: if the voter does not replace, then the incumbent reverts to $a = \underline{a}$ regardless of his type. These "punishments" are thus akin to those found in repeated prisoners' dilemma games of reverting to one–shot Nash equilibrium behavior upon observing a defection, the difference here being it is only the currently employed candidate, and no other candidate, who reverts to "myopic" play.

On the other hand, since we have assumed that candidates select actions based only on their ***own*** history of rewards, the same sort of trigger cannot be used to keep the voter retaining incumbents when relatively high rewards are witnessed, since a newly chosen candidate does not condition on the voter's responses to previous incumbents. What gives the voter the incentive to retain the incumbent upon observing such rewards is that the voter's belief shifts, placing relatively greater weight on higher types. That is, as long as the incumbent has continually generated rewards greater than r^*, the voter's belief about the incumbent will ***strongly stochastically dominate*** that of an untried candidate. This dominance implies the voter receives a higher one–period payoff from the incumbent than from an untried candidate, and will continue to so as long as the incumbent generates rewards greater than r^*, where the incumbent is in addition more likely to generate rewards greater than r^*. Therefore the voter ***wants*** to retain the incumbent, and thus does not need the behavior of subsequent incumbents to generate this incentive.

A number of issues are worthy of mention. The first is that, while the proof of the Proposition assumed $r^* > \hat{r}$, this is not necessary. All that is required for the proof is r^* be such that for all $r > r^*$, $f(r,a^*(\omega_k)) > f(r,a^*(\omega_j))$ for $k > j$, i.e.

on the region $[r^*,\infty)$, the reward densities are "stacked" with densities from higher types strictly above those of lower types. Thus, for instance, as long as r^* is greater than the (unique) mode of $f(.;a^*(\omega_n))$, the proof goes through without modification.

Second, we could have dispensed with the assumption of a bounded action space for the incumbent, in place of additional assumptions elsewhere. For instance, if we assume an upper bound the candidates' discount factor of $\bar{\rho} < 1$, then for any $\omega \in \Omega$ the value function $V(\omega)$ will be bounded above by $u(\underline{a},\omega)/(1-\bar{\rho}) < \infty$. Therefore, we can let $A = [0,\infty)$, and as long as the Inada condition $u_1(a,\omega) \to -\infty$ as $a \to \infty$ continues to hold, we will effectively have an upper bound on the incumbent's actions.

Third, there is a continuum of equilibria of the form outlined in the above Proposition, parametrized by the critical value of rewards r^*. In terms of the players' preferences over these equilibria, it is easily shown that for all $\omega \in \Omega$ and for r^* high enough, $da^*(\omega)/dr^* < 0$, so lowering the cut–off increases the action choices by all types of candidate. This obviously increases the voter's immediate payoff; however the long–term effects are not so transparent. It may well be that these new action choices are closer together than the old, implying the voter "learns" an incumbent's type more slowly, thereby lowering the voter's utility if she is sufficiently patient. Conversely, it is clear that candidates prefer lower values of r^* to higher regardless of type while in office. Yet, if candidate i has not yet been employed, lowering the value of r^* pushes this higher payoff farther into the future (in an expectational sense).

Fourth, we can address the issue of the "cost" borne by the voter to create the incentive for the candidate of type ω_j to take action $a^*(\omega_j)$ as long as previous rewards have been sufficiently high. Suppose the candidates' strategies were simply to play $a^*(\omega)$ regardless of history, i.e. $\alpha(\omega) = a^*(\omega)$. Now with the voter learning about the incumbent's type, the best response for the voter would not be to maintain a time–invariant cut–off rule, but rather the retain/replace decision would be a function of the updated belief about the incumbent. For example, if $N = 2$, so there are exactly two types of candidate, Banks and Sundaram (1990) show that the best response by the voter to the above candidate strategy is to retain or replace the incumbent depending on whether the incumbent generates a higher or lower expected one–period reward relative to an untried candidate, i.e. the voter's optimal decision rule is ***myopic***. Therefore, if enough high rewards have been observed from the current incumbent, then the voter would "forgive" an occasional low reward and continue to employ the incumbent, albeit lowering her belief the incumbent is a hard working type. On the other hand, if the voter adopted this "forgiving" rule in the current context, the candidates would obviously alter their action rule, possibly taking lower actions when confronted with a more lax standard. Hence in the above equilibrium the voter will occasionally replace an incumbent even when the voter believes the incumbent is a relatively good type, to maintain the incentives for the candidates to take good actions.

Next, consider the probability of an incumbent's re–election, and how this probability might change over time. Clearly, this probability is simply equal to $[1 - F(r^*,a^*(\omega)] \equiv \pi(\omega)$ for a type–ω candidate, where $\pi(\omega)$ is increasing in ω. Thus, from a candidate's perspective the probability of his re–election is constant over time. From the ***voter's*** perspective, on the other hand, if the current belief about the incumbent is b^t, then the probability of the incumbent being retained for the next period is $\sum_{j\in N} b_j^t.\pi(\omega_j)$. Now if the incumbent is ***in fact*** retained for period $\tau+1$, then the voter's belief will shift to $b^{t+1}(r)$, where the dependence of this belief on $r \in R$ is through Bayes' Rule. And by the above discussion, we know that for all "acceptable" rewards, i.e. all $r \in R$ such that the incumbent is retained, it is the case that $b^{t+1}(r)$ strongly stochastically dominates b^t, implying in particular that

$$\sum_{j \in N} b_j^{t+1}(r).\pi(\omega_j) \quad > \sum_{j \in N} b_j^t.\pi(\omega_j) \quad .$$

Therefore, from the voter's perspective, the probability of retaining the incumbent is increasing with an incumbent's tenure regardless of the actual history of rewards realized. Note that this increase occurs not because the incumbent works harder as his tenure increases, for in fact an incumbent's actions do not vary over time. Rather, it is due to the learning process of the voter, who's beliefs over time place greater weight on "better" (from the voter's perspective) types as long as the incumbent remains in office.

Finally, we note that the retrospective voting rule adopted by the voter is highly non–stationary in the voter's beliefs about the incumbent. That is, the voter's belief about the incumbent upon observing t rewards above r^* followed by a single reward below r^* could be the same as that from observing t' rewards above r^* and ***no*** rewards below r^*, and yet in the first instance the voter replaces the incumbent while in the second the incumbent is retained. Thus, it may be enlightening to characterize stationary simple equilibria, where the voter's strategy and candidate i's strategy are only functions of the voter's current belief about i's type. Note that existence of such equilibria is not at issue, since the simple equilibria in which all candidates adopt $a = \underline{a}$ and the voter always replaces the incumbent trivially satisfies the stationarity requirement. Whether there exist more interesting stationary equilibria, and what the characteristics of such behavior might be, are as yet unanswered questions.

References

Alesina, A. (1988) Credibility and policy convergence in a two–party system with rational voters, ***American Economic Review*** 78:796–806.

Alesina, A. and S. Spear (1988) An overlapping generations model of electoral competition, ***Journal of Public Economics*** 37:359–379

Austen–Smith, D. and J. Banks (1989) Electoral accountability and incumbency, in ***Models of Strategic Choice in Politics***, ed. by Peter Ordeshook. Ann Arbor: University of Michigan Press.

Banks, J. and R. Sundaram (1990) Incumbents, challengers, and bandits: Bayesian learning in a dynamic choice model, RCER working paper no. 235, University of Rochester.

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

Ferejohn, J. (1986) Incumbent performance and electoral control, ***Public Choice*** 50:5–25.

Fiorina, M. (1981) ***Retrospective Voting in American National Elections***. New Haven: Yale University Press.

Kreps, D. and R. Wilson (1982) Sequential Equilibria, ***Econometrica*** 50:863–894.

Ledyard, J. (1989) Information aggregation in two–candidate elections, in ***Models of Strategic Choice in Politics***, ed. by Peter Ordeshook. Ann Arbor: University of Michigan Press.

Maitra, A. (1968) Discounted dynamic programming in compact metric spaces, ***Sankhya Ser A*** 30:211–216.

McKelvey, R. and R. Reizman (1990) Seniority in legislatures, working paper, California Institute of Technology.

Reed, R. (1991) Retrospective and prospective voting in a political market with finite–lived politicians, working paper, Texas A&M University.

Rogoff, K. (1990) Equilibrium political budget cycles, ***American Economic Review*** 80:21–37.

Rogoff, K. and A. Sibert (1988) Elections and macroeconomic policy cycles, ***Review of Economic Studies*** 55:1–16.