

A welfarist critique of social choice theory: interpersonal comparisons in the theory of voting

AKI LEHTINEN

TINT, University of Helsinki

Abstract: This paper provides a philosophical critique of social choice theory insofar as it deals with the normative evaluation of voting and voting rules. I will argue that the very method of evaluating voting rules in terms of whether they satisfy various conditions is deeply problematic because introducing strategic behaviour leads to a violation of any condition that makes a difference between voting rules. I also argue that it is legitimate to make interpersonal comparisons of utilities in voting theory. Combining a realistic account of voters' behaviour with a utilitarian evaluation of the outcomes then leads to the judgment that strategic voting is beneficial. If it is, then Arrow's theorem does not have far-reaching consequences for democracy because one of its conditions is not normatively acceptable.

Keywords: economic methodology, strategic voting, interpersonal comparisons of utility, foundations of social choice

JEL Classification: B41, D72, D81

Insofar as social choice theory deals with voting, it aims to provide answers to two main questions. First, how well does voting perform in collecting information on individual preferences? Second, what is the relative performance of various voting rules? And which rule is the best?

The theory aims to answer these questions by imposing conditions on aggregation rules that are intended to represent the voting rules.

AUTHOR'S NOTE: Earlier versions of this paper have been presented at the public choice conference in Turku 2006, at a conference on theoretical and applied ethics in St. Petersburg 2013, and at the moral and political philosophy seminar at University of Helsinki in 2014. In addition to engaging me in fruitful conversations about the topics, Donald Saari has commented on an earlier unpublished version of this article (see Saari 2010). Eerik Lagerspetz and Hannu Nurmi have also made useful comments on various versions of the manuscript. Given that the aforementioned scholars may well disagree with some of my conclusions, it is clear that they could not be held responsible for the end result. Finally, two anonymous referees also helped me to make this a better paper.

The conditions are thought to be normatively compelling in the sense that they should be satisfied. They express how given inputs in terms of individual preferences are to be translated into collective judgments or decisions. For example, various monotonicity conditions specify that becoming more popular should increase the chances that a candidate is selected (or at least that the chances do not decrease); and the independence of irrelevant alternatives (IIA) condition requires, roughly, that the social choice between any two alternatives depends only on the individual preference orderings concerning those two alternatives.

The voting rules are then evaluated in terms of the set of conditions that the corresponding aggregation rules satisfy. The best aggregation rule is the one that satisfies the largest number of normatively important conditions. This is how the second question is answered. The famous theorems of Arrow (1951; 1963) and of Gibbard and Satterthwaite (Gibbard 1973; Satterthwaite 1975) provide one answer to the first question by showing that there are no aggregation rules that satisfy sets of conditions that are thought to be normatively compelling.

Social choice theory is only concerned with aggregating preferences or votes. In other words, it does not take behavioural assumptions into account, and voters are assumed to vote *sincerely* (i.e., non-strategically). In most voting rules, strategic voting means that a voter gives a vote to an alternative that she does not consider to be the most preferred one. One may consider the assumption of sincere behaviour as merely a consequence of division of labour among extant approaches: social choice theory evaluates voting procedures *normatively* by imposing conditions on aggregation rules, while various decision and game theoretical approaches study the *positive* properties of voting rules. Finally, implementation theory combines the *normative* and *positive* approaches. This division of labour seems sensible because strategic voting itself is widely thought to be self-evidently harmful, and thus unacceptable in a normative evaluation of voting.

However, I have recently shown that strategic behaviour increases the frequency with which the *utilitarian winner* is chosen compared to sincere behaviour in all commonly used voting rules studied thus far (Lehtinen 2007a; 2007b; 2008). The utilitarian winner is the candidate with the highest sum of utility for all voters. The basic reason for these results is that broadly supported candidates and alternatives are more likely to obtain strategic votes than those that have a number of vehement supporters, but also many enemies. To put it briefly,

these results show that strategic voting allows voters to express their *preference intensities*, and it is thus typically beneficial rather than harmful.

This paper endeavours to sort out the methodological, ethical, and normative consequences of my results for social choice and implementation theory. Although the argument is written in a self-contained way, some of the issues are already discussed in more detail in a previously published article (Lehtinen 2011). In that paper, I argued that even though an aggregation rule satisfies IIA, this does not mean that preference intensities or third alternatives do not affect the results, and it does not guarantee the observability of preference orderings or the absence of strategic voting. In this paper I generalise such argumentation to all central *difference-making* conditions. A difference-making condition is one that is satisfied by some, but not all, commonly used aggregation rules.

For the purposes of this paper, it is necessary to distinguish between an *aggregation rule* and a *voting rule*. The former refer to rules that operate on the basis of the assumption of sincere behaviour. The latter refer to voting rules in which strategic voting may affect the results.¹ My main claim can now be put as follows: while there are difference-making conditions for aggregation rules, there are no difference-making conditions for voting rules. Every voting rule violates every condition that is difference-making for the corresponding aggregation rule. If this claim is correct, it follows that comparing and evaluating voting rules in terms of difference-making conditions that the corresponding aggregation rules should satisfy is deeply problematic. This means that the conditions fail to provide reliable knowledge about the kind of information voting results really depend on.

Evaluating voting rules in terms of the conditions on the corresponding aggregation rules is the only way in which one can proceed in evaluating voting rules if interpersonal comparisons of utilities are not admissible. This is why I will argue against the view of some social choice theorists about the role of interpersonal

¹ There is thus a difference between what I call a *voting rule* and a *voting procedure*, as defined, for example by Blin and Satterthwaite (1977). The latter takes the vote ballots as arguments and yields social choices as output, while the former takes the utility profile as well as behavioural assumptions and beliefs as arguments, and yields social choices as output. A voting procedure does not make any behavioural assumptions because the procedure merely counts the votes irrespective of their relation to the underlying preferences of the voters. A voting procedure is thus an aggregation rule according to the terminology adopted in this paper.

comparisons in voting theory. In particular, according to this view, different voting (or aggregation) rules incorporate different interpersonal comparisons (or worse, that some but not all voting rules incorporate interpersonal comparisons). This view then leads to the idea that some particular voting rules are problematic because interpersonal comparisons are problematic. I will argue against this idea: Interpersonal comparisons rather concern the preference or utility profile, and thus they are completely independent of voting rules. This point is important for establishing the legitimacy of employing interpersonal utility comparisons in voting theory: it does not matter that it is impossible to obtain accurate information about interpersonal comparisons if the same comparison is made for every voting rule. At least voting theorists need not favour one rule over another by making interpersonal comparisons.

Insofar as social choice theorists refrain from making interpersonal comparisons, they never make normative judgments about the best alternative on the basis of the preference profile alone. All the normative work is done by the conditions imposed on the aggregation rules. My proposal for a more reasonable alternative for social choice theory is to define the best candidate independently of the voting rule, and then study voting rules in terms of how often and under what kind of circumstances the different rules find this candidate.

The paper is organised in the following way. I provide a synopsis of the overall argument in section 1. In section 2, I discuss the difference between votes and preferences. I then argue, in section 3, that the very idea that interpersonal comparisons are incorporated into an aggregation rule is not tenable. In Section 4, I make a case for a utilitarian evaluation of alternatives. Sections 2 to 4 thus describe how social choice theorists' views concerning interpersonal comparisons arise in voting theory, and why they are flawed.

In section 5, I show how to generate a violation for any difference-making condition with strategic voting. This is done by considering examples from a recent debate between Donald Saari and Mathias Risse. In section 6, I explain the philosophical implications of the fact that strategic voting typically increases welfare. Sections 5 and 6 thus concentrate on evaluating aggregation rules in terms of normative conditions. In section 7, I discuss some constraints on the kinds of interpersonal comparisons that voting theorists could use.

1. AGGREGATION RULES AND VOTING RULES

Let \mathbf{p} denote a profile of individual preferences; it assigns a preference ordering for each voter. Let \mathbf{U} similarly denote a profile of individual utilities. An aggregation rule yields a social choice or a social ordering for a given preference or utility profile. For my purposes, it does not matter whether an aggregation rule yields an ordering, a single choice, or a set; and it also does not matter whether the rule takes the preference profile or the utility profile as its argument or whether the rule computes the outcomes from pairwise comparisons or orderings. I will thus refer to preference and utility profiles simply as Input \mathbf{I} , and the values under aggregation rule A_j simply as Output O_j . An *aggregation rule* can now be described as follows:

$$O_j = A_j(\mathbf{I})$$

Let ω denote a theory of how voters behave when they vote. In particular, if voters are strategic, it specifies *how* they vote strategically when they do so. All realistic models of strategic voting assume that voters' behaviour depends on their belief profile \mathbf{b} as well as their utility profile \mathbf{U} so that we can write $\omega = \omega(\mathbf{b}, \mathbf{U})$. Let V_j denote voting rule j . Since voters' behaviour also depends on the voting rule, we can write $\omega_j(\mathbf{b}, \mathbf{U})$. A *voting rule*, then, can be represented as follows.

$$O_j = V_j(\omega_j(\mathbf{b}, \mathbf{U})), \text{ or } O_j = V_j(\omega_j(\mathbf{b}, \mathbf{I}))$$

The distinction between an aggregation rule A_j and the *corresponding voting rule* V_j is thus simply that only the latter takes strategic voting and other behavioural assumptions into account. Aggregation rule A_j corresponds to voting rule V_j , and vice versa, if and only if the two types of rules process the incoming information in the same way. For example, the Borda aggregation rule corresponds to the Borda voting rule, and the Plurality aggregation rule corresponds to the Plurality voting rule, and so on. I will not provide specifications for ω_j here. This was done in my earlier studies of strategic voting and in various papers in voting theory. The overall argument in this paper can now be put as follows:

- 1) Preference intensities are normatively important.
- 2) The best alternative is the utilitarian winner.
[Social choice theory is based on the methodological injunction against making interpersonal comparisons. Thus, even if the utilitarian winner were to be admitted as the best alternative, if one has no way of observing which alternative is the utilitarian winner, it cannot be used in voting theory.]
- 3) However, interpersonal comparisons can be legitimately used in voting theory because the arguments against using interpersonal comparisons are not cogent:
 - Different voting rules do not incorporate different comparisons.
 - Interpersonal comparisons concern the utility profile rather than the relationship between the profile and the voting rule.
 - Although it is impossible to obtain reliable information about interpersonal comparisons, they can be made in a methodologically acceptable way in evaluating the performance of voting rules if the same comparison is made under every voting rule.
 - If one studies all voting rules with the same profile, one can make the same interpersonal comparisons under every rule.
- 4) [From 1, 2, and 3] The utilitarian criterion can and should be used in a normative evaluation of voting outcomes.
- 5) Strategic voting may well occur under incomplete information, and it does not require that voters have particularly precise information. Hence, it is not merely a logical possibility, but occurs reasonably often.
- 6) Every difference-making condition can be violated if voters are strategic.
- 7) Even if aggregation rule A_i satisfies condition C, but aggregation rule A_j does not, [from 6] there is no guarantee that C is less frequently violated in voting rule V_i than in voting rule V_j .
- 8) [From 7] Social choice theory fails to provide a satisfactory normative comparison of voting rules because the very idea of evaluating voting rules in terms of conditions on aggregation rules is misguided.
- 9) My models (in Lehtinen 2007a; 2007b; 2008) show that strategic voting is typically beneficial [this claim uses point 4].

- 10) [From 9] Social choice theory also fails to provide a satisfactory analysis of the first question because there *cannot be* an acceptable set of conditions that a voting rule should satisfy.
- 11) [From 8 and 10] The part of social choice theory dealing with voting is fundamentally flawed.

It is well known that Kenneth Arrow ignored the “game aspects” (1963, 7) of voting by assuming that people’s preferences can be aggregated into a social choice or ordering without considering strategic voting. My critique could thus be taken to be disingenuous. However, taking those game aspects into account has more far-reaching epistemological consequences than has been previously acknowledged. My arguments for the steps above are all based on taking such consequences into account.

Steps 8, 10, and 11, are strong claims, and at this point might seem arbitrary, unclear, or unjustified. Some social choice theorists might attack 5 (i.e., that strategic voting does not require complete or perfect information) and then conclude that while 6 may be a logical possibility, 7 does not hold because strategic voting is too rare to be relevant. Then 8 would not hold either. In this paper, I will try to demonstrate 6, and then, taking 5 as given, argue for 7 and 8. Thus I will not try to argue for 5. It is primarily an empirical endeavour (see Mackie 2003), and it is not suitable to be discussed in a philosophical paper.²

If step 4 holds, then 9 is demonstrated in my earlier writings. The crucial steps are thus 1 (i.e., that preference intensities are normatively important), 2 (i.e., that the best alternative is the utilitarian winner), and especially 3 (i.e., that interpersonal comparisons can be legitimately used in voting theory). I presume that most social choice theorists acknowledge that 1 and 2 could be true, but they might claim that step 4 (i.e., that the utilitarian criterion can and should be used in a normative evaluation of voting outcomes) simply does not follow from 2. Even if the best alternative were the utilitarian winner, this information is useless if it cannot be identified. In particular, the sum of utility requires information on interpersonal utility comparisons but it is impossible to obtain reliable information on them. Given this epistemic problem, social choice theorists typically interpret voting

² For an explanation of why strategic voting does not require particularly precise information on the part of the voters, see Lehtinen 2007b. On how to measure the extent of strategic voting, see Kawai and Watanabe 2013, and the references therein.

rules as procedures for finding the best alternative. Rather than trying to define the best alternative in terms of the preference profile, the theory characterizes voting rules in terms of the conditions that the aggregation rules satisfy, and the conditions then define the best outcomes indirectly. Any outcome that could emerge from an aggregation rule that satisfies certain conditions is normatively acceptable in terms of satisfying those conditions. There is an aggregation rule that guarantees the choice of the utilitarian winner under the standard assumption that the ballots correspond to individual utilities: the zero-one rule (or range voting). However, this voting rule is so highly manipulable that very few social choice theorists believe that it could possibly be guaranteed to find the utilitarian winner.

On the one hand, various descriptive models of strategic voting (e.g., McKelvey and Ordeshook 1972; Enelow 1981) show how preference intensities affect the outcomes in many voting rules which seemingly only provide an opportunity to express preference orderings or mere top-preferences. On the other hand, as many authors have acknowledged, if we are only interested in defining the best alternative in terms of individual preferences, without regard to practical questions such as how to ascertain information about the socially most preferred alternative, preference intensities are obviously relevant (e.g., Dummett 1984, 54). I will thus not argue for steps 1 or 2. I take them to be obviously true. If step 1 holds, the only options seem to be the sum or the product of utilities. Intensities are thus acknowledged to be normatively and positively relevant, but if we are to believe social choice theorists, they should not be employed because interpersonal comparisons of intensities are meaningless (Arrow 1951; 1963) and epistemically suspect (Arrow 1977).³

What needs to be established is thus not that preference intensities are normatively or positively relevant (steps 1 and 5), but rather why the arguments against interpersonal comparisons fail to support the conclusion that a normative evaluation of voting rules is best conducted without making interpersonal comparisons. In this paper, I do not distinguish between interpersonal comparisons of preference intensities and of utilities, and I always mean what is referred to as 'full comparability' in the literature on social welfare functionals (SWFLs). Many social choice theorists think that, in contrast to intensities, it is

³ Arrow (1978) acknowledged the possibility of ordinal interpersonal comparisons.

possible to obtain reliable information on preference orderings, and that different voting rules make different interpersonal comparisons.

I will argue that the arguments against making interpersonal comparisons of intensities would be cogent if the assumption of sincere behaviour characterized voting. But it does not. I will thus provide an argument for making interpersonal comparisons of preference intensities in voting theory. Since making interpersonal comparisons is necessary only in a normative evaluation of voting rules, but not in predicting voters' choices, this is tantamount to providing an argument for evaluating the candidates or alternatives in utilitarian terms. However, although I am defending a utilitarian evaluation of outcomes in evaluating voting rules, I do not intend to defend utilitarianism in general. Neither can I be defending or criticising any particular voting rules, especially the utilitarian ones, because, as I will show, interpersonal comparisons concern the utility profile.

If voting theorists are not allowed to make interpersonal comparisons, they are not allowed to pass normative judgments on the alternatives directly on the basis of the utility (or preference) *profile*.⁴ Once this requirement is adopted, examining the conditions that the various aggregation rules satisfy (and the paradoxes associated with violating them) becomes the only way in which voting rules can be normatively evaluated. It has long been acknowledged that the normative force and importance of various conditions is a matter of controversy, and the problem of weighing the importance of different conditions easily explains why social choice theorists do not agree on which voting rules are the best.

There is a much more serious problem, however, with the very idea of evaluating voting rules in terms of conditions on aggregation rules. Although there are differences between the sets of conditions that the various aggregation rules satisfy, due to strategic voting, the corresponding voting rules never satisfy any conditions that make a difference to their comparative evaluation. It is thus impossible to make meaningful comparative normative assessments on voting rules with such conditions, because we ought to be interested in whether voting rules, rather than aggregation rules, satisfy various conditions. This argument establishes that eschewing interpersonal comparisons forces

⁴ As mentioned at the beginning of this section, an ordinal preference profile is a list of preference orderings, one for each voter. A utility profile is a listing of all utilities in an electorate, one utility number for each individual and each alternative.

upon us a highly dissatisfactory method of comparative normative evaluation.

Step 10 (i.e., that there cannot be an acceptable set of conditions that a voting rule should satisfy) can be explicated as follows. We know that strategic voting cannot be prevented in any real voting rule (Gibbard 1973; Satterthwaite 1975). If strategic voting is beneficial, there cannot be a set of normatively appealing conditions that any voting rule could ever satisfy, because strategic voting implies that all other difference-making conditions will be violated. It follows that the set of conditions for Arrow's theorem are normatively unacceptable, or, to put it more strongly, that there cannot be a set of normatively acceptable conditions in the first place.

2. DISTINGUISHING BETWEEN PREFERENCES AND VOTES

When Kenneth Arrow founded the modern version of social choice theory in the 1950s, only ordinal utilities were considered epistemically acceptable. Interpersonal comparisons of utilities were considered even more problematic than cardinal utilities because there are no individual choices that could provide reliable information on such comparisons. Myerson (1985) gives an often cited expression of such epistemic arguments in pointing out that it is not possible to *choose to be* person A in state Y rather than person B in state X. Interpersonal comparisons are epistemically problematic because it is impossible to obtain choice-based information which is similar to information obtained from an elicitation procedure such as the reference lottery technique (see, e.g., Hirshleifer and Riley 1992), and which would allow definition of interpersonal differences and similarities in utilities. Roy Harrod links this epistemic argument with the scientific respectability of economics (Harrod 1938, 395-396; see also Robbins 1938; 1952).

To motivate the discussion, consider some recent views about interpersonal comparisons. Mathias Risse argues as follows:

Often, for instance, we would like to be able to use information about the strength of preferences, or use more fine-grained judgments, say, as provided by a point system. Yet what matters for the choice of an aggregation rule is not merely what information should be used given the purpose of the decision process, but also what information one can *reliably* elicit. There are two concerns. On the one hand, the more fine-grained information we admit, the more

interpersonal comparability among individuals is problematic [...] The other problem is *manipulability* (Risse 2009a, 797; see also Risse 2004, 58).

Laslier (2012) reports the results of a poll conducted at a meeting of voting theory specialists concerning what they think about the performance of various voting rules. The participants were given the chance of providing their reasons for the choices they made. Here is one comment:

The Condorcet Principle takes into account only the ordinal preferences of every voter between any pair of alternatives because attempting to take into account also voters' cardinal preferences (as under the Range Voting procedure) would not only imply that a Condorcet winner⁵ may not be elected or, worse, that a Condorcet loser may be elected, but also that inter-personal comparisons of utility are possible and acceptable—which they are not! (Dan Felsenthal, in Laslier 2012, 339).

Risse and Felsenthal are thus arguing that some voting rules incorporate interpersonal comparisons, particularly ones that aim to collect information on preference intensity, and some others do not.⁶ I will now argue that such a view is based on a rather elementary confusion between voters' underlying preferences and expressed votes.⁷ I will proceed by discussing an example that Arrow used in arguing for the IIA condition. My concern here is not so much with IIA, but rather with showing how the failure to distinguish between vote ballots and preferences leads to the mistaken view that some but not all voting rules incorporate interpersonal comparisons.

The assumption of sincere voting makes it impossible to formally distinguish between preferences and vote ballots within the theory. The theory thus proceeds on the assumption that preferences directly affect the outcomes of voting rules. This, in turn, leads to the idea

⁵ A Condorcet winner is a candidate that has a majority against each of the other candidates, so that it beats any other candidate in a pairwise contest.

⁶ Risse and Felsenthal do not represent the views of all social choice theorists. Saari reports the following in a footnote: "Some readers of earlier drafts of this paper claim that the Borda Count has nothing to do with interpersonal comparisons, others claim it does" (Saari 1998, 258).

⁷ It is not that votes and preferences are never distinguished in social choice theory. Saari (2003) says that the former describe what *happens* with a particular rule, not what *should happen* based on the voters' preferences. See also Reynolds and Paris 1979; Austen-Smith and Banks 1999.

that different voting rules collect different information on preferences, and that the differences have to do with how voters may express their preferences. Thus the idea arises that different voting rules make different interpersonal comparisons. I will argue that this assumption is misguided because strategic voting implies that we cannot assume that the votes correspond to preferences. When we compare the outcomes of voting rules, insofar as the outcomes are evaluated by employing interpersonal comparisons, the comparisons are independent of the voting rule. By the end of section 4 below, I will come to the conclusion that interpersonal comparisons are methodologically legitimate if they concern the utility profile because the same profile is used to study different voting rules.

Arrow introduced IIA in order to restrict social-welfare judgements to those that may be the result of voting processes (e.g., Arrow 1973). Insofar as an aggregation rule is assumed to represent a voting rule, this is tantamount to identifying such judgements with the outcomes of voting choices. Arrow's theorem is thus concerned with evaluating the performance of voting rules.

Arrow's discussion concerns the so-called *zero-one rule*.⁸ Each person's utility of the most preferred alternative is normalised to one and the least preferred alternative to zero. The rule may be interpreted as a version of utilitarianism understood as an ethical doctrine, as well as a decision-making procedure.⁹ Interpreting the zero-one rule as a decision-making mechanism means that if there are three alternatives each voter 'i' will report his or her preference ordering by assigning one *point* to the best, zero points to the worst, and $v_i \in (0, 1)$ to the middle alternative. The outcome is determined by adding these points. The alternative with the largest number of points is selected. If there are only two alternatives, voters assign one point to the better one and zero points to the worse.

Arrow's example is shown in Table 1 below. Three individuals: 1, 2, and 3, have the following preferences for alternatives A, B, and C:

⁸ Isbell (1959) was the first to formalise this rule and to suggest that it could be used for interpersonal comparisons. Dhillon and Mertens (1999) provide an axiomatic justification for it. See also Sheng 1987; and Hausman 1995.

⁹ If it is interpreted as a decision-making mechanism, its functioning is equivalent to "utilitarian voting" (Hillinger 2005) and "range voting" (Smith 2000). Tangian (2000) calls it "Laplace's method".

Table 1: Arrow's example

1	2	3
A (1)	A (1)	B (1)
B (0.9)	B (0.9)	A (0.5)
C (0)	C (0)	C (0)

Arrow states that this zero-one assignment of utilities is “designed to make individual utilities interpersonally comparable” (1963, 32). He thus views the numbers in parentheses both as points that voters may express and as interpersonally comparable utilities.

Arrow presents two related arguments against this rule for making social choices. They are also arguments for IIA, and both are based on a modification of voter preferences. According to the first, deleting the worst alternative, C, from the set of available alternatives should not affect the final outcome. The utilities would be as in Table 2.

Table 2: Arrow's example with C deleted

1	2	3
A (1)	A (1)	B (1)
B (0)	B (0)	A (0)

Under the second modification, the first two individuals are indifferent between A and C, and the third now finds C indifferent to B, while the relative positions of A and B are unchanged in all individual orderings. “Then the assignment of utilities to A and B becomes the same as it became in the case of blotting out C entirely, so that again the *choice* between A and B is altered, contrary to Condition 3 [IIA]” (Arrow 1963, 32). The utility assignments after these changes are shown in Table 3.

Table 3: Arrow's example with changes in preferences

1	2	3
A C (1)	A C (1)	B C (1)
B (0)	B (0)	A (0)

IIA is violated because ordinal preferences between A and B remain unchanged, but social choice changes from B in Table 1 to A in Table 2,

and the social ranking between A and B changes between Tables 1 and 3. In Arrow's account, the utilities for the alternatives are first as follows:

$$A: 1 + 1 + 0.5 = 2.5$$

$$B: 0.9 + 0.9 + 1 = 2.8$$

$$C: 0 + 0 + 0 = 0.$$

Blotting C out of existence then changes them to:

$$A: 1 + 1 + 0 = 2$$

$$B: 0 + 0 + 1 = 1.$$

Changing the relative position of C changes them to:

$$A: 1 + 1 + 0 = 2$$

$$B: 0 + 0 + 1 = 1$$

$$C: 1 + 1 + 1 = 3.$$

Here is how Arrow makes his point:

If C were blotted out of existence, it should not make any difference to the final *outcome*; yet, under the proposed rule for assigning utilities to alternatives, doing so would cause the first two individuals to *have utility* 1 for A and 0 for B, while the third individual *has utility* 0 for A and 1 for B, so that the ordering by the sum of utilities would cause A to be *preferred* to B (Arrow 1963, 32).

It is customary to think that IIA is mainly concerned with ruling out information on preference intensities (or cardinal utility). However, it is clear that it “does not rule out the *existence* of preference intensities [...] The *construction* of these relative preference intensities may involve irrelevant alternatives, but IIA is still satisfied” (Hammond 1987, 200). Hammond's point is that it may be possible to *find out* about intensities for the various alternatives only by using information on “irrelevant” alternatives, but this has nothing to do with the existence of such intensities. Hammond suggests that preference intensities *underlie* the choices of voters (see also Mongin and d'Aspremont 1998; Ng 2000, 68). It is clear that such an interpretation treats intensities as existing independently of the voting rule, and that the latter is not interpreted as a preference-elicitation procedure. It is instructive to think of intensity as referring to some substantial notion of utility such as desire

satisfaction or feelings of pleasure. It seems to make sense to say that the substantive utilities of voters would be different under different outcomes, and that they are independent of their choices in the counterfactual sense that a voter's substantive utility could be assumed to be the same if A (or B or C) was selected, irrespective of whether he or she voted for it.

If 'utility' refers to something substantive, the numbers we use to describe voter preferences remain fixed irrespective of the choices that they may be able to make. In contrast, according to the modern conception of utility (such as von Neumann and Morgenstern's), utility numbers represent a person's preferences or choices, and are not connected to any substantive notions. Ultimately, the utility numbers are only based on choices because they are supposedly constructed in choice experiments such as the reference lottery. It is consistent with this viewpoint to say that voters 1 and 2 have utility zero for outcome B in Table 2, even though they had utility 0.9 in Table 1. (Since the utility numbers are not unique, however, one could just as well use any two numbers such that the number for A is larger than that for B.) However, if the utility number 0.9 refers to substantive preference intensity, it makes no sense to say that it disappears when we move from Table 1 to Table 2: it is just that voters are no longer able to *express* this preference intensity by directly reporting a cardinally measurable number.

If a voter gives her vote to B when A and C are available, we cannot even say, merely on the basis of this choice, that she prefers B to A and C, because she may have voted strategically. If such a choice were taken to define her preferences, it would define them falsely. Whatever plausibility revealed-preference arguments have elsewhere, in the voting context they are hardly compelling, and the whole idea of utility as a representation of preference cannot be meaningfully applied. I do not intend to claim that social choice theory is committed to revealed preferences. My point here is rather that if one were to appeal to such an idea so as to defend the identification of preferences and votes, it would be a highly unsatisfactory defence.

Arrow analyses the example as if there were first an intensity of preference between A and B which, however, suddenly disappears when C is deleted. It is obvious, however, that he did not intend to claim that the underlying preference intensities had actually changed. Had they done so, there would have been no reason to worry about the change in

the outcome. If the utilities in Tables 1, 2, and 3 are interpreted as substantive preference intensities, a utilitarian would accept the change in outcome from B to A, precisely *because* they change. It is true that, if we take Arrow literally, he seems to be saying that cardinal preferences between A and B change while ordinal preferences remain the same. However, this would be misreading his obvious intention. His point is clearly to show that this method of voting yields unsatisfactory outcomes, because voters are first able to express cardinal preferences for three alternatives, but after the change they may only express their ordinal preferences for two alternatives.

Let us now see how the example should be analysed if it is assumed that the utilities in Table 1 represent intensities of desire satisfaction that underlie voters' choices. Table 4 displays the preference intensities for the two remaining alternatives after the elimination of C.

Table 4: Arrow's example modified

1	2	3
A (1)	A (1)	B (1)
B (0.9)	B (0.9)	A (0.5)

Table 5 displays similar intensities following the second modification.

Table 5: Arrow's example modified and remodelled

1	2	3
A C (1)	A C (1)	B C (1)
B (0.9)	B (0.9)	A (0.5)

These tables show that B is more preferred than A in utilitarian terms irrespective of whether C is included or not in the set of alternatives. It is an altogether different matter that voters would be able to *express* only their ordinal preference if there were only two alternatives. A utilitarian would then say that B *should* be selected but that A *is* selected.

Arrow's example does prove what it was intended to prove, viz., that the zero-one rule violates IIA if it is interpreted as a decision-making

mechanism.¹⁰ The problem with this analysis is rather that Arrow seems to assume that, since voters may express points with a cardinal significance under the zero-one rule, these choices simultaneously provide information about their *interpersonally comparable utilities*.

3. CAN INTERPERSONAL COMPARISONS OF UTILITY BE INCORPORATED INTO VOTING RULES?

Social Welfare Functionals (SWFLs) are the standard tool for investigating interpersonal comparisons in social choice theory (see, e.g., d'Aspremont and Gevers 2001; Bossert and Weymark 2004; Blackorby and Bossert 2006). In Sen's proof of the impossibility theorem with cardinal utility functions (1970, 128-130), non-comparability of utility means that the ordering provided by the social welfare functional is the same irrespective of which linear transformation of the utility profile is used. This means that even though the SWFL approach discusses interpersonal comparability, no judgement is ever formulated on the candidates (or the choice alternatives) in terms of a *given* utility or preference *profile* alone, but rather only insofar as intensities or interpersonal comparisons affect the SWFL. The essence of Sen's proof is that, since the SWFL yields the same ordering irrespective of which linear transformation of the utility profile is used, cardinal utilities provide the same information for it as ordinal utilities. It follows that cardinal utilities do not provide any information additional to that derived from preference orderings. An impossibility result can then be proven with cardinal utilities, because it can be proven with ordinal ones.

Hammond and Fleurbaey (2004) criticise this approach for its failure to provide an unambiguous procedure for *embodying* or *incorporating* the interpersonal comparisons into the SWFL. I will now argue, in the spirit of Hammond and Fleurbaey, that voting rules cannot incorporate interpersonal comparisons even in principle. In other words, voting rules can never be represented by SWFLs. Incorporating interpersonal comparisons in a voting rule could be taken to mean one of two things: that people may somehow express interpersonal comparisons when they vote, or that when the outcomes of a voting rule are evaluated

¹⁰ This claim must be qualified slightly. Bordes and Tideman (1991) argue that IIA is defined for a situation in which the set of alternatives is fixed and where the preference profile changes. Arrow's first modification thus does not qualify as a violation of IIA because it is undefined, but the second one does.

normatively, it is necessary to make interpersonal comparisons. I will show that the first interpretation is false, and that while the second interpretation is sensible if the normative evaluation is to be based on the utility profile, the interpersonal comparisons are not different under different voting rules.

Voting choices cannot express interpersonal comparisons. The argument for this claim is quite simple. It follows from Myerson's argument that it is not possible to choose to have alternative A selected in a voting rule with *i*'s preferences rather than alternative B with *j*'s preferences.¹¹ Hence, interpersonal comparisons cannot be expressed in terms of individual voting choices, even in principle (see Sen 1999, 355), because each individual's choices only provide information on his or her preferences, but cannot yield any interpersonal information.

Arrow's initial idea was presumably that the zero-one rule violates IIA and incorporates interpersonal comparisons, but that the so-called *method of majority decision* does not incorporate interpersonal comparisons and does satisfy IIA.¹² This method is defined as follows. If a majority of voters prefer alternative A to B, then A is higher in the social-preference relation than B. Clifford Hildreth argued that the method of majority decision involves an interpersonal comparison "because it chooses between A and B by comparing the number who would gain utility in passing from A to B with the number who would lose" (Hildreth 1953, 91). He concluded that the issue was not whether we could do without interpersonal comparisons, but rather what sorts of comparisons we would be willing to make. Hildreth thus argued that all voting rules involved them, and that different rules incorporated different comparisons.

Different voting rules could be taken to incorporate different interpersonal comparisons if the utilities expressed by voters could be aggregated by assuming that the method of adding points or scores would incorporate the relevant comparability assumption. The zero-one rule, conceived of as a decision-making rule, would thus incorporate

¹¹ Even though I apply Myerson's argument here, I do not wish to deny the possibility of making such comparisons by claiming that they cannot be elicited. After all, everybody—including Lionel Robbins (1981)—admits that we make interpersonal comparisons in our everyday lives. Furthermore, our knowledge of interpersonal comparisons is in fact based on choices (Little 1957; see also List 2003). It is just that these choices are more difficult to interpret than those made in a reference lottery experiment.

¹² He never states this explicitly, but it seems to me to be a fair interpretation of his intention.

interpersonal comparisons of intensities if the points were identical with the preference intensities. However, the example discussed above shows that we have reason to believe that this is often not the case if intensities are interpreted in a substantive way.

Hildreth seems to be arguing that, since voters can only express their preference orderings (or rather, pairwise preferences) under the majority-decision method, but can also express preference intensities under the zero-one rule, then utilities are cardinally comparable in the latter and ordinally comparable in the former. And indeed, if one assumes that the vote ballots correspond to utilities, then the zero-one rule incorporates interpersonal comparability of preference intensities because voters are allowed to express their *intrapersonal* intensities “directly”, that is, without misrepresenting their preferences.

We only have knowledge on the expressed points in the voters’ ballots. But then, given that the vote ballots are supposed to reflect the underlying preferences, social choice theorists assume that they correspond to one another. Hence they arrive at the judgment that the zero-one rule is problematic as concerns interpersonal comparisons because the theorist must assume that the preference intensities are interpersonally compared when the voting rule computes their sum: the weight of each individual voter is assumed to be the same.

However, voters do not express their interpersonal or even intrapersonal intensities of preference directly because they only report *points* or *numbers* that may or may not have a one-to-one relation to them. The problem with Hildreth’s argument is thus that it blurs the distinction between the kind of information voters are able to express sincerely with their votes, and the kind of interpersonal comparisons that are required for interpreting or evaluating the outcomes under various voting rules. It may be natural to conclude that different voting rules come with different comparability assumptions if one adopts the standard social-choice assumption that the votes or points correspond to individual utilities. After all, if the points do correspond one-to-one to voters’ utilities, it seems as if the rule is making an interpersonal comparison because it is adding points across individuals. Furthermore, given that there is no other way of obtaining information on interpersonal comparisons of utilities in the voting context, some social choice theorists have assumed that the rules are making the comparisons.

The main problem with this view is that if the rules were to be making the comparisons, they would be wildly inaccurate in every rule rather than just in the zero-one rule. If strategic voting is taken into account, every voting rule is likely to reflect at least some intensity information at the aggregate level. Hildreth's conclusion that different voting rules make different interpersonal comparisons is false if the expressed votes do not necessarily correspond to their utilities.

I will now show that the votes do not necessarily correspond to utilities. Consider how voters might act under the zero-one rule. Let us assume—applying the logic of Ordeshook and McKelvey (1972)—that voters rank the candidates in terms of expected gains. Let p_{jk} denote the probability of being pivotal between candidate j and k , and let U_j denote the utility for candidate j . The expected gains for j are given by $E_j = \sum_{j \neq k} p_{jk}(U_j - U_k)$. It seems safe to assume that in Arrow's example, the pivot probabilities p_{AC} and p_{BC} are equal to zero because C never has a chance to win. For voters 1 and 2, the expected gains are then $E_A = p_{AB}(U_A - U_B)$ and $E_B = p_{AB}(U_B - U_A)$. As long as p_{AB} is positive, E_A is positive but E_B is negative. If voters give positive points only to those candidates for whom the expected gains are positive, they might end up giving a *bullet vote* to A so that the vector of points is $(1, 0, 0)$ for (A, B, C) . Bullet voting means giving maximum points to one candidate and zero to the rest even though the voting rule allows for expressing more fine-grained intensity information. If this is how voters behave, the outcomes from the zero-one rule end up depending on the same kind of information that the plurality rule would take as input under the assumption that voters are sincere.

Voting rules such as range voting or the zero-one rule, which might yield precise information on preference intensities under sincere behaviour, may end up providing only top-preference information under strategic behaviour, and voting rules such as plurality rule that only collect top-preference information under sincere behaviour may end up providing rich intensity information under strategic behaviour.¹³ Even though each individual voter's choice only indicates ordinal preferences

¹³ There are other voting rules that allow voters to provide at least some intensity information such as *Approval voting* (Brams and Fishburn 1983) and *Majority judgment* (Balinski and Laraki 2007; 2010). I have shown (Lehtinen 2010) that there can be bullet voting in the former which is problematic under behavioural heterogeneity. I have thus far not been able to construct a model of strategic voting under majority judgment, and I have not been able to study whether there could be bullet voting in this rule. Whether its relative resistance to manipulation is an argument in its favour or against it is thus an open question.

in many voting rules (except under the aforementioned forms of utilitarian voting), preference intensities affect the results in all voting rules. This, in turn, follows from the fact that the conditions for strategic voting always depend on preference intensities (e.g., McKelvey and Ordeshook 1972; Enelow 1981). Thus, even though individuals provide their information in ordinal form, and even though the voting rule provides only a ranking or a single winner, the outcomes do depend on preference intensities through strategic voting. No voting rule is thus able to only depend on ordinal information. If the kind of information voters are able to express with their strategic votes were to determine the kind of interpersonal comparisons that a voting rule makes, this would mean that every voting rule makes interpersonal comparisons of preference intensities. The absurdity of this claim should make it easier to understand that the kind of information that voters are able to express, or the form in which vote information enters or exits the voting rule, have nothing to do with interpersonal comparisons of utilities.

It follows that one cannot reasonably define interpersonal comparisons on the basis of the kind of information that voters are allowed to express, or on the basis of the kind of information that ends up affecting the results in any voting rule. If voting theorists make interpersonal comparisons, they do so in order to be able to interpret voting outcomes in welfarist terms. We have seen that there is no reason to think that they are different under different voting rules, merely because voters may provide different information about their preferences sincerely under different rules. Interpersonal comparisons thus have nothing to do with the various voting rules, and different voting rules do not incorporate different interpersonal comparisons. It follows that neither IIA, nor any other condition could *rule out* interpersonal comparisons because it would be a category mistake to claim that it does: IIA concerns the relationship between individual preferences and the outcomes of aggregation rules but interpersonal comparisons have nothing to do with the latter.

What, then, do interpersonal comparisons concern if they cannot concern the voting rule or the relationship between the utility profile and the voting rule? I take it to be obvious that they must concern the utility profile itself. Now that we have put the object of interpersonal comparisons in its proper place, let us consider why the impossibility of obtaining reliable information on interpersonal comparisons does not

provide an insuperable obstacle to using them in evaluating voting rules normatively.

4. WHAT FOLLOWS FROM THE EPISTEMIC ARGUMENT AGAINST INTERPERSONAL COMPARISONS OF UTILITY?

As we have seen, the epistemic argument (Myerson 1985) against interpersonal comparisons posits that there is no choice-based procedure for obtaining information on interpersonal comparisons. Furthermore, it is often argued that it is more difficult to obtain information on preference intensities than on orderings, and I presume that this is the reason why social choice theorists consider Risse and Felsenthal's view plausible. Furthermore, the literature on social welfare functionals has studied various comparability assumptions and shown how different social choice functionals require different comparability assumptions. The rationale for studying SWFLs—in other words different notions of comparability—then hinges on there being such epistemic differences between the various aggregation rules.

Just like many other writers on interpersonal comparability, I do not think that it is possible to obtain reliable information on interpersonal comparisons of preference intensities, and that interpersonal comparisons of intensities are epistemically more difficult than, say, comparisons of levels of utility. However, I will now argue that the injunction against making interpersonal comparisons (of intensities) in evaluating voting rules simply does not follow from the epistemic argument.

We could interpret Arrow's treatment of the zero-one rule by suggesting that the social choice theorist makes the interpersonal comparison once she assumes that the points correspond to voters' utilities. One way of making sense of the idea that some voting rules require more precise information on interpersonal comparisons than others would be to argue that in order to use the zero-one rule, we would have to obtain information on interpersonal comparisons of preference intensities. But since the zero-one rule is highly manipulable, the rule does not provide reliable information about preference intensities. This is the basis for Risse's claim—I believe—that we cannot obtain reliable information on preference intensities. I agree with Risse that trying to elicit information on preference intensities with rules like range voting or the zero-one rule may be very difficult. But the reason for this is not that such rules may be more susceptible to manipulation

than other rules. I believe they are, but that is not the main shortcoming of such rules. It is rather that strategic voting leads to bullet voting, and such voting rules may thus end up giving us less intensity information than other rules which ostensibly only collect information on preference orderings or top preferences!

Arrow thought that IIA guarantees the observability of preference orderings, and indeed it does this in his original framework because it assumes sincere behaviour. Thus, there would be a case against preference intensities and interpersonal comparisons if people never voted strategically. Note that the majority rule is not manipulable if it is used to choose from two alternatives. However, in order to apply it (or “method of majority decision”) when there are more than two alternatives, one has to specify how exactly one is to make all the pairwise comparisons, and how exactly one is to use the information from such comparisons. This is why the majority rule translates into a large number of aggregation rules when there are more than two alternatives (e.g., Borda, Copeland, amendment agendas, Kemeny). I have shown (see Lehtinen 2011) that the outcomes depend on preference intensities under amendment agendas.¹⁴ Such agendas are a common way of implementing majoritarian voting, and they satisfy IIA. This means, however, that if voters act strategically, satisfying IIA does not mean that a voting rule precludes preference intensity information or that satisfying IIA guarantees observability of preference orderings. Moreover, IIA itself does not preclude strategic voting (see Saari 2008, 60; and Lehtinen 2011 for a detailed argument). Risse thus proceeds on the false presupposition that majoritarian voting rules are not affected by preference intensities. Although these rules do not explicitly ask voters about their individual intensities, they do affect the outcomes. Of course, if *all* of the conditions for Arrow’s theorem were satisfied, preference orderings would be observable. But they are not satisfied, even when the preference profile is not cyclical (whether or not IIA is satisfied) as soon as any voter votes strategically, and we know from the Gibbard-Satterthwaite theorem that this is always possible if there are at least three alternatives.

Social choice theory is able to entertain the illusion of observability of orderings only by assuming the problem away by assuming that voters never vote strategically. This leads to the assumption that there

¹⁴ With three alternatives, for example, such an agenda might first pit A and B against each other in a pairwise majority contest, and then the winner against C.

is a one-to-one relationship between a voter's preferences and votes. This assumption, when combined with the use of aggregation rules, gives rise to the idea that different voting rules embody different interpersonal comparisons.

Let me conclude this section with a reflection on how interpersonal comparisons can legitimately be used in studying voting rules. Suppose that a voting theorist studies the relative performance of different voting rules, and makes interpersonal comparisons so as to be able to judge the alternatives normatively. Such comparisons are made in a *methodologically satisfactory* way if they do not favour any voting rules as compared to others. The same interpersonal comparisons can be applied in all voting rules if the same utility profiles are used to study different voting rules. Hence, they can be made in a methodologically satisfactory way.

Consider now a possible way of interpreting this proposal to use interpersonal comparisons of preference intensities in evaluating voting rules. It might be taken to imply that one should make interpersonal comparisons in studying real elections. Suppose, for example, that voting theorists tried to find the *best candidate* in, say, the U.S.A. presidential elections in 2000 by postulating specific cardinal and fully comparable utilities for the millions of voters. It is obvious that such an enterprise would be preposterously misguided because we do not have such information on individual intensities, and we cannot have such information. But then again, we do not have information on preference orderings either. The reason is—as we have seen—that the possibility of strategic voting makes it impossible to obtain reliable information on individual preferences from voting choices: votes are likely to differ from the underlying preferences.

Finding the underlying preferences *from* the votes would require quite some knowledge: since preference intensities have been shown to affect strategic voting, we would need to know what they are, but we already agreed that this is even more difficult than finding out about the preference orderings. But this would not suffice either because we would also need to know the exact distribution of behavioural dispositions because some people refuse to vote strategically even when they have incentives to do so. And as Riker (1982) argued, we have no way of finding out who votes strategically. Finding out about the preference orderings in real elections is thus just as impossible as finding out about interpersonal comparisons. From this perspective,

it is easy to understand why social choice theorists occasionally note that they do not know what the preference profile is in any real election.

These considerations explain why social choice theorists do not use real-world data to evaluate their theories.¹⁵ Thus, for example, when Saari (2003) talks about preference profiles as “data”, he refers to profiles that the social choice theorists consider as arguments for an aggregation rule, but these “data” are constructed or postulated by the theorist rather than given by observations. What can be given in observations are the votes.

Whatever is the cogency of observability arguments for ordinal utility in other contexts, the epistemic argument should not have much weight against making interpersonal comparisons of preference intensities *in voting theory*, because the field is, and will be, dominated by approaches that do not use any kind of *real* data on actual individual preferences. Social choice theorists thus apply observability criteria selectively: they argue that interpersonal comparisons should not be made on observability grounds, even though they never use any real-world evidence and there cannot be such evidence,¹⁶ because there are no reliable observations on preference orderings in the voting context.

In my work on strategic voting (Lehtinen 2007a; 2007b; 2008), the utility profiles are randomly generated, and the interpersonal comparisons are thus also random but the same under every voting rule because they concern the utility profile. Such random interpersonal comparisons are of course methodologically legitimate in the sense indicated above. If one is concerned over the justification for using this or that particular interpersonal comparison, it is surely some consolation that making different interpersonal comparisons of preference intensities with a given ordinal profile does not really affect the results: thus far they have always been robust to different interpersonal comparisons (see Lehtinen 2007b; 2008). Such robustness results indicate that whether one uses this or that particular comparison is unlikely to make a difference to the normative evaluation of voting rules.¹⁷

¹⁵ Attempts to measure the amount of strategic voting in some real elections constitute an exception. However, the fact that various authors disagree about this amount is testimony to the underlying observability problem.

¹⁶ Tideman and Plassmann (2014), however, use a dataset of actual votes.

¹⁷ I do not wish to discourage those interested in voting theory from experimenting with different interpersonal comparisons. It would be an important contribution to show that some voting rule yields satisfactory results only under some peculiar interpersonal comparisons. However, given that interpersonal comparisons do not

5. DEBATING ABOUT CONDITIONS

Let us now have a closer look at the idea of evaluating aggregation rules by the set of conditions that they satisfy. The conditions determine how the outcomes under aggregation rules should depend on individual preferences under *any possible* configuration of individual preferences, and violations of conditions are demonstrated by way of giving theoretical counterexamples. Given that social choice theorists never (except in the case described below) pass judgment on which alternative or candidate would be the best on the basis of a *given* utility or preference profile, i.e., independently of the aggregation rules that represent various voting rules, they never have to make interpersonal comparisons. Indeed, if interpersonal comparisons are not to be made, evaluating conditions satisfied by the aggregation rules becomes the only way in which voting can be normatively evaluated.

In contrast, defining the best alternative in terms of a given preference or utility profile requires making interpersonal comparisons. Arrow argues that evaluating the alternatives is tantamount to using a decision-making process: welfare judgement means “an evaluation of the consequences to all individuals based on their evaluations [...] The process of formation of welfare judgements is logically equivalent to a social decision process or constitution” (Arrow 1983 [1967], 68). Since the conditions concern the relationships between individual preferences and the outcomes of aggregation rules, there is no need to pass judgement on which alternative is the best, in any given situation, in terms of the preference profile alone, independently of how the profile affects the aggregation rule. In other words, one never evaluates *the alternatives* on the basis of their utility (or preference) characteristics in the profile. All the normative work is done, and must be done, by the conditions that specify how the characteristics of any abstract set of individual preferences should translate into social orderings or choices via the aggregation rule.

In social choice theory, the conditions never specify the relationship between the profile of individual preferences and the actual outcomes of voting rules when behavioural assumptions are taken into account. They rather specify the conditions between the preference profile and the aggregation rule under the assumption of sincere voting. The theory

have anything to do with how voters vote or how the voting rules compute the outcomes, I would be rather surprised if someone could establish such results.

proceeds on the assumption that satisfying or failing to satisfy various conditions tells us about the kind of information that is being collected by the different rules. Some rules (e.g., the Borda count) are supposed to obtain information on preference orderings, some others (e.g., agenda voting) on pairwise comparisons, some on approval (approval rule), some on top preferences (e.g., plurality rule), and some on preference intensity (e.g., range voting, the zero-one rule, and perhaps majority judgment).

I will now show that, since strategic voting breaks the connection between preferences and choices, it is no longer valid to infer that a voting rule reflects a certain kind of information if the corresponding aggregation rule satisfies a given condition. The problem is that different voting rules reflect different kinds of information only on the assumption of sincere behaviour. The conditions are imposed on aggregation rules rather than voting rules, but we should be more interested in conditions imposed on the latter because they take behavioural assumptions into account. While different aggregation rules may satisfy different conditions, no voting rule satisfies any condition that makes a difference to the comparative normative evaluation of such rules.

In order to show the relevance of the problem, I will discuss a recent debate between Donald Saari (2000; 2003; 2006) and Mathias Risse (2001; 2004; 2005; 2009a; 2009b) which concerns—at least seemingly—whether the Borda count or the Kemeny rule (or as Risse calls it, the “Condorcet proposal”) is a better voting rule. My main interest does not lie in discussing the relative merits of the two voting rules, but rather in showing with a contemporary normative debate on voting rules how problematic it is to evaluate voting rules on the basis of whether or not the corresponding aggregation rules satisfy various conditions.

Saari’s main line of argument is that one has to be able provide conditions that make a difference between voting rules, and that this can only be done with an appeal to the “data”, the preference profile: “We should try to use the information from the full profile to determine the will of the people” (Saari 2003, 342). If the profile of voters’ preferences has symmetrical elements called *Condorcet n-tuples*, the voting rule should yield a complete tie with regard to those preferences. Then, removing from the profile voters that are involved in such n-tuples shows the difference-making voters. In other words, any voting rule should provide the same choice or ranking with and without such

voters. Saari presents two new conditions that any voting rule should satisfy. Let ABC denote the ordering A is better than B, and B is better than C. The Neutral Reversal Requirement (NRR) states that pairs of voters with orderings ABC and CBA should be removed from the profile.

A *Condorcet triplet* (i.e., an n-tuple in the case of three candidates) is constructed as follows. Take the most preferred candidate from the ordering ABC, and place it last so as to get BCA, then do the same to this ordering so as to get CAB. Note that each candidate is at each position exactly once in the triplet {ABC, BCA, CAB}, so that in this sense, no candidate is more preferred than another. The Neutral Condorcet Requirement (NCR) states that such cyclic elements should be removed from the profile. NCR is the condition that makes a difference to the evaluation of Borda vs. Kemeny: The Borda count satisfies it, but the Kemeny rule does not (and on this point Saari and Risse agree).

Consider now the profile that Saari and Merlin (2000) employed to demonstrate that the Kemeny rule does not satisfy the NCR:

Table 6: Example from Saari and Merlin (2000)

voter type	1	2	3	4	5	6
number of voters	6	0	3	5	0	5
	A (1)	A	B (1)	B (1)	C	C (1)
	B (.4)	C	A (.1)	C (.5)	B	A (.9)
	C (0)	B	C (0)	A (0)	A	B (0)

The outcomes are determined as follows under the Kemeny rule. If there is a Condorcet winner, the rule selects it. If not, the rule aims to find the best ranking by looking at how many pairwise reversals of preferences would be required to yield a transitive social ranking. In this example, 11 voters prefer A over B, and 8 have the reverse preference. Similarly, there are 9 voters who prefer A to C, and 10 with the reverse preferences; and 14 voters who prefer B over C, and 5 voters with reverse preferences. These pairwise preferences constitute a cycle {AB, BC, CA}, and the social ordering is ABC because it requires that only one type 6 voter reverses her ordering of A and C.

There are two Condorcet triplets. {ABC, BCA, CAB} constitute one cycle and {ACB, CBA, BAC} another. The number of voters with orderings ABC, BCA and CAB are 6, 5, and 5, respectively. Since the smallest of these numbers is 5, we remove 5 voters from these three voter types so that the remaining profile is the following.

Table 7: Applying the NCR

1	2	3	4	5	6
1	0	3	0	0	0
A (1)	A	B (1)	B	C	C
B (.4)	C	A (.1)	C	B	A
C (0)	B	C (0)	A	A	B

BAC is thus the social ordering after removing the cyclic elements, but this means that the NCR is violated. However, consider now strategic voting. If type 3 voters realise that sincere voting is likely to yield ABC, they may vote strategically by reporting BCA rather than BAC so as to lower A's position. If they do, the reported rankings are:

Table 8: Strategic orderings

1	2	3	4	5	6
6	0	0	8	0	5
A	A	B	B	C	C
B	C	A	C	B	A
C	B	C	A	A	B

Then the pairwise comparisons are:

A over B: 11 vs. 8

A over C: 6 vs. 13

B over C: 14 vs. 5

The social ranking is now BCA because this ordering only requires reversing two voters' AB preference, the social choice would be B, and type 3 voters would be satisfied. In other words, the Kemeny rule does not violate the NCR with this profile if voters of type 3 act strategically in this way.

Consider now the Borda count. The Borda scores are computed as follows. With n candidates, the most preferred candidate obtains $n-1$ points, the second-best $n-2$, ..., and the least preferred 0 points. The Borda scores are 20, 22, and 15 for A, B, and C, respectively, so that B is selected. However, if, say, three type 6 voters vote strategically by reporting ACB rather than CAB, the Borda scores are 25, 22, and 10 for A, B, and C, respectively, and A is selected instead of B. But now the

Borda count no longer satisfies the NCR because removing the voters within the Condorcet triplet {ABC, BCA, CAB} yields B as the winner. Note, however, that type 3 voters could thwart such a strategy because they could report BCA instead of BAC:

Table 9: Strategic orderings under the Borda count

1	2	3	4	5	6
6	3	0	8	0	2
A	A	B	B	C	C
B	C	A	C	B	A
C	B	C	A	A	B

If they do, the Borda scores are again 20, 22, and 15, for A, B, and C, respectively. Suppose, however, that they do not realise that they should do this because all type 6 voters seem to speak in favour of C.

To summarize, we have seen that the Kemeny aggregation rule violates the NCR, that the Kemeny voting rule does not violate the NCR, that the Borda aggregation rule does not violate the NCR, and that the Borda voting rule violates the NCR, and all these results were derived from a single profile. What difference does it make, then, that the Borda aggregation rule satisfies the NCR, but the Kemeny aggregation rule does not? This is not an argument for Kemeny or against Borda because, as far as I understand, Risse's main argument is that the Kemeny rule finds the "best ranking". If ABC is the "best ranking", the Kemeny voting rule does not find it in our example with strategic voting.

Breaking the connection between preferences and choices through strategic behaviour means that satisfying a condition does not necessarily make a voting rule normatively more attractive than another. If the conditions that are currently imposed on aggregation rules were to be formulated in terms of the relationship between the preference profile and the actual voting outcomes, *all* of the conditions that *make a difference* between different voting rules could be violated in any voting rule.

There are conditions like anonymity, neutrality, and universal domain, which cannot be violated through strategic voting, and which most voting rules satisfy. It is clear that my claim does not concern such conditions. Tideman (2006, 151) provides a classification of five different kinds of condition: domain, consistency, responsiveness, stability, and qualitative attractiveness. My claim concerns conditions in

the consistency (e.g., Condorcet, majority, transitivity), responsiveness (e.g., positive responsiveness), and stability (e.g., monotonicity and various independence conditions, perhaps strategy-proofness) categories.

Social choice theory does provide some useful information about the functioning of voting rules because, for example, some rules may only yield satisfactory results with single-peaked preferences. Violations of neutrality and anonymity may well be important for some questions. However, I claim that social choice theorists' disagreements over the most common voting rules can be traced back to different views concerning the relative importance of various difference-making conditions. My point is that even if one could agree on their relative importance, whether a given condition is satisfied by some aggregation is irrelevant because the corresponding voting rule will always violate it.

Although I have not provided an exhaustive list of examples in which *all* the various conditions are violated, such examples can easily be generated for any given rule and condition. Violating a given condition under given voting rule with a specific preference profile always requires that at least two alternatives are at least somewhat popular. If a profile of ordinal preferences is proposed and some alternative, say A, is shown to be selected under sincere behaviour, one can always find a way in which another alternative B becomes selected by introducing high (intrapersonal) preference intensities for B and low intensities for A. The reason for this is that in all voting rules, alternatives for which intensities are high on average (i.e., for which the sum of utility is large) are more likely to obtain strategic votes than other alternatives (Lehtinen 2007a; 2007b; 2008).

It is impossible to provide a general proof of my claim because the conditions concerning utilities and beliefs under which voters act strategically are different under each rule.¹⁸ Yet I am not willing to withdraw from the strong formulation of the claim that involves all voting rules and all difference-making conditions because in every rule that has been studied thus far, the utilitarian winner obtains a lot of strategic votes. Furthermore, given that strategic voting always depends on preference intensities and on voters' incomplete information, it is to be expected that similar results hold under all voting rules.

¹⁸ Suppose that there are 20 genuinely different voting rules and 20 genuinely different difference-making conditions. It would then take 400 examples of the kind described in this paper to prove that all difference-making conditions can be violated.

These results mean that, as long as the conditions are formulated in terms of ordinal preferences, one can always generate a violation of any condition (which is formulated between real preferences and outcomes of voting rules) by tweaking voters' intensities or behavioural dispositions. I thus challenge social choice theorists to come up with just one difference-making condition and a voting rule in which that condition is not violated.¹⁹

One might argue that evaluating voting rules on the basis of conditions that the corresponding aggregation rules satisfy makes sense despite my criticism because there are cases in which a condition is violated only under strategic behaviour, but also under sincere behaviour under another voting rule. Thus, a given condition may be more frequently violated under one rule than under another. However, even this argument is questionable because a voting rule may violate a condition under sincere behaviour with a given preference profile, but satisfy it under strategic behaviour with exactly the same preference profile. We have already seen an example of this with the Kemeny rule. For another example consider a violation of monotonicity from Tideman (2006, 191). 11 voters have the following preferences among three alternatives A, B, and C. I have supplemented the preference intensities in parentheses so as to provide a plausible account of strategic voting.

Table 10: Example from Tideman (2006)

3	4	3	1
A (1)	B (1)	C (1)	C (1)
B (.9)	C (.9)	A (.9)	B (.9)
C (0)	A (0)	B (0)	A (0)

Under the runoff rule, two alternatives obtaining the most votes in a first round of voting meet in a second-round pairwise contest (unless

¹⁹ It may be tricky for me to provide examples for some rules under which defining the strategic voting strategy is particularly complex. To the best of my knowledge, there are no models of strategic voting under the *Single transferable vote* (STV) and the *Majority judgment*. One way to deal with STV is to use a "plurality heuristic": assume that voters behave as if the voting rule were the *Plurality rule*, and use McKelvey and Ordeshook's (1972) pivot probability model to define expected gains for each alternative. The strategic rankings are then given by the ranking of expected gains. I have obtained some simulation results (under STV) which indicate that such behaviour is rational in the sense that, on average, voters that give such rankings are better off than they would be under sincere behaviour. At present, I do not know whether it makes sense to use the plurality heuristic in the majority judgment.

one candidate has a majority in the first round). This aggregation rule violates monotonicity:²⁰ A is eliminated in the first round of voting, and B wins the second round against C. However, if the voter in the last column puts B ahead of C, A and C tie for fewest votes, and if the tie is resolved in A's favour, it also wins B in the second round. Note carefully, however, that if this voter were to realise the consequences of her change of mind, she would see that changing her vote is against her own interests: instead of the tolerable B, the social choice will be the abhorred A, and she should vote strategically for C. If she does, B continues to win. In other words, strategic voting implies that monotonicity is no longer violated. The runoff aggregation rule thus violates monotonicity with this profile, but the runoff voting rule does not. Note that B beats C rather clearly 7-4 in the second round, and that the race between A and C is close in the first round. This voter's strategic choice is thus quite sensible.

Consider now the Borda aggregation rule which satisfies monotonicity. The Borda scores are 9 for A, and 12 for B and C. B and C thus tie for victory. If one of the four voters in the second column puts C ahead of B by reporting CBA, C becomes the sole winner and monotonicity is satisfied. What happens under strategic behaviour? If the three voters in the leftmost column vote strategically by reporting BAC, they avoid the abhorred choice of C. The Borda scores would then be 6, 14, and 12, for A, B, and C, respectively. B would thus become the sole winner and the Borda voting rule violates monotonicity. Again, the irony lies not in the fact that an alternative that has become more popular is no longer chosen, but rather in the fact that this happens in a voting rule that corresponds to an aggregation rule that *satisfies monotonicity*. It is true that the voters could vote strategically in this way even before the BCA voter changes her preference ordering. Note, however, that they have *systematically* stronger incentives to do so after the change: rather than a tie between B and C, the social choice would now be C under sincere behaviour. What difference does it make that the Borda social choice rule satisfies monotonicity and the runoff does not?

²⁰ A ranked voting system is *monotonic* if it is neither possible to prevent the election of a candidate by ranking them higher on some of the ballots, nor possible to elect an otherwise unelected candidate by ranking them lower on some of the ballots (while nothing else is altered on any ballot). In single winner elections that is to say no winner is harmed by up-ranking and no loser can win by down-ranking.

A voting rule V_1 cannot be better than V_2 on the grounds that the corresponding aggregation rule A_1 , unlike A_2 , satisfies a normatively compelling condition C if voting rule V_1 may violate C and voting rule V_2 satisfy C with exactly the same profile. As a result, social choice theory is unable to provide useful normative information on the relative merits of various voting rules.

Given the ease with which such examples can be generated, whether or not an aggregation rule satisfies a given condition does not really seem to matter. We do not seem to have any reason to assume that if a condition is violated under aggregation rule A_2 and satisfied under aggregation rule A_1 , it is more frequently violated under the corresponding voting rule V_2 than under V_1 . Comparative normative advice which is derived from social choice theory is always uncertain. All we can say is that no real voting rule ever satisfies the intuitive demands behind any condition that makes a difference. Social choice theory is able to entertain the illusion of being able to provide normative advice only by assuming away what prevents it, viz., by assuming that voters never vote strategically. We do not know which rule is more likely to violate this condition in real voting rules because nobody has studied the issue with a plausible model of strategic voting.²¹ I conclude that it is deeply problematic to evaluate voting rules in terms of the sets of normative conditions that they satisfy under the assumption of sincere behaviour because the inferences from social choice rules to real voting rules are highly unreliable due to strategic voting.

Different voting rules may be normatively compared as soon as we are willing to pass judgements on the alternatives merely on the basis of the utility profile. The best voting rules are then found simply by seeing which ones select the best alternative, i.e., the utilitarian winner, most frequently. This kind of evaluation, however, requires rejecting two important principles of social choice theory. First, in order to evaluate the alternatives on the basis of the utility profile, one must be willing to make interpersonal comparisons of utility that concern the choice alternatives. Secondly, one must distinguish between preference and choice so as to be able to define different voter behaviours in different rules while retaining the same utility profile. The normative standard

²¹ It can be done, however. I have studied path-dependence, see Lehtinen 2015. Felsenthal provides such results under several rules, but he assumes that strategic voters have complete information, see Felsenthal 1990.

must be behaviour-independent because otherwise outcomes from either different voting rules or different behavioural assumptions cannot be compared. Comparability also requires that one must keep the utility profile fixed across voting rules.

In social choice theory, the standard for evaluating whether or not a given voting rule violates a given condition is whether one can find an *example* of a preference profile that violates it. There is no concern over how alarmingly unlikely it is that a condition is violated.²² An obvious remedy for this is to study many utility profiles. Studying thousands or millions of profiles across voting rules is possible, however, only with a computer simulation, and this will require rejecting another principle of social choice theory: that only analytical results are considered acceptable.²³

6. STRATEGIC VOTING IS BENEFICIAL

We have seen that taking strategic voting into account destroys the normative force of social choice theory in making comparative assessments of voting rules. Consider now the second main topic of social choice theory, namely, evaluating the general performance of voting rules in collecting information on individual preferences. Arrow's theorem and the aggregation functions have been interpreted as being concerned with *social decision mechanisms* such as voting rules, and with *social welfare judgements* (see Buchanan 1954; Sen 1977; 1995). Accordingly, there are two main interpretations of the concept of an aggregation rule and of Arrow's theorem. The first is that there is no optimal voting rule because cyclic preferences make path-dependent social choices, strategic voting, and agenda manipulation possible (e.g., Bordes and Tideman 1991). The theorem has thus been interpreted to mean that there is no ideal democratic system of voting. The second interpretation is that there is no such thing as a well-defined "social optimum" (e.g., Riker 1982, 137; Ordeshook 1986, 57). Sen's 1970 proof of an impossibility theorem with cardinal utilities is commonly taken to

²² There may be quantitative differences in how often various rules violate the various conditions, and some contributions have tried to compute violation probabilities from a variety of different profiles. For example of this kind of work, see Mbih and Moyouwou 2008. See also Nurmi 2012.

²³ Admittedly, the impartial culture assumption, which is used in many models that study strategic voting, is very likely to generate more strategic voting opportunities than there is in reality. However, since some studies that measure the extent of strategic voting indicate somewhat extensive strategic voting (e.g., Blais, et al. 2001), it would be odd to deny its importance on the basis of the data that we have.

imply that it is impossible to define a social optimum without making interpersonal comparisons.

Riker (1982) argued that the possibility of strategic voting renders voting results meaningless. More generally, Arrow's theorem and the Gibbard-Satterthwaite theorem are often interpreted as implying that there is no democratic voting rule that satisfies a set of intuitively acceptable conditions, and that all methods of aggregating individual preferences are unsatisfactory. The justification of democracy thus rests on a shaky conceptual foundation if this is the right interpretation of the theorems.

If there is anything that social choice theorists agree on, it is that strategic voting should be minimized if possible. Furthermore, it appears that social choice theorists have also managed to convince other political scientists and philosophers (e.g., Miller 1992; Dryzek and List 2002), and even those who are not sympathetic to social choice theory (Christiano 1995), that only distortion and chaos can come from strategic voting. The Gibbard-Satterthwaite theorem states that any single-valued strategy-proof voting procedure that considers at least three alternatives is dictatorial. A voting procedure satisfies strategy-proofness if no individual has an incentive to vote strategically.

While there is disagreement concerning which particular conditions are the most important, the general welfarist principle that the best voting rules should generate outcomes that reflect individual preferences is universally accepted. Yet, the arguments that have been presented in favour of strategy-proofness have been non-welfarist.²⁴ The arguments are non-welfarist because evaluating the welfare consequences of strategic voting explicitly requires making interpersonal comparisons. Any such evaluation must postulate two different behavioural assumptions, sincere and strategic, and a behaviour-independent utility-based normative standard with respect to which the outcomes may be compared. It is impossible to evaluate the welfare consequences of strategic voting without making interpersonal comparisons of voter utilities because the normative standard must be independent of voters' behaviour and rather based on the utility profile.

Cohen (1986) points out that Bentham distinguished between the utility principle and the majority opinions, believing that "general consent" provides the surest visible sign and immediate evidence of the

²⁴ For these arguments, see Kelly 1988, 103; and for an extensive critique, see van Hees and Dowding 2008.

sum of utility. I also take the social optimum to be given by the utilitarian winner and adopt an epistemic populist notion of voting: voting provides fallible evidence of the best candidate. Note, however, that Coleman and Ferejohn (1986) as well as Cohen argue that a unique general does not always exist. I have assumed that the utilitarian winner always defines the best outcome, and that the interesting question is whether, how often, and under which conditions the various voting rules find it. I promote adopting this stronger interpretation of epistemic populism in evaluating the performance of voting rules.

Studying the welfare consequences of strategic behaviour requires a kind of counterfactual comparison. One must deduce sincere and strategic behaviour from a single utility profile, and compare the outcomes against the normative standard. Such an exercise cannot be conducted in a traditional social-choice framework in which individual utilities or preferences are aggregated because aggregation does not allow the aforementioned disentangling of utilities from choices. In other words, studying the welfare consequences of strategic voting in a welfarist manner requires incorporating epistemic populism into the very structure of the model with which the evaluation is made.

As we have seen, strategic voting increases the frequency with which the utilitarian winner is chosen in all major voting rules (Lehtinen 2007a; 2007b; 2008). Trying to minimize strategic voting might lead to worse outcomes in terms of individual utilities. The most unanimously cherished condition among social choice theorists is then not normatively acceptable from the utilitarian point of view. The results of social choice theory with respect to strategy-proofness provide a fundamentally misleading picture of the performance of various voting rules in finding satisfactory alternatives in terms of individual utilities, both comparatively and absolutely.

Note also that the methods of social choice theory could not have produced the result that strategic voting is beneficial because it is an *emergent* outcome: it makes sense to impose conditions on the relationships between inputs and outputs in an *aggregation* only if we know that there are no emergent properties (see, e.g., Wimsatt 1997), but the very idea of aggregation is misguided in the context of voting because an invisible hand affects the outcomes (see Lehtinen 2009).

Consider now the consequences of the fact that strategic voting is beneficial from the point of view of Arrow's theorem. Given that all social choice results assume sincere behaviour, *there cannot be an*

acceptable set of conditions which is satisfied by any real-world voting rule if strategic voting is beneficial: Strategic voting implies that no difference-making condition can be satisfied by voting rules, but at the same time, the results of voting would be worse in utilitarian terms in the absence of strategic voting. Arrow's set of conditions, or any set of conditions that has been proposed in social choice theory cannot be normatively acceptable because it always includes the assumption that voters act sincerely.²⁵ This means that while Arrow's theorem and the Gibbard-Satterthwaite are logically impeccable, they fail to have the devastating consequences for democracy that have sometimes been attributed to them.

Donald Saari's (1995; 1998; 2010) arguments against interpreting Arrow's theorem too strongly are similar to mine in spirit but he targets a different condition. He argues that IIA is not normatively acceptable because it overrides transitivity and thus individual rationality. Satisfying IIA is problematic because it means that an aggregation rule must ignore information on the transitivity of preferences even though transitivity has been explicitly required, and thus that aggregation rules that satisfy IIA fail to recognise rational voters. Saari's conclusion is powerful if one accepts the standard aggregation rule framework. However, I do not think this argument has much force when considering voting rules. First, fortunately, rational voters can resort to strategic voting, and since rational voting always requires evaluating at least three alternatives, even voting rules that correspond to IIA-satisfying aggregation rules may indirectly take at least some transitivity information into account. Secondly, strategic voting under voting rules that correspond to aggregation rules not satisfying IIA may often mean that the orderings reported do not correspond to the actual preference orderings. Thus, even if a voting rule explicitly collects information in terms of orderings, there is no guarantee that it takes the correct transitivity information into account.

Note that if there is any condition that will be violated as a consequence of strategic voting, it is individual rationality (see Lehtinen 2011 for details). If individual i 's choice from a set of alternatives S is $C_i(S)$, and xR_iy denotes her preferences, individual rationality can be expressed as follows: $C_i(S) = \{x | x \in S: \forall y \in S: xR_iy\}$. It is this condition that is always violated, and if one agrees with me that strategic voting is very

²⁵ This is usually expressed in terms of *individual rationality*.

likely to be rather commonly beneficial, it is this condition that is not normatively acceptable.

Social choice theory has evolved into *implementation theory*.²⁶ This theory studies which normative conditions various institutions satisfy in various game-theoretical equilibria. The theory thus relaxes the assumption of sincere behaviour while retaining the endeavour to study normative issues by way of examining whether institutions (such as voting rules) satisfy various normative conditions.

Insofar as the theory is used to study voting, the aggregation rules²⁷ are taken to characterize the desirable properties of voting institutions. Therefore, just like in social choice theory, interpersonal comparisons need not be made because the very same conditions are studied and the theory never passes any normative judgments on a particular profile. The theory rather defines the desirable outcomes in terms of aggregation rules which satisfy various conditions. The basic idea is to study whether one can find a *game form* which always yields the same outcomes as those that sincere voting would generate under the aggregation rule in question. A game form, which is also sometimes called a 'mechanism', is a game without the players' preferences: it specifies the strategies available for the players and the consequences of various combinations of strategies. An aggregation rule is said to be implementable in β -equilibrium if, for any profile of preferences, there is a game form in which the strategic behaviour of individuals leads to the same outcomes as those that would ensue from sincere voting in the aggregation rule. Here β refers to the various possible solution concepts. For example, if the Borda aggregation rule were implementable with dominant strategies, then this would mean that the conditions satisfied by this rule would always be satisfied in a game form in which the players have dominant strategies.

The most celebrated result in this field is the so-called *revelation principle* (Gibbard 1973; Myerson 1979) which posits something like this: If an institution implements an aggregation rule, then there is a *direct mechanism* which also implements that rule. A direct mechanism is a game form in which truth-telling is equilibrium behaviour. To put it simply, and relating the point to voting institutions, the revelation principle posits that if there is an aggregation rule that is implemented

²⁶ For introductions to this theory, see Jackson 2001; or Austen-Smith and Banks 2005, chapter 3.

²⁷ These are often called 'social choice rules' in implementation theory.

by a game form (voting institution), then that aggregation rule can be implemented in sincere strategies. In other words, if some set of normative conditions can be satisfied in equilibrium, then those normative conditions can be satisfied in a game form in which the players' truth-telling is equilibrium behaviour.

The revelation principle is similar in spirit to what I have said about violating conditions in that it indicates that sincere voting at least can be compatible with satisfying various conditions, and I have argued that strategic voting implies that all other relevant conditions will be violated. It is this incompatibility of strategic voting with any condition together with the beneficial nature of strategic voting that generates the impossibility of a set of normatively acceptable conditions. Given that strategic voting implies that all other difference-making conditions will be violated, it is perhaps not so surprising that *none* of the aggregation rules that are commonly used in political elections are implementable in Nash equilibrium (see, e.g., Jackson 2001, 672). This is the reason why implementation theory is just as incapable of providing comparative normative advice on voting rules as social choice theory. This impasse in implementation theory may explain in part why scholars working on voting and welfare may have chosen to bite the bullet and continue studying normative issues related to voting with the traditional tools of social choice theory. Ultimately, this impasse derives from the same problem that afflicts social choice theory: strategic voting is incompatible with every other normatively desirable condition that aggregation functions might satisfy.

Note that the result that strategic voting increases the frequency with which the utilitarian winner is selected is intuitively in conflict with the revelation principle because it means that strategic voting may be necessary for obtaining normatively desirable results. Implementation theory would study a similar question by looking at whether the zero-one rule is implementable under some solution concept. It is clear that it is not implementable under any solution concept because full implementation would require that the utilitarian winner is always selected. This reminds us once again that satisfying just about any condition is too much to ask, and that obtaining genuinely useful results for a comparative normative evaluation of voting rules may require computer simulations.

One might also argue that any given set of conditions that precludes taking intensity of preference into account is not normatively

acceptable. I agree with social choice theorists that there is no optimal voting rule because there is no rule that is guaranteed to find the utilitarian winner. The correct conclusion from this is that trying to propose a set of conditions that would take preference intensities into account is also doomed to failure: no voting rule will satisfy a set of conditions that take intensities into account because strategic voting implies bullet voting under the zero-one rule or range voting, and this means that it is impossible to obtain intensity information reliably from those rules as well. Intensities are thus something that cannot be properly treated by imposing conditions on aggregation rules or on voting rules. However, they are out there, and the only way in which we can give them their proper normative importance is by making interpersonal comparisons.

7. WHAT KIND OF INTERPERSONAL COMPARISONS SHOULD ONE MAKE?

Let us finally consider what the interpersonal comparisons should be like in voting theory. In order to do so, let us go back to the dispute between Saari and Risse. Table 6 is reproduced here for ease of reference.

Table 6: Example from Saari and Merlin (2000)

voter type	1	2	3	4	5	6
number of voters	6	0	3	5	0	5
	A (1)	A	B (1)	B (1)	C	C (1)
	B (.4)	C	A (.1)	C (.5)	B	A (.9)
	C (0)	B	C (0)	A (0)	A	B (0)

Consider first whether violations of NCR are desirable in the example we considered in section 5. The sums of utilities are $6+(3 \cdot 1)+(5 \cdot .9)=10.8$ for A, $(6 \cdot .4)+3+5=10.4$ for B, and $(5 \cdot .5)+5=7.5$ for C in the initial situation, and 1.1 for A, 3.4 for B, and 0 for C in the final one. In other words, according to the utilitarian criterion, A should be selected in the initial situation and B in the final one. Violating the NCR through strategic voting in the Borda voting rule is desirable because the utilitarian winner A becomes selected. The utilitarian criterion thus conflicts with the NCR, and for obvious reasons: the NCR neglects the preference intensity information by treating all identical orderings the same. Saari argues for the NCR on the grounds that profile settings where no candidate has an advantage over another should yield

a complete tie because it “reflects the principle that each voter’s preferences are treated equally” (Saari 2006, 112). This is an interpersonal comparison because it concerns the preference profile. This means that Saari’s conditions are based on interpersonal comparisons of ordinal utilities: if the voter with ordering ABC has a much greater stake in the issue than the other two voters with orderings BCA and CAB, it would seem inappropriate to declare that the social choice should be a complete tie: there would be an argument for selecting A (see, e.g., Brighthouse and Fleurbaey 2010).

In sum, if one agrees that preference intensities matter, violating NCR seems to be perfectly justified, and the fact that the Borda voting rule violates NCR is not meant to be an argument against the rule. I am somewhat diffident to present these remarks on the NCR as a criticism of Saari, because he is clearly aware of the fact that the whole debate becomes different if preference intensities are taken into account (2003, 349). Furthermore, as far as I know, he has been silent on preference intensities—except “intensities for rankings” (see, e.g., Saari 1995)—and their interpersonal comparisons.

Saari’s admonition to treat each voter’s preferences equally may be compared to Hammond’s claim for the zero-one rule. He argues that “it seems almost perverse not to make the weights equal” (Hammond 1987, 197). Indeed, that is what we have done thus far by assuming that each voter’s maximum utility is unity and the minimum zero. However, James Griffin argues that the zero-one rule is, in fact, empirically false: “It is not the case that we all reach the same peaks and valleys” (Griffin 1986, 120). In other words, if zero-one normalisation is considered an empirical assumption about people’s utilities, it is false. If Griffin’s empirical claim is correct (remember, we do not have any scientific means for determining whether it is), insofar as each person still has only one vote, the conclusion is that voting rules cannot take empirical interpersonal differences in utilities into account (see Mackie 2003, 144–145). This example could be analysed by assuming that voters’ utility numbers are drawn randomly from the (0, 1) interval. If interpersonal comparisons are introduced in this way, the conclusions of the analysis are likely to be the same as with the uniform zero-one assumption. How about assuming that type 3 voters have double the stakes compared to types 4 and 6, and that type 1 voters have half the stakes? The utilities would then be as follows:

Table 11: Modified example

1	2	3	4	5	6
6	0	3	5	0	5
A (.5)	A	B (2)	B (1)	C	C (1)
B (.2)	C	A (.2)	C (.5)	B	A (.9)
C (0)	B	C (0)	A (0)	A	B (0)

If we knew that individual utilities are like this, I presume that we would not want to apply the NCR. It would seem sensible to give type 1 voters one half the weight of the other types in the Condorcet triplet {ABC, BCA, CAB}. Then, it would be as if there were three type 1 voters with a weight of one. Similarly, it would be as if there were six type 3 voters. The profile that would follow from deleting the voter types in the Condorcet triplet would be as follows:

Table 12: Deleting Condorcet triplets

1	2	3	4	5	6
0	0	6	2	0	2
A (1)	A	B (1)	B (1)	C	C (1)
B (.4)	C	A (.1)	C (.5)	B	A (.9)
C (0)	B	C (0)	A (0)	A	B (0)

B would now clearly be the best candidate. It is clear that we could not come up with such exact knowledge of interpersonal differences in utilities. My point is just that deviating from the assumption of uniform interpersonal utility scales would be problematic for applying the NCR, but a utilitarian evaluation does not depend on such an assumption. Nevertheless, I do not intend to criticise the NCR for making such a uniformity assumption. No voting rule can take interpersonal differences in utility into account, and normative evaluations of voting rules should employ the same, or at least roughly the same scales.

The one-man-one-vote principle posits that each voter should have only one vote. It is possible to interpret this principle merely in terms of equality of opportunity (e.g., McGann 2006), i.e., in such a way that it does not imply anything about the definition of a social optimum. However, if voting outcomes are interpreted as reflecting individual utilities, they need to be normatively evaluated, and thus the need to

justify the principle in terms of claims about the utility profile arises. In other words, if we are primarily interested in voting results from the welfarist point of view, we also need to justify the one-man-one-vote principle by welfarist means.

If the principle is justified by appealing to interpersonal comparisons, the weight of each individual in determining the social optimum ought to be the same. The justification of the one-man-one-vote principle derives here from the claim that, *a priori*, each individual ought to have equal weight in determining the will of the people. It follows that each individual ought to have the same opportunity to affect the outcome of a voting process. Another way to look at the issue is to note that since the one-man-one-vote principle is violated only if we know that some voters have a legitimate claim to more than equal influence on the voting outcome, when there are no such reasons to violate the principle, we should also assume that each voter's utility is measured with the same scale. We thus impose restrictions on the utility profile by requiring that every voter has the same scale. Even though the individual utilities are usually considered empirically given facts under utilitarianism, here they are also determined by normative considerations. This viewpoint is not entirely new to utilitarianism (see Harsanyi 1975; and Nunan 1981).

Interpersonal comparisons are made in a normatively satisfactory way if the comparisons correspond to acceptable ethical judgements concerning how the utilities of different individuals should be weighed. In particular, if a model is based on interpersonal comparisons, they should be impartial in the right way: particular persons should not be favoured unless there are convincing arguments for doing so. Normatively satisfactory interpersonal comparisons require that the weight of each individual voter in the welfare function that evaluates the outcomes must be equal or at least roughly equal. It follows from the normative justification of the one-man-one-vote principle that the differences should not be too large. No voter's utility should count more than twice as much as that of any other voter. If it were otherwise, one would wonder why each voter is given one vote. Interpersonal comparisons are thus normatively satisfactory if they take some differences in utility into account, but at the same time, restricting the size of the differences can be justified by a welfarist reading of the one-man-one-vote principle.

8. CONCLUSIONS

I have argued that social choice theory does not distinguish between choices and utilities (or preferences), and that this leads them to the mistaken idea that some but not all voting rules incorporate interpersonal comparisons. Strategic voting means, however, that one cannot assume that the ballots cast correspond to voter utilities. Interpersonal comparisons are different in different voting rules only under the questionable assumption that voters are sincere. Once we distinguish between votes and utilities, we can see that different voting rules depend on different kinds of information if voters act sincerely, but once strategic behaviour is taken into account, they all depend on at least some intensity information. Strategic voting also implies that social choice theorists are merely fooling themselves if they claim to have reliable information on ordinal preferences. Strategic voting demotes the epistemic argument against interpersonal comparisons, because it shows that social choice theorists do not and cannot have reliable knowledge of preference orderings either.

If we are not allowed to make interpersonal comparisons then imposing conditions on aggregation rules is the only way in which we can evaluate voting and voting rules. However, strategic voting also has another devastating consequence for social choice theory: it renders the theory unable to provide comparative normative advice on which voting rules are the best. While aggregation rules differ with respect to the conditions they satisfy, no difference-making condition is satisfied by any voting rule. One may provide normative advice if one is willing make normative judgements about the alternatives in terms of the preference profile alone, rather than in terms of how the profile may affect some function that is supposed to describe a voting rule. Doing so in a welfarist manner, however, requires making interpersonal comparisons. This should not be such a big problem, however, because interpersonal comparisons can be made in a perfectly satisfactory way.

REFERENCES

- Arrow, Kenneth J. 1951. *Social choice and individual values*, vol. 12. New York; London: Wiley; Chapman & Hall.
- Arrow, Kenneth J. 1963. *Social choice and individual values*, vol. 12 [2nd ed.]. New Haven: Yale University Press.
- Arrow, Kenneth J. 1973. Formal theories of social welfare. In *Dictionary of the history of ideas*, ed. Philip Wiener. New York: Charles Scribner's Sons, 276-284.

- Arrow, Kenneth J. 1977. Current developments in theory of social choice. *Social Research*, 44 (4): 607-622.
- Arrow, Kenneth J. 1978. Extended sympathy and possibility of social choice. *Philosophia*, 7 (2): 223-237.
- Arrow, Kenneth J. 1983 [1967]. Values and collective decision making. In *Collected papers of Kenneth Arrow, vol. 1: Social Choice and Justice*, ed. Kenneth J. Arrow. Cambridge Massachusetts: Belknap Press of Harvard University Press, 59-77.
- Austen-Smith, David, and Jeffrey S. Banks. 1999. *Positive political theory I: collective preference*. Ann Arbor: University of Michigan Press.
- Austen-Smith, David, and Jeffrey S. Banks. 2005. *Positive political theory II: strategy and structure*. Ann Arbor: University of Michigan Press.
- Balinski, Michel L., and Rida Laraki. 2007. A theory of measuring, electing and ranking. *Proceedings of the National Academy of Sciences of the United States of America*, 104 (21): 8720-8725.
- Balinski, Michel L., and Rida Laraki. 2010. *Majority judgment*. Cambridge (MA): MIT Press.
- Blackorby, Charles, and Walter Bossert. 2006. Interpersonal comparisons of well-being. In *The Oxford handbook of political economy*, eds. Barry R. Weingast, and Donald A. Wittman. New York: Oxford University Press, 408-425.
- Blais, André, Richard Nadeau, Elisabeth Gidengil, and Neil Nevitte. 2001. Measuring strategic voting in multiparty plurality elections. *Electoral Studies*, 20 (3): 343-352.
- Blin, Jean-Marie, and Mark A. Satterthwaite. 1977. On preferences, beliefs, and manipulation within voting situations. *Econometrica*, 45 (4): 881-888.
- Bordes, Georges, and Nicolaus T. Tideman. 1991. Independence of irrelevant alternatives in the theory of voting. *Theory and Decision*, 30 (2): 163-186.
- Bossert, Walter, and John A. Weymark. 2004. Utility in social choice. In *Handbook of utility theory*, eds. Salvador Barberà, Peter J. Hammond, and Christian Seidl. Dordrecht: Kluwer Academic Publishers, 1099-1178.
- Brams, Steven J., and Peter C. Fishburn. 1983. *Approval voting*. Boston: Birkhäuser.
- Brighouse, Harry, and Marc Fleurbaey. 2010. Democracy and proportionality. *Journal of Political Philosophy*, 18 (2): 137-155.
- Buchanan, James M. 1954. Social choice, democracy, and free markets. *Journal of Political Economy*, 62 (2): 114-123.
- Christiano, Thomas. 1995. Social choice and democracy. In *The idea of democracy*, eds. David Copp, Jean Hampton, and John E. Roemer. Cambridge: Cambridge University Press, 173-195.
- Cohen, Joshua. 1986. An epistemic conception of democracy. *Ethics*, 97 (1): 26-38.
- Coleman, Jules, and John A. Ferejohn. 1986. Democracy and social choice. *Ethics*, 97 (1): 6-25.
- d'Aspremont, Claude, and Louis Gevers. 2001. Social welfare functionals and interpersonal comparability. In *Handbook of social choice and welfare*, eds. Kenneth J. Arrow, Amartya K. Sen, and Kotaro Suzumura. Amsterdam: Elsevier Science, 459-541.
- Dhillon, Amrita, and Jean-Francois Mertens. 1999. Relative utilitarianism. *Econometrica*, 67 (3): 471-498.
- Dryzek, John S., and Christian List. 2002. Social choice theory and deliberative democracy: a reconciliation. *British Journal of Political Science*, 33 (1): 1-28.

- Dummett, Michael A. 1984. *Voting procedures*. Oxford: Clarendon Press.
- Enelow, James M. 1981. Saving amendments, killer amendments, and an expected utility theory of sophisticated voting. *Journal of Politics*, 43 (4): 1062-1089.
- Felsenthal, Dan S. 1990. *Topics in social choice: Sophisticated voting, efficacy, and proportional representation*. New York: Praeger.
- Gibbard, Allan F. 1973. Manipulation of voting schemes: a general result. *Econometrica*, 41 (4): 587-601.
- Griffin, James. 1986. *Well-being: its meaning, measurement and moral importance*. Oxford: Clarendon Press.
- Hammond, Peter J. 1987. On reconciling Arrow's theory of social choice with Harsanyi's fundamental utilitarianism. In *Arrow and the foundations of the theory of economic policy*, ed. George R. Feiwel. London: Macmillan, 179-221.
- Hammond, Peter J., and Marc Fleurbaey. 2004. Interpersonally comparable utility. In *Handbook of utility theory*, eds. Salvador Barberà, Peter J. Hammond, and Christian Seidl. Dordrecht: Kluwer Academic Publishers, 1181-1285.
- Harrod, Roy. 1938. Scope and method of economics. *Economic Journal*, 48 (191): 383-412.
- Harsanyi, John C. 1975. Can the maximin principle serve as a basis for morality: critique of Rawls' theory. *American Political Science Review*, 69 (2): 594-606.
- Hausman, Daniel M. 1995. The impossibility of interpersonal utility comparisons. *Mind*, 104 (415): 473-490.
- Hildreth, Clifford. 1953. Alternative conditions for social orderings. *Econometrica*, 21 (1): 81-94.
- Hillinger, Claude. 2005. The case for utilitarian voting. *Homo Oeconomicus*, 22 (3): 295-321.
- Hirshleifer, Jack, and John G. Riley. 1992. *The analytics of uncertainty and information*. Cambridge: Cambridge University Press.
- Isbell, John. 1959. Absolute games. In *Contributions to the theory of games*, eds. A. W. Tucker, and Duncan R. Luce. Princeton: Princeton University Press, 357-396.
- Jackson, Matthew. 2001. A crash course in implementation theory. *Social Choice and Welfare*, 18 (4): 655-708.
- Kawai, Kei, and Yasotura Watanabe. 2013. Inferring strategic voting. *American Economic Review*, 103 (2): 624-662.
- Kelly, Jerry S. 1988. *Social choice theory: an introduction*. Berlin: Springer-Verlag.
- Laslier, Jean-François. 2012. And the loser is... plurality voting. In *Electoral systems*, eds. Dan S. Felsenthal, and Moshe Machover. Berlin: Springer, 327-351.
- Lehtinen, Aki. 2007a. The Borda rule is also intended for dishonest men. *Public Choice*, 133 (1-2): 73-90.
- Lehtinen, Aki. 2007b. The welfare consequences of strategic voting in two commonly used parliamentary agendas. *Theory and Decision*, 63 (1): 1-40.
- Lehtinen, Aki. 2008. The welfare consequences of strategic behaviour under approval and plurality voting. *European Journal of Political Economy*, 24 (3): 688-704.
- Lehtinen, Aki. 2009. Intentions in invisible-hand accounts. *Journal of Economic Methodology*, 16 (4): 409-416.
- Lehtinen, Aki. 2010. Behavioural heterogeneity under approval and plurality voting. In *Handbook on approval voting*, eds. Jean-François Laslier, and M. Remzi Sanver. Heidelberg: Springer, 285-310.

- Lehtinen, Aki. 2011. A welfarist critique of social choice theory. *Journal of Theoretical Politics*, 23 (3): 359-381.
- Lehtinen, Aki. 2015. Strategic voting and the degree of path-dependence. *Group Decision and Negotiation*, 24 (1): 97-114.
- List, Christian. 2003. Are interpersonal comparisons of utility indeterminate? *Erkenntnis*, 58 (2): 229-260.
- Little, Ian M. D. 1957. *A critique of welfare economics* [2nd ed.]. Oxford: Clarendon Press.
- Mackie, Gerry. 2003. *Democracy defended*. Cambridge: Cambridge University Press.
- Mbih, Boniface, and Issouf Moyouwou. 2008. Violations of independence under amendment and plurality rules with anonymous voters. *Group Decision and Negotiation*, 17 (4): 287-302.
- McGann, Anthony. 2006. *The logic of democracy: reconciling equality, deliberation, and minority protection*. Ann Arbor: University of Michigan Press.
- McKelvey, Richard D., and Peter C. Ordeshook. 1972. A general theory of the calculus of voting. In *Mathematical applications in political science*, eds. James F. Herndon, and Joseph L. Bernd. Charlottesville: University Press of Virginia, 32-78.
- Miller, Daniel. 1992. Deliberative democracy and social choice. *Political Studies*, 40 (s1): 54-67.
- Mongin, Philippe, and Claude d'Aspremont. 1998. Utility theory and ethics. In *Handbook of utility theory*, eds. Salvador Barberà, Peter J. Hammond, and Christian Seidl. Dordrecht: Kluwer Academic Publishers, 371-481.
- Myerson, Roger B. 1979. Incentive compatibility and the bargaining problem. *Econometrica*, 47 (1): 61-73.
- Myerson, Roger B. 1985. Bayesian equilibrium and incentive-compatibility: an introduction. In *Social goals and social organization: essays in memory of Elisha Pazner*, eds. Leonid Hurwicz, David Schmeidler, and Hugo Sonnenschein. Cambridge: Cambridge University Press, 229-259.
- Ng, Yew-Kwang. 2000. *Efficiency, equality and public policy: with a case for higher public spending*. Basingstoke: Macmillan.
- Nunan, Richard. 1981. Harsanyi vs. Sen: does social welfare weigh subjective preferences? *The Journal of Philosophy*, 78 (10): 586-600.
- Nurmi, Hannu. 2012. On the relevance of theoretical results to voting system choice. In *Electoral systems*, eds. Dan S. Felsenthal, and Moshe Machover. Heidelberg: Springer, 275-303.
- Ordeshook, Peter C. 1986. *Game theory and political theory: an introduction*. Cambridge: Cambridge University Press.
- Reynolds, James F., and David C. Paris. 1979. The concept of 'choice' and Arrow's theorem. *Ethics*, 89 (4): 354-371.
- Riker, William H. 1982. *Liberalism against populism: a confrontation between the theory of democracy and the theory of social choice*. San Francisco: W.H. Freeman.
- Risse, Mathias E. 2001. Arrow's theorem, indeterminacy, and multiplicity reconsidered. *Ethics*, 111 (4): 706-734.
- Risse, Mathias E. 2004. Arguing for majority rule. *Journal of Political Philosophy*, 12 (1): 41-64.
- Risse, Mathias E. 2005. Why the count de Borda cannot beat the Marquis de Condorcet. *Social Choice and Welfare*, 25 (1): 95-113.

- Risse, Mathias E. 2009a. On the philosophy of group decision methods I: the nonobviousness of majority rule. *Philosophy Compass*, 4 (5): 793-802.
- Risse, Mathias E. 2009b. On the philosophy of group decision methods II: alternatives to majority rule. *Philosophy Compass*, 4 (5): 803-812.
- Robbins, Lionel. 1938. Interpersonal comparisons of utility: a comment. *Economic Journal*, 48 (192): 635-641.
- Robbins, Lionel. 1952. *An essay on the nature and significance of economic science* [2nd ed.]. London: Macmillan.
- Robbins, Lionel. 1981. Economics and political economy. *American Economic Review*, 71 (2, Papers and Proceedings of the Ninety-Third Annual Meeting of the American Economic Association): 1-10.
- Saari, Donald G. 1995. Inner consistency or not inner consistency: a reformulation is the answer. In *Social choice, welfare, and ethics*, eds. William A. Barnett, Hervé Moulin, Maurice Salles, and Norman J. Schofield. Cambridge: Cambridge University Press, 187-212.
- Saari, Donald G. 1998. Connecting and resolving Sen's and Arrow's theorems. *Social Choice and Welfare*, 15 (2): 239-261.
- Saari, Donald G. 2003. Capturing the 'will of the people'. *Ethics*, 113 (2): 333-349.
- Saari, Donald G. 2006. Which is better: the Condorcet or Borda winner? *Social Choice and Welfare*, 26 (1): 107-129.
- Saari, Donald G. 2008. *Disposing dictators, demystifying voting paradoxes: social choice analysis*. Cambridge: Cambridge University Press.
- Saari, Donald G. 2010. From Black's advice and Arrow's theorem to the Gibbard-Satterthwaite result. In *Collective decision making*, eds. Adrian Van Deemen, and Agnieszka Rusinowska. Berlin: Springer, 1-16.
- Saari, Donald G., and Vincent R. Merlin. 2000. A geometric examination of Kemeny's rule. *Social Choice and Welfare*, 17 (3): 403-438.
- Satterthwaite, Mark A. 1975. Strategy-proofness and Arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10 (2): 187-217.
- Sen, Amartya K. 1970. *Collective choice and social welfare*, vol. 5. San Francisco: Holden-Day.
- Sen, Amartya K. 1977. On weights and measures: informational constraints in social welfare analysis. *Econometrica*, 45 (7): 1539-1572.
- Sen, Amartya K. 1995. Rationality and social choice. *American Economic Review*, 85 (1): 1-24.
- Sen, Amartya K. 1999. The possibility of social choice. *American Economic Review*, 89 (3): 349-378.
- Sheng, C. L. 1987. A note on interpersonal comparisons of utility. *Theory and Decision*, 22 (1): 1-12.
- Smith, Warren. 2000. *Range voting* [article published online]. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.32.9150> (accessed Nov. 2015).
- Tangian, Andranik S. 2000. Unlikelihood of Condorcet's paradox in a large society. *Social Choice and Welfare*, 17 (2): 337-365.
- Tideman, Nicolaus T. 2006. *Collective decisions and voting*. Aldershot: Ashgate.
- Tideman, Nicolaus T., and Florenz Plassmann. 2014. Which voting rule is most likely to choose the "best" candidate? *Public Choice*, 158 (3): 331-357.

Van Hees, Martin, and Keith M. Dowding. 2008. In praise of manipulation. *British Journal of Political Science*, 38 (1): 1-15.

Wimsatt, William C. 1997. Aggregativity: reductive heuristics for finding emergence. *Philosophy of Science*, 64 (4): 372-384.

Aki Lehtinen is a professor (temporary) of Social and Moral Philosophy at the University of Helsinki, and a university researcher at the Finnish Centre of Excellence in the Philosophy of the Social Sciences (TINT). He has written on modelling and robustness, game theory, and other miscellaneous topics in the philosophy of economics. He has also constructed computer simulation models of strategic voting and published the results in economics and political science journals.

Contact e-mail: <aki.lehtinen@helsinki.fi>

Website: <<http://www.mv.helsinki.fi/home/alehtine/akilehtinen.htm>>