Robustness against inefficient manipulation

Author(s): Eyal Baharad and Zvika Neeman

Source: *Social Choice and Welfare*, July 2007, Vol. 29, No. 1 (July 2007), pp. 55-67

Published by: Springer

Stable URL: https://www.jstor.org/stable/41106844

Soc Choice Welfare (2007) 29:55–67
DOI 10.1007/s00355-006-0197-9

ORIGINAL PAPER

Robustness against inefficient manipulation

Eyal Baharad · Zvika Neeman

Received: 25 March 2005 / Accepted: 10 August 2006 /
Published online: 16 September 2006

Abstract This paper identifies a family of scoring rules that are robust against coalitional manipulations that result in inefficient outcomes. We discuss the robustness of a number of Condorcet consistent and "point runoff" voting rules against such inefficient manipulation and classify voting rules according to their potential vulnerability to inefficient manipulation.

1 Introduction

It is well known that with sincere voting, some voting rules might select a Pareto dominated alternative. For example, under either Approval Voting[1] or the "vote for t candidates" rule with $t > 1$,[2] if the choice set is restricted to be a singleton, then the application of a neutral tie-breaking rule among the

[1] For an extensive discussion on Approval Voting, see Brams and Fishburn (1978) and Fishburn (1978).

[2] Under the vote for t candidates rule every voter assigns one point to t out of the k candidates and zero points to the remaining k-t candidates. The candidate with the largest number of points is selected (see Baharad and Nitzan 2005).

E. Baharad (✉)
Department of Economics, University of Haifa, Haifa 31905, Israel
e-mail: baharad@econ.haifa.ac.il

Z. Neeman
Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215, USA

Z. Neeman
Center for Rationality and Department of Economics, Hebrew University of Jerusalem, Jerusalem 91904, Israel
e-mail: zvika@BU.edu
URL: http://people.bu.edu/zvika/

"qualified" alternatives might result in the selection of a Pareto dominated alternative. In contrast, some other voting rules such as plurality rule and Borda rule never admit a Pareto dominated alternative under sincere voting. The purpose of this study is to examine which voting rules ensure that a Pareto dominated alternative cannot be chosen also when voters behave strategically, and try to manipulate the voting rule in their favor.

Indeed, a priori, it is not clear that a manipulation that results in a Pareto inferior outcome is necessarily unattractive to coalitions of agents. For example, suppose that a group of agents have to choose one of four alternatives a, b, c, and d. Suppose that all the agents strictly prefer alternative a to b, and that agent 1's preferences are such that she prefers a to b, b to c, and c to d. Suppose further that alternative c is selected by some voting rule. It may be the case that agent 1 cannot manipulate the voting rule so that alternative a is chosen instead of c, but she may be able to manipulate so that alternative b, which she prefers but is Pareto dominated, is selected.

We obtain the following main results. We show that a family of scoring rules, traditionally called "positional voting",[3] are robust against coalitional manipulation that results in Pareto inefficient outcomes. Thus, positional voting rules guarantee Pareto efficient outcomes, regardless whether manipulation is attempted or not. We also show that two Condorcet consistent rules, the Copeland and Black mechanisms, are robust against inefficient manipulation. (Counter-) examples showing that other "reasonable" voting rules are subject to such inefficient manipulation are provided. In particular, we show that plurality rule, which is a non-positional scoring rule, and two Condorcet consistent voting rules, namely Simpson rule and the top cycle rule, are susceptible to inefficient manipulation. A classification of voting rules according to their vulnerability to inefficient manipulation is presented in Sect. 6 below.

Our result thus mitigates previous results that suggested that scoring rules are highly manipulable.[4] It highlights the need to refine the notion of robustness against manipulation to allow for a consideration of the type of manipulation that is attempted. Not all manipulations are equally undesirable from the perspective of the maximization of social welfare. Some manipulations reduce aggregate social welfare, while others only affect the distribution of surplus.

The rest of the paper is organized as follows. The general setting is presented in the next section. The result that positional voting rules are robust against inefficient manipulation appears in Sect. 3. In Sect. 4 we characterize the set of tie-breaking rules that might make plurality rule susceptible to inefficient manipulation. The manipulation of Condorcet consistent rules is discussed in Sect. 5. Finally, in Sect. 6, we examine the dimension of the space in which inefficient manipulation may occur, which allows us to distinguish among voting rules according to their vulnerability to inefficient manipulation.

[3] See Riker (1982).

[4] See Nitzan et al. (1980), Nitzan (1985), Saari (1990a), Lepelley and Mbih (1994), Smith (1999), and Lepelley and Valognes (2003).

2 Assumptions and general setting

We employ the following notation. Let $N = 1,\ldots, n$ denote a finite set of agents, and A denote a finite set consisting of k distinct alternatives (candidates). The preference relation of agent $i, i \in N$, is denoted by L_i. Preferences are represented by linear orderings. That is they are assumed to be complete, transitive, and asymmetric—agents are never indifferent between any two alternatives. The space of linear orderings over the set of alternatives A is denoted by L. Denote the utility that agent $i \in N$ derives from the selection of alternative $a \in A$ by $u^i\ (a) \in \Re$. We assume that the agents are expected utility maximizers. That is, the utility that agent $i \in N$ derives from a lottery in which alternative $a_j \in A$ is selected with probability $p_j \geq 0$, $\sum_{j\in\{1,\ldots,k\}} p_j = 1$, is given by $\sum_{j\in\{1,\ldots,k\}} p_j u^i(a_j)$.

A voting rule is a mapping $V : \mathbf{L}^n \to 2^A \backslash \emptyset$ from the space of linear orderings of the agents into a nonempty subset of alternatives. Sometimes, the objective is to select a unique alternative, in which case a tie breaking rule has to be applied in order to pick just one of the alternatives that were chosen by the voting rule. A tie breaking rule is said to be *neutral* if a permutation of the names of the alternatives leads to the corresponding permutation of the name of the chosen alternative. A tie breaking rule is said to be *anonymous* if a permutation of the names of the agents does not affect the choice made by the voting rule. Neutrality thus imposes symmetry among the alternatives, while anonymity imposes symmetry among the agents.

We make the assumption that the agents‘ preferences are common knowledge. This assumption obviously facilitates manipulation. The study of manipulation under incomplete information is considerably more difficult and is beyond the scope of this paper.

3 Positional voting rules are robust against inefficient manipulation

We show that positional voting rules are robust against inefficient manipulation by a coalition of rational agents. We begin by defining the more general family of scoring rules[5].

Definition *Fix a non-decreasing sequence of real numbers $S_1 \leq S_2 \leq \cdots \leq S_k$, such that $S_1 < S_k$. Suppose that each agent ranks the alternatives, giving a score of S_1 to the alternative that is ranked last, a score of S_2 to the alternative that is ranked next to last, and so on. A scoring rule is a rule that selects the alternative that received the highest total score (see, e.g., Moulin* 1988, *p.* 231*). In case several alternatives all received the highest total score, an arbitrary tie breaking rule may*

5 Scoring rules are vulnerable to some other “paradoxes” of voting. See, for example, Fishburn (1974a), Nurmi (1999), and Saari (2000).

be used to determine the selected alternative from among those that received the highest total score.[6]

Definition *A Positional Voting Rule is a special scoring rule, in which the sequence of scores $S_1, S_2, \ldots, S_k$ is strictly increasing.*[7]

Proposition 1 *Suppose that all agents are rational—among several possible manipulations, they always use the one that maximizes their expected utility under the assumption that all other agents report truthfully. Then, a manipulation by a coalition of agents that results in the selection of a Pareto inferior alternative is impossible under any positional voting rule.*

Proof Fix a profile of utility functions $u \in \mathbf{L}^n$. For every agent $i \in N$, let $s_{u^i}(a) : A \mapsto \{S_1, \ldots, S_k\}$, denote the sincere score given by agent i with utility u^i to the alternative $a \in A$.

For any two different alternatives $a,b \in A$, and a (non empty) coalition of agents $M \subseteq N$, let

$$D_u^M(a,b) = \sum_{i \in M} (s_{u^i}(a) - s_{u^i}(b))$$

denote the difference between the total score of alternatives a and b for the agents in the coalition M.

Suppose that alternative a Pareto dominates alternative b. That is, $s_{u^i}(a) > s_{u^i}(b)$ for every agent i. It follows that,

$$D_u^M(a,b) > 0$$

for every non empty coalition $M \subseteq N$.

Suppose that some alternative c is receiving the highest total score, and that every member in some non empty coalition of agents, $M \subseteq N$, prefers alternative a to b to c. Otherwise, the coalition would surely not want to manipulate in favor of b against c. Our assumption about the rationality of the agents implies that we may assume that the coalition M cannot manipulate such that alternative a is selected instead of alternative c. (If the coalition's members could manipulate in favor of alternative a in such a way, they would, but this would *not* result in a Pareto inferior alternative being chosen with a positive probability.) If the coalition consists of $1 \leq m \leq n$ agents, this implies that:

$$D_u^{N \setminus M}(c,a) \geq m(S_k - S_1).$$

6 Under the plurality rule $S_1 = 0, S_2 = 0, \ldots, S_{k-1} = 0$ and $S_k = 1$. On this rule see, e.g., Richelson (1978) and Ching (1996).

7 A well known Positional Voting Rule is the Borda rule, where $S_1 = 0, S_2 = 1, \ldots,$ and $S_k = k-1$. On the Borda rule, see, e.g., Young (1974) and Saari (1990b).

By definition of D,

$$D_u^{N\setminus M}(c,b) = D_u^{N\setminus M}(c,a) + D_u^{N\setminus M}(a,b) > m(S_k - S_1).$$

So the coalition M cannot manipulate so that alternative b is chosen with a positive probability either. □

Proposition 1 requires that a scoring rule be increasing in order for it to be robust against inefficient manipulation. In the next section we show that the result of Proposition 1 cannot be extended to apply to all scoring rules. In fact, it can be shown that any scoring rule that is not strictly increasing may be subject to inefficient manipulation.[8]

4 Inefficient manipulation of plurality rule

In this section we show that a neutral tie-breaking rule ensures that the plurality rule cannot be inefficiently manipulated. As demonstrated in Example 2 below, this is not the case for some other non-increasing scoring rules.

Proposition 2 *Suppose that all agents are rational—among several possible manipulations, they always use the one that maximizes their expected utility under the assumption that all other agents report truthfully. Then, a manipulation by a coalition of agents that results in the selection of a Pareto inferior alternative is impossible under plurality rule with a neutral tie-breaking rule.*

Proof There are two ways in which an alternative can be selected by plurality rule: (i) as a clear "plurality winner" where a tie-breaking is not needed; and (ii) as one of the alternatives that is included among the set of winners that is selected by the tie breaking rule.

In case (i), if a coalition of agents can manipulate such that a Pareto dominated alternative is selected directly without recourse to the tie breaking rule, then it can obviously manipulate such that the Pareto dominating alternative will be selected as well. It therefore follows that a Pareto dominated alternative cannot be directly selected, without recourse to a tie breaking rule, by plurality rule.

Consider now case (ii) where the alternative that is selected by plurality rule is selected by a tie breaking rule. Suppose that alternative k is dominated by alternative h. Suppose that $\overline{J} > 1$ alternatives all receive the same largest

[8] Inspection of the proof of Proposition 1 reveals that if a scoring rule is not increasing, then it is possible that although alternative a Pareto dominates alternative b, $D_u^M(a,b) = 0$, and so $D_u^{N\setminus M}(c,b) = D_u^{N\setminus M}(c,a) + D_u^{N\setminus M}(a,b) = m(S_k - S_1)$. A coalition M would obviously not manipulate in favor of alternative b if it can manipulate in favor of alternative a, as would be the case under a neutral tie breaking rule, but it may well manipulate in favor of the Pareto dominated alternative b if the tie breaking rule ranks alternative b above alternative c above alternative a. It should be emphasized that sometimes a coalition may manipulate in favor of a Pareto dominated alternative even when the tie breaking rule is neutral. See Example 4 below for an example.

number of votes. Denote the set of these alternatives by J. It follows that there are $\bar{J}$ types of agents whose most favorite alternative may be chosen by the tie breaking rule. If the tie breaking rule is neutral, then the expected utility of any such agent i is given by $u^i = \frac{1}{\bar{J}}\sum_{a_j \in J} u^i(a_j)$. We show that no agent can benefit from manipulating in favor of alternative k. If all the agents report their preferences truthfully, then alternative k cannot be included in the set of most favorite alternatives J because k is Pareto dominated. The fact that we restricted our attention to case (ii) implies that an agent who manipulates in favor of alternative k can, at most, only bring to that alternative k is included in the set of most favorite alternatives J. Because the only way in which an agent can do that is at the expense of its own most favorite alternative, which as a consequence of such a manipulation will be taken out of the set J, a rational agent would not want to engage in such a manipulation. □

As demonstrated by the next example, plurality rule can be inefficiently manipulated if it employs a non neutral tie-breaking rule.

Example 1 Inefficient manipulation under plurality rule with a non-neutral tie-breaking rule.

Let the set of alternatives be given by $\{a,b,c,d\}$. There are three agents whose preferences are given by the linear orderings $a \succ b \succ d \succ c, c \succ a \succ b \succ d$, and $d \succ c \succ a \succ b$, respectively. Suppose that ties are decided according to the non-neutral priority rule: $b \succ c \succ a \succ d$. Note that alternative b is Pareto dominated by alternative a. If all the agents report their preferences truthfully, then plurality rule selects alternative c. The first agent may manipulate so that alternative b is chosen by reporting the preferences: $b \succ a \succ d \succ c$. This is the best manipulation that is available to agent 1: since she cannot bring to the selection of her first best alternative a, the best she can do is to ensure that her second most favorite alternative, b, is selected.

The analysis extends to the case of plurality rule with runoff,[9] which is formally defined as follows:

Plurality with runoff: In the first round each voter casts a vote for one alternative. If an alternative wins a strict majority of votes, then it is selected. Otherwise, a runoff by majority voting is called between the two alternatives that won the largest number of votes in the first round (see, e.g., Moulin 1988, p. 235). As in every other voting rule, ties are decided according to some tie breaking rule.

Suppose that alternative a Pareto dominates alternative b. If alternative a does not qualify for the runoff, then alternative b surely cannot, even if manipulation is attempted (because the set of those who vote for b as their favorite alternative is included in the set of those who voted for a). If alternative a does qualify to the runoff but does not win (assuming a neutral tie breaking rule), then the same argument that was used in the proof of Proposition 2 implies that no one would want to manipulate in favor of alternative b. However, under

9 For an analysis of scoring runoffs see Smith (1973) and Richelson (1980).

a non-neutral tie-breaking rule the plurality rule with runoff could be vulnerable to inefficient manipulation.[10] Clearly, other runoff mechanism could be constructed. For example, different scoring rule can be used in each round to sequentially eliminate the alternatives (see Richelson 1980). Such a runoff mechanism is robust to inefficient manipulation if a positional rule is used at each round.

Finally, the next example shows that Proposition 1 cannot be generalized to the class of all scoring rules even if one restricts attention only to scoring rules with neutral tie-breaking rules:

Example 2 Inefficient manipulation under a non increasing scoring rule with a neutral tie-breaking rule. Suppose that the set of alternatives is given by $\{a, b, c, d, e\}$ and that the scores are given by: $S_1 = 0, S_2 = S_3 = 6$, and $S_4 = S_5 = 10$. Ties are decided according to a neutral tie-breaking rule. That is, in case several alternatives all received the highest total score, every such alternative is selected with equal probability. There are 14 agents. The preferences of the first five agents are given by: $c \succ d \succ a \succ b \succ e$; of the next five agents by: $c \succ e \succ a \succ b \succ d$; and of the last four agents by: $a \succ e \succ b \succ d \succ c$. Note that alternative b is Pareto dominated by alternative a. If all the agents report their preferences truthfully, then alternatives a and c tie with a score of a 100 each, and are each selected with probability $\frac{1}{2}$. Observe that the coalition of last four agents may manipulate so that alternatives a, c, and b, are chosen with a probability $\frac{1}{3}$ each, by reporting the preferences: $a \succ b \succ e \succ d \succ c$.

5 Inefficient manipulation of condorcet consistent voting rules

We consider four Condorcet consistent rules,[11] Simpson rule, the Top Cycle Rule, Black rule, and Copeland rule.[12] We show that while the former two are vulnerable to inefficient manipulation, the latter two are not.[13]

Example 3 Inefficient manipulation of Simpson rule. Simpson rule selects the alternative that has the smallest number of agents against it in pairwise comparisons with other alternatives. It is formally defined as follows:

Simpson rule. For any two different alternatives $a, b \in A$, let $N(a,b)$ denote the number of agents who prefer alternative a to b. The Simpson score of alternative a is the minimum of $N(a,b)$ over all b, $b \neq a$. The alternative with the

[10] The argument is similar to the one given in footnote 8.

[11] A rule is Condorcet consistent (see, e.g., Moulin 1988, p. 229) if it selects the Condorcet winner (an alternative that is preferred to every other alternative by a majority of the agents) when it exists. On Condorcet consistent functions see Fishburn (1977).

[12] On Black's Condorcet-consistent mechanism see Black (1958) and Fishburn (1973, 1974b, 1977). On Copeland rule see Fishburn (1973, 1977).

[13] Scoring rules are not Condorcet consistent, since there exists a preference profile in which the Condorcet winner is not selected by any scoring rule—see Fishburn (1973, 1984).

highest Simpson score, called the Simpson winner, is selected (see, e.g., Moulin 1988, p. 233).

In case several alternatives all received the highest Simpson score, every such alternative is selected with equal probability (that is a neutral tie-breaking rule).

Let $a,b,c,d \in A$. Suppose that the preferences of agents 1, 2, and 3 are given by $d \succ a \succ c \succ b, b \succ d \succ a \succ c$, and $c \succ b \succ d \succ a$, respectively. Note that alternative a is Pareto dominated by alternative d. It can be verified that if the agents all report their preferences truthfully, then alternatives b, c, and d, are all chosen with an equal probability, and agent 1's expected utility is given by:

$$\frac{1}{3}u^1(b) + \frac{1}{3}u^1(c) + \frac{1}{3}u^1(d).$$

Agent 1 may increase her expected utility by reporting the preferences: $a \succ d \succ c \succ b$. Such a manipulation results an expected utility of

$$\frac{1}{4}u^1(a) + \frac{1}{4}u^1(b) + \frac{1}{4}u^1(c) + \frac{1}{4}u^1(d)$$

for agent 1. If $3u^1(a) > u^1(b) + u^1(c) + u^1(d)$, then such manipulation increases agent 1's expected utility. Observe that although the new outcome might be preferred by agent 1 to the previous one, it allows the choice of alternative a that is Pareto dominated by alternative d.

Moreover, if a non-neutral tie-breaking rule is employed, then it is possible to construct an example in which the Pareto dominated alternative is chosen with probability one under some preference profiles. For example, suppose that ties are decided according to the following priority rule, $a \succ b \succ c \succ d$. Agents' preferences are the same as above. If the agents report their preferences truthfully, alternative b is selected as the Simpson winner. As before, agent 1 may manipulate by reporting the preferences: $a \succ d \succ c \succ b$. Such a manipulation results in alternative a being selected.

Example 4 Inefficient manipulation of the Top Cycle rule. The Top Cycle rule selects the alternative that is preferred by a majority of the agents to all other alternatives if such an alternative exists. If such a Condorcet winner does not exist, then the Top Cycle rule selects randomly from among the top-cycle (the transitive closure of majority rule). It is formally defined as follows:

Top Cycle rule. For any two alternatives a_v, $a_w \in A$, let $a_v \succeq_T a_w$ if and only if there is an integer q and a sequence $a_v = a_1, \ldots, a_q = a_w$, such that a_j is preferred to a_{j+1} by at least half of the agents for every $j \in \{1,\ldots, q\text{-}1\}$. The Top Cycle is defined as the non empty set of maximal elements of $\succeq_T$. That is, an alternative a_v belongs to the top cycle if and only if $a_v \succeq_T a_w$ for every alternative $a_w \neq a_v$ (see, e.g., Moulin 1988, p. 253). The Top Cycle rule selects each alternative in the top cycle with equal probability.

Let $a,b,c,d,e \in A$. Suppose that there are nine agents. The preferences of the first four agents are given by: $a \succ b \succ d \succ c \succ e$; the preferences of the next three agents are given by: $c \succ b \succ e \succ a \succ d$; and the preferences of the last two agents are given by: $e \succ a \succ c \succ d \succ b$. Note that alternative d is Pareto dominated by alternative a. If all the agents report their preferences truthfully, then the top cycle ($a \succeq_T c \succeq_T b \succeq_T e \succeq_T a$) rule generates an expected utility of:

$$\frac{1}{4}u^i(a) + \frac{1}{4}u^i(b) + \frac{1}{4}u^i(c) + \frac{1}{4}u^i(e)$$

for the first four agents. The coalition of the first four agents can manipulate the outcome in its favor by reporting the preferences: $a \succ d \succ b \succ c \succ e$. In this case the top-cycle consists of alternatives $a,b,c,d,$ and e, ($a \succeq_T c \succeq_T d \succeq_T b \succeq_T e \succeq_T a$) and therefore applying the top cycle rule generates an expected utility of:

$$\frac{1}{5}u^i(a) + \frac{1}{5}u^i(b) + \frac{1}{5}u^i(c) + \frac{1}{5}u^i(d) + \frac{1}{5}u^i(e)$$

for the first four agents. If $4u^i(d) > u^i(a) + u^i(b) + u^i(c) + u^i(e)$ for every one of the first four agents, then these four agents increase their expected utility through this manipulation. Under this manipulation the Pareto inferior alternative d is selected with a positive probability.

As in the previous example, a non-neutral tie-breaking rule may lead to the selection of the Pareto dominated alternative with probability one under some preference profiles. Suppose that ties are decided according to the priority rule: $d \succ e \succ c \succ b \succ a$. Suppose that the nine agents have the same preferences as before. With the new tie-breaking rule, if all the agents report their preferences truthfully, then alternative e is selected. The first four agents can manipulate as before and guarantee the choice of alternative d.

Not every Condorcet consistent voting rule is vulnerable to inefficient manipulation. The Copeland rule, for example, which is defined as follows, is not.

The Copeland rule. Compare candidate a with every other candidate x. Score +1 if a majority prefers a to x, −1 if a majority prefers x to a, and 0 if a and x tie. Summing up those scores over all $x \neq a$ yields the Copeland score of a. The candidate with the highest Copeland score, called the Copeland winner, is elected (see Moulin 1988, p. 233).

The Copeland rule is robust against inefficient manipulation; even if agents try to manipulate the outcome in their favor, a Pareto dominated alternative cannot be selected. The proof proceeds as follows. Suppose that alternative a dominates alternative b, and the coalition of agents M has the following preferences: $x_1 \succ x_2 \succ \cdots \succ x_l \succ a \succ y_1 \succ y_2 \succ \cdots \succ y_m \succ b \succ v_1 \succ v_2 \succ \cdots \succ v_n$. Hence, the alternatives that are preferred to a by the coalition M are denoted by $x_1, \ldots, x_l$, the ones that are preferred to b and are inferior to a by $y_1, \ldots, y_m$, and those that are inferior to b by $v_1, \ldots, v_n$. A coalition that wants to manipulate in favor of alternative b can report either one of the two following preferences: (i)$x_1 \succ x_2 \succ \cdots \succ x_l \succ a \succ b \succ y_1 \succ y_2 \succ \cdots \succ y_m \succ v_1 \succ v_2 \succ \cdots \succ v_n$ or (ii) $x_1 \succ \cdots \succ x_j \succ b \succ x_{j+1} \succ \cdots \succ x_l \succ a \succ y_1 \succ y_2 \succ \cdots \succ y_m \succ v_1 \succ v_2 \succ \cdots \succ v_n$. (i) would not result in alternative b being selected among the winners, since every alternative defeated by b is defeated by a as well, and a defeats b. (ii) might result in b being qualified to the set of winners, but this implies that the coalition is powerful enough to manipulate in favor of alternative a, which Pareto dominates b. Note that such a coalition cannot be a majority coalition because a majority coalition can ensure the selection of its most favorable alternative without performing any manipulation. It therefore follows that a Pareto dominated alternative b cannot be chosen by the Copeland rule.

The reason that some Condorcet consistent rules might allow for inefficient manipulation is that a Pareto dominated alternative might be included in the choice set, and so might be chosen by the tie breaking rule that is employed by the voting rule. It therefore follows that Proposition 1 implies that if a Condorcet consistent rule employs a positional voting rule as a tie-breaking rule, then inefficient manipulation becomes impossible. Black's Condorcet consistent mechanism, which selects the Condorcet winner if it exists and otherwise selects the Borda winner, is an example for such a Condorcet consistent rule.

6 The dimension of inefficient manipulation

Suppose that there are k alternatives. Saari (1994) defines voting rules as mappings from the set of rational points in the simplex $S(k!)$ into the set of alternatives. The set $S(k!)$ is defined as follows:

$$S(k!) = \left\{ x \in R^{k!} : \sum_{t=1}^{k!} x_t = 1, \quad x_t \geqslant 0 \right\}$$

where x_t is the fraction of voters having preferences of type t (individual preferences are assumed to be linear orderings).

Consider the subset $S_{a,b}$ $(k!/2)$ of the simplex $S(k!)$ in which alternative b is Pareto dominated by alternative a. The set $S_{a,b}$ (k!/2) has dimension $\frac{k!}{2} - 1$. If the dimension of the subset of $S_{a,b}$ $(k!/2)$ on which inefficient manipulation is possible is strictly smaller than $\frac{k!}{2} - 1$, then it follows that as the number of agents becomes large, the proportion of preference profiles that permit inefficient manipulation, and hence also the conditional probability of inefficient manipulation, converges asymptotically to zero both under the assumptions of impartial culture and impartial anonymous culture (see Kuga and Nagatani 1974; Gehrlein and Fishburn 1976; and Gehrlein 2002).

We can thus distinguish among four classes of voting rules according to their susceptibility to inefficient manipulation as the number of agents becomes large as follows.

1. Voting rules that are robust against inefficient manipulation (e.g., Positional voting rules, Plurality rule with neutral tie breaking, Copeland and Black rules).
2. Voting rules that are susceptible to inefficient manipulation on a set of dimension strictly smaller than $\frac{k!}{2}-1$. The proportion of preference profiles under which such rules are susceptible to inefficient manipulation decreases to zero as the number of agents becomes large.
3. Voting rules that are susceptible to inefficient manipulation on a set of full dimension. The proportion of preference profiles under which such rules are susceptible to inefficient manipulation remains bounded away from zero as the number of agents becomes large.
4. Voting rules that may generate Pareto inefficient outcome even without any manipulation.

Propositions 1 and 2 imply that positional voting rules and plurality rule with a neutral tie breaking rule belong to class (1). As shown in Sect. 5, Copeland rule and Black rule also belong to this class. As explained in the introduction, Approval Voting and the "vote for t candidates" rule with $t > 1$, belong to class (4). The next two examples are of voting rules that belong to classes (2) and (3), respectively.

Example 5 We show that for plurality rule with a non-neutral tie breaking rule inefficient manipulation occurs on a set of dimension $\frac{k!}{2}-2$. We illustrate the argument for the case where $k = 5$. Denote the five alternatives by a,b, c,d, e, respectively. Consider a profile $x \in S_{a,b}$ $(k!/2)$ and a number $1/3 < \alpha < 1/2$ such that:
A proportion α of the voters hold the preferences $a \succ b \succ c \succ d \succ e$ (Group A); A proportion α of the voters hold the preferences $c \succ a \succ b \succ d \succ e$ (Group B); and a proportion $1-2\alpha$ of the voters hold the preferences $d \succ a \succ b \succ e \succ c$ (Group C).

Suppose that ties are determined according to the non-neutral rule $b \succ c \succ a \succ e \succ d$. Thus if all the voters report their preferences truthfully, then plurality rule selects alternative c. However, by reporting, say, the preference profile $b \succ a \succ c \succ d \succ e$, the coalition of agents in Group A can and would manipulate the outcome in favor of the Pareto dominated alternative b. By varying the preferences held by the voters in Group C it is possible to derive new independent profiles in which similar inefficient manipulation is possible. The voters in Group C may have any preferences that are different from the preferences held by the voters in Groups A and B subject to the following two qualifications:

1. If the voters in Group C rank alternative a as the top one, which means that the Pareto dominant alternative a can be chosen without any manipulation, then the preferences of a proportion $1-2\alpha$ of the voters in Group A need to be changed to $d \succ a \succ b \succ e \succ c$. Under this new profile of preferences, if voters report their preferences truthfully, then because of the tie-breaking rule, alternative c is selected. But (at least) the coalition that consists of all

the voters who originally belonged to Group A who now hold either the preferences $a \succ b \succ c \succ d \succ e$ or the preferences $d \succ a \succ b \succ e \succ c$ can and would manipulate in favor of alternative b.

2. If the voters in Group C rank alternative c as the top one, which means that the Pareto dominant alternative c can be chosen without any manipulation, then the preferences of a proportion $1-2\alpha$ of the voters in Group B need to be changed to $d \succ a \succ b \succ e \succ c$. Under this new profile of preferences, if voters report their preferences truthfully, then because of the tie-breaking rule, alternative c is selected. But (at least) the coalition that consists of voters in Group A can and would manipulate in favor of alternative b.

It thus follows that there exist $5!/2 - 2$ independent profiles, and by changing $1/3 < \alpha < 1/2$, $5!/2 - 2$ dimensions of the set $S_{a,b}(5!/2)$ in which inefficient manipulation is possible.

Example 6 We show that the Top Cycle rule permits inefficient manipulation on a subset of full measure in $S_{a,b}$ $(k!/2)$. Consider the normalization of the profile of preferences described in Example 4 in the simplex $S_{a,d}(5!/2)$. Denote this profile by y. Consider a ball B of radius ϵ around the profile $I = (2/k!, \ldots, 2/k!)$. For any profile $x \in B$, the results of the pairwise comparisons between any two alternatives are close to zero except for the pair a and d because, by assumption, a Pareto dominates d. There exists a sufficiently small λ such that inefficient manipulation is possible for every profile $z = \lambda x + (1 - \lambda)y$ where $x \in B$. This is because at the profile z, the comparison between any two alternatives except for a and d is determined by the results of the respective pairwise comparison for the profile y. It follows that inefficient manipulation is possible in an open ball of profiles in the simplex $S_{a,d}(5!/2)$.

For voting rules in class (2), using an adequate tie breaking rule (such as Borda rule) might eliminate inefficient manipulation altogether, as it indeed does for Plurality Rule. For voting rules in class (3), an appropriate tie breaking rule might also help make the rule less manipulable, but because the tie breaking rule would have to be applied more frequently under such rules, this would come at the expense of changing the voting rule itself.

Acknowledgements We thank Drew Fudenberg and Eric Maskin for useful discussions, and an anonymous referee for his or her perceptive and helpful comments. We also thank seminar participants at Harvard, Haifa University, and the Technion for their comments. Neeman gratefully acknowledges financial support from the National Science Foundation under grant No. SBR 9806832. We are grateful to the referee for suggesting the ideas contained in Sect. 6.

References

Baharad E, Nitzan S (2005) The inverse plurality rule—an axiomatization. Soc Choice Welfare 25(1):173–178
Black D (1958) The theory of committees and elections. Cambridge University Press, Cambridge
Brams SJ, Fishburn PC (1978) Approval voting. Am Polit Sci Rev 72:831–847
Ching S (1996) A simple characterization of plurality rule. J Econ Theory 71:298–302
Fishburn PC (1973) The theory of social choice. Princeton University Press, Princeton
Fishburn PC (1974a) Paradoxes of voting. Am Polit Sci Rev 68:537–546

Fishburn PC (1974b) On the sum-of-ranks winner when losers are removed. Discret Math 8:25–30
Fishburn PC (1977) Condorcet social choice functions. SIAM J Appl Math 33:469–489
Fishburn PC (1978) Axioms for approval voting: direct proof. J Econ Theory 19:180–185
Fishburn PC (1984) Discrete mathematics in voting and group choice. SIAM J Algebr Discret Meth 5:263–275
Gehrlein WV (2002) Condorcet's paradox and the likelihood of its occurrence. Theory Decis 52:171–199
Gehrlein WV, Fishburn PC (1976) Condorcet's paradox and anonymous preference profiles. Public Choice 26:1–18
Kuga K, Nagatani H (1974) Voter antagonism and the paradox of voting. Econometrica 42:1045–1067
Lepelley D, Mbih B (1994) The vulnerability of four social choice functions to coalitional manipulation of preferences. Soc Choice Welfare 11:253–265
Lepelley D, Valognes F (2003) Voting rules, manipulability and social homogeneity. Public Choice 116:165–184
Moulin H (1988) Axioms of cooperative decision making, Econometric Society Monographs. Cambridge University Press, Cambridge
Nitzan S (1985) The vulnerability of point-scoring schemes to preference variation and strategic manipulation. Public Choice 47:349–370
Nitzan S, Paroush J, Lampert S (1980) Preference expression and misrepresentation in point voting schemes. Public Choice 35:421–436
Nurmi H (1999) Voting paradoxes and how to deal with them. Springer, Berlin Heidelberg New York
Richelson JT (1978) A characterization result for the plurality rule. J Econ Theory 19:548–550
Richelson JT (1980) Running off empty: run-off point systems. Public Choice 35:457–468
Riker WH (1982) Liberalism against populism: a confrontation between the Theory of Democracy and the Theory of Social Choice. Freeman, San Francisco
Saari DG (1990a) Susceptibility to manipulation. Public Choice 64:21–41
Saari DG (1990b) The Borda dictionary. Soc Choice Welfare 7:279–317
Saari DG (1994) Geometry of voting. Springer, Berlin Heidelberg New York
Saari DG (2000) Mathematical structure of voting paradoxes II: positional voting. Econ Theory 15:55–102
Smith DA (1999) Manipulability measures of common social choice functions. Soc Choice Welfare 16:639–661
Smith JH (1973) Aggregation of preferences with variable electorate. Econometrica 6:1027–1041
Young HP (1974) An axiomatization of Borda's rule. J Econ Theory 9:43–52