



LUND UNIVERSITY

Democracy - Two Models

Rabinowicz, Wlodek

Published in:

Neither/Nor - Philosophical Essays dedicated to Erik Carlson on the Occasion of His 50th Birthday

2011

[Link to publication](#)

Citation for published version (APA):

Rabinowicz, W. (2011). Democracy - Two Models. In R. Sliwinski (Ed.), *Neither/Nor - Philosophical Essays dedicated to Erik Carlson on the Occasion of His 50th Birthday* (pp. 219-241). Uppsala Philosophical Studies.

Total number of authors:

1

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Democracy – two models*

Wlodek Rabinowicz

Department of Philosophy, Lund University

This is not a paper, but rather an idea of one. What I am going to present is a skeleton of an argument, which needs to be fleshed out in various ways before its merits can be judged. Some of the questions posed below are merely stated, but the answers are not yet available.¹

The point of departure in my story is the contrast between two models of democratic voting process: *popular democracy*, as exemplified by popular elections and referenda, and what might be called *committee democracy*, i.e., voting in smaller bodies of experts or specially appointed laymen. What is the difference between these two models? On one interpretation, voting in popular democracy is a procedure whose function is to aggregate the individuals' *preferences* to something like a collective preference, while in committee democracy what is being aggregated are committee members' opinions, or *judgments*, and the outcome is the collective judgment of the committee as a whole. The relevant judgments on the agenda often address a normative or an evaluative question: What is to be done? Or, what is the best alternative? Or, how are the alternatives to be ranked from the best to the worst? But, in some cases, the question before the committee might instead be factual in nature: Will the bridge that is being planned withstand heavy traffic? What will be the noise level in the vicinity of the railway tracks if the number of tracks is doubled? Etc..²

* This is a small birthday offering for my good friend Erik Carlson. The issues I here take up relate to some of the problems he and I discussed during our stay at the Swedish Collegium of Advanced Study in Uppsala, in 2008.

¹ Some of the answers might be soon be forthcoming. Stephan Hartmann and I collaborate on a project in which the issues I take up in the penultimate section of this paper will hopefully be examined in more detail. I have discussed with Stephan several of the questions considered below. The ideas of this paper were presented at a meeting of the Tampere Club, at the Swedish Collegium of Advanced Study in Uppsala, at the department of philosophy in Lund, and at a moral philosophy conference in Copenhagen. I am indebted to the participants in these events for useful comments and suggestions.

² On another interpretation, even in popular democracy voters are expressing their judgments rather than preferences. But, on this interpretation, while in a committee all members are supposed to answer the same

Preferring one alternative to another is *not* the same, I assume, as judging it to be a better alternative. Judgments of betterness, and in general value judgments, often accompany preferences and the latter often are based on the former. But it is perfectly conceivable that one could prefer *a* to *b*, even though one lacks a clear view about their relative value. Indeed, it is even conceivable to consider *b* to be better than *a* and still prefer *a* to *b* (perhaps because one thinks that *a* is better *for oneself*, even though one considers *b* to be better overall; or perhaps because one is simply irrational). Consequently, aggregation of preferences is not reducible to aggregation of judgments.

Needless to say, contrasting aggregation of preferences with aggregation of judgments is an extremely simplified way of describing the difference between popular democracy and committee democracy. Real life democratic processes are much more complicated than this. Thus, for example, in a committee vote, some of the members might give expression to their personal preferences rather than to their impartial judgments on the matter at hand. And in a popular election or a referendum some voters might well think of the process in terms of judgment aggregation: their vote is an expression of opinion on the issue that faces the society. Also, the multi-tiered structure of representative democracy additionally complicates the nature of the aggregation procedure. And the focus on aggregation means that one ignores other essential elements of the democratic process: the stage of deliberation and negotiations, setting-up the agenda on which the vote is to be made, etc.

While bearing this in mind, still my focus here will be on the simple contrast between aggregation of preference rankings and aggregation of judgments. What I want to look at, in particular, is the case in which the two aggregation scenarios exhibit a far-reaching structural similarity: more precisely, the case in which, in the judgment aggregation scenario, the individual judgments that are to be aggregated are *value rankings*. This means that, formally, the individual judgments in this case have the same structure as preference rankings over a given set of alternatives, but while in a preference ranking the alternatives are ordered in accordance with one's preferences, the order in a value ranking expresses one's comparative evaluation of the alternatives: say, this alternative is best, those two alternatives are second-best, that alternative is third-best, etc. I will suggest that, despite of their formal similarity as

question (ideally, at least), the questions posed to the voters in a popular election and a referendum vary from one voter to another. Each voter is meant to answer something like the question: Is this proposal good *for you*? And the outcome of aggregation is then a judgment concerning whether the proposal is good for the collective as a whole. In what follows, I shall mostly pay no attention to this 'welfarist' interpretation and instead assume the preferentialist account.

rankings, this difference in the nature of individual inputs in two aggregation scenarios has important implications for the task of aggregation: the kind of procedure that looks fine for the aggregation of judgments turns out to be inappropriate for the aggregation of preferences. The kind of procedure I have in mind consists in similarity maximization, or – more precisely – in minimization of the average distance from individual inputs. When applied to judgment aggregation, this procedure can also be approached from the epistemic standpoint: the questions will be posed concerning its advantages as a truth-tracker. In this context, what matters is not only the probability of the outcome of the procedure being true, but also the expected verisimilitude of the outcome: its expected distance from truth.

Impossibility theorems

In recent years, much work has been done on judgment aggregation. It is an area that has generated an intensive activity over the last decade. For a survey, see List and Puppe (2009). Here it is perhaps enough if I mention some of the people who have participated in this very successful research program: Christian List, Philip Pettit, Franz Dietrich, Luc Bovens and myself, Elad Dokov and Ron Holzman, Peter Gärdenfors, Marc Pauly, Martin van Hees, Philippe Mongin, Klaus Nehring, Clemens Puppe and Gabriella Pigozzi. This list is by no means complete, but the interested reader can find an extensive bibliography on Christian List's LSE website <http://personal.lse.ac.uk/list/doctrinalparadox.htm>

A typical set-up for judgment aggregation involves a finite set of individuals and an agenda - a finite set of propositions that may or may not be logically interconnected. Individuals are supposed to come up with their judgments, i.e. to specify which propositions on the agenda they accept and which they reject. It is assumed that each such individual input is logically consistent. The goal is to aggregate these individual inputs into a collective output – a set of collective judgments. More precisely, the goal is to specify which propositions of the agenda are accepted by the collective and which are rejected.

Many of the contributions to this discussion concern the existence problem: Does there exist a general procedure for judgment aggregation that satisfies a number of reasonable requirements? Here are some examples of such requirements:

consistency: the collective outcome should be logically consistent;

universal domain: the procedure should deliver a definite collective outcome for every possible configuration of individual inputs;

non-dictatorship: there should be no individual whose vote is decisive for the collective outcome given every configuration of individual inputs, i.e., independently of how the other individuals vote;

anonymity: the outcome should be invariant under permutations on individuals, i.e., all individuals should be given equal influence (this is of course a much stronger requirement than non-dictatorship);

unanimity: if all individuals accept (reject) a certain judgment, then that judgment should be accepted (rejected) in the collective outcome;

neutrality: all propositions on the agenda should be treated equally, i.e., if we permute the propositions, the collective outcome should be permuted in the same way;

independence: the collective judgment on each proposition should only depend on the individual judgments on that proposition.

Now, it has been proved by several researchers that different lists of such plausible requirements give rise to *impossibility theorems* to the effect that there is no aggregation procedure that satisfies all the requirements on the list in question.

Clearly, there is an obvious analogy here with Arrow's famous impossibility theorem for preference aggregation. The requirements that have been shown to spell trouble for judgment aggregation exhibit striking similarities to the postulates that Arrow and his followers imposed on aggregation of preferences. There we also have such postulates as universal domain, non-dictatorship, anonymity, unanimity, neutrality and independence (with the latter condition stating that the collective preference regarding any two alternatives only depends on the individual preferences with respect to the two alternatives in question). And we also have an analogue of consistency: the collective preferences should be transitive. It is normally also required that, like individual preferential inputs, the collective preferences should be complete, which together with transitivity implies that they form a well-behaved ranking.³

In fact, there are close analogies in the very points of departure for these two 'impossibility programs': In the case of preference aggregation, the point of departure was Condorcet's famous paradox for majority voting. In that paradox, which involves three

³ The corresponding requirement of completeness for individual judgments would be that every proposition on the agenda is either accepted or rejected by the individual in question; i.e., that there is no room for withholding judgment. Such a requirement is quite problematic. For more on this issue, see below.

alternatives, a , b , and c , and at least three voters, there is a majority, M_1 , that prefers a over b , and another majority, M_2 , that prefers b over c ; but there is no majority that prefers a over c . In fact, there is a majority with the opposite preference - for c over a . The reason is that, as it happens, a is preferred to c only by voters who belong to both M_1 and M_2 , but the overlap between majorities M_1 and M_2 is too small to form a majority itself.⁴

The point of departure for the impossibility program with respect to judgment aggregation was the so-called ‘doctrinal paradox’ in legal theory (cf. Kornhauser and Sager 1986, 1993). Nowadays, following Philip Pettit, one often refers to this paradox as ‘the discursive dilemma’ (cf. Pettit 2001). In a dilemma of this kind, there is a majority M_1 for a proposition p , another majority, M_2 , for a proposition q , but there is no majority for a proposition r , which is the conjunction of p and q . (In some versions, r instead is a proposition that all voters consider to be equivalent to that conjunction.) In fact, there is a majority against r . Just as in Condorcet’s paradox, the source of the dilemma is that the overlap between the majorities M_1 and M_2 is too small to form a majority itself.⁵

Finessing impossibility results

What I want to examine is a way to finesse these impossibility theorems. Instead of discussing various requirements on the aggregation procedure one by one and trying to question some of them, we could try to re-think the very nature of the procedure and its ‘point’, so to speak. An attractive idea is to look at aggregation as an *optimization* task: from this perspective, aggregation is a goal-driven activity and the right procedure for aggregation should promote that goal as much as possible. So, what could be the proper goal for aggregation? A plausible suggestion is to look at aggregation as a process in which we endeavour give individuals as much influence on the collective output as possible, in the following sense: the output should reflect the individual inputs to a maximal extent. To put it in a different way, this means that we endeavour to reach an outcome that is *as similar as possible* to the individual inputs. On

⁴ If we add a fourth alternative, d , we could even have a situation in which there are majorities for a over b , for b over c and for c over d , while at the same time the overlap between these majorities is nil and the voters unanimously prefer d over a .

⁵ Two comments: (i) The analogy with the Condorcet’s paradox would be even closer if r weren’t identical or equivalent with the conjunction of p and q , but instead would be some proposition entailed by this conjunction. What we then would have to add in order to get a dilemma is the assumption that r is accepted only by the voters who accept both a and b . (Analogously, in the Condorcet’s paradox, only the voters who prefer a over b and for b over c are assumed to prefer a over c .) (ii) If we add a third conjunct to p and q , we could even have a situation in which there are majorities for each conjunct, while at the same time the overlap between these majorities is nil and the voters unanimously reject the conjunction of these three propositions.

one way of making this task of similarity maximization more precise, the goal could be to reach an outcome that maximizes the *average* similarity to the individual inputs that we have started with. Or, to put it in other words, to reach an outcome that *minimizes the average distance* to inputs.⁶

This distance minimization approach to aggregation is not new, of course. In the case of the aggregation of preference rankings it can be traced as far back as to Kemeny (1959), where it is presented as a way of disarming Arrow's impossibility result. I shall say more about Kemeny's proposal below. In the case of judgment aggregation, distance minimization was suggested by Pigozzi (2006) and it has been further studied by Miller & Osherson (2009) and Duddy & Piggins (2011).

Obviously, distance minimization need not deliver a *unique* outcome: There may be several outcomes that all minimize average distance. In such a case, this optimization procedure will deliver a set of outcomes as its output, rather than a single outcome. To this extent, then, distance minimization violates the condition of universal domain: it does not necessarily deliver a *definite* collective outcome for every configuration of individual inputs. However, this violation is rather innocuous. The optimization procedure still delivers a definite *set* of admissible collective outcomes for every such configuration.^{7 8}

The distance minimization procedure can be expected to violate some other standard requirements as well. In particular, it is to be expected that the *independence* condition is not going to hold for plausible distance measures. But I suppose that we can look upon such violations with equanimity. After all, the requirements on the list are not untouchable. If we

⁶ An alternative could be to adopt some form of a 'prioritarian' approach, giving a higher negative weight to the larger distances between an outcome and individual inputs. This would involve subjecting the outcome's distances to inputs to some convex transformation (such as raising them to the power of k , for some $k \geq 2$) and then minimizing the average of these transforms. Another and more extreme approach on these lines is 'leximin', on which we first try to make the outcome's longest distance to inputs as short as possible and then minimize the number of inputs that lie at this maximal distance, then do the same with the distance that is second in length (minimize its length and then minimize the number of inputs lying at that distance), and so on. In what follows, however, I shall mostly ignore these possibilities.

⁷ This point was already made by Kemeny (1959)

⁸ For the case of *judgment* aggregation, however, it might be argued that the collective should withhold judgment in the case of a tie. I.e., the collective should be neutral with respect to the outcomes that minimize the distance to the inputs. This might mean either that the collective judgment should be the intersection of the outcomes that minimize distance, or – what might be more reasonable – that the collective standpoint should be thought of as an indeterminate epistemic state that is representable as a set of judgmental outcomes (the set of all those outcomes that minimize the distance to the inputs). On this approach, though, none of the outcomes in this set is itself admissible: The only admissible outcome is the indeterminate state that is represented by the set. Note that the set representation is more informative than the intersection approach, since different sets of outcomes can have the same intersection.

This idea of the collective withholding judgment when faced with a tie will only be marginally touched upon in what follows, but it certainly deserves further study.

can provide an explanation why they shouldn't hold if the objective is to reach an outcome that is as similar as possible to the individual inputs, then this explanation as such might justify our rejection of these requirements, *if* we think that this objective is the one that we should pursue. But I shall say more on this issue below.

Determination of similarities or distances is much facilitated if inputs and outcomes are objects of the same category. Thus, if inputs are consistent sets of judgments with respect to a certain agenda, an outcome should also be a consistent set of judgments, with respect to the same agenda. Or, if the inputs are rankings of a certain set of alternatives, an outcome should preferably be a ranking as well, over the same set of alternatives. If it is a matter of *preferential* rankings on the input side, the same should apply to the output side: an outcome should specify the collective preference. If, on the other hand, the inputs are *value rankings* of the alternatives, the same should apply to outcomes: an outcome should then be a collective value ranking, of the same alternative set.

Distance minimization

Since aggregation of rankings by way of distance minimization will be my main topic here, I should say more about how to measure distance between rankings. Needless to say, one might propose several different metrics in this context, but the one that probably is most widely known is the measure proposed in Kemeny (1959) and Kemeny & Snell (1962). In what follows, I shall refer to it as the *KS-measure* or the *KS-distance*. The idea of the KS-distance is simple: You count the number of comparisons between the alternatives with respect to which the two rankings differ. To put it more precisely, a ranking, x , can be represented as a set of ordered pairs of alternatives, with a pair (a, b) belonging to x if and only if x ranks a at least as highly as b . The KS-distance between two rankings, x and y , is then the cardinality of the symmetrical difference between x and y , i.e., the number of ordered pairs that belong to either x or y but not to both these rankings. The so-called Kemeny rule then enjoins us to opt for an outcome that minimizes the sum of the KS-distances to the individual inputs, or – what amounts to the same – that minimizes the average KS-distance to individual outputs.⁹

Kemeny (1959) was fully aware that his rule violates some of Arrow's requirements on the aggregation procedure. What is violated is not only the requirement that the procedure

⁹ See Kemeny (1959). For a study of the Kemeny rule and its properties, see Saari & Merlin (2000). In Kemeny (1959), there is also a 'prioritarian' rule that is being mentioned: On that rule, one minimizes the sum of the *squared* distances to inputs.

should deliver a unique outcome (on this matter, see above), but also the famous requirement of the “independence of irrelevant alternatives”. We might well have a case in which the collective outcome picked out by the Kemeny rule ranks alternatives a and b differently vis-à-vis each other partly depending on how some third alternative is ranked by the individuals. Thus, the collective ranking of two alternatives does not exclusively depend on how these alternatives are mutually ranked by the individuals in question.¹⁰ But Kemeny didn’t consider it to be a weighty objection against his proposal. If maximization of similarity violates independence, then so be it: this shows that independence is not a reasonable requirement.

If individual rankings that are the inputs in the aggregation process are interpreted as value judgments, or sets of such judgments, then the Kemeny rule may be seen as a form of judgment aggregation. But what are then the propositions on the agenda in this case? We might think of these propositions in a holistic fashion, as the competing value rankings of the alternatives, with the voter’s task being to accept one of the rankings and to reject all the others. But we might also think of them in a more piecemeal fashion: For each pair (a, b) of the alternatives that are being compared, we might assume that the agenda contains the proposition that a is at least as good as b . Then the voter’s task is to accept some of these propositions and to reject the others, in such a way that the pattern of acceptances and rejections gives rise to a well-formed ranking. The voter’s judgments are then of the form “ a is/is not at least as good as b ”. Given this second interpretation, we can connect Kemeny’s rule as applied to value rankings to Pigozzi’s (2006) general account of judgment aggregation in terms of distance minimization. Essentially, her idea was to let an outcome be a consistent and complete set of judgments that, as compared with other such consistent and complete sets

¹⁰ The example he used to show this involved two different ways in which three individuals might rank three alternatives, a , b , and c :

Profile A

| <u>Ind 1</u> | <u>Ind 2</u> | <u>Ind 3</u> | | <u>Outcome</u> |
|--------------|--------------|--------------|---------------|----------------|
| a | a | b | Kemeny rule → | a |
| c | c | a | | b, c |
| b | b | c | | |

Profile B

| <u>Ind 1</u> | <u>Ind 2</u> | <u>Ind 3</u> | | <u>Outcome</u> |
|--------------|--------------|--------------|---------------|----------------|
| a | a | b | Kemeny rule → | a, b |
| b | b | c | | c |
| c | c | a | | |

In the outcome for profile A, a is ranked above b , while in the outcome for profile B a and b are tied. But the two profiles exhibit exactly the same pattern as far the mutual ranking of a and b is concerned: individuals 1 and 2 rank a above b , while individual 3 ranks b above a .

of judgments, minimizes the average distance to the sets of judgments that constitute individual inputs. The distance measure she used was the so-called *Hamming distance*: For sets X and Y , the Hamming distance between X and Y is the cardinality of their symmetric difference, i.e., the number of items that belong either to X or to Y , but not to both.¹¹ Now, if we apply this to value rankings and interpret value rankings as sets of judgments of the form “ a is/is not at least as good as b ”, then it is easy to see that the Hamming distance between rankings is just the KS-distance multiplied by two. (To every ordered pair (a, b) that belongs to one of the rankings x and y , but not to both, correspond two judgments in the symmetrical difference between the judgment sets associated with x and y : “ a is at least as good as b ” and “ a is not at least as good as b ”.) Consequently, minimization of the average Hamming distance between rankings is equivalent to the minimization of the average KS-distance.^{12 13}

¹¹ Here’s an example of how this measure works in the discursive dilemma. Consider the case with three individuals and the agenda consisting of four propositions $\{p, q, r, p \wedge q \wedge r\}$. Suppose that the individuals’ judgment sets are, respectively, $\{p, q, \neg r, \neg(p \wedge q \wedge r)\}$, $\{p, \neg q, r, \neg(p \wedge q \wedge r)\}$, $\{\neg p, q, r, \neg(p \wedge q \wedge r)\}$. In this case, the so-called “premise-based procedure” would yield $\{p, q, r, p \wedge q \wedge r\}$ as the outcome, despite the fact that the individuals unanimously reject the conjunctive conclusion $p \wedge q \wedge r$. That outcome is nevertheless selected in view of the fact that each of the ‘premises’ p, q, r has a majority support. However, $\{p, q, r, p \wedge q \wedge r\}$ does not minimize the average Hamming distance to the individual inputs: Its Hamming distance to each individual judgment set equals 4. The outcomes that do minimize the average Hamming distance to individual inputs are those inputs themselves: each of them has the average Hamming distance of $8/3$ to the set of individual inputs.

¹² Using the Hamming distance as the measure of distance between sets of judgements is problematic. The obvious objection is that this kind of measure abstracts from the content of the judgments that are being compared. Thus, for atomic propositions, the Hamming distance is the same between, say, a judgment that the value of a certain parameter is 1 and the judgment that this value is 0 as between the judgments that this value is 1 and $1/2$, respectively. Another kind of criticism has been levelled by Duddy & Piggins (2011). These authors argue that the Hamming metric will sometimes involve double-counting if the propositions on the agenda are allowed to be logically interconnected. Thus, to use their example, if two individuals both accept a proposition p , then they disagree on the conjunction $p \wedge q$ iff they disagree on q . The Hamming metric, which counts *both* their disagreement on $p \wedge q$ and their disagreement on q , seems guilty of double-counting in such a case. In view of the relationship between the Hamming metric and the KS-measure, this objection might have implications for the use of the latter measure as well. And in fact, Duddy and Piggins argue that it does. For their own proposal of a measure of distance between rankings, see below.

¹³ The general idea that aggregation of rankings can be reduced to aggregation of judgments sets has been used by List & Pettit (2004) and by Dietrich & List (2007) in their reduction of Arrow’s theorem to an impossibility theorem for judgment aggregation. However, I would argue that this reduction is problematic if what is being reduced in this way is not a value ranking, but a preference ranking, as these authors assume. The statement of an individual i ’s weak preference for a over b is on their proposal interpreted as the claim “from i ’s perspective, a is at least as good as b ”, while the corresponding statement concerning the preference of the collective is interpreted as the claim “from the group’s perspective, a is at least as good as b ” (cf List & Pettit 2004). But if “from X ’s perspective” means “according to X ”, then this interpretation ignores the difference between preferences and value judgments, i.e., precisely the difference that we here focus on.

Measuring distance

Let $d(x, y)$ stand for the distance between x and y . Kemeny and Snell (1978) have shown that the KS-metric is the only measure of distance between rankings that satisfies their set of axioms, one of which is *the axiom of betweenness*:

If y lies between x and z , then $d(x, y) + d(y, z) = d(x, z)$.¹⁴

This uniqueness result, however, is based on a rather strict definition of betweenness, which perhaps might be questioned:

y lies *between* x and z iff x , y and z are distinct rankings such that (i) y contains every ordered pair that belongs to both x and z , and (ii) every ordered pair in y belongs to x or to z (or to both).

Apart from the betweenness axiom, other axioms imposed by Kemeny and Snell seem rather unproblematic. Their metric is supposed to satisfy the standard conditions on a distance measure: $d(x, y)$ is a non-negative real value; each x has distance 0 only to itself, the distance between x and y is the same as that between y and x (symmetry), the sum of the distances between x and y and between y and z is at least as large as the distance between x and z (triangle inequality). In addition, Kemeny and Snell assume:

Neutrality: d is invariant under all permutations of alternatives;

Reduction: If x and y agree in their top (bottom) alternatives, then $d(x, y)$ is the same as the distance between these rankings after all the top (bottom) alternatives have been removed.

Minimum: The minimal positive distance is 1.

Is the KS-metric a satisfactory measure of distance between rankings? It is not that easy to tell. One worry about this measure is that it seems to be insufficiently favourable to compromises. Condorcet's voting paradox may be used to illustrate this point. Thus, suppose that the individual inputs are the following three rankings, x , y and z , of three alternatives, a , b and c :

¹⁴ This axiom is a further specification of the standard triangle inequality axiom for distance, according to which $d(x, y) + d(y, z) \geq d(x, z)$, for all x , y , and z . The latter inequality becomes an equality according to the axiom of betweenness when y is located 'between' x and z , with betweenness being intuitively understood as a location on the straight line connecting x with z (or on a straight line connecting x with z , if the relevant geometry allows that several such lines could exist between two points in a given space).

| | | |
|-----|-----|-----|
| x | y | z |
| a | b | c |
| b | c | a |
| c | a | b |

An attempt to arrive to the collective ranking using the simple majority rule ends up with a cycle in this case: a is ranked above b , which is ranked above c , which is ranked above a . Intuitively, the natural compromise would be to opt for the equal ranking instead

$$a, b, c$$

as the collective outcome.

But, as is easy to calculate, the average KS-distance from the equal ranking to the rankings in the set $\{x, y, z\}$ is larger than from each of the rankings in this set to the set as a whole. While the former is 3, the latter equals $8/3$.¹⁵ Thus, the compromise solution – the equal ranking – is rejected by the Kemeny rule. In fact, this rule yields as the output the original set of inputs: $\{x, y, z\}$. Each input ranking minimizes in this case the average KS-distance to the Condorcetian set of input rankings.¹⁶

¹⁵ The KS-distance of the equal ranking to each ranking in $\{x, y, z\}$ is 3. On the other hand, each ranking that belongs to this set has the KS-distance 4 to each of the other two rankings in the set, which means that its average KS-distance to $\{x, y, z\}$ is $(0 + 4 + 4)/3 = 8/3$.

¹⁶ If one thinks that the right approach for the collective is to *withhold judgment* with respect to optimal options, if there are more than one, and if withholding judgment is interpreted as accepting only what is common to all the optimal options, i.e., as accepting only their intersection (see fn 8 above), then in Condorcet's paradox the collective would end up with the empty set of ordered pairs as the output. In other cases, though, the intersection of options chosen by the Kemeny rule might not be empty but still need not be a complete ranking. A natural question to ask is what would be the status of such a partial ranking from the point of view of distance-minimization? Would it still minimize the average distance to the individual inputs? To answer this question we would first need to generalize KS-distance to partial rankings: For any two such rankings, the KS-distance between them can still, I suppose, be understood as the number of ordered pairs with respect to which the two rankings differ. Note that, on this definition, the KS-distance from the empty set to each individual input ranking in Condorcet's paradox is 3, just as the KS-distance from the equal ranking, which means that the empty set of pairs does *not* minimize the average KS-distance. In other words, by withholding judgment, we suffer a loss in terms of distance minimization. The intersection of the optimal options need not itself be optimal. (For a generalization of distance to partial rankings, see Cook, Cress, and Seiford 1986.)

On the other hand, if withholding judgment with respect to optimal options is instead interpreted as accepting the *class* of these options as the collective standpoint (cf. fn. 8 above), then – in order to answer the question whether this class solution itself is optimal in terms of distance minimization – we need to determine the distances from such a class to different individual inputs. But this means that we need to determine how to measure the distance between a class of rankings and a single ranking. The natural answer seems to be that the choice here depends on the rule we are using: If it is the Kemeny rule, i.e. the rule of minimization of the average distance to individual inputs, then we should interpret the distance between a class of rankings C and a single ranking x in the corresponding way – as the average distance from the elements of C to x . (Analogously if we use the maximin rule, then we should interpret the distance between C and x correspondingly, as the maximal distance from the elements of C to x . And so on.) Given this interpretation, it is easy to prove that *if several rankings are optimal, then the class of these rankings will also be optimal*. If several rankings have the same (minimal) average distance to individual inputs, then the average distance from the class of these rankings to individual inputs will be the same. Proof: As is easy to see, the average of the average distances from C to the

The compromise solution is rejected, because, by the KS-definition of betweenness, the equal ranking is *not* located between any two of the rankings in the set $\{x, y, z\}$. Thus, consider, for example, rankings x and y . In both of them b is ranked above c , but in the equal ranking these two alternatives are tied. Thus, clause (ii) of the KS-definition of betweenness is not satisfied by the equal ranking: the ordered pair (c, b) belongs to that ranking but it doesn't belong to either x or y . This location of the equal ranking explains why the distance from x to y is shorter than the sum of the distances from x to the equal ranking and from the equal ranking to y . Similarly, the distance from x to z is shorter than the sum of the distances from x to the equal ranking and from the equal ranking to z . Given the symmetry of the distance measure, both these inequalities taken together imply that the sum of the distances from x to the other rankings in $\{x, y, z\}$ is shorter than the sum of the distances from the equal ranking to $\{x, y, z\}$. And this in turn explains the difference in the average distances.

Thus, if we want to have a measure that favours compromises, we might be well-advised to redefine the notion of betweenness. This is, in essence, what was done by Cook & Seiford (1978).¹⁷ They define their measure as follows: We start with assigning numbers, call them the *CS-numbers*, to the alternatives in a ranking, starting with 1 for the top alternative and continuing with 2 for the next alternative, etc. (Alternatively, we could assign 1 to the bottom alternative and start from below. Also, we could start with 0 instead of 1. It doesn't matter what convention we choose.) In case of a tie, we assign the average number to the tied alternatives. Thus, to take an example, if there are two alternatives at the top, each gets the number $(1+2)/2 = 1.5$.¹⁸ To avail ourselves of some formalism, let \mathbf{A} be the set of alternatives that are being ranked. \mathbf{A} is supposed to be the same for the rankings whose distances to each

rankings in the class X of individual rankings = the average of the average distances from the elements of C to X . To put it formally, letting μ stand for the average:

$$\mu\{\mu\{d(y, x): y \in C\}: x \in X\} = \mu\{\mu\{d(y, x): x \in X\}: y \in C\}.$$

Now, since all the rankings in C are optimal, their average distances to X are the same, i.e.,

for all y in C , $\mu\{d(y, x): x \in X\} = k$, for some k ,

and so the average of these averages is equal to k .

¹⁷ Another alternative would be to replace minimization of the average distance with minimization of the average convex transform of the distance between the outcome and the inputs or with a leximin-type approach (see above). These kinds of rules put a premium on solutions that do not lie too far away from any of the inputs. As a matter of fact, Cook and Seiford themselves suggested using minimization of the average of squared distances as a method of reaching the consensus ranking. Cf. Cook & Seiford (1982).

¹⁸ This kind of numbering was suggested by Kendall (1962). If one starts with 0 and assigns numbers from below upwards, the CS-numbering is a variant of the well-known Borda count with ties taken into consideration: Each alternative below a gives a one point; each alternative tied with a gives a half a point; the sum of the points received by a is its Borda number in a given ranking.

other we consider. Then, for every alternative a in \mathbf{A} and every ranking x of the alternatives in \mathbf{A} , let a^x be the CS-number assigned to a in x . The *CS-distance* between rankings is the sum of the absolute differences between the CS-numbers of the alternatives in the rankings in question:

$$d(x, y) = \sum_{a \in \mathbf{A}} |a^x - a^y|$$

On this proposal, the equal ranking (a, b, c) does minimize the average distance to the rankings in Condorcet's paradox: In particular, the average CS-distance to $\{x, y, z\}$ from each ranking in that set is $8/3$, while the average CS-distance of the equal ranking to $\{x, y, z\}$ is 2.

Cook and Seiford (1978) prove that their measure is the only one that satisfies the set of axioms that are very similar to the KS-axioms, but with a different interpretation of the betweenness relation. On their definition of betweenness,

y lies *between* x and z iff x, y and z are distinct rankings such that, for every alternative a in \mathbf{A} , its CS-number in y is between its CS-numbers in x and z . I.e., $a^x \leq a^y \leq a^z$ or $a^x \geq a^y \geq a^z$.

On this definition of betweenness, in contrast to the KS-definition, the equal ranking of a, b and c does lie between any two input rankings in Condorcet's paradox. Thus, to illustrate, consider rankings x and y in that paradox:

| | |
|-----------------|-----------------|
| \underline{x} | \underline{y} |
| a | b |
| b | c |
| c | a |

In the equal ranking, each alternative gets 2 as its CS-number (i.e., $(1+2+3)/3$), while the CS-numbers for the alternatives in x and y are, respectively, 1 and 3 (for a), 2 and 1 (for b), and 3 and 2 (for c). Thus, for each alternative, its CS-number in the equal ranking is between its CS-numbers in x and y .

This explains why the equal ranking, which doesn't minimize the average KS-distance to the rankings in Condorcet's paradox, does minimize the average CS-distance to these rankings.

Needless to say, there have been other attempts to replace the KS-metric with competing measures of distance between rankings.¹⁹ But, for my purposes in what follows, it is not necessary to discuss this issue in any further detail.²⁰

¹⁹ One such alternative measure is proposed in Duddy and Piggins (2011). Essentially, their idea is to look at the distance between two rankings as the smallest number of steps needed to transform one ranking into the other. You need just one step to move from one ranking x to another ranking y iff you can reach one from the other by raising or lowering the position of some alternative a with respect to some set X of alternatives that are equal-ranked in x and making no changes otherwise (i.e. making no other changes in the relative positions of the alternatives). This is possible only in two cases: if a is equal-ranked in x with the alternatives in X or if a in x is immediately above or below X . If a is equal-ranked in x with the alternatives in X , then you can raise a to a position immediately above X or lower it to a position immediately below X . If a in x is immediately above or below X , then you can move a to X 's level. Thus, to give an example that Duddy and Piggins use themselves, there is just one step between these two rankings:

| | |
|-----------------|-----------------|
| \underline{x} | \underline{y} |
| a | |
| b c | a b c |

So, on their approach, the distance between these rankings is 1, while on both the KS- and the CS-approach, this distance is 2. KS, CS and DP all order distances between rankings in different ways. Thus, on the CS-approach, this distance between x and y is the same as the distance from the equal ranking to any linear ranking of a , b and c , while on the DP-approach the latter distance is longer; it requires two steps. If one now compares KS with DP and CS, then it is easy to see that on both the DP- and the CS- approach the distance between x and the linear ranking of a , b , and c (in this order) is the same as the distance between x and y , while on the KS-approach the latter distance is longer than the former.

It might be noted that, just like the CS-approach, the DP-approach favours compromises: It picks out the equal ranking as the outcome in Condorcet's paradox.

²⁰ An interesting but still undeveloped suggestion has been made by Gustaf Arrhenius: The KS-measure only looks at ordered *pairs* with respect to which two rankings differ. But why not also look at triples, quadruples, etc? (A triple (a, b, c) can be said to belong to a ranking x iff x ranks a at least as highly as b and c and ranks b at least as highly as c . Similarly for quadruples, etc.) If x and y differ with respect to the same number of pairs as x and z , but the latter differ with respect to more triples, for example, then one might think that the distance between x and y should be shorter than the one between x and z . Examples of this kind have been independently constructed by Erik Carlson and Arrhenius. Here is one such case:

| | | | |
|-----------------|-----------------|-----------------|-----------------|
| \underline{x} | \underline{y} | \underline{z} | \underline{u} |
| a | b | b | b |
| b | c | d | a |
| c | d | a | d |
| d | a | c | c |

Rankings x and y differ from each other with respect to six pairs and so do rankings x and z . But x and y differ with respect to fewer triples, since they have the triple (b, c, d) in common.

However, it is unclear how such an alternative measure, which not only looks at pairs, should look like. In particular, how does it deal with trade-offs? For example: How should one weigh a disagreement in a triple as compared with a disagreement in a pair? To illustrate, in the example above x and y differ with respect to two more pairs than x and u (six pairs versus four). But x and u differ with respect to two more triples: (b, c, d) and (b, d, c) . Is the distance between x and u shorter or longer than the distance between x and y ?

If one generalizes the KS-definition of betweenness from pairs to all n -tuples, we get the following result: y lies *between* x and z iff x , y and z are distinct rankings such that, for every n , (i) every n -tuple that belongs to both x and z belongs to y as well, and (ii) every n -tuple that belongs to y belongs to x or to z (or to both).

In this definition, clause (ii) is extremely demanding: it implies that there seldom exists a ranking that lies between two rankings. ("Seldom" doesn't mean "never", however. Thus, for example, the definition implies that the ranking of a , b , and c , with a on top, and b and c tied for the second place, lies between the equal ranking of these alternatives, and the linear ranking of a , b , and c , in that order.)

Pareto condition

Are there any important differences in the formal requirements on the aggregation procedure depending on which model of democracy that is applicable in a given case? As we have seen, many standard requirements seem to be the essentially the same for preference aggregation and for the aggregation of judgments: non-dictatorship, unanimity, universal domain, etc. But if we focus on the case in which the judgmental inputs are *value rankings*, we discover one striking formal difference between the two aggregation exercises. This difference concerns the status of the Pareto condition.

Pareto: If everyone ranks a at least as highly as b and some individuals rank a higher than b , then a is ranked higher than b in the collective outcome.

This condition is intuitively plausible for preference aggregation, if we think of collective preferences as guides to choice and in addition if we take it as important that the collective in its choices endeavors to satisfy individual preferences: If some individuals prefer a to b and everyone else is indifferent, then it does seem reasonable for the collective to prefer a to b , since in the choice between a and b , it should opt for a in order to maximize preference satisfaction. That way it will satisfy the preferences of some and it won't frustrate the preferences of anyone else.²¹ One might also argue for the Pareto condition in welfaristic terms, if the degree of preference satisfaction is identified with the degree of welfare. On this interpretation, if some individuals prefer a to b and everyone else is indifferent, then the total welfare is increased if one moves from b to a .²²

When it comes to the aggregation of value rankings, things are different, however. In this aggregation process it is important to require that all individual value judgments be taken into consideration, to equal extent. Judgments to the effect that the alternatives that being compared are equally good should be given the same consideration as the competing value judgments. Therefore, if some individuals consider a to be better than b , but the

²¹ Here, I ignore the well-known objections against the Pareto condition that have to do (i) with the impossibility of the Paretian liberal (Sen 1970), and (ii) with spurious unanimity in preferences that is due to differences in beliefs (see, for example, Mongin 2005). (Note that both these objections apply not just to the Pareto condition, but also to its weaker variant: the unanimity condition.) Here, however, I assume that the cases I consider are not of the kind in which the objections (i) and (ii) are applicable.

²² In private communication, Teddy Seidenfeld has questioned the legitimacy of the Pareto condition for preference aggregation. His argument goes like this: If the collective prefers a to b , then it should be willing to sacrifice something in order to get a rather than b . But, if only few members of the collective prefer a to b , while the overwhelming majority is indifferent, then why should the collective be willing to make any sacrifice?

The preference-utilitarian answer to this objection is clear, I think: If only few members prefer a to b , there's still a reason for the collective to sacrifice something in order to get a rather than b , but the size of the sacrifice should be correspondingly small.

overwhelming majority considers *a* and *b* to be equally good, then – it would seem – the collective value judgment should follow the majority view: *a* and *b* should be considered by the collective to be of equal value. Thus, it is to be expected that the Pareto condition will be violated by any reasonable procedure for the aggregation of value rankings.^{23 24}

Now, as it turns out, the Pareto condition *is* violated by any method of aggregation that consists in distance minimization. Intuitively speaking, if we look for an outcome that is as similar to the individual inputs as possible, and in a large majority of those inputs *a* and *b* are equally ranked, then it is only to be expected that these two alternatives will be equally ranked in the outcome that maximizes similarity to the inputs.

That minimization of average distance violates the Pareto condition is something that can be shown quite generally, for all possible distance measures. More precisely, we can prove the following:

Observation 1: Suppose the set **A** of alternatives consists of just *a* and *b*. Consider any distance measure on the rankings of **A**. If a minority of individuals ranks *a* higher than *b*, while everyone else (a majority) ranks *a* and *b* equally, then the latter ranking has a shorter average distance to the individual rankings than the former.

²³ Note, however, that if one changes the nature of the task, the Pareto-type considerations might become applicable even if the task is to arrive at value judgments on the basis of individual value rankings. Thus, suppose that the task for the collective is to pick out *a best alternative* (just one of them, if there are several), rather than to deliver a complete value ranking as an output. Then, in the choice between *a* and *b*, if all the other alternatives are inferior according to everyone, the collective should, it seems, come up with *a* as its proposal, if some members of the collective consider *a* to be better than *b* and all the other members (perhaps even the overwhelming majority) take *a* and *b* to be equally good. The reason is that the collective is *unanimous* about *a* being one of the best alternatives, but not about *b* being one of them. I am indebted to Gustaf Arrhenius for pressing this point.

²⁴ There might be other interesting formal differences between aggregation of judgments and aggregation of preferences. One difference was mentioned in passim above: For judgment aggregation it might seem reasonable or perhaps even mandatory to withhold a collective judgment in case of a tie. I.e., it is not satisfactory for the collective to opt for one particular judgmental outcome, if there are other outcomes that are equally good. The solution is to accept only what's common to all such outcomes. This restriction doesn't seem to apply in the case of preference aggregation, so far as I can see. Another, related difference concerns the issue of *individual* abstaining from judgment. An individual might well withhold judgment, if the evidence isn't conclusive. It is less clear whether withholding preference is as easily conceivable. Maybe it is possible to lack a preferential attitude with respect to two alternatives, where this absence of preference is something else than indifference (cf. Rabinowicz 2008a). But such abstaining from preference is in any case a much more problematic phenomenon than abstaining from judgment. (Note, though, that it is perfectly unproblematic to abstain from *declaring* a preference. But that's something else, of course.)

A further difference that sometimes is mentioned in this context is that judgments are *logically interrelated* in various ways. I think, however, that the same goes for preferences. If, as I would argue, preference is not so much a dyadic comparative attitude, but rather a relation between the degrees of monadic attitudes, then preferences are logically related as well. Thus, if preferring one item to another consists in favouring the former to a higher degree than the latter, then transitivity, for example, is a logical relation between preferences: preferring *a* to *b* and *b* to *c* logically implies preferring *a* to *c*.

Proof: If d is a distance measure, then, for all x and y , $d(x, y)$ is a non-negative real number that equals 0 iff $x = y$. Also, symmetry holds: $d(x, y) = d(y, x)$. Consequently, if the number of individuals who rank a above b is m , while the remaining $n - m$ individuals rank a and b equally, the average distance from the unequal ranking (a over b) to the individual rankings is $(m0 + (n - m)k)/n$, i.e., $(n - m)k/n$, where k is the distance from the unequal ranking to the equal one. By the symmetry of d , the average distance from the equal ranking to the individual rankings is $(mk + (n - m)0)/n$, i.e., mk/n . Now, since $k > 0$,

$mk/n < (n - m)k/n$ if and only if $m < n - m$,

i.e., the average distance to the individual rankings is shorter from the equal ranking than from the ranking that places a above b iff the individuals who rank a and b equally are in majority.²⁵

It should be mentioned, though, that if the number of alternatives is increased, it will no longer always be the case that the equal ranking of a and b is going to be favoured by the average distance minimization in a Pareto-type case in which the majority ranks a at the same level as b . Other alternatives might come in between a and b in the minority rankings, which complicates the picture. Thus, consider the following case that involves three alternatives, a , b , and c :

| | | |
|-----------------|-----------------|-----------------|
| \underline{x} | \underline{y} | \underline{z} |
| a | c | $a\ b$ |
| c | $a\ b$ | c |
| b | | |

Note that in this example we have a situation in which the Pareto condition is applicable: in one ranking, x , a is placed above b , while the remaining two voters rank a and b equally. If we now calculate distances in accordance with the proposal of Kemeny & Snell, it is easy to see that the average KS-distance from x to $\{x, y, z\}$ is 2, which is shorter than the corresponding distance to $\{x, y, z\}$ from any other ranking, and in particular from any ranking in which a and b are placed at the same level. Their average KS-distance to $\{x, y, z\}$ is $7/3$ in each case. Thus, in this particular case, the Pareto condition is satisfied by the KS-measure.

²⁵ Note that this result will still hold if we apply a ‘prioritarian’ rule of distance minimization, i.e. replace distance with its convex transform in the minimization of the average. Such a change would only mean that we replace 0 and k with their convex transforms. Since $k > 0$, the same applies to their convex transforms, which is sufficient for the conclusion we have been after. Clearly, the result we have proved will also hold for the ‘leximin’ approach, since the number of individual rankings to which the unequal ranking has the maximal distance, k , is smaller than the corresponding number for the equal ranking.

Still, this is just a marginal point. The main lesson is this: *The Pareto condition marks an important dividing line between aggregation of preferences and aggregation of value rankings. As a result, distance minimization as an aggregation method seems fine for judgments, but not for preferences.*

It might be mentioned that there is another condition, closely related to Pareto, with respect to which aggregation of preferences differs from aggregation of value rankings. The intuition behind the Pareto condition of preference aggregation is that the voters who are indifferent can be safely ignored: their preferences won't be frustrated anyway, whatever ranking of the alternatives the collective decides upon. The principle that expresses this intuition can be formulated as follows:

Indifference: The collective ranking of the alternatives doesn't change if voters who rank all the alternatives equally are removed from consideration (as long as some voters still remain to be considered).

As is easy to see, the Pareto condition can be derived from Indifference conjoined with Unanimity, if Independence of Irrelevant Alternatives is assumed. Here is the proof: Suppose that everyone ranks a at least as highly as b and some individuals rank a higher than b . By Independence of Irrelevant Alternatives, in order to determine the mutual standing of a and b , we can exclude from consideration all the other alternatives, i.e. we can reduce the set of alternatives to $\{a, b\}$, while keeping the individuals' rankings of a and b *vis-à-vis* each other unchanged. By Indifference, we can then also exclude from consideration all the voters who rank a and b equally. Since all the remaining voters by hypothesis rank a higher than b , Unanimity entails that the collective ranks a higher than b .

Admittedly, this proof is not worth very much, since Independence of Irrelevant Alternatives is a highly questionable principle. But at least the proof shows that Pareto is related to Indifference, in terms of the underlying intuition.

Now, while Indifference, as we have seen, is a very reasonable principle for preference aggregation, it is clearly intuitively implausible for the aggregation of value rankings. Judgments of voters who rank the alternatives equally should surely be given just as much consideration as judgments of the other voters. And, using the same case as the one that we have made use of in Observation 1, it is easy to show that Indifference, just as Pareto, is violated by minimization of average distance, for all possible distance measures.

Observation 2: Suppose the set \mathbf{A} of alternatives consists of just a and b . Consider any distance measure on the rankings of \mathbf{A} . If a minority of individuals ranks a higher than b , while everyone else (a majority) ranks a and b equally, then (i) the former ranking doesn't have the minimal average distance to the individual rankings. But (ii) if the individuals who rank a and b equally are excluded from consideration, then the ranking in which a comes above b does have the minimal average distance to the remaining individual rankings.²⁶

This means then that removing indifferent voters from consideration does change the outcome of the distance minimization procedure, in violation of Indifference.

Proof of Observation 2: (i) follows from Observation 1. As for (ii), if the individuals who rank a and b equally are excluded from consideration, then all the remaining individuals rank a higher than b . But then, by the definition of a distance measure, the ranking in which a comes above b has a shorter average distance to these remaining individual rankings (namely, zero) than any other ranking.

Aggregation of value rankings from the epistemic perspective

This section is just a sketch. It poses some questions, but leaves them unanswered.²⁷

As is well-known, judgment aggregation can be seen as an epistemic device: as a way of arriving to an opinion that has good chances of being correct. The classical result in this area is *Condorcet's jury theorem* for majority voting: If the voters are relatively competent with respect to a proposition A , i.e., if they are at least better than random devices as far as the probability of correctly judging the truth of A is concerned, if their competence is the same and their judgments concerning A are independent of each other, then the majority's assessment of A is more reliable than a single voter's (i.e., the probability of the majority's judgment being true exceeds the corresponding probability for a judgment of a single voter) and this reliability converges to 1 as the number of voters goes to infinity. The underlying intuition is very simple: If the voters can be seen as relatively reliable independent witnesses, then – clearly - the more such witnesses we consult, the better.

What one wonders is whether this epistemic approach could be used for the aggregation of value judgments. The affirmative answer presupposes (i) that such judgments can be

²⁶ Just as it was the case with Observation 1, this result will still hold if we replace distance with its convex transform in the minimization of the average. And it will also hold for the 'leximin' approach.

²⁷ Some of the answers, I hope, can be provided by a joint work with Stephan Hartmann.

independently true or false (or, more cautiously, correct or incorrect), and (ii) that epistemic competence with respect to such judgments is possible. In fact, Condorcet himself thought that both these conditions could well be satisfied and he formulated his jury theorem precisely in order to deal with the aggregation of value rankings:²⁸ He proved it as a part of his argument in favor of using the pairwise majority rule as the aggregation method for value rankings. His idea was that each voter could be seen as a relatively competent judge with respect to each pair of alternatives. The question he then posed was: What ranking would most probably correspond to the true value ranking, given the rankings of the voters?

The assumptions he made were the following:

- Only linear rankings are allowed as individual inputs.
- Each voter has the same competence as any other voter.
- Each voter has the same competence with respect to each pair of alternatives.
- The probability of a voter's judgment regarding a pair of alternatives being correct is independent of the correctness of her judgments regarding other pairs.
- The probability of a voter's judgment being correct is probabilistically independent of the correctness of other voters' judgments.

Under these conditions, he argued, the judgments of the majority would be the ones to follow for each pair of alternatives: The pairwise majority ranking has the highest probability of being correct, apart from the cases in which the majority method leads to cycles (i.e., cases in which we encounter Condorcet's voting paradox). The latter had to be treated in a more complicated way. (Condorcet's particular proposal how to deal with these complications wasn't satisfactory, however. Young (1988) suggested a needed improvement.).

Condorcet's assumption that individual inputs are linear orderings, i.e., do not contain any ties, is unmotivated and should be rejected. Also, it might be of interest to look at individual value rankings holistically, rather than piecemeal, as Condorcet has done. In other words, instead of ascribing to individuals epistemic competence with respect to each pair of alternatives, separately, we might want to ascribe to them competence regarding the ranking as a whole. From this perspective, it would be interesting to pose questions about the epistemic value of distance minimization as an aggregation procedure: How does such a procedure fare from the epistemic standpoint? How good is it as a truth-tracker? Are some distance measures preferable from this epistemic viewpoint to other distance measures?

²⁸ Cf. Condorcet (1785). The presentation that follows is based on Young (1988).

Note that, for distance-based aggregation procedures, we can try to compute not only the *probability of truth* for an outcome, i.e. the probability of an aggregation outcome being the true ranking, but also the *expected verisimilitude* of that outcome, or, equivalently, its *expected truth-distance*. By this I mean the sum of the outcome's distances to different possible rankings, with the distances in question being weighted with the probabilities for each of those rankings of being the true ranking of the alternatives. (Minimal expected truth-distance is zero, while the maximal expected truth-distance equals the maximal distance between two rankings in the set of possible rankings of a given alternative set.) A reasonable question about a procedure that minimizes average distance to individual inputs is how good it is, as compared with a single voter, not only in increasing the probability of truth, but also in decreasing the expected truth-distance.

It is not quite clear, though, how to deal with the epistemic issues in the case under consideration. In the standard Condorcetian set-up, the voters face a binary choice: to vote for or against a given proposition. But in the case we are interested in, each voter instead chooses a ranking out of a set of several possible rankings. So the choice is not binary.

List and Goodin (2001) extended Condorcet's theorem to the case of choice among several options. A rule they proved to have Condorcetian features was *plurality voting*: the option that gets the largest amount of votes wins. That option is more likely to be correct than any other option on the table. (Which doesn't mean, of course, that it is more likely than not that this option is correct.) But List and Goodin did not consider potential similarities and dissimilarities between the options. Such similarity relations play an important role when options are *structured* objects, such as rankings. For this reason, minimization of the average distance might well yield as an outcome a ranking that no voter has proposed. For example, if half of the voters rank four alternatives, *a*, *b*, *c*, and *d*, in this descending order, while the other half rank them in the opposite order, the equal ranking of the four alternatives will be the unique optimal choice according to the Kemeny rule.

For the Condorcetian approach in which we view the problem as the case of choice between rankings and the purpose is to increase *the probability of truth*, a reasonable idea would be to ascribe to each voter epistemic competence understood along the following lines: The probability of the voter picking the true ranking should be higher than some threshold. A plausible probability threshold appears to be 1/the number of possible rankings of a given alternative set. (This value is the probability that the ranking picked at random will be the correct one.) Also, for any ranking *x*, the probability that the voter chooses that ranking might

possibly be assumed to increase with the decrease of the distance between x and the true ranking.²⁹ A simpler alternative would be to assume, instead, that for any false rankings x and y the probability of picking x by the voter is the same as the probability of his picking y .

What should the analogue of Condorcet's theorem establish for the case of an aggregation procedure that need not deliver a unique result? Remember that there might be several rankings that are optimal in the sense that each minimizes the average distance to the input rankings. So, what should we expect from a good truth-tracking procedure in such a case? Possibly, only this: The probability of each of the optimal rankings being correct should be higher than the corresponding probability for any non-optimal ranking. Also, the probability of one of the optimal rankings being correct (i) should be higher than the probability of a single voter picking out the correct ranking and (ii) should converge to 1 when the number of (relatively competent and independent) voters goes to infinity.

If the epistemic objective for an aggregation procedure is minimization of the expected truth-distance rather than maximization of the probability of truth, then a voter's competence should instead be specified as the expected truth-distance of her ranking. What is a reasonable competence threshold in this case is not clear. But perhaps something like this would fit the bill: The expected truth-distance of the voter's ranking should be higher than the expected truth-distance that the equal ranking would have under the uniform probability distribution among rankings. It can be shown that, under such probability distribution, the equal ranking minimizes the expected truth-distance, as measured by KS,³⁰ and it's a fair conjecture that a similar result can be established for other plausible distance measures. In this sense, then, a voter is more reliable than an ignorant person who chooses an option that minimizes expected truth-distance under the state of total ignorance.

When is the aggregation procedure satisfactory from the perspective of the minimization of expected truth-distance? I suppose that the expected truth distance of each ranking that is optimal according to a given procedure should be shorter than the corresponding distance of the non-optimal rankings. Also, the average expected truth-distance of the optimal rankings should be higher than the expected truth-distance that characterizes the competence of a single voter. Finally, the average expected distance of the optimal rankings should converge to 0 when the number of voters goes to infinity.

²⁹ Should the increase in probability be proportional to the decreases in distance? I am not sure.

³⁰ See Rabinowicz (2008b)

All these are unproven and shaky conjectures. Some of them, I hope, can be tested in future work.

Summing up

The focus of this paper was on the contrast between aggregation of preferences and aggregation of judgments, with a particular attention to value judgments represented as rankings. Distance-based aggregation methods, which presuppose that aggregation is treated as an optimization task, seem to provide an attractive way of finessing the standard impossibility theorems. Such methods exhibit a large spectrum of alternative options: they can differ from each other depending on the distance measure they assume, but also depending on the particular use to which they put the measure in question (minimization of the average distance, of the average squared distance, leximin, etc). However, as has been argued here, the distance-based methods only appear to be appropriate for judgment aggregation and *not* for the aggregation of preferences. The reason is that they violate the Pareto condition and the condition of Indifference, which represent a watershed between preference aggregation and the seemingly analogous task of aggregating value rankings.

As methods for judgment aggregation, the distance-based approaches invite an evaluation from the epistemic perspective. How good are they in increasing the probability of truth and how do they perform in increasing the expected verisimilitude? Some tentative thoughts on this matter have been presented in the last, sketchy section of this paper.

References

- Condorcet, Marquis de, 1785, *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la probabilité des voix*, Paris, De l'imprimerie royale.
- Cook, W. D. and L. D. Seiford, 1978, "Priority Ranking and Consensus Formation", *Management Science* 24: 1721-1732.
- Cook, W.D. and L.M. Seiford (1982): "On the Borda-Kendall Consensus Method for Priority Ranking Problems," *Management Science* 28, 621-637.
- Cook, W.D., M. Kress, and L.M. Seiford (1986): "Information and Preference in Partial Orders: A Bimatrix Representation", *Psychometrika* 51: 197-207
- Dietrich, F. and C. List, 2007, "Arrow's Theorem in Judgment Aggregation", *Social Choice and Welfare* 29: 19-33.
- Duddy, C. and A. Piggins, 2011, "A Measure of Distance between Judgment Sets", working paper.
- Kemeny, J. G., 1959, "Mathematics Without Numbers", *Daedalus* 88: 577-591.

- Kemeny, J. G. and J. L. Snell, 1962, "Preference rankings", in J. G. Kemeny and J. L. Snell, *Mathematical Models in the Social Sciences*, New York: Blaisdell, 9-23.
- Kendall, M., 1962, *Rank Correlation Methods*, 3rd. ed., Hafner, New York.
- Kornhauser, L. A., and L. G. Sager, 1986. "Unpacking the Court", *Yale Law Journal* 96: 82-117.
- Kornhauser, L. A., and L. G. Sager, 1993. "The One and the Many: Adjudication in Collegial Courts", *California Law Review* 81: 1-51.
- List, C. and R. Goodin, 2001, "Epistemic Democracy: Generalizing the Condorcet Jury Theorem", *Journal of Political Philosophy* 9: 277-306.
- List, C. and P. Pettit, 2004, "Aggregating Sets of Judgments: Two Impossibility Results Compared", *Synthese* 140: 207-235.
- List, C. and C. Puppe, 2009, "Judgment aggregation: a survey", in P. Anand, P. Pattanaik and C. Puppe (eds.), *Oxford Handbook of Rational and Social Choice*, Oxford: Oxford University Press, 457-482.
- Miller, M. K. and D. Osherson, 2009, "Methods for Distance-Based Judgment Aggregation", *Social Choice and Welfare* 32: 575-601.
- Mongin, P., 2005, "Spurious Unanimity and the Pareto principle", The Centre for Philosophy of Natural and Social Science (CPNSS) working paper, vol. 1, no. 5. London School of Economics, London.
- Pettit, P., 2001. "Deliberative democracy and the discursive dilemma", *Philosophical Issues* 11: 268-299.
- Pigozzi, G., 2006, "Belief merging and the discursive dilemma: an argument-based account to paradoxes of judgment aggregation", *Synthese* 152: 285-298.
- Rabinowicz, W., 2008a, "Value Relations", *Theoria* 74: 18-49.
- Rabinowicz, W., 2008b, "Presumption of Equality", in Martin L. Jönsson, *Proceedings of the 2008 Lund-Rutgers Conference*, Lund Philosophy Reports: Lund: 109-155.
- Saari, D. G. and V. R. Merlin, 2000, "A Geometric Examination of Kemeny's Rule", *Social Choice and Welfare* 17: 403-438.
- Sen, A. K., 1970, "The Impossibility of a Paretian Liberal", *Journal of Political Economy* 78: 152-157.
- Young, H. P., 1988, "Condorcet's Theory of Voting", *The American Political Science Review* 82: 1231-1244.