


DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES
CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

STRATEGIC MANIPULABILITY IS INESCAPABLE: GIBBARD-SATTERTHWAITE
WITHOUT RESOLUTENESS

John Duggan
California Institute of Technology

Thomas Schwartz
University of California, Los Angeles



SOCIAL SCIENCE WORKING PAPER 817

November 1992

STRATEGIC MANIPULABILITY IS INESCAPABLE: GIBBARD-SATTERTHWAITE WITHOUT RESOLUTENESS

John Duggan

Thomas Schwartz

Abstract

The Gibbard-Satterthwaite Theorem on the manipulability of collective-choice procedures treats only of *resolute* procedures. Few real or reasonable procedures are resolute. We prove a generalization of Gibbard-Satterthwaite that covers the nonresolute case. It opens harder questions than it answers about the prediction of behavior and outcomes and the design of institutions.

STRATEGIC MANIPULABILITY IS INESCAPABLE: GIBBARD-SATTERTHWAITE WITHOUT RESOLUTENESS*

John Duggan

Thomas Schwartz

A voting rule, market economy, management hierarchy, or other *collective choice procedure* turns the professed preferences of several individuals into one collective choice — one chosen alternative or one choice set of two or more “tied” alternatives. It is *strategy free* or *nonmanipulable* if it ensures that no one acting alone can ever gain a preferred outcome by misrepresenting his preferences — voting for a second-best candidate, for example, to stop a worse one from winning, or introducing a disliked amendment to kill a disliked bill, or dumping many shares of a prized stock to depress its price. The celebrated theorem of Gibbard (1973) and Satterthwaite (1975) is often said to show that any strategy-free procedure for choosing among three or more alternatives must be dictatorial. In one way the theorem is stronger than this summary: it shows that a nondictatorial procedure must be manipulable, not only for some set of three or more feasible alternatives, but for any given set. In another way the theorem is weaker: it is not about all collective-choice procedures but only *resolute* ones. They choose single alternatives in every possible circumstance, never allowing ties.

Most collective-choice procedures of any importance are *nonresolute*, we argue (§1), and the Gibbard-Satterthwaite Theorem does not show that they are manipulable except in special circumstances. Some published manipulability theorems do not assume resoluteness, but in other ways their generality is limited. We contrast these theorems with a new one (§2), a generalization of Gibbard-Satterthwaite that covers the nonresolute case while avoiding the limitations of earlier results. Before stating (§4) and proving (§5) this theorem, we explain it in detail but informally (§3). It opens harder questions than it answers about the incidence of strategic behavior, the existence and predictive value of equilibria, and the design of institutions (§6).

*We thank Peter Ordeshook, Thomas Palfrey, Charles Plott, and Martin van Hees for helpful discussions. Schwartz thanks the UCLA Senate for research support.

1 Nonresolute Procedures and the Gibbard-Satterthwaite Theorem

Outside the two-alternative case, few collective-choice procedures found in practice or in the imaginations of theorists and reformers are resolute: most allow ties, or multi-member choice sets. Ties are allowed by all familiar election rules (Plurality, Runoff, Borda, List PR, STV, SNTV, etc.), welfare criteria (maximum average welfare, maximin, and whatnot), and preference-sensitive taxation schemes (Lindahl, Clarke, Groves-Ledyard). Market economies and other exchange mechanisms also allow ties (multiple competitive equilibria, multi-member contract sets and cores). So must any procedure that fulfills the democratic ideals of anonymity (individuals count equally) and neutrality (alternatives count equally).

Parliamentary procedure is an apparent exception. It turns any multi-alternative contest into a series of *yes-or-no* votes. At each vote, the *no* option wins by default if the *yes* option lacks the required majority. That ensures a unique final winner. But to see this procedure as resolute is to see the agenda on any occasion — the series of *yes-or-no* votes — as part the procedure, making “the procedure” a host of *ad hoc* procedures that differ between any two occasions of choice. If “the procedure” is seen instead as one that endures from occasion to occasion of choice, then the choice set on any occasion comprises every alternative choosable under some agenda then permissible, each in effect a tie-breaker. Seen that way, parliamentary procedure is especially irresolute — it often yields large choice sets — even when the alternatives to appear on the agenda are specified in advance (Ordeshook and Schwartz 1987).

Which is the right way to see parliamentary procedure? Neither is uniquely right. We climb a ladder of resolution when we start with general parliamentary procedure and add reporting committees, scheduling and recognition rules, and finally a specific agenda: each addition yields a more resolute procedure, one that narrows the choice sets of its predecessor. We do the same when we start with the constitutional requirements for some elective office and add statutes governing the details of balloting, nominations, constituency divisions, and the like. And we do the same when we start with a whole constitution and add legislative rules, ministries, standing programs, and so forth, or when we start with a code of property law and add contracts, courts, and money, or again when we start with a corporate charter and add a table of organization, hiring policies, and whatnot. Climb to the top of any ladder of resolution and you will see either a resolute procedure or one that resolves some residual ties by chance or historical accident. But you are free to stop before the top and see the part climbed as a collective-choice procedure and the part above as a tie-resolving mechanism for that procedure.

The question of manipulability can be posed at any rung, but often it is more interesting at lower rungs, where the choice sets that might be changed by manipulation are bigger, and changes in them more consequential. Take the February-to-December procedure for electing a U.S. president: state primaries (rung 1), national conventions (rung 2), general election (rung 3), formal vote of the Electoral College (rung 4). A manipulation of this four-rung procedure might be more discomfiting, or anyway more surprising, than a manipulation of the procedure (rung 5) by which the House of Representatives resolves Electoral College ties.

Gibbard emphasizes that his theorem says nothing about procedures involving chance. That is an understatement: the theorem It also says nothing about low-rung procedures. Does it imply more than it says? Yes, but not very much.

Standard election rules are *near* resolute: electoral ties are rare and rarely large. Although resoluteness bans all possible ties, these procedures fall so little short of resoluteness, one might contend, that nothing much can hang on the shortfall. But see how much is hung on little shortfalls by the Gibbard-Satterthwaite Theorem: to block all possible opportunities for manipulation, a nondictatorial procedure for choosing among three or more alternatives must deviate at least a bit from perfect resoluteness. Just a little bit may be enough, and the opportunities that have to be blocked may be very few, so far as we know from the theorem.

When multi-member sets are chosen, it is tempting to count them as “alternatives” on all fours with their members (Barberá, Sonnenschein, and Zhou 1991). But Gibbard-Satterthwaite still does not apply because it assumes unrestricted preference orderings of whatever qualify as “alternatives,” in this case allowing someone who prefers $\{x\}$ to $\{y\}$ to $\{z\}$ also to prefer $\{y, z\}$ to $\{x\}$.

Yes, a nonresolute procedure must be combined with some tie-resolving mechanism. But suppose that this mechanism is stochastic or partly so: it picks lotteries, not all of them degenerate. Then Gibbard-Satterthwaite does not apply to the combined procedure, much less the original one, unless we assume unrestricted preference orderings of lotteries, and that is no more reasonable than unrestricted preference orderings of sets. So suppose instead that the tie-resolving mechanism is completely deterministic. Let there be three or more feasible alternatives and no dictator. Now Gibbard-Satterthwaite implies that the combined procedure is manipulable. But it is the original, lower-rung procedure whose manipulability may be of interest. Does Gibbard-Satterthwaite imply that it is manipulable? Not if the tie-resolving mechanism depends on the professed preferences of individuals: for all we know from the theorem, a manipulator of the combined procedure can change the final outcome from one to another member of the pre-resolution choice set but can never change that set. If, on the other hand, the tie-resolving mechanism does not depend at all on professed preferences, then a manipulator of the combined procedure must change the pre-resolution choice set to change the final selection from

that set, so the original procedure is manipulable once the tie-resolving mechanism is in place. But that follows from Gibbard-Satterthwaite only in this rather special case: the tie-resolving mechanism is completely independent of professed preferences yet completely deterministic.

2 What is Known and What is to be Shown

Beginning in the late eighteenth century, students of voting have learned that all familiar voting rules, none of them resolute, are manipulable when the feasible alternatives are three or more. Do we know anything more general, anything like the Gibbard-Satterthwaite theorem shorn of its resoluteness assumption?

This question raises another. Strategy freedom bans all possible cases of manipulation. When collective choices are single alternatives and preferences are represented by orderings of alternatives, the banned cases are of this type:

A change in Mr. i 's professed ordering of alternatives, all else remaining the same, changes the collective choice from an alternative x to an alternative y though y is preferred to x according to i 's original ordering.

But when collective choices are *sets* of alternatives, what does it mean for one to be preferred to another according to an ordering of *alternatives*? Four approaches have been taken to this problem.

Where X and Y are potential choice sets, the Heroic Approach ascribes a preference for Y over X to individual i only when such a preference is plainly necessitated by i 's ordering of alternatives. For example, someone who prefers x to y to z must prefer $\{x\}$ to $\{y, z\}$ and $\{x, y\}$ to $\{z\}$ but not necessarily $\{y\}$ to $\{x, z\}$ or $\{x, z\}$ to $\{y\}$. Kelly (1977) and Barberá (1977a,b) prove manipulability theorems based on this approach, but they pay for weak preferential assumptions with strong procedural ones. Barberá (1977b) assumes the acyclicity and Kelly the transitivity of (strict) social preference, and Barberá (1977a) assumes "strict monotonicity," a property exemplified by few voting rules — by Borda, but not, for example, by Plurality, Runoff, or parliamentary procedure.

The Maximin Approach assumes that i prefers Y to X only if i prefers the worst alternative in Y (according to his preference ordering) to the worst in X . Pattanaik (1978) uses this approach to prove the manipulability of a number of important types of procedure. But those types are quite specific, defined by strong "democratic" requirements. Also Pattanaik shares with Kelly a formal framework which differs from

Gibbard-Satterthwaite's and Barberá's in a way that weakens his results: instead of assuming a fixed set of feasible alternatives, he lets the feasible set vary over all or many subsets of some given set, leaving open the possibility that manipulability afflicts few feasible sets.

The Set Preference Approach represents preferences to begin with by orderings that rank *sets* of alternatives, though it does not count all such orderings as "admissible": an admissible ordering might rank $\{x\}$ over $\{x, y\}$ over $\{y\}$, for example, but not $\{x\}$ over $\{y\}$ over $\{x, y\}$. Schwartz (1982) uses this approach to deduce manipulability from mild procedural assumptions. But he makes a host of opaque assumptions about admissible orderings of sets, and he too lets the feasible set vary.

Finally, the Possible Preference Approach represents preferences by orderings of alternatives but assumes that strategy freedom bans all cases of the following type:

A change in i 's professed ordering, all else remaining the same, changes the collective choice from a set X to a set Y though a preference for Y over X is *compatible* with i 's original beliefs and ordering — though it would be possible if not compulsory for a rational individual who had i 's original beliefs and ordering to prefer Y to X .

What makes a procedure manipulable, after all, is not that it is actually manipulated but that it fails to make manipulation impossible. Taking this approach, Zeckhauser (1973) and Gibbard (1977) hit nearer our target of Gibbard-Satterthwaite without resoluteness. They prove manipulability for nonresolute procedures and fixed feasible sets, assuming that a preference for Y over X is compatible with i 's original beliefs and ordering if a certain lottery over Y has a greater expected utility than a certain lottery over X for some utility function compatible with i 's ordering. But besides adding some hard-to-interpret procedural assumptions, Zeckhauser and Gibbard assume that lotteries are procedurally determined: in effect, the procedure picks both a choice set and the probability distribution used by all individuals to evaluate that set. That is reasonable in most electoral contexts, where ties are resolved randomly. It is not so reasonable in constitutional, parliamentary, market, management, and other low-rung contexts, where typical choice sets are large and individuals differ in their beliefs about the resolution of ties, about what happens higher on the ladder, sometimes making strategic use of those differences. Think of legislators who profit from knowing better than their colleagues what motions will be recognized in what order. Or think of voters in a U.S. presidential election who have different beliefs about how the House of Representatives would resolve an Electoral College tie. Zeckhauser and Gibbard cover top-rung procedures that use chance to resolve residual ties, but they descend the ladder of resolution only in worlds marked by an unusual coincidence of beliefs.

We too take the Possible Preference Approach and prove a generalization of Gibbard-Satterthwaite that covers the nonresolute case without assuming any common beliefs about the resolution of ties. On the procedural side, we follow Gibbard and Satterthwaite by assuming a fixed set of three or more alternatives, each feasible in the “citizens’ sovereignty” sense that there is some way to secure its choice, and a fixed population of individuals, none a dictator, with variable preference orderings of alternatives. On the preferential side, we follow Zeckhauser and Gibbard by assuming that an individual compares potential choice sets as if comparing lotteries according to their expected utilities for some utility function compatible with his ordering of alternatives. But we generalize their framework by letting each individual i have his own lottery over every potential choice set X . And we let this lottery depend, not only on i and X , but on i ’s true preference ordering and the professed preferences of others. Given these factors, the corresponding lottery is *arbitrary but fixed*: our theorem shows that manipulation is possible, not merely for *some* set of lotteries (or beliefs) of the fancied sort, but for *any* given set.

3 Explanation of Theorem

Like Gibbard and Satterthwaite (G&S), we assume a fixed population of individuals, Messrs. $1, 2, \dots, n$, and a fixed set A of three or more feasible alternatives. Unlike G&S, we explicitly assume that A is finite. Actually it is enough that choice sets be finite, as G&S assume, but then we might as well make the simpler assumption that A is finite. The realism of this assumption would be hard to contest: the candidates for any office are finite in number, as are the policy options of any sort if each is formulable from some given finite stock of symbols (those on your keyboard, say) in a string no longer than some given physical constraint (the number of quarks in the Milky Way, say).

We copy G&S’s version of citizens’ sovereignty: every member of A is the unique collective choice in some possible circumstance. The weaker assumption that every member of A belongs to the choice set in some possible circumstance is nigh impotent: it is automatically true if there is no real choice at all in some odd circumstance, the choice set being A itself. Our G&S version captures the idea that every so-called feasible alternative really is feasible, not only in the weak sense that Messrs. $1, 2, \dots, n$ can somehow *permit* its choice, but also in the strong sense that they can somehow *secure* its choice.

G&S’s version of nondictatorship says there is no individual whose professed favorite alternative in every possible circumstance is the unique collective choice. For nonresolute procedures, that is vacuous, a consequence of nonresoluteness itself. Our version says there is no individual whose professed favorite in every possible circumstance belongs to the choice set. This is stronger than it may look. Besides dictatorships of the usual sort, it bans the Collective Hamlet Rule, whose choice set in any circumstance contains

everyone's professed favorite. More important, both versions ban dictators of some sort for the given set A , yet some reasonable procedures allow dictators for certain sets, those representing issues that affect one individual's rights. Thus, we must think of A as not being such a set (if it were then strategy freedom would be assured).

It is not always clear what to count as an expression of someone's preference, nor when a given expression is candid (Gibbard 1973, Pattanaik 1978). But for us as for G&S, it is enough that every possible expression be uniquely determined by some possible preference ordering of alternatives, that there exist a mapping from possible preference orderings to possible individual strategies. Unlike G&S, we assume that all such orderings are *linear*: they never rank two alternatives at the same level. Like G&S, we assume that these orderings are otherwise unrestricted, singly and in combination. The "possible circumstances" of the previous two paragraphs are all the *profiles*, or ordered n -tuples of linear orderings of A . While barely strengthening nondictatorship and citizens' sovereignty, our linear restriction greatly strengthens our conclusion. If manipulation is possible in the universe of linear preferences, then of course it is possible in any larger universe which contains that one. But its possibility in the linear universe shows that it is possible under those preferential voting systems that require ballots to express linear orderings (most do) and also those nonpreferential systems that require ballots to designate single "favorite" alternatives, either once or at each of several stages of voting (most do): only linear orderings uniquely determine such ballots. It shows as well that strategic misrepresentation is not just the advantageous but arbitrary resolution of subjective "ties."

Our version of strategy freedom bans all cases of the following type:

A change in Mr. i 's professed ordering of alternatives, all else remaining the same, changes the collective choice from a set X to a set Y though the lottery that i associates with Y and the original profile has a greater expected utility than the lottery that i associates with X and the original profile for some utility function compatible with i 's original ordering.

The lottery, or probability distribution, that i associates with a set X and profile \mathbf{v} summarizes his beliefs about how the tie would be resolved if the choice set were X , the i th ordering in \mathbf{v} were i 's true preference ordering, and the others were the professed orderings of the other $n - 1$ individuals. This lottery assigns positive probabilities only to members of X , of course, but maybe not to all members. It is enough that the best and worst members of X have positive probabilities: i is not so pessimistic or optimistic that he ignores these possible outcomes.

This version of strategy freedom is quite general in that individual lotteries are min-

imally constrained but fixed: our theorem holds for *any given* set of lotteries of the assumed sort. It is quite general as well in that it allows a potential strategist's beliefs about the resolution of ties to be similar or not to those of other individuals, procedurally determined or not, well-informed or not, dependent or not on his own preference ordering, and sensitive or not to the professed preferences of others (he cannot directly observe their *true* preferences), hence based or not on his assessment of the behavior of others in the tie-resolving process. Naturally there are limitations. Mr. i 's lotteries may evince some risk aversion or risk acceptance: because they depend on i 's own preference ordering, he can assign probabilities that rapidly increase or rapidly decrease down his ordering. But beyond that, i 's lotteries cannot depend on his utilities, or preference intensities, or preferences between lotteries themselves: i cannot simultaneously pick lotteries and utility functions to "fit" each other, as in the classical Ramsey-Savage framework, and we exploit this fact to prove the first lemma of §5.

The promised theorem says that strategy freedom is inconsistent with the assumptions sketched earlier in this section. In proving it, we first gain control over comparisons between potential choice sets by deducing, in effect, that an individual has a possible preference (one compatible with his preference ordering and beliefs) for one such set over another whenever he prefers the worst alternative in the one set to the worst in the other or the best in the one to the best in the other. Beyond that, our proof is somewhat like G&S's. Where they deduce their theorem from Arrow's (1963), we similarly exploit a variant of Arrow's Theorem that strengthens nondictatorship to nonblocker (or nonvetoer) and weakens transitivity of social preference-or-indifference to that of strict social preference.

4 Formal Statement of Theorem

Formally, our theorem is about a set A , positive integer n , unary function C , and ternary functions p_1, \dots, p_n . An *alternative* is any member of A . A *utility function* is any real-valued function on A . A *profile* is any ordered n -tuple of linear orderings of A , each a binary relation on A that is asymmetric, transitive, and connected in A (borne by one to the other of any two members of A). Denote alternatives by x, y , etc., nonempty sets of them by X, Y , etc., integers $1, 2, \dots, n$ by i, j , etc., utility functions by u, u' , etc., and profiles by $\mathbf{v} = (v_1, \dots, v_n)$, $\mathbf{v}' = (v'_1, \dots, v'_n)$, etc. Profiles \mathbf{v} and \mathbf{v}' are *i -variants* if $v_j = v'_j$ for all $j \neq i$. A is to be interpreted, of course, as the set of feasible alternatives, n as the number of individuals, C as the function that turns every \mathbf{v} into a choice set $C(\mathbf{v})$, and p_i as the function that turns every \mathbf{v} , X , and x into i 's assessment $p_i(\mathbf{v}, X, x)$ of the probability that x would be the final choice if X were the choice set, the i th ordering in \mathbf{v} were i 's true preference ordering, and the others were the professed orderings of the other $n - 1$ individuals.

Theorem. The following six conditions are inconsistent:

F3A A is a finite set of three or more objects.

CH C associates with every \mathbf{v} a nonempty subset $C(\mathbf{v})$ of A , called the *choice set* in \mathbf{v} .

PROB p_i associates an element $p_i(\mathbf{v}, X, x)$ of $[0, 1]$ with every \mathbf{v} , X , and x so that $\sum_{x \in A} p_i(\mathbf{v}, X, x) = 1$ and $p_i(\mathbf{v}, X, y) > 0 = p_i(\mathbf{v}, X, z)$ whenever y is the v_i -least member or the v_i -first member of X and $z \notin X$.

CiSov $C(\mathbf{v}) = \{x\}$ for some \mathbf{v} (citizens' sovereignty).

\mathcal{D} No i is such that, for all \mathbf{v} , x , if x ranks first in v_i then $x \in C(\mathbf{v})$ (nondictatorship).

\mathcal{S} If \mathbf{v} and \mathbf{v}' are i -variants then $\bar{u}_i^{\mathbf{v}'}(C(\mathbf{v}')) > \bar{u}_i^{\mathbf{v}}(C(\mathbf{v}))$ for no representative u of v_i (strategy freedom),

where u is a *representative* of v_i if and only if, for all x, y , $u(x) > u(y)$ just when xv_iy ,

and $\bar{u}_i^{\mathbf{v}}(X) = \sum_{x \in X} p_i(\mathbf{v}, X, x)u(x)$ (i 's expected u -value of X in \mathbf{v}).

5 Proof

To prove the inconsistency of these conditions, we first reduce our task to one of deducing from them that a certain function must meet six other conditions proved inconsistent by Fishburn (1973:128, or Mas-Colell and Sonnenschein 1972, or Schwartz 1986:59):

- 3A** A has three or more members.
- SoPREF** \mathcal{P} is a function that associates with every \mathbf{v} an asymmetric binary relation $\mathcal{P}^{\mathbf{v}}$ on A (a strict “social preference” relation).
- IIA** If $x\mathcal{P}^{\mathbf{v}}y$ and \mathbf{v}' is an xy -twin of \mathbf{v} then $x\mathcal{P}^{\mathbf{v}'}y$ (independence of irrelevant alternatives),
- where \mathbf{v}' is an xy -twin of \mathbf{v} if and only if, for all i , $xv'_iy \Leftrightarrow xv_iy$.
- PARETO** If xv_iy for all i then $x\mathcal{P}^{\mathbf{v}}y$.
- B** No i is such that, for all \mathbf{v} , x , y , if xv_iy while yv_jx for every $j \neq i$ then not $y\mathcal{P}^{\mathbf{v}}x$ (nonblocker).
- TRANS** If $x\mathcal{P}^{\mathbf{v}}y\mathcal{P}^{\mathbf{v}}z$ then $x\mathcal{P}^{\mathbf{v}}z$.

Because these conditions are inconsistent, no function \mathcal{P} of profiles can satisfy all of them, so *this* one cannot:

$x\hat{\mathcal{P}}^{\mathbf{v}}y$ if and only if $x \neq y$ and $\{x\} = C(\mathbf{v}')$ for every xy -twin \mathbf{v}' of \mathbf{v} in which $\{x, y\}$ is a top set,

where X is a *top set* in \mathbf{v} if and only if every member of X ranks above every member of $A - X$ in every v_i .

Hence, it suffices to deduce from our own six conditions that $\hat{\mathcal{P}}$ satisfies the six conditions just above.

From the definition of $\hat{\mathcal{P}}$ it follows immediately that $\hat{\mathcal{P}}$ satisfies **SoPREF** (the Definition obviously ensures asymmetry of $\hat{\mathcal{P}}^{\mathbf{v}}$) and **IIA** (by definition, $x\hat{\mathcal{P}}^{\mathbf{v}'}y$ holds for all or no xy -twins \mathbf{v}' of \mathbf{v}). And our **F3A** is obviously stronger than **3A**. Hence, it suffices to assume our own six conditions and prove that $\hat{\mathcal{P}}$ satisfies **PARETO**, **B**, and **TRANS**.

To this end we first milk **PROB** and **S** of all their useful content in a lemma.

Strategy Lemma. If $x \in C(\mathbf{v})$ and \mathbf{v}' is an i -variant of \mathbf{v} , then (1) x or something ranked lower in v_i belongs to $C(\mathbf{v}')$, and (2) x or something ranked higher in v'_i belongs to $C(\mathbf{v}')$.

Proof of (1). Suppose not. Then the v_i -least member y' of $C(\mathbf{v}')$ must rank higher

in v_i than the v_i -least member y of $C(\mathbf{v})$. Let $p = p_i(\mathbf{v}, C(\mathbf{v}), y)$ and y^H = the highest-ranked alternative in v_i . Then $p > 0$ by **PROB**, and $u(y^H) \geq u(y') > u(y)$ for every representative u of v_i . So some such u must make $u(y^H) - u(y')$ small enough and $u(y') - u(y)$ big enough to ensure this:

$$u(y') > pu(y) + (1 - p)u(y^H).$$

But

$$\bar{u}_i^{\mathbf{v}}(C(\mathbf{v}')) \geq u(y')$$

because $u(y') \leq u(w)$ for all $w \in C(\mathbf{v}')$, and

$$\bar{u}_i^{\mathbf{v}}(C(\mathbf{v})) \leq pu(y) + (1 - p)u(y^H)$$

because $pu(y)$ is a summand of $\bar{u}_i^{\mathbf{v}}(C(\mathbf{v}))$ and $u(y^H) \geq u(w)$ for all w . It follows that $\bar{u}_i^{\mathbf{v}}(C(\mathbf{v}')) > \bar{u}_i^{\mathbf{v}}(C(\mathbf{v}))$, contrary to \mathcal{S} .

Proof of (2). Suppose not. Then the v'_i -first member z of $C(\mathbf{v})$ must rank higher in v'_i than the v'_i -first member z' of $C(\mathbf{v}')$. Let $p = p_i(\mathbf{v}', C(\mathbf{v}), z)$ and z^L = the lowest-ranked alternative in v'_i . Then $p > 0$ by **PROB**, and $u(z) > u(z') \geq u(z^L)$ for every representative u of v'_i . So some such u must make $u(z') - u(z^L)$ small enough and $u(z) - u(z')$ big enough to ensure this:

$$u(z') < pu(z) + (1 - p)u(z^L).$$

But

$$\bar{u}_i^{\mathbf{v}'}(C(\mathbf{v}')) \leq u(z')$$

because $u(z') \geq u(w)$ for all $w \in C(\mathbf{v}')$, and

$$\bar{u}_i^{\mathbf{v}'}(C(\mathbf{v})) \geq pu(z) + (1 - p)u(z^L)$$

because $pu(z)$ is a summand of $\bar{u}_i^{\mathbf{v}'}(C(\mathbf{v}))$ and $u(z^L) \leq u(w)$ for all w . It follows that $\bar{u}_i^{\mathbf{v}'}(C(\mathbf{v})) > \bar{u}_i^{\mathbf{v}'}(C(\mathbf{v}'))$, contrary to \mathcal{S} .

We shall make repeated use of three more lemmata.

Singleton-Monotonicity Lemma. If $C(\mathbf{v}) = \{x\}$ and every alternative that ranks above x in any v'_i also ranks above x in the corresponding v_i , then $C(\mathbf{v}') = \{x\}$.

Proof. It is enough to show this when \mathbf{v} and \mathbf{v}' are i -variants for some i ; the full lemma then follows by $n - 1$ repetitions. If $C(\mathbf{v}')$ contained any y that ranked higher than x

in v_i , then y or something ranked even higher in v_i would have to belong to $C(\mathbf{v})$ by Strategy Lemma (2) (with the roles of \mathbf{v} and \mathbf{v}' reversed). But that is impossible because $C(\mathbf{v}) = \{x\}$. If $C(\mathbf{v}')$ contained any y that ranked lower than x in v_i , then y would also rank lower than x in v'_i by hypothesis, and y or something ranked even lower in v'_i would have to belong to $C(\mathbf{v}) = \{x\}$ by Strategy Lemma (1), again an impossibility. Hence, no alternative but x can belong to $C(\mathbf{v}')$. But $\phi \neq C(\mathbf{v}') \subseteq A$ by **CH**. So $C(\mathbf{v}') = \{x\}$.

$\hat{\mathcal{P}}$ -Sufficiency Lemma. If $x \neq y$ and $C(\mathbf{v}) = \{x\}$ then $x \hat{\mathcal{P}}^{\mathbf{v}} y$.

Proof. We must show that $C(\mathbf{v}') = \{x\}$ for every xy -twin \mathbf{v}' of \mathbf{v} in which $\{x, y\}$ is a top set. But at most y ranks above x in any v'_i , in which case y also ranks above x in v_i . So $C(\mathbf{v}') = \{x\}$ by Singleton-Monotonicity Lemma.

Top Lemma. If X is a top set in \mathbf{v} then $\phi \neq C(\mathbf{v}) \subseteq X$.

Proof. $\phi \neq C(\mathbf{v}) \subseteq A$ by **CH**. To show that $C(\mathbf{v}) \subseteq X$, take some $x \in X$ and some \mathbf{v}^x in which every ordering ranks x first. By **CiSov**, $\{x\} = C(\mathbf{v}')$ for some \mathbf{v}' . So $\{x\} = C(\mathbf{v}^x)$ by Singleton-Monotonicity Lemma. Now change \mathbf{v} to \mathbf{v}^x one ordering at a time. By Strategy Lemma (1), each change from v_i to v_i^x preserves membership in the choice set by any given member or something ranked lower in v_i . But every member of $A - X$ ranks below every member of X in every v_i . Hence, if any member of $A - X$ belonged to $C(\mathbf{v})$, then some member of $A - X$ would belong to $C(\mathbf{v}^x)$, which is impossible because $C(\mathbf{v}^x) = \{x\}$ and $x \in X$. So no member of $A - X$ belongs to $C(\mathbf{v})$, and thus $C(\mathbf{v}) \subseteq X$.

With these lemmata in hand, we can now complete the proof of our theorem by deducing that $\hat{\mathcal{P}}$ satisfies **PARETO**, **B**, and **TRANS**.

Proof that $\hat{\mathcal{P}}$ satisfies **PARETO**. Suppose xv_iy for all i (so $x \neq y$). To prove that $x \hat{\mathcal{P}}^{\mathbf{v}} y$, we must show that $C(\mathbf{v}) = \{x\}$ whenever \mathbf{v}' is an xy -twin of \mathbf{v} and $\{x, y\}$ is a top set in \mathbf{v}' . But in that case $\{x\}$ too is a top set in \mathbf{v}' , so $C(\mathbf{v}') = \{x\}$ by Top Lemma.

Proof that $\hat{\mathcal{P}}$ satisfies **B**. Suppose on the contrary that i is a *blocker* in this sense: for all x, y, \mathbf{v} , if xv_iy while yv_jx whenever $j \neq i$ then not $y \hat{\mathcal{P}}^{\mathbf{v}} x$. We shall deduce, contrary to **D**, that for all x, \mathbf{v} , if x ranks first in v_i then $x \in C(\mathbf{v})$. For suppose $x \notin C(\mathbf{v})$. Let y rank second in v_i . Take some \mathbf{v}' in which $v'_i = v_i$ while every other v'_j ranks y first and x last. It suffices to show that $C(\mathbf{v}') = \{y\}$, which implies, by $\hat{\mathcal{P}}$ -Sufficiency Lemma, that $y \hat{\mathcal{P}}^{\mathbf{v}'} x$, contrary to our hypothesis that i is a blocker.

If the change from v_j to v'_j for any $j \neq i$ let x enter the choice set, then the change from v'_j back to v_j would keep x in by Strategy Lemma (1) because nothing ranks below x in v'_j . Hence, $x \notin C(\mathbf{v}')$. Now change \mathbf{v}' to \mathbf{v}^y by putting y above x in the i th ordering, leaving all else the same. Then if anything that ranked lower than y in v'_i belonged to $C(\mathbf{v}')$, it or something that ranked even lower in v'_i would have to belong to $C(\mathbf{v}^y)$

by Strategy Lemma (2). But $\{y\}$ is a top set in \mathbf{v}^y , so $C(\mathbf{v}^y) = \{y\}$ by Top Lemma. Therefore, nothing lower than y in v'_i can belong to $C(\mathbf{v}')$, so only x and y can belong. But $x \notin C(\mathbf{v}')$. Hence, $C(\mathbf{v}') = \{y\}$.

Proof that $\hat{\mathcal{P}}$ satisfies **TRANS**. Suppose $x\hat{\mathcal{P}}^{\mathbf{v}}y\hat{\mathcal{P}}^{\mathbf{v}}z$. Then $x \neq y$, $y \neq z$, and $x \neq z$ by **SoPREF** (asymmetry). To show that $x\hat{\mathcal{P}}^{\mathbf{v}}z$, first change \mathbf{v} to \mathbf{v}^{xyz} by moving x , y , and z above all other alternatives in every ordering while preserving their positions relative to each other. Then $\{x, y, z\}$ is a top set in \mathbf{v}^{xyz} , and \mathbf{v}^{xyz} is an xy -twin, a yz -twin, and an xz -twin of \mathbf{v} . By **IIA**, therefore, $x\hat{\mathcal{P}}^{\mathbf{v}^{xyz}}y\hat{\mathcal{P}}^{\mathbf{v}^{xyz}}z$, and it suffices to show that $x\hat{\mathcal{P}}^{\mathbf{v}^{xyz}}z$.

We first show that $y \notin C(\mathbf{v}^{xyz})$. For suppose $y \in C(\mathbf{v}^{xyz})$. Change \mathbf{v}^{xyz} to \mathbf{v}^{xy} one ordering at a time by moving z immediately below x and y (below them but above all other alternatives) unless it is already there, leaving all else the same. Each change from v_i^{xyz} to v_i^{xy} must preserve y 's membership in the choice set. That is trivial if $v_i^{xy} = v_i^{xyz}$. So suppose not. Then z ranks last among x, y , and z in v_i^{xy} but not in v_i^{xyz} . Thus, y must rank first in v_i^{xy} or last among x, y , and z in v_i^{xyz} . If y ranks first in v_i^{xy} , then y stays in the choice set by Strategy Lemma (2). And if y ranks last among x, y , and z in v_i^{xyz} , then y or something even lower in v_i^{xyz} must belong to the post-change choice set by Strategy Lemma (1). But $\{x, y, z\}$ remains a top set after the change, so nothing but x, y , or z can belong to the post-change choice set by Top Lemma, and thus y must belong. Hence, $y \in C(\mathbf{v}^{xy})$, so $C(\mathbf{v}^{xy}) \neq \{x\}$. But that is impossible because $x\hat{\mathcal{P}}^{\mathbf{v}^{xyz}}y$, $\{x, y\}$ is a top set in \mathbf{v}^{xy} , and \mathbf{v}^{xy} is an xy -twin of \mathbf{v}^{xyz} . Consequently, $y \notin C(\mathbf{v}^{xyz})$ after all.

By a similar argument (move x below y and z), $z \notin C(\mathbf{v}^{xyz})$. But $\phi \neq C(\mathbf{v}^{xyz}) \subseteq \{x, y, z\}$ by Top Lemma. So $C(\mathbf{v}^{xyz}) = \{x\}$. It follows by $\hat{\mathcal{P}}$ -Sufficiency Lemma that $x\hat{\mathcal{P}}^{\mathbf{v}^{xyz}}z$, as desired.

6 Open Questions

Resolute or not, any multi-alternative collective-choice procedure must be manipulable or dictatorial. That is true regardless of the population of individuals or the set of feasible alternatives, regardless of how ties are resolved, regardless of the content, accuracy, or variety of beliefs about the resolution of ties, and regardless of what counts as an expression of preferences — regardless of the hypothesized mapping from preference orderings to individual strategies.

How important is all this? That depends on three open questions, harder than any we have answered.

Question 1. How common is strategic misrepresentation? Manipulability is one thing, manipulation another. The one is unavoidable. The other may be rare: a manipu-

lator must be in a position to change the choice set all by himself, and of course he must wish to do so and know enough about the (perhaps yet-to-be) professed preferences of others to know that he can. But manipulation is not the only kind of strategic misrepresentation. If the opportunity to gain from misrepresentation is rare, the opportunity to adopt a *dominant strategy* of misrepresentation may be greater: a misrepresentation by Mr. i that has no effect given the *actual* combination of acts by others might be advantageous under *some* such combination and disadvantageous under none. Dominant strategies of misrepresentation are quite common under parliamentary procedure with fixed agendas (Farquharson 1969, McKelvey and Niemi 1978, Moulin 1979). Here is one subject that merits further investigation. More important for us now, if the opportunity for *individuals* to gain by misrepresentation is rare, the opportunity for *groups* to do so may be greater.

What groups? A small enough group may be scarcely more potent than a single individual, and a big enough group may be empowered to get its way without need of misrepresentation. Among intermediate groups, an arbitrary one is not likely to be organized enough in its actions or cohesive enough in its preferences to strategize. But specific groups may be. Let our population be partitioned into groups organized and cohesive enough to strategize, given the chance. Call them *factions*. At least one factional partition must exist because individuals (their unit sets, that is) are factions in our sense. The rub is that nothing stops us from reinterpreting $1, 2, \dots, n$ as factions, however big they may be: our theorem applies to *any* factional partition of the population, coarse or fine.

Roughly speaking, the chance of strategic behavior is greatest when the factional partition is coarsest, and an institutional designer bent on minimizing strategic behavior should seek to minimize the coarseness of the most likely factional partitions. For small and mid-size organizations, that goal might be achieved by the continual shuffling of personnel among potential factions (or their continual elimination, as Stalin appreciated). For democratic polities, factional manipulation might be minimized by anti-careerist policies — term limits, weak tenure for civil servants, no collective bargaining. Ancient Athens perfected these devices: office holders were chosen by lot and frequently changed. For large democratic polities, James Madison argued in his celebrated “Federalist 10” that factional manipulation of low-rung procedures (constitutions) can be minimized by establishing numerous cross-cutting territorial and functional jurisdictions and electoral constituencies. How specific arrangements of these sorts affect the chances for manipulating specific procedures is another subject for further investigation: we suspect there are theorems to be proved that would complement the ideas and findings of constitutional and organizational theory.

Question 2. Under what conditions can an observer predict an *equilibrium*, a professed profile that no one has an incentive to change? Our theorem shows only that the “true” profile is not necessarily an equilibrium. Equilibriumhood of any profile is

relative to some given profile, itself or another, of supposed true preferences: \mathbf{v}' is an equilibrium relative to \mathbf{v} if and only if $\bar{u}_i^{\mathbf{v}}(C(\mathbf{v}'')) > \bar{u}_i^{\mathbf{v}}(C(\mathbf{v}'))$ for no i , i -variant \mathbf{v}'' of \mathbf{v}' , and representative u of v_i . \mathcal{S} says that every profile is an equilibrium relative to itself. But neither \mathcal{S} nor the equilibriumhood of the true profile is necessary for the existence of equilibria, of which there may be many, most of them — possibly all of them — predictively useless. Neither is the existence or uniqueness of equilibria, even useful ones, necessary for all predictive purposes. We illustrate the range of possibilities with six toy examples.

Example 1. Under Plurality Rule (choose the professed favorites of the most voters), the first of these two profiles, the “true” one, yields the choice set $\{x\}$, the second $\{z\}$:

1	2	3	1	2	3
x	x	x	z	z	z
y	y	y	y	y	y
z	z	z	x	x	x

Although Plurality Rule is manipulable, both profiles are equilibria (relative, as always, to the true one). But the second is predictively useless because it is *inaccessible*, foreign to any *path* of manipulations that starts from the true profile.

Example 2. Under Plurality Rule, the true profile

1	2	3	4	5
x	x	z	y	y
y	y	x	x	x
z	z	y	z	z

yields the choice set $\{x, y\}$. Mr. 3 can advantageously change this to $\{x\}$ by raising x above z . He alone can manipulate, and no further manipulations are possible. So the true profile is not an equilibrium, but the new one is, and of course it is accessible. Besides many inaccessible equilibria, there is one more accessible one, also yielding $\{x\}$. In it, Mr. 3 ranks x above y above z .

Example 3. Under Plurality Rule, the true profile

1	2	3
x	y	z
y	z	x
z	x	y
w	w	w

yields $\{x, y, z\}$. Ties are resolved randomly, and everyone knows it. With preference intensities represented by vertical distance in our picture, Mr. 1 can advantageously change $\{x, y, z\}$ to $\{y\}$ by raising y above x . Mr. 2 or Mr. 3 could have manipulated in a similar way, creating an equilibrium that yielded $\{z\}$ or $\{x\}$. So different paths end in different equilibria and different choices. There is a coordination problem: Messrs. 1, 2, and 3 might try to manipulate simultaneously, creating a nonequilibrium profile that yielded the original $\{x, y, z\}$. Despite the multiplicity of accessible equilibria yielding different choice sets and despite the coordination problem, we can still predict the rejection of w : it is not chosen along any path.

Example 4. The true profile is

1	2
x	z
y	x
z	y
w	w

The operative procedure says the choice set is $\{w\}$ unless one or both individuals rank w last, in which case the choice set is $\{x\}$ if both rank x above y , $\{y\}$ if both rank y above x , and $\{z\}$ otherwise. So the true profile yields $\{x\}$. But Mr. 2 can advantageously change this to $\{z\}$ by raising y above x , after which Mr. 1 can advantageously change $\{z\}$ to $\{y\}$ by raising y above x , after which Mr. 2 can advantageously change $\{y\}$ to $\{z\}$ by raising x above y , and so on. Although no profile is an equilibrium, we can still predict the rejection of w : as before, w is not chosen along any path.

Example 5. This is like Example 4 except that the choice set is $\{w\}$ unless *both* individuals rank w last. Now, all profiles in which *neither* individual ranks w last are equilibria, and they are the only ones. But all of them are inaccessible.

In Examples 1-3 and 5, the inaccessible equilibria can be eliminated by successively eliminating dominated professions of preference. But that is not always true, as witness:

Example 6. Of these two profiles, the first is the true one:

1	2	1	2
x	x	y	y
y	y	x	x
z	z	z	z

Under the operative procedure, the first profile yields $\{x\}$, the second $\{y\}$, all others $\{z\}$. So both displayed profiles are equilibria. The second is inaccessible, but the orderings therein are undominated.

A manipulator “changes” the choice set from X to Y by “changing” his professed ordering from his true one to a new one. His manipulation is *contractive* if $Y \subset X$, *disruptive* if not. It might be thought that all paths would end in equilibria if all possible manipulations were contractive, if our procedure satisfied

W/S If \mathbf{v} and \mathbf{v}' are i -variants and $\bar{u}_i^{\mathbf{v}'}(C(\mathbf{v}')) > \bar{u}_i^{\mathbf{v}}(C(\mathbf{v}))$ for some representative u of v_i , then $C(\mathbf{v}') \subset C(\mathbf{v})$ (weak strategy freedom).

Any series of changes that merely contract the choice set must end, after all, because A is finite.

But it does not immediately follow from **W/S** that all strategic *paths* must end. Imagine a path from \mathbf{v}^1 to \mathbf{v}^2 to \mathbf{v}^3 that changes the choice set from $\{x, y, z\}$ to $\{x, y\}$ to $\{x\}$. Suppose Mr. i can unilaterally change the choice set from $\{x\}$ back to $\{x, y\}$ by changing \mathbf{v}^3 to \mathbf{v}^4 . Assuming **W/S**, $\{x, y\}$ cannot be better than $\{x\}$ according to v_i^3 : $\bar{u}_i^{\mathbf{v}^3}(\{x, y\}) > \bar{u}_i^{\mathbf{v}^3}(\{x\})$ for no representative u of v_i^3 . However, if $\mathbf{v}_i^3 \neq \mathbf{v}_i^1$ then it is not Mr. i 's currently professed v_i^3 but his true v_i^1 that motivates his behavior, and for all we know, $\{x, y\}$ is better than $\{x\}$ according to v_i^1 : maybe $\bar{u}_i^{\mathbf{v}^1}(\{x, y\}) > \bar{u}_i^{\mathbf{v}^1}(\{x\})$ for some representative u of v_i^1 . For all we know, indeed, a path of strategic contractions could create the strategic opportunity, not originally present in \mathbf{v}^1 , to choose an alternative outside $\{x, y, z\}$. Maybe **W/S** does ban this sort of thing, but that is not obvious: a proof is needed.

To formalize all this, denote any sequence of profiles by \mathbf{s} , and any ordered n -tuple of utility functions by $\mathbf{u} = (u_1, \dots, u_n)$. Those functions are representatives of the corresponding orderings in one profile. Call it $\mathbf{v}^{\mathbf{u}}$. Define:

\mathbf{v}' *u-succeeds* \mathbf{v} if and only if, for some i , \mathbf{v}' is an i -variant of \mathbf{v} and $\bar{u}_i^{\mathbf{v}^{\mathbf{u}}}(C(\mathbf{v}')) > \bar{u}_i^{\mathbf{v}^{\mathbf{u}}}(C(\mathbf{v}))$,

and \mathbf{s} is a *u-path* if and only if $\mathbf{v}^{\mathbf{u}}$ is the first profile in \mathbf{s} and every successive profile in \mathbf{s} *u-succeeds* its immediate predecessor.

Then \mathbf{v} is a *u-equilibrium* if and only if no profile *u-succeeds* \mathbf{v} ,

and \mathbf{v} is *u-accessible* if and only if \mathbf{v} belongs to some *u-path*.

As we know, there may be a u -accessible u -equilibrium different from v^u , and there may be more than one: each is then *path-dependent*. Also a u -path may end with a u -equilibrium or cycle endlessly (u -paths of both sorts sometimes coexist, we have found). The obvious open questions demand conditions governing the existence, nonexistence, uniqueness, and path-dependence of u -accessible u -equilibria for any given u , and the effects of $W\mathcal{S}$ on such matters. An open question of a different sort is whether anything like our theorem remains true when \mathcal{S} is weakened to $W\mathcal{S}$.

Regardless of the answers, our imagined observer might predict, given u , that the chosen alternative will belong to

$$PRED(u) = \{x \mid \text{for some } u\text{-path } s \text{ and } v \text{ in } s, x \in C(v) \text{ and either } v \text{ is a } u\text{-equilibrium or } v \text{ recurs infinitely often in } s\}.$$

A more elaborate treatment would allow p_i to depend on finite sequences of profiles rather than single profiles. Also it would allow simultaneous strategic moves, with and without coordinating signals. And it would pare down the range of profiles by (perhaps among other ways) successively eliminating dominated professions of preference. A greater elaboration would impute greater sophistication to individuals by letting them see the branches of extensive-form games instead of seeking the momentary advantages of short steps down foggy paths. Coordination problems would then show up as information sets, and endless cycles would give way to mixed-strategy equilibria.

Question 3. When can an institutional designer achieve his goals? Never, if one of his goals is to make manipulation impossible. Maybe quite often, if he seeks merely to make manipulation unlikely. Why would he care about manipulation?

He might deplore manipulation because he deplores the “dishonesty” of manipulators or the information costs they bear or the “unfair” advantages they gain. He would specifically deplore *contractive* manipulation if he sought to keep choice sets as big as possible once certain constraints were met. Maybe the intended tie-resolving process is completely decentralized, driven by voluntary individual behavior, and our designer seeks to maximize liberty by making pre-resolution choices as permissive as can be. Or perhaps he is designing a constitution for a sovereign state, understands that it must be self-enforcing for want of any higher authority, and appreciates that a self-enforcing constitution cannot be too specific or constraining in its policy requirements lest opponents of particular policies pursue their ends outside rather than inside the constitution. However, if his sole goal is “good” choices, and if $C(v)$ always comprises “good” alternatives when v is the true profile, then contractive manipulation by itself is not objectionable to him.

This suggests that our designer would be happy to settle for $W\mathcal{S}$, which bans only

disruptive manipulations. But for all we know, $\mathcal{W} \setminus \mathcal{S}$ is neither necessary nor sufficient for “good” choices. Strategic contractions may be unobjectionable *by themselves*. As noted earlier, however, a path of strategic contractions might conceivably create the opportunity, not originally present, for departures from the unmanipulated choice set. Also a path of manipulations, some disruptive, might conceivably lead outside the unmanipulated choice set but then terminate inside that set or settle down in an infinite cycle within that set.

If $C(\mathbf{v})$ always comprises “good” alternatives when \mathbf{v} is the true profile, and if our PRED function makes satisfactory predictions, then “good” choices are ensured by

$$PRED(\mathbf{u}) \subseteq C(\mathbf{v}^{\mathbf{u}}).$$

But even this is stronger than necessary. For we have unnecessarily assumed that C represents both the operative procedure and the criterion for “good” choices absent manipulation. Let C continue to represent the former, but now let G represent the latter. While we are at it, we may as well let G depend on utility functions (though G may be invariant under monotonic changes in those functions). So $G(\mathbf{u})$ is the set of “good” alternatives when $u_1 \dots u_n$ are the true utility functions of Messrs. 1, 2, \dots , n . To ensure “good” choices, it is not necessary that $C(\mathbf{v}^{\mathbf{u}}) \subseteq G(\mathbf{u})$ or $PRED(\mathbf{u}) \subseteq C(\mathbf{v}^{\mathbf{u}})$. Again assuming the adequacy of PRED, *this* is necessary and sufficient for “good” collective choices:

$$\mathcal{S}^* \qquad PRED(\mathbf{u}) \subseteq G(\mathbf{u}).$$

The problem is to find plausible conditions on C and G which imply \mathcal{S}^* , or failing that, to find plausible conditions inconsistent with \mathcal{S}^* .

The former problem is akin to that of finding plausible conditions for *Nash-implementability* of G (Dasgupta, Hammond, and Maskin 1979, Maskin 1985). One difference is that \mathcal{S}^* does not require the existence of Nash equilibria. Implementation theory does contemplate different solution concepts, but all are refinements of Nash. We instead let strategic maneuvers cycle endlessly, so long as they never lead outside the G -set, or anyway so long as the cycle eventually settles down for good inside that set. If our goal is to ensure that collective choices belong to some target set, the demand for Nash equilibria is unwarranted.

Another difference is that we equate individual strategies with preference orderings. Implementation theory is more abstract: strategies can be objects of any sort. Our equation is more than the legacy of the strategy-freedom literature: it lets us define *accessibility* in terms of strategic *paths* from true profiles. Accessibility is more important, we think, than equilibriumhood, both because equilibria tend to abound when they exist

at all and because our substantive problem (to ensure “good” choices) does not require their existence. If, however, we wish to let strategies be objects of any sort, we can then define accessibility by naming one “naive” strategy N_i^u for every i and u and requiring (N_1^u, \dots, N_n^u) to be the first step of every u -path.

The biggest difference from our point of view is that we do not require resoluteness: a profile (or vector of strategies) produces a *set* of alternatives. In a sense, games must yield single outcomes. But in that sense, those outcomes are vectors of strategies. They can be mapped into further consequences of any sort in any number of ways, depending on intended applications. Implementation theory applies game theory to social-choice theory by starting with a mapping of outcomes into collective choices. But it assumes that collective choices are single chosen alternatives rather than choice sets. That assumption is pointless on its face, unrealistic because real collective-choice procedures rarely are resolute, and unwarranted by the stated substantive problem because multi-member choice sets are compatible with the goal of achieving “good” choices.

References

- ARROW, KENNETH J. (1963): *Social Choice and Individual Values*. 2nd ed. New Haven: Yale University Press.
- BARBERÁ, SALVADOR (1977): "Manipulation of Social Choice Mechanisms that Do not Leave 'Too Much' to Chance." *Econometrica* 45: 1573-1588.
- _____ (1977): "Manipulation of Social Decision Functions." *Journal of Economic Theory* 15: 266-278.
- BARBERÁ, SALVADOR, HUGO SONNENSCHN, AND LIN ZHOU (1991): "Voting by Committees." *Econometrica* 59: 595-609.
- DASGUPTA, PARTHA, PETER HAMMOND, AND ERIC MASKIN (1979): "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility." *Review of Economic Studies* 46: 185-216.
- FARQUHARSON, ROBIN (1969): *Theory of Voting*. New Haven: Yale.
- FISHBURN, P.C. (1973): *The Theory of Social Choice*. Princeton: Princeton University Press.
- GIBBARD, ALLAN (1973): "Manipulation of Voting Schemes: A General Result." *Econometrica* 41: 587-601.
- _____ (1977): "Manipulation of Schemes that Mix Voting with Chance." *Econometrica* 45: 665-681.
- KELLY, JERRY (1977): "Strategy-Proofness and Social Choice Functions without Singlevaluedness." *Econometrica* 45: 439-446.
- MASKIN, ERIC S. (1985): "The Theory of Implementation in Nash Equilibrium: A Survey." In Leonid Hurwicz, David Schmeidler, and Hugo Sonnenschein, eds. *Social Goals and Social Organization*. New York: Cambridge.
- MCKELVEY, RICHARD D., AND RICHARD NIEMI (1979): "A Multi-Stage Game Representation of Sophisticated Voting for Binary Procedures." *Journal of Economic Theory* 18: 3-22.
- MOULIN, HERVÉ (1979): "Dominance-Solvable Voting Schemes." *Econometrica* 47: 1337-1351.
- ORDESHOOK, PETER, AND THOMAS SCHWARTZ (1987): "Agendas and the Control of Political Outcomes." *American Political Science Review* 81: 179-199.

- PATTANAIK, PRASANTA K. (1978): *Strategy and Group Choice*. New York: North-Holland.
- SATTERTHWAITE, MARK ALLEN (1975): "Strategy-Proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions." *Journal of Economic Theory* 10: 187-217.
- SCHWARTZ, THOMAS (1982): "No Minimally Reasonable Collective-Choice Procedure can be Strategy-Proof." *Mathematical Social Sciences* 3: 57-72.
- _____ (1986): *The Logic of Collective Choice*. New York: Columbia.
- ZECKHAUSER, RICHARD (1973): "Voting Systems, Honest Preferences, and Pareto Optimality." *American Political Science Review* 67: 934-946.