

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

A SOCIAL CHOICE LEMMA ON VOTING OVER LOTTERIES WITH APPLICATIONS TO A CLASS OF DYNAMIC GAMES

Jeffrey S. Banks
California Institute of Technology

John Duggan
University of Rochester



SOCIAL SCIENCE WORKING PAPER 1163

May 2003

A Social Choice Lemma on Voting over Lotteries with Applications to a Class of Dynamic Games

Jeffrey S. Banks

John Duggan

Abstract

We prove a lemma characterizing majority preferences over lotteries on a subset of Euclidean space. Assuming voters have quadratic von Neumann-Morgenstern utility representations, and assuming existence of a majority undominated (or “core”) point, the core voter is decisive: one lottery is majority-preferred to another if and only if this is the preference of the core voter. Several applications of this result to dynamic voting games are discussed.

JEL classification numbers: C73, D71, D72

Key words: lotteries, dynamic games, representative voter, decisive voter

A Social Choice Lemma on Voting over Lotteries with Applications to a Class of Dynamic Games[‡]

Jeffrey S. Banks

John Duggan

1 Introduction

The prominence of electoral systems based on majority rule has prompted much research on the nature of majority preferences. McGarvey (1953) has shown that, when the set of alternatives is finite, asymmetry is the only property displayed by strict majority preferences across electorates with varying size and preferences. When the size and preferences of the electorate are fixed, more structure must be imposed on the social choice environment (the set of alternatives and the form of individual preferences) to deduce a sharper characterization. The spatial model, where the set of alternatives is a convex subset of finite-dimensional Euclidean space and individual preferences are subject to various regularity conditions, provides such an environment.

Work initiated by Plott (1967) and McKelvey (1976, 1979, 1986) has, for example, established strong necessary conditions required for the existence of a majority-undominated (or “core”) point in multiple dimensions, and it has provided characterizations of several sets (the top cycle and uncovered set) derived from the majority preference relation. Davis, DeGroot, and Hinich (1972), in contrast, begin with the assumption that the core is non-empty and completely characterize the majority preference relation when individual preferences are Euclidean, i.e., the preferred of two alternatives is the one closer to a voter’s “ideal point,” and when an additional regularity condition holds.¹ Their Theorem 4 establishes that the majority weak preference relation is transitive and, in fact, coincides with the weak preference relation of the “core voter.” That is, the majority-preferred of any two alternatives is the one closer to the core. While this is strictly a social choice theory result, it has proved useful in game-theoretic models of elections as well (cf. Calvert (1985) and Duggan and Fey (2001)).

*Support from the National Science Foundation, grant numbers SES-9975173 and SES-0213738, is gratefully acknowledged.

[‡]This paper was completed after Jeff Banks’s death. I am deeply indebted to him for his friendship and his collaboration on this and many other projects.

¹They assume that, for every direction, there is a unique “median hyperplane,” a condition that holds if the number of individuals is odd or if ideal points are distributed according to a positive density. It implies that there is at most one core point.

In this paper, we strengthen Davis, DeGroot, and Hinich's (1972) assumption on individual preferences and extend their conclusion to majority preferences over lotteries. We assume not only that individual preferences over pure alternatives are Euclidean, but also that they have quadratic numerical representations that extend individual preferences to the space of lotteries over alternatives. If the core is non-empty and the regularity condition of the latter authors holds, then, once again, the core voter is decisive: the majority-preferred of two alternatives is the one that yields the highest expected utility, evaluated using the quadratic utility function with ideal point at the core. The proof of the result is rather simple, relying on mean-variance analysis, a property of quadratic utilities. Because majority preferences over pure alternatives are unaffected by arbitrary monotonic transformations of quadratic utilities, Davis, DeGroot, and Hinich's (1972) result for Euclidean preferences follows as a corollary.

While our assumption of quadratic preferences is restrictive, of course, it is widely used in theoretical and empirical work. We demonstrate the usefulness of our lemma with a result for a class of dynamic games, in which play proceeds in stages and determines a sequence of outcomes over time evaluated by the players according to discounted quadratic utility. We assume a common discount factor; we allow for incomplete information, so that players actually must evaluate lotteries over sequences of outcomes; and we allow for a continuum of players, to capture models of large electorates. The result is that majority preferences over action profiles in any stage coincide with the preferences of the "core player," appropriately defined. In the context of a binary voting stage, in which the continuation of the game depends only on which of two alternatives receives majority support, and assuming players eliminate stage-dominated voting strategies, this means that the alternative preferred by the core player will win.

We illustrate with several theoretical applications. Among them, we consider a one-dimensional version of the infinite-horizon bargaining model of Baron and Ferejohn (1989) with an arbitrary status quo policy, as in Banks and Duggan (2003). Assuming quadratic utilities, our results imply that the approval of the median voter is necessary and sufficient for a proposal to pass, an observation used in Banks and Duggan (2000) to prove "core equivalence" under the assumption of perfect patience. We also consider a model of repeated elections related to the family of models analyzed in Duggan (2000), Bernhardt, Hughson, and Dubey (2002), and Banks and Duggan (2002). Applied here, the lemma generalizes observations in the latter papers that the median voter is decisive, in the sense that an incumbent is reelected if and only if the median voter prefers it. Thus, the lemma delivers a type of centrality of the median voter in these models close to (but distinct from) the median voter theorem of Black (1958) and Downs (1957).

Technically, by characterizing the set of proposals that can pass (in the bargaining model) and the policies that lead to reelection (in repeated elections models), the lemma allows us to reformulate the optimization problem of the proposer or officeholder, thereby simplifying the task of working through analytic examples, the numerical computation of stationary equilibria, and the determination of comparative statics in quadratic models. The lemma has applications to other models of repeated elections (Campuzano (2001),

Duggan and Fey (2002), Meirowitz (2002)), and it offers to facilitate the investigation of topics such as uniqueness of stationary equilibria in the one-dimensional bargaining model (Cho and Duggan (2002)) and the possibility of multiple parties in models of repeated elections.

2 A Social Choice Lemma

Suppose a set of individuals must choose from a subset X of Euclidean space of finite dimension d . Each individual i has a Euclidean preference weak relation \succsim_i on X , i.e., there exists an ideal point $\tilde{x}_i \in \mathbb{R}^d$ such that, for all $x, y \in X$, $x \succsim_i y$ if and only if $\|x - \tilde{x}_i\| \leq \|y - \tilde{x}_i\|$. Because an individual's preference relation is determined by his or her ideal point, we write $\succsim_{\tilde{x}}$ for the preference relation corresponding to ideal point \tilde{x} . We use $\succ_{\tilde{x}}$ and $\sim_{\tilde{x}}$ to denote the associated strict preference and indifference relations. (This notation for strict preference and indifference is applied to other relations throughout the paper.) Suppose that individual ideal points are distributed across \mathbb{R}^d according to the probability measure π , which may be non-atomic to capture a continuum of individuals. We say x is *weakly majority-preferred* to y , written $x \succsim y$, if

$$\pi(\{\tilde{x} \mid x \succsim_{\tilde{x}} y\}) \geq \frac{1}{2}.$$

The *majority core* consists of the points $x \in X$ that are weakly preferred to all others, i.e., $x \succsim y$ for all $y \in X$.

Let Λ denote the set of lotteries, or Borel probability measures, on X , and let \succsim_i^* denote the extension of i 's preference relation to this space. Suppose the preferences of an individual with ideal point \tilde{x} are extended to lotteries on X by the quadratic von Neumann-Morgenstern utility function $u_{\tilde{x}}(x) = -\|x - \tilde{x}\|^2$. Following the above convention, let $\succsim_{\tilde{x}}^*$ denote the preferences over lotteries corresponding to ideal point \tilde{x} . Thus, for all $\lambda, \mu \in \Lambda$, $\lambda \succsim_{\tilde{x}}^* \mu$ if and only if

$$\int u_{\tilde{x}}(x) \lambda(dx) \geq \int u_{\tilde{x}}(x) \mu(dx).$$

We say λ is weakly majority-preferred to μ , written $\lambda \succsim^* \mu$, if

$$\pi(\{\tilde{x} \mid \lambda \succsim_{\tilde{x}}^* \mu\}) > \frac{1}{2}.$$

Our goal in this section is to characterize the majority preference relation on lotteries in terms of the preferences of the “core voter.”

A special property of the quadratic extension is that expected utility from a lottery λ can be decomposed into two components, one depending on the mean of λ and the other depending on its variance. Let $m = \int x \lambda(dx)$ denote the mean of λ , and let

$$v = \int (x - m) \cdot (x - m) \lambda(dx)$$

denote the variance. Then the expected utility of λ for an individual with ideal point \tilde{x} is

$$\int u_{\tilde{x}}(x) \lambda(dx) = -\|m - \tilde{x}\| - v.$$

Note that the disutility from the variance of λ is independent of the individual's ideal point.

Given $x, t \in \mathbb{R}^d$, the open half-space at x in direction t is

$$H_t^+(x) = \{y \in \mathbb{R}^d \mid y \cdot t > x \cdot t\}.$$

We say x is a *median in all directions* if, for every direction t , $\pi(H_t^+(x)) \leq 1/2$. Every median in all directions, if contained in X , is a majority core point. Furthermore, if $x \in \text{int}X$ is a majority core point, then it is a median in all directions (cf. Davis, DeGroot, and Hinich, 1972): if the proportion of individuals in $H_t^+(x)$ were greater than one half, then a small enough move from x in the t direction would result in a majority-preferred point, an impossibility.

We say π is *resolute* at $x \in \mathbb{R}^d$ if, for every direction t with $\pi(H_t^+(x)) \leq 1/2$ and for every $\epsilon > 0$,

$$\pi(H_t^+(x + \epsilon t)) < \frac{1}{2}.$$

In words, given any hyperplane through x that evenly divides the ideal points of the individuals, an arbitrarily small shift of the hyperplane will put strictly more than half of the individuals to one side.² This condition clearly holds at every $x \in \mathbb{R}^d$ if the set of individuals is finite and odd in number.

If π is non-atomic with strictly positive density on some open set containing x , then the following holds: for every open set G around x and for every direction t , we have $\pi(G \cap H_t^+(x)) > 0$. Furthermore, the latter condition implies that π is resolute at x : given direction t with $\pi(H_t^+(x)) \leq 1/2$ and given $\epsilon > 0$, let G be any open set around x with $G \cap H_t^+(x + \epsilon t) = \emptyset$ to fulfill the definition of resoluteness. Thus, if π is supported on an open set around x , then it is resolute at x .

If x is a median in all directions and π is resolute at x , then x is the unique median in all directions; in fact,

$$\pi(\{\tilde{x} \mid x \succ_{\tilde{x}} y\}) > \frac{1}{2}$$

for all $y \neq x$. To see this, take any y distinct from x . Then $\pi(H_{y-x}^+(x)) \leq 1/2$, but then, because π is resolute at x , we have $\pi(H_{y-x}^+(z)) < 1/2$ for $z = (1/2)(x + y)$, which implies $\pi(H_{x-y}^+(y)) > 1/2$. Thus, y is not a median in all directions.

²Assuming existence of a median in all directions, x , resoluteness at x is equivalent to uniqueness of a median hyperplane in all directions, which is used by Davis, DeGroot, and Hinich (1972).

The main result of this paper is the next lemma. It is stated in terms of a median in all directions for maximum strength, but the primary interest is in the application to (interior) majority core points. Part (i) shows that, if one lottery is weakly preferred to another by the core voter, then the first is weakly majority-preferred to the second. Moreover, any two lotteries that are equally desirable to the core voter are majority-indifferent. This result, which generalizes Davis, DeGroot, and Hinich's (1972) Theorem 3, does not show that a lottery strictly preferred by the core voter is strictly majority preferred, but it does not assume π is resolute. It therefore holds even if there are multiple core points. Part (ii) strengthens the conclusion of the first part: an implication is that, if $\pi(\{x\}) = 0$ or there is a finite, odd number of individuals, then any two lotteries that are equally desirable to the core voter must have equal proportions of voters with opposing strict preferences over them. It is closely related to Theorem 2 of McKelvey, Ordeshook, and Ungar (1980), which generalizes Plott's (1967) necessary symmetry conditions at a core point.³ Finally, part (iii) gives a full characterization of majority preferences over lotteries when there is a majority core point at which π is resolute (and is, therefore, the unique core point): majority preferences are identical to the core voter's.

Lemma 1 *Let x be a median in all directions. Then (i) $\lambda \succsim_x^* \lambda'$ implies $\lambda \succsim^* \lambda'$, (ii) if $\lambda \sim_x^* \lambda'$, then*

$$|\pi(\{\tilde{x} \mid \lambda \succ_{\tilde{x}} \lambda'\}) - \pi(\{\tilde{x} \mid \mu \succ_{\tilde{x}} \lambda\})| \leq \pi(\{x\}),$$

and (iii) if π is resolute at x , then $\lambda \succsim_x^ \lambda'$ if and only if $\lambda \succsim^* \lambda'$.*

To prove part (i) of the lemma, take any $\lambda, \lambda' \in \Lambda$ such that $\lambda \succsim_x^* \lambda'$. Letting m and m' denote the means of these lotteries and letting v and v' denote their variances, we then have,

$$u_x(m) - u_x(m') \geq v - v'. \quad (1)$$

Since x is a median in all directions, we have $\pi(H_{m'-m}^+(x)) \leq 1/2$. Note that, for all $\tilde{x} \notin H_{m'-m}^+(x)$, we have

$$\tilde{x} \cdot (m - m') \geq x \cdot (m - m'),$$

implying

$$u_{\tilde{x}}(m) - u_{\tilde{x}}(m') \geq u_x(m) - u_x(m'). \quad (2)$$

Then, by (1) and (2), we have

$$u_{\tilde{x}}(m) - u_{\tilde{x}}(m') \geq v - v'$$

for all $\tilde{x} \notin H_{m'-m}^+(x)$. Therefore, $\lambda \succsim^* \lambda'$, as required.

³In contrast to McKelvey, Ordeshook, and Ungar (1980), we allow the distribution of ideal points to be discrete, in order to capture a finite number of individuals.

To prove part (ii), take any $\lambda, \lambda' \in \Lambda$ such that $\lambda \sim_x^* \lambda'$. Using the notation above, let m and m' denote the means of these lotteries and v and v' the variances. We then have

$$u_x(m) - u_x(m') = v - v'.$$

Note that $H_{m-m'}^+(x) = \{\tilde{x} \mid \lambda \succ_{\tilde{x}} \lambda'\}$ and $H_{m'-m}^+(x) = \{\tilde{x} \mid \lambda' \succ_{\tilde{x}} \lambda\}$, let

$$\epsilon = \pi(H_{m-m'}^+(x)) - \pi(H_{m'-m}^+(x)) - \pi(\{x\}), \quad (3)$$

and suppose $\epsilon > 0$. For any $t \in \mathbb{R}^d$, let

$$H_t(x) = \{y \in \mathbb{R}^d \mid y \cdot t = x \cdot t\}$$

denote the hyperplane through x with normal t . Take any $t \in \mathbb{R}^d$ such that

$$\pi(H_{m-m'}(x) \cap H_t(x) \setminus \{x\}) < \frac{\epsilon}{2}. \quad (4)$$

(The existence of such a t follows from an argument similar to that used in the proof of Theorem 2 in McKelvey, Ordeshook, and Ungar (1980, pp.164-165).) Without loss of generality, assume that

$$\pi(H_{m-m'}(x) \cap H_t^+(x)) \geq \pi(H_{m-m'}(x) \cap H_{-t}^+(x)). \quad (5)$$

Define $y(n) = m + (1/n)t$, let f_n denote the indicator function of $H_{y(n)-m'}^+(x)$, and note that the sequence $\{f_n\}$ of indicator functions converges pointwise to the indicator function of $H_{m-m'}(x)$. Thus, Egoroff's Theorem (cf. Kingman and Taylor (1966), Theorem 7.1) yields n such that

$$\pi(H_{m-m'}^+(x) \setminus H_{y(n)-m'}^+(x)) < \frac{\epsilon}{4}. \quad (6)$$

It is an identity that

$$\begin{aligned} 1 &= \pi(H_{m-m'}^+(x)) + \pi(H_{m'-m}^+(x)) + \pi(H_{m-m'}(x) \cap H_t^+(x)) \\ &\quad + \pi(H_{m-m'}(x) \cap H_{-t}^+(x)) + \pi(H_{m-m'}(x) \cap H_t(x) \setminus \{x\}) + \pi(\{x\}). \end{aligned}$$

Solving for $\pi(H_{m-m'}^+(x))$ in (3) and substituting, and using (5), we have

$$\begin{aligned} 1 &\leq 2\pi(H_{m'-m}^+(x)) + 2\pi(H_{m-m'}(x) \cap H_t^+(x)) + 2\pi(\{x\}) \\ &\quad + \pi(H_{m-m'}(x) \cap H_t(x) \setminus \{x\}) + \epsilon. \end{aligned}$$

Using (4) the latter inequality yields

$$\pi(H_{m'-m}^+(x)) + \pi(\{x\}) + \pi(H_{m-m'}(x) \cap H_t^+(x)) + \frac{3\epsilon}{4} > \frac{1}{2}. \quad (7)$$

From (6) and (3), respectively, we have

$$\begin{aligned} \pi(H_{y(n)-m'}^+(x) \cap H_{m-m'}^+(x)) &> \pi(H_{m-m'}^+(x)) - \frac{\epsilon}{4} \\ &= \pi(H_{m'-m}^+(x)) + \pi(\{x\}) + \frac{3\epsilon}{4}. \end{aligned} \quad (8)$$

Given $\tilde{x} \in H_{m-m'}(x) \cap H_t^+(x)$, note that

$$(\tilde{x} - x) \cdot (y(n) - m') = (\tilde{x} - x) \cdot (m + (1/n)t - m') = (1/n)(\tilde{x} - x) \cdot t > 0.$$

Therefore, $H_{m-m'}(x) \cap H_t^+(x) \subseteq H_{y(n)-m'}^+(x)$, which implies

$$\pi(H_{y(n)-m'}^+(x) \cap H_{m-m'}(x)) \geq \pi(H_{m-m'}(x) \cap H_t^+(x)). \quad (9)$$

Finally, combining (7), (8), and (9), we have $\pi(H_{y(n)-m'}^+(x)) > 1/2$, but then x is not a median in all directions, a contradiction. A symmetric argument addresses the case $\pi(H_{m'-m}^+(x)) > \pi(H_{m-m'}^+(x)) + \pi(\{x\})$, as required.

To prove part (iii), take any $\lambda, \lambda' \in \Lambda$ such that $\lambda \succ_x^* \lambda'$, i.e.,

$$u_x(m) - u_x(m') > v - v'. \quad (10)$$

Since x is a median in all directions, we have $\pi(H_{m'-m}^+(x)) \leq 1/2$. Because π is resolute at x , we have $\pi(H_{m'-m}^+(x(\epsilon))) < 1/2$ for all $\epsilon > 0$, where $x(\epsilon) = x + \epsilon(m' - m)$. Note that, for all $\tilde{x} \notin H_{m'-m}^+(x(\epsilon))$, we have

$$\tilde{x} \cdot (m - m') \geq x(\epsilon) \cdot (m - m'),$$

implying

$$u_{\tilde{x}}(m) - u_{\tilde{x}}(m') \geq u_{x(\epsilon)}(m) - u_{x(\epsilon)}(m'). \quad (11)$$

Pick ϵ small enough, using continuity and (10), so that

$$u_{x(\epsilon)}(m) - u_{x(\epsilon)}(m') > v - v'. \quad (12)$$

Then, by (11) and (12), we have

$$u_{\tilde{x}}(m) - u_{\tilde{x}}(m') > v - v'$$

for all $\tilde{x} \notin H_{m'-m}^+(x(\epsilon))$. Because $\pi(H_{m'-m}^+(x(\epsilon))) < 1/2$, this implies that $\lambda \succ^* \lambda'$, as required. This completes the proof.

While we have stated Lemma 2.1 for majority rule, it is easily extended to the class of weighted majority rules by applying the lemma to the “weighted” distribution of ideal points. Specifically, suppose that (N, Σ) is a measurable space of voters; suppose that $\tilde{x}: N \rightarrow X$ is a measurable mapping such that \tilde{x}_i is the ideal point of individual i ; and suppose that ω is a probability measure on Σ assigning weights to measurable coalitions $C \in \Sigma$ such that $x \succsim y$ if and only if

$$\omega(\{i \in N \mid x \succsim_i y\}) \geq \frac{1}{2}.$$

Then, defining $\pi = \omega \circ \tilde{x}^{-1}$ as the weighted distribution of ideal points, the weighted majority rule is formally equivalent to majority rule with distribution π . The lemma extends to arbitrary simple voting rules, those generated by a “monotonic” and “proper” class of winning coalitions, though parts (ii) and (iii) require that the rule be “strong” and continuous in a certain sense. (See Banks, Duggan, and Le Breton (2003) for this framework.) We omit the details of this extension.

3 A Class of Dynamic Games

Suppose a (possibly finite) measurable set $N \subseteq [0, 1]$ of individuals play an extensive form game of incomplete information with the following structure. Each individual has a measurable set $T \subseteq \mathbb{R}^e$ of types. The set of type profiles, the measurable mappings $\rho: N \rightarrow T$, is endowed with a sigma algebra, which may vary depending on the application, and a common prior probability measure.

The extensive form is a stage game, where, given the initial history h_0 , a measurable subset $N(h_0)$ of players simultaneously choose actions in the set $A(h_0)$, while all other players choose a null action \bar{a} . Denoting the profile of actions by active players as $a \in A(h_0)^{N(h_0)}$, we have the history $h_1 = (h_0, a)$ of length one. After h_1 , a measurable set $N(h_1)$ of players then simultaneously choose actions in $A(h_1)$ while others play the null action, producing a history of length two, and so on. Because the strategy sets of the players are arbitrary, we may assume without loss of generality that this process continues *ad infinitum*. Let H^k denote the set of histories of length k ; let $H = \bigcup_{k=1}^{\infty} H^k$ denote the set of all finite histories of the game; and let H^{∞} denote the set of infinite histories. Finally, let $A = \bigcup_{h_k \in H} A(h_k) \cup \{\bar{a}\}$ denote the set of all possible actions. The sets H^k must be endowed with appropriate sigma algebras, which may vary with the application and which we leave undefined. See Application 4.3 for an example.

Each type t corresponds to an ideal point \tilde{x}_t in \mathbb{R}^d . We assume the mapping $t \mapsto \tilde{x}_t$ is measurable and one-to-one, and we simply identify types with ideal points in the sequel. We assume every finite history h_k generates a sequence $f^k(h_k)$ of length l_k in a bounded set $X \subseteq \mathbb{R}^d$, e.g., $f^k(h_k) = (x_1, x_2, \dots, x_{l_k})$, where each f^k is a measurable mapping from H^k to X^{l_k} , the set of sequences in X of length l_k . We impose the consistency condition that the continuation of a history simply extends the sequence of outcomes in X : given histories h_k and $h_{k+1} = (h_k, a)$, we have $f_n^k(h_k) = f_n^{k+1}(h_{k+1})$ for all $n = 1, 2, \dots, l_k$. Without loss of generality, we assume that $l_k \rightarrow \infty$. Thus, every infinite history corresponds to an infinite sequence of points in X , denoted $\mathbf{x} = (x_1, x_2, \dots)$. Given an infinite history and an associated sequence of outcomes, a type t player's payoff in the game is

$$U_t(\mathbf{x}) = (1 - \delta) \sum_{k=1}^{\infty} \delta^{k-1} u_t(x_k),$$

where u_t is the quadratic utility function with ideal point \tilde{x}_t and $\delta \in [0, 1)$ is a common discount factor.

A *strategy profile* is a mapping $\sigma: N \times T \times H \rightarrow A$ such that $\sigma_i(t, h_k) \in A(h_k)$ for all $i \in N(h_k)$ and $\sigma_i(t, h_k) = \bar{a}$ for all $i \notin N(h_k)$. Here, of course, $\sigma_i(t, h_k)$ is the action taken by type t of player i after history h_k . In applications, σ_i must also respect a player's information sets, not specified above. That is, if (h_k, a) and (h_k, a') are not distinguishable to player i , and so lie in the same information set, then we must have $\sigma_i(t, (h_k, a)) = \sigma_i(t, (h_k, a'))$.

A *belief assessment* is a mapping β from $N \times T \times H$ to probability measures on type profiles, where $\beta_i(t, h_k)$ represents the beliefs of type t of player i about the types of all players after history h_k . In applications, beliefs must be consistent with Bayes rule and must respect the players' information sets. For each i, t , and h_k , the beliefs β , together with a strategy profile σ , induces a probability measure $\gamma_i^{k+1}(t, h_k)$ on histories of length $k + 1$ defined by

$$\gamma_i^{k+1}(t, h_k)(G) = \beta_i(t, h_k)(\{\rho \mid (h_k, (\sigma_i(\rho(i), h_k))_{i \in N(h_k)}) \in G\})$$

for all measurable $G \subseteq H^{k+1}$. We maintain throughout the joint measurability assumption on H, σ , and the space of type profiles that sets of type profiles of the above form are measurable.⁴ We define the probability measure $\gamma_i^{k+2}(t, h_k)$ on histories of length $k + 2$ by

$$\begin{aligned} & \gamma_i^{k+2}(t, h_k)(G) \\ &= \beta_i(t, h_k)(\{\rho \mid (h_k, (\sigma_i(\rho(i), h_k))_{i \in N(h_k)}, (\sigma_i(\rho(i), h_{k+1}))_{i \in N(h_{k+1})}) \in G\}) \end{aligned}$$

for all measurable $G \subseteq H^{k+2}$, where we use the shorthand

$$h_{k+1} = (h_k, (\sigma_i(\rho(i), h_k))_{i \in N(h_k)}).$$

We define the probability measure $\gamma_i^{k+m}(t, h_k)$ on histories of length $k + m$, for any m , similarly. Again, we maintain the measurability assumptions required for this.

Given t and h_k , we can use these probability measures to define player i 's beliefs about the sequence of future outcomes in X . For measurable subsets X_1, X_2, \dots, X_n of X , define $\xi_i(t, h_k)$ as

$$\begin{aligned} & \xi_i(t, h_k)(X_1 \times \dots \times X_n \times X \times X \dots) \\ &= \gamma_i^{k+m}(t, h_k)(\{h_{k+m} \in H^{k+m} \mid f_j^{k+m}(h_{k+m}) \in X_j \text{ for all } j = 1, 2, \dots, n\}), \end{aligned}$$

where $l_{k+m} = n$. This defines $\xi_i(t, h_k)$ on the initial finite cylinder sets, and we can extend $\xi_i(t, h_k)$ to the sigma algebra on $X^{\mathbb{N}}$ generated by sets of that form. We can then define i 's *continuation value* at h_k , conditional on t , by

$$v_i(t, h_k) = (1 - \delta) \int \sum_{m=k+1}^{\infty} \delta^{m-k-1} u_t(x_m) \xi_i(t, h_k)(d\mathbf{x}), \quad (13)$$

which is well-defined as $U_{\rho(i)}$ is measurable with respect to the sigma algebra on $X^{\mathbb{N}}$. We say beliefs are *symmetric* at type profile ρ and history h_k if $\xi_i(\rho(i), h_k) = \xi_j(\rho(j), h_k)$ for all $i, j \in N(h_k)$. Note that, in this case, players of the same type have the same continuation value, so we may drop the subscript i from ξ_i and v_i .

The following result characterizes majority preferences over action profiles following an arbitrary history, when an arbitrary probability measure on individuals is used to

⁴This measurability holds, of course, if H is endowed with the trivial sigma algebra, $\{\emptyset, H\}$, but then the f^k 's must be history-independent.

“count votes.” Moreover, the characterization holds for an arbitrary type profile, which may not be common knowledge. The main assumption, other than quadratic utilities and existence of a majority core point, is that beliefs about future outcomes in X are symmetric. Using Lemma 1, we show that majority preferences over action profiles are identical to the “core” player’s.

Proposition 2 *Let ρ be an arbitrary type profile, let h_k be an arbitrary finite history, let ν be an arbitrary probability measure on $N(h_k)$, and assume beliefs are symmetric at ρ and h_k . Define $\pi = \nu \circ \rho^{-1}$, let $x \in X$ be a median in all directions, and suppose π is resolute at x . Then, for all $a, a' \in A(h_k)^{N(h_k)}$,*

$$\pi(\{t \mid v(t, (h_k, a)) \geq v(t, (h_k, a'))\}) \geq \frac{1}{2}$$

if and only if $v(x, (h_k, a)) \geq v(x, (h_k, a'))$.

By belief symmetry, the beliefs of each $i \in N(h_k)$ about the sequence of future outcomes in X following a are given by the same probability measure on $X^{\mathbb{N}}$. Let λ_m denote the marginal of this probability measure on the m th component. By additive separability, we may write

$$v(t, (h_k, a)) = (1 - \delta) \sum_{m=1}^{\infty} \int \delta^{m-1} u_t(x_m) \lambda_m(dx_m)$$

Letting

$$\lambda = (1 - \delta) \sum_{m=1}^{\infty} \delta^{m-1} \lambda_m,$$

we then have

$$v(t, (h_k, a)) = \int u_t(x) \lambda(dx).$$

That is, we can express a type t player’s continuation value as the expected utility from the lottery λ , which is itself independent of t . Similarly, we can express continuation values following a' as the expected utility from some lottery λ' . The result then follows directly from Lemma 2.1, completing the proof.

Given type profile ρ , we call h_k a *voting stage* if (i) beliefs are symmetric at ρ and h_k , (ii) $A(h_k)$ consists of two elements, say 0 and 1, and (iii) there exists a probability measure ν on $N(h_k)$ such that each player’s continuation value following $a \in \{0, 1\}^{N(h_k)}$ takes one value if the proportion of “one-votes” is above one half and takes another value if the proportion is less than or equal to one half. Formally, by the latter condition we mean that there exists ν such that, for each $t \in T$, there exist $q_t, r_t \in \mathbb{R}$ such that

$$v(t, (h_k, a)) = \begin{cases} q_t & \text{if } \nu(\{i \mid a_i = 0\}) \geq 1/2 \\ r_t & \text{if } \nu(\{i \mid a_i = 1\}) > 1/2, \end{cases}$$

where a_i is the action taken by i in the profile a . Note that this definition of a voting stage restricts not only the extensive game form but also the strategy profile σ , because individual continuation values depend on future stages only through the “winner” at h_k . Thus, our result, below, on voting stages would not apply to equilibria in which any voter is subject to punishment contingent specifically on his/her vote.

An issue that arises in the analysis of voting stages is the multiplicity of Nash equilibria: any profile of votes such that no voter is pivotal (i.e., no voter can unilaterally move the proportion of zero-votes above and below one half) forms a Nash equilibrium of the voting subgame. Indeed, in models with a continuum of voters, voters can never be pivotal in elections. In such cases, it makes sense to refine the possible equilibria by assuming each voter takes the action that offers the higher continuation value. Let $\mathbf{0}$ denote the action profile consisting of all zeroes for members of $N(h_k)$, and let $\mathbf{1}$ denote the action profile consisting of all ones. Formally, we might require that, given type profile ρ , each i votes for 1 ($a_i = 1$) if

$$v_i(\rho(i), (h_k, \mathbf{1})) > v_i(\rho(i), (h_k, \mathbf{0}))$$

and votes for 0 ($a_i = 0$) if the reverse inequality holds. In the context of a finite number of voters, this refinement amounts to the elimination of strategies dominated in the stage game. In the context of a continuum of voters, because no strategy is dominated in the voting stage, it amounts to a sincere voting requirement (but does not imply any voter myopia). If σ satisfies this condition, we say it is *stage-undominated* for type profile ρ at history h_k . This does not restrict i 's strategy when the above holds with equality. In applications with a finite number of voters, it is often assumed that the voter flips a coin in this case. Another common assumption is that one alternative is the “default” choice and receives i 's vote. If equality implies $\sigma_i(\rho(i), h_k) = 0$, then we say σ is *stage-undominated with default bias*.

Our last result characterizes the outcomes of voting stages for stage-undominated strategy profiles: essentially, voting produces a continuation value (q_t or r_t) if and only if it is higher for the core voter's type. Again, the characterization holds even if the type profile is not common knowledge.

Proposition 3 *Let ρ be an arbitrary type profile, let h_k be a voting stage with probability measure ν on $N(h_k)$, let σ be stage-undominated for ρ at h_k , and assume beliefs are symmetric at ρ and h_k . Define $\pi = \nu \circ \rho^{-1}$, let $x \in X$ be a median in all directions, and suppose π is resolute at x . Then*

- $\nu(\{i \mid a_i = 1\}) > 1/2$ if $v(x, (h_k, \mathbf{1})) > v(x, (h_k, \mathbf{0}))$
- $\nu(\{i \mid a_i = 0\}) > 1/2$ if $v(x, (h_k, \mathbf{0})) > v(x, (h_k, \mathbf{1}))$.

If σ is stage-undominated with default bias, then

- $\nu(\{i \mid a_i = 0\}) \geq 1/2$ if and only if $v(x, (h_k, \mathbf{0})) \geq v(x, (h_k, \mathbf{1}))$.

If, for example, $v(x, (h_k, \mathbf{1})) > v(x, (h_k, \mathbf{0}))$, then Proposition 3.1 implies

$$\pi(\{t \mid v(t, (h_k, \mathbf{1})) > v(t, (h_k, \mathbf{0}))\}) > \frac{1}{2},$$

or equivalently,

$$\nu(\{i \mid v(\rho(i), (h_k, \mathbf{1})) > v(\rho(i), (h_k, \mathbf{0}))\}) > \frac{1}{2}.$$

Because σ is stage-undominated, this implies that $\nu(\{i \mid a_i = 1\}) > 1/2$, as required. The rest of the proposition is proved similarly, completing proof.

4 Applications

We end with three applications of the above results.

4.1 Campaigning

Suppose a finite, odd number of voters have quadratic von Neumann-Morgenstern utilities over a one-dimensional policy space $X \subseteq \mathbb{R}$ and must vote between an incumbent, with known policy position $x \in X$ and an unknown challenger. The challenger's position, y , is distributed according to the common prior distribution function F with density f . Conditional on y , the voters observe a common signal z , distributed according to $G(\cdot|y)$ with density $g(\cdot|y)$. After updating using Bayes rule, the voters' common beliefs about the challenger's position are given by the density

$$b(y) = \frac{g(z|y)f(y)}{\int g(z|s)f(s) ds}.$$

Assume the voters eliminate weakly dominated strategies, so that a voter with ideal point \tilde{x} votes for the incumbent if

$$u_{\tilde{x}}(x) > \int u_{\tilde{x}}(y)b(y) dy, \tag{14}$$

and votes for the challenger if this inequality is reversed. The winner is the candidate with the most votes.

This model can be mapped into the framework of the previous section by including the two candidates as players, the challenger with type y distributed according to f , and adding a dummy player n , "nature," whose type w is distributed according to the continuous distribution H and whose action determines the signal observed by voters: n takes an action $z = \sigma_n(w, y)$, where σ_n is continuous and strictly increasing in w . We require that H and σ_n satisfy

$$G(z|y) = H(\sigma_n(z|y)^{-1})$$

for all $y, z \in \mathbb{R}$, where $\sigma_n(z|y)^{-1}$ is the unique w satisfying $\sigma_n(w, y) = z$. After the choice of z , we give the voters beliefs over y with the above density b , consistent with Bayesian updating when voters observe z but not y or w . Obviously, this is a voting stage with the counting measure on the set of voters. Because X is one-dimensional and the number of voters is odd, there is a unique median in all directions, x^* , and the distribution of ideal points is resolute at x^* . In fact, x^* is simply the median of the distribution of ideal points.

A direct implication of Proposition 3.2 is that *the incumbent will win if the strict inequality in (14) holds for the median, and the challenger will win if the opposite holds*. An implication of part (ii) of Lemma 2.1 is that, when equality obtains in (14), the number of voters who strictly prefer the incumbent will equal the number of voters who strictly prefer the challenger. If indifferent voters flip coins in the original model, then each candidate will win with probability one half, conditional on the median being indifferent. If indifferent voters vote for the incumbent, as is sometimes assumed, he/she will win with probability one, again conditional on median indifference.

In the sketch of the above model, we did not specify the voters' information about each other's ideal points. In fact, because Proposition 3.2 holds for an arbitrary type profile, even if there is incomplete information, the results described above hold for every realization of voter ideal points. The difference made by incomplete information is that the identity of the median voter may be unknown (even to the median).

4.2 Bargaining

Suppose an odd number of agents, $1, 2, \dots, n$, have quadratic von Neumann-Morgenstern utilities over a one-dimensional space $X \subseteq \mathbb{R}$ of alternatives, and suppose the agents decide on an alternative in each of an infinite number of periods according to the following protocol: in any period t , if no alternative has previously been agreed upon, then an agent is drawn randomly to make a proposal, x , which is followed by a vote; if the proposal passes, then the outcome in period t and all subsequent periods is x , and each agent i receives utility $u_i(x)$ in periods $t, t+1, t+2, \dots$; if the proposal is rejected, then the outcome in period t is a status quo policy $q \in X$, each agent i receives utility $u_i(q)$ in period t , and bargaining moves to period $t+1$, where this process is repeated. If outcome x is passed in period t , then each agent i 's discounted utility is

$$(1 - \delta^{t-1})u_i(q) + \delta^{t-1}u_i(x),$$

where $\delta \in [0, 1)$ is a common discount factor. (Here, we normalize the discounted sum by $(1 - \delta)$.) Let p_i denote the history-independent probability that i is drawn to make a proposal. Complete information is assumed.

A class of strategy profiles of much interest in this model are the *stationary* strategy profiles, in which the proposal of an agent i is history-independent and the vote of an agent depends only on the alternative proposed. Letting A_i denote the "acceptance set"

for i , the proposals that agent i would vote for, and letting \mathcal{M} denote the collection of coalitions C containing a majority of agents, we define

$$A = \bigcup_{C \in \mathcal{M}} \bigcap_{i \in C} A_i$$

as the “social acceptance set.” These are the proposals that would be passed if made. A *no-delay* strategy profile is one for which $x_i \in A$ for all agents. Given a no-delay profile, it is easy to see that each agent i ’s continuation value, v_i , is history-independent and, moreover,

$$v_i = \sum_{j \in N} p_j u_i(x_j).$$

A *no-delay stationary equilibrium* is a subgame perfect equilibrium in which (i) each agent i proposes $x_i \in \arg \max\{u_i(x) \mid x \in A\}$, and (ii) each agent i votes for any proposal that is weakly preferred to continuing the bargaining process, i.e.,

$$A_i = \{x \in X \mid u_i(x) \geq (1 - \delta)u_i(q) + \delta v_i\}.$$

Note that this equilibrium condition on voting strategies incorporates the refinement of stage-undominated strategies in voting stages with default bias in favor of proposals. Existence of no-delay stationary equilibria follows from a general result, allowing for multiple dimensions and arbitrary concave utility functions, in Banks and Duggan (2003).

This model can be mapped into the framework of the previous section by adding a countable set of dummy players, indexed d_1, d_2, \dots , each with type set $\{1, 2, \dots, n\}$, to incorporate the random selection of proposers. The types of the dummy players are independently and identically distributed according to the distribution $p = (p_1, p_2, \dots, p_n)$. We give the agents singleton sets of types, reflecting the assumption of complete information in the original model. In each period t , the dummy player d_t designates an agent to make a proposal; that agent then proposes; and then all agents vote. We suppose that all agents use stationary strategies as described above, while type i of any dummy player simply designates player i to propose. The agents’ beliefs about the dummy players’ types are simply iid with marginal p throughout the game. This transformation of the original model preserves the continuation values of the agents and leaves the strategy profiles satisfying (i) and (ii) unaffected. Obviously, a stage following the proposal of any agent is a voting stage with the counting measure on the set of agents, and any no-delay stationary equilibrium of the original model corresponds to a strategy profile that is stage-undominated with default bias (in favor of the proposal) after every history. Because X is one-dimensional and the number of voters is odd, there is a unique median in all directions, x^* , and the distribution of ideal points is resolute at x^* .

Thus, Proposition 3.2 implies that, given a no-delay stationary equilibrium and given any history, *a proposal will pass if and only if the median voter weakly prefers passage to continuing the bargaining process*. That is, the social acceptance set is exactly equal to the median voter’s acceptance set in every no-delay stationary equilibrium. This provides a

game-theoretic foundation for a partial median voter result on bargaining in committees, one that differs from Black (1958) in that the median voter dictates the outcomes of each vote but, because other agents may make proposals, his/her ideal point may not be the unique outcome. This consideration arises, of course, because we explicitly model the proposal process and the incentives of agents in their roles as proposer and voter. Elsewhere (Banks and Duggan (2003)), we have shown that the set of outcomes passed in no-delay stationary equilibria converges to the median as agents become arbitrarily patient, even if utilities are not quadratic.

Whereas the above model is investigated in Banks and Duggan (2003), a related model is taken up in Banks and Duggan (2000). A special case is the model in which utility functions are quadratic plus a common constant c such that $u_i(x) + c \geq 0$ for all $x \in X$, and in which each agent's status quo payoff is zero (so the status quo is "bad" for all agents, as in Baron and Ferejohn (1989)). It can be checked that, in that model too, the median voter is decisive in the above sense.

4.3 Repeated Elections

Suppose a continuum of voters, $N = [0, 1]$, must choose between an incumbent and a challenger for political office in each of an infinite number of periods. The officeholder in any period t chooses a policy $y_t \in X = [0, 1]$, and a voter with ideal point x receives quadratic utility $u_x(y_t)$ in that period. In the subsequent period, a challenger is drawn uniformly from the population of voters to run against the incumbent in a majority-rule election. Voters vote simultaneously by secret ballot, and the incumbent wins if and only if he/she receives at least half of the votes; otherwise, the challenger takes office. In either case, the winner chooses the policy y_{t+1} for period $t + 1$. This sequence of challenger draws, elections, and policy choices determines an infinite sequence y_1, y_2, \dots of policies, yielding utility

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_x(y_t),$$

to a voter with ideal point x , where $\delta \in [0, 1)$ is a common discount factor.

Voter ideal points are private information, with the marginal distribution of each voter's ideal point given by the continuous, strictly increasing distribution function F . Because there is a continuum of voters, we do not assume voter ideal points are independently distributed, but we assume *finite-dimensional independence*, i.e., conditional on any finite number of other voters' ideal points, the ideal point of each voter i is distributed according to F . Moreover, we assume that a law of large numbers holds: the distribution of type profiles is such that, for almost every type profile, the set of voters with ideal points less than or equal to x is measurable, and the fraction of such voters is $F(x)$, i.e.,

$$\lambda(\{i \mid \rho(i) \leq x\}) = F(x),$$

where λ here is the Lebesgue measure.⁵ We refer to the profiles for which these conditions hold as *regular*. Thus, F represents not only the marginal distribution of any voter's ideal point, but it also represents (with probability one) the distribution of ideal points across the electorate and, therefore, the distribution of challenger ideal points. Note that, because F is continuous, the probability that any one voter is ever selected is zero. We assume that a voter's ideal point is never directly observed, even if drawn as a challenger and elected as officeholder, though an officeholder may reveal information about his/her ideal point through policy choices while in office. By finite-independence, updating about the officeholder's ideal point does not affect a voter's beliefs about other voters' ideal points.

A class of strategy profiles of interest in this model are the *simple* strategy profiles, in which (i) a voter with ideal point x always chooses the same point p_x if ever elected to office, where p_x is measurable with respect to x , and (ii) a voter with ideal point x votes to reelect the incumbent if and only if the incumbent's policy choice y in the previous period satisfies a history-independent utility standard \bar{u}_x , which is measurable as a function of x . Thus, a voter with ideal point x votes for the incumbent if and only if $u_x(y) \geq \bar{u}_x$. Letting

$$A_x = \{y \in [0, 1] \mid u_x(y) \geq \bar{u}_x\}$$

denote the acceptance set for a voter with ideal point x , letting $\Pr(C) = \int_C dF$ denote the proportion of voters with ideal points in a measurable set $C \subseteq X$, and letting \mathcal{M} denote the collection of all measurable sets C satisfying $\Pr(C) \geq 1/2$, we define

$$A = \bigcup_{C \in \mathcal{M}} \bigcap_{x \in C} A_x$$

as the social acceptance set, the set of policies that, if chosen, would lead to reelection of the incumbent. Because \bar{u}_x is measurable, A is measurable for every regular type profile. To see this, note that, if ρ is regular, then $x \in A$ if and only if

$$\int \phi(x, z) F(dz) \geq \frac{1}{2},$$

where

$$\phi(x, z) = \begin{cases} 1 & \text{if } x \in A_z \\ 0 & \text{else,} \end{cases}$$

a jointly measurable function. Then measurability of $\int \phi(\cdot, z) dF$ follows from Fubini's theorem (cf. Aliprantis and Border (1999), Theorem 11.26). Thus, A is measurable.

Given a simple strategy profile, we define beliefs throughout the game as follows. In period t , following elections of officeholders i_1, i_2, \dots, i_k , let y_1, y_2, \dots, y_k denote the last policy choices of these politicians. For each politician m , let $T_m = p^{-1}(y_m)$ if this

⁵Theorem 2 in Judd (1985) establishes the existence of such a distribution over type profiles.

set is non-empty, and otherwise let $T_m = \{0\}$.⁶ Then the updated beliefs of a type x voter are given by the prior over ρ conditioned on $\rho(i_m) \in T_m$, $m = 1, 2, \dots, k$ and $\rho(i) = x$. By finite independence, the conditional beliefs about the types of all other voters are still given by F . Let $W = \{x \in X \mid p_x \in A\}$ denote the set of “winning” ideal points, and let $L = X \setminus W$ denote the set of “losing” ideal points, both measurable under our assumptions. Assuming it exists, it is then straightforward, using the law of large numbers, to solve for the continuation value of electing a challenger for a voter with ideal point x as

$$v_x = \frac{\int_W u_x(p_z) F(dz) + (1 - \delta) \int_L u_x(z) F(dz)}{1 - \Pr(L)\delta},$$

where the first term in the numerator accounts for challengers who choose policies in A and the second, appropriately discounted, for those who do not.

A *simple equilibrium* is a simple strategy profile in which (iii) an officeholder with ideal point x chooses $y \in \arg \max\{u_x(z) \mid z \in A\}$ if

$$\max\{u_x(z) \mid z \in A\} \geq \delta v_x$$

and otherwise chooses $y = x$, and (iv) a voter with ideal point x votes to reelect an incumbent who chose $y \in A \cup L$ in the previous period (and so followed the path of play) if and only if

$$u_x(y) \geq v_x.$$

That is, an officeholder chooses between remaining in office by satisfying a majority of voters (choosing the best point in A) and leaving office (choosing his/her ideal point). Voters compare the continuation value of retaining an incumbent and the continuation value of an unknown challenger, and vote for the incumbent if he/she offers at least as high a continuation value as the challenger. Note that this equilibrium condition on voting strategies incorporates the refinement of stage-undominated strategies in voting stages with default bias in favor of the incumbent. Existence of simple equilibria is established as a special case of a general result in Banks and Duggan (2002) that allows for multiple policy dimensions and arbitrary concave utility functions, but we do assume a finite number of possible types there. In that case, we must allow officeholders with ideal points $x \notin A$ to mix over the best policies in the acceptance set (thereby winning) and their ideal point (thereby losing) when indifferent. This is not an issue when F is continuous and strictly increasing, as we assume here.

To map this model into the framework of the previous section, we use a countable set, $M \subseteq [-1, 0)$, of dummy players, as in Application 4.2, to incorporate the random selection of challengers. We index these players d_1, d_2, \dots . We will assume that the types of the dummy players are independently and identically distributed according to

⁶Without loss of generality, we will assume that the median of F is greater than one half. Thus, in equilibrium, officeholders who choose policies off the equilibrium path will not be reelected.

the distribution F . In each period t , the dummy player d_t designates a voter to run for office as the challenger; all voters other than the candidates (for technical reasons, below) then cast their votes; and the winner then chooses a policy for that period. We assume that all voters use simple strategies as described above, while a type x dummy player simply chooses a voter of type x . Beliefs of the voters are as described above. This transformation of the original model preserves the continuation values of the voters, if the expression in (13) is well-defined, and leaves the strategy profiles satisfying (i) and (ii) unaffected.

To construct continuation values for the voters, we extend the definition of a type profile ρ so that it maps every $i \in M \cup N$ to a type in X , and we extend the definition of regular profile to this space. We endow the space of type profiles with the sigma algebra generated by sets of the form

$$\{\rho \mid \rho(i) \in C\},$$

where $i \in M \cup N$ and C is a Lebesgue measurable subset of X . We then extend this sigma algebra so that the space of regular profiles is measurable and assume the regular profiles have probability one (cf. Judd (1985), Theorems 1 and 2). Letting 0 denote a vote for the incumbent and 1 a vote for the challenger, and letting R denote the set of measurable functions $r: N \rightarrow \{0, 1\}$, representing election returns in any given period, the infinite histories in this electoral game lie in the set

$$X \times N \times R \times X \times N \times R \times X \times \dots,$$

where a sequence $(x_1, i_2, r_2, x_2, i_3, r_3, x_3, \dots)$ represents the period one policy (chosen by an arbitrary initial officeholder, i_1), the period two challenger, the period two election return, the period two policy, and so on. We endow all finite histories with the sigma algebra generated by rectangles of the form

$$X_1 \times N_2 \times R_2 \times X_2 \times \dots,$$

where N_1, N_2, \dots are finite or cofinite subsets of $[0, 1]$, and where X_1, X_2, \dots are Lebesgue measurable subsets of X . As for election returns, we let R_I denote the set of returns r such that $\int r d\lambda \geq 1/2$; and we let R_C denote the set of returns r such that $\int r d\lambda < 1/2$. We then impose that R_2, R_3, \dots lie in $\{\emptyset, R_I, R_C, R\}$, indicating either that neither candidate won, that the incumbent won, that the challenger won, or that one of the candidates won.

Given a strategy profile as specified above, the set of type profiles for which the period one policy lies in X_1 is just

$$\mathcal{P}_1 = \{\rho \mid \rho(i_1) \in p^{-1}(X_1)\},$$

which is measurable by measurability of policy choice strategies and by choice of sigma algebra on the space of type profiles. The set of type profiles for which $(x_1, i_2) \in X_1 \times N_2$ is

$$\mathcal{P}_2 = \mathcal{P}_1 \cap \{\rho \mid \rho(d_1) \in N_2\},$$

which is also measurable. The set of (regular) type profiles for which $(x_1, i_2, r_2) \in X_1 \times N_2 \times R_I$ is

$$\mathcal{P}_3 = \mathcal{P}_2 \bigcap_{i \in N_2} \{\rho \mid \rho(i) \in W\},$$

which is measurable because W is measurable, and similarly for other collections of election returns.⁷ Measurability of profiles corresponding to histories of greater length can be established by an induction argument based on these observations. This allows us to define the probability measures $\gamma_i^{k+m}(t, h_k)$ on histories using σ in the obvious way. Furthermore, the outcome functions for finite histories are simply defined as projections onto the policy choices of officeholders, e.g.,

$$f^{k3k-2}(x_1, i_2, r_2, x_2, \dots, i_k, r_k, x_k) = (x_1, x_2, \dots, x_k),$$

which are measurable with respect to the above sigma algebras on the H^k 's. This allows us to define the probability measures $\xi_i(t, h_k)$ on outcome sequences required for (13), as required.

Note that beliefs over outcome sequences are symmetric at every vote: though the non-candidates have private information about their own types, and therefore how they would vote in the future, the probability that any one is selected to run for office is zero. The only voters with private information about future policy outcomes are the candidates, who do not cast a vote. (Indeed, along the equilibrium path, the incumbent's future policy choices are known, so he/she does not have private information about policy outcomes.) Thus, following the selection of a challenger, we have a voting stage with the Lebesgue measure on the set of voters, and any simple equilibrium of the original model corresponds to a strategy profile that is stage-undominated with default bias (for the incumbent) at every vote. Because X is one-dimensional, there is a unique median in all directions, and the distribution of ideal points is resolute at that point.

Thus, Proposition 3.2 implies that, given a simple equilibrium, given any history, and given any type profile, *an incumbent will be reelected if and only if the median voter weakly prefers the incumbent to a randomly chosen challenger*. That is, the social acceptance set is exactly equal to the median voter's acceptance set in every simple equilibrium. This partially extends the median voter result of Downs (1957) from a static model of party competition, in which two parties adopt policy platforms prior to election, to a very different model of infinitely repeated elections with incomplete information, in which political actors are unable to compete for votes by position-taking. Here, of course, the median voter does not dictate policy outcomes, but rather electoral outcomes. From a technical point of view, this property has proven to drastically simplify the computation of equilibria and the theoretical analysis of this model.

⁷Non-regular type profiles may produce reelection when $x \notin W$, or may give victory to a challenger when $x \in W$, because the distribution of types in the electorate may disagree with F . The voters' updated posteriors place outer probability zero on these profiles throughout the game, so we may ignore them.

References

- [1] C. Aliprantis and K. Border (1999) *Infinite Dimensional Analysis: A Hitchhiker's Guide* New York: Springer-Verlag.
- [2] J. Banks and J. Duggan (2000) "A Bargaining Model of Collective Choice," *American Political Science Review* 94:73-88.
- [3] J. Banks and J. Duggan (2002) "A Multidimensional Model of Repeated Elections," mimeo., University of Rochester.
- [4] J. Banks and J. Duggan (2003) "A Bargaining Model of Legislative Policy-making," mimeo., University of Rochester.
- [5] J. Banks, J. Duggan, and M. Le Breton (2003) "Social Choice in the General Spatial Model of Politics," mimeo., University of Rochester.
- [6] D. Baron and J. Ferejohn (1989) "Bargaining in Legislatures." *American Political Science Review* 83:1181-1206.
- [7] D. Bernhardt, E. Hughson, and S. Dubey (2002) "Term Limits and Pork Barrel Politics," *Journal of Public Economics*, forthcoming.
- [8] D. Black (1958) *The Theory of Committees and Elections* Cambridge: Cambridge University Press.
- [9] R. Calvert (1985) "Robustness of the Multidimensional Voting Model: Candidate Motivations, Uncertainty, and Convergence," *American Journal of Political Science* 29:69-95.
- [10] L. Campuzano (2001) "Parties in a Repeated Elections Model," mimeo., University of Rochester.
- [11] S. Cho and J. Duggan (2002) "Uniqueness of Stationary Equilibria in a One-dimensional Model of Bargaining," *Journal of Economic Theory*, forthcoming.
- [12] O. Davis, M. DeGroot, and M. Hinich (1972) "Social Preference Orderings and Majority Rule," *Econometrica* 40:147-157.
- [13] A. Downs (1957) *An Economic Theory of Democracy* New York: Harper and Row.
- [14] J. Duggan (2000) "Repeated Elections with Asymmetric Information," *Economics and Politics* 12:109-135.
- [15] J. Duggan and M. Fey (2001) "Electoral Competition with Policy-motivated Candidates," W. Allen Wallis Institute of Political Economy Working Paper No.19, University of Rochester.
- [16] J. Duggan and M. Fey (2002) "A Downsian Model of Repeated Elections," mimeo., University of Rochester.

- [17] K. Judd (1985) “The Law of Large Numbers with a Continuum of IID Random Variables,” *Journal of Economic Theory* 35:19-25.
- [18] J. Kingman and S. Taylor (1966) *Introduction to Measure and Probability* Cambridge: Cambridge University Press.
- [19] D. McGarvey (1953) “A Theorem on the Construction of Voting Paradoxes,” *Econometrica* 21:608-610.
- [20] R. McKelvey (1976) “Intransitivities in Multidimensional Voting Models, and Some Implications for Agenda Control,” *Journal of Economic Theory* 2:472-482.
- [21] R. McKelvey (1979) “General Conditions for Global Intransitivities in Formal Voting Models,” *Econometrica* 47:1086-1112.
- [22] R. McKelvey (1986) “Covering, Dominance, and Institution-Free Properties of Social Choice,” *American Journal of Political Science* 30:283-314.
- [23] R. McKelvey, P. Ordeshook, and P. Ungar (1980) “Conditions for Voting Equilibria in Continuous Voter Distributions,” *SIAM Journal of Applied Mathematics* 39:161-168.
- [24] A. Meirowitz (2002) “Accountability and Representation in Repeated Elections with Uncertain Policy Constraints,” mimeo., Princeton University.
- [25] C. Plott (1967) “A Notion of Equilibrium and its Possibility under Majority Rule,” *American Economic Review* 57:787-806.