Wlodek Rabinowicz

Aggregation of value judgments differs from aggregation of preferences

Book section

Original citation: Originally published in: Kuźniar, Adrian and Odrowąż-Sypniewska, Joanna, (eds.) Uncovering Facts and Values: Studies in Contemporary Epistemology and Political Philosophy. Poznań Studies in the Philosophy of the Sciences and the Humanities(107). Leiden, The Netherlands : Brill | Rodopi, 2016, pp. 9-40.

Published in *A. Kuźniar and J. Odrowąż-Sypniewska (eds.), Uncovering Facts and Values. Studies in Contemporary Epistemology and Political Philosophy*, Brill - Rodopi, Leiden / Boston 2016, 9-40.

AGGREGATION OF VALUE JUDGMENTS DIFFERS FROM AGGREGATION OF PREFERENCES

Wlodek Rabinowicz

Lund University and London School of Economics

Abstract: This paper focuses on the contrast between aggregation of individual preference rankings to a collective preference ranking and aggregation of individual value judgments to a collective value judgment. The targeted case is one in which the two aggregation scenarios exhibit a far-reaching structural similarity: more precisely, the case in which the individual judgments that are to be aggregated are *value rankings*. This means that, formally, the individual judgments are isomorphic to individual preference rankings over a given set of alternatives. The paper suggests that, despite of their formal similarity as rankings, the difference in the nature of individual inputs in two aggregation scenarios has important implications: the kind of procedure that looks fine for aggregation of judgments turns out to be inappropriate for aggregation of preferences. The relevant procedure consists in similarity maximization, or – more precisely – in minimization of average distance from individual inputs. It is shown that, whatever measure is chosen, distance-based procedures violate the (strong) Pareto condition. This seems alright as value judgment aggregation goes, but would be unacceptable for preference aggregation.

When applied to judgment aggregation, distance-based procedures might also be approached from the epistemic perspective: questions might be posed concerning their advantages as truth-trackers. From that perspective, what matters is not only the probability of the outcome being true, but also its expected verisimilitude: its expected distance from truth.

Key words: preference aggregation, judgment aggregation, preference, value judgment, distance-based methods, Pareto, Condorcet's jury theorem, distance measures, verisimilitude, truth-tracking, Kemeny.

The point of departure in my story is the contrast between two models of democratic voting: *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, viewed as ideal types? On one interpretation, voting in popular democracy aggregates the individuals' *preferences* to something like a collective preference. In committee democracy, on the other hand, what is being aggregated are the committee members' *judgments* and the outcome is the collective judgment of the committee as a whole. In some cases, the question before a committee might be of an empirical nature: Will the bridge that is being planned hold? Or, what will be the noise level in the vicinity of railway if the number of tracks is doubled? But, very often, the relevant question *facing* the committee is normative or evaluative: What is to be done? What is the best alternative? Or, how are the alternatives to be ranked from the best to the worst?[1]

Preferring one alternative to another is *not* the same as judging it to be better. Judgments of betterness, and in general value judgments, often accompany preferences and the latter might often be based on the former. But it is possible to prefer *a* to *b* even though one lacks a clear view about their relative value. Indeed, it is even possible to judge *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 value judgments.

Needless to say, contrasting aggregation of preferences with aggregation of judgments is a highly idealized and simplified way of describing the difference between popular democracy and committee democracy. Real life democratic processes are more complicated than this. 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: instead of expressing a preference, their vote might express a value judgment regarding the options between which the choice is to be made. The multi-tiered structure of representative democracy additionally complicates the nature of the aggregation procedure. The focus on aggregation means that one ignores such essential elements of the democratic process as deliberation and negotiations, setting-up the agenda of issues on which the vote is to be made, etc.

---

[1] On another interpretation, even in popular democracy voters are expressing their judgments rather than preferences. But, on that interpretation, while in a committee all the members answer the same question (ideally, at least), the question posed in a popular election or a referendum varies from voter to voter. Ideally at least, each voter answers something like the question: Is this proposal good *for me*? And the outcome of aggregation is then a judgment concerning whether the proposal is good for the collective of voters taken as a whole.

If one, instead, thought of voters in popular democracy as expressing their views on the same question, say, on whether a given proposal is good *for the collective as a whole*, then – on *this* idealized picture – popular democracy would be like committee democracy: In both cases the function of the democratic procedure would be to aggregate the voters' judgments on the issue at hand.

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 consider is the case in which the two aggregation scenarios exhibit a far-reaching structural similarity: more precisely, the case in which the individual judgments that are to be aggregated are *value rankings*. This means that, formally, the individual judgments are isomorphic to individual 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, from the best at the top to the worst at the bottom. I will suggest that, despite of their formal similarity as rankings, this difference in the nature of individual inputs in two aggregation scenarios has important implications: the kind of procedure that looks fine for aggregation of judgments turns out to be inappropriate for aggregation of preferences. The procedure I have in mind consists in similarity maximization, or – more precisely – in minimization of average distance from individual inputs. When applied to judgment aggregation, this procedure can also be approached from the epistemic perspective: the questions will be posed concerning its advantages as a truth-tracker. From that perspective, 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. For a survey, see List and Puppe (2009) and List (2012). 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 and complete (in the sense that each proposition on the agenda is either accepted or denied). 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: Is there any general procedure for judgment aggregation that satisfies reasonable requirements? Here are some examples of such requirements:

Universal Domain: The procedure should deliver a definite collective outcome for every possible profile of individual inputs;

Consistency: The collective outcome should be logically consistent;

Completeness: Each proposition on the agenda should be either accepted or denied in the collective outcome.[2]

Non-Dictatorship: There should be no individual whose vote is decisive for the collective outcome, independently of how the other individuals vote;

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

Unanimity: If all individuals agree on a certain judgment, then that judgment should be part of the collective outcome;

Neutrality: If in a given profile of individual judgments two propositions are treated equally by each individual (one proposition is accepted if and only if the other is accepted), the collective outcome should also treat these propositions equally;

Independence: The collective judgment regarding each proposition should only depend on the individual judgments regarding that proposition.

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.[3] 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

---

[2] If the number of voters who accept a proposition is the same as the number of voters who deny it, then requiring the collective to either deny or to accept the proposition in question seems implausible. This shows that Completeness is a problematic requirement, especially if the procedure is required to deliver a definite (single) collective outcome for each profile, but one needs something in the vicinity of Completeness if the conditions on the aggregation procedure are to have any real bite.

[3] To be more precise, whether some axioms lead to an impossibility depends on the agenda,: the more the propositions on the agenda are interconnected, the more it is likely that we will run into impossibility. Furthermore, the logically stronger the axioms are, the larger is the class of agendas for which we get impossibility. I am indebted to the anonymous referee for pressing these points.

collective preference over any two alternatives only depends on the individual preferences over the two alternatives in question). And we also have an analogue of consistency and completeness: just as individual preferences, the collective preferences should form a ranking, i.e. they should be transitive and complete.

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

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 problem as ‘the discursive dilemma’ (cf. Pettit 2001). In the discursive dilemma, there is a majority M_1 for a proposition *p*, another majority, M_2, for a proposition *q*, but there is no majority for the conjunction $p \wedge q$. In fact, there is a majority against $p \wedge q$. 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.

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 undermine some of them, we could re-think the very nature of the procedure and its point. 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 to 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 aim to reach an outcome that is as *similar* as possible to the individual inputs. On one way of making this task of similarity maximization more precise, the goal could be to reach an outcome whose *distance* to inputs is as short as possible, on the average.

An alternative to simple averaging could be to adopt some form of a 'prioritarian' approach, which would mean that we overweight in aggregation longer distances between an outcome and individual inputs. I.e., that we give such longer distances disproportionately larger negative weight, so to speak. This could be done by subjecting the outcome's distances to inputs to a convex transformation (such as, say, raising the distances to the power of k, for some $k > 1$) and then minimizing the average of the transformed distances. Another and more extreme approach on these lines is 'leximin', on which we first try to make the outcome's maximal distance to inputs as short as possible and then minimize the number of inputs that lie at this maximal distance, secondly we do the same with the distance that is second in length (i.e., we minimize its length and then minimize the number of inputs lying at that distance), and so on. In what follows, I shall mostly focus on the minimization of average distance but I do this primarily for simplicity's sake.

The distance minimization approach to aggregation isn't new, of course. In the case of preference aggregation, it can be traced back to Kemeny (1959), where it was 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. Optimization procedure will then 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 deliver a *definite* collective stance for every profile of individual inputs. However, this violation is rather innocuous. The optimization procedure still delivers a definite *set* of admissible collective stances for every such profile (cf. Kemeny 1959). The condition of universal domain will thus be satisfied by minimization of average distance if we let the aggregation procedure be a multifunction, which to every profile of individual inputs assigns a non-empty set of alternative collective outputs – a set which need not be a singleton.[4]

[4] At least in the case of *judgment* aggregation, however, it might be argued that the collective should suspend judgment in the case of a tie, rather than be permitted to simply pick one of the tied outcomes as its stance. Personally, I am in favour of doxastic permissivism, but there are philosophers who find it abhorrent. For recent

The distance minimization procedure can however be expected to violate some other standard requirements on aggregation. 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 standard requirements on are not unassailable. Supposing we can show that they shouldn't hold if the proper objective is to reach an outcome that is as similar as possible to the individual inputs, then this would justify our rejection of those requirements. 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 type. Thus, if inputs are consistent and complete sets of judgments with respect to a certain agenda, an outcome is of the same type if it is a consistent and complete set of judgments with respect to the same agenda. Analogously, if the inputs are rankings of a certain set of alternatives, an outcome is of the same type if it is a ranking of the same 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 individual *value rankings*, the same should apply to outcomes: an outcome should then be a collective value ranking.

Distance minimization

Since aggregation of rankings by way of distance minimization will be my main topic, I should say more about how one can measure distance between rankings. Needless to say, several different metrics might be used in this context; see Appendix for a presentation of a number of metrics of this kind. However, probably the 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*. To define it, note first that a ranking, x, can be represented as a set of ordered pairs of alternatives, with a pair (a, b) belonging to x if and

---

criticisms of permissivism (though not in the particular context of aggregation of individual judgments but instead with respect to the more general relation between evidence and judgment), see White (2005, 2013), Hedden (forthcoming) and Horowitz (forthcoming). For defences of permissivism (again, with respect to the relation between evidence and judgment), see Rosen (2001), Kelly (2013), Schoenfield (1014), Meacham (forthcoming) and Mayo-Wilson & Wheeler (forthcoming).

The anti-permissivist position might require either that the collective judgment should be the intersection of the outcomes that minimize average distance to inputs, or that the collective standpoint should be an indeterminate doxastic state that can be represented as the set of all those judgmental outcomes that minimize the average distance to the inputs. On the latter approach, none of the outcomes in in the doxastic state is admissible: The only admissible stance is the indeterminate state itself. 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 suspending judgment when faced with a tie will only be marginally touched upon in what follows, but it certainly deserves further study.

only if *x* ranks *a* at least as highly as *b*. Now, the KS-distance between two rankings, *x* and *y*, is simply the number of ordered pairs that belong to either *x* or *y* but not to both these rankings.

The *Kemeny rule* enjoins us to choose a ranking that minimizes the sum of its KS-distances to the individual inputs, or – what amounts to the same – that minimizes the average KS-distance to individual inputs.[5]

Kemeny (1959) was fully aware that his rule violated some of Arrow's requirements on the aggregation procedure: not only the requirement that the procedure should deliver a unique outcome (on this matter, see above), but also the requirement of "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 depending on how some third alternative is ranked by the individuals. Thus, in violation of the independence of irrelevant alternatives, the collective ranking of two alternatives does not exclusively depend on how these alternatives are mutually ranked by the individuals in question.[6] But Kemeny did not consider it to be a weighty objection against his proposal. If distance minimization 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 as sets of value judgments, then the Kemeny rule may be seen as a form

---

[5] See Kemeny (1959). For a study of the Kemeny rule and its properties, cf. Saari & Merlin (2000). Kemeny (1959) also considers a 'prioritarian' rule as an alternative: On that rule, one minimizes the sum of the *squared* distances to inputs, which means that longer distances are given a disproportionate weight as compared with shorter distances.

[6] The example he used to show this involved two different ways in which three individuals, 1, 2 and 3, might rank three alternatives, *a*, *b*, and *c*:

Profile A

1	2	3		Outcome
a	*a*	*b*		*a*
c	*c*	*a*	Kemeny rule →	*b*, *c*
b	*b*	*c*		

Profile B

1	2	3		Outcome
a	*a*	*b*		*a*, *b*
b	*b*	*c*	Kemeny rule →	*c*
c	*c*	*a*		

In the outcome for profile A, *a* is ranked above *b* if the Kemeny rule is used, 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 rankinga of *a* and *b* are concerned: individuals 1 and 2 rank *a* above *b*, while individual 3 ranks *b* above *a*.

of judgment aggregation. But what are then the propositions on the agenda in? We might think of these propositions in holistic terms, 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 piecemeal fashion: For each ordered pair (*a*, *b*) of the alternatives that are being compared, we might assume that the agenda contains the proposition that <u>*a* is at least as good as *b*</u>. 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. She considered a case in which individual inputs were consistent and complete sets of judgments with respect to a given agenda (i.e. took a stand on every proposition on that agenda). 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 of judgments, minimized the average distance to individual inputs. The metric 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. 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.[7] [8]

---

[7] Using the Hamming distance as the measure of distance between sets of judgements is problematic. The obvious objection is that this kind of metric abstracts from the content of the judgments that are being compared. Thus, 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 .9, respectively. Another kind of criticism has been levelled by Duddy & Piggins (2012), who argue that the Hamming metric will sometimes involve double-counting if the propositions on the agenda are allowed to be logically interconnected. 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 indeed, Duddy and Piggins (2012) argue that it does. For their own proposal of a measure of distance between rankings, see Appendix.

Pareto condition

Are there any important differences in the formal requirements on the aggregation procedure depending on what is being aggregated? As we have seen, many standard requirements seem to be 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 judgments to be aggregated exhibit the same formal structure as preferences, i.e. on the case when 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 every individual ranks *a* at least as highly as *b* and some individuals rank *a* higher than *b*, then *a* is ranked higher than *b* by the collective.

This condition is intuitively plausible for preference aggregation, if we think of collective preferences as primarily guides to choice and if we in addition take it to be 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 it should opt for *a* in the choice between *a* and *b* in order to maximize individual preference satisfaction. By opting for *a* rather than *b*, it will satisfy the preferences of some and frustrate the preferences of no one.[9]

To be sure, one might question the validity of the Pareto condition for preference aggregation. And indeed it has been questioned by Teddy Seidenfeld (in private communication). His objection 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

---

[8] 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 (unless it is meant as a purely formal exercise) if what is being aggregated are not value rankings, but preference rankings, 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 *ibid*.). 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.

[9] Here, I ignore the well-known objections against the Pareto condition that have to do with (i) the impossibility of the Paretian liberal (Sen 1970), and with (ii) spurious unanimity in preferences which is due to differences in individual 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.) I assume that the cases we consider are not of the kind in which objections (i) and (ii) are applicable.

collective be willing to make any sacrifice? Note that it may well be the case that the majority of the members would prefer that *no* sacrifice be made. That is, while they are indifferent between *a* and *b*, they would prefer *b* to *a*-with-sacrifice, however small.[10]

It is a striking and thought-provoking objection. However, what it shows, I think, is that it is conceivable and not inconsistent to prefer one option to another without being willing to make any sacrifice to get the former rather than the latter.[11] This seems to be the predicament of the collective if most of its members are indifferent between *a* and *b* and only few prefer *a* to *b*. That preferences do not always come together with willingness to make sacrifices may be surprising, but - on reflection – we should accept this implication. Especially so, when it comes to collective preferences, which are a sort of a theoretical artefact. As such, they need not have quite the same features as standard individual preferences.[12] It seems, then, that we do have strong reasons to accept the Pareto condition for the preference aggregation scenario.

When it comes to the aggregation of value rankings, things are different. In this aggregation process it is important to require that the collective judgment as far as possible approximates the judgments of the individuals. Needless to say, all individual value judgments should then be taken into consideration, to equal extent. Individual judgments to the effect that the alternatives that are being compared are equally good should thus be given the same consideration as the competing value judgments. Therefore, if some individuals believe *a* to be better than *b*, but the overwhelming majority believes *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

[10] But couldn't one argue that there is room for a collective sacrifice in this case, since the collective could use for this purpose the resources that would be made available by those of its members who prefer *a* to *b*? This response, however, would not help in the absence of private resources. Then every sacrifice would have to draw on the collective's common resource.

[11] In terms of choice guidance, a preference for *a* over *b* guides one to choose *a* rather than *b* (if confronted with these two alternatives), while willingness to sacrifice something to get *a* rather than *b* guides one to choose *a*-with-sacrifice rather than *b* (if these are the alternatives available and the sacrifice is small enough). Clearly, if one makes the former choice but not the latter, one need not be logically confused.

[12] The picture changes if we think of preference aggregation in welfarist terms. If the collective's goal is to maximize the welfare of its members and the degree of preference satisfaction is identified with the degree of welfare, then, on this interpretation, if some individuals prefer *a* to *b* and everyone else is indifferent, the total welfare is increased by a move from *b* to *a*. On this picture, sacrifices can be justified, even if those who prefer *a* to *b* are few. Even then the collective still has a reason to sacrifice something in order to get *a* rather than *b*, but the size of the sacrifice should be correspondingly small, so that it can be outweighed in its welfare effects by the gains of the individuals who prefer *a* to *b*.

that the Pareto condition will be violated by any reasonable procedure for the aggregation of value rankings.[13]

A worry concerning this view has been raised by Franz Dietrich:[14] On his suggestion, a profile of individual ordinal evaluations might be seen as an incomplete representation of a profile of interpersonally comparable cardinal evaluations of the items that are being compared. This ordinal representation is incomplete to the extent that it admits of a broad range of different cardinal extensions. Now, when it comes to the aggregation of interpersonally comparable cardinal evaluations, it is very plausible to use as the aggregation rule some form of weighted average, in which the evaluation of each individual is given a positive weight. But this means that if some individuals rank *a* above *b*, while all the others rank these two items equally, then on any cardinal extension of this profile of rankings, the weighted-average aggregation rule will deliver a collective evaluation in which *a* is ranked above *b*. In other words, Pareto will be satisfied on every cardinal extension, which suggests that Pareto should be satisfied even when it is not determined which of the cardinal extensions is the right one.

I am inclined to respond to this worry by insisting that what I focus on are cases of aggregation in which what is being aggregated are fundamentally ordinal individual evaluations, i.e. evaluations that are made merely in terms of "better", "equally good" and "worse". They shouldn't be seen as expressions of underlying cardinal comparisons between options.[15] Or, more cautiously, to the extent that cardinal comparisons might be implicit in the individual evaluations, we cannot assume that individual cardinal evaluations are fully interpersonally comparable. This makes such rules as weighted average inapplicable.

In his comments to this paper, Franz Dietrich suggests that there might be another way to defend the rejection of the Pareto condition for aggregation of value rankings. I quote:

---

[13] Note, however, that if one changes the nature of the judgment aggregation task, then Pareto-type considerations might become applicable. 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, if all the other alternatives are according to everyone inferior to *a* and *b*, 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 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.

[14] In his comments at an LSE workshop on deliberation, in June 2011.

[15] Franz Dietrich is not impressed with this reply, however. He writes (in his referee comments): "Should the rejection of the Pareto principle rely on people's (current) inability to develop finer (i.e., cardinal rather than ordinal) judgments?" Well, I don't see why not. What we aggregate are people's current value judgments, after all.

> To me, the best defense of your claim [that the Pareto condition isn't valid as as general principle for the aggregation of value judgements] would be a third route. Suppose value levels are objectively "discrete". For instance, the value of something is either "high"/"good" or "low"/"bad", and nothing in between makes any sense. Or, exactly seven value levels might [make] sense. In such a case, a compromise between two neighbouring value levels is meaningless. If the overwhelming majority of people rank *a* equally valuable as *b* and the remaining people rank *a* over *b*, then (since *a* cannot be "very very slightly" better than *b*, as there are objective value steps) *a* should be socially judged equally valuable as *b*, against the Pareto principle. Note that this defense of your claim doesn't come from the (epistemic) problem that people are unable to form nonordinal value judgments, but from the (metaphysical) problem that value is discrete rather than continuous.

This argument is ingenious and plausible, but – obviously – its presupposition that value orderings are basically discrete might be questioned.

Let me assume, anyway, that I am right in my suggestion that the Pareto condition, which is valid for preference aggregation, doesn't hold for the aggregation of value rankings.[16] What lessons can we draw from this? Now, as it turns out, the Pareto condition *is* violated by any method of aggregation that consists in distance minimization. Intuitively speaking, if in a large majority of individual 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 is maximally similar to the inputs.

[16] There might be other interesting formal differences between aggregation of judgments and aggregation of preferences. One difference was mentioned above: For judgment aggregation it might seem reasonable or perhaps even mandatory to suspend a collective judgment in case of a tie between optimal outcomes of the aggregation procedure. It might not be permissible for the collective to opt for one particular judgmental outcome, if there are other outcomes that are equally satisfactory. The solution is to accept only what's common to all such outcomes. This restriction doesn't seem to apply or at least is not equally compelling in the case of preference aggregation, so far as I can see. Another, related difference concerns the issue of *individual* suspension of judgment. An individual might well suspend judgment, if the evidence isn't conclusive. It is less clear whether suspending preference is equally natural. 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 2008). But abstaining from preference is in any case more problematic 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 of favouring or disfavoring, i.e., if preferring one item to another consists in favouring it to a higher degree or disfavouring it to a lesser degree, then preferences are logically related as well. Transitivity, for example, is on this picture a logical relation: preferring *a* to *b* and *b* to *c* logically implies preferring *a* to *c*. (Cf. Rabinowicz 2012.)

Yet another difference has to do with anonymity. While this condition is nearly always assumed for preference aggregation, it appears less obvious as far as judgment aggregation is concerned: After all, some of the individuals might have more expertise than others, in which case it might seem justified to give more weight to their judgments. (I am indebted to Franz Dietrich for pressing this point.) Still, I think that in many contexts such differences in expertise are deliberately bracketed in the aggregation procedure: Insisting on some memebers of the collective being more competent than others is considered to be inappropriate. Instead, the differences in expertise often play an important role in determining who the relevant members are going to be: When committees are formed, their members are chosen on the basis of competence.

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 (i.e. a majority) ranks a and b equally, then the latter ranking has a shorter average distance to the individual rankings than the former.

Proof: If d is a distance measure, then, for all x and y,

(i) $d(x, y)$ is a non-negative real number that equals 0 iff $x = y$.

Also, symmetry holds:

(ii) $d(x, y) = d(y, x)$.

Consequently, if the number of individuals who rank a above b is m, while the remaining n - m individuals rank a and b equally, (i) implies that 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 $k > 0$ is the distance from the unequal ranking to the equal one. By (ii), 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 if and only if the individuals who rank a and b equally are in majority.[17]

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. Even if only a minority ranks a over b, while everyone else ranks a and b equally, other alternatives might come in between a and b

[17] Note that this result does not presuppose triangle inequality. This means that it holds not only for all distance measures but in fact for the larger class of all cardinal dissimilarity measures. (Cardinality must be assumed if minimization of the average dissimilarity is to make sense).

Note also that this Observation 1 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 the convex transforms of k and 0, which is sufficient for the conclusion we have been after. The result we have proved will also hold for the ‘leximin’ approach, since the number n-m of individual rankings to which the unequal ranking has the maximal distance, k, is larger than the corresponding number (m) for the equal ranking. This result will not hold, however, if we use the simple maximin, because the maximal distance to individual rankings is the same for both the equal and the unequal rankings. But even maximin will violate Pareto, since it will pick out equal ranking as one of the optimal solutions in this case.

in the minority rankings, which complicates the picture. Thus, consider the following case with three alternatives, *a*, *b*, and *c*, and three individual rankings, *x*, *y* and *z*:

<u>*x*</u>	<u>*y*</u>	<u>*z*</u>
a	*c*	*a b*
c	*a b*	*c*
b		

In this example one voter *a* placed above *b*, while the remaining two voters rank *a* and *b* equally. If we use the KS-measure of distance, 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 if one follows the Kemeny rule.

Still, this is just a marginal point. The main lesson is the following: *The Pareto condition marks an important dividing line between aggregation of preferences and aggregation of value rankings. As a result, distance minimization, which violates Pareto, seems fine as an aggregation method 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 for preference aggregation is that the voters who are indifferent may be safely ignored, since their preferences won't be frustrated anyway, whatever ranking the collective decides upon. In case of Pareto, this applies to voters who are locally indifferent, i.e. whose indifference concerns a given pair of alternatives. But "irrelevance of indifferent voters" can be extended to global indifference as well. The principle that expresses the latter 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* 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 isn't worth very much, since Independence of Irrelevant Alternatives is a highly questionable principle. But the proof at least shows that Pareto is related to Indifference, in terms of the underlying intuitions.

Now, while Indifference, as we have seen, is a reasonable principle for preference aggregation, it is 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 **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 (i) the unequal 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.[18]

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, distance zero) than any other ranking.

[18] Just as it was the case with Observation 1, this result will still hold if we replace distance with a dissimilarity measure that does not obey triangle inequality. And it will hold if we minimize the average convex transform of distance or if we apply 'leximin'.

Aggregation of value rankings from the epistemic perspective

This section is just a rough sketch. It poses questions, but leaves them unanswered.[19]

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 (i) the voters are relatively competent with respect to a proposition *A*, i.e., if they are at least better than purely random devices as far as the probability of correctly judging the truth of *A* is concerned, if (ii) their competence is the same, and if (iii) 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 the reliability of the majority judgment converges to 1 as the number of voters goes up to infinity. The underlying intuition behind the jury theorem is very simple: If the voters can be seen as relatively reliable independent sources of information concerning the question at hand, then – clearly - the more such sources 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 value judgments can be 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 for the aggregation of value rankings:[20] 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 being correct is probabilistically independent of the correctness of other voters' judgments.

[19] Some preliminary answers have been reached in a joint work with Stephan Hartmann.
[20] Cf. Condorcet (1785). The presentation that follows is based on Young (1988).

Under these conditions, he argued, the judgments of the majority would be the ones one should 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 cycling cases had to be treated in a more complicated way. (Condorcet's particular proposal how to do it 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 restrictive and should be rejected. Also, it might be of interest to look at individual value rankings holistically, rather than piecemeal, i.e. pair by pair, as Condorcet has done. Instead of ascribing to individuals epistemic competence with respect to each pair of alternatives, separately, we might want to ascribe to them competence with respect to the ranking as a whole. From this holistic 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?

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 distance from truth*. 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. 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 distance to truth.

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 the set of all possible rankings of the available alternatives. 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 all possible rankings of the available alternatives. This value is the probability that a ranking picked at random will be the correct one.

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. What should we expect from a good truth-tracking procedure in such a case? Possibly, at least the following: 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 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 (maximization of expected verisimilitude) 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 unclear. But perhaps something like this could 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 (cf. Rabinowicz 2011b), and it is a fair conjecture that a similar result can be established for other plausible distance measures. In this sense, then, a voter is more reliable that an ignorant person who chooses an option that minimizes expected truth-distance under the state of total ignorance.[21]

When is the aggregation procedure satisfactory from the perspective of the minimization of expected truth-distance? I suppose that at least the following must be required: The expected truth-distance of each optimal ranking (i.e. 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 converge to 0 when the number of voters goes to infinity.

That distance-based methods of aggregating value rankings are satisfactory in terms of increasing the probability of truth and decreasing the expected distance from truth are as yet just unproven conjectures. I hope they 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 that have the form of 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 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 appear to be appropriate for judgment aggregation but *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 of judgment aggregation, the distance-based approaches invite an evaluation from the epistemic perspective. How good are they in increasing the probability of truth and in increasing the expected verisimilitude? Some tentative thoughts on how to go about these epistemic issues were presented above.

[21] Another way of interpreting a voter's competence regarding expected verisimilitude would be to assume that, for any ranking *x*, the probability that the voter chooses that ranking decreases with the increase of the distance between *x* and the true ranking. Should this decrease in probability be proportional to the increases in distance? I am not sure.

Acknowledgements

This paper is a significantly revised version of Rabinowicz (2011a). It was presented at seminars in the Tampere Club, the Swedish Collegium of Advanced Study in Uppsala, departments of philosophy in Lund, London and Edinburgh, and at moral philosophy conferences in Copenhagen and Stockholm. I am grateful to the participants of these events and to other people with whom I have discussed this paper, especially Gustaf Arrhenius, Luc Bovens, Geoffrey Brennan, Erik Carlson, Christian List and Teddy Seidenfeld. Very special thanks are due to Stephan Hartmann, with whom I have done preliminary work on the epistemic aspects of value judgment aggregation, and to Franz Dietrich who has kindly provided me with extremely useful oral and written comments. I also wish to acknowledge the generous financial support from the Swedish Research Council.

Appendix: Measuring distance between rankings

Let **A** be the set of alternatives and let $d(x, y)$ stand for the distance between rankings x and y of the alternatives in **A**. Kemeny and Snell (1962) have shown that the KS-metric, which takes this distance to be the number of ordered pairs of alternatives in **A** with respect to which x and y disagree with other, is the only measure of distance between rankings that satisfies their set of axioms. One of these axioms is the axiom of betweenness:

Betweenness: If a ranking y lies between rankings x and z, then

$$d(x, y) + d(y, z) = d(x, z).$$

Intuitively, being between x and z can be understood as being located on the straight line connecting x with z (or on one of such straight lines, if the relevant geometry allows for several straight lines existing between two points in a given space). On this picture, it is obvious that Betweenness must hold.[22]

Kemeny and Snell's uniqueness result for the KS- measure is based, however, on their particular definition of betweenness, which might well be questioned:

[22] The standard triangle inequality axiom for distance states that $d(x, y) + d(y, z) \geq d(x, z)$, for all x, y, and z. Betweenness adds that this inequality becomes an equality when y is located between x and z.

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 *z* (or to both).

Apart from the betweenness axiom, the other axioms imposed by Kemeny and Snell seem rather innocuous. They stipulate that *d* satisfies 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 a least as large as the distance between *x* and *z* (triangle inequality). In addition, Kemeny and Snell assume that the distance measure is not sensitive to the identity of the alternatives that are being ranked:

Neutrality: *d* is invariant under all permutations of the alternatives.

They also impose two additional axioms:

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.[23]

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 is that this measure 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*:

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 instead to opt for the equal ranking

[23] In fact, this axiom is redundant. As has recently been shown by Can and Storcken (2013), Reduction follows from the other axioms.

a, *b*, *c*

as the collective outcome.

But it is easy to calculate that the average KS-distance from the equal ranking to the rankings in the set {*x*, *y*, *z*} is larger than the average distance from each of the rankings in this set to the set as a whole. While the former is 3, the latter equals 8/3.[24] 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 the average KS-distance to the Condorcetian set of input rankings.[25]

The compromise equal-ranking solution is rejected, because, by the KS-definition of betweenness, the equal ranking is *not* located between any two rankings in the set {*x*, *y*, *z*}. Thus, consider, for example, *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 does not 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

---

[24] 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.

[25] If one thinks that the right approach for the collective is to *suspend judgment* among optimal options, and interprets suspending judgment as accepting only what's common to all the optimal options, i.e., only their intersection (see fn. 4 above), then in the Condorcet paradox the collective output would be the empty set of ordered pairs. In other cases, the intersection of the 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 suspending 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 suspending judgment with respect to the class of optimal options is interpreted as moving to an indeterminate state in which the collective undecided between options in this class (cf. fn. 4 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 see that *if several rankings are optimal, then the class of these rankings will also be optimal*. The reason is that if several rankings have the same (minimal) average distance to individual inputs, then this will also be the average distance from the class of these rankings to individual inputs.

ranking to *y*. The corresponding result holds for the distances form *x* to *z* and from *y* to *z*. Which explains why the equal ranking doesn't come out well on the KS-measure.

Thus, if we want to have a measure that favours compromises, we might be well-advised to redefine the notion of betweenness. This, in essence, is what was done by Cook & Seiford (1978).[26] They defined 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, 2 for the next-best alternative, etc. In case of a tie, we assign the average number to the tied alternatives. Thus, to take an example, if there is one alternatives at the top and two alternatives are tied just below, each of the latter gets the average of 2 and 3, i.e. 2,5.[27] To avail ourselves of some formalism, let a^x be the CS-number of alternative *a* in ranking *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:

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

The equal ranking (*a*, *b*, *c*) does minimize the average CS-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 the betweenness relation interpreted in a new way. 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 **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 $a^x \geq a^y \geq a^z$.

[26] Another alternative would be to keep the KS-measure but replace the Kemeny rule (= minimization of the average distance) with some rule that puts a premium on solutions that do not lie too far away from any of the inputs. This would favour compromise solutions. For example, we could opt for minimization of the average convex transform of distance or for a leximin-type approach (see above). As a matter of fact, Cook and Seiford in another paper suggested using minimization of the average of squared distances as a method of reaching the consensus ranking. Cf. Cook & Seiford (1982).

[27] This kind of numbering was suggested by Kendall (1962). Note that one might just as well reverse the ordering and start from the bottom instead of the top. In addition, one might start with 0 instead of 1. If one then assigns higher numbers as one moves upwards in the ranking, then this version of 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; and the sum of the points received by *a* is its Borda number in a given ranking.

Obviously, a disadvantage of this definition from the intuitive point of view is that it might be accused of begging the issue: It is very closely tied to Cook and Seiford's own method of calculating distance.

Still, 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 example:

<u>*x*</u>	<u>*y*</u>
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. One such proposal is due to Duddy and Piggins (2012). 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. 'Steps' are defined as follows: You need just one step to move from one ranking *x* to another ranking *y* iff you can reach one from the other just by raising or lowering the position of a single alternative *a* with respect to some set *X* of alternatives that are equal-ranked in *x* and making no other changes in the relative positions of the alternatives. This is possible only in two cases: (i) if *a* is equal-ranked in *x* with the alternatives in *X* or (ii) if *a* in *x* is immediately above or below *X*. If (i) holds, then you can raise *a* to a position immediately above *X* or lower it to a position immediately below *X*. If (ii) holds, then you can move *a* to *X*'s level. Thus, to give an example, there is just one step between these two rankings:

<u>*x*</u>	<u>*y*</u>
a	
b c	*a b c*

Thus, on their approach, the distance between these rankings equals 1, while on both the KS- and the CS-approach, this distance is longer (equals 2). Duddy and Piggins (2012) characterize the difference between the Kemeny metric and their own metric as follows: "Under the Kemeny metric, [in each step] we can raise or lower the position of one alternative relative to *just one other* alternative. Under our metric, at each step we can raise or lower the position of one alternative relative to multiple other alternatives, provided those other alternatives are together in a single equivalence class [i.e., provided they are tied]."

KS, CS and DP all order distances between rankings in different ways. Thus, on the CS-approach, the distance between *x* and the equal ranking *y* in the example that we have just given is the same as that from *y* 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 the Condorcet paradox.[28]

---

[28] Might there be some other distance measures that are worth considering? An interesting but still undeveloped suggestion has been made by Gustaf Arrhenius (in private communication) in connection with KS: The KS-measure only considers the ordered *pairs* with respect to which two rankings differ. But why not also look at the 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 ranks *b* at least as highly as *c*. Similarly for quadruples, etc.) If *x* and *y* differ from each other with respect to the same number of pairs as *x* and *z*, but as *x* and *z* 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:

<u>*x*</u>	<u>*y*</u>	<u>*z*</u>	<u>*u*</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 considers 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*?

References

Can, B., and Storcken, T. (2013). "A Re-characterization of the Kemeny Distance", working paper RM/13/009, Graduate School of Business and Economics, Maastricht.

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, 2012, "A Measure of Distance between Judgment Sets", *Social Choice and Welfare* 39, 855-867.

Hedden, B. (forthcoming), "Time-slice Rationality", *Mind*.

Horowitz, S. (forthcoming), "Immoderately Rational", forthcoming in *Philosophical Studies*.

Kelly, T. (2013),"Evidence Can Be Permissive", in M.Steup, J. Turri and E. Sosa, *Contemporary Debates in Epistemology*, 2nd edition, Wiley-Blackwell, 298-311.

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.

If one generalizes the KS-definition of betweenness from pairs to n-tuples, for all *n*, 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) each *n*-tuple that belongs to both *x* and *z* belongs to *y* as well, and (ii) each *n*-tuple that belongs to *y* belongs to *x* or *z* (or to both).

In this definition, especially clause (ii) is very demanding: it implies that there seldom exists a ranking that lies between two rankings. ("Seldom" doesn't mean "never", though. 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.)

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., 2012, "The theory of judgment aggregation: an introductory review", *Synthese* 187, 179-207.

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.

Mayo-Wilson, C. and Wheeler, G., (forthcoming), "Scoring Imprecise Credences: A Mildly Immodest Proposal", forthcoming in *Philosophy and Phenomenological Research*.

Meacham, Ch. (forthcoming), "Impermissive Bayesianism", *Erkenntnis*.

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., 2008, "Value Relations", *Theoria* 74: 18-49.

Rabinowicz, W., 2011a, "Democracy – Two Models", in R. Sliwinski and F. Svensson (eds.), *Neither/Nor - Philosophical Papers Dedicated to Erik Carlson on the Occasion of His Fiftieth Birthday*, Uppsala Philosophical Studies 58, 2011, 219-41.

Rabinowicz, W., 2011b, "Presumption of Equality as a Requirement of Fairness", in E. N. Dzhafarov and L. Perry (eds.), *Descriptive and Normative Approaches to Human Behavior,* Advanced Series on Mathematical Psychology, v. 3, New Jersey: World Scientific 2011, 203-24.

Rabinowicz, W. (2012), "Value Relations Revisited", *Economics and Philosophy* 28: 133-64.

Rosen, G. (2001), "Nominalism, Naturalism, Epistemic Relativism", *Philosophical Perspectives* 15: 69-91.

Saari, D. G. and V. R. Merlin, 2000, “A Geometric Examination of Kemeny’s Rule”, *Social Choice and Welfare* 17: 403-438.

Schoenfield, M. (2014), “Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief”, *Nous* 48: 193-218.

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.

White, R. (2005), “Epistemic Permissivism”, *Philosophical Perspectives* 19: 445-459.

White, R. (2013), “Evidence Cannot Be Permissive”, in M.Steup, J. Turri and E. Sosa, *Contemporary Debates in Epistemology*, 2nd edition, Wiley-Blackwell, 312-323.