

Jeffrey S. Banks · John Duggan

A social choice lemma on voting over lotteries with applications to a class of dynamic games

Received: 22 April 2003 / Accepted: 22 March 2004 / Published online: 4 April 2006
© Springer-Verlag 2006

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.

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, however, more structure can 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.

This paper was completed after Jeff Banks’s death. John Duggan is deeply indebted to him for his friendship and his collaboration on this and many other projects.

J. S. Banks
Division of Humanities and Social Sciences, California Institute of Technology,
Pasadena, CA 91125, USA

J. Duggan (✉)
Department of Political Science and Department of Economics, University of Rochester,
Rochester, NY 14627, USA
E-mail: dugg@troi.cc.rochester.edu

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 et al. (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; Duggan and Fey 2005).

In this paper, we reinforce Davis et al.’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 von Neumann–Morgenstern utility functions 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 simplest version of this result is proved straightforwardly using mean-variance analysis, a property of quadratic utilities. Because majority preferences over pure alternatives are unaffected by arbitrary monotonic transformations of quadratic utilities, Davis et al.’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 the 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 (2005). Assuming quadratic utilities, our results imply that the approval of the median voter is necessary and sufficient for a proposal to pass in equilibrium, an observation used

¹ They 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 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 et al. (2004), and Banks and Duggan (2006). 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 the reformulation of 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 in other models of repeated elections (Campuzano 2001; Duggan and Fey 2006; Meirowitz 2003), and it offers to facilitate the investigation of topics such as uniqueness of stationary equilibria in the one-dimensional bargaining model (Cho and Duggan 2003) 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 weak preference 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 $\|x - \tilde{x}_i\| \leq \|y - \tilde{x}_i\|$. We write $\succsim_{\tilde{x}}$ for the preference relation corresponding to ideal point \tilde{x} , and we extend this notation even to points in Euclidean space that do not correspond to ideal points, i.e., $x \succsim_z y$ means $\|x - z\| \leq \|y - z\|$. We use \succ_z and \sim_z to denote the associated strict preference and indifference relations. Let us 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

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

and we employ the above convention to denote strict preference and indifference. 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 $u_{\tilde{x}}(x) = -\|x - \tilde{x}\|^2$, and, more generally, let u_z be the quadratic function defined with respect to the arbitrary point z . Paralleling the above convention, let \succsim_z^* denote the preferences over lotteries corresponding to z . Thus, for all $\lambda, \mu \in \Lambda$, $\lambda \succsim_z^* \mu$ if and only if

$$\int u_z(x) \lambda(dx) \geq \int u_z(x) \mu(dx).$$

We say λ is *weakly majority-preferred* to μ , written $\lambda \succ_{\tilde{x}}^* \mu$, if

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

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. Conversely, if x is a majority core point interior to X , then it is a median in all directions (cf. Davis et al. 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. Interiority can be dropped, if X is convex and contains all ideal points. To see this for the case in which π has finite support,² suppose x is a majority core point but, for some direction t , $\pi(H_t^+(x)) > 1/2$; let Y denote the convex hull of the ideal points $\tilde{x}_i \in H_t^+(x)$, so that $Y \subseteq X$; and note that the unique solution to $\min_{y \in Y} \|x - y\|$ is strictly majority-preferred to x , a contradiction. That something like interiority or convexity is needed can be seen from the following example: let $X = \{(0, 0), (0, 1), (1, 0)\}$, and suppose there are just two individuals, with ideal points $\tilde{x}_1 = (0, 1)$ and $\tilde{x}_2 = (1, 0)$. Then $(0, 0)$ is a majority core point, but it is not a median in all directions.

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

² A slightly more complicated argument is required for the general case.

than half of the individuals on one side.³ This condition clearly holds at every $x \in \mathbb{R}^d$ if the set of individuals is finite and odd in number. Resoluteness also holds if π is non-atomic with strictly positive density on some open set containing x . In that case, for every open set G around x and for every direction t , we have $\pi(G \cap H_t^+(x)) > 0$; then, 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.

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.

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 it applies equally to majority core points when the two notions coincide, as they do under the weak conditions mentioned above. Then part (i) of the lemma shows that if one lottery is weakly preferred over another by the core voter, then the first lottery is weakly majority-preferred over the second. Moreover, any two lotteries that are equally desirable to the core voter are majority-indifferent. This result, which generalizes Davis et al.'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 there is a finite, odd number of individuals with at the most one ideal point at the core, then any two lotteries that are equally desirable to the core voter must have equal proportions of voters with opposing strict preferences over them. The same is true in a continuum model if the set of individuals with ideal points at the core has measure zero. This part is closely related to Theorem 2 of McKelvey et al. (1980), which generalizes Plott's (1967) necessary symmetry conditions at a core point.⁴ Finally, part (iii), which corresponds to Davis et al.'s (1972) Theorem 4, 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 that of the core voter.

Lemma 2.1 *Let x be a median in all directions. Then (i) $\lambda \succ_x^* \lambda'$ implies $\lambda \succ^* \lambda'$, (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 $\lambda \succ_x^ \lambda'$ if and only if $\lambda \succ^* \lambda'$.*

³ Assuming 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 et al. (1972).

⁴ In contrast to McKelvey et al. (1980), we allow the distribution of ideal points to be discrete, in order to capture a finite number of individuals.

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

$$u_x(m) - u_x(m') \geq v - v'. \tag{2.1}$$

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

$$u_{\tilde{x}}(m) - u_{\tilde{x}}(m') \geq u_x(m) - u_x(m'). \tag{2.2}$$

Then, by (2.1) and (2.2), we have

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

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

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

$$\epsilon = \pi(H_{m-m'}^+(x)) - \pi(H_{m'-m}^+(x)) - \pi(\{x\}), \tag{2.3}$$

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

$$\pi(H_{m-m'}(x) \cap H_t(x) \setminus \{x\}) < \frac{\epsilon}{2}. \tag{2.4}$$

(The existence of such a t follows from an argument similar to that used in the proof of Theorem 2 in McKelvey et al. (1986, pp. 164–165).) Without loss of generality, assume that

$$\pi(H_{m-m'}(x) \cap H_t^+(x)) \geq \pi(H_{m-m'}(x) \cap H_t^-(x)). \tag{2.5}$$

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

$$\pi(H_{m-m'}^+(x) \setminus H_{y(n)-m'}^+(x)) < \frac{\epsilon}{4}. \tag{2.6}$$

It is an identity that

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

Solving for $\pi(H_{m-m'}^+(x))$ in (2.3) and substituting, and using (2.5), we have

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

Using (2.4) the latter inequality yields

$$\pi(H_{m'-m}^+(x)) + \pi(\{x\}) + \pi(H_{m-m'}(x) \cap H_t^+(x)) + \frac{3\epsilon}{4} > \frac{1}{2}. \quad (2.7)$$

From (2.6) and (2.3), respectively, we have

$$\pi(H_{y(n)-m'}^+(x) \cap H_{m-m'}^+(x)) > \pi(H_{m-m'}^+(x)) - \frac{\epsilon}{4} \\ = \pi(H_{m'-m}^+(x)) + \pi(\{x\}) + \frac{3\epsilon}{4}. \quad (2.8)$$

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

$$\pi(H_{y(n)-m'}^+(x) \cap H_{m-m'}(x)) \geq \pi(H_{m-m'}(x) \cap H_t^+(x)). \quad (2.9)$$

Finally, combining (2.7), (2.8), and (2.9), 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 $\lambda >^*_x \lambda'$, i.e.,

$$u_x(m) - u_x(m') > v - v'. \quad (2.10)$$

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

$$u_{\tilde{x}}(m) - u_{\tilde{x}}(m') \geq u_{x(\epsilon)}(m) - u_{x(\epsilon)}(m'). \quad (2.11)$$

Pick ϵ small enough, using continuity and (2.10), so that

$$u_{x(\epsilon)}(m) - u_{x(\epsilon)}(m') > v - v'. \quad (2.12)$$

Then, by (2.11) and (2.12), 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.

The assumption of quadratic utility in Lemma 2.1 cannot be weakened, even to the assumption that utility functions are Euclidean and identical up to translations, even in one-dimensional settings. To see this, suppose there are three individuals, suppose $d = 1$ with $X = [-2, 2]$, and let ideal points be $\tilde{x}_1 = -1$, $\tilde{x}_2 = 0$, and $\tilde{x}_3 = 1$. Let $u_i(x) = v(|x - y|)$, where $v: [0, 4] \rightarrow \mathbb{R}$ is a strictly concave, decreasing function. Let $\epsilon > 0$ be small enough that

$$\frac{1}{3}[v(2) + v(0) + v(2)] > v(2 - \epsilon). \tag{2.13}$$

We leave $v(3 - \epsilon) - v(3) = \Delta$ as a variable, which, by letting v decrease rapidly on $(3 - \epsilon, 3]$, can be chosen to be arbitrarily large without affecting the value of v on $[0, 3 - \epsilon]$. Now let λ be the lottery with probability $1/3$ each on $-2, 0$, and 2 , and let λ' be the lottery with probability $1/2$ each on $-2 + \epsilon$ and $2 - \epsilon$. By (2.13), the core voter, 2, strictly prefers λ to λ' . Individuals 1 and 3 have the opposite strict preference if

$$\frac{1}{2}[v(1 - \epsilon) + v(3 - \epsilon)] > \frac{1}{3}[v(1) + v(1) + v(3)],$$

which holds for sufficiently large Δ . Thus, Lemma 2.1 does not extend appreciably beyond quadratic utilities.

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 also 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 et al. (2006) 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 with the following structure. The extensive form is a stage game, where a state ω is initially drawn from a probability space Ω . We write $h_0 = \omega$

for this initial history. Given the initial history h_0 , a measurable subset $N(h_0)$ of players simultaneously choose actions in the measurable space $A(h_0)$, while all other players choose a null action \bar{a} . Denoting the profile of actions by active players as $a \in A(h_0)^{N(h_0)}$, which is assumed measurable, 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 , and let $H = \bigcup_{k=0}^{\infty} H^k$ denote the set of all finite histories of the game. 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.

We assume every finite history h_k generates a sequence $f^k(h_k)$ of length l_k (independent of h_k) in a bounded set $X \subseteq \mathbb{R}^d$ of *outcomes*, e.g.,

$$f^k(h_k) = (f_1^k(h_k), f_2^k(h_k), \dots, f_{l_k}^k(h_k)) = (x_1, x_2, \dots, x_{l_k}),$$

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 conditions that l_k is weakly increasing in k and 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 $n = 1, 2, \dots, l_k$. Thus, we can drop the superscript on f . Without loss of generality, we assume that $l_k \rightarrow \infty$. Thus, every infinite history corresponds to an infinite sequence, denoted by $\mathbf{x} = (x_1, x_2, \dots)$, of outcomes in X . We endow the space $X^{\mathbb{N}}$ of infinite sequences of outcomes with the sigma algebra generated by finite initial cylinder sets, i.e., sets of the form $X_1 \times \dots \times X_n \times X \times X \dots$, where X_1, \dots, X_n are measurable subsets of X .

Each state ω determines a *type*, specifically $\rho_i(\omega) \in \mathbb{R}^d$, for each individual i , where we assume $\rho : N \times \Omega \rightarrow \mathbb{R}^d$ is measurable. Here, we interpret a type, say \tilde{x} , as the ideal point of an agent. Given an infinite history and an associated sequence \mathbf{x} of outcomes, a type \tilde{x} player's payoff in the game is

$$U_{\tilde{x}}(\mathbf{x}) = (1 - \delta) \sum_{l=1}^{\infty} \delta^{l-1} u_{\tilde{x}}(x_l),$$

where $u_{\tilde{x}}$ is the quadratic utility function with ideal point \tilde{x} and $\delta \in [0, 1)$ is a common discount factor. Since $\delta < 1$, it follows that $U_{\tilde{x}}$ is measurable as a function of \mathbf{x} .

A *strategy profile* is a measurable mapping $\sigma : N \times H \rightarrow A$ such that $\sigma_i(h_k) \in A(h_k)$ for all $i \in N(h_k)$ and $\sigma_i(h_k) = \bar{a}$ for all $i \notin N(h_k)$.⁵ Here, of course, $\sigma_i(h_k)$ is the action taken by player i after history h_k . In applications, σ must also respect the players' 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 $\sigma_i(h_k, a) = \sigma_i(h_k, a')$. In the remainder of this section, we take a strategy profile σ as given.

⁵ For simplicity, we omit consideration of mixed strategies.

Given this strategy profile, prior beliefs on Ω determine probability measures $\xi_i(h_k)$ on $X^{\mathbb{N}}$ representing player i 's beliefs about the sequence of outcomes, conditional on realizing history h_k . In applications, these beliefs must be consistent with Bayes rule and the players' information sets. We can then define i 's *continuation value*, conditional on history h_k in which i 's type is \tilde{x} , by

$$v_i(h_k) = (1 - \delta) \int \sum_{l=l_k+1}^{\infty} \delta^{l-l_k-1} u_{\tilde{x}}(x_l) \xi_i(h_k)(d\mathbf{x}), \tag{3.1}$$

which is well-defined given our sigma algebra on $X^{\mathbb{N}}$. We assume that v is jointly measurable on $N \times H$. We say beliefs are *symmetric* at history h_k if $\xi_i(h_k) = \xi_j(h_k)$ for all $i, j \in N(h_k)$, in which case we can simply write $v_{\tilde{x}}(h_k)$ for the continuation value of an active player of type \tilde{x} . In the remainder of this section, we take beliefs as given.

The next result characterizes majority preferences over action profiles following an arbitrary history, when an arbitrary probability measure on individuals is used to "count votes." Note that the payoff to active player i from action profile a , given history h_k with initial state ω , is

$$\phi_i(h_k, a) = (1 - \delta) u_{\rho_i(\omega)}(f_{l_{k+1}}(h_k, a)) + \delta v_i(h_k, a),$$

where the first term on the righthand side denotes the payoff from the outcome corresponding to the current choice of a , and the second term is the continuation value. If beliefs are symmetric at (h_k, a) , we may write $\phi_{\tilde{x}}(h_k, a)$ for the payoff of an active player of type \tilde{x} . The following characterization holds for an arbitrary history, which, depending on the information structure of the game, may not be common knowledge. The main assumption, other than quadratic utilities and existence of a majority core point, is that beliefs are symmetric. Using Lemma 2.1, we show that majority preferences over action profiles are identical to that of the core player.

Proposition 3.1 *Let h_k be an arbitrary finite history with initial state ω , and let ν be an arbitrary probability measure on $N(h_k)$. Assume, for some $a, a' \in A(h_k)^{N(h_k)}$, that beliefs are symmetric at (h_k, a) and at (h_k, a') . Define $\pi = \nu \circ \rho^{-1}(\cdot|\omega)$,⁶ let $j \in N(h_k)$ be such that $\rho_j(\omega) = x^* \in X$ is a median in all directions, and assume π is resolute at x^* . Then*

$$\pi(\{\tilde{x} \mid \phi_{\tilde{x}}(h_k, a) \geq \phi_{\tilde{x}}(h_k, a')\}) \geq \frac{1}{2}$$

if and only if $\phi_{x^}(h_k, a) \geq \phi_{x^*}(h_k, a')$.*

By belief symmetry, the beliefs of each $i \in N(h_k)$ about the sequence of outcomes given history (h_k, 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

$$v_{\tilde{x}}(h_k, a) = (1 - \delta) \sum_{m=l_{k+1}+1}^{\infty} \int \delta^{m-l_{k+1}-1} u_{\tilde{x}}(x_m) \lambda_m(dx_m)$$

⁶ Here, we intend $\rho^{-1}(Y|\omega)$ as the set of individuals i such that $\rho_i(\omega) \in Y$.

for a type \tilde{x} active player. Letting $\lambda_{l_{k+1}}$ denote the degenerate measure with point mass on $f_{l_{k+1}}(h_k, a)$, and letting

$$\lambda = (1 - \delta) \left[\lambda_{l_{k+1}} + \delta \sum_{m=l_{k+1}+1}^{\infty} \delta^{m-1} \lambda_m \right],$$

we then have

$$\phi_{\tilde{x}}(h_k, a) = \int u_{\tilde{x}}(x) \lambda(dx).$$

That is, we can express a type \tilde{x} player’s payoff as the expected utility from the lottery λ , which is itself independent of \tilde{x} . 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.

We call history h_k a *voting stage* if (1) $A(h_k)$ consists of two elements, say 0 and 1, (2) for all action profiles $a \in \{0, 1\}^{N(h_k)}$, beliefs are symmetric at (h_k, a) , and (3) there exists a probability measure ν on $N(h_k)$ such that each player’s payoff from any measurable $a \in \{0, 1\}^{N(h_k)}$ 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 $\tilde{x} \in \mathbb{R}^d$, there exists $q_{\tilde{x}}, r_{\tilde{x}} \in \mathbb{R}$ such that

$$\phi_{\tilde{x}}(h_k, a) = \begin{cases} q_{\tilde{x}} & \text{if } \nu(\{i \mid a_i = 0\}) \leq 1/2 \\ r_{\tilde{x}} & \text{if } \nu(\{i \mid a_i = 1\}) > 1/2, \end{cases}$$

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, a single voter can never be pivotal in an election. In such cases, it makes sense to refine the possible equilibria by assuming each voter takes the action that offers the highest “potential” payoff. 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 each individual i of type \tilde{x} votes for 1 ($a_i = 1$) if

$$\phi_{\tilde{x}}(h_k, \mathbf{1}) > \phi_{\tilde{x}}(h_k, \mathbf{0})$$

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* 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 $\sigma_i(h_k) = 0$, 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 the payoff (q or r) that is higher for the core voter’s type. Again, the characterization holds even if the type profile and history are not common knowledge.

Proposition 3.2 *Let h_k be a history with initial state ω . Assume that h_k is a voting stage, with probability measure ν on $N(h_k)$, and that σ is stage-undominated at h_k . Define $\pi = \nu \circ \rho^{-1}(\cdot|\omega)$, let $j \in N(h_k)$ be such that $\rho_j(\omega) = x^* \in X$ is a median in all directions, and assume π is resolute at x^* . Then*

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

If σ is stage-undominated with default bias, then

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

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

$$\pi(\{\tilde{x} \mid \phi_{\tilde{x}}(h_k, \mathbf{1}) > \phi_{\tilde{x}}(h_k, \mathbf{0})\}) > \frac{1}{2},$$

or equivalently,

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

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 . The elected politician is assumed to implement his/her position, which voters evaluate according to quadratic utility. Conditional on challenger position 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

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

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

$$u_{\tilde{x}}(x) > \int u_{\tilde{x}}(y)b(y|z) dy \quad (4.1)$$

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 to determine the outcome of the election, assuming undominated voting strategies, following any signal realization z . The initial state $\omega = (y, z)$ specifies the position of the challenger and the signal observed by the voters; assume that the marginal on y is given by F and, conditional on y , the distribution of z is given by $G(\cdot|y)$. In the first stage, only the voters are active, each casting a ballot for the incumbent or the challenger, and in the second stage, the elected candidate is the only active player and simply picks his/her position, which is then the outcome.⁷ Voters' beliefs upon electing a candidate are symmetric: if the incumbent is elected, then the outcome in the second stage is x with probability one, and if the challenger is elected, then the outcome is given by the density $b(\cdot|z)$. 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, say 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 (4.1) 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 (4.1), 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 does not presume common knowledge of the history, the above characterization holds even if the identity of the median voter is 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, say x , which is followed by a vote; if the proposal passes, then the outcome in period t and all

⁷ For technical reasons (since we assumed an infinite number of stages in Sect. 3), we define players and action sets arbitrarily in later stages, and we specify a fixed outcome in every later stage. This does not affect the analysis.

subsequent periods is x , and each agent i receives quadratic 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. We assume that information is complete and that all prior proposals and votes are observed.

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. Given such a strategy profile, let x_i denote the alternative agent i would propose, if selected to do so. Also let A_i denote the ‘‘acceptance set’’ for i , the proposals that agent i would vote for, and let \mathcal{M} denote the collection of coalitions C containing a majority of agents. Then we define

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

as the ‘‘social acceptance set,’’ the proposals that would pass 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 (1) each agent i proposes $x_i \in \arg \max\{u_i(x) \mid x \in A\}$, and (2) each agent i votes for any proposal that is weakly preferred to continuing the bargaining process, i.e.,

$$A_i = \{x \in X \mid u_i(x) \geq (1 - \delta)u_i(q) + \delta v_i\}. \quad (4.2)$$

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 (2005).

This model can be mapped into the framework of the previous section to characterize the social acceptance set in any stationary equilibrium. We let the initial state determine the sequence of proposers over time, not observed by the agents, where marginals on proposers are iid, with probabilities given by p_1, \dots, p_n . Ideal points of the agents are exogenously fixed. In the first stage, the first proposer is the only active player, and he/she chooses an alternative in X ; in the second stage, all players are active and vote either to accept or reject the proposal; and so on. Outcomes are generated every time a proposal is voted on, as follows: if a proposal

has been passed, then that alternative is the outcome;⁸ otherwise, the outcome is the status quo. Assuming all agents use stationary strategies and since information is complete (so beliefs are symmetric), a stage following a proposal by any agent is a voting stage with the counting measure on the set of agents. Furthermore, any no-delay stationary equilibrium will be 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, say 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 ending in the selection of a proposer, 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 2005), 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 (2005), 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 \tilde{x} receives quadratic utility $u_{\tilde{x}}(y_t)$ 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

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} u_{\tilde{x}}(y_t)$$

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

⁸ Thus, once a proposal is passed, the game is essentially over. We focus on histories such that no proposal has already been passed.

Voter ideal points are private information. Borrowing from earlier notation, we use $\rho: N \rightarrow X$ to denote a type profile, where ρ_i is the ideal point of voter i . We assume that type profiles are measurable and that the distribution over type profiles is such that the marginal of each voter's ideal point is 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.⁹ 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 also (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 (1) a voter with ideal point \tilde{x} always chooses the same point $p_{\tilde{x}}$ if ever elected to office, where $p_{\tilde{x}}$ is measurable with respect to \tilde{x} , and (2) a voter with ideal point \tilde{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 $\bar{u}_{\tilde{x}}$, which is measurable as a function of \tilde{x} . Thus, a voter with ideal point \tilde{x} votes for the incumbent if and only if $u_{\tilde{x}}(y) \geq \bar{u}_{\tilde{x}}$. Letting

$$A_{\tilde{x}} = \{y \in [0, 1] \mid u_{\tilde{x}}(y) \geq \bar{u}_{\tilde{x}}\}$$

denote the acceptance set for a voter with ideal point \tilde{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

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

as the social acceptance set, the set of policies that, if chosen, would lead to reelection of the incumbent. Because $\bar{u}_{\tilde{x}}$ 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

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

⁹ Theorem 2 in Judd (1985) establishes the existence of such a distribution over type profiles.

where

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

a jointly measurable function. Then measurability of $\int \psi(\cdot, z)dF$ 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 most recent policy choices of these politicians. For each politician m , let $Y_m = p^{-1}(y_m)$ if this set is non-empty, and otherwise let $Y_m = \{0\}$.¹⁰ Then the updated beliefs of a type \tilde{x} voter are given by the prior over ρ conditioned on $\rho_{i_m} \in Y_m, m = 1, 2, \dots, k$, and $\rho_i = \tilde{x}$. By finite-independence, the conditional beliefs about the types of all other voters are still given by F . Let $W = \{\tilde{x} \in X \mid p_{\tilde{x}} \in A\}$ denote the set of ideal points of winning office-holders, and let $L = X \setminus W$ denote the set of ideal points of losing office-holders, 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 \tilde{x} as

$$v_{\tilde{x}} = \frac{\int_W u_{\tilde{x}}(p_z) F(dz) + (1 - \delta) \int_L u_{\tilde{x}}(z) F(dz)}{1 - \Pr(L)\delta}, \tag{4.3}$$

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 (3) an officeholder with ideal point \tilde{x} chooses $y \in \arg \max\{u_{\tilde{x}}(z) \mid z \in A\}$ if

$$\max\{u_{\tilde{x}}(z) \mid z \in A\} \geq \delta v_{\tilde{x}}$$

and otherwise chooses $y = \tilde{x}$, and (4) a voter with ideal point \tilde{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

$$u_{\tilde{x}}(y) \geq v_{\tilde{x}}. \tag{4.4}$$

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 they 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 (2006) that allows for multiple policy dimensions and arbitrary concave utility functions, but we do assume for simplicity a finite number of possible types there. In that case, we must allow officeholders with

¹⁰ Without 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.

ideal points $\tilde{x} \notin A$ 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.

We can map this model into the framework of the previous section to characterize the social acceptance set of any simple equilibrium. We let the initial state $\omega = (\gamma, \rho)$ consist of two components. First, $\gamma = (i_1, i_2, \dots)$ is the sequence of challengers over time, not observed by voters. We endow the space $[0, 1]^{\mathbb{N}}$ of sequences of challengers with the sigma algebra generated by rectangles of the form $N_1 \times N_2 \times \dots$, where N_1, N_2, \dots are measurable subsets of N , and we assume the marginals on challengers are distributed uniformly and independently in each period. Second, a state determines the type profile of the voters. We endow the space of type profiles with the sigma algebra generated by sets of the form

$$\{\rho \mid \rho_i \in C\},$$

where $i \in N$ and C is a measurable subset of X . We then extend this sigma algebra so that the space of regular profiles is measurable, and we assume the regular profiles have probability one (cf. Judd 1985; Theorems 1 and 2). We assume that sequences of challengers and type profiles are distributed independently.

In the first stage, the initial officeholder is the only active player and chooses a policy in X ; in the second stage, all voters other than the candidates are active and vote between the incumbent and challenger, with the winner taking office next period;¹¹ and this process is repeated. Letting 0 denote a vote for the incumbent and 1 a vote for the challenger, we represent election returns by a measurable function $r: N \rightarrow \{0, 1\}$, and we let R denote the set of such election returns. A finite history in this electoral game is then a sequence $(x_1, i_2, r_2, x_2, i_3, r_3, x_3, \dots)$, which represents the period one policy choice, the period two challenger, the period two election return, the period two policy choice, 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 measurable 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, respectively, the events that neither candidate won, that the incumbent won, that the challenger won, and that at least one of the candidates won. The outcome functions for finite histories are simply defined as projections onto the policy choices of officeholders, e.g.,

$$f(x_1, i_2, r_2, x_2, \dots, i_k, r_k, x_k) = (x_1, x_2, \dots, x_k),$$

which are measurable with respect to the above sigma algebras on the H^k s.

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

$$\mathcal{S}_1 = \{\omega = (\gamma, \rho) \mid \rho_{i_1} \in p^{-1}(X_1)\},$$

¹¹ We exclude previously selected politicians to maintain belief symmetry in voting stages. Since they are finite in number, this leaves the strategic aspects of the game unaffected.

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

$$\mathcal{S}_2 = \mathcal{S}_1 \cap \{\omega = (\gamma, \rho) \mid i_1 \in N_2\},$$

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

$$\mathcal{S}_3 = \mathcal{S}_2 \cap \{\omega = (\gamma, \rho) \mid \rho_{i_1} \in W\},$$

which is measurable because W is measurable, and similarly for other collections of election returns.¹² Measurability of profiles corresponding to histories of greater length can be established by an induction argument based on these observations. With beliefs as specified above, this allows us to define probability measures $\xi_i(h_k)$ on outcome sequences; continuation values are as in (3.1); and a stage following the choice of a policy by any officeholder is a voting stage. Because X is one-dimensional and F is continuous and strictly increasing, 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 any simple equilibrium and any regular history ending with selection of a policy, the 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.

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

References

- Aliprantis C, Border K (1999) *Infinite dimensional analysis: a Hitchhiker's guide*. Springer, Berlin Heidelberg New York
- Banks J, Duggan J (2000) A bargaining model of collective choice. *Am Polit Sci Rev* 94:73–88
- Banks J, Duggan J (2005) A general bargaining model of legislative policy-making. *Q J Polit Sci* 1:49–85
- Banks J, Duggan J (2006) A dynamic model of democratic elections in multidimensional policy spaces. Mimeo, University of Rochester
- Banks J, Duggan J, Le Breton M (2006) Social choice and electoral competition in the general spatial model. *J Econ Theory* 126:194–234

¹² Non-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.

-
- Baron D, Ferejohn J (1989) Bargaining in legislatures. *Am Polit Sci Rev* 83:1181–1206
- Bernhardt D, Dubey S, Hughson E (2004) Term limits and pork barrel politics. *J Public Econ* 88:2383–2422
- Black D (1958) The theory of committees and elections. Cambridge University Press, Cambridge
- Calvert R (1985) Robustness of the multidimensional voting model: candidate motivations, uncertainty, and convergence. *Am J Polit Sci* 29:69–95
- Campuzano L (2001) Parties in a repeated elections model. Mimeo, University of Rochester
- Cho S, Duggan J (2003) Uniqueness of stationary equilibria in a one-dimensional model of bargaining. *J Econ Theory* 113:118–130
- Davis O, DeGroot M, Hinich M (1972) Social preference orderings and majority rule. *Econometrica* 40:147–157
- Downs A (1957) An economic theory of democracy. Harper and Row, New York
- Duggan J (2000) Repeated elections with asymmetric information. *Econ Polit* 12:109–135
- Duggan J, Fey M (2005) Electoral competition with policy-motivated candidates. *Games Econ Behav* 51:490–522
- Duggan J, Fey M (2006) Repeated Downsian electoral competition. *Int J Games Theory* (forthcoming)
- Judd K (1985) The law of large numbers with a continuum of IID random variables. *J Econ Theory* 35:19–25
- Kingman J, Taylor S (1966) Introduction to measure and probability. Cambridge University Press, Cambridge
- McGarvey D (1953) A theorem on the construction of voting paradoxes. *Econometrica* 21:608–610
- McKelvey R (1976) Intransitivities in multidimensional voting models, and some implications for agenda control. *J Econ Theory* 2:472–482
- McKelvey R (1979) General conditions for global intransitivities in formal voting models. *Econometrica* 47:1086–1112
- McKelvey R (1986) Covering, dominance, and institution-free properties of social choice. *Am J Polit Sci* 30:283–314
- McKelvey P, Ordeshook P, Ungar P (1980) Conditions for voting equilibria in continuous voter distributions. *SIAM J Appl Math* 39:161–168
- Meirowitz A (2003) Accountability and representation in repeated elections with uncertain policy constraints. Mimeo, Princeton University
- Plott C (1967) A notion of equilibrium and its possibility under majority rule. *Am Econ Rev* 57:787–806