The Constrained Instability of Majority Rule: Experiments on the Robustness of the Uncovered Set

William T. Bianco

Department of Political Science, Indiana University,
Bloomington, IN 47405
e-mail: wbianco@indiana.edu

Michael S. Lynch

Department of Political Science, University of Kansas,
504 Blake Hall, Lawrence, KS 66044
e-mail: mlynch@ku.edu

Gary J. Miller and Itai Sened

Department of Political Science, Washington University in St. Louis,
Campus Box 1063, One Brooking Drive, St. Louis, MO 63130
e-mail: gjmiller@wustl.edu
e-mail: sened@wustl.edu (corresponding author)

The uncovered set has frequently been proposed as a solution concept for majority rule settings. This paper tests this proposition using a new technique for estimating uncovered sets and a series of experiments, including five-player computer-mediated experiments and 35-player paper-format experiments. The results support the theoretic appeal of the uncovered set. Outcomes overwhelmingly lie in or near the uncovered set. Furthermore, when preferences shift, outcomes track the uncovered set. Although outcomes tend to occur within the uncovered set, they are not necessarily stable; majority dominance relationships still produce instability, albeit constrained by the uncovered set.

1 Introduction

After decades of research, there is a clear divide between theory and empirical evidence on majority rule voting. On one hand, there is a theoretical argument that majority rule voting is inherently unstable (Riker 1980). Indeed, the so-called Chaos Theorem implies that in the absence of a core, majority rule voting in an institution-free environment may not yield unambiguous predicted outcomes (McKelvey 1976, 1979; Schofield 1978; McKelvey and

Schofield 1987). On the other hand, majority rule voting, both in the real world and in experiments, often (but not always) yields centrally located outcomes, even when no core exists (Fiorina and Plott 1978; McKelvey et al. 1978; Tullock 1981). Even when outcomes are not in the center of voters' preferences, they are often (but again, not always) located within a fairly small region of the Pareto set. Both results suggest that expectations of inherent chaos are overdrawn, and that some yet unknown feature of majority rule voting is at work when this procedure is used in the real world.

Theories that have attempted to explain the seemingly anomalous behavior of majority rule have generally consisted of a solution concept that coincides with the core when it exists, but which can predict the observed patterns in settings without the core. Early contenders were the Von-Neumann solution, the bargaining set, and McKelvey and Ordeshook's competitive solution. Unfortunately, these solution concepts did not fare well in experimental testing (McKelvey et al. 1978; McKelvey and Ordeshook 1983). To date, no theory has consistently predicted the outcomes of majority rule voting in situations where there is no core.

Without an explanation for majority rule outcomes based in preferences and the particulars of the decision rule, the focus for explaining such outcomes has largely shifted to the institutions under which voting takes place. Shepsle's (1979) structurally induced equilibrium argues that institutions induce equilibria on voting. In the U.S. Congress, for example, agenda control, the committee system, and the germaneness rule are thought to combine to force complex multidimensional issues into a series of one-dimensional, single-issue votes (Cox and McCubbins 2005). In this way, the search for explanations of observed outcomes in Congress and elsewhere has focused on these and other institutions, including partisan organizations, leaving the most fundamental institution, majority rule itself, largely ignored.

In this paper, we go "back to the future" by revising expectations of outcomes under majority rule. Our focus is on the uncovered set, a well-known but little-used predictor of majority rule outcomes. Although the uncovered set was formulated in the 1980s, until recently, it was impossible to estimate uncovered sets but for the simplest cases.[1] However, in recently published work, Bianco et al. (2004) offer a grid search procedure that allows the computation of the uncovered set. This technique moves the uncovered set from a theoretically appealing, but difficult to implement concept, to a potentially useful tool for predicting the outcome of majority rule in real-world settings.[2]

The purpose of this paper is to use this new technology to test the robustness of the uncovered set as a solution concept in majority rule settings. This is done through a series of experiments: five-person, computer-mediated experiments and a 35-person pencil-and-paper version. We find that the uncovered set performs well as a solution concept. Its success helps us to address the question of "why so much stability?" (Tullock 1981) in many majority rule settings. Our conclusion is that there is strong empirical evidence

---

[1]For special cases, see Miller (1980), Feld et al. (1987), Hartley and Kilgour (1987), Epstein (1998), and Penn (2006). Many intuitions derived from these special cases are not generally true. For example, while the uncovered set is a subset of the Pareto set (Miller 1980), the percentage of the Pareto set occupied by the uncovered set can vary substantially (Bianco et al. 2004). Thus, while the Pareto set is easily determined, it does not provide much insight into the uncovered set as previously argued (Epstein 1998). The same is true for the yolk (Bianco et al. 2006).

[2]The uncovered set has previously been implemented in a variety of settings considering a limited set of alternatives. The grid-search procedure used in this article allows for the implementation of the uncovered set for a large number of finite alternatives, effectively allowing for the approximation of an infinite alternative space (Bianco et al. 2004).

that majority rule decision processes are constrained by the boundaries of the uncovered set.

2 The Uncovered Set: Theoretic Rationale and Empirical Relevance

Research on majority rule has at its core a fundamental question: which outcomes are possible? Contemporary spatial modeling offers two important predictions. In a one-dimension spatial model, the expected outcome of majority rule voting is the ideal point of the median voter. This result does not generalize to cases where multiple dimensions characterize preferences and outcomes, implying that such outcomes are sensitive to agendas, voting rules, and other constraints (Shepsle 1979, 1986). "Chaos Theorems" (McKelvey 1976, 1979; Schofield 1978; McKelvey and Schofield 1987) state that the outcomes of majority rule voting unchecked by institutions can be indeterminate.

Further work showed that if voters consider the ultimate consequences of their actions across an agenda of votes, rather than choosing myopically between alternatives at each decision point, majority rule voting will yield an outcome in the *uncovered set* (Miller 1980; McKelvey 1986).[3] "If one accepts. . .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 of the opponent's possible strategies, then one can conclude that candidates will confine themselves to strategies in the uncovered set" (Cox 1987, 419). While Cox's argument focuses on elections, its logic also applies to legislatures. Covered outcomes are unlikely to be offered or supported by sophisticated decision makers who know these proposals cannot win. Moreover, supporters of uncovered outcomes can secure these outcomes using simple agendas and defend them against opponents who want something else (Shepsle and Weingast 1984).

The work presented in this paper utilizes a new grid-search computational method (Bianco et al. 2004) for estimating the uncovered set.[4] Using this technique, Bianco et al. (2006) calculated the uncovered set for all previously published, relevant, two-dimensional majority rule experiments. These experiments have two common characteristics: none of them were designed as a means of testing the uncovered set, and none of them show systematically convincing evidence for any of the solution concepts that they were designed to test. The analysis finds that the uncovered set does very well across the range of previously published experiments—better than the solution concepts that the experiments were designed to test. Overall, 93% of the outcomes fell in the uncovered set. This result emerged despite a remarkable dispersion of institutional designs—formal and informal majority rule decision-making processes, limited or full communication, forward, and backward agendas.

2.1 *Further Questions*

Despite the success of the uncovered set in postdicting the results of previous spatial majority rule experiments, a series of important questions remain. For one, many previous experiments happened to be designed so that the uncovered set is relatively large. This

---

[3]Formally, let N be the set of n voters or legislators. Assume n is odd. For any agent, $i \in N$, (Euclidian) preferences are defined by an *ideal point* ρ_i. Let x, y, z be elements of the set X of all possible outcomes. A point x *beats* another point y by majority rule if it is closer than y to more than half of the ideal points. A point x is *covered* by y if y beats x and any point that beats y beats x. The uncovered set includes all points not covered by other points.

[4]Space considerations preclude a detailed discussion of the algorithm. For details, see Bianco et al. (2004).

means that even though the uncovered set did well in predicting the results of the pre-viously published experiments, those experiments constitute a relatively weak test of the uncovered set. The experiments presented here are designed so that the uncovered set is a relatively small subset of the Pareto set.

Another question is stability. Many interpret the largely centrist results of previous majority rule experiments as indicating a kind of unexpected stability of majority rule decision making. However, it could be that majority rule decision making is unstable. Cycles may exist in a small, centrist subset of the policy space, and lead to centrally located, but unstable, outcomes. Our experiments are designed to assess the amount of stability present in majority rule decision making.

Finally, the ideal experiments should examine the robustness of the uncovered set to a variety of payoffs, procedures, and communication conditions under majority rule. The experiments reported here incorporate these institutional differences.

2.2 *Why Experiments?*

The study of majority rule has been one area of political science in which the symbiosis between theory and laboratory experimentation has been most productive, and progress has been marked and cumulative. Laboratory experiments can be explicitly designed to create critical tests between alternative theoretical predictions. And theorists have regarded the results of laboratory experiments as definitive—embracing solution concepts that have survived in the laboratory and repeatedly abandoning those that have failed careful laboratory experimentation.[5]

As compelling as empirical research on real-world legislative bodies is, a controlled laboratory experiment is still an essential part of a complete program of empirical testing of any solution concept. If the House of Representatives, for example, chooses a different package of policy choices when the partisan majority shifts, in a manner consistent with a shift in the uncovered set, it is impossible to determine, with any satisfactory level of certainty, whether that correlation is due to the shift in preferences, institutional changes, a change in historical context, or even to random variation.[6] Tests with naturally occurring data from the House of Representatives or elsewhere rarely constitute a definitive test by themselves due to difficulties inherent in estimating legislators' preferences and the exact spatial location of final policy decisions and the subtlety of capturing the impact of institutional settings.

With laboratory experimentation, however, a large number of decisions can be observed with two or more treatments that correspond to different sets of preferences. A large number of subjects may be randomly assigned to preferences in each of the treatments. All possible controllable differences are held constant across the treatments, and those that are not controllable (peculiarities of the subjects, for instance) are neutralized by random assignment to treatments. As a result, when we observe statistically different variations in outcomes under two treatments, we can be confident that the treatment variables are causal. When the preferences in the treatments are designed in such a way as to induce sharply different predicted uncovered sets, then a failure of the outcomes in the two treatments to track the uncovered set would be a convincing falsification of the relevance of the uncovered set as a solution concept.

---

[5]See McKelvey and Ordeshook (1990) for a review of this literature.
[6]Indeed, Ordeshook and Schwartz (1987) demonstrate that certain House institutional rules make it possible for agendas to lead to outcomes outside of the uncovered set.

3 The Experiments

The experimental results reported here involved two relatively different decision contexts, labeled here as the small-N and large-N designs:

3.1 *Small-N Design*

The small-N experiments were conducted at one medium-sized private university in the Midwest. Each experiment consisted of five players. The experiments were computerized, allowing for full player communication through the use of an unlimited (but anonymous) messaging system and a computer-mediated system of randomly assigned recognition of agenda setters. Their task was to choose an outcome using majority rule. Players were given payoffs that were a linear function of the distance between their ideal point and the final outcome of the majority rule decision. Details on the payoff functions are in Appendix.

The top two plots in Fig. 1 show the two configurations in this design, with player ideal points denoted as diamonds and the uncovered set as a gray shape. We refer to the Small-N Configuration One design as S1 and the Small-N Configuration Two design as S2. The ideal points of players 2–5 remained constant in the two configurations, while players 2 and 3 have ideal points close together, making them in some sense natural allies. The same is true for players 4 and 5. The only difference in the two configurations was the location of player 1's ideal point, making her the natural ally of one cluster or the other in S1 and S2. The uncovered set shifts dramatically with the movement of player 1. When player 1 is located at the top of the figure, the five players form a five-side Pareto-optimal hull and the uncovered set is rather large—approximately 64% of the Pareto set for the five players. In S2, player 1's ideal point is inside the Pareto set of the other four players, and the uncovered set shrinks dramatically to 18% of the five-player Pareto set.

The experiments were conducted in sessions of 10 or 15 students, with two or three voting groups per session. Participants were paid a show-up fee of $5 plus their earnings in the experiment. Participants read the instructions as they were read aloud by the experimenter. They had to answer several questions correctly about payoffs and procedure before going on to a practice session (see Appendix for complete player instructions). The computer screen showed the exact payoff of each point, as the cursor landed on it. It also showed the ideal points of all five players and the status quo (see Fig. 6 for a sample screen image).

After a practice session, subjects were allowed 20 min to make decisions. Each participant had an equal chance to be recognized as the first agenda setter. The agenda setter had 90 s to either make a proposal, propose to adjourn, or pass. A proposal point appeared on the screen, and the participants' payoff for that point and the status quo point were readily available. After a proposal was made, agents had 30 s to make the voting decision that appeared on the screen. A simple majority was required to move the status quo from the original point (75, 75) to the point proposed. After a vote result was announced and shown on the screen, another participant was randomly recognized.[7] Agenda setters would continue to be randomly recognized until a move to adjourn was made and approved by majority vote.

Half the sessions made an S1 decision first, resulting in a final outcome. Participants were told that they would be paid the cash payoff associated with the final outcome—the status quo at the time of adjournment. They were then given the opportunity to make additional money by participating in a second similar decision. Every single group

---

[7]Participants were randomly recognized without replacement. After all five players were randomly selected, the process was repeated.
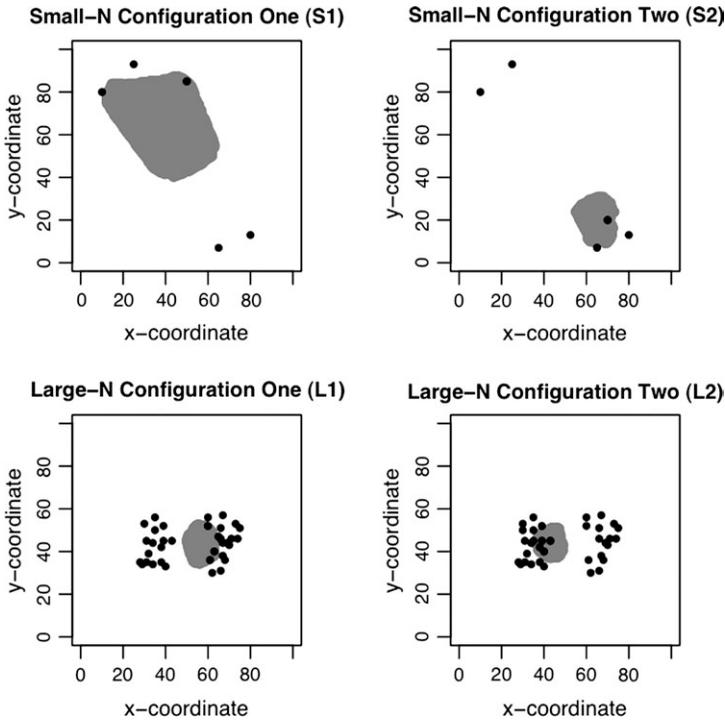
Fig. 1 Large and small-N experimental designs.

unanimously elected to proceed with a second decision. If the group participated in a S1 decision the first time, the second decision was S2 and vice versa.

3.2 *Large-N Design*

The large-N design was a 35-participant, paper-and-pencil format. There were two configurations, L1 and L2, as shown in the bottom two plots in Fig. 1 along with the associated uncovered sets. Like the estimated ideal points of members in Congress, the ideal points in the L1 and L2 were divided into two "partisan" clusters. The difference between L1 and L2 may be thought of as resulting from the switch of a few pivotal voters from one partisan cluster to another, as if a few competitive seats had changed partisan hands.

Notice that the uncovered set largely occupies the "empty" area between the partisan clusters. This is particularly challenging, in that there were few "inhabitants" of the uncovered set to advocate outcomes close to their ideal point. It would seem that difficult multilateral negotiations would be necessary to arrive at an outcome in the uncovered set; however, such negotiations were impossible under the large-N procedure, which forbade open discussion and consisted largely of proposals being made by randomly selected subjects.

Each experimental run began by randomly assigning each subject an ideal point. Subjects were then given a figure that depicted their ideal point and indifference curves that enabled them to calculate their payoff from a particular (x, y) proposal. (Fig. 7 shows one such figure.) A period of instruction, a practice decision, and a question-and-answer session followed.

The actual experiment proceeded as follows. The status quo outcome at the beginning of a run was (40, 60). One subject was randomly chosen to propose an outcome—an (x, y) pair. A majority rule vote followed. If a majority of subjects voted for the proposal, it

became the new status quo. If the proposal did not receive a majority, the old status quo remained in place. Following the vote, the moderator asked for a motion to adjourn.[8] If the motion carried, the experiment ended and subjects received payoffs based on their ideal point and the current status quo. If a motion to adjourn was not proposed or did not carry, another subject was chosen to propose a new outcome, which was then voted on, and so on until a motion to adjourn receives a majority.

The number of experimental runs per group of subjects was initially set at two, first L1 then L2 or first L2 then L1, with subjects receiving new ideal points for each decision. However, midway through data collection, the number of decision per experimental group was expanded to three (either L1–L2–L1 or L2–L1–L2), and then to four (L1–L2–L1–L2 or L2–L1–L2–L1). Again, ideal points were randomly selected for each run. As there appears to be no differences in outcomes or votes as a function of the number of runs, the results will be presented together.

3.3 *Comparison of the Experimental Designs*

These two experimental designs represent polar opposites in terms of communication between subjects and the difficulty of the task before them. These differences were chosen to stress test the predictive power of the uncovered set—to see if predictive power varies with communication, the number of participants, and the other differences between the two designs. Subjects in the small-N designs S1 and S2 can communicate freely, allowing them to form coalitions to secure outcomes. In contrast, the large-N participants (L1 and L2) cannot communicate; their only actions during the experiment are to propose outcomes (if they are selected), to vote on outcomes, and to propose and vote on motions to adjourn. Small-N participants can use their computer screens to easily calculate payoffs from different outcomes; in contrast, large-N participants must interpolate this information from the ideal point and indifference curve diagrams distributed to them at the beginning of each experimental run. Large-N participants must also choose strategies given a much larger number of participants, which may increase the difficulty of predicting the consequences of different actions.

The two designs share the characteristic that the location of the uncovered set varies across configurations. This variation provides a rough test of the uncovered set: our null hypothesis was that the outcomes were not affected by the preference-induced shift in the uncovered set. The working hypothesis was that outcomes selected will be within the uncovered set in both configurations.

3.4 *Hypothesis Tests*

Our analysis of the experimental data proceeds as follows. First, given a set of ideal points and experimental outcomes for each design and configuration, we plot the estimated uncovered sets, overlay all proposals that win by majority rule at some point in a decision process, as well as all proposals chosen as final outcomes, and assess whether these outcomes are contained within the estimated uncovered set. Our analysis accounts for outcomes that are inside the uncovered set as well as close misses—outcomes that are within one unit of the uncovered set calculated for each configuration.[9]

---

[8]This is different than the procedure used in the small-N design. In the small-N experiments, the recognized proposer had to choose between making a proposal and moving to adjourn. This likely made the cost of moving to adjourn higher in the small-N experiments and likely helps to explain the fact the groups, on average, made more proposals prior to adjournment in the small-N experiments.

[9]We include close misses because previous analysis shows that our algorithm slightly (~1%) underestimates the size of the uncovered set (Bianco et al. 2004).

Table 1 Basic data on the experiments

	Small-N design		Large-N design	
	S1	S2	L1	L2
Average N of proposals	9.32	6.61	5.91	4.41
Average N of winning proposals	5.25	3.25	1.91	1.77
Average time to decision	8:23	6:23	—	—
N sessions	28	28	25	22

We also calculate the expected distribution of proposals and final outcomes in each experimental configuration. The aim is to describe the aggregate distribution of outcomes, in order to determine whether the distribution of final outcomes is measurably different from the distribution of proposals, as well as to see whether there is variation in the location of final outcomes across different configurations. This analysis assumes that these data are distributed bivariate normal.[10] Our expectation is that the expected distribution of final outcomes should be centered on and roughly the same size as the uncovered set, whereas the distribution of all proposals (which includes those that did not gain a majority) should be significantly larger, approaching the limits of the Pareto set in each configuration.[11]

We utilize a binomial test to assess whether outcomes are occurring in the uncovered set more often than we would expect by chance alone, given a process that is constrained by the Pareto set. We compare the percentage of outcomes occurring in the uncovered set (observed π) with the percentage of outcomes that one would expect to find in the uncovered set given a uniform distribution of outcomes throughout the Pareto set (hypothesized-π). If the observed π is significantly higher (one-tailed test) than the hypothesized π, we reject the null hypothesis that outcomes in the uncovered set are as likely as outcomes in the Pareto set.

4 Results

Table 1 gives basic information about the four sets of experimental runs, including the number of trials, the average number of proposals offered in each configuration, the average number of winning proposals, and the average time to completion (the last is only available for S1 and S2).

4.1 *Example One: S1*

To illustrate the dynamics of the decision-making processes in these experiments, we review the major decisions made by a particular five-person group in S1. All five members of this group were very active in the process. Each of the five members made at least two successful proposals, and three of them (players 1, 3 and 4) made four successful proposals each. Each participant's successful proposals were approved by more than one coalition. Thirteen of the 16 winning proposals were supported by minimal winning coalitions.

---

[10]This distribution is not a prediction of the uncovered set per se but seems a plausible characterization of the outcomes that are likely to result from the experimental runs. Assume that our x and y data have means μ_x, μ_y, standard deviation σ_x, σ_y, and correlation ρ. If so, $((x - \mu_x)/\sigma_x)^2 - 2\rho((x - \mu_x)/\sigma_x)((y - \mu_y)/\sigma_y) + ((y - \mu_y)/\sigma_y)^2 = -2(1 - \rho^2)\log \alpha$ contains $100(1 - \alpha)\%$ of the bivariate distribution (Johnson and Kotz 1972, 87). We are indebted to Tse-Min Lin and James Granato for their help on this matter.

[11]A point x is unanimously preferred to a point y if x is closer than y to all voters' ideal points. The Pareto set is the set of points such that there is no point that is unanimously preferred to any point in the Pareto set.
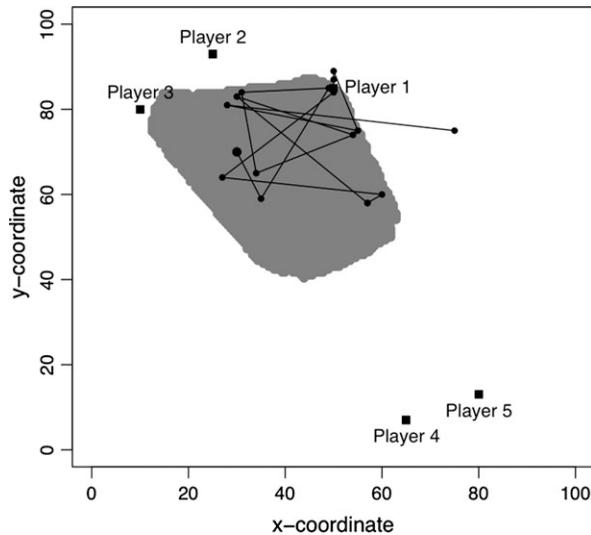
Fig. 2 Sample S1 experimental session. *Note.* This figure shows the winning proposals for Group I from the September 24, 1:00 p.m. session. Square dots are ideal points. The gray field is the estimated uncovered set. Small dots are winning proposals. The large dot is the final outcomes of the session. The lines show the progression of the winning proposals.

Fourteen of the winning proposals were in the uncovered set, and the other two were almost exactly on player 1's ideal point. Figure 2 depicts the proposals, votes, and ultimate outcome.

The first successful proposal was made by player 2. It was an initial bid to the partisan coalition of 1, 2, and 3. It was very close to the midpoint of the contract curve between 1 and 3 and close to 2's ideal point. After this first successful proposal, there were five losing proposals. Then player 4 proposed a point (55, 71) and was backed by players 1 and 5. Then player 1 re-enlisted the other members of the partisan coalition to vote for (50, 89). This proposal is an interesting strategic move as it moved the status quo outside the Pareto set, just 4 units away from the player's ideal point—a small subsequent adjustment might well bring the policy exactly back to the proposer's ideal point. This is what happened, with a four-person coalition including 3, 4, and 5.

Unfortunately for player 1, this result was not the end of the game. Player 3 broke with his partisan coalition members to mobilize players 4 and 5 to move across the uncovered set toward his ideal point to (27, 64). The policy was moved across the uncovered set to (27, 65) by the coalition of 3, 4, and 5. Then four enlisted players 1 and 5 for a move farther away from the partisan coalition to (60, 60). With this change and the subsequent two moves, players 4 and 5 were playing different wings of the partisan coalition off against each other to move the Y coordinate lower to (58, 58). This collapsed when player 3 regrouped the partisan coalition behind (30, 83); but player 1 immediately broke with her fellow partisans to form another coalition with 4 and 5.

In this group, the cycle of appealing proposals for players 4 and 5, followed by a collapse toward the "natural" coalition, occurred two more times. The final outcome was at (30, 70)—a point not particularly near the partisan coalition and reflecting the bargaining power of players 4 and 5. At the end of the 16 successful moves, players 4 and 5 had been just as active in successful coalitions (voting for ten successful moves each) as players 1, 2, and 3 who voted in 11, 8, and 12 successful coalitions each.
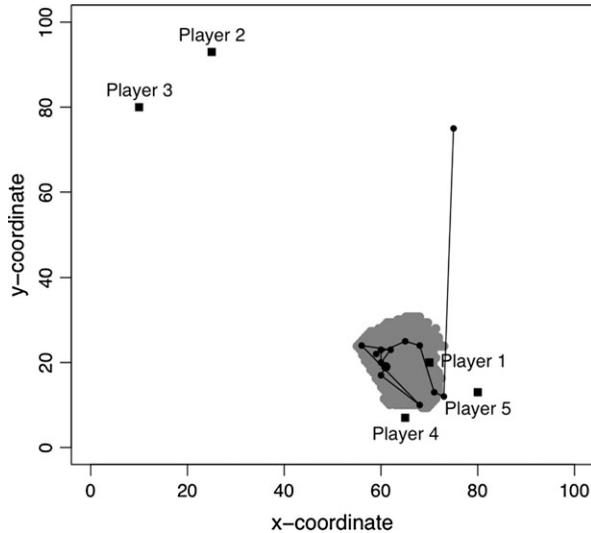
Fig. 3 Sample S2 experimental session. *Note*. This figure shows the winning proposals for Group 2 from the September 17, 2:30 p.m. session. Square dots are ideal points. The gray field is the estimated uncovered set. Small dots are winning proposals. The large dot is the final outcome of the session. The lines show the progression of the winning proposals.

This process is one that the uncovered set seems to capture quite well as the cycles toward and away from 4 and 5 virtually fill the space of the uncovered set, whereas moving out of the uncovered set occurred only twice when player 1 got a temporary majority for a policy very near her ideal point. In other words, with a vigorous multilateral negotiation process, the uncovered set captured the trajectory and the ultimate result of majority rule cycling.

4.2 *Example Two: S2*

An additional example from an S2 session shows how sessions in which the coalition formation was most vigorous and multilateral seemed to produce the outcomes in the uncovered set. One session produced 12 successful proposals all over the uncovered set before finally adjourning. No region of the uncovered set was immune to further coalition formation. The trajectory of proposals, winning proposals, and final outcomes is shown in Fig. 3. In this session, all five players made between 1 (for player 2) and 4 (for player 5) successful proposals. The first proposal was the partisan majority coalition—1, 4, and 5 to a point close to these three players, but outside the uncovered set. Coalition 2, 3, and 4 supported a small adjustment that moved the policy inside the uncovered set, and it never left the uncovered set again, although it moved around almost the entire space of the uncovered set.

The subsequent moves were up and to the left, involving minority players 2 and 3. Although all the successful proposals except the first were in the uncovered set and close to the ideal points of players 1, 4, and 5, the minority players 2 and 3 played an active and indeed essential part in the negotiations. Indeed, players 2 and 3, along with player 5, all voted in favor of 9 of the 12 successful proposals, compared to eight and seven successful proposals for players 4 and 1, respectfully. The partisan majority coalition 1, 4, and 5 formed only in the first, seventh, and tenth successful moves. As in the S1 experiments, active participation by the partisan minority members seemed to be sufficient to generate

Table 2 Proposal locations: Small-N design

	In uncovered set	Not in uncovered set	Total
S1			
Proposal in winset	155 (59.4%)	20 (7.7%)	175 (67.0%)
Proposal not in winset	70 (26.1%)	16 (6.1%)	86 (33.0%)
Total	225 (86.2%)	36 (13.8%)	261 (100%)
S2			
Proposal in winset	38 (20.5%)	57 (30.8%)	95 (51.4%)
Proposal not in winset	23 (12.4%)	67 (36.2%)	90 (48.6%)
Total	61 (33.0%)	124 (67.0%)	185 (100%)

Note. Numbers indicate the raw number of proposals made in the various possible locations. Numbers in parentheses are the percent of total outcomes present in the respective locations. This analysis does not include close misses.

an outcome in the uncovered set and normally outside the Pareto set of the majority coalition.

4.3 *Analysis of Proposals*

The two examples described above demonstrate the strategic considerations that players seem to be making when proposing alternatives and voting for those alternatives. Sophisticated players should confine themselves to proposing alternatives that are located in the uncovered set because covered points can be easily defeated and are unlikely to be enacted as a final outcome. Further, sophisticated players are likely to propose alternatives that are in the winset of the current status quo. This assures that the alternative is attractive to enough players to allow it to defeat the current status quo. Table 2 provides the location of proposals made during the small-N experiments and Table 3 reports the win rates of proposals by location.

As expected, the vast majority of proposals for configuration S1 were located in the uncovered set (86.2%) and in the winset of the current status quo (67.0%). In total, 59.4% of proposals made were in the union of the uncovered set and the winset of the current status quo. Covered proposals and proposals made outside the winset of the current status

Table 3 Proposal win rates by location: small-N design

	In uncovered set	Not in uncovered set	Total
S1			
Proposal in winset	116/155 (74.8%)	11/20 (55.0%)	127/175 (72.6%)
Proposal not in winset	17/70 (24.3%)	3/16 (18.8%)	20/86 (23.3%)
Total	133/225 (59.1%)	14/36 (38.9%)	147/261 (56.3%)
S2			
Proposal in winset	32/38 (84.2%)	43/57 (75.4%)	75/95 (78.9%)
Proposal not in winset	5/23 (21.7%)	11/67 (16.4%)	16/90 (17.8%)
Total	37/61 (60.7%)	54/124 (43.5%)	91/185 (49.2%)

Note. Ratios indicate the number of winning proposals by the number of total proposals for each possible location. Numbers in parentheses are the percent of proposals that were winning. This analysis does not include close misses.

Table 4 Measuring predictive power (1): Hits and misses

	Small-N design		Large-N design	
	S1	S2	L1	L2
Final outcomes in Pareto set	28 (100.0%)	26 (92.9%)	25 (100.0%)	22 (100.0%)
(including close misses)	28 (100.0%)	26 (92.9%)	25 (100.0%)	22 (100.0%)
Final outcomes in uncovered set	28 (100.0%)	17 (60.7%)	10 (40.0%)	10 (45.5%)
(including close misses)	28 (100.0%)	21 (75.0%)	16 (64.0%)	13 (59.1%)

quo were much more likely to be defeated than other proposals. While 74.8% of proposals in the union of the uncovered set and the winset of the current status quo received enough votes to become the new status quo, only 38.9% of covered proposals passed. Only 23.3% of proposals outside the winset of the current status quo were able to win.

Far fewer proposals (33.0%) made were located in the uncovered set in the smaller S2 configuration. The number of proposals made in the winset of the current status quo was also far lower (51.4%). In total, only 20.5% of proposals made were located in both the uncovered set and the winset of the current status quo. This limited number of proposals did win at very high rates (84.2%). And despite the large number of covered proposals and proposals made outside winset of the current status quo, the number of these proposals that were able to muster majority support was low—43.5% and 17.8%, respectively.

These results indicate that while players did not always propose alternatives in a way that implies sophisticated behavior, unsophisticated proposals were unlikely to survive the voting process. And those proposals that did manage to gain majority support for a single iteration of voting were seldom able to continue to defeat new proposals and, as such, very infrequently survived to be a group's final outcome. Indeed, as the next section reports, the final outcomes of these voting games are overwhelmingly located in the uncovered set.

4.4 *Analysis of Outcomes*

Table 4 gives summary data on the predictive power of the uncovered set across the two designs and four configurations used in our analysis. For each configuration, the table reports the percentage of final outcomes that were in the respective uncovered set—both the percentage inside and the percentage including close misses (outcomes that were within one grid square of the uncovered set). In one configuration (S1), all the experimental outcomes were inside the uncovered set. In the other three configurations, the percentage of outcomes in the uncovered set were between 59% and 75%.

These findings provide impressive support for the uncovered set's predictive power. The participants in these experiments were typical college students recruited from all majors and backgrounds, with no experience in majority rule decision making. Nevertheless, their collective decisions found the uncovered set a relatively high percentage of the time.

Moreover, even when outcomes for a given configuration were not in the uncovered set or a close miss, they were typically in the vicinity of the set. Consider Fig. 4, which shows the distribution of all proposals and of final outcomes for S1 and S2. The figures also show the 2-sigma–expected distribution of proposals and final outcomes, calculated using the bivariate normal distribution discussed earlier.

The two plots in Fig. 4 confirm the predictive power of the uncovered set. For one thing, the location of final outcomes shifts profoundly across the two configurations—from the
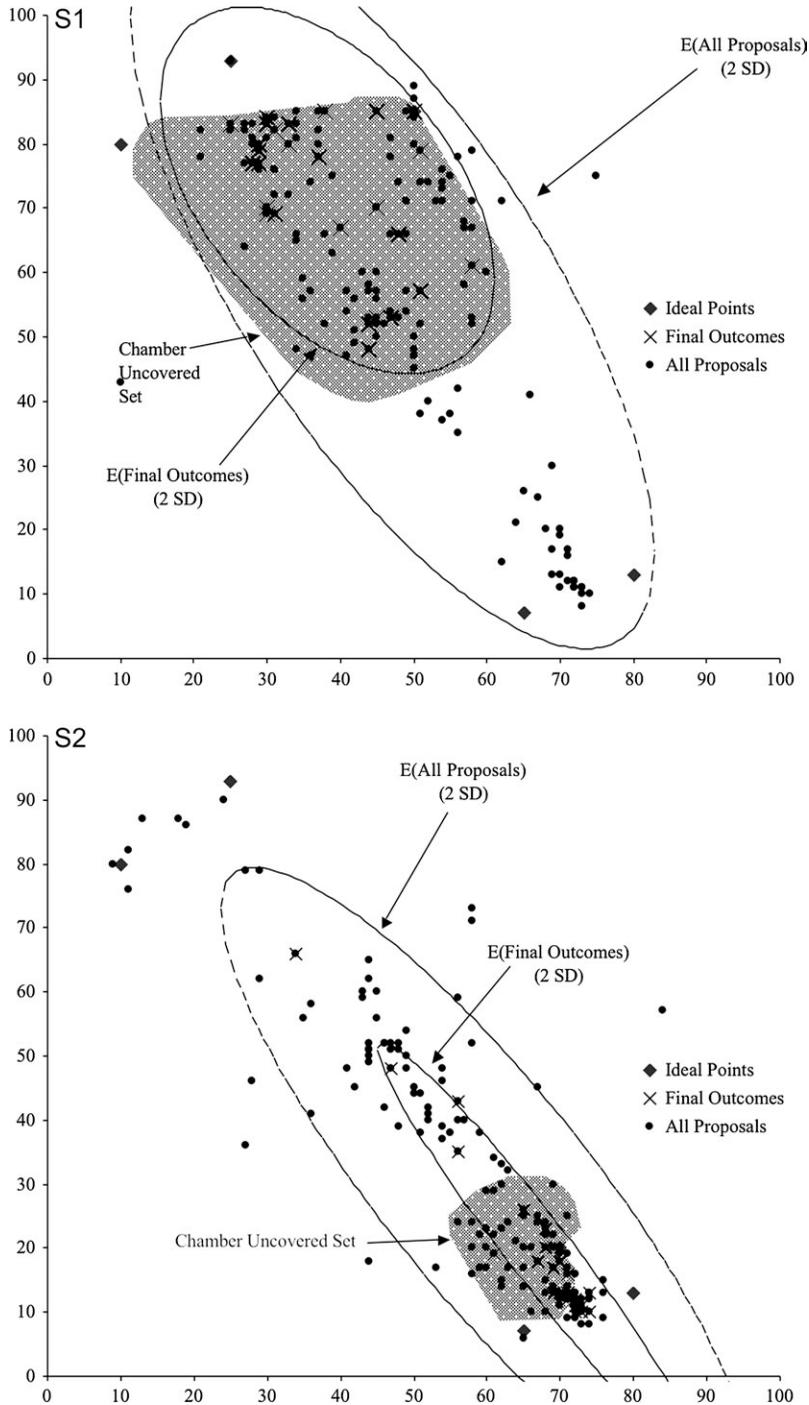
Fig. 4 Outcomes in the small-N experiments. Configuration 1 (S1) and Configuration 2 (S2).

upper-left quadrant in S1 to the lower-right quadrant in S2, matching the shift in the uncovered set. Thus, final outcomes in these experiments were highly sensitive to the location of the uncovered set, even when the decision process did not yield an uncovered
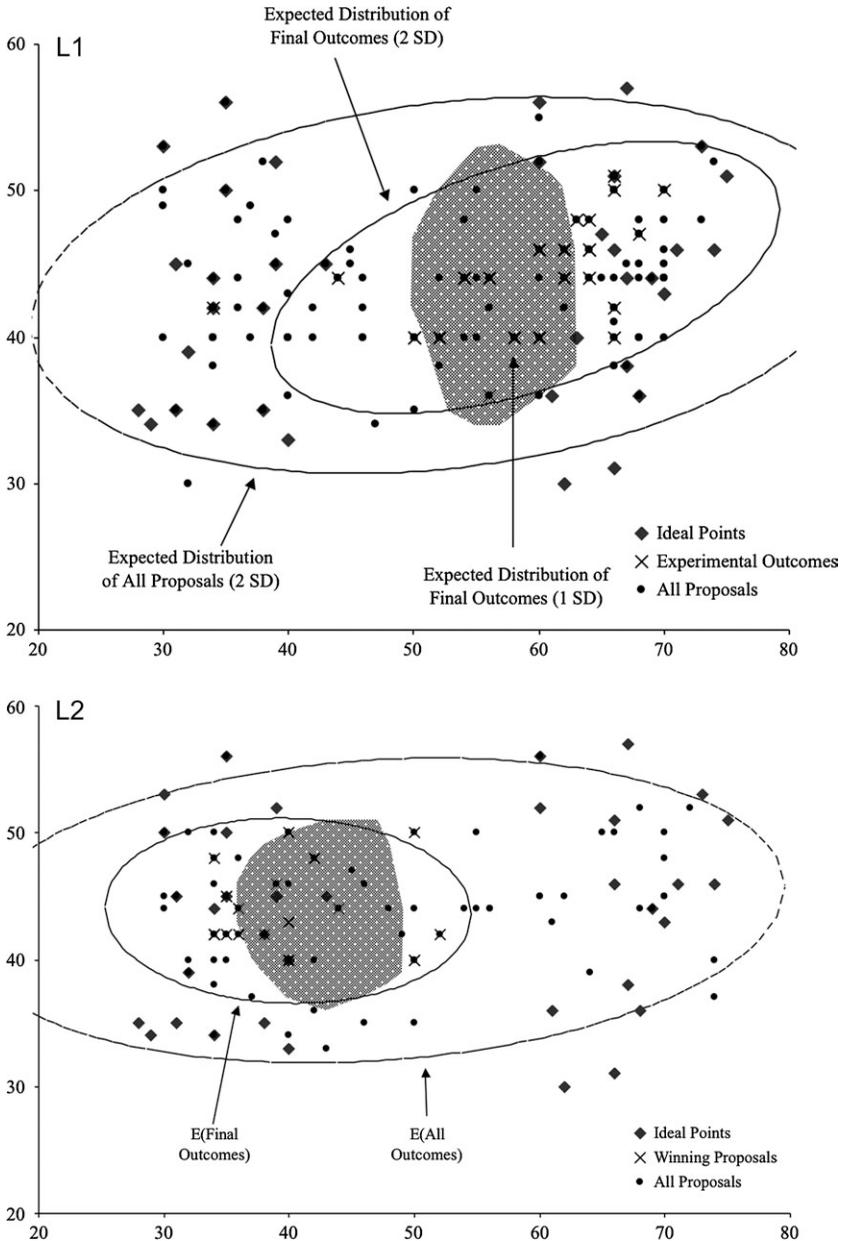
Fig. 5 Outcomes in the large-N experiments. Configuration 1 (L1) and Configuration 2 (L2).

outcome in S2. Similarly, the size and location of the 2-sigma ellipses in the two config-
urations also match our expectations, with the distribution of all proposals covering almost
the entire Pareto set, but the distribution of final outcomes being roughly centered on and
the same size as the uncovered set in each configuration.

Next, Fig. 5 provides final outcome and distribution of expected outcome information
for the two large-N configurations, L1 and L2. As with the small-N experiments, the plots
show that the final outcomes shift with the change in the location of the uncovered set

Table 5 Measuring predictive power (2): Binomial test

	Small-N design		Large-N design	
	S1	S2	L1	L2
Hypothesized % hits	64.1	16.7	16.8	13.4
Actual % hits	100.0	75.0	64.0	59.1
p Value	.0000	.0000	.0012	.0000

Note. Close misses are included as hits.

across the two treatments. Moreover, the distribution of expected outcomes matches our expectations.

The final demonstration of the uncovered set's predictive power is our binomial test, as reported in Table 5. As the table shows, in all four configurations, the uncovered set's predictive power far exceeds the Pareto set baseline. For example, in S1, while the uncovered set constituted only 64.1% of the S1 Pareto set, the uncovered set performs significantly better than is expected by chance for final outcomes—the observed π of 1.000 is significantly higher than the hypothesized π of 0.641 ($p = .000$).

5 Conclusion

Since Shepsle (1979, 1986), the prominent explanation for the empirical stability of majority rule outcomes has been institutional. Institutions of agenda control, germaneness rules, or committee veto power have seemed to constrain majority rule in a predictable way. This result justifiably causes concern since institutional constraints on majority rule may privilege those with greater institutional power; strict monopolistic agenda control, for example, transforms a formal majority rule democracy into something very much like a dictatorship (McKelvey 1976). Yet institutional constraints have appeared to be the only alternative to majority rule instability and arbitrariness. An important implication of our results is that institutions are not the only binding constraint on majority rule outcomes, and that simple majority rule, with an open agenda, is a largely predictable process.

Majority rule instability is certainly present in these experiments. Many experimental sessions experienced the formation of multiple coalitions and approved many different policy outcomes. There seemed to be no area of the uncovered set that was unbeatable by another majority coalition. Nevertheless, the instability of majority rule voting was largely limited by the boundaries of the uncovered set.

These experiments offer striking evidence of the ability of the uncovered set to predict outcomes in majority rule settings, even when the uncovered set is small. Just as importantly, majority rule demonstrated positive responsiveness to individual preference changes. In particular, majority rule outcomes shift significantly with the movements of pivotal players from one partisan cluster to another, and the shifts are consistent with the shift predicted by the uncovered set. The treatment effect is a striking demonstration of both the predictability of majority rule and the responsiveness of majority rule decision making to individual preference changes.

The predictability of outcomes by the uncovered set is significant in light of the pessimism about majority rule voting. In their critique of rational choice theory, Green and Shapiro (1994) placed a good deal of emphasis on the expanding role of the uncovered set, unsupported by empirical investigation. Up until recently, they note, the uncovered set
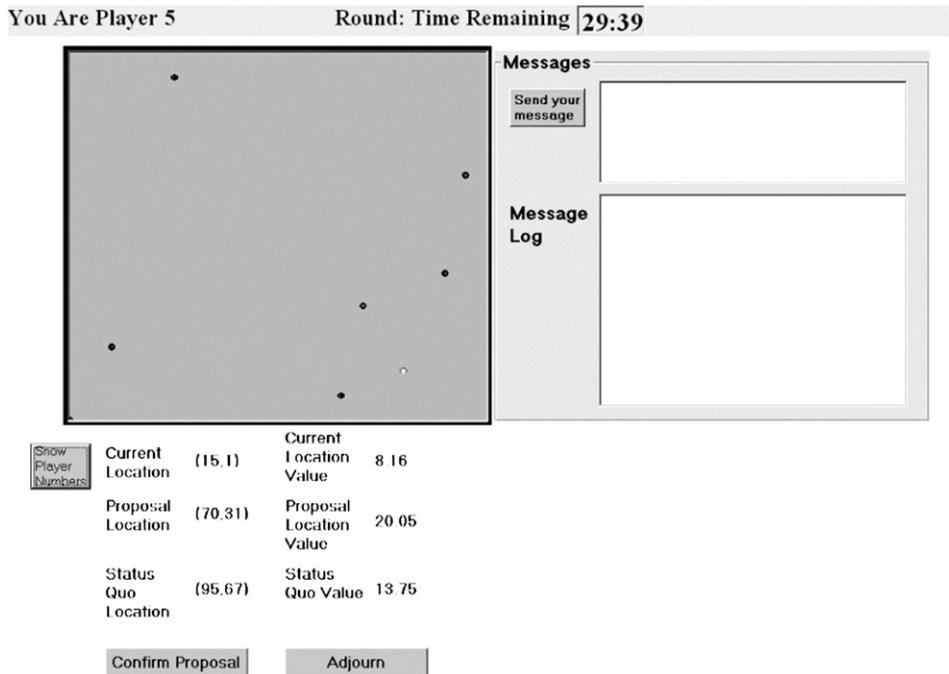
Fig. 6 Sample work screen for small-N experiments.

had not been tested against the experimental data "because it is extremely difficult to identify the region encompassed by the uncovered set, even in simple cases where just five legislators with circular indifference curves evaluate policies in two dimensions" (1994: 134). Green and Shapiro clearly regard this as a very telling argument against what they call the "spurious formal precision" of rational choice theory (1994:134).

Not only is the uncovered set testable, the evidence in these experiments is largely supportive of the uncovered set. Furthermore, we can say something about *why* majority rule tends to lead to outcomes in the uncovered set. The predictability in the majority rule decision process is built into the negotiation process itself. The borders of the uncovered set seem to denote the limits of how far the negotiation process can profitably be pursued in the direction of one or another pivotal player. The uncovered set thus describes how the negotiation process constrains the outcomes generated by majority rule instability.

The predictability of the uncovered set seems, on the basis of these experiments, relatively robust. In the small-N experiments, the results were more tightly distributed in S2, with a smaller uncovered set. The success of the uncovered set in both large-N and small-N experiments is striking, given the relative absence of communication and strict limitations on proposal making in the large-N experiments.

Finally, our experiments indicate that majority rule, in the absence of formal institutional constraints on agenda control, is not inevitably chaotic. An open agenda leads to outcomes that are relatively centrist and responsive to individual preference changes. The strategic process of formulating proposals, and forming coalitions in support of those proposals, is sufficient to generate predictable and moderate outcomes, as captured in the uncovered set.
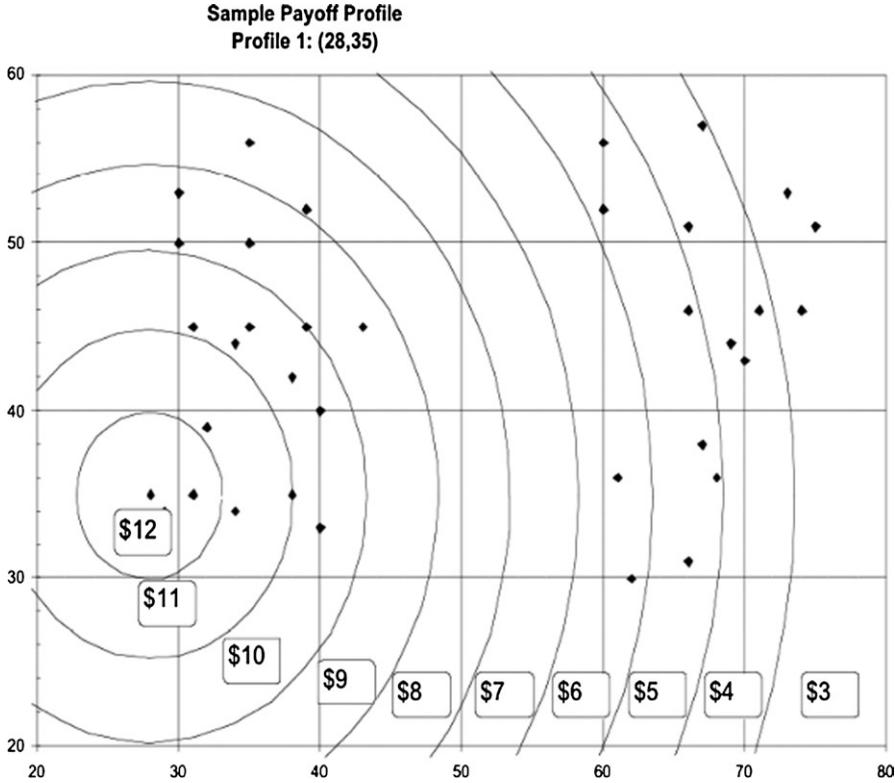
Sample Payoff Profile
Profile 1: (28,35)

Fig. 7 Sample worksheet for large-N.

Appendix: The Experiments

Payoff Functions

Payoffs for the players in the S1 design are as follows:

Player 1: $18.00 − $0.18[(50 − X)^2 + (85 − Y)^2]$
Player 2: $15.75 − $0.13[(25 − X)^2 + (93 − Y)^2]$
Player 3: $15.75 − $0.13[(10 − X)^2 + (80 − Y)^2]$
Player 4: $15.75 − $0.13[(65 − X)^2 + (7 − Y)^2]$
Player 5: $15.75 − $0.13[(80 − X)^2 + (13 − Y)^2]$

Payoffs for the S2 are as follows:

Player 1: $19.25 − $0.18[(70 − X)^2 + (20 − Y)^2]$
Player 2: $15.75 − $0.13[(25 − X)^2 + (93 − Y)^2]$
Player 3: $15.75 − $0.13[(10 − X)^2 + (80 − Y)^2]$
Player 4: $15.75 − $0.13[(65 − X)^2 + (7 − Y)^2]$
Player 5: $15.75 − $0.13[(80 − X)^2 + (13 − Y)^2]$

For the large-N design, payoffs for a player are computed as 0.20 times the distance (in the units shown in Fig. 1) between the player's ideal point and the realized outcome.

Instructions to Participants in Small-N Experiment

You are about to participate in a decision-making experiment in which one of numerous competing alternatives will be chosen by majority rule. The purpose of this experiment is to gain insight into certain features of complex political processes. The instructions are simple. If you follow them carefully and make good decisions, you might earn a considerable amount of money. You will be paid in cash.

The alternatives are represented by a 100×100 grid, as shown on the blackboard. Each alternative has an *X* coordinate, indicating how far to the left or right it is on the blackboard. Each point has a *Y* coordinate indicating how high or low the point is on the blackboard. The point ($X = 0, Y = 0$) would be as far down and to the left possible on the blackboard. The point ($X = 100, Y = 100$) is as far up and to the right as possible.

Using majority rule, each five-person committee will adopt as their decision one and only one point. Your compensation depends on the particular point chosen by a majority of your committee. For example, suppose your payoff chart is that given on the blackboard. Your highest compensation would be given by the point at the center of the rings, called your "highest payoff point." You would be paid less, the farther the chosen point is from your highest payoff point. You can imagine that there are concentric rings around your highest payoff point, representing equally valued points. Larger rings are worth less than the smaller rings closer to your highest payoff point. For instance, the point ($X = 70, Y = 50$) can be seen to be exactly on the circular line labeled $8000. If your compensation was the same as that shown, and a majority of the participants voted for that point, you would earn $8000. These payoffs are just illustrative.

Your own compensation chart will be different than the one on the blackboard. Each participant's compensation chart will be different. This means that the patterns of preferences differ and the monetary amounts may not be comparable. The point that would result in the highest payoff to you may or may not result in the highest payoff to someone else. You should decide what choice you want the players to make and do whatever you wish within the confines of the rules to get things to go your way. The experimenters, however, are not primarily concerned with whether or how you participate so long as you stay within the confines of the rules. Under no circumstances, may you mention anything *quantitative* about your compensation. You are free, if you wish, to indicate which ones you like best, etc., but you cannot mention anything about the actual monetary amounts. Under no circumstances may you mention anything about activities that might involve you and other players after the experiment; that is, no deals to split up afterward and no physical threats.

[The practice period will be started here.]

At this point, we will turn on the computer for a practice decision. You will not be paid for the decision made during the practice period. You will be notified when the practice period is over, and when the real decision period begins.

Please turn to your computer. Your computer screen shows your *sample* compensation chart for the practice exercise. (Your real payoff chart will be different when the practice period is over.) The computer shows your highest payoff point as a *white* dot. Points that are farther from this white dot will result in a lower cash payoff to you than points that are closer to this white dot.

By using the mouse, you may move the cursor to any point in the space. When you do so, the table below the payoff space will show you the coordinates of the point where the cursor is; this is shown as "current location." Next to it, the table on the screen indicates

"current location value"; this is the monetary payoff to you, of the point on the screen where your cursor is currently located. That is, if the committee were to choose that point, you would be paid that amount in cash. In this way, you can readily determine the exact value of each alternative to you. Please practice this by moving your cursor close to the point ($X = 10$, $Y = 10$) and seeing how much that alternative would be worth if the committee was to choose that point.

Write what the payoff to you (for this practice period) if your committee were to select ($X = 10$, $Y = 10$).

_____

Now find out how much your highest payoff point for the practice period would be worth to you, by moving the cursor to the white dot. *Write here the approximate coordinates of your highest payoff point for this practice session:*

$X =$ _____ $Y =$ _____

The computer screen also shows as black dots the highest payoff points that would give the most possible money to each of the other four members of the committee. By clicking the button at the left, you can see which of these points are the highest payoff points for which player numbers. Using the mouse, please click on the box on your screen labeled "show player numbers"; then click again so that the player numbers disappear. You may have the computer show the other players' player numbers or not, as you choose. I recommend that you leave it on, so you will know the "location" of anyone who messages you.

All the players' highest payoff points, including yours, will change when the practice period is over, and the real decision period begins.

You may communicate with the other members of the committee at any time. You can do this *only* through the messaging segment of your screen. You may compose a message by typing on your keyboard and using the mouse to click on the "send your message" button. Please do *not* hit the enter key, as you would on an instant messaging system. Send the message only by means of the mouse and the "send message" button on the screen. Please practice now by sending the statement, "I am player #____."

Has everyone had a chance to send a message with your player number? You may send messages at any time.

During the meeting, decisions are made by simple majority rule among the five members of your committee. Debate is open to all participants. Any player may make informal suggestions or comments at any time.

The process begins with an existing status quo. For the practice session, the initial status quo is the point ($X = 95$, $Y = 95$). This point is shown as a red dot on your screen. If no majority agrees to change this, that will be the point that will determine your compensation. That is, you will be paid the amount of money that is associated with the red dot on your screen. Please use your mouse to move the cursor to the red dot on the screen to see how much the status quo would be worth if the committee adjourns without changing the status quo.

How much would you be paid if your committee makes no successful changes in the status quo?

_____

At any given time, *one* member of each committee will be authorized to make a formal proposal to change the status quo. The person who is authorized will have a blue policy screen. Others will have a pink screen. If you have a pink screen, you can make informal

proposals or discuss alternatives by your messaging system. But only the person with a blue screen may make a formal proposal.

One person from each committee will be randomly assigned to have the first blue screen. That person's screen will remain blue for 90 s. During that time, this person may make a proposal or not, as he or she chooses. During that time, the person with the blue screen (whom we will call the *proposer*) can click his cursor on any point. If the proposer wants to have the committee vote on that point, the proposer can then use his cursor to click on the "confirm proposal" button. The proposer's computer screen will indicate the coordinates of this proposed point and the value of the proposed point to the proposer. At this point, the computer will immediately ask the proposer whether or not he wants to confirm his proposal. The proposer can change his mind or he can confirm that he wants to make the proposal.

If the proposer makes a proposal, the proposal will appear on everyone's computer screen as a *blue* dot. As soon as a proposal is made, there will be an immediate voting opportunity. Everyone in the committee will be asked to vote either for the existing *status quo* (the red dot) or the *new proposal* (the blue dot). You can see from the computer screen how much payoff will be associated with each of these outcomes.

If you abstain, that will be counted as a vote against the new proposal. If there are three votes for the new proposal, the new proposal passes and becomes the new status quo. That means that if the exercise were to end at this point, you would be paid on the basis of the successful new proposal, now called the *new* status quo.

The computer will randomly give everyone an opportunity to be the proposer, one after the other. After everyone has had a turn, the computer will create a new random order for proposers, and so on until adjournment.

Someone should now have a blue screen. Will that person move their cursor to outcome (10, 10) and propose that point? You will do so by using the mouse to move the cursor to the (10, 10), then clicking the mouse. The coordinates and payoff value to you of the proposed location will show up on the screen. To formally propose that point, you should click on the "propose" box and then *confirm* that you wish to propose it. If you see the blue screen and have questions about how to make this proposal, please feel free to raise your hand.

When you see the voting box, please vote either for the status quo or the new proposal. We will continue with the practice session until everyone has the opportunity to make at least one proposal. When you see the blue screen, either make a proposal or raise your hand to ask for assistance.

Players may propose and pass as many proposals as they wish. Each time the committee votes for a new proposal to replace the prior status quo, the status quo changes to the newly approved proposal. But the members of the committee will be paid only on the basis of the status quo at the time of adjournment.

The proposal process will continue until a motion to adjourn the meeting is made by a proposer and approved by a majority of the committee. A motion to adjourn may be made by any player while he has the blue screen. When a motion to adjourn is made, then it will be voted on immediately. You will have 30 s to vote. An abstention will count as a vote *against* adjournment. The motion to adjourn will pass if at least three players vote to adjourn.

If the motion to adjourn passes, then the voting exercise is over. It is important to remember that you will be paid only on the basis of the status quo (the last successful proposal) at the time of adjournment**.** Payoffs associated with earlier successful proposals do not affect your compensation for this experiment.

If you have the blue screen now, will you please make a proposal to adjourn? When you get an opportunity to vote, please vote to adjourn. That will end the practice session and we will be able to begin the real experiment, using your real compensation charts.

To review, the person with the blue screen may make a proposal to change the status quo, may propose to adjourn, or do neither. Your compensation will be determined by the status quo *at the time of adjournment*. That is, the last successful proposal will determine how much you are paid. *Everyone* will be paid based on the status quo, *not just the members of the coalition who supported that alternative.*

Are there any questions before we begin the real decision exercise?

The real decision exercise is now beginning with a status quo at the point ($X = 75, Y = 75$). As before, one person at a time will have a blue screen, for 90 s, indicating proposal powers. Proposals may be approved by a majority rule of the committee, and an approved proposal will become the new status quo. You will be paid a cash amount equal to the value of the status quo at the time of adjournment.

The committee will have up to 20 min to make and vote on as many proposals as it wishes and to adjourn. If the committee has not adjourned by the time 20 min is up, then the members of the committee will be paid on the basis of the status quo when time ran out.

If your committee is the first to adjourn, we will ask you to sit quietly until all other committees have adjourned as well.

Large-N Design Instructions

You are about to participate in a decision-making experiment in which one of numerous competing alternatives will be chosen by majority rule. The purpose of this experiment is to gain insight into certain features of complex political processes. The instructions are simple.

The alternatives that the committee can vote on are represented by points on the sheet in front of you. Each point has an *X* coordinate indicating how far to the left or right it is on the sheet. Each point has a *Y* coordinate indicating how high or low the point is on the sheet. The point ($X = 20, Y = 20$) would be as far down and to the left possible on the sheet. The point ($X = 80, Y = 60$) is as far up and to the right as possible.

Take a look at the figure on the overhead. It shows your highest payoff point as a black dot in the center of the rings, at the point ($X = 28, Y = 35$). Points that are farther from this black dot will result in a lower cash payoff to you than points that are closer to this black dot. We have drawn rings on the figure that show how your payoff declines as you move farther and farther away from your highest payoff point. The dollar figures give the payoffs for outcomes that are located in each ring. Note that the larger the ring, the lower the payoff. The figure also shows the highest payoff points that would give the most possible money to each of the other members of the committee.

For example, suppose your payoff chart is that given in the top figure in your packet.

- Your highest compensation, a payoff of $12.00 would be given by the point at the center at the rings, called your highest payoff point, which is ($X = 28, Y = 35$).
- You would be paid less, the farther the chosen point is from your highest payoff point. Points that are in between circles have payoffs that are in between the payoffs for each circle.
- If the outcome (60, 40)—the starred outcome on your figure—is the final result, your payoff would be about $6.50.
- Any point that is outside the largest circle in the figure, such as 80, 60, will give you a payoff of $2.00.

Remember that this payoff profile, highest payoff point, and concentric circles are just illustrative. Your own compensation chart may be different than the one in the figure. Each participant's compensation chart will be different. In other words, the point that would result in the highest payoff to you will not result in the highest payoff to someone else, although some players have highest payoff points that are extremely close to each other.

Using majority rule, the committee (all of you) will adopt as their decision one and only one point (X, Y) on the figure. Your compensation depends on the particular point chosen by a majority of the participants. The process begins with an existing status quo, which is the point $(X = 50, Y = 60)$. Find this point on your payoff profile. If no majority agrees to change this proposal, that will be the point that will determine your compensation. That is, you will be paid the amount of money that is associated with this point on your payoff figure.

Following the discussion, one of the participants will be allowed to make a formal proposal to change the status quo. We will select this person randomly. After being selected, the proposer has a minute to make a proposal. The proposal will be a new point (X, Y) that if enacted will determine the payoffs to committee members based on their payoff profiles.

If the proposer offers a proposal, there will be an immediate vote. Everyone in the committee will be asked to vote either for the existing *status quo* or the *new proposal*. You can see from your payoff figure what your payoff would be if either of these proposals is the final result of committee deliberations.

For example, suppose the initial status quo is $X = 50, Y = 60$, and the new proposal is $(X = 50, Y = 40)$. Based on the figure in front of you, your payoff from the status quo is about \$6.30, while your payoff from the new proposal is \$7.40.

The vote will be by a show of hands. If more people vote for the proposal than against it, the new proposal passes and becomes the new status quo. If you abstain, that will be counted as a vote against the new proposal. That means that if the exercise were to end at this point, you would be paid on the basis of the successful new proposal, now called the status quo.

After voting on a proposal (or if the proposer does not offer a proposal), there will be an opportunity to adjourn the proceedings. Any participant can raise their hand and offer a motion to adjourn. There will be an immediate vote using the same procedure as for proposals. If a majority votes yes, the proceedings will end. Your payoff will be based on your payoff profile and the status quo proposal.

If the motion to adjourn passes, then the voting exercise is over. It is important to remember that you will be paid only on the basis of the status quo (the last successful proposal) at the time of adjournment. Payoffs associated with earlier successful proposals do not affect your compensation for this experiment. So please remember that your payoff is determined by the status quo at the time that the committee votes to adjourn. *Everyone* will be paid based on the status quo, *not just the members of the coalition who supported that alternative.*

If no one offers a motion to adjourn, or if a motion is offered but does not gain majority support, we will select another proposer, and so on until there is a motion to adjourn that gains majority support.

At this point, we will have a practice decision. You will not be paid for the decision made during the practice period. You will be notified when the practice period is over, and when the real decision period begins.

Please turn to the next figure. It shows your *sample* compensation chart for the practice exercise. (Your real payoff chart will be different when the practice period is over).

Funding

References

Bianco, William T., Ivan Jeliazkov, and Itai Sened. 2004. The uncovered set and the limits of legislative action. *Political Analysis* 12:256–76.

Bianco, William T., Michael S. Lynch, Gary J. Miller, and Itai Sened. 2006. 'A theory waiting to be discovered and used': A reanalysis of canonical experiments on majority rule decision-making. *Journal of Politics* 68:837–50.

Cox, Gary W. 1987. The uncovered set and the core. *American Journal of Political Science* 31:408–22.

Cox, Gary W., and Mathew D. McCubbins. 2005. *Setting the agenda: Responsible party government in the US House of Representatives*. Cambridge: Cambridge University Press.

Epstein, David. 1998. Uncovering some subtleties of the uncovered set: Social choice theory and distributive politics. *Social Choice and Welfare* 15:81–93.

Feld, Scott L., Bernard Grofman, Richard Hartley, Marc Kilgour, Nicholas R. Miller, and Nicolas Noviello. 1987. The uncovered set in spatial voting games. *Theory and Decision* 23:129–55.

Fiorina, Morris P., and Charles R. Plott. 1978. Committee decisions under majority rule: An experimental study. *American Political Science Review* 72:575–98.

Green, Donald P., and Ian Shapiro. 1994. *Pathologies of rational choice theory: A critique of applications in political science*. New Haven: Yale University Press.

Hartley, Richard, and D. Marc Kilgour. 1987. The geometry of the uncovered set. *Mathematical Social Sciences* 14:175–83.

Johnson, N. L., and S. Kotz. 1972. *Distributions in statistics: Continuous multivariate distributions*. New York: Wiley.

McKelvey, Richard D. 1976. Intransitivities in multidimensional voting models and some implications for agenda control. *Journal of Economic Theory* 12:472–82.

———. 1979. General conditions for global intransitivities in formal voting models. *Econometrica* 47:1085–112.

———. 1986. Covering dominance and institution free properties of social choice. *American Journal of Political Science* 30:283–314.

McKelvey, Richard D., and Peter C. Ordeshook. 1983. Some experimental results that fail to support the competitive solution. *Public Choice* 40:281–91.

———. 1990. A decade of experimental research in spatial models of elections and committees. In *Advances in the spatial theory of voting*, ed. James M. Enelow and Melvin J. Hinich, 99–144. Cambridge: Cambridge University Press.

McKelvey, Richard D., and Norman Schofield. 1987. Generalized symmetry conditions at a core. *Econometrica* 55:923–33.

McKelvey, Richard D., Peter C. Ordeshook, and Mark D. Winer. 1978. The competitive solution for N-person games without transferable utility, with an application to committee games. *American Political Science Review* 72:599–615.

Miller, Nicholas. 1980. A new solution set for tournament and majority voting. *American Journal of Political Science* 24:68–96.

Ordeshook, Peter C., and Thomas Schwartz. 1987. Agendas and the control of political outcomes. *American Political Science Review* 81:179–200.

Penn, Elizabeth M. 2006. Alternatives definitions of the uncovered set and their implications. *Social Choice and Welfare* 27:83–7.

Riker, William H. 1980. Implications from the disequilibrium of majority rule for the study of institutions. *American Political Science Review* 74:432–46.

Schofield, Norman. 1978. Instability of simple dynamic games. *Review of Economic Studies* 45:575–94.

Shepsle, Kenneth A. 1979. Institutional arrangements and equilibria in multidimensional voting models. *American Journal of Political Science* 23:27–59.

———. 1986. Institutional equilibrium and equilibrium institutions. In *Political science: The science of politics*, ed. Herbert F. Weisberg, 51–81. New York: Agathon Press.

Shepsle, Kenneth A., and Barry R. Weingast. 1984. Uncovered sets and sophisticated voting outcomes with implications for agenda institutions. *American Journal of Political Science* 28:49–74.

Tullock, Gordon. 1981. Why so much stability? *Public Choice* 37:189–202.