

Satisficing and Selection in Electoral Competition¹

Jonathan Bendor² Dilip Mookherjee³ Debraj Ray⁴

November 15, 2005

¹We thank Alberto Diaz-Cayeros, Jim Fearon, Tim Groseclose, Sunil Kumar, Dave Laitin, Nolan McCarty, Adam Meirowitz, Ken Shotts, and seminar participants at Berkeley, Stanford, U.C. Irvine, UCLA, the University of Chicago, the 2001 APSA meetings and the 2002 MPSA meetings for their helpful comments, and our research assistants, past and present—Bill Hauk, Adam Meirowitz, Dave Siegel, Jon Woon and Muhamet Yildiz—for their work on the computational model.

²Graduate School of Business, Stanford University, Stanford CA 94305-5015 (corresponding author).

³Professor, Department of Economics, Boston University, Boston MA 02215.

⁴Professor, Department of Economics, NYU, 269 Mercer St., New York City NY 10012.

Abstract

We model political parties as adaptive decision makers who compete in a sequence of elections. The key assumptions are that *winners satisfice* (the winning party in period t keeps its platform in $t + 1$) while *losers search*. Under fairly mild assumptions about losers' search rules, we show that the sequence of winning platforms is absorbed into the top cycle of the (finite) set of feasible platforms with probability one. This implies that if there is a majority rule winner then ultimately the incumbent party will espouse it. However, our model, unlike Downs-Hotelling or Kollman-Miller-Page, does *not* predict full convergence: we show, under weak assumptions about the out-party's search, that losing parties do not stabilize at the majority rule winner (should it exist). We also establish, both analytically and computationally, that the adaptive process is robust: if a majority rule winner “nearly” exists then the trajectory of winning platforms tends to be “close” to the trajectory of a process which actually has such a winner.

When *An Economic Theory of Democracy* (1957) was published, behavioralism was already a force in American political science. More a mood (Dahl 1961) than a research program, its major impulse was to make the discipline more scientific (ibid., p.766). In this respect *Economic Theory*, more rigorous than most books on elections, fit in well.

But the book's foundation—the idea of rational choice—differed sharply from the intellectual tendencies of most leading behavioralists, who were trained in or influenced by social psychology and sociology. And though no detailed behavioral model of choice then existed, there *were* theoretical ideas, mostly social psychological in character. As works such as *The American Voter* revealed, decision makers described via psychological notions look quite different from the rational actors of *Economic Theory*.

Given these differences, mainstream behavioralists could have reacted to the book ambivalently, praising its rigor but questioning its micro-premises. This did not happen, at least not initially. Indeed, aside from Lindblom's prescient review, it is hard to find *any* written reactions to *Economic Theory* from major political scientists, behavioralist or otherwise.⁵

More importantly, behavioralism did not generate an alternative theory of electoral competition—certainly none that has mounted a serious challenge to Downs' ideas. Instead, behavioralists who studied parties and elections mostly ignored both Downs' book and the impressive research program that it spawned. (Kelley (1965) and Schlesinger (1966) were notable exceptions.) Hence, the rational choice program and a behavioral one have not competed head-to-head in this field.

Such a contest is possible, however. Behavioralism has valuable intellectual resources that could generate a coherent alternative: in particular, Herbert Simon's now-famous essay on satisficing (1955) contains key elements of a behavioral theory of choice. We propose to construct a behavioral model of elections based on Simon's paper, coupled to the Schattschneider–Schumpeter–Downs macrohypothesis that in vigorous democracies major parties are organized to win elections. We model political parties as adaptive decision makers who compete in a sequence of elections. Our central premises about decision making closely follow Simon's analysis: *winners satisfice* (the winning party in period t keeps its platform in $t + 1$) while *losers search*. Simon's general notion of an agent's aspiration level is thus represented here by the domain-specific hypothesis that winning an election is satisfying while losing isn't.

⁵Stokes published a thoughtful critique of the book in 1963; this was, we believe, the first critical, behaviorally oriented assessment by an eminent political scientist. Though it would not be the last—Green and Shapiro (1994) devote a chapter to the Downsian program—there have not been many.

A key motivation for this approach is that politicians usually are uncertain about voter preferences. To be sure, parties conduct numerous polls; yet uncertainty often persists throughout campaigns and sometimes even *after* an election has been decided. (E.g., the fierce debates among Democrats over Kerry’s loss indicate that even the past can be cloudy.) Many Downsian models simply ignore this uncertainty. Others incorporate it via a standard game theoretic formulation, i.e., as a game of incomplete information, with each party’s uncertainty depicted by subjective priors over the distribution of voter preferences. This approach’s logic requires the parties to think about their uncertainties precisely and consistently, i.e., they have common knowledge about their respective prior distributions. This presumes a cognitive capacity for dealing with complexity and level of coordination of their information bases that strikes even pure game theorists (Kreps 1990; Rubinstein 1998) as unrealistic. Positing that in such contexts agents experiment with different policies and learn from experience is, we believe, a more plausible account of their behavior.⁶

To lay bare the logic of satisficing-and-search in two-party competition, we present a stripped-down model. It assumes, as do most models of repeated elections, that voters’ preferences stay put. Empirically, of course, these preferences do change over time, but posing the problem this way allows us to focus our analytical attention on how parties adapt and adjust their policies as they try to win office. Understanding how policies evolve given fixed voter preferences is a necessary first step for understanding their evolution in more dynamic settings. Hence, we do not expect our model to be the last behavioral formulation in this area. Instead, we see Simon’s work as a foundation for a research program on electoral competition that hopefully will address many topics, including turnout and how citizens vote.⁷ Because we want to show what a behavioral model can do on the turf defined by the incumbent program, this paper emphasizes standard questions, e.g., Do the two parties’ platforms converge to the same policy?

The rest of the paper is organized as follows. The second section reviews relevant literature. Section three presents the model and several implications. Proposition 2 shows that if winners satisfice, then experimentation by losers is *necessary* for a well-defined type of electoral “progress”. Proposition 3 demonstrates that if experimentation has certain weak properties then it and satisficing-by-winners are *sufficient* to ensure that the sequence of

⁶Indeed, satisficing-and-search and the related rules of reinforcement learning are sometimes considered algorithms for learning how to maximize complicated payoff functions in complex dynamic games with uncertainty (Arthur 1993; Fershtman and Pakes 2004; Narendra and Thathachar 1989).

⁷This research program will develop its own internal logic, as they are wont to do (Ferejohn 1995). This will include its own list of questions and problems, e.g., When do people learn to vote sophisticatedly? What problems do they find cognitively difficult (Kotovsky et al. 1985)?

winning policies converges to the policy space's top cycle set with probability one. Hence, if there is a majority rule winner then ultimately the incumbent party will espouse it. However, Propositions 4 and 5 show, given weak assumptions about the out-party's search, that when a median voter exists both parties do *not* stabilize at her bliss point. Thus, in contrast to both the Hotelling-Downs rational choice theory *and* Kollman, Miller and Page's adaptive model (1992), full convergence is not predicted.

The fourth section investigates alternative specifications of the challenger's search behavior by endowing him with different degrees of sophistication and certain kinds of knowledge about the political terrain, following Kramer (1977), Miller (1980) and Ferejohn et al. (1980, 1984). Section five analyzes whether our results are sensitive to small changes in key assumptions. Proposition 6 shows that they are robust: e.g., if a majority rule winner nearly exists then the trajectory of winning platforms tends to be close to a trajectory of a process that *does* have a generalized median. We then present a computational model that provides further results for ill-structured electoral environments. Computational results show that in the steady state winning policies are centrally located, and their dispersion is strongly correlated with the size of the uncovered set. The last section concludes.

RELATED WORK

Kollman, Miller and Page (1992, 1998; henceforward KMP) pioneered work on adaptive parties. In their simulation model winners satisfice; challengers generate platforms via adaptive search algorithms, and office-oriented ones then select the vote-maximizing platform. KMP showed that the distribution of winning platforms in a two-dimensional policy space tends to be centrally located. When the simulation is re-run in a unidimensional setting (the papers don't present results on this case), office-oriented parties converge to the median voter's ideal point (Page, personal communication, 1999).⁸ Hence, *this* adaptive model yields the same long-run prediction as Hotelling-Downs.

Unfortunately, for several reasons KMP's impact on the field has been less than their papers deserved. First, in one sense KMP were too ambitious: they analyzed multidimensional policy spaces without examining the unidimensional setting. Second, in another respect, however, KMP were not ambitious enough: their papers settled for simulation results rather than analytical ones. This took their work out of the Downsian mainstream, which emphasizes mathematical models and analytical results. The combination of unorthodox substantive premises and a deviant method of analysis may have marginalized it.

⁸If the set of platforms is finite then one can prove (simulations are unnecessary) that two of KMP's search rules yield convergence to the median voter with probability one.

Third, there are many plausible ways to model bounded rationality. Hence, to get robust results we should posit *general properties of adaptation* (e.g., successful actions are more likely than unsuccessful ones to be repeated) rather than specific functional forms. Simulation is ill-suited to the general approach: a computer must be told exactly what kind of search rule to use. Sensitivity testing (which KMP did) can alleviate but not eliminate this problem. Indeed, because there are many types of bounded rationality, an efficient form of sensitivity testing is to establish results analytically.

This paper complements KMP by using a different research strategy. First, we establish results that hold for the canonical Downsian setting: a unidimensional policy space and single-peaked preferences. Only then do we move to the less well understood multidimensional environment. Second, we push analytical results as far as we can, turning to simulation only when necessary. (We agree with KMP that when analytical results are elusive, computing beats giving up.) Third, in the interests of generality and robustness we specify a few general properties of adaptation rather than relying on detailed heuristics. These features are interdependent—for modelling general properties of adaptation math is better than simulation—so they form a coherent approach with its own distinctive set of tradeoffs.

Unorthodox Downsian work

Although KMP were probably the first to make limited rationality *central* to models of party competition, they were not the first to incorporate aspects of bounded rationality in such models. Kramer (1977) and the matched pair of Ferejohn, Fiorina and Packel (1980) and Ferejohn, McKelvey and Packel (1984) [henceforth FFMP] were mostly Downsian in nature. They focused (Kramer, p.311-315; Ferejohn et al. 1980, p.140-141; Ferejohn et al. 1984, p.45-46) on the so-called “chaos problem”, which arises out of the internal logic of the Downsian project (Ferejohn 1995). None of these papers made rationality a central concern. But in several ways they were unorthodox members of the program, and some of these respects involved nonoptimizing behavior.⁹ Because these papers provide some interesting contrasts to the present work, we take a brief look at their features now.

First, all involve *status quo based processes*: in Kramer’s electoral model, today’s winner keeps the same platform tomorrow; in FFMP’s committee model, today’s winning alternative is tomorrow’s status quo. Though the authors do not give a cognitive interpretation for this¹⁰, we think there is a natural one: winners don’t fix what isn’t broken. However, this satisficing

⁹We think that a behavioral program could embrace Kramer (1977) and FFMP as pioneering papers. On how a new program appropriates theories once part of an earlier one, see Laudan (1977, p.93-95).

¹⁰Kramer is especially terse about this assumption, saying only that “In each period one of the parties is elected, enacts the policy it advocated, and in the next election must defend this same policy” (p.317).

interpretation of status quo based processes is ours, not theirs. Indeed, this assumption fits awkwardly with the Downsian perspective.

Second, agents are myopic. FFMP make this explicit (Ferejohn et al. 1980, p.144), but it also squares with Kramer’s model in two ways. (a) That the winner keeps yesterday’s platform is clearly myopic. (b) Kramer also assumes that the out-party selects a platform that maximizes vote-share against the status quo policy. This is an example of what is now called myopic best response: an alternative is chosen today without regard for the long run.

Third, nonequilibrium (i.e., non-Nash) solution concepts are used. This, we believe, follows naturally from the agents’ myopic qualities.

Because we will compare our results with Kramer’s and FFMP’s in section four, we defer a description of their findings until then.

THE MODEL AND ITS IMPLICATIONS

We study a standard electoral game: a contest between two candidates or parties. A few modifications have been introduced to enhance analytical tractability. There is a finite set of citizens, $N = \{1, \dots, n\}$, with n odd, and a finite set of policies or platforms, $X = \{x_1, \dots, x_m\}$, with $m > 1$. (X may be huge, but finiteness is analytically useful for several results.) Each citizen has a strict preference ordering over policies. If the candidates adopt the same policy then voters break their indifference by independently tossing nondegenerate coins. A voter’s coin need not be fair; it might, e.g., be biased toward the incumbent. Voters may have different probabilities of voting for the incumbent or for either party. For simplicity we assume that all these random procedures are Markovian and stationary: they are independent of the electoral history prior to t and of t itself.

Since n is odd the majority preference relation is also strict: for any two policies, x_i and x_j , either a majority of citizens strictly prefer x_i to x_j (written $x_i \succ x_j$) or x_j to x_i . Further, since the above assumption of indifference-breaking means that individual voters always reach a decision, elections are conclusive even when the candidates adopt the same policy. (In these convergent outcomes, vote-shares are random variables.)

We do not require spatial policies or preferences, but we can recover spatial settings (of different dimensionalities) when doing so is desirable. The following description of the majority preference relations induced by the above assumptions enables us to make this translation to a spatial framework. We partition the policy set into z disjoint subsets, $\{L_1, \dots, L_z\}$, where $1 \leq z \leq m$, by iteratively applying the idea of the top cycle set.¹¹ Let

¹¹The top cycle set can be defined constructively. In our setting, a policy is in the top cycle set if and only if it is reachable from every other policy by a chain of strict majority preference.

L_1 be X 's top cycle set. (As is known, L_1 cannot be empty.) Consider the reduced set, X/L_1 . If this too is nonempty (it may not be), let L_2 denote the top cycle of X/L_1 . Proceed in this way until all $x \in X$ are assigned to an L . This procedure must terminate since X is finite.

We call these subsets of X ‘levels’ to suggest a mental picture: one can see the electoral environment as a series of levels or plateaus (if $z > 1$). Each plateau electorally dominates those below it: each policy at a given level is majority-preferred to *all* policies at lower levels.¹² Further, every policy at a level *covers* (Miller 1980) all policies at lower elevations.¹³ Within a level, however, no policy beats all others, nor does any lose to all others at its level. Thus in nonsingleton levels, every two policies are joined by a majority rule cycle. Hence adaptive parties may get “hung up” on the cycles that pervade nonsingleton levels. (Figure 1 gives an example of a policy set with three levels; each level has a cycle.)

[figure 1 about here]

Partitioning the policy set into levels is useful: it allows us to analyze concisely how the parties hill-climb—or “plateau-climb”—in the electoral landscape.

The partitioning also allows us to describe different majority preference relations. For example, let us use it to examine the classic spatial context: citizens with single-peaked preferences defined over a unidimensional policy space. For simplicity, assume that preferences are symmetric (e.g., quadratic utility). Symmetric preferences, plus our assumption that each voter’s preference ordering is strict, together imply that the m policies are strictly ordered in Euclidean distance from the median voter’s ideal point, which we call x_{mv}^* . Because the median voter is decisive here, the majority preference relation is identical to her preference ordering: the policy closest to x_{mv}^* is the Condorcet winner, the next closest policy loses to the Condorcet winner but beats everything else, and so on. Hence, in this classic setting the L -sets have a very simple structure. First, each L contains exactly one policy. Second, the policy in L_1 is the Condorcet winner, L_2 's policy is the median voter’s second most preferred option, and so on all the way down to L_m , whose policy is the farthest from x_{mv}^* . Thus here the L 's form a staircase (one policy per step) that leads up to the Condorcet winner.

The picture is naturally more complicated in most multidimensional policy spaces, since

¹²This follows from two facts. First, no policy outside the top cycle set can beat anything in it (Austen-Smith and Banks 1999, p.169.) Second, our assumptions rule out ties, so any x in the top cycle must beat any y outside it.

¹³This implies that X 's uncovered set is a subset of L_1 , the uncovered set of X/L_1 is a subset of L_2 , etc. But some policies in a given L_r may cover other policies in L_r (Bendor, Mookherjee and Ray 2004).

these generically lack a generalized median. In such contexts the L 's will often not be singletons. Figure 2 depicts a two-dimensional space with three voters (again with symmetric preferences for simplicity) and six policies. The relatively centrist platforms, x_1 , x_2 and x_3 , form a cycle; each of them beats all of the more distant policies, x_4 , x_5 and x_6 , which also form a cycle. Hence L_1 is composed of the more centrist policies and L_2 , the less centrist ones.

[figure 2 about here]

Consistent with idealizations that are standard in models of electoral competition, we assume universal turnout and sincere, error-free voting. (Proposition 6 shows that our results are robust regarding these idealizations.) In short, it is assumed that if $x_i \succ x_j$ then x_i will defeat x_j in an election.

There is an indefinitely long sequence of elections, $1, 2, \dots$, with one election per period. (At the start an incumbent party and its platform are randomly picked.) In every election the candidates simultaneously announce platforms; then citizens vote. The party that gets more votes wins the election and is the incumbent at the start of the next period.

The state variables of the associated stochastic process are the platforms of the incumbent, I_t , and the challenger, C_t . Thus, I_t is the policy that won the election in $t-1$. We focus on how I_t —what Kramer (1977) called the *trajectory of winning policies*—evolves. (Because of I_t 's importance we sometimes abuse terminology and refer to it as if it were *the* state variable.)

The parties' adaptive behavior

The heart of the model is how the candidates respond to winning and losing. As already noted, we assume that *winners satisfice* (Simon 1955).

(A1): the winner of the election in t stands for re-election in $t+1$ by keeping the platform that won him office in t .

In effect, the candidates have an implicit aspiration level in-between the payoffs of winning and losing. Hence winning is satisfying, and incumbents don't fix what isn't broken.¹⁴ (Apart

¹⁴Exogenously fixed implicit aspirations in-between the payoffs of winning and losing are equivalent to explicit aspirations that adjust endogenously as follows. Let w be the payoff of winning and l , losing ($w > l$). Use the standard assumption (Cyert and March 1963) that tomorrow's aspirations are a weighted average of today's aspiration and today's payoff: $a_{i,t+1} = \lambda a_{i,t} + (1 - \lambda)\pi_{i,t}$, where $a_{i,t}$ is party i 's aspiration in t , $\pi_{i,t}$ is its electoral payoff (w or l), and the adjustment parameter λ is in $(0, 1)$. If initially aspirations are "conventional"—winning is satisfying, losing isn't—then the endogenous sequence of aspirations will always be in (l, w) . So for all possible sequences of elections, winning will be satisfying but losing won't.

from proposition 5, we do not explicitly model aspirations or their dynamics. See Bendor, Mookherjee and Ray (2001) for an overview of such models.) Similarly, losing is dissatisfying, and we follow Simon (1955) and Cyert and March (1963) in positing that dissatisfaction triggers search. Hence our *challengers* (sometimes) *search*: in at least some elections they try policies that differ from their losing one from the previous election.¹⁵

Before examining how challengers search, we report an important implication of satisficing which will help us understand which assumptions drive which conclusions. Proposition 1 refers to a stochastic process produced by the behavior of incumbents and challengers in the electoral environment of our z levels and (A1). So far, only part of this process has been defined: incumbents' (deterministic) behavior.¹⁶ Randomness arises from the challenger's search, as we will see shortly.

Proposition 1: Suppose (A1) holds. If I_t is in L_r , then thereafter it must be at that level or higher.

Thus, assuming that incumbents satisfice ensures that electoral outcomes never slip downhill to lower plateaus.¹⁷ The government's policy must either stay where it is or climb higher. The proof is simple. Suppose $I_t \in L_r$. The challenger either proposes a policy from a lower level or not. If it is lower then the construction of the level sets plus sincere voting imply that the challenger must lose. Because winners satisfice, $I_{t+1} = I_t \in L_r$. If he proposes $x \in L_r$ then no matter who wins, $I_{t+1} \in L_r$. Finally, if the challenger proposes a higher platform then he wins and satisficing-by-winners implies $I_{t+1} \in L_q$, where $q < r$. Induction does the rest.

Proposition 1 presumes *nothing* about the challenger. The result follows from the combination of the static electoral stage described by the z levels and satisficing by winners. Recall, however, that we also assume error-free voting. Thus, the parties generate options that the voters test (per Simon's [1964] adaptive "generate-and-test" process), *and the test phase is flawless*. Error-free tests plus satisficing imply that the process cannot slip downhill.

So, half of the adaptive picture—how incumbents behave—acts as a *brake* on the process. The other half, how challengers search, supplies uphill momentum.

¹⁵This is such a weak premise that we rarely label it a formal assumption. (We use it explicitly only in proposition 4.) If it didn't hold then *both* parties would always keep the platforms they championed in period one. This is neither realistic nor interesting.

¹⁶Proposition 1 can therefore analyze specific *sample paths*. It is not confined to probability distributions over a population of sample paths.

¹⁷Proposition 1 has bite only when there are multiple policy levels ($z > 1$).

A key aspect of search is discovering new alternatives, as defined below.

Definitions. A party *experiments* in period t if it selects a platform that it has never used before t . It *innovates* in t if it espouses a platform that neither party has ever used.

Because incumbents satisfy they don't experiment. Only challengers can generate novel policies. The importance of this responsibility is underscored by the next result, which shows that experimentation's strong form—innovation—is *necessary* for upward progress, i.e., for the trajectory of winning policies to move uphill.¹⁸

Proposition 2: Suppose (A1) holds. For all periods s and t where $s < t$, if I_s is in L_r and no challenger in $[s, \dots, t]$ innovates, then with probability one I_t is also in L_r .

The proof is simple. By proposition 1, winning policies never slip downhill. Thus, if $I_s \in L_r$, all prior winning platforms (hence losing ones as well) have come from level r or lower. Therefore, if no challenger innovates in $[s, \dots, t]$, then all of them must have chosen platforms from level r or lower. But if so and no innovation occurs in $[s, \dots, t]$, then the set of available policies in that era is in $(L_r \cup \dots \cup L_z)$. So, $I_{s+1} \in (L_r \cup \dots \cup L_z)$, whence by induction all winning policies in $[s, \dots, t]$ must also be in $(L_r \cup \dots \cup L_z)$. Hence innovation is necessary for progress.¹⁹

The electoral dynamic need not grind to a halt if no one innovates: an incumbent can be beaten even if his opponent does not experiment, much less innovate. For example, suppose $L_r = \{a, b, c\}$ and $a \succ b \succ c \succ a$. Then the status quo policy can change endlessly without experimentation, as the parties follow the cycle. But this behavior cannot kick the trajectory of winning policies up to a higher plateau.

When does experimentation *suffice* for upward progress? Consider this condition.

(A2): There is an $\epsilon > 0$ such that for every history and in every election in which the challenger hasn't already tried everything, the probability s/he experiments is at least ϵ .

The challenger has good reason to try something other than his last platform: it lost in the last election and the incumbent has stayed put. Further, experimentation is reasonable because *everything* that the challenger has tried is electorally flawed.

¹⁸As noted earlier, it is trivially true that search by challengers is necessary for upward progress. Proposition 2 is stronger: search *that yields systemic novelty* is required for progress. (By searching its memory for policies that have worked in the past but which it hasn't tried recently, a party could search without experimenting. And it could experiment without innovating.)

¹⁹Neither proposition 1 nor 2 require that the set of policies be finite.

Remark 1: Every platform that the out-party of $t > 1$ has used before t has lost at least once.

The proof is by contradiction. If remark 1 didn't hold, then the challenger in t must have espoused a policy, x , that won in an earlier period s . Satisficing implies that that party would have continued to uphold x in $(s + 1, \dots, t - 1)$, triumphing in all those elections. But then in t that party would be the incumbent, not the challenger.

If we consider platforms that have lost in the past to be electorally damaged goods, then we see why (A2) would hold. Further, (A2) is a weak assumption about challengers: it says nothing specific about what they know or how sophisticated they are. A challenger might, e.g., know the entire policy set, have beliefs about voters' preferences about platforms, and sophisticatedly update those beliefs. Or she might grope about blindly. (A2) is only a summary statement about the effects of the out-party's knowledge and strategic sophistication: it always has *some* chance of experimenting (when that is possible). Moreover, (A2) does not presume any specific stochastic model or even any familiar *class* of stochastic models. In particular, (A2) is not Markovian: what happens today could depend on events in the distant past.²⁰

Although (A2) is a weak assumption about experimentation, it and satisficing-by-incumbents ensure that the trajectory of winning policies converges to X 's top cycle. Satisficing by incumbents and experimentation by challengers work well together: the former prevents the process from slipping downhill and the latter provides upward momentum.

Proposition 3: If (A1) and (A2) hold, then the trajectory of winning policies converges to and is absorbed by L_1 with probability one.

The proof (in the appendix) relies on the fact that any level below L_1 is *transitory*: if I_t ever goes to such a level, (A2) ensures that eventually it must leave there and, by proposition 1, never return. Since *all* lower levels are transitory, long-run convergence to the top level is guaranteed. And proposition 1 implies that L_1 is absorbing: once a trajectory of winning policies reaches the top, it stays there.

Proposition 3 implies that if L_1 is a singleton then the trajectory of winning policies converges to that unique platform. (The challenger may keep moving around; more on this

²⁰For example, the chance of experimenting could rise the more often the out-party loses, as (e.g.) desperation to reclaim the White House sets in. And the magnitude of the increases could depend on *when* the losses occurred: e.g., the more recent the losses, the more the challenger wants to experiment. Because the party's memory is not itself a state variable in our model, this isn't Markovian.

shortly.) Thus when a Condorcet winner exists, *the government's policy must converge to it.*²¹

Convergence to the top cycle can be established without presuming that the electoral dynamic belongs to a special class of stochastic processes. In particular, the process could have multiple limiting distributions.²² (For an example, see Bendor, Mookherjee and Ray 2004.) However, although convergence to a unique long run distribution is not *necessary* for trajectories of winning policies to be absorbed into L_1 , this property is *sufficient* for this convergence: as we show elsewhere (Bendor, Mookherjee and Ray 2004), assuming that the process converges to a unique probabilistic steady state substitutes for positing that the challenger experiments.

It is important to note that the convergence of the I_t 's into L_1 is not equivalent (even when L_1 is a singleton) to the famous Hotelling-Downs prediction that the parties adopt identical policies. In view of the facts—candidates in two-party systems rarely champion identical platforms (Ansolabehere, Snyder and Stewart 2001; Levitt 1996)—that proposition 3 does not imply the median voter result is good for the model's empirical standing.

However, though this result does not imply complete convergence, it does not preclude it either. Proposition 3 is silent on the matter because it uses (A2), which does not say what challengers do once they can no longer experiment because they've tried all feasible platforms. This is unlikely if X is large, but it is not ruled out. To tie up this loose end, consider the following weak formalization of Simon-Cyert-March's idea that dissatisfaction triggers search. (It also incorporates the domain-specific assumption that losing an election is dissatisfying. For a justification of this via endogenous aspirations, see footnote 15.)

(A3): There is an $\epsilon > 0$ such that for every history, the challenger in $t+1$ adopts a platform that differs from the one it lost with in t with a probability of at least ϵ .

This weak specification of problem-driven search ensures that the electoral process will not settle down into a tweedledum-tweedledee pattern. Note that (A3) is sufficiently weak—e.g., it is not Markovian—so that the next result cannot pin down the properties of the

²¹Determining the speed of convergence requires search-assumptions that are stronger than (A2). We explore this issue in the next section. For now, note that even unsophisticated search can push winning platforms uphill rapidly. E.g., if the challenger searches blindly then the chance of moving up in one period equals the fraction of policies that are in plateaus above the incumbent's. Thus, if I_t is at a low level then the probability of hill-climbing in one election is substantial.

²²In one important respect we needn't worry about the number of limiting distributions. By proposition 3, the electoral dynamic goes into L_1 for sure. So no matter where in L_1 it winds up, eventually the government's policy must be electorally undominated, beating all policies in lower levels.

process's limiting distribution(s). But it can say what the process will *not* do. (Its proof, being obvious, is omitted.)

Proposition 4: If (A3) holds then no state in which the parties adopt identical platforms is absorbing, nor is any set of such states.

Proposition 4 implies that *if* the process converges to a unique limiting distribution, then that distribution must put positive weight on states in which the parties adopt different platforms. Further, if it has multiple stationary distributions, then *all* of them must put weight on circumstances in which the parties offer distinct policies. That parties don't converge to the same position emerges naturally from two weak behavioral premises: (1) losing can be dissatisfying and (2) dissatisfaction triggers search.²³

How Robust is Divergence? Proposition 4 presumes that candidates are purely office-oriented. But many scholars have argued that they probably also care about policy and, since passing legislation is easier with a big mandate, vote-share as well. Is the divergence result robust with respect to assumptions about candidate motivation? We address this question now.

Let $u_i(w, v, x)$ denote candidate i 's utility when she wins and obtains v votes, where x is the winning platform. Similarly, $u_i(l, v, x)$ is her utility when she loses. We then represent the standard assumptions about candidates' payoffs as follows. First, politicians prefer winning to losing: $u_i(w, v, x) > u_i(l, v, x)$ for any given $v \in [0, 1]$ and $x \in X$. Second, they prefer more votes to less, *ceteris paribus*: $u_i(w, v, x) > u_i(w, v', x) \Leftrightarrow v > v'$, and similarly for $u_i(l, \cdot, \cdot)$. Third, they prefer policies closer to their bliss points, all else equal: $u_i(w, v, x) > u_i(w, v, x') \Leftrightarrow d(x, x_i^*) < d(x', x_i^*)$, and similarly for $u_i(l, \cdot, \cdot)$, where $d(x, x_i^*)$ is the distance between policy x and i 's bliss point x_i^* .

Given multiple goals, one cannot assume that winning is always satisfying or that losing is dissatisfying. (E.g., a winner with high aspirations who espoused a platform far from his ideal point may find the victory bitter.) Hence, we must make aspirations explicit, which entails replacing (A1) and (A3) by their more fundamental counterparts (A1') and (A3'), below. These more basic assumptions are defined by a comparison of utility to aspirations, instead of (A1)'s and (A3)'s dependence on specific events (i.e., on winning and losing). (A1') is thus a more general definition of satisficing; (A3'), a more general definition of search. In both, $a_{i,t}$ denotes candidate i 's aspiration level at date t and $u_{i,t}$, his utility.

(A1') (satisficing): If $u_{i,t} \geq a_{i,t}$ then in $t + 1$ candidate i espouses the same platform she

²³The finiteness of X plays no essential role in proposition 4.

used in t .

(A3') (search): There is an $\epsilon_i > 0$ such that for all t and every history leading up to t , if $u_{i,t} < a_{i,t}$ then in $t + 1$ candidate i espouses a platform that differs from her platform in t with a probability of at least ϵ_i .

Since aspirations are now explicit we must stipulate how they are formed.

(A4:) There is an $\epsilon \in (0, 1)$ such that i 's aspiration level adjusts by a rule that satisfies the following conditions with probability one, for all t and all histories leading up to t :

- **(i)** If $u_{i,t} > a_{i,t}$ then $a_{i,t} + \epsilon(u_{i,t} - a_{i,t}) \leq a_{i,t+1} < u_{i,t}$.
- **(ii)** If $u_{i,t} = a_{i,t}$ then $a_{i,t+1} = a_{i,t}$.
- **(iii)** If $u_{i,t} < a_{i,t}$ then $u_{i,t} < a_{i,t+1} \leq a_{i,t} - \epsilon(a_{i,t} - \pi_{i,t})$.

(A4) is quite general: it specifies only the direction of aspiration-adjustment and some weak restrictions on speed—in particular, adjustment cannot become arbitrarily sluggish. We then obtain the following result, which shows that Proposition 4's divergence conclusion is robust with respect to assumptions about candidates' goals.

Proposition 5 presumes that candidates have 'qualitatively similar' motives, which means that they pursue the same goals: e.g., either both care about how many votes they get or neither does. Their preference intensities, however, may differ sharply. (Indeed, their utilities may be described by different functional forms.)

Proposition 5: Suppose (A1'), (A3') and (A4) hold. Candidates' motives are qualitatively similar. Then convergent outcomes are absorbing if and only if candidates care only about policy and not at all about winning or vote-share.

Proposition 5 implies that if candidates care about all three goals (winning, votes and policy) then the classical median voter outcome is unstable. Further, it is stable only if politicians put *zero* weight on winning—a knife-edged possibility that seems far-fetched.

This result may appear vulnerable to the following criticism: the parties might eventually become sophisticated enough to interpret a close loss (after, e.g., they've adopted similar or identical platforms) as resulting from chance factors that have nothing to do with their chosen policy. They may therefore infer that their platform-choice was a good one. Thus, because the losers policy is not really 'broken', the party doesn't see the need to 'fix' it, whence convergence may be an absorbing outcome. However, this argument is not compelling.

When parties pick the same platform voters will be indifferent between them, and the election outcome will be determined by stochastic factors such as arbitrary indifference-breaking rules, random voter turnout or noisy vote-counts. These will produce randomness in observed vote shares: the loser will end up with a vote share of less than half. Even a sophisticated party will find it hard to disentangle the impact of these stochastic variables from a genuine failure of its platform to appeal to a majority of voters. Because the temptation to second-guess its electoral strategy might be overpowering, the losing party is likely to engage in continued policy experimentation to improve its perceived chances of success in the next election.

Note that KMP's model and ours yield conflicting predictions when a median voter exists. Because KMP's challenger is trying (albeit myopically) to maximize his vote-share, in the standard Downsian environment both parties end up espousing the median voter's bliss point, and so eventually take up identical platforms. Once that happens, the loser—however determined (chads in Florida? the Supreme Court?)—*will never move*. In our model losing must eventually be dissatisfying even if the loser maximized his vote share; hence, the tweedledum-tweedledee pattern is unstable. Thus, just as with rational choice models, different behavioral models can generate distinct empirical claims.

Thus we have here an intriguing situation: two behavioral theories produce conflicting predictions, but one of the behavioral theories (KMP) agrees with the main rational choice theory (Hotelling-Downs). This complicates the evaluation of the competing programs: if Stokes' summary (1999) of the empirical literature—the two main parties rarely adopt the same platform—stands up, then demerits must be given to *both* research programs. Similarly, brownie points should be given to those theories in *both* programs (e.g., Wittman- or Calvert-type models in the rational choice program and the present aspiration-based theory in the bounded rationality framework) that make the more accurate prediction of nonconvergence.

INFORMED AND SOPHISTICATED CHALLENGERS

One can extend this model by making the challenger more informed and/or more sophisticated, hence sharpening his search for winning platforms. Regarding information, the extreme case would be to assume that the challenger knows the majority preference relation for each pair of platforms. Then, given varying degrees of strategic sophistication, the challenger might choose platforms that vote-maximize against the incumbent (Kramer 1977) or those that are in X 's uncovered set (Miller 1980). Or one might posit an intermediate degree of information, e.g., the challenger has some chance of knowing what vote-maximizes against the status quo, and combine that with a degree of sophistication. Obviously, many extensions are possible. We can examine only a few prominent ones. Some are taken directly from the literature; others involve modifications.

Kramer's 1977 model

Recall that incumbents in Kramer's model must retain the platform that won them office, while challengers choose policies that vote-maximize against the status quo. His main result, theorem 1, says that the trajectory of winning platforms converges to the minmax set. It need not stay there, but "Theorem 1 does ensure that a trajectory which jumps outside must immediately return toward the minmax set..." (1977, p.324). In contrast, our propositions 1 and 3 state that the trajectory is absorbed into the top cycle.

His model and ours yield different implications partly because they make different assumptions about what challengers know and/or can implement. Our challenger knows what he has tried in the past and has some chance of experimenting; Kramer's knows what vote-maximizes against the incumbent's policy. One might object to the latter because it gives challengers an unrealistic amount of information: to know *exactly* what policy vote-maximizes against the status quo is asking a lot of any decision maker or advisor. And presuming that the out-party is so well-organized, so immune to internal squabbles, that it can always *choose* a vote-maximizing option can also be questioned. What, then, happens if challengers try to vote-maximize but occasionally "tremble" and mistakenly pick a platform that is not vote-maximizing? (For simplicity we make the standard assumption that following a tremble the challenger plays a fixed and totally mixed strategy: anything in X can be picked with positive probability.) The following result, which follows from proposition 3, gives the answer.

Remark 2: Suppose (A1) holds. With probability $1 - \epsilon$ the challenger selects a platform that vote-maximizes against the incumbent's policy; with probability $\epsilon > 0$ he trembles and plays a strategy that is totally mixed over X . Then the trajectory of winning policies converges to and is absorbed by L_1 with probability one.

Hence, if the challenger can err in his effort to vote-maximize against the incumbent then we recover the conclusion of Proposition 3, even if the chance of error is arbitrarily small. Thus, though the challenger is *trying* to vote-maximize against the incumbent and usually does so, the process is led toward the top cycle, *not* to the minmax set (unless the two coincide). Why?

The reason is that the challenger's errors are filtered by the electorate, making the trajectory of winning policies drift toward higher plateaus in the short run and the top cycle in the long run. Proposition 1 tells us that given this electoral filtering, no mistake by the out-party can shove the dynamic down to lower levels. Further, since any kind of mistake is possible, the challenger has *some* chance of stumbling onto policies on higher levels. The

voters approve of such mistakes, thus pushing the trajectory uphill—whether or not the minmax set and the top cycle coincide.²⁴

Thus, regardless of the challenger’s *intentions*, the electoral environment ensures that the dynamic is driven by *winning* per se rather than by the magnitude of victory. Consistent with Satz and Ferejohn’s argument that “[when] we are...interested in explaining...the general regularities that govern the behavior of all agents...it is not the agents’ psychologies that primarily explain their behavior, but the environmental constraints they face” (1994, p.74), the selection environment trumps the agent’s intentions.

This trumping holds quite generally; the Kramerian challenger’s specific objective—to maximize votes against the status quo—was inessential in remark 2. As long as the challenger has some chance of trembling and playing a strategy that is totally mixed over X , the conclusions of remark 2 hold, *regardless* of the challenger’s objectives. Thus the parties could have different goals. For example, when the Democrats challenge they could vote-maximize against the incumbent, but when the Republicans are the out-party they select a policy that maximizes some ideological criterion (as in, e.g., Chappell and Keech 1986, p.884). Or both parties could pursue a mix of office-seeking and ideology, as in Wittman- or Calvert-type models. In the long run those goals do not matter. One could even allow for parties that suffer from Arrowian problems and so lack coherent preferences. All that counts is that the pattern of error gives the electoral mill enough grist to work on.²⁵ So long as this condition is met, the challenger could be as sophisticated and informed as one likes, with any kind of preferences; the end result is the same.

Miller (et al.) and the uncovered set

Miller (1980, p.93) and others (McKelvey 1986; Cox 1987; Epstein 1998) have argued that if x covers y then x electorally dominates y . As Cox put it,

If one accepts the extremely mild assumption that candidates will not adopt a spatial strategy y if there is another available strategy x which is at least as good as y against any strategy the opponent might take and is better against some

²⁴This victory of the top cycle is hollow, strictly speaking, when it is the entire policy space. But proposition 6 will show that if the top cycle is “almost” a strict subset of X then I_t will spend “most” of its time in a strict subset of the policy space. Because this subset of X and the minmax set can be disjoint (for an example see Bendor, Mookherjee and Ray 2004), the thrust of Remark 2 can hold even when the top cycle is everything.

²⁵Thus, one can regard this as an evolutionary theory: “blind” variation is produced by error; selection is the electoral environment. We thank John Padgett for this interpretation.

of the opponent’s possible strategies, then one can conclude that candidates will confine themselves to strategies in the uncovered set (1987, p.420).

If we are to use the uncovered set as a solution concept, we must assume that candidates are both well-informed and relatively sophisticated. But expecting candidates to *invariably* pick policies in the uncovered set may be unrealistic.²⁶ Yet even a bit of information can help the challenger search, as the next result shows. (The proof is in the appendix.)

Remark 3: If (A1) is satisfied and for every history the challenger alights on X ’s uncovered set with a probability of at least $\epsilon > 0$, then the following hold.

- (i) I_t is in L_1 with positive probability for all $t > 1$.
- (ii) If $\Pr(I_t \in L_1)$ is less than one then $\Pr(I_1 \in L_1) < \dots < \Pr(I_t \in L_1)$.
- (iii) $I_t \rightarrow L_1$ with probability one as $t \rightarrow \infty$.

Thus even fragmentary information about the uncovered set’s location and even crude understanding about the strategic value of uncovered policies can have substantial impacts, in both the short-run (parts (i) and (ii)) and the long (part (iii)).

Now consider a less demanding possibility: the challenger might not know all of the uncovered set but may know policies that cover *what he must try to beat today*—the incumbent’s platform. (This presumes that *some* alternative covers I_t . If not, then I_t is in X ’s uncovered set and so is already in L_1 .)

First we establish the importance of the challenger finding something that covers the incumbent’s platform. It is necessary: electoral hill-climbing cannot occur without it.

Remark 4: Suppose (A1) holds. Consider any $r = 1, \dots, z$. If I_t is in L_r and the probability that C_t covers I_t is zero, then I_{t+1} must also be in L_r .

The logic is straightforward. Any policy at higher levels, say any $x_i \in L_1 \cup \dots \cup L_r$, covers any policy at lower ones, i.e., any $x \in L_{r+1} \cup \dots \cup L_z$. Hence if $I_t \in L_r$ and today’s challenger has no chance of finding a platform that covers the incumbent’s, then he has no chance of finding anything in $L_1 \cup \dots \cup L_{r-1}$, since anything there would in fact cover I_t . Hence hillclimbing cannot occur. Since proposition 1 ensures that the process cannot slip downhill, it must stay at the same plateau.

²⁶If they did then the process would jump to L_1 in one period, since X ’s uncovered set is a subset of L_1 .

Hence, now consider elections in which the challenger *does* have some chance of finding platforms that cover the incumbent's. The following assumption formalizes this idea.

(A5): For every history and in any election in which I_t is covered by some $x \in X$, with probability of at least $\epsilon > 0$ the challenger finds an option that covers the status quo.

(A5) neither implies nor is implied by (A2), which stipulates the possibility of experimentation.²⁷ But as the next result shows, their long-run effect is the same. (The proof is in the appendix.)

Remark 5: If (A1) and (A5) hold, then the trajectory of winning policies converges to and is absorbed by L_1 with probability one.

Although the challenger is sophisticated enough to have a chance of finding a platform that covers the incumbent's, the process is *not* guaranteed to be absorbed into X 's uncovered set (unless that set and L_1 coincide), though it will visit that set infinitely often. The reason: L_1 may have policies that are not in the uncovered set, and because any two policies in the same level are connected by a cycle, the trajectory of winning policies can leave the uncovered set.

FFMP (1980, 1984)

FFMP assume that new platforms come from a uniform distribution over the status quo's win set. This is an intermediate degree of information and sophistication: less demanding than assuming that new options must cover the status quo but more demanding than assuming experimentation. However, FFMP posited a *uniform* distribution for computational reasons: they (1984) calculated bounds on the limiting distribution of winning platforms. This is unnecessary for qualitative results and we shall disregard it. For our purposes, the key part of FFMP's premise is that the challenger's search puts positive probability on anything that beats the incumbent's platform. As usual, this assumption need not be forced into a Markovian mold.

(A6): Following every history, the challenger's search has a probability of at least $\epsilon > 0$ of finding any option that beats the status quo.

Because the set of policies that beat any x must include some in L_1 , (A6) implies stochastic hill-climbing in the short-run and convergence to L_1 in the long-run, just as remark 3

²⁷(A2) does not imply (A5) because the challenger might experiment but his set of possible new policies may not include anything that covers the status quo. For an example that shows why (A5) does not imply (A2), see Bendor, Mookherjee and Ray (2004).

did. (The proof is similar to remark 3's and so is omitted.)

Remark 6: If (A1) and (A6) are satisfied then the following hold.

- (i) I_t is in L_1 with positive probability for all $t > 1$.
- (ii) If $\Pr(I_t \in L_1)$ is less than one then $\Pr(I_1 \in L_1) < \dots < \Pr(I_t \in L_1)$.
- (iii) $I_t \rightarrow L_1$ with probability one as $t \rightarrow \infty$.

Thus if challengers are as informed and sophisticated as FFMP posit, then long-term convergence to the top level is ensured, as is short-term progress.

In general, this section's findings show that endowing the challenger with more information and/or more strategic sophistication has a quantitative effect—convergence to L_1 is sped up—but does not affect the model's qualitative conclusions.

ROBUSTNESS ISSUES

The price for analytical results is stylized assumptions. This means reshaping vague-but-plausible ideas (e.g., incumbents are often content with the platforms that won them office) into crisper but less plausible ones (incumbents *always* satisfy). To be sure, to theorize one must simplify: as Jonathan Swift observed long ago, the most realistic model of a phenomenon is the phenomenon itself. But it would be troubling if our results turned out to be knife-edge findings—if changing a premise a little altered the conclusions a lot.

Several features of our model might cause concern in this regard. As noted, assuming that incumbents invariably keep winning platforms exaggerates the plausible scenario it is meant to capture. Similarly, that x is majority-preferred to y may not *guarantee* that x will beat y : variations in, e.g., turnout or voters' errors may change the outcome.

A more subtle concern, conceptually more serious than the above simplifications, is that our model seems to predict little in “ill-structured” environments where the Plott conditions fail. In such situations the top cycle may be the entire policy space. (It is well-known that an n -dimensional spatial voting setting is especially vulnerable to this problem.) But then Proposition 3 has no bite, and our model apparently loses all predictive power.

Yet this conclusion is too hasty: much depends on how badly the Plott conditions are violated. Take the most drastic perturbation: a y from the lowest level now beats an x from L_1 , collapsing all levels into one big set. Now the top cycle is X , so technically our convergence results are vacuous. Yet this perturbation's substantive bite could be minimal: it will matter only when the incumbent's platform is that particular x *and only if the challenger*

can find that particular y . Challengers might face a needle in the haystack problem: in the perturbed electoral environment they *might not discover y very often*, if X is large and search is crude.²⁸ Then the probability that the sequence of winning policies will go from x to y will also be low, so the process will spend most of its time at the top plateau of policies, even though the top cycle is the entire policy space in this perturbed environment.

Below we state a robustness proposition that handles all these (and possibly other perturbations) in the same general setup. The generality exacts a price—additional abstraction—but we will link the abstract structure to the substantive robustness questions raised here.

Consider a *family* of models, indexed by a parameter $\theta \in [0, 1]$. Each model θ has a policy set $X(\theta)$, partitioned into level sets $L_1(\theta), L_2(\theta), \dots, L_z(\theta)$, where z may also depend on θ . In each model θ a typical state is defined as before: as a pair (I, C) , called σ for convenience. Abusing terminology slightly, we say that the state $\sigma = (I, C)$ lies in some L if the incumbent’s platform $I \in L$. A state $\sigma = (I, C)$ is *in the top set* if $I \in L_1(\theta)$.

For each $t \geq 1$, let h_t be the t -history at date t : all states that transpired up to t , including the current state at t . Let $\sigma(h_t)$ denote the state at date t under the t -history h_t . For each t and each t -history, let $\pi^\theta(h_t, i)$ be the probability that the system enters $L_i(\theta)$ at the next date $(t + 1)$, starting from h_t . This is information that a particular model would give us. For instance, Proposition 1 says that given our stylized assumptions, if the state at history h_t lies in some L_j then this transition probability $\pi(h_t, i)$ is zero whenever $i > j$; the process doesn’t move from electorally higher plateaus to lower ones. Now we want to allow for these “perverse drifts”, but with low probability.

(A7): There exists a function $\psi(\theta)$, with $\psi(\theta) \rightarrow 0$ as $\theta \rightarrow 0$, such that whenever $\sigma(h_t)$ lies in $L_j(\theta)$ for some $j < i$, $\pi^\theta(h_t, i)$ may be positive but is less than $\psi(\theta)$.

This assumption acknowledges that movements to lower plateaus can occur in our family of models. This might happen if, e.g., voters sometimes punched the ballot incorrectly.

But the most subtle interpretation of (A7) is that preference cycles destroy the L ’s multi-layered structure. Under this interpretation, $L_1(\theta), L_2(\theta), \dots, L_z(\theta)$ can no longer be viewed as the *true* electoral plateaus, but as some underlying counterfactual plateau-structure were the problematic cycles artificially removed. Here the possible drift from higher to lower levels is interpreted not as a failure of the model’s assumptions but as a genuine majority-based defeat of an x in one level by a y in another. Thus, under this last interpretation the true ill-structured environment has only a few levels (perhaps just one), but our formulation strips

²⁸E.g., if search is completely blind then the chance that the challenger will land on y is $\frac{1}{m}$.

away the ill-structure by placing the burden on “wrong” movements in the state under a well-structured environment.

(A7) states that for θ close to zero, the ill-structure “almost” vanishes (or equivalently, under the other interpretations, that the model’s other assumptions “almost” hold). This is captured by the bound function $\psi(\theta)$ on “wrong” movements, which goes to zero as $\theta \rightarrow 0$.

A second assumption guarantees that the usual upward movements, such as those guaranteed by (A2), continue to exist throughout:

(A8): The infimum value of $\pi^\theta(h_t, i)$, over every θ , every t -history, and every i is strictly positive, provided $\sigma(h_t) \in L_j(\theta)$ for some $j > i$.

Although we need to state (A8) formally for this family of models, we have already seen that it is an implication of our other assumptions (e.g., experimentation). Since no further comment about (A8) is required, we now state our robustness result.

Proposition 6: Assume (A7) and (A8).

- **(i)** For every $\epsilon > 0$, there exist integers $T(\epsilon)$ and $\theta(\epsilon)$ such that for every $T \geq T(\epsilon)$ and $\theta \leq \theta(\epsilon)$, and every initial state, the system lies in the top set at date T with probability of at least $1 - \epsilon$.
- **(ii)** Suppose that the system enters the top set at date t . Pick any date $s \geq t$. Then, if $\theta \geq \theta(\epsilon)$, the system continues to lie in the top set at date s with (conditional) probability of at least $1 - \epsilon$.

Part (i) says that if the perturbations (represented by θ) are small enough, then after “sufficiently” many periods the system must be in the top set $L_1(\theta)$ with very high probability. Hence, if the perturbations are due to ill-structured majoritarian preferences but the “destructive” cycles are sparse enough, then by artificially stripping away these cycles one can significantly boost the framework’s predictive power. One interpretation: as θ becomes small the number of feasible policies grows without bound, while the policies causing the ill-structure increase more slowly. If the challenger’s search is crude, the chance of alighting on the latter becomes very small; then (A7) holds. Proposition 6 follows.²⁹

²⁹Proposition 6 implies that our conclusion about Kramer’s trajectory is robust. Suppose that the minmax set and the “top set” are disjoint, and the top cycle is everything but the Plott conditions nearly hold. Thus, because the minmax set is outside L1 and because proposition 6 implies that the dynamic must eventually live mostly in L1, Kramer’s conclusion is fragile even when the top cycle is everything.

Part (ii) states that for any sample path that enters the top set, the *conditional* probability of staying there is also high. Informally, the system enters the top set *and stays there* with high probability. This rules out certain perverse dynamics (e.g., entry into the top set being positively correlated with swift departure from that set).

Note that this result used *no* Markovian assumptions. This is striking: it shows that the model’s robustness does not rest on any special stochastic features.

Ill-structured majoritarian preferences: a computational model

Proposition 6 studies only “small” changes in the model’s key assumptions; it does not tell us what happens when there are big changes in, e.g., preference-profiles. In particular, it does not say what happens to the trajectory of winning platforms when the preference-profile is far from having a generalized median. This is a difficult question; analytical results are hard to come by. Hence we resort to a computational model, which we now briefly describe. (For a description of the computer program, see [https://faculty-gsb.stanford.edu/bendor/.](https://faculty-gsb.stanford.edu/bendor/)) Because the simulation is a special case of our mathematical model—the policy space is finite, (A1) and (A2) hold, etc.—we focus on its distinctive properties: policies are set in a two-dimensional space and voters have quadratic loss functions.³⁰

To ensure that the simulation results are meaningful and interpretable we make the search Markovian. Hence, the probability distribution of (I_{t+1}, C_{t+1}) depends only the parties’ current platforms and the transition rules created by the incumbent’s satisficing and the challenger’s search. We stipulate time-homogeneous search rules, so the (I_t, C_t) process is stationary. Hence we can invoke powerful theorems for finite-state stationary Markov chains (e.g., Kemeny and Snell 1960) which tell us when such processes are ergodic. All of our computational results arise from ergodic processes.³¹ Thus we will be scrutinizing the steady state distributions of the winning platforms. (More precisely, the output—for every set of parametric values, 1,000 sample paths run for 1,000 periods—will closely approximate such steady states.)

Results

We examine two types of results: (1) how different preference-profiles affect the limiting distribution of winning platforms and (2) how different search rules affect this distribution.

(1) To measure a profile’s symmetry, we use a standard metric: the size of the uncovered

³⁰The following ensures that citizens have *strict* preference orderings over policies (as the analytical model requires): if policies x and y are equidistant from voter i ’s ideal point then they are randomly and independently given different “valence” (nonspatial) values.

³¹Establishing that these processes are ergodic is straightforward, so the proofs are omitted.

set. (At one extreme, if the uncovered set is a singleton then a generalized median exists; at the other, it is the entire policy space. So the measure ranges from $\frac{1}{m}$ to 1.) The challenger’s search rule is represented by a probability distribution over the policy space; here the distribution is single-peaked (a truncated normal). Thus in t the challenger is more likely to choose a platform close to the one he espoused in $t - 1$ than something far away.

The size of the uncovered set and the distribution of winning platforms are strongly related (figure 3). This reflects how the strength of centripetal forces varies across electoral environments. When the uncovered set is small these centripetal forces are strong, so in the steady state winning platforms are centrally located; when this set is big the centripetal forces are weak, so winning platforms are scattered throughout the policy space.

[figure 3 about here]

This pattern complements proposition 6, which showed that the process is well-behaved for small perturbations to voter profiles. The computational results of figure 3 suggest that the process is well-behaved *globally*: the dispersion of winning platforms increases steadily as the uncovered set expands. But the electoral environment can mold the steady state distribution of winning platforms only if the out-party’s search yields enough variety for the voters’ selective forces to work on. So we now turn to the effect of different search rules.

(2) The output in the first part of table 1 is based on naive search: challengers search blindly, putting probability of $\frac{1}{m}$ on every platform. Yet the size of the uncovered set and the long-run dispersion of winning platforms remain highly correlated: the main pattern—centrist platforms tend to win when the uncovered set is small—continues to hold.

[table 1 about here]

Hence, this pattern does *not* require search to be prospectively attuned to winning. Instead, what suffices is that challengers generate enough grist (variety) for the electorate’s mill.

Table 1’s second part shows that the relation between the uncovered set’s size and the dispersion of winning platforms is reduced if the out-party’s search is keyed to its policy preferences. This makes sense: that search creates a centrifugal force—the challenger being tugged back toward his ideal policy—that is *independent of the size of the uncovered set*.

Yet, although the pattern is weakened it is still present, even in this extreme case when the challenging party is completely dominated by ideologues. (We have also studied search that is a mix of the above pure types: search centers on a policy that is a weighted average of the party’s ideal point and its last platform. Results [unreported here] show that the system’s long run tendencies are intermediate between those of the two pure search rules whose outcomes are reported by table 1.) Even ideological parties cannot completely ignore

the strong electoral forces that are present when the uncovered set is small.

The shaping power of the electoral environment: analytics once more

Our computational results support the claim that the system is well-behaved even when the voters’ preference profile is far from having a generalized median. But computational models must use specific assumptions—here, a two-dimensional policy space, quadratic utility and a small electorate—so we now supplement these findings with an analytical one.

Since we computed when we couldn’t derive general results analytically, we must simplify our analytical model *somehow*. But this should be consistent with our substantive objective: to examine the electoral environment’s centripetal forces. Therefore we should *not* constrain voters’ preferences. Instead, we make our mathematical model tractable by simplifying the challenger’s search: we assume that it is blind—uniform over X . This not only helps to ensure tractability; it is also substantively useful: since challengers don’t learn, we know that conclusions about the steady-state probabilities of winning platforms depend only on the selective forces inherent in the electorate (and on satisficing-by-incumbents).

Our last result shows that even if the top cycle is the entire policy space and nothing is assumed about the Plott conditions, the electoral environment can still impart some stochastic order to outcomes. (The proof is in the appendix.)

Proposition 7: Assume (A1) and blind search. If $L_1 = X$ and x covers y , then the steady-state probability that x is the incumbent’s platform exceeds y ’s steady-state probability.

Thus, if the electoral environment has long strings of covering relations (e.g., a covers b which covers c which...), then the steady state probabilities of winning policies will be structured by monotonicity (e.g., a is more likely than b which is more likely than c which...).³² Proposition 7 also implies that if a platform’s long-run probability of being the government’s policy is maximal, then *it must be in the uncovered set*. These properties buttress the claims of Miller et al. that if x covers y then x electorally dominates y .

CONCLUSIONS

Our results emphasize how powerfully certain electoral landscapes shape the behavior of two competing, boundedly rational candidates. In highly structured electoral environments—

³²Proposition 7 *cannot* be generalized by replacing “ x covers y ” with “ x beats more policies than y does”. Though the resulting conjecture—platforms that beat more rivals should be more likely to be the government’s policy in the steady state—is intuitively plausible, it is not true in general. This is so even if one rules out spatially bizarre “preferred-to” relations, e.g., x beats y even though the former loses to thousands of other platforms while the latter loses only to a handful. (We will provide, upon request, a non-bizarre counterexample to the conjecture.)

those with many levels—electoral competition constrains the process greatly, even if challengers are ignorant and unsophisticated. Thus our paper is consistent with work in economics on zero-intelligence agents (e.g., Gode and Sunder 1993) which analyzes how much of market performance is due to the market environment rather than the intelligence of agents. Gode and Sunder concluded, “Adam Smith’s invisible hand may be more powerful than some may have thought: when embodied in market mechanisms such as a double auction, it may generate aggregate rationality not only from individual rationality but also from individual irrationality” (p.136). When the median voter exists his/her hand is similarly powerful, in guiding the trajectory of winning platforms.³³

We also investigated the effects of endowing challengers with more information and/or sophistication. The results show that their effect is quantitative, not qualitative: they tend to speed up hill-climbing, making I_t converge faster to the top level.

A concern about our model is that slight perturbations of voters’ preferences can create an ill-structured electoral environment and make many of our results vacuous. But proposition 7 shows that small perturbations have only a minor effect on winning platforms in the steady-state. And our computational results and proposition 6 indicate that the process is well-behaved even when the preference profile is far from having a Condorcet winner.

This is part of a larger project on behavioral models of elections. The study of elections encompasses more than party competition; the choices of voters, especially turnout and vote-choice, are obviously vital. (For an behavioral model of turnout, see Bendor, Diermeier and Ting 2003.) We hope that models of bounded rationality will compete with rational choice models across all major electoral topics. When they do, the two research programs will be in a real horse race.

³³The parallel with models of zero-intelligence agents is incomplete: in a pure zero-intelligence model *both* candidates would choose platforms blindly, whereas our winners satisfy, which is fairly sensible behavior. We have not pursued that limiting case here. But even a cursory examination of the pure zero-intelligence model would reveal that the electoral environment strongly shapes the trajectory of winning policies. (At the opposite extreme of zero-intelligence agents—fully rational and completely informed candidates—we recover the McKelvey-Schofield-Cohen world where even very extreme policies can be electorally viable. Hence, if one’s normative democratic theory implies that extreme policies are bad, then one might conclude that a polity is better off with politicians who aren’t perfectly rational or fully informed. [We thank an anonymous referee for this point.])

Appendix

Proof of Proposition 3

If $I_1 \in L_1$ with certainty then the result is immediate (given proposition 1), so assume that $\Pr(I_1 \in L_1) < 1$. Define $U_t(i)$ as the number of platforms untried by party i ($i = D, R$) at the start of period t . $U_t(i)$ is a weakly decreasing process, with $U_0(i) = m$. Note that if $\min(U_t(D), U_t(R)) = 0$ then $I_t \in L_1$. (This holds because if one of the parties has tried all platforms then it must have tried those in L_1 , and once something in L_1 was tried, proposition 1 implies that the process never leaves L_1 thereafter.) So we need to show only that the probability of the event that $\min(U_t(D), U_t(R)) > 0$ goes to zero as $t \rightarrow \infty$.

Since there is a challenger in every period, at least one party (say D) must be a challenger infinitely often. Suppose, for date t , $U_t(D) = n$, where $0 < n$. Consider the event that $U_t(D)$ stays at n . Suppose D 's experimentation probability were exactly ϵ , whenever D challenges and $U_t(D) > 0$. Then the counting process $U_t(D)$ is just a geometric process with probability ϵ , whence $\lim_{t \rightarrow \infty} [\Pr(U_t(D) = n)] = 0$. Since D experiments with a probability of at least ϵ , this must continue to hold. Hence with probability one $U_t(D)$ will hit $n - 1$ eventually. If $n - 1 = 0$ we are done; if not, just repeat the above argument. So with probability one $U_t(D) = 0$ eventually, implying that the probability of $\min(U_t(D), U_t(R)) > 0$ goes to zero as $t \rightarrow \infty$. QED.

Proof of Proposition 5

Sufficiency. Suppose that the candidates care only about policy. Then in any the convergent outcome, say of (x, x) , candidate i gets $u_i(x)$ no matter who wins. Thus, if $a_{i,t} = u_i(x)$ both sides will be happy and by (A1') both will again espouse policy x in $t + 1$. Further, by (A4) both will continue to have aspirations equal to their respective payoffs. Hence, actions and aspirations are self-replicating, i.e., the state is absorbing. QED.

Necessity. This is by contradiction: assume that the politicians want either to win or to get more votes yet some (x, x) is absorbing. Here we take up the case where the politicians want to win but don't care about votes; the proofs for the other two cases (i.e., (1) they want more votes but don't care about winning or (2) they care about both) are similar.

Since the parties adopt the same platform, each voter will with positive probability vote for either party. Hence two outcomes arise with positive probability for each side: (w, x) and (l, x) , where $u_i(w, x) > u_i(l, x)$ for $i = D, R$. Without loss of generality, normalize $u_i(l, x)$ to zero; let $\bar{u}_D \equiv u_D(w, x)$ and $\bar{u}_R \equiv u_R(w, x)$.

The logic of the proof is to show that at any convergent outcome, at least one side's

aspirations must rise high enough so that disappointment and, therefore, search for new platforms are inevitable. First we show aspirations will reach such levels.

If (x, x) is absorbing and the process enters that state at a date t then it never leaves it. So assume that the state at t is (x, x) . It is convenient to analyze by three (disjoint and collectively exhaustive) cases.

Case 1: $\min(a_{D,t}, a_{R,t}) \geq 0$. By (A4), the new aspiration of the winner of the t^{th} election moves toward \bar{u} ; since it was already at least 0, it must be strictly positive in $t + 1$. The loser in t gets 0, so his aspiration in $t + 1$ must continue to be at least zero. So by induction, $\max(a_{D,t'}, a_{R,t'}) > 0$ and $\min(a_{D,t'}, a_{R,t'}) \geq 0$ for all $t' > t$.

Case 2: $\min(a_{D,t}, a_{R,t}) < 0 \leq \max(a_{D,t}, a_{R,t})$. For convenience and w.l.o.g., assume that $\max(a_{D,t}, a_{R,t}) = a_{D,t}$. (A4) and this case's premise that $a_{D,t} \geq 0$ together imply that $a_{D,t'} \geq 0$ for all $t' > t$. Regarding $a_{R,t}$, define a_R^* such that it solves the equation $a_R^* + \epsilon(\bar{u}_R - a_R^*)$, where ϵ is the parameter defined in (A4). Hence, if ever $a_{R,t} > a_R^*$ and R wins the election in t , then $a_{R,t+1} > 0$.

But we know that for any finite $a_{R,t} < a_R^*$, the “no arbitrary sluggishness” property of (A4) ensures that there exists a finite positive integer s such that even if R loses s elections consecutively, starting in t , then $a_{R,t+s} \geq a_R^*$ with probability one. Further, since (A4) implies that the sequence of $a_{R,t}, a_{R,t+1}, \dots$, must move monotonically toward 0, if $a_{R,t+s} < 0$ then $a_{R,t'} > a_{R,t+s}$ for all $t' > t + s$. Hence, by the construction of a_R^* , with a single victory at any time after $t + s$, R 's aspiration level will exceed 0. Since voters break their indifference (in the face of convergent platforms) via nondegenerate and stationary coins, the probability that R wins eventually goes to 1 as $t \rightarrow \infty$. Because D 's aspiration must have remained weakly positive, the system will eventually be in case 1 with probability one, whence that case's logic takes over.

Case 3: $\max(a_{D,t}, a_{R,t}) < 0$. Again fix $\max(a_{D,t}, a_{R,t}) = a_{D,t}$. By case 2, there are positive integers s_D and s_R such that if candidate i lost s_i consecutive elections, starting in t , then $a_{i,t+s_i} \geq a_i^*$ with probability one. Hence, the latter must hold once $\max(s_D, s_R)$ elections have occurred. In the next election one side must win, so by construction of the a_i^* the winner's aspiration must then exceed zero. Then case 2 applies, and we can proceed from there.

Together, cases 1-3 imply that for any pair of $a_{D,t}$ and $a_{R,t}$, within finitely many periods after t at least one of the candidates will have a strictly positive aspiration level with probability one. Under the hypothesis that (x, x) is absorbing, this will continue to be so forever.

But since voters break indifference (given convergent outcomes) with nondegenerate, stationary and independently tossed coins, each side can lose with a probability that's bounded away from zero uniformly in t , so the chance that the candidate with the higher aspiration level never loses goes to 0 as $t \rightarrow \infty$. Since the probability of search, given dissatisfaction, is also bounded away from zero uniformly in t and for all histories, the chance that (a) the loser is dissatisfied yet (b) doesn't search also $\downarrow 0$ as $t \rightarrow \infty$. So the probability that the process leaves (x, x) goes to one as $t \rightarrow \infty$. QED.

Proof of Remark 3

(i) Proposition 1 implies that if a sample path of the winning policies is in L_1 in t then it stays in L_1 thereafter. Hence if $\Pr(I_t \in L_1) > 0$ then $\Pr(I_s \in L_1) > 0, \forall s > t$, and in particular if $\Pr(I_1 \in L_1) > 0$ then $\Pr(I_s \in L_1) > 0, \forall s > 1$. Alternatively, suppose that $\Pr(I_t \in L_1) = 0$. Let $UC(X)$ denote X 's uncovered set. Since $UC(X) \subseteq L_1$ it follows immediately that $\Pr(I_2 \in L_1) > 0$, and we can then use the preceding argument.

(ii) If $\Pr(I_t \in L_1) < 1$ then by proposition 1, $\Pr(I_s \in L_1) < 1, \forall s < t$. Because $UC(X) \subseteq L_1$, it follows that $\Pr(I_{s+1} \in L_1 | I_s \notin L_1) \geq \Pr(I_{s+1} \in UC(X) | I_s \notin L_1)$. Since $\Pr(I_{s+1} \in UC(X) | I_s \notin L_1) = \Pr(C_s \in UC(X) | I_s \notin L_1)$, and the latter term is bounded away from zero, the result follows.

(iii) Let the bound on hitting $UC(X)$ be some $\epsilon > 0$. If the challenger's probability of hitting $UC(X)$ were exactly ϵ , then the process would have a geometrically distributed waiting time, whence the probability that it stays out of L_1 would go to 0 as $t \rightarrow \infty$. Since the process's probability hitting $UC(X)$ is at least ϵ , the result follows. QED.

Proof of Remark 5

Consider the event that the challenger fails to espouse an alternative that covers I_t . Since the probability of this event is at most $1 - \epsilon < 1$, the probability that it will recur infinitely often is zero. Hence with probability one the challenger will eventually find something that covers I_t . Because I_t is arbitrary here, this holds for every status quo platform. Hence, since the covering relation is transitive, the trajectory of winning policies must land in $UC(X)$ with probability one. Because $UC(X) \subseteq L_1$, the result is established. QED.

As the proof of proposition 6 is rather technical, it is omitted. (It is available upon request, and has been provided to the referees.)

Proof of Proposition 7

Given the proposition's hypotheses, it is easily established that the process is ergodic.

Hence the proposition's conclusion is a meaningful claim.

Given blind search, the steady state probability of policy i (denoted π_i) equals $\sum_{j \in l(i)} \pi_j + \pi_i(1 - |w(i)|)$, where $l(i)$ denotes the set of policies that lose to i and $w(i)$ is the set that wins against i . Thus

$$|w(i)| \cdot \pi_i = \frac{1}{m} \sum_{j \in l(i)} \pi_j.$$

As $\frac{1}{m}$ is a constant in this set of simultaneous equations, it drops out; hence

$$\pi_i = \frac{1}{|w(i)|} \sum_{j \in l(i)} \pi_j.$$

Because either $x \succ j$ or $j \succ x$ for all $i \neq j$, it follows that $|l(i)| + |w(i)| + 1 = m$. So

$$w(i) \subset w(j) \Rightarrow l(j) \subset l(i)$$

Hence, if x covers y then $l(y) \subset l(x)$. And since $\pi_x = \frac{1}{|w(x)|} \sum_{j \in l(x)} \pi_j$ whereas $\pi_y = \frac{1}{|w(y)|} \sum_{j \in l(y)} \pi_j$, the result follows by simple algebra. QED.

References

- Ansolabehere, Stephen, James Snyder and Charles Stewart. 2001. "Candidate Positioning in U.S. House Elections." *American Journal of Political Science* 45:136-59.
- Arthur, Brian. 1993. "On Designing Economic Agents that Look Like Human Agents." *Journal of Evolutionary Economics* 3:1-22.
- Austen-Smith, David, and Jeffrey Banks. 1999. *Positive Political Theory I: Collective Preference*. Ann Arbor: University of Michigan Press.
- Bendor, Jonathan, Daniel Diermeier and Michael Ting. 2003. "A Behavioral Model of Turnout." *American Political Science Review* 97:261-80.
- Bendor, Jonathan, Dilip Mookherjee and Debraj Ray. 2001. "Aspiration-Based Reinforcement Learning in Repeated Games." *International Game Theory Review* 3:159-74.
- Bendor, Jonathan, Dilip Mookherjee and Debraj Ray. 2004. "Satisficing and Selection in Electoral Competition." Typescript, Stanford University.
- Chappell, Henry, and William Keech. 1986. "Policy Motivation and Party Differences in a Dynamic Spatial Model of Party Competition." *American Political Science Review* 80:881-99.
- Cox, Gary. 1987. "The Uncovered Set and the Core." *American Journal of Political Science* 31:408-22.
- Cyert, Richard, and James March. 1963. *A Behavioral Theory of the Firm*. Englewood Cliffs NJ: Prentice-Hall.
- Dahl, Robert. 1961. "The Behavioral Approach in Political Science." *American Political Science Review* 55:763-72.
- Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper & Row.
- Epstein, David. 1998. "Uncovering Some Subtleties of the Uncovered Set: Social Choice Theory and Distributive Politics." *Social Choice and Welfare* 15:81-93.
- Ferejohn, John. 1995. "The Development of the Spatial Theory of Elections." In J. Farr and J. Dryzek (eds.), *Political Science and History*. Cambridge: Cambridge University Press.
- Ferejohn, John, Morris Fiorina and Edward Packel. 1980. "Nonequilibrium Solutions for Legislative Systems." *Behavioral Science* 25:140-48.
- Ferejohn, John, Richard McKelvey and Edward Packel. 1984. "Limiting Distributions for

- Continuous State Markov Voting Models.” *Social Choice and Welfare* 1:45-67.
- Fershtman, C. and Ariel Pakes. 2004. “Finite State Dynamic Games with Asymmetric Information: A Computational Framework.” Mimeo, Harvard Economics Department.
- Gode, Dhananjay, and Shyam Sunder. 1993. “Allocative Efficiency of Markets with Zero-Intelligence Traders.” *Journal of Political Economy* 101:119-37.
- Green, Donald and Ian Shapiro. 1994. *Pathologies of Rational Choice Theory*. New Haven: Yale University Press.
- Kelley, Stanley. 1965. “Forward.” In Anthony Downs, *Economic Theory of Democracy*. New York: Harper & Row.
- Kotovsky, Kenneth, John Hayes and Herbert Simon. 1985. “Why Are Some Problems Hard?” *Cognitive Psychology* 17:248-94.
- Kemeny, John, and J. Laurie Snell. 1960. *Finite Markov Chains*. Princeton: Van Nostrand.
- Kollman, Ken, John Miller and Scott Page. 1992. “Adaptive Parties in Spatial Elections.” *American Political Science Review* 86:929-37.
- Kollman, Ken, John Miller and Scott Page. 1998. “Political Parties and Electoral Landscapes.” *British Journal of Political Science* 28:139-58.
- Kramer, Gerald. 1977. “A Dynamical Model of Political Equilibrium.” *Journal of Economic Theory* 16:310-34.
- Kreps, David. 1990. *Game Theory and Economic Modelling*. Oxford: Clarendon Press.
- Laudan, Larry. 1977. *Progress and Its Problems*. Berkeley: University of California Press.
- Levitt, Steven. 1996. “How Do Senators Vote? Disentangling the Role of Voter Preferences, Party Affiliation, and Senator Ideology.” *American Economic Review* 86:425-41.
- McKelvey, Richard. 1986. “Covering, Dominance, and Institution-free Properties of Social Choice.” *American Journal of Political Science* 30:283-315.
- Miller, Nicholas. 1980. “A New Solution Set for Tournaments and Majority Voting.” *American Journal of Political Science* 24:68-96.
- Narendra, Kumpati and Mandaym Thathachar. 1989. *Learning Automata*. Englewood Cliffs, NJ: Prentice-Hall.
- Rubinstein, Ariel. 1998. *Modeling Bounded Rationality*. Cambridge MA: MIT Press.
- Satz, Debra, and John Ferejohn. 1994. “Rational Choice and Social Theory.” *Journal of*

Philosophy 91:71-87.

Schlesinger, Joseph. 1966. *Ambition and Politics*. Chicago: Rand McNally.

Simon, Herbert. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics* 69:99-118.

Simon, Herbert. 1964. "The Concept of Organizational Goal." *Administrative Sciences Quarterly* 9:1-22.

Stokes, Susan. 1999. "Political Parties and Democracy." In Nelson Polsby (ed.), *Annual Review of Political Science*, volume 2. Palo Alto: Annual Reviews Inc.

Table 1: Relationship Between Size of Uncovered Set and Dispersion of Winning Platforms

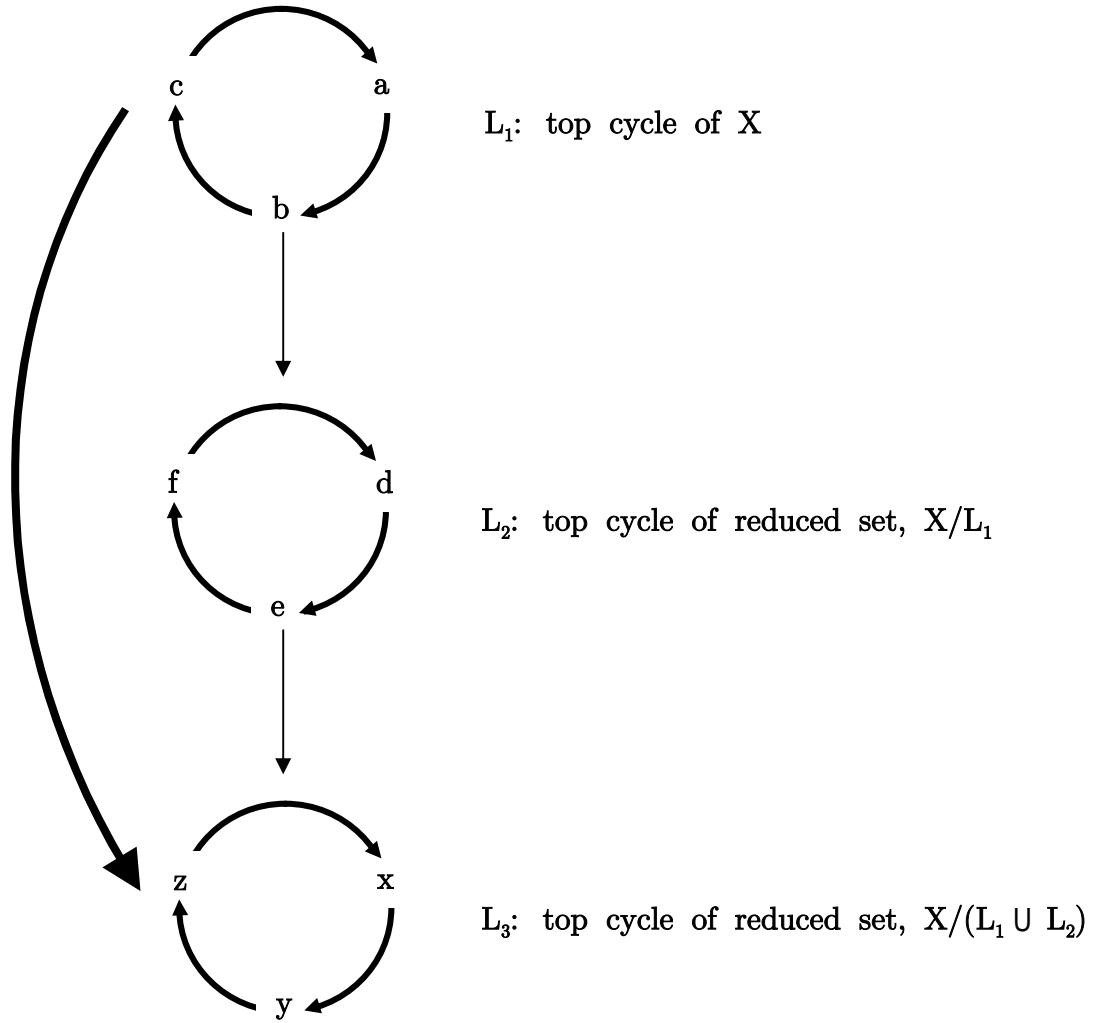
Search Rule	Mean Uncovered Set Ratio	Mean Variance of Winning Policies	β	R^2
Blind (uniform)	0.098 (0.063)	18.369 (18.575)	265.822 (14.92)**	0.82
Ideological	0.088 (0.056)	10.512 (10.495)	98.253 (4.25)**	0.27

Note: Standard deviations and absolute value of t-statistics in parentheses.

**significant at 1%

Uncovered Set Ratio is the ratio of the size of the uncovered set to the total policy space.

Figure 1:
An electoral landscape with three plateaus



Key:

a → b: a beats b.



Figure 2

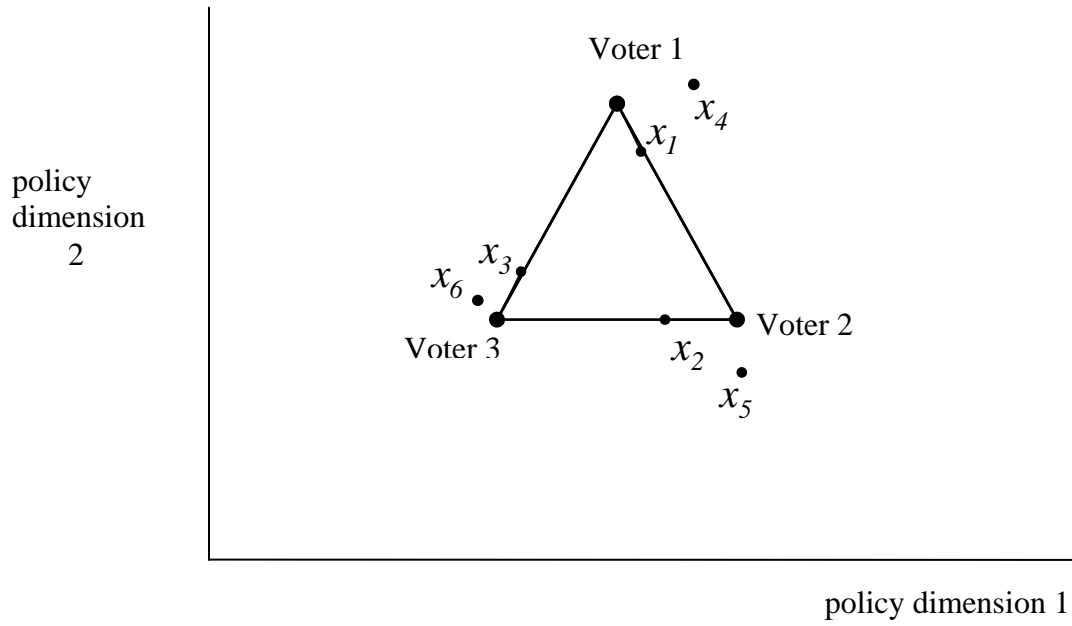


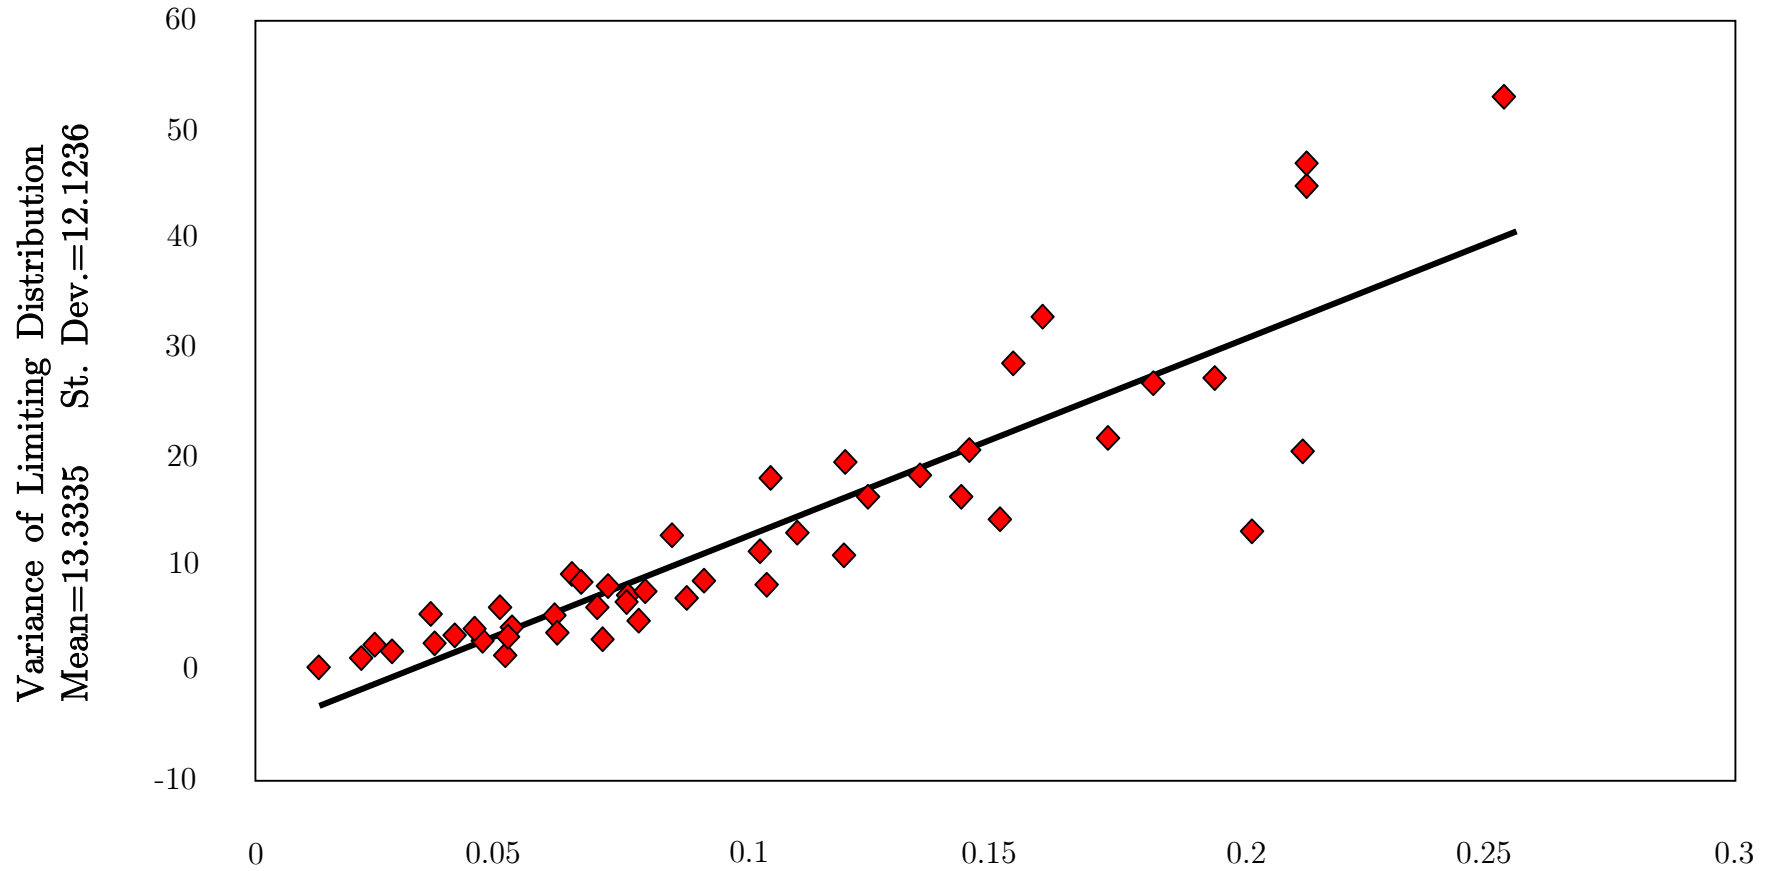
Figure 2

There are 2 L 's in this electoral landscape:

$$L_1 = \{x_1, x_2, x_3\} \text{ and } L_2 = \{x_4, x_5, x_6\}.$$

Note: voters are assumed to have symmetric, single-peaked preferences (circular indifference curves).

Figure 3: Normal search rule around last policy



Ratio of Size of the Uncovered Set to Size of Top Cycle
Mean=0.1000 St. Dev.=0.0594

Regression Statistics: Slope=187.6 T-statistic=16.17 $R^2=0.84$