A Dynamic Model of Democratic Elections in Multidimensional Policy Spaces*

Jeffrey S. Banks[1] and John Duggan[2,†]

[1]*Division of Humanities and Social Sciences, California Institute of Technology, USA*
[2]*Department of Political Science and Department of Economics, University of Rochester, Rochester, NY 14627; dugg@troi.cc.rochester.edu*

ABSTRACT

We propose a general model of repeated elections. In each period, a challenger is chosen from the electorate to run against an incumbent politician in a majority-rule election, and the winner then selects a policy from a multidimensional policy space. Individual policy preferences are private information, whereas policy choices are publicly observable. We prove existence and continuity of equilibria in "simple" voting and policy strategies; we provide examples to show the variety of possible equilibrium patterns in multiple dimensions; we analyze the effects of patience and office-holding benefits on the persistence of policies over time; and we identify relationships between equilibrium policies and the core of the underlying voting game. As a byproduct of our analysis, we show how equilibrium incentives may lead elected representatives to make policy compromises, even when binding commitments are unavailable. We provide an informational story for incumbency

advantage. Finally, we give an asymptotic version of the median voter theorem for the one-dimensional model as voters become-arbitrarily patient.

Elections occupy a central position in the determination of public policies in representative democracies. By selecting the individuals whose subsequent decisions determine final policy outcomes, elections resolve conflicts among competing majorities, and transform the preferences of voters into collective choices. It is well-known that, when the policy space is one-dimensional and voters have single-peaked preferences, a single policy outcome, the ideal point of the median voter, is majority-preferred to all others. In the canonical model of Downs (1957), in which two candidates commit to policy platforms before a single election, this drives the candidates to the median and yields a unique Nash equilibrium of the electoral game. When the policy space is multidimensional, however, majority undominated (or "core") points rarely exist, and in the absence of a core point, the entire space of policy alternatives will be contained in a majority preference cycle, suggesting to some authors (e.g., Riker 1980) the instability of policies over time. In contrast to that literature, where coalitions are assumed to form fluidly irrespective of institutional constraints, we explicitly model electoral institutions and the incentives of voters and politicians, creating the potential to restrict the formation of coalitions and limit the potential instability of collective choices.

In this paper, we consider an infinite-horizon model of electoral accountability in which a representative is elected in every period and chooses a policy from a multidimensional space. As in the citizen–candidate literature, we view campaign promises as cheap talk (and therefore omit them from the model), and elected representatives choose policy unconstrained by past commitments. We consider farsighted and rational voters, who must calculate expected streams of payoffs conditional on all available information to vote for their preferred candidates, and we assume that elected representatives similarly anticipate the future consequences of policies and choose optimally. The challenge in proving existence of equilibrium, in contrast to the Downsian model where candidates compete for majority support, lies in resolving these decision problems simultaneously. In contrast to the usual result in citizen–candidate models, where elected officials simply choose their ideal policies, we find that under quite general conditions, some (even all) types of representative optimize by choosing a suitable compromise, balancing their interest in a desirable policy in the current period with concerns for re-election in the future. Thus, we provide a theory of endogenous compromise by elected representatives, relying on asymmetric information and repetition, rather than on the Downsian assumption of commitment.

Our framework is general with respect to the policy space and voter preferences: we allow policies to lie in a subset of any finite-dimensional Euclidean space, and borrowing weak assumptions from the spatial modeling literature, we assume that voter utilities are continuous and strictly quasi-concave. In each period, a challenger is drawn from the electorate to run against the incumbent in a majority-rule election, with the winner choosing the policy for that period. The process then moves to the next period, and the above sequence of events is repeated *ad infinitum*. Voters observe the policies chosen by

the representatives but not their preferences. Thus, incomplete information in the form of adverse selection is present, and elections confront voters with a non-trivial problem: they must update their beliefs about the incumbent based on her past policy choices and compare this to the expected policy outcomes upon electing a challenger. Representatives, being chosen from the electorate at large, have well-defined policy preferences of their own and face a trade-off in choice of policy: they have short-term incentives to choose policies in their personal interest, but they have long-term interests in staying in office to capture "non-policy" benefits of office and to obtain better policy outcomes than expected from a challenger. The key to this trade-off, and the role of private information in the model, is that pursuit of short-term policy interests may reveal information to voters that damages the representative's chances of re-election.

We prove the existence of "simple" equilibria in which voters use strategies that are retrospective (Fiorina 1981) in the following sense: an individual votes for re-election if and only if her utility in the previous period was at or above a fixed critical level, this level being determined endogenously as the expected value of an untried challenger. Thus, voters are also prospective in that they vote as though pivotal in the current election, reconciling the usual notions of retrospective and prospective voting. Because an untried challenger is inherently risky and voters are risk averse, an elected representative has a degree of leverage in choosing policies to achieve re-election, creating a form of incumbency advantage in equilibrium: an elected representative may be able to obtain a satisfactory policy outcome and win re-election by placating a majority of voters. We show through a series of examples that a wide variety of policy and re-election patterns can emerge in equilibrium, particularly in multiple dimensions. It is possible that no representative is ever re-elected, each choosing her ideal policy while in office and failing to gain the support of a majority of voters. With different parameter values, it is possible that all types of representative receive majority support. In such examples, some representatives may choose a compromise policy sufficient to ensure re-election but not too far from her ideal, or it may be that all representatives can win by simply choosing their ideal policies.

The dynamics of our model are relatively simple. In order to bring out the logic of equilibrium, which must resolve the expectations of voters and optimal policy choices of representatives, we assume a stationary environment. This allows us to focus on the class of simple equilibria in which voters use a fixed cut-off, described above, and the policy choices of representatives are history-independent. Thus, in case all types are re-elected in equilibrium, the first individual to hold office will remain there, choosing the same policy in every period, demonstrating that an extreme form of "policy persistence" can occur in the model. We interpret the extent of policy persistence as a measure of incumbency advantage in the model and as indicative of stability of policy outcomes were to introduce an element of noise in the environment.[1] We prove that if non-policy benefits of holding office are sufficiently high or individuals are sufficiently

---

[1] As we remark in section "Extensions," our results, suitably adapted, would go through even if we added an exogenous and time-invariant probability of removal from office. This would complicate the dynamics of the model but would not affect the spirit of our findings.

patient (and non-policy benefits are positive), then all simple equilibria exhibit such policy persistence. Patience on the part of the voters and representatives can produce this stability in any number of policy dimensions, even in the absence of a core point. If non-policy benefits are zero, then it turns out that patience leads to policy persistence *unless* there is a core point: the presence of a core point can actually be a destabilizing force. Even then, however, we are able to show that the set of policies acceptable to a majority of voters collapses to the core as patience increases. When patience is great enough, therefore, either policy persistence obtains, or the long run distribution of policies is concentrated arbitrarily close to the core.

We then characterize simple equilibria in connection to the core, and for the one-dimensional special case, we provide a game-theoretic foundation for the median voter theorem. It is known that the Downsian result of convergence at the median fails when the commitment assumption is dropped in the standard, one-shot model of elections. Our results establish that if we imbed the Downsian model in a more realistic framework that explicitly models repeated elections and incorporates private information about policy preferences, we in fact re-obtain the median voter result. We show it is possible that all representatives choose the same policy in equilibrium, a phenomenon we call "policy coincidence," only if non-policy office benefits are sufficiently high, individuals are sufficiently patient, and a core point exists. In that case, all representatives must choose the core point, and we say that the equilibrium exhibits "core equivalence." In one dimension, the core is always non-empty and consists of the median voter's ideal point, and we can show that if there are sufficient benefits of office (with positive discount factors) or sufficient patience (with positive non-policy benefits), then there is a *unique* simple equilibrium. In it, all representatives choose the median, giving us full core equivalence. If holding office confers no non-policy benefit, then core equivalence need not obtain, but we show that, as voters become more patient, the set of policies that ensure re-election, and the long run distribution of equilibrium policies along with it, still collapse to the median.

In multiple dimensions, where the core is typically empty, it follows from above that policy coincidence will be the exception. Thus, in equilibrium, some representatives choose distinct policies. Then, when voters are sufficiently patient or non-policy benefits of office are sufficiently high, our policy persistence result implies that multiple policies can be sustained in equilibrium. Such a conclusion comes not from a multiplicity of equilibria, but rather from the possibility that representatives with different policy preferences have the willingness and ability to attract and maintain different majority coalitions within a single equilibrium. In this way, when the policy space is multidimensional, two electorates with identical voter preferences can be associated with distinct stable policies.

LITERATURE REVIEW

Most analyses of elections follow the Downsian tradition in highlighting the pre-election campaign aspects of the competition for the role as representative. In the basic model,

each of two otherwise identical candidates simultaneously announces a policy to be implemented if elected, with voters then casting their ballots for the candidate offering the preferred policy. While originally presented by Downs as a model of a single election in a one-dimensional policy space with office-motivated candidates and complete information, subsequent research has analyzed repeated elections, policy-motivated candidates, and probabilistic voting. All of this work, however, has retained the important underlying assumption of the Downsian model that the winning candidate will faithfully carry out her announced policy. This commitment assumption is often rationalized on the grounds that if a candidate broke a campaign promise, then there would be some (unmodelled) electoral punishments inflicted in the future. This maneuver effectively "black boxes" a principle component of the public policy process, namely, why representatives behave as they do while in office. An alternative approach is to drop the commitment assumption and model political candidates as citizens, with policy preferences of their own (Osborne and Slivinski 1996, Besley and Coate 1997). But in the citizen–candidate literature, the focus is on one-shot elections, and voters have complete information about candidate preferences. As a consequence, an elected official simply chooses her ideal policy in every period.

A third approach, beginning with the work of Barro (1973) and Ferejohn (1986) and sometimes referred to as models of "electoral accountability," drops commitment in the context of repeated elections. In contrast to the Downsian model, voters base their decisions on the past performance of incumbents, rather than their current promises. In selecting policies, representatives typically care not only about winning, but also about their actions while in office, either through their own policy preferences or else in terms of the "effort" expended on their constituents' behalf. With the exception of Barro (1973), electoral accountability models assume some form of incomplete information is present: either the motivations of the representatives are known but their influence over policy, and hence over voter utility, is not (Ferejohn 1986, Austen-Smith and Banks 1989), or their influence over policy is known but their motivations are not (Reed 1994, Duggan 2000, Bernhardt *et al.* 2004, Meirowitz 2007), or neither is known (Rogoff 1990, Banks and Sundaram 1993, 1998, Coate and Morris 1995, Fearon 1999). To date, all of this work has maintained the original Downsian assumption of a unidimensional policy space, conceptualized either as a space of effort levels or as an ideological dimension. In fact, many of these models are further simplified by the assumption that there is just one voter.

Of the papers in the electoral accountability literature, the structure of our model is most similar to that of Duggan (2000), with the key differences being that the latter assumes a one-dimensional policy space, "tent-shaped" Euclidean distance utilities, and symmetry of the distribution of challenger ideal points. The existence of simple equilibria is proved, and it is shown that in all such equilibria, the median voter is decisive: a policy choice by an officeholder secures re-election if and only if it gives the median voter a payoff at least equal to the median's expected payoff from electing a challenger. This model has been extended to allow for term limits (Bernhardt *et al.* 2004), for parties (Bernhardt *et al.* 2005, Campuzano 2005), and for costly signaling in campaigns (Kang 2005).

Before proceeding, we mention several papers on dynamic elections that are otherwise dissimilar to ours. Duggan and Fey (2006) imbed the Downsian model with office-motivated candidates within an infinite-horizon model of repeated elections and show that under weak conditions on discount factors, arbitrary paths of policies can be supported in equilibrium using various constructions. Alesina (1988) assumes a one-dimensional policy space in a repeated elections setting, but in a two-candidate, simultaneous-move model without commitment. The preferences of the candidates are known to the voters, and include both policy and non-policy components. He shows that, when discount factors are high enough, a range of policy outcomes can be sustained in equilibrium when voters and candidates employ trigger strategies of a certain form. Kramer (1977) studies a two-candidate model of repeated elections in multiple dimensions such that in any period, the challenger may commit to a policy, while the incumbent is bound to her previous policy choice. Challengers maximize their margin of victory, and politicians and voters are myopic. He shows that, when voters have Euclidean preferences, equilibrium policies converge to the "minmax" set, a set that coincides with the core when the latter is non-empty.

Finally, the working paper version of this paper (Banks and Duggan 2001, 2008) contains several results removed from this economized version and a discussion of the connections between our equilibrium concept and that of perfect Bayesian equilibrium.

THE ELECTORAL MODEL

Let $X \subset \Re^d$ denote a compact and convex set of policies, let $N = [0, 1]$ be a continuum of individuals, and let the possible preferences of voters be indexed by a finite set T of types, denoted t. Each individual i's type t_i is drawn from the distribution $\rho = (\rho_1, \ldots, \rho_{|T|})$, where ρ_t is the probability of type t. To preclude trivial instances of the model, we assume that are at least two types and each type has positive probability, i.e., $|T| \geq 2$ and $\rho_t > 0$ for all $t \in T$. We extend the idea of independent types to the current model as follows: the distribution of an individual's type, conditional on the types of any finite number of other individuals, remains ρ. We assume that the law of large numbers holds, so that, with probability one, the fraction of type t individuals is ρ_t for all $t \in T$.[2] Any one individual's type is private information, but the distribution ρ is common knowledge. The preferences of type t individuals are represented by a utility function u_t on X, assumed to be continuous, concave, and strictly quasi-concave. We normalize payoffs so that $u_t(x) \geq 0$ for all $t \in T$ and all $x \in X$. Let $x_t = \arg\max\{u_t(x) : x \in X\}$ denote the unique ideal policy for type t individuals. Assume that ideal policies are distinct, i.e., $x_t \neq x_{t'}$ for all $t, t' \in T$, and note that, by strict quasi-concavity, $u_t(x_t) > 0$ for all $t \in T$.

---

[2]   We also assume that, for each type t, the set $\{i \in N : t_i = t\}$ is Lebesgue measurable with probability one. Judd (1985) establishes the existence of a joint distribution of voter types for which these conditions are satisfied for almost all realizations of voter types. See Banks and Duggan (2006) for rigorous foundations of this model.

Of interest later is the weighted majority voting game among the types in T, with weights given by the proportions $(\rho_1, \ldots, \rho_{|T|})$ of types present in the electorate. Let

$$\mathcal{D} = \left\{ C \subseteq T : \sum_{t \in C} \rho_t > 1/2 \right\}$$

denote the decisive coalitions of types. We impose the condition that there is no coalition of types $C \subset T$ such that $\sum_{t \in C} \rho_t = 1/2$, i.e., no coalition of types has precisely half of the population. This implies that the voting game is *strong,* in the following sense: for all $C \subseteq T$, either $C \in \mathcal{D}$ or $T \setminus C \in \mathcal{D}$. The *core* is the set, K, of policies that are undominated in this voting game, i.e.,

$$K = \left\{ x \in X : \begin{array}{l} \text{there do not exist } y \in X \text{ and } C \in \mathcal{D} \\ \text{such that, for all } t \in C, u_t(y) > u_t(x) \end{array} \right\}.$$

Because X is convex and utility functions are strictly quasi-concave, it follows that K, if non-empty, will be a singleton. Denote this core policy by x^c. In addition, the core satisfies the following *external stability* condition: for all $y \neq x^c$, $\{t \in T : u_t(x^c) > u_t(y)\} \in \mathcal{D}$. It is known that, because $\mathcal{D}$ is strong, the core is typically empty when X is multidimensional, but in *one* dimension, the core is always non-empty and is equal to the ideal policy of the weighted median type. Defining m as the unique element of T satisfying

$$\{t \in T : x_t \leq x_m\} \in \mathcal{D} \quad \text{and} \quad \{t \in T : x_t \geq x_m\} \in \mathcal{D},$$

we therefore have $x^c = x_m$.

Elections proceed as follows. In period 1, an individual is randomly chosen as representative and selects a policy in X. In each period $\tau = 2, 3, \ldots$, an individual is selected as representative as follows. A challenger is randomly drawn from a density on N to run against the incumbent, the representative from period $\tau - 1$. Let $\gamma = (\gamma_1, \ldots, \gamma_{|T|})$ denote the distribution of challenger types, where $\gamma_t > 0$ for all $t \in T$. We do not assume that the distribution of challenger types mirrors the distribution of types in the electorate. Individuals observe the "name" of the challenger, but not her type. Once the challenger is determined, each individual casts a vote in $\{In, Ch\}$, where In denotes a vote for the incumbent and Ch a vote for the challenger. If the proportion of individuals voting for the incumbent is at least one half, then the incumbent wins the election and becomes the period τ representative. Otherwise, the challenger wins. The period τ representative selects any policy in X, this selection is observed by the voters, and the game moves to the next period, where this process is repeated. Note that because the challenger is drawn from N according to a density, the probability any given individual is chosen to run as challenger in any given period is zero.

A *public history of length* τ, denoted h^τ, describes the publicly observed events in the first τ periods, namely, the individuals chosen as representatives, those chosen as challengers, vote tallies from elections, and policies selected by winners. An infinite

public history, h^∞, is an infinite sequence of these variables. In particular, let $\{i^\tau\}$ denote the corresponding sequence of representatives and $\{x^\tau\}$ the sequence of policies. An individual i's payoff from an infinite public history h^∞ is then defined as

$$(1 - \delta) \sum_{\tau=1}^{\infty} \delta^{\tau-1} [u_{t_i}(x^\tau) + \omega_i(i^\tau)\beta],$$

where $\delta \in [0, 1)$ is a common discount factor, $\beta \geq 0$ is a common non-policy benefit from being representative, and ω_i is the indicator function on N taking on the value of one if $i = i^\tau$ and zero otherwise.

A strategy for $i \in N$ describes, for each time period τ, a vote $v_i^\tau \in \{In, Ch\}$ and a policy $p_i^\tau \in X$ if selected as representative, both functions of the public history of length $\tau - 1$. Because types are private information, we follow Harsanyi (1967–1968) in modeling votes and policy choices as also depending on an individual's type. We focus on equilibria in which the individuals' strategies are especially simple. First, individuals employ *retrospective* voting rules: for all $i \in N$, there exists (measurable) $\underline{u}_i \colon T \to \Re$ such that for all $t \in T$, all $\tau \geq 1$, and all $h^{\tau-1}$,

$$v_i^\tau(h^{\tau-1}, t) = In \quad \text{if and only if} \quad u_t(x^{\tau-1}) \geq \underline{u}_i(t).$$

That is, i votes to retain the incumbent if and only if the incumbent's most recent policy choice satisfied the utility standard, or "cut-off," $\underline{u}_i(t)$. This cut-off is time-invariant, consistent with a "What have you done for me lately?" attitude on the part of the voters. Second, individuals' policy choices are *history-independent:* for all $i \in N$, there exists $p_i \colon T \to X$ such that for all $t \in T$, all $\tau \geq 1$, and all $h^{\tau-1}$,

$$p_i^\tau(h^{\tau-1}, t) = p_i(t).$$

Thus, i chooses the same policy any time she is elected as representative. Note that these two requirements are mutually re-enforcing: if voter strategies depend on history only through the incumbent's last chosen policy, then an incumbent's policy decision problem looks the same in all periods she is selected. Hence, if an individual has an optimal policy strategy, then she necessarily has an optimal strategy that is history-independent. Similarly, if representatives adopt history-independent policies, then knowledge of the last policy chosen by an individual is sufficient for a voter to accurately predict that individual's policy choices in all future periods.

To resolve equilibrium existence issues, however, we must complicate our description of policy choice strategies by allowing for "mixing" by representatives, i.e., the arbitrary choice of policies over which the representative is indifferent, in the first term of office. To preserve the idea of history-independence, we look for equilibria in which, after that initial policy choice, the individual then chooses the same policy in every subsequent term of office. In general, we represent mixing over policies of a representative i, newly elected in period τ, as a Borel probability measure π_i^τ, which again is a function of public history and i's type. Let $\mathcal{P}(X)$ denote the set of Borel probability measures on

X, endowed with the topology of weak convergence.[3] We focus on simple equilibria in which policy choices by individual i, newly elected in period τ, can be described by a (measurable) mapping $\pi_i \colon T \to \mathcal{P}(X)$ such that, for all $t \in T$, all $\tau \geq 1$, and all $h^{\tau-1}$,

$$\pi_i^{\tau}(h^{\tau-1}, t) = \pi_i(t).$$

Here, $\pi_i(t)(Y)$ is the probability that type t of individual i initially chooses a policy in the (measurable) subset $Y \subseteq X$, that policy being chosen by i whenever she is re-elected.

A *simple strategy* for i consists of a pair $\sigma_i = (\pi_i, \underline{u}_i)$. A simple strategy *profile*, denoted $\sigma = (\pi_i, \underline{u}_i)_{i \in N}$, specifies a simple strategy for every individual with the added restriction of type-symmetry: $\underline{u}_i(t) = \underline{u}_j(t)$ and $\pi_i(t) = \pi_j(t)$ for all $i, j \in N$ and all $t \in T$. Note that this is a "within-type" symmetry condition, and does not restrict policy choices across different types of an individual. Abusing notation slightly, let $\underline{u}_t$ denote the cut-off and π_t the mixed policy choice strategy used by all type t voters. We will also use the notation $\pi = (\pi_1, \ldots, \pi_{|T|})$ for a profile of mixed policy choice strategies.

Each strategy profile σ induces a probability distribution over infinite histories from the beginning of the game (prior to selecting the first representative) and, with it, an expected utility $v_i(\sigma, t)$ for every $i \in N$ and $t \in T$.[4] Because challengers are drawn from a density on N, in almost all histories a challenger will not have held office previously. By our independence assumption, therefore, the voters' beliefs about a challenger's type are given by ρ after almost all histories. By our restriction to simple strategies, then, $v_i(\sigma, t)$ is also i's expected utility, or *continuation value,* of replacing the current incumbent with an untried challenger, after almost every history.[5] Further, because individuals of the same type, say t, have a common per-period utility function, a common discount factor, and common beliefs about challengers, they will have the same continuation value, which we henceforth express as $v_t(\sigma)$. Informally, a simple strategy profile σ^* constitutes a *simple equilibrium* if for all $t \in T$, π_t^* is a "best response" whenever a type t representative makes a policy choice and $\underline{u}_t^*$ is a "best response" in every vote.

SIMPLE EQUILIBRIA

In this section, we give conditions on a simple strategy profile σ formalizing the idea that voting and policy choice strategies are best responses for all individuals. Our optimality condition on voting strategies is, essentially, that individuals decide to retain or replace the current incumbent based on which candidate offers the higher payoff. That is,

---

[3] A sequence $\{\pi^n\}$ of probability measures on X *weakly converges* to a probability measure π if, for all (bounded) continuous functions $f \colon X \to \Re$, we have $\int f d\pi^n \to \int f d\pi$. Because the policy space X is a compact metric space, the set $\mathcal{P}(X)$ will be compact in this topology. See Aliprantis and Border (1994, Theorem 12.10).

[4] See Banks and Duggan (2006) for an explicit construction of this distribution.

[5] We do not consider the probability zero set of histories in which a challenger has previously held office. After such histories, continuation values would be defined to reflect updating based on all relevant information.

voters act as though "pivotal" in the current election, voting for the incumbent if the expected utility from re-electing her is at least as great as the expected utility from electing an untried challenger. The latter, for a type t of individual i, is simply $v_t(\sigma)$. As for retaining the incumbent, suppose $x \in X$ is the incumbent's policy choice in the previous period. Because individuals are adopting history-independent policy choice strategies, the incumbent will continue to select x in the current period if retained. If σ determines that the incumbent subsequently be replaced, then the expected utility to i from retaining the current incumbent is $(1 - \delta) u_t(x) + \delta v_t(\sigma)$, which is greater than $v_t(\sigma)$ if and only if $u_t(x)$ is greater than $v_t(\sigma)$. If σ determines that the incumbent be forever retained, then the expected utility to i from retaining the current incumbent is simply $u_t(x)$, and so again retaining the incumbent is preferred by i if and only if $u_t(x)$ is greater than $v_t(\sigma)$. Thus, the cut-off $\underline{u}_t = v_t(\sigma)$ captures the decision of a pivotal voter.[6] *Our best response condition for voting strategies is therefore that, for all $t \in T$,*

$$\underline{u}_t = v_t(\sigma).$$

Note that, while we have described the voters' strategies as "retrospective" because votes are determined by simple cut-off rules, they are actually "prospective" as well in equilibrium: an individual votes for an incumbent only when retaining the incumbent generates a higher expected future payoff than that generated by replacing her.

Given that individuals of the same type adopt common cut-off rules and that ρ_t is the actual proportion of type t voters, the voting stage, from the perspective of the candidates, is simply a weighted voting game among the types in T, with decisive coalitions $\mathcal{D}$. This simplifies the statement of the best response condition on the policy strategies, because an incumbent is retained if and only if the set of types voting for the incumbent is in $\mathcal{D}$. For each $t \in T$, let

$$A_t(\sigma) = \{x \in X : u_t(x) \geq \underline{u}_t\}$$

denote the *acceptance set* for type t individuals, i.e., those policies satisfying the cut-off $\underline{u}_t$ and inducing all type t individuals to vote for the incumbent. By the compactness and convexity of X and the continuity and concavity of u_t, this set is compact and convex. For each coalition $C \subseteq T$ of types, define the set

$$A_C(\sigma) = \bigcap_{t \in C} A_t(\sigma)$$

of those policies inducing all types $t \in C$ to vote for the incumbent. As the intersection of compact and convex sets, $A_C(\sigma)$ is compact and convex as well. Finally, define

$$A(\sigma) = \bigcup_{C \in \mathcal{D}} A_C(\sigma)$$

---

[6] Strictly speaking, the prediction of x here is justified by Bayesian updating about the incumbent's type only following histories consistent with σ. See Banks and Duggan (2008) for a detailed discussion of these issues.

as those policies that receive majority support and will, therefore, lead to re-election of the incumbent. This *social acceptance set* is compact but not necessarily convex (see Example 2).

Suppressing for the moment the dependence of the set A on the profile σ, the choice for the type t of individual i when selected as representative is to either select a policy $x \in A$ (if non-empty), in which case she is retained for the next period, or select a policy $x \notin A$ and subsequently be replaced. Choosing any $x \neq x_t$ from outside of A is dominated by simply choosing x_t. Additionally, when $x_t \in A$, then she will optimally select this as her policy in all periods and remain as incumbent forever. Otherwise, i.e., when $x_t \notin A$, the representative faces a trade-off: select x_t in the current period and be replaced or choose a u_t-maximizing policy from A and be retained. The payoff from choosing $x_t \notin A$ is equal to

$$(1 - \delta)\left[u_t\left(x_t\right) + \beta\right] + \delta v_t\left(\sigma\right),$$

reflecting the one-time payoff from the representative's ideal point, followed by the continuation value of an untried challenger thereafter. Further, if choosing from A is optimal in the current period, then it will remain so in all future periods, and any u_t-maximizing policy from A will remain u_t-maximal in all future periods. Let

$$M_t(A) = \arg\max\{u_t(x) : x \in A\}$$

denote the set of best socially acceptable policies for a type t individual. *Our best response condition for policy choice strategies is therefore that for all $t \in T$,* (i) *when a type t representative prefers to remain in office, i.e.,*

$$\sup\{u_t(x) : x \in A\} + \beta > (1 - \delta)[u_t(x_t) + \beta] + \delta v_t(\sigma),$$

she choose from the best policies that ensure re-election, i.e., $\pi_t(M_t(A)) = 1$, (ii) *when the inequality is reversed, $\pi_t(\{x_t\}) = 1$ (and the representative is replaced in the next period), and* (iii) *when equality holds, $\pi_t(M_t(A) \cup \{x_t\}) = 1$.* This completes our definition of simple equilibrium.

The forgoing shows how representatives, themselves members of the electorate, take into consideration the future policy consequences — even after being removed from office — of their current policy decisions. By choosing her best available policy from the social acceptance set A, a representative can guarantee that this policy remains in effect forever. Alternatively, she can choose from outside of A, with the future policy consequences of such an act summarized by σ. Which of these two options is preferred then depends on the location of her best policy in A relative to her ideal policy (i.e., her best policy in X) and her continuation value, as well as the value of future policies relative to those of the present (represented by the discount factor δ) and the non–policy benefits of remaining in office (given by β).

Given a simple strategy profile σ, a type t individual's continuation value satisfies

$$v_t(\sigma) = \sum_{t' \in T} \gamma_{t'} \left[[1 - \pi_{t'}(A(\sigma))][(1 - \delta)u_t(x_{t'}) + \delta v_t(\sigma)] + \int_{A(\sigma)} u_t(x)\pi_{t'}(dx) \right].$$

The first term in the brackets is the probability that the current representative chooses from outside $A(\sigma)$ multiplied by t's expected payoff in that case, which is simply one period of the representative's ideal policy followed by her removal and subsequently "starting over." The second (integral) term gives t's expected payoff if the current representative selects from $A(\sigma)$, in which case, by history-independence, the latter will make the same decision and be re-elected in all future periods. Manipulating this equation to get an explicit solution, we have

$$v_t(\sigma) = \frac{\sum_{t' \in T} \gamma_{t'}\big[[1 - \pi_{t'}(A(\sigma))](1 - \delta)u_t(x_{t'}) + \int_{A(\sigma)} u_t(x)\pi_{t'}(dx)\big]}{1 - \delta \sum_{t' \in T} \gamma_{t'}[1 - \pi_{t'}(A(\sigma))]},$$

which is a convex combination of the one-period payoffs to t conditional on representatives choosing from outside $A(\sigma)$ (i.e., $u_t(x_{t'})$) and from inside $A(\sigma)$ (i.e., $\int_{A(\sigma)} u_t(x)\pi_{t'}(dx)/\pi_{t'}(A(\sigma))$). Thus, $v_t(\sigma)$ can be written as the expectation of u_t with respect to a probability distribution over X, where elements in $X \backslash A(\sigma)$ receive relatively less weight (by a factor of $1 - \delta$) as these policies are "temporary," whereas policies in $A(\sigma)$ are "permanent."[7] Now define

$$x(\sigma) = \frac{\sum_{t' \in T} \gamma_{t'}\big[[1 - \pi_{t'}(A(\sigma))](1 - \delta)x_{t'} + \int_{A(\sigma)} x\pi_{t'}(dx)\big]}{1 - \delta \sum_{t' \in T} \gamma_{t'}[1 - \pi_{t'}(A(\sigma))]},$$

which is a similarly weighted average of equilibrium *policies.* Thus, $x(\sigma)$ is the expected outcome associated with the probability distribution over X induced by σ. Concavity implies $u_t(x(\sigma)) \geq v_t(\sigma)$ for all $t \in T$, and by strict quasi-concavity and our assumption that $\gamma_{t'} > 0$ for all $t' \in T$, this inequality strict unless all individuals of all types choose the same policy when in office. Therefore, $x(\sigma) \in A_t(\sigma)$ for all $t \in T$ whenever σ satisfies the best response condition for voters, and so the set $A(\sigma)$ of policies that lead to re-election will always include at least $x(\sigma)$ and will, therefore, be non-empty.

We note, further, that if the core is non-empty, so $K = \{x^c\}$, then $x^c \in A(\sigma)$ whenever σ satisfies the best response condition for voters because $x(\sigma) \in A_t(\sigma)$ for all $t \in T$ (as just argued) and $u_t(x^c) \geq u_t(x(\sigma))$ for a weighted majority of types (by external stability), we therefore have $u_t(x^c) \geq u_t(x(\sigma)) \geq v_t(\sigma)$, implying $x^c \in A_t(\sigma)$, for a weighted majority of types. Thus, the core policy x^c, whenever it exists, is an acceptable policy in every equilibrium. An implication, of course, is that in one dimension, the ideal point of the weighted median type is always sufficient to ensure re-election in a simple equilibrium.

In equilibrium, because the social acceptance set is non-empty, there will always exist policies representatives could choose to ensure re-election. The question is whether they

---

[7]   Formally, we define the "continuation distribution" of σ, denoted ψ, as follows: for measurable $Y \subseteq X$,

$$\psi(Y) = \frac{\sum_{t \in T} \gamma_t[(1 - \pi_t(A))(1 - \delta)\mu_{x_t}(Y) + \pi_t(Y \cap A)]}{1 - \delta \sum_{t \in T} \gamma_t(1 - \pi_t(A))},$$

where μ_{x_t} is the point mass on x_t.

find it optimal to do so. With this in mind, given a simple equilibrium σ^*, partition the set T of types into three subsets, W ("winners"), L ("losers"), and C ("compromisers") as follows:

$$W(\sigma^*) = \{t \in T : x_t \in A(\sigma^*)\}$$
$$L(\sigma^*) = \{t \in T : x_t \notin A(\sigma^*) \text{ and } \pi_t^*(\{x_t\}) > 0\}$$
$$C(\sigma^*) = \{t \in T : x_t \notin A(\sigma^*) \text{ and } \pi_t^*(M_t(A(\sigma^*))) = 1\}.$$

Thus, winning types find their ideal policy acceptable to a majority, and so implement this policy in all periods. Compromising types are not so fortunate, but they still find *some* acceptable policy as good as choosing their ideal policy and subsequently being replaced, and they always choose such a policy. Finally, losing types have the opposite preference, in that *no* acceptable policy is better than simply choosing their ideal policy and subsequently being replaced, and a positive fraction of these types do choose the latter option. In the next section, we show by way of a series of examples that any one of these sets, or even two, may be empty in equilibrium.

   This is of interest because the emptiness or non-emptiness of these sets largely determines the equilibrium dynamics of elections in our model. In particular, if $L(\sigma^*) = \emptyset$, then all representatives choose policies in the social acceptance set. The first individual to hold office is therefore re-elected, and, by history-independence, remains in office forever, implementing the same policy in each period. We refer to this as *perfect policy persistence*. On the other hand, if $L(\sigma^*) \neq \emptyset$, then the first representative and any newly elected challenger will, with positive probability, choose a losing policy and be replaced in the following period. As long as it is not the case that $\pi_t^*(A^*) = 0$ for all types, however, a representative will (with probability one) eventually be elected and choose a policy in the social acceptance set, where again this policy remains in place forever. We call this *eventual policy persistence*. When it obtains, the long run distribution of policy outcomes puts probability one on the social acceptance set, though the short run distribution may put positive probability on policies outside the social acceptance set chosen by losing types.

EXAMPLES

Example 1 "*All losers*" *equilibrium.*

Let there be two dimensions, $d = 2$; let there be three types of individuals, $|T| = 3$, with quadratic utilities, $u_t(x) = 1 - (\|x_t - x\|)^2$, and ideal points at the vertices of a simplex, $\|x_t - x_{t'}\| = 1$ for all $t, t' \in T$. Assume no office benefit, $\beta = 0$, and types are equally represented in the population and in the pool of possible candidates, $\rho_t = \gamma_t = 1/3$ for all t. Assuming all individuals propose their ideal policy and subsequently are replaced, the continuation value for any individual is given by $v_t = (1/3)(1) + (2/3)(0) = 1/3$, i.e., in all periods there is a $1/3$ chance of having their ideal policy being chosen, generating a utility of 1, and a $2/3$ chance of some other type's ideal policy being chosen, generating a utility of 0. What needs to be checked is that individuals in their role as representative

prefer this losing strategy to compromising. For a type t individual, the closest point in A^* to x_t is $(1 - \sqrt{2/3})$ away, because all individuals have continuation value equal to $1/3$ and utility is quadratic. Thus, t can either lose and receive $(1 - \delta)(1) + \delta(1/3)$, or compromise and receive $1 - (1 - \sqrt{2/3})^2$. Grinding through the algebra, we see that losing is preferred as long as $\delta < (5/2) - 3\sqrt{2/3} \approx 0.05$. Further, because losing is *strictly* preferred, the equilibrium is unaffected if β is positive and small enough.

Example 2 "*All compromisers*" *equilibrium.*

Let the parameter values be the same as in Example 1, except for δ. Consider Figure 1, from Baron's (1991) model of spatial bargaining. Here the points a, c, and e are chosen so that $u_1(c) = (u_1(a) + u_1(e))/2$, with symmetric equalities holding for types 2 and 3. We claim that the following constitutes an equilibrium for δ sufficiently large: all type 1 individuals select policy a and set $\underline{u}_1 = u_1(c)$, all type 2 individuals select policy c and set $\underline{u}_2 = u_2(e)$, and all type 3 individuals select e and set $\underline{u}_3 = u_3(a)$. Given these cut-offs, each type is optimizing conditional on choosing from A^*, and further, if individuals adopt these policy strategies, then their cut-offs are indeed equal to their continuation values. Thus, what remains to be checked is whether representatives are optimizing by selecting from A^*, rather than choosing their ideal points. By symmetry, we need only to check this condition for one type, say type 1. The relevant comparison is between choosing $p = a$ and remaining in office forever, and choosing $p = x_1$ and being replaced in the following period. The utility of the former is equal to $u_1(a)$, while the utility of the latter is $(1 - \delta)(1) + \delta u_1(c) = 1 - \delta(1 - u_1(c))$. Thus, a type 1 individual prefers to compromise whenever
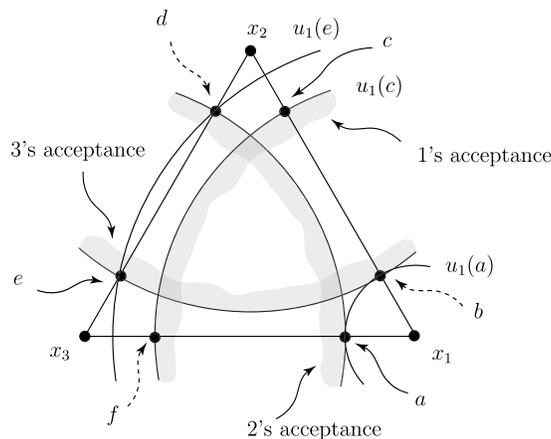
$$u_1(a) \geq 1 - \delta(1 - u_1(c)),$$

Figure 1. "All compromisers" equilibrium.

or equivalently,

$$\delta \geq \frac{1 - u_1(a)}{1 - u_1(c)}.$$

Because $1 > u_1(a) > u_1(c) > 0$, the right-hand side of the above expression lies in $(0, 1)$, and when δ is above this amount, we have an equilibrium.[8]

Example 3 "*All winners*" equilibrium.

Let there be two dimensions, $d = 2$; let there be fives types, $|T| = 5$, with highly risk averse utilities, $u_t(x) = 16 - (\|x_t - x\|)^4$. Arrange ideal points in a square with one at the center, $x_1 = (1, 0)$, $x_2 = (0, 1)$, $x_3 = (-1, 0)$, $x_4 = (0, -1)$, $x_5 = (0, 0)$, and note that the ideal point of type 5 is the core point. Again assume equal representation, $\rho_t = \gamma_t = 1/5$ for all t. Assuming all individuals propose their ideal policy and have it accepted, the continuation value for types 1–4 is given by $v_t = (1/5)(16) + (1/5)(15) + (2/5)(12) + (1/5)(0) = 11$, while the continuation value for a type 5 is $v_5 = (1/5)(16) + (4/5)(15) = 15.8$. Hence, type 5 individuals only vote to re-elect their own, and so an individual of, e.g., type 1 must secure the votes of types 2 and 4 to be re-elected. Because $u_2(x_1) = u_4(x_1) = 12 > 11$, type 2 and type 4 individuals indeed vote to re-elect type 1 representatives even when they propose their ideal policy. Similarly, types 2 and 4 vote to re-elect type 3 representatives when they propose their ideal policy, and types 1 and 3 vote to re-elect type 2 and type 4 representatives. Finally, $p_5 = x_5$ is acceptable to all types. Because all individuals are implementing their ideal policy when chosen as representative *and* remaining as incumbent forever, the policy strategies are clearly optimal, and we therefore have an equilibrium. A distinguishing feature of this strategy profile is that it constitutes an equilibrium for *every* value of δ and β, regardless of time preferences or non–policy benefits. We note in the next section that when $\beta \geq (1 - \delta)/\delta$, there is a second simple equilibrium in which all representative choose the core point $(0, 0)$.

Example 4 "*Mixed*" equilibrium.

Let the policy space be an interval, $X = [-1, 1]$; let there be five types, $|T| = 5$, with quadratic utilities, $u_t(x) = 4 - \|x - x_t\|^2$, and ideal points $x_1 = -1$, $x_2 \in (-1, 0)$, $x_3 = 0$, $x_4 \in (0, 1)$, $x_5 = 1$. Assume no office benefit, $\beta = 0$, and equal representation, $\rho_t = \gamma_t = 1/5$ for all t. We construct an equilibrium in which type 1 and type 5 individuals lose, type 2 and type 4 individuals compromise at $-c$ and c, respectively (where $c \in (0, 1)$), and type 3 individuals win. With quadratic utilities and a single dimension, one can show that type 3 individuals are decisive, in the sense that a proposal will satisfy a majority if and only if it satisfies the median voter (see Banks and Duggan

---

8 Note that by the symmetry of the environment we actually have, as in Baron (1991), another equilibrium where type 1's select b, type 2's select d and type 3's select f, as well as an equilibrium where half the type 1's select a and the other half b, half the type 2's c and the other half d, and half the type 3's e and the other half f.
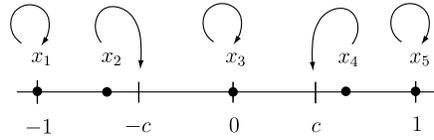
Figure 2. "Mixed" equilibrium.

2006, Lemma 2.1), so we only check this continuation value: $v_3 = [(2/5)(1 - \delta)(3) + (1/5)(4) + (2/5)(4 - c^2)]/(1 - (2/5)\delta)$. For a type 3 individual to be indifferent between accepting and rejecting c, we set v_3 equal to $4 - c^2$. We find that the desired value for c is

$$c(\delta) = \sqrt{\frac{2 - 2\delta}{3 - 2\delta}}.$$

Note that $c(1) = 0$, $c(0) = \sqrt{2/3}$, and $c' < 0$. Because $c(\delta)$ is bounded away from 1, there exists a positive δ, say δ^+, for which type 1 and type 5 individuals prefer to lose rather than compromise. Now set x_2 slightly to the left of $-c(\delta^+)$ and x_4 slightly to the right of $c(\delta^+)$, so that type 2 and type 4 individuals prefer to compromise. See Figure 2.

The equilibria in Examples 2 and 3 exhibit perfect policy persistence, in that the first representative remains as incumbent forever by choosing the same acceptable policy in every period. In contrast, the equilibria in Example 4 exhibits eventual policy persistence: only types 2, 3, and 4 choose acceptable policies, and so there will exist policy variability until such a type is elected.

EQUILIBRIUM ANALYSIS

Existence and Continuity

In this subsection, we address two fundamental theoretical issues: existence and continuity properties of equilibria. While the main contribution of the paper are our results on policy persistence and connections to the core, and implications for policy compromise and incumbency advantage, these theoretical results inform us that our characterizations are not vacuous or fragile.

Theorem 1 *There exists a simple equilibrium.*

A further desirable property of equilibria is uniqueness. Indeed, Duggan (2000) establishes uniqueness in a one-dimensional version of the model. But Example 3 demonstrates that there may be multiple simple equilibria when the policy space is multidimensional, so no general uniqueness result is available. In proving existence, we need to allow for the possibility that individuals of the same type adopt different policies while in office due to mixing. This comes about because, as seen in Example 2, the social acceptance set $A(\sigma)$ need not be convex, and so we may have a situation in which two distinct policies

$x, y \in A(\sigma)$ are optimal for type t individuals and yet no convex combination of x and y is in $A(\sigma)$. In addition, even if $A(\sigma)$ *is* convex, type t individuals may be indifferent between choosing (optimally) from $A(\sigma)$ and choosing $x_t \notin A(\sigma)$, with no convex combination giving as high a payoff. Allowing some type t individuals to choose one policy when in office while others choose a different policy, and then having these proportions determined in equilibrium, effectively smooths out, or "convexifies," representative behavior from the perspective of the voters.

We next show that the set of equilibrium policies changes in a nice way as one varies the underlying parameters of the model, which include the type distribution ρ, the discount factor δ, the non-policy office benefit β, and the distribution γ of challenger types. In addition, so that we may analyze the effect of varying voter preferences, we parameterize utility functions by λ, which lies in a subset $\Lambda \subseteq \Re^k$ of finite-dimensional Euclidean space. Technically, we expand the domain of the utility function of a type t to $u_t \colon X \times \Lambda \to \Re_+$, and we assume that each u_t is jointly continuous in (x, λ), that each $u_t(\cdot, \lambda)$ is concave and strictly quasi-concave in x, and that ideal points are distinct across types. As an example of such a parameterization, it could be that λ is a vector $(\lambda_1, \ldots, \lambda_{|T|})$ with each λ_t representing the ideal point of a quadratic utility function for type t. For parameters ρ, δ, β, and γ, let $E(\rho, \delta, \beta, \gamma)$ denote the set of profiles of simple equilibrium policy choice strategies. Our next result formalizes the idea that small variations in $(\rho, \delta, \beta, \gamma)$ cannot lead the set $E(\rho, \delta, \beta, \gamma)$ of simple equilibria to expand discontinuously.[9]

Theorem 2 *The correspondence E of simple equilibria is upper hemicontinuous in the parameters of the model.*

One of the important consequences of Theorem 2 is the following. If we can solve for all of the equilibria at some particular parameter values, then we know that for values suitably close to this, all equilibria will be close (in the sense of weak convergence) to the original set: though policies far from this set may occur with positive probability, that probability must go to zero as we approach the original parameter values of the model. Hence, when we fully characterize the equilibria in specific situations, we can be confident that these results are not "knife-edge" and that they provide an upper bound for the equilibria in that region of the parameter space.

Policy Persistence

We now analyze the dynamics of simple equilibria. Our focus is on the likelihood that some or all types choose losing policies, for it is this aspect of equilibrium strategies

---

[9] A correspondence $\Phi \colon X \rightrightarrows Y$ is *upper hemicontinuous* if given each point x in the domain and each open set G containing $\Phi(x)$, the set inclusion $\Phi(x) \subseteq G$ is maintained by arbitrarily small perturbations of x. See the Appendix for a formal definition for the equilibrium correspondence E.

that determines the persistence of policies over time.[10] The first result of this subsection shows that if either voters are sufficiently patient (and $\beta > 0$) or the non-policy benefits from office are sufficiently large (and $\delta > 0$), then all types are winners or compromisers. In that case, we have perfect policy persistence, i.e., the first representative remains as incumbent forever by choosing the same policy in every period, and a strong incumbency advantage. Our analysis also identifies a weaker constraint on our parameters under which not all types are losers, implying eventual policy persistence, i.e., a type of representative who selects a policy from A^* and implements it in all remaining periods is eventually chosen. The last result of the subsection is that even when non-policy benefits of office are zero, sufficiently high δ implies perfect policy persistence, *unless* the core is non-empty — and in that case, the social acceptance set, and therefore the long run distribution of policies, must collapse to the core.

As argued earlier, in any simple equilibrium σ^*, we must have $u_t(x(\sigma^*)) \geq v_t(\sigma^*)$ for all $t \in T$, and because $x(\sigma^*) \in A^*$, it follows that $\hat{u}_t(\sigma^*) \equiv \max\{u_t(x) : x \in A^*\}$ satisfies $\hat{u}_t(\sigma^*) \geq v_t(\sigma^*)$. A type t incumbent will prefer to compromise whenever

$$\hat{u}_t(\sigma^*) + \beta > (1 - \delta)[u_t(x_t) + \beta] + \delta v_t(\sigma^*).$$

For every $\delta < 1$, the first term on the left-hand side is strictly greater than the last term on the right-hand side (recall utilities are non-negative), and therefore, if $\beta \geq (1 - \delta)[u_t(x_t) + \beta]$, then it must be that t is not a losing type, i.e., $t \notin L(\sigma^*)$. Rewriting this inequality,

$$u_t(x_t) \leq \frac{\beta\delta}{1 - \delta}.$$

Now define

$$\overline{\alpha} = \max\{u_t(x_t) : t \in T\} \quad \text{and} \quad \underline{\alpha} = \min\{u_t(x_t) : t \in T\},$$

which are well-defined and positive, because each $u_t(x_t)$ is strictly positive and T is finite. The next result is an immediate consequence of the foregoing observations.

Theorem 3 *Let σ^* be a simple equilibrium.* (i) *If $\beta\delta/(1 - \delta) \geq \overline{\alpha}$, then $L(\sigma^*) = \emptyset$ in every simple equilibrium σ^*.* (ii) *If $\beta\delta/(1 - \delta) \geq \underline{\alpha}$ then $T \backslash L(\sigma^*) \neq \emptyset$ in every simple equilibrium σ^*.*

Theorem 3 has two important implications. First, (i) implies that for every positive level $\beta > 0$ of office benefit, there exists a level $\overline{\delta} \in [0, 1)$ of patience such that, when $\delta \geq \overline{\delta}$, all simple equilibria exhibit perfect policy persistence. And for every $\delta > 0$, there

---

[10]   We are grateful to a referee for the observation that policy persistence in our model can be viewed as a form of gridlock arising from a representative's incentives for re-election, rather than from conflict of interest among representatives from different parties. See Krehbiel (1998) for an in-depth discussion of partisan gridlock.

exists $\overline{\beta} > 0$ such that when $\beta \geq \overline{\beta}$, every equilibrium exhibits perfect policy persistence. Second, using (ii), we can get lower bounds than these, at the cost of replacing "perfect" with "eventual." That the result of Theorem 3 does not hold for arbitrary parameter values follows from Example 1, which shows the possibility of "all losers" in multiple dimensions when δ and β are both sufficiently small.

Because $\underline{\alpha}$ is positive, the above result is silent when either β or δ (or both) are equal to zero. Indeed, it is clear that when δ equals zero, representatives will simply choose their ideal policy while in office regardless of the value of β, implying that in certain situations (e.g., Example 1), no representative is ever re-elected. Hence, eventual policy persistence fails to hold. On the other hand, even when β equals zero, the next result shows that for sufficiently high values of δ, eventual policy persistence must hold in every equilibrium (e.g., Example 2).

Theorem 4 *Let σ^* be a simple equilibrium with social acceptance set A^*. There exists $\overline{\delta} \in (0, 1)$ such that if $\delta \in [\overline{\delta}, 1)$, then $\pi_t^*(A^*) > 0$ for some $t \in T$.*

Theorems 3 and 4 give us two results on the properties of simple equilibria as the discount factor δ approaches one: perfect policy persistence must occur if β is positive, and eventual policy persistence must occur even if β is zero. The final result of this subsection completes the analysis. We show that in any environment where $\beta = 0$ and where high δ does not imply perfect policy persistence, the core must be non-empty. Further, in the absence of perfect policy persistence, the social acceptance sets must converge to the core point, x^c. Convergence here is with respect to the Hausdorff metric (see Aliprantis and Border 1994), which for our purposes can be simplified to the following: for any compact set $Y \subseteq X$ and element $x \in X$, define the Hausdorff distance between Y and x as $h(Y, x) = \max\{\|y - x\| : y \in Y\}$. Then a sequence $\{Y^k\}$ of compact sets is said to "converge to x" if the sequence $\{h(Y^k, x)\}$ converges to zero.

Theorem 5 *Let $\{\delta^k\}$ converge to one, and let $\{\sigma^k\}$ be a corresponding sequence of simple equilibria, with social acceptance sets $\{A^k\}$. If $\min_{t \in T} \pi_t^k(A^k) < 1$ for all k, then the core is non-empty, and $\{A^k\}$ converges to x^c.*

Thus, greater patience leads to perfect policy persistence, except under rather specific conditions. When the core is empty, the typical case in multiple dimensions, perfect policy persistence necessarily obtains for discount factors close enough to one. When the core is non-empty, perfect policy persistence may not obtain: for discount factors arbitrarily close to one, there may be equilibria in which some types choose their ideal points and fail to be re-elected. In this case, however, the social acceptance sets corresponding to these equilibria, and the long run distribution of policies, must converge to the core.

Core Equivalence

It can be shown quite generally (Banks and Duggan 2008) that if the core is non-empty, then there is some type of representative that chooses the core policy and is continually re-elected. Thus, the long run distribution of policy outcomes puts positive probability on

the core point. As in Example 3, however, there may be other policy outcomes that occur with positive probability in the long run. In this subsection, we investigate the conditions under which the core point is the *only* policy selected in equilibrium, i.e., $\pi_t(\{x^c\}) = 1$ for all $t \in T$, a phenomenon we call "core equivalence." Note the implication, in particular, that all representatives must choose the same policy, which we call "policy coincidence." Our first result shows that policy coincidence, while conceptually weaker, is actually equivalent to core equivalence in equilibrium, and it gives a necessary and sufficient condition for policy coincidence to hold. Going to one dimension, we then show that, assuming sufficient patience or non-policy benefits of office (with $\delta > 0$ and $\beta > 0$), core equivalence obtains in *every* simple equilibrium, giving a strong version of the median voter theorem for repeated elections. Finally, when non-policy benefits are zero, we prove an asymptotic median voter result for patient electorates.

Suppose that in a simple strategy profile σ, all representatives choose the same policy, say, $\pi_t(\{\hat{x}\}) = 1$ for all $t \in T$. In this case, $v_t(\sigma)$ is simply equal to $u_t(\hat{x})$, and so individuals always vote to retain the incumbent and unanimity prevails. Clearly, it cannot be an equilibrium for all individuals to adopt a common policy $\hat{x}$ other than the core point: there would then be a policy y and a decisive coalition C of types such that $u_t(y) > u_t(\hat{x}) = v_t(\sigma^*)$ for all $t \in C$, and, hence, any time a member of C is elected, she would not select $\hat{x}$ as her policy. Conversely, if $\hat{x} = x^c$, then we may have an equilibrium, depending on the values of δ and β. Because x^c is the unique core point and $v_t(\sigma) = u_t(x^c)$, it follows from external stability that $A^* = \{x^c\}$, and, therefore, we need only to check whether representatives prefer compromising at x^c to choosing their ideal points. If

$$u_t(x^c) + \beta \geq (1 - \delta)(u_t(x_t) + \beta) + \delta u_t(x^c) \tag{1}$$

for all $t \in T$, then $\pi_t^*(\{x^c\}) = 1$ for all $t \in T$ is an equilibrium. If this inequality fails to hold for some $t \in T$, then this is *not* an equilibrium. Re-arranging (1), we have

$$\frac{\delta\beta}{1 - \delta} \geq u_t(x_t) - u_t(x^c).$$

Define

$$\alpha^c = \max\{u_t(x_t) - u_t(x^c) : t \in T\},$$

and note that, because $x_t \neq x_{t'}$ for all $t, t' \in T$, we have $\alpha^c > 0$. Note also that α^c is only defined when x^c exists, whereas $\overline{\alpha}$ and $\underline{\alpha}$ are always defined. Because $u_t(x^c) \geq 0$ for all $t \in T$, it must be that $\alpha^c \leq \overline{\alpha}$ when the core is non-empty. Thus, we have proved the following result.

Theorem 6 *There is a simple equilibrium σ^* such that $\pi_t^*(\{\hat{x}\}) = 1$ for all $t \in T$ if and only if the core is non-empty, $\hat{x} = x^c$, and $\delta\beta/(1 - \delta) \geq \alpha^c$.*

As an application, return to Example 3, and note that x^c exists and is at the origin. For all $t \in \{1, 2, 3, 4\}$, we have $u_t(x_t) - u_t(x^c) = 16 - 15 = 1$, and because $x_5 = x^c$,

this difference for $t = 5$ is zero. Thus, $\alpha^c = 1$. By Theorem 6, therefore, whenever $\beta \geq (1 - \delta)/\delta$, we will have a second equilibrium in which all representatives select the core point $(0, 0)$. A further consequence of Theorem 6 is that if non-policy benefits of holding office are zero or if the core is empty, then policy coincidence cannot occur in equilibrium. Because the core is generically empty in two or more dimensions, we have a negative result for policy coincidence in multiple dimensions. Of course, the core is non-empty and equal to the ideal point of the weighted median type whenever the policy space is one-dimensional, and Theorem 6 yields the following.

Corollary 1 *Assume $d = 1$. There is a simple equilibrium σ^* such that $\pi_t^*(\{x_m\}) = 1$ for all $t \in T$ if and only if $\delta\beta/(1 - \delta) \geq \alpha^c$.*

Henceforth, we focus on the one-dimensional setting, $d = 1$. From Corollary 1, it follows that as long as the individuals are sufficiently patient and non-policy benefits are sufficiently high, we can support the Downsian prediction of convergence to the median in at least *one* simple equilibrium of the one-dimensional model. In that case, clearly all types but the median compromise, and the first representative chosen remains as incumbent forever, continually implementing x_m. We next take up the issue of when core equivalence obtains in *all* simple equilibria. Example 3 shows that in multiple dimensions, other equilibria may exist for all β and δ, but we will show that a strengthening of the condition in Corollary 1 is sufficient for a unique equilibrium outcome at the core in one dimension. As a step in that direction, the next lemma shows that in one dimension, we can partition the set of equilibria into two distinct classes: either all representatives choose the median, or else some representatives do not compromise at all. Thus, in one dimension, perfect policy persistence implies core equivalence.

Lemma 1 *Assume $d = 1$, and let σ^* be a simple equilibrium. If $L(\sigma^*) = \emptyset$, then $\pi_t^*(\{x_m\}) = 1$ for all $t \in T$.*

We can now state our sufficient condition, a direct consequence of Theorem 3 and Lemma 1, for core equivalence in every simple equilibrium. It delivers a median voter result based not on competition between candidates, as in the Downsian model, but on the expectations of voters and the incentives of representatives in a dynamic electoral model with asymmetric information. The prominence of the median is sustained even with privately informed politicians' and in the absence of commitment to policy platforms.

Theorem 7 *Assume $d = 1$. If $\delta\beta/(1 - \delta) \geq \overline{\alpha}$, then there is a unique simple equilibrium, and in equilibrium, $\pi_t^*(\{x_m\}) = 1$ for all $t \in T$.*

In sum, Corollary 1 shows that policy coincidence at the core constitutes *an* equilibrium in one dimension when $\delta\beta/(1 - \delta) \geq \alpha^c$, and Theorem 7 establishes that policy coincidence at the core constitutes *the* equilibrium when δ and β satisfy the stronger restriction that $\delta\beta/(1 - \delta) \geq \overline{\alpha}$. An implication of Theorem 7 is that when $\beta > 0$, the social acceptance set A^* will be equal to the median voter's ideal point x_m when the electorate is sufficiently patient. In the stylized case of $\beta = 0$, however, no such conclusion

can be drawn: results in Banks and Duggan (2008) imply that the equilibrium acceptance set A^* is always a strict superset of x_m when $\beta = 0$. Our final result closes this issue by showing that an asymptotic version of the median voter theorem obtains even in the absence of any exogenous benefit from holding office.

Theorem 8 *Assume $d = 1$. Let $\{\delta^k\}_{k=1}^\infty$ converge to 1, and let $\{\sigma^k\}$ be a corresponding sequence of simple equilibria, with social acceptance set A^k. Then $\{A^k\}$ converges to x_m.*

Even when there are no non-policy benefits from holding office, and so policies other than the median may be sufficient to ensure re-election of the incumbent, the social acceptance set A^* does indeed converge to the weighted median x_m as δ approaches one, and the long run distribution of policies will, therefore, be close to the weighted median. The results of this section are illustrated in a numerical example in Duggan (2000). There, it is shown for particular parametric specifications that when office benefit is positive, all types of representative compromise at the median ideal point when δ is sufficiently large; when office benefit is zero, such perfect compromise is not achieved, but appears to be met in the limit. Our results inform us that these patterns are indeed regularities, which hold for completely general specifications of the one-dimensional model.

EXTENSIONS

The analysis of our model would be complicated, but our results largely unaffected, by any of several generalizations and extensions. First, Theorems 1–4, on existence and policy persistence, would continue to hold for arbitrary voting rules, such as a quota rule or even an arbitrary collection, say $\mathcal{D}$, of decisive coalitions of types. This would change the definition of $A(\sigma)$ slightly without changing its fundamental continuity properties. Our other results would continue to hold as long as the voting rule were strong in the sense defined in section "Electoral Model". This extension is perhaps uninteresting in a model of representative democracy, the main focus of our paper, but it permits a much broader interpretation of the model that captures dictatorial political systems in which the current dictator must maintain the support of a winning coalition of political elite in order to stay in power.[11] Assuming private information about preferences, the logic of our model extends in a straightforward way: the dictator faces a tradeoff between choosing a personally favorable outcome, possibly losing the support of the elite, or choosing a policy at least as good for a decisive coalition as a draw from the pool of potential dictators. Here, the benefit from holding office, β, can incorporate punishment of dictators who are removed from office (in which case, β is high).

Second, though Example 3 shows that core equivalence need not obtain in multiple dimensions (even when a core point exists and δ and β are arbitrarily high), our core equivalence results do extend to multiple dimensions for a restricted class of preferences,

---

[11]   We are grateful to Ken Shotts for pointing out this interpretation of the model.

namely, when utility functions are quadratic: $u_t(x) = k_t - ||x - x_t||^2$. Using Lemma 2.1 of Banks and Duggan (2006), we can show that the core type, if one exists, is decisive in equilibrium, in the following sense: a policy is socially acceptable if and only if it is acceptable to the core voter type. With this result, and the assumption that the core is nonempty, we can prove results exactly analogous to Lemma 1 and Theorem 7. Of course, the core is generically empty in multiple dimensions, so the interest in this version of our core equivalence result is limited. But, from Theorem 2, we know that, when the core is "close" to nonempty, equilibrium policies will be "close" (in the sense of weak convergence) to being in the core. See our working paper (Banks and Duggan 2008) for these results and an extended example.

Third, our results would hold if we added an exogenous and time-invariant positive probability of an incumbent being removed from office (through death, impeachment, etc.), though the results on policy persistence would obviously have to be re-interpreted in terms of expected duration of tenure in office. In such a model, even winners and compromisers, through no fault of their own, would eventually be replaced, generating richer dynamics for the model. The current formulation, which admittedly leads to very stark dynamics, was chosen for its simplicity and because the possibility of turnover does not essentially change our message about voter patience and non-policy benefits of office.

Finally, all of our results would hold in a version of the model with a finite number of voters and a separate, countably infinite pool of potential challengers (who do not vote) with types identically and independently distributed according to γ. Suppose a new challenger is drawn in every period to run against the incumbent. The continuation values of voters and representatives would be unchanged: what is essential is that no voter perceives a chance that she will be drawn as a challenger, and no representative perceives a chance that she will be re-drawn as a challenger after losing an election. The main advantage of this reformulation, aside from avoiding some technical complexities that arise in a model with a continuum of players, is that voters are now conceivably pivotal in elections, so that our equilibrium condition on voting strategies can be justified directly in terms of weak dominance. The obvious disadvantage is that we must treat voters as essentially different from candidates (who cannot vote). Our philosophical preference, given this trade-off, is to model candidates exactly as voters.

CONCLUSION

Our objective in this paper has been to improve our general understanding of electoral processes, with particular interest in their dynamic and informational aspects. Thus, we have proposed a model of repeated elections in which politicians determine policies in a multidimensional issue space and in which preferences are private information. Our framework captures the strategic incentives of politicians, whose private preferences and concern for re-election confronts them with a trade-off in choosing policies, and it captures the strategic calculus of voters, who must anticipate the future policy choices of incumbent politicians and challengers. We have focused on foundational issues, such as the existence of equilibria, the stability of policies over time, and the relationship

between equilibrium policies and the core. As a byproduct of our equilibrium analysis, we have shown that concern for reputation, combined with an infinite horizon, can lead to endogenous compromise by elected representatives, even without the constraints of previous commitments. Further, due to the informational asymmetry between untried challengers and incumbents, a strong form of incumbency advantage can arise in equilibrium. Finally, we have proved a version of the median voter theorem for the one-dimensional model, showing that the long run equilibrium policies of the model collapse to the median voter's ideal point as voters become patient.

APPENDIX

Theorem 1 *There exists a simple equilibrium.*

Proof: We first prove existence of an equilibrium in a modified version of the above game, and we then argue that any equilibrium of the modified game corresponds to an equilibrium in the original game. Augment the set of options available to a representative to include a "shirk" option, s, interpreted as choosing her ideal point and then sitting out the next election. If the current incumbent uses the shirk option, therefore, the voters *must* choose the challenger in the next period. We focus on equilibria in which a representative chooses the shirk option whenever her optimal choice would lose the next election, i.e., if a representative would choose her ideal point and that policy is not in the social acceptance set, then she chooses to shirk. These equilibria are distinguished from others in that representatives foresee the result of choosing $x_t \notin A^*$, taking the initiative by choosing s and declining to run, instead of choosing x_t and forcing voters to replace them.

A policy strategy for type t individuals is now a Borel probability measure $\tilde{\pi}_t$ on $\tilde{X} = X \cup \{s\}$.[12] Given a profile $\tilde{\pi} = (\tilde{\pi}_1, \ldots, \tilde{\pi}_{|T|})$, and assuming that all future representatives who do not shirk are re-elected, the continuation value of electing a challenger for a type t voter can be expressed as a function of $\tilde{\pi}$ only:

$$v_t(\tilde{\pi}) = \sum_{t' \in T} \gamma_{t'} \left[\tilde{\pi}_{t'}(\{s\})[(1-\delta)u_t(x_{t'}) + \delta v_t(\tilde{\pi})] + \int_X u_t(x)\tilde{\pi}_{t'}(dx) \right],$$

implying

$$v_t(\tilde{\pi}) = \frac{\sum_{t' \in T} \gamma_{t'} \left[\tilde{\pi}_{t'}(\{s\})(1-\delta)u_t(x_{t'}) + \int_X u_t(x)\tilde{\pi}_{t'}(dx) \right]}{1 - \delta \sum_{t' \in T} \gamma_{t'}\tilde{\pi}_{t'}(\{s\})}.$$

Note that v_t is a continuous function of $\tilde{\pi}$ with the topology of weak convergence on $\mathcal{P}(\tilde{X})$, the Borel probability measures on $\tilde{X}$. We look for an equilibrium in terms of policy strategies only, because individuals vote for the incumbent if and only if the continuation

---

[12]   We define $\tilde{Y} \subseteq \tilde{X}$ to be open if $Y \subseteq \tilde{Y} \subseteq Y \cup \{s\}$ for some open $Y \subseteq X$.

value of the incumbent is at least that of a challenger. That is, a type t individual votes to re-elect if and only if the incumbent chose a policy in the set

$$A_t(\tilde{\pi}) = \{x \in X : u_t(x) \geq v_t(\tilde{\pi})\}.$$

For each $t \in T$, the set $A_t(\tilde{\pi})$ is non-empty, compact, and convex by the continuity and concavity of u_t. If $\tilde{\pi}_t(\{s\}) < 1$, then let

$$y(\tilde{\pi}_t) = \frac{\int_X x\tilde{\pi}_t(dx)}{1 - \tilde{\pi}_t(\{s\})}$$

denote the expected outcome associated with a type t incumbent conditional on the incumbent not shirking. If $\tilde{\pi}_t(\{s\}) = 1$, then let $y(\tilde{\pi}_t)$ be defined arbitrarily. By the concavity of u_t, the policy

$$x(\tilde{\pi}) = \frac{\sum_{t' \in T} \gamma_{t'} [\tilde{\pi}_{t'}(\{s\})(1 - \delta)x_{t'} + (1 - \tilde{\pi}_{t'}(\{s\}))y(\tilde{\pi}_{t'})]}{1 - \delta \sum_{t' \in T} \gamma_{t'}\tilde{\pi}_{t'}(\{s\})}$$

therefore satisfies $u_t(x(\tilde{\pi})) \geq v_t(\tilde{\pi})$, and hence $x(\tilde{\pi}) \in A_t(\tilde{\pi})$, for all t. As in section "Simple Equilibria," for all $C \in \mathcal{D}$, define $A_C(\tilde{\pi}) = \bigcap_{t \in C} A_t(\tilde{\pi})$, also non-empty, compact, and convex. And define $A(\tilde{\pi}) = \bigcup_{C \in \mathcal{D}} A_C(\tilde{\pi})$, non-empty and compact but not necessarily convex. By arguments in the appendix of Banks and Duggan (2000), it follows that $A(\cdot)$ is a continuous correspondence on $[\mathcal{P}(\tilde{X})]^T$, the set of profiles of policy choice strategies over $\tilde{X}$.

Given $\tilde{\pi}$, an incumbent chooses a policy or shirks so as to maximize her discounted expected payoff. Thus, define $U_t(\cdot ; \tilde{\pi}) \colon \tilde{X} \to \Re$ by

$$U_t(x; \tilde{\pi}) = \begin{cases} (1 - \delta)[u_t(x_t) + \beta] + \delta v_t(\tilde{\pi}) & \text{if } x = s, \\ u_t(x) + \beta & \text{otherwise,} \end{cases}$$

and note that U_t is jointly continuous in $(x, \tilde{\pi})$. Let

$$M_t(\tilde{\pi}) \equiv \arg\max\{U_t(x; \tilde{\pi}) : x \in A(\tilde{\pi}) \cup \{s\}\}.$$

Because $A(\cdot) \cup \{s\}$ is a continuous correspondence, the Maximum Theorem implies that the correspondence $M_t \colon [\mathcal{P}(\tilde{X})]^T \rightrightarrows \tilde{X}$ has non-empty and compact values, and it is upper hemicontinuous. It is not necessarily convex-valued, however, because $A(\tilde{\pi}) \cup \{s\}$ is not convex. Let $B_t(\tilde{\pi}) = \mathcal{P}(M_t(\tilde{\pi}))$ denote the set of probability measures over optimal choices, which defines a non-empty, compact-, and convex-valued correspondence. Moreover, by Aliprantis and Border's (1994) Theorem 14.14, B_t is upper hemicontinuous. Define the correspondence $B \colon [\mathcal{P}(\tilde{X})]^T \rightrightarrows [\mathcal{P}(\tilde{X})]^T$ by

$$B(\tilde{\pi}) = B_1(\tilde{\pi}) \times B_2(\tilde{\pi}) \times \cdots \times B_{|T|}(\tilde{\pi}),$$

which inherits these properties. Because $[\mathcal{P}(\tilde{X})]^T$ is compact and convex, Glicksberg's (1952) theorem yields a fixed point of B, say $\tilde{\pi}^* = (\tilde{\pi}_1^*, \ldots, \tilde{\pi}_{|T|}^*)$. Then $\tilde{\pi}^*$, together with cut-off rules

$$\underline{u}_t^* = v_t(\tilde{\pi}^*), \quad t = 1, \ldots, |T|,$$

constitutes an equilibrium of the augmented game in which individuals either shirk or are re-elected. Finally, it is easy to see how equilibria in the augmented game translate into equilibria of the original game: for all $t \in T$ and for all measurable $Y \subseteq X$, set

$$\pi_t^*(Y) = \tilde{\pi}_t^*(\{s\})\mu_{x_t}(Y) + \tilde{\pi}_t^*(Y),$$

where μ_{x_t} is the point mass on x_t.   ∎

Let Δ denote the unit simplex in $\Re^{|T|}$, and define

$$\Delta_\circ = \left\{ \rho \in \Delta : \forall\, C \subseteq T, \sum_{t \in C} \rho_t \neq 1/2 \right\}.$$

Then define the correspondence $E \colon \Delta_\circ \times [0,1) \times \Re_+ \times \Delta \times \Lambda \rightrightarrows [\mathcal{P}(X)]^T$ such that $E(\rho, \delta, \beta, \gamma, \lambda)$ consists of the profiles of simple equilibrium policy choice strategies for model parameters $(\rho, \delta, \beta, \gamma, \lambda)$. We say that E is *upper hemicontinuous* if, for every $(\rho, \delta, \beta, \gamma, \lambda)$ in this space and for every open set $Y \subseteq [\mathcal{P}(X)]^T$ with $E(\rho, \delta, \beta, \gamma, \lambda) \subseteq Y$, there exists an open set $Z \subseteq \Delta_\circ \times [0,1) \times \Re_+ \times \Delta \times \Lambda$ with $(\rho, \delta, \beta, \gamma, \lambda) \in Z$ such that, for all $(\rho', \delta', \beta', \gamma', \lambda') \in Z$, $E(\rho', \delta', \beta', \gamma', \lambda') \subseteq Y$.

Theorem 2 *The correspondence E of simple equilibria is upper hemicontinuous in the parameters of the model.*

Proof: We first consider the augmented game defined in the proof of Theorem 1. Given parameters $(\rho, \delta, \beta, \gamma, \lambda)$, with $\rho \in \Delta_\circ$, and strategy profile $\tilde{\pi} = (\tilde{\pi}_1, \ldots, \tilde{\pi}_{|T|})$, let θ denote the vector $(\rho, \delta, \beta, \gamma, \lambda, \tilde{\pi})$, and define

$$A_t(\theta) = \{x \in X : u_t(x, \lambda) \geq v_t(\tilde{\pi}, \rho, \delta, \lambda)\},$$

where v_t is a type t individual's continuation value as defined above but using $u_t(\cdot, \lambda)$. By arguments in the appendix of Banks and Duggan (2000), it follows that

$$A(\theta) \equiv \bigcup_{C \in \mathcal{D}(\rho)} \left[\bigcap_{t \in C} A_t(\theta) \right]$$

is continuous as a correspondence at θ, where the collection $\mathcal{D}(\rho)$ of decisive coalitions of types generated by ρ is constant on an open set containing ρ, because $\rho \in \Delta_\circ$. Define $U_t(\cdot; \theta) \colon \tilde{X} \to \Re$ by

$$U_t(x; \theta) = \begin{cases} (1 - \delta)[u_t(x, \lambda) + \beta] + \delta v_t(\tilde{\pi}, \rho, \delta, \lambda) & \text{if } x = s, \\ u_t(x, \lambda) + \beta & \text{otherwise.} \end{cases}$$

Because $U_t(x; \theta)$ is continuous in (x, θ), the Maximum Theorem implies that

$$M_t(\theta) \equiv \arg\max\{U_t(x; \theta) : x \in A(\theta) \cup \{s\}\}$$

is upper hemicontinuous at θ, and therefore so is $B_t(\theta) \equiv \mathcal{P}(M_t(\theta))$. Because B_t has closed values and regular range as well, it has closed graph at θ (Aliprantis and Border 1994, Theorem 14.17). Now let $(\rho^m, \delta^m, \beta^m, \gamma^m, \lambda^m) \to (\rho, \delta, \beta, \gamma, \lambda) \in \Delta_\circ \times [0, 1) \times \Re_+ \times \Delta \times \Lambda$, take any sequence $\{\pi^m\}$ of policy choice profiles in the original game such that $\pi^m \in E(\rho^m, \delta^m, \beta^m, \gamma^m, \lambda^m)$ for all m, and suppose $\pi^m \to \pi$. Transform these into policy choice profiles, $\{\tilde{\pi}^m\}$ and $\tilde{\pi}$, in the augmented game in the obvious manner, e.g., if $x_t \notin A(\pi^m)$, then define $\tilde{\pi}_t^m(\{s\}) = \pi_t^m(\{x_t\})$. Thus, $\tilde{\pi}_t^m \in B_t(\theta^m)$ for all m and $\tilde{\pi}^m \to \tilde{\pi}$ weakly. Because B_t has closed graph at θ, we have $\tilde{\pi}_t \in B_t(\theta)$ for all $t \in T$. Therefore, $\pi \in E(\rho, \delta, \beta, \gamma, \lambda)$, and we conclude that E has closed graph. Because it has compact Hausdorff range as well, it is upper hemicontinuous (Aliprantis and Border 1994, Theorem 14.12). ∎

Theorem 4 *Let σ^* be a simple equilibrium with social acceptance set A^*. There exists $\bar{\delta} \in (0, 1)$ such that if $\delta \in [\bar{\delta}, 1)$, then $\pi_t^*(A^*) > 0$ for some $t \in T$.*

Proof: The case $\beta > 0$ is covered by Theorem 3. Assume $\beta = 0$, and suppose that we can find a sequence $\{\delta^k\}$ with $\delta^k \to 1$ and a corresponding sequence $\{\sigma^k\}$ of simple equilibria with acceptance sets $\{A^k\}$ such that $\pi_t^k(A^k) = 0$ for all $t \in T$ and all k. Hence, given any k, each type of officeholder chooses her ideal point and fails to gain re-election, and so $v_t(\sigma^k)$ is simply $\sum_{t' \in T} \gamma_{t'} u_t(x_{t'})$, which is independent of k. Denote this amount $\hat{v}_t$. Thus, the equilibrium social acceptance sets, A^k, are also independent of k, which implies $\max\{u_t(x) : x \in A^k\}$ is independent of k. Denote this amount $\hat{u}_t$. Because ideal points are distinct, utilities are strictly concave, and $\gamma_{t'} > 0$ for all $t' \in T$, we have $\hat{u}_t > \hat{v}_t$. Therefore, for all $t \in T$ and for k high enough, we have

$$\hat{u}_t > (1 - \delta^k) u_t(x_t) + \delta^k \hat{v}_t.$$

But then the optimal policy choice for type t representatives is to compromise by choosing a point in A^k, a contradiction. ∎

Theorem 5 *Let $\{\delta^k\}$ converge to one, and let $\{\sigma^k\}$ be a corresponding sequence of simple equilibria, with social acceptance sets $\{A^k\}$. If $\min_{t \in T} \pi_t^k(A^k) < 1$ for all k, then the core is non-empty, and $\{A^k\}$ converges to x^c.*

Proof: Take any sequence $\{\sigma^k\}$ of simple equilibria such that $\min_{t \in T} \pi_t^k(A^k) < 1$ for all k. We first show that the core is non-empty. From Theorem 3, we know that β must equal zero. And because T is finite, there must exist a type $t' \in T$ and a subsequence (also indexed by k) such that $\pi_{t'}^k(A^k) < 1$ for all k. It follows that representatives of type t' are willing to shirk for all k:

$$(1 - \delta^k) u_{t'}(x_{t'}) + \delta^k v_{t'}(\sigma^k) \geq \hat{u}_{t'}^k,$$

where $\hat{u}_{t'}^k = \max\{u_{t'}(x) : x \in A^k\}$. Because $\{v_{t'}(\sigma^k)\}$ and $\{\hat{u}_{t'}^k\}$ lie in compact sets, we may go to a subsequence (also indexed by k) along which these sequences converge. It follows that

$$\lim v_{t'}(\sigma^k) \geq \lim \hat{u}_t^k. \tag{A.1}$$

Because $\hat{u}_{t'}^k \geq v_{t'}(\sigma^k)$ for all k, (A.1) actually holds with equality. Now let ψ^k be the distribution on X associated with the kth equilibrium, so that

$$v_{t'}(\sigma^k) = \int u_{t'}(x)\psi^k(dx)$$

for all k. Because X is compact, $\{\psi^k\}$ has a subsequence (also indexed by k) that converges weakly to some probability measure ψ on X. By weak convergence,

$$\lim v_{t'}(\sigma^k) = \int u_{t'}(x)\psi(dx). \tag{A.2}$$

Let

$$\overline{x}(\psi^k) = \int x\psi^k(dx) \quad \text{and} \quad \overline{x}(\psi) = \int x\psi(dx),$$

and note that $\overline{x}(\psi^k) \in A^k$ for all k by concavity of voter utility functions. Hence,

$$\hat{u}_{t'}^k \geq u_{t'}(\overline{x}(\psi^k)) \tag{A.3}$$

for all k. Using $\overline{x}(\psi^k) \to \overline{x}(\psi)$ and the continuity of $u_{t'}$, (A.1), (A.2) and (A.3) yield

$$\int u_{t'}(x)\psi(dx) = \lim v_{t'}(\sigma^k) \geq \lim \hat{u}_{t'}^k \geq \lim u_{t'}(\overline{x}(\psi^k)) = u_{t'}(\overline{x}(\psi)).$$

From strict quasi-concavity, we conclude that ψ is concentrated on some point $\hat{x}$, i.e., $\psi(\{\hat{x}\}) = 1$. We claim that $\hat{x}$ is a core point, i.e., $\hat{x} = x^c$. If not, then there exist $y \in X$ and $C \in \mathcal{D}$ such that, for all $t \in C$, $u_t(y) > u_t(\hat{x})$. Because

$$\lim v_t(\sigma^k) = \int u_t(x)\psi(dx) = u_t(\hat{x}),$$

we have $y \in A^k$ for high enough k. Also, we have

$$u_t(y) > (1 - \delta^k)u_t(x_t) + \delta^k v_t(\sigma^k)$$

for all $t \in C$ when k is high enough. This is implies that, for all $t \in C$, $\pi_t^k(A^k) = 1$ when k is high enough. Let $Y \subseteq X$ be any open set such that $\hat{x} \in Y$ and for all $t \in C$ and all $z \in Y$, $u_t(y) > u_t(z)$. For each $t \in C$, clearly $\pi_t^k(A^k) = 1$ implies $\pi_t^k(A^k \setminus Y) = 1$. But then ψ^k does not converge weakly to the point mass at $\hat{x}$, contradicting the above result. We now show that $\{A^k\}$ converges to the core. If not, then there is an open set Y with $x^c \in Y$ and a subsequence (also indexed by k) $\{x^k\}$ such that, for all k, $x^k \in A^k \cap (X \setminus Y)$. Because $X \setminus Y$ is compact, there is a subsequence (also indexed by k) and a policy $\tilde{x} \in X \setminus Y$ such that $x^k \to \tilde{x}$. Because $\mathcal{D}$ is finite, we may suppose (going to a subsequence if necessary) there is some $C \in \mathcal{D}$ such that, for all k, $x^k \in A_C(\sigma^k)$. Thus, for all $t \in C$, $u_t(x^k) \geq v_t(\sigma^k)$. Now, by our first argument, we may choose a subsequence (also indexed by k) with continuation distributions $\{\psi_k\}$ converging to the point mass on x^c, and, therefore, $v_t(\sigma^k) \to u_t(x^c)$ for all $t \in T$. Then, by continuity, we have $u_t(\tilde{x}) \geq u_t(x^c)$ for all $t \in C$. But then strict quasi-concavity implies $u_t((1/2)\tilde{x} + (1/2)x^c) > u_t(x^c)$ for all $t \in C$, a contradiction. Therefore, $A^k \to \{x^c\}$. ∎

Lemma 1 *Assume $d = 1$, and let σ^* be a simple equilibrium. If $L(\sigma^*) = \emptyset$, then $\pi_t^*(\{x_m\}) = 1$ for all $t \in T$.*

Proof: Suppose $L(\sigma^*) = \emptyset$, i.e., all types propose policies in A^* and, hence, are re-elected. Let $S(\sigma^*)$ denote the support of the policy strategies in σ^*, i.e., the smallest closed subset of X with probability one under π_t for all $t \in T$, and define

$$\underline{p} = \min_{\geq} S(\sigma^*) \quad \text{and} \quad \bar{p} = \max_{\geq} S(\sigma^*).$$

If $\underline{p} = \bar{p}$, i.e., $S(\sigma^*) = 1$, then we know from Theorem 6 that $S(\sigma^*) = \{x_m\}$, and so it must be that $\pi_t^*(\{x_m\}) = 1$ for all $t \in T$. If $\underline{p} < \bar{p}$, i.e., $S(\sigma^*) > 1$, then by strict quasi-concavity of u_t, we have

$$\arg\min\{u_t(x) : x \in S(\sigma^*)\} \subseteq \{\underline{p}, \bar{p}\}$$

for all $t \in T$. Define $\underline{T} \subseteq T$ as those types for which this minimizer is unique and equal to $\underline{p}$, $\bar{T} \subseteq T$ as the types for which this is unique and equal to $\bar{p}$, and $I \subseteq T$ as the types for which $u_t(\underline{p}) = u_t(\bar{p})$. Because $x(\sigma^*) \in A^*$, we have

$$\max\{u_t(x) : x \in S(\sigma^*)\} \geq v_t(\sigma^*) \geq \min\{u_t(x) : x \in S(\sigma^*)\} \tag{A.4}$$

for all $t \in T$. By strict quasi-concavity and our assumption that $\gamma_{t'} > 0$ for all $t' \in T$, we then have $u_t(x(\sigma^*)) > v_t(\sigma^*)$, and because $x(\sigma^*) \in A(\sigma^*)$, this implies

$$\max\{u_t(x) : x \in A(\sigma^*)\} > v_t(\sigma^*)$$

for all $t \in T$. And this, with $\gamma_t > 0$, implies that both of the inequalities are strict in (A.4) for all $t \in T$. In particular, $v_t(\sigma^*) > u_t(\underline{p}) = u_t(\bar{p})$ for all $t \in I$. Therefore, all $t \in I$ vote against incumbents choosing policies $\underline{p}$ and $\bar{p}$, and so the only types voting in favor of the policy $\underline{p}$ are those in $\bar{T}$, and the only types voting in favor of $\bar{p}$ are those in $\underline{T}$. Each of these policies is accepted, so it must be that $\bar{T} \in \mathcal{D}$ and $\underline{T} \in \mathcal{D}$, i.e., $\sum_{t \in \bar{T}} \rho_t > 1/2$ and $\sum_{t \in \underline{T}} \rho_t > 1/2$. However, $\bar{T} \cap \underline{T} = \emptyset$, a contradiction. ∎

Theorem 8 *Assume $d = 1$. Let $\{\delta^k\}_{k=1}^{\infty}$ converge to 1, and let $\{\sigma^k\}$ be a corresponding sequence of simple equilibria, with social acceptance set A^k. Then $\{A^k\}$ converges to x_m.*

Proof: Suppose that there exists a subsequence of $\{A^k\}$ (also indexed by k) and some $\varepsilon > 0$ such that for all k, $d(A^k, \{x^c\}) \geq \varepsilon$. It follows from Theorem 7 that $\beta = 0$, and Lemma 1 then implies that $L(\sigma^k) \neq \emptyset$ for all k. By Theorem 5, we have $A^k \to \{x^c\}$, a contradiction. ∎

REFERENCES

Alesina, A., 1988. "Credibility and Convergence in a Two-Party System with Rational Voters." *American Economic Review* 78: 796–805.

Aliprantis, C., and K. Border. 1994. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Berlin: Springer-Verlag.

Austen-Smith, D., and J. Banks. 1989. "Electoral Accountability and Incumbency." In *Models of Strategic Choice in Politics*, ed. P. Ordeshook, Ann Arbor: University of Michigan Press.

Banks, J., and J. Duggan. 2000. "A Bargaining Model of Collective Choice." *American Political Science Review* 94: 73–88.

Banks, J., and J. Duggan. 2001. A Multidimensional Model of Repeated Elections, W. Allen Wallis Institute of Political Economy, *Working Paper No. 24.*

Banks, J., and J. Duggan. 2006. "A Social Choice Lemma on Voting Over Lotteries with Applications to a Class of Dynamic Games." *Social Choice and Welfare* 26: 285–304.

Banks, J., and J. Duggan. 2008. "A Dynamic Model of Democratic Elections in Multidimensional Policy Spaces." W. Allen Wallis Institute of Political Economy, *Working Paper No. 24A.*

Banks, J., and R. Sundaram. 1993. "Moral Hazard and Adverse Selection in a Model of Repeated Elections." In *Political Economy: Institutions, Information, Competition and Representation*, eds. W. Barnett *et al.*, New York: Cambridge University Press.

Banks, J., and R. Sundaram. 1998. "Optimal Retention in Agency Problems." *Journal of Economic Theory* 82: 293–323.

Baron, D. 1991. "A Spatial Bargaining Theory of Government Formation in Parliamentary Democracies." *American Political Science Review* 85: 137–164.

Barro, R., 1973. "The Control of Politicians: An Economic Model." *Public Choice* 14: 19–42.

Besley, T., and S. Coate. 1997. "An Economic Model of Representative Democracy." *Quarterly Journal of Economics* 112: 85–114.

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

Bernhardt, D., L. Campuzano, and F. Squintani. 2005. "On Benefits of Party Competition." *Working Paper.*

Campuzano, L. 2005. "Parties in a Repeated Elections Model." *Working Paper.*

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

Downs, A. 1957. *An Economic Theory of Democracy*. New York: Harper and Row.

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

Duggan, J., and M. Fey. 2006. "Repeated Downsian Electoral Competition." *International Journal of Game Theory* 35: 39–69.

Fearon, J. 1999. "Electoral Accountability and the Control of Politicians: Selecting Good types versus Sanctioning Poor Performance." In *Democracy, Accountability, and Representation*, (eds. B. Manin *et al.*), New York: Cambridge University Press.

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.

Glicksberg, I. 1952. "A Further Generalization of the Kakutani Fixed Point Theorem, with Applications to Nash Equilibrium." *Proceedings of the American Mathematical Society* 3: 170–174.

Judd, K. 1985. "The Law of Large Numbers with a Continuum of IID Random Variables." *Journal of Economic Theory* 35: 19–25.

Kang, I. 2005. "Strategic Challengers and the Incumbency Advantage." *Working Paper.*

Kramer, G. 1977. "A Dynamical Model of Political Equilibrium." *Journal of Economic Theory* 16: 310–334.

Krehbiel, K. 1998. *Pivotal Politics: A Theory of U.S. Lawmaking.* Chicago: University of Chicago Press.

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 Political Competition with Citizen–Candidates." *Quarterly Journal of Economics* 111: 65–96.

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

Riker, W. 1980. "Implications from the Disequilibrium of Majority Rule for the Study of Institutions." *American Political Science Review* 74: 432–446.

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