

Reaching a Consensus

Richard Bradley
 Department of Philosophy, Logic and Scientific Method
 London School of Economics
 Houghton Street
 London WC2A 2AE

February 10, 2006

Abstract

This paper explores some aspects of the relation between different ways of achieving a consensus on the judgemental values of a group of individuals; in particular, aggregation and deliberation. We argue firstly that the framing of an aggregation problem itself generates information that individuals are rationally obliged to take into account. And secondly that outputs of the deliberative process that this initiates is in tension with constraints on consensual values typically imposed by aggregation theory, at least when deliberation is modelled as process of learning from others compatible with Bayesian updating principles.

1 Introduction

1.1 Conflicting Opinion

Suppose that a decision or series of decisions must be made which depend on the values taken by some set of variables and that there exists a number of different perspectives or opinions on these values. Variables might include quantities of money to be assigned to particular projects, probabilities of events, expected utilities of prospects, or truth values of propositions, while the perspectives might be those of individuals, social classes, times or states of the world. Any situation of this kind can be represented in a table as follows, with the x_i^j denoting the value of perspective i on variable X^j :

		<i>Variables</i>				
		X^1	X^2	X^m
<i>Perspectives</i>	I_1	x_1^1	x_1^2	x_1^m
	I_2	x_2^1
	x_i^j	...

	I_n	x_n^1	x_n^m

When the different perspectives disagree on the correct values of the variables in question, the question arises as to which values should be used in making the decisions: those of the ‘best’ perspective, some amalgam of the different ones, or some third set altogether? Or to put it somewhat differently on what basis might a consensus be achieved on the values to be used in decision making?

There are broadly speaking three different ways of settling the question, or at least trying to: by aggregation or amalgamation, by deliberation or discussion, and by inquiry or learning. The three

operate on somewhat of a different time scale, with inquiry and learning typically occurring over long stretches of time (ontogenetic learning can take millennia, for instance), deliberation and discussion over somewhat shorter ones and aggregation being close to instantaneous. One might thus conveniently think of these processes as occurring sequentially, with inquiry yielding inputs into deliberation which in turn yields inputs into aggregation. This is something of a simplification of course - inquiry is often guided by the deliberations of the enquirers, for instance - but on the face of it it offers a helpful way of subdividing the problem in such a way as to allow each aspect to be examined separately. Such a subdivision is, in any case, widespread, with aggregation theorists taking as their point of departure a deliberative equilibrium in which all exchanges of information and accommodation to others' viewpoints has already occurred, and deliberation theorists assuming that there are no inputs for processes of enquiry separate from the deliberation being studied.

The three methods of achieving a consensus differ in more substantial respects too. Aggregation leaves individual opinion unchanged and if it produces a consensus on some question of common concern, it is a consensus of a derivative kind; one that depends on the existence a prior or more basic consensus amongst individuals concerning the legitimacy of the aggregation procedure employed. Individuals consent to the adoption of the aggregate judgements because they endorse the method that produces it, but there is no further requirement that they adopt the collective judgements as their own.¹ Herein lies the normative significance of the characterisation results of Social Choice theory: By showing what mechanisms are consistent with the various general conditions that it imposes on aggregation functions, they allow agreement on these conditions to form the basis for a commitment to the mechanisms picked out by them, and hence to their outputs.

In contrast to this view of consensus as a by-product of a shared commitment to an aggregation method, is that of consensus as the outcome of an actual convergence of individual judgements as the result of some process; most notably either enquiry and learning or rational deliberation and discussion. The contrast must not be overdone of course, since it is true that the legitimacy of the outputs of both deliberation and enquiry also depends on the acceptance of norms of one kind or another: Norms of free and fair discussion on the one hand and those of scientific method on the other are the salient examples. But to a large extent those involved in a discussion or enquiry aimed at settling some question are bound to accept the outcomes of these processes in a substantive way. This is because these activities are regulated by ideals of objectivity, whose acceptance is a precondition for genuine involvement in the activities. Of course, people do enter into discussions, and even enquiries, in 'bad faith' in the sense that they do not intend to change their opinions, whatever the outcome, but there is a clear sense in which they are exploiting, rather than participating in, these activities.

In this paper, I will examine a number of aspects of the relationship between these ways of securing consensus. In the rest of this introductory section it is argued that a moderate strengthening of the usual rationality conditions that aggregation theorists place on the individual perspectives implies that the aggregation problem is, in a certain sense, unstable and hence that some attempt must be made to accommodate deliberation and learning; in particular learning from others. To assess the prospects for such an accommodation I look in particular at a model developed separately by Morris DeGroot and by Keith Lehrer and Carl Wagner. In section 2 we study the branch of aggregation theory that serves as the backdrop for their model, while in section 3 the model itself is introduced and criticised. I conclude that their model is unsatisfactory from the point of view of the rational belief revision and hence that the accommodation may be less than straightforward.

1.2 Aggregation

Social Choice theorists are apt to treat the problem of finding consensual values as an aggregation problem; that is, as a problem of finding an acceptable mapping from the set of individual perspectives

¹To some it has seemed that this implies that individuals both endorse and reject the judgements in question: this is the heart of 'Wollheim's Paradox'.

to a joint or aggregate one, taking the former as given. Since Arrow at least, they have proceeded by seeking reasonable general conditions on any such aggregation mechanism or function, attempting thereby to constrain the set of mappings worthy of consensus. Three kinds of conditions are worth noting even at this point as they will turn up in our discussion on various occasions.

1. Unanimity (or Pareto) conditions which require that unanimous individual judgements be preserved at the aggregate level.
2. Rationality, and especially consistency, conditions on both individual and aggregate judgements.
3. Domain conditions which require that the aggregation method be applicable over some specified class of possible profiles of individual judgements

That the aggregation problem is framed in Social Choice theory in terms of a relation between a *given* (and often presumed to be fixed) set of attitudes or perspectives and some joint or aggregate attitude, does not imply that the social choice theorist must deny the importance of inquiry and deliberation. Rather, as we noted before, her starting point is the assumption that these forces have played themselves out without producing consensus. It is true, of course, that Arrow and many others since have in fact required of aggregation functions that they should be able to handle any profile of rational individual judgements – the so called Universal Domain condition. The thought here seems to be that the aggregation method itself should impose no constraints on judgements of individuals; individuals are in this sense sovereign. The justification for the condition is thus normative and unaffected by the fact that the judgements of individuals belonging to the same society or culture are typically correlated or even that this correlation may be partially constitutive of the group to which they belong.

This being said, the requirement of a universal domain implicitly commits the social choice theorist to the view that rational deliberation and enquiry need not constrain the possible judgements of individuals in any particular way. Implicitly the aggregation theorist regards any input from completed processes of deliberation and enquiry to be compatible with her basic assumptions and, in particular, the assumption of individual rationality. The deliberation theorist, on the other hand, may willingly accept the applicability of the principle of individual sovereignty to the views that might be brought to the table for discussion, but deny that all combinations of individual's viewpoints are possible outputs of deliberation. Some views may not be rationally co-tenable, given norms of rational dialogue or requirements of rational response to the expressed views of others. A similar concern may of course be expressed about a universal domain condition for deliberation on the grounds that views contrary to established scientific knowledge or to moral norms should not be allowed into discussion, but this possibility will not occupy us here.

1.3 Deliberation

Deliberation as a means of securing consensus has many advocates, especially in democratic political theory: see, for instance, Dryzek [9], Fishkin [10] and Bohman [4]. But the mechanisms by which deliberation produce consensus or more generally correlations in judgements have not been given much attention by this literature which often seems to rest on confidence in the power of reason to bring people together. But there are both empirical and normative grounds for taking the issue seriously. Regarding the former, there is a good deal of evidence to suggest that people not only conform to social norms of judgement, but that homogeneity is often (though not always) increased by deliberation and exposure to the opinions of others: see, for example, the evidence cited in Sunstein [32].

More importantly for the normative point of view, there are situations in which others' judgements, or the expressions of them, provide grounds for modifying one's own; situations in which others speak with a certain authority. The most obvious examples are those in which somebody has information that one does not hold. For instance, if someone is able to make an observation concerning the value

of some variable X , then their testimony to the effect that $X = x$ should lead one to adopt the same value. But someone's authority on a question may be more complicated than merely a matter of additional information; it might, for instance, derive from some special expertise that they have or special training or method. Doctors may be able to make better judgements about one's condition because their diagnostic abilities have been honed by experience, even though they may have no special information about one's condition.²

Cases in which other people's judgements are grounds for revising one's own are not confined to expression of belief or knowledge, but extend to value and preference judgements too. Furthermore, an expression by someone of their preferences may be informative in more than one way:

1. *Interdependence of preferences.* If you have a general desire to see the preferences of someone you care for fulfilled, then their expressions of preference for some outcome give you reason to wish for that outcome too, and possibly to try and bring it about.
2. *Informational content.* Someone's expressions of preference carry information about their beliefs which in turn reflect the information available to them. Inferring what this information is could lead one to revise one's own beliefs.
3. *Evaluative content.* Someone may have authority in the domain of value judgements by virtue of their 'taste' or special capacity for judgement. When a fashion guru declares something to be in style, for instance, others will rush out to buy.

An example may serve to illustrate the distinction between these cases. Suppose that Alice has visited a number of restaurants, ranked them in accordance with her preferences and informed Bob of her ranking. Case (1): Bob revises his preferences over restaurants because he wants to go with Anne to the one which she will find most congenial. Case (2): Bob knows that Alice cares only about price and infers that her ranking reflects how expensive each restaurant is. His revised estimate of the cost of visiting each restaurant founds new relative preferences for eating at each one. Case (3): Bob considers Alice's tastes to be exemplary and so adopts her ranking forthwith.

The observation that expressions by others of their beliefs and preferences can be informative suggests that a central tenet of scientific methodology, the Principle of Total Evidence, applies in cases of conflicting opinion. The principle says that my beliefs must be consistent with all the evidence available to me and, further, if I acquire new evidence, I should revise my beliefs to accommodate it. So it follows, in particular, that if I believe that someone is an authority on question A , then I should revise my beliefs in the light of his or her expressed judgement about A , just as I should revise my opinion in the face of any reliable evidence concerning A .

The implication is that, in a certain sense, the problem that the aggregation theorist takes to have been defined at the beginning of the paper may not be a stable one, given the strengthened rationality condition introduced here. For the very statement of the original aggregation problem generates information about the judgements of the various individuals which the Principle of Total Evidence requires each to take into account. And only in very special cases will this not imply that individuals are rationally obliged to revise their judgements. These revisions need not produce a consensus, of course, so a new aggregation problem will emerge. But this one may not be stable either, for it is possible that the manner in which others revise their judgements generate information about what they know about the relative reliability of the judgements of some other person. Ann may know that Bob knows whether Cara is reliable of matters X . If she observes Bob to revise his judgements on X so as to bring them in line with Cara's, then Anne may surmise that Cara's judgements on X can be expected to be closely correlated to the truth. Anne may now wish to revise her own judgements once again in the light of this. In practice no doubt such revising will come to an end at some point; in principle it need not.

²See Jeffrey [15, chapter 1] for a discussion of this point.

The upshot is that it is not open to the aggregation theorist to both require the kind of judgemental rationality expressed by the Principle of Total Evidence and to assume that the judgements of individuals are in an equilibrium state in which the forces of deliberation and learning have played themselves out. What difference does it make to the consensual values we are inclined to regard as plausible, if we pay attention to deliberative forces? It is not impossible that the answer is ‘none’, that the consensual judgements determined by an aggregation function on the initial problem are just those determined by the function on each of the problems derived by successive revisions of the individual judgements. But it would be surprising if this were so, for it would suggest that all the learning that goes on at the individual level matters not a jot at the aggregate level. And if it does make a difference, then I think we must accept the need to review our attitude to the results of aggregation theory, asking ourselves in particular whether they are robust with respect to the possible effects of deliberation and learning from others. It is difficult, however, to answer the question in any kind of general way, since the precise relation will depend on what we take as our theories of aggregation and deliberation for the purposes of comparison. What I propose to do therefore is to study the relationship between the two for rather salient instances of both in the belief that they are sufficiently representative to support the general claim that aggregation theory is not easily reconciled with the basic lessons drawn from consideration of deliberation.

2 Aggregation and Allocation

The table with which the paper began defines an aggregation problem of the following kind: can we determine for any given m perspectives on n variables, an aggregate or collective perspective on these variables? More formally, can we define a mapping from any $m \times n$ table to a m -place vector representing the collective perspective on the variables? In its different instantiations, the question has been extensively studied in social choice theory, statistics and elsewhere. In one rather general form the problem is at the heart of the theory of judgement aggregation developed by List, Dietrich and others: see for instance [23], [7] and [8]. The problem of aggregation of diverse probability judgements has mainly occupied the attention of statisticians: Genest and Zidek [11] contains a very useful overview. The problem of aggregating utilities, qua values of a function representing a ranking of some kind, received its classic treatment in the work of Sen [30] and has since been studied by amongst others Gevers [12], D’Aspremont and Gevers [1] and Roberts [28]. The connection between this literature and our problem is somewhat disguised by the fact that in it the problem is typically posed, not in terms of aggregating utilities, but of finding a mapping from sets of utilities to the social ranking of prospects that such an aggregation would induce (in Sen’s vocabulary, of characterising social welfare functionals). But the overlap is clear enough. Finally the problem of simultaneously aggregating probability and utility judgements has been studied by, amongst others, Mongin [26], Broome [2], Seidenfeld et al [29] and Bradley [5].

In this section we focus on linear averaging as a solution to a particular class of aggregation problem, that we call allocation problems. Linear averaging functions have a certain saliency in the literature on probability aggregation and are quite common in the preference and utility aggregation literature too. But I do not by any means intend to claim that they are the only aggregation functions of interest. Geometric averaging is a well established alternative form of probability aggregation, for instance; while minimax and leximin rules are at least as salient as Utilitarian ones in the Social Choice literature. What I do claim, however, is that the study of linear averaging in the context of allocation problems is illustrative of quite general difficulties facing the integration of theories of aggregation and deliberation. Since the axioms which characterise it in this case are both familiar and compelling, I will concentrate on supporting this claim by arguing for the applicability of these results to a number of domains of interest and showing their connection to the literature referred to above.

2.1 Characterising Linear Averaging

There exists a number of different axiomatic characterisations of linear averaging functions, but here we draw mainly on the treatment of the problem by Wagner (see especially [34] and [19]). First we define allocation problems as the class of aggregation problems for which the values of both the individuals' judgements and that of the collective sum to a constant. We then state a pair of theorems for allocations problems, the first due to Wagner and the second a generalisation of it to cases where judgemental values need not be non-negative. The essence of these results is that the only aggregation functions on the set of allocation problems that both satisfy an independence of irrelevant alternatives condition and preserve unanimous individual judgements are the linear averaging functions i.e. those functions that determine the aggregate judgement to be a weighted sum of the individual ones.

Let I be a set of n individual perspectives, $V = \{V^1, \dots, V^m\}$ be a set of m variables and $X \subseteq \mathfrak{R}$ be a set of permissible judgemental values for these variables. The contents of X can, of course, vary with the kind of judgement involved, but we assume here that $\exists S \in X$, such that $\forall x \in X$, $S - x \in X$. Sets with this property include the set of real numbers, the open interval $[0, 1]$ and the set $\{0, 1\}$ (in the latter two cases, $S = 1$). Let $\Gamma(m; S) \subseteq X^m$ be the set of vectors $v = \langle v^1, \dots, v^j, \dots, v^m \rangle$ of values such that $\sum_{j=1}^m v^j = S$.

An **allocation problem** $\Gamma(m, n; S)$ is formally identified by the set of all $m \times n$ matrices with rows in $\Gamma(m; S)$. For any $A \in \Gamma(m, n)$ let A^j be its j th column and A_i its i th row. A solution to an aggregation problem, an **allocation function** ϕ , is a total mapping from $\Gamma(m, n; S)$ to $\Gamma(m; S)$. We denote $\phi(A)$ by $a = \langle a^1, \dots, a^m \rangle$. Note that by definition an allocation function respects a universal domain condition. The following additional conditions on such functions will be referred to in the discussion.

Axiom 1 *Independence of Irrelevant Alternatives (IIA):* $A^j = B^j \implies a^j = b^j$

Axiom 2 *Zero Unanimity (ZU):* $A^j = \langle 0, 0, \dots, 0 \rangle \implies a^j = 0$

Axiom 3 *Unanimity (U):* $A^j = \langle a, a, \dots, a \rangle \implies a^j = a$

Theorem 4 (Wagner [19, p. 121-2]) *Let $\Gamma(m, n; S)$ be an allocation problem and ϕ be an allocation function for it. Suppose that $\forall A \in \Gamma(m, n; S)$, $a_i^j \geq 0$ and that $m \geq 3$. Then if ϕ satisfies IIA and ZU, there exists $w_1, \dots, w_n \in \mathfrak{R}$ such that $w_i \geq 0$, $\sum_i w_i = 1$ and $\forall A \in \Gamma(m, n; S)$:*

$$a^j = \sum_{i \in I} w_i a_i^j$$

Corollary 5 *Let $\Gamma(m, n; 1)$ be an allocation problem based on the set $X = \{0, 1\}$ and ϕ be an allocation function for it. If $m \geq 3$ and ϕ satisfies IIA and ZU, then ϕ is dictatorial in the sense that there exists $i^* \in I$ such that $\forall A \in \Gamma(m, n; 1)$ and $\forall V^j \in V$, $a^j = a_{i^*}^j$.*

Proof. Note that since $X = \{0, 1\}$ and every row sums to 1, it follows that every m -place vector $v \in \Gamma(m; 1)$ contains a unique value v^{j^*} equal to 1 with every other value equal to 0. Now since ϕ satisfies IIA and ZU, it follows by Theorem 4, that $\exists w_1, \dots, w_n \in \mathfrak{R}$ such that $w_i \geq 0$, $\sum_i w_i = 1$ and $\forall A \in \Gamma(m, n; 1)$:

$$a^j = \sum_{i \in I} w_i a_i^j$$

Let I_D be the set of all individual perspectives $i \in I$ such that $w_i > 0$. We now show that I_D is a singleton set. Suppose to the contrary that I_D can be split into two disjoint non-empty subsets I_1 and I_2 . Then consider the matrix B and variables V^1 and V^2 , such that $\forall i \in I_1$, $b_i^1 = 1$ and $b_i^2 = 0$ and $\forall i \in I_2$, $b_i^1 = 0$ and $b_i^2 = 1$. Now $b^1 = \sum_{i \in I_1} w_i b_i^1$ and $b^2 = \sum_{i \in I_2} w_i b_i^2$. But either $b^1 = 0$ or $b^2 = 0$. Hence either $\forall i \in I_1, w_i = 0$ or $\forall i \in I_2, w_i = 0$. But this contradicts the assumption that $\forall i \in I_D, w_i > 0$. Hence I_D must be a singleton set. ■

Theorem 6 Let $\Gamma(m, n; S)$ be any allocation problem and ϕ an allocation function for it. If $m \geq 3$ and ϕ satisfies IIA and U, then $\exists w_1, \dots, w_n \in \mathbb{R}$ such that $w_i \geq 0$, $\sum_i w_i = 1$ and $\forall A \in \Gamma(m, n; S)$:

$$a^j = \sum_{i \in I} w_i a_i^j$$

Proof. Let A be any member of $\Gamma(m, n; S)$ and B be an $m \times n$ matrix such that $b_i^j = a_i^j + \beta$, where $\beta = |\min(a_i^j)|$. Since $b_i^j \geq 0$ and the row values of the B_i sum to $S + m\beta$, it follows that B is an allocation matrix. Now let $\bar{\phi}$ be a mapping on the allocation problem $\Gamma(m, n; S + m\beta)$ such that $\bar{b}^j = a^j + \beta$, where $\bar{\phi}(B) = \langle \bar{b}^1, \dots, \bar{b}^j, \dots, \bar{b}^m \rangle$. Clearly $\bar{\phi}$ satisfies IIA because ϕ does. Note that if $B^j = \langle 0, 0, \dots, 0 \rangle$ then $A^j = \langle -\beta, -\beta, \dots, -\beta \rangle$. Hence by Unanimity, $a^j = -\beta$ and so $\bar{b}^j = 0$. It follows that $\bar{\phi}$ satisfies ZU as well. Furthermore, note that $\sum_j \bar{b}^j = \sum_j (a^j + \beta) = S + m\beta$ and so $\bar{\phi}$ is an allocation function on $\Gamma(m, n; S + m\beta)$. Hence by Theorem 4 $\exists w_1, \dots, w_n \in \mathbb{R}$ such that $w_i \geq 0$, $\sum_i w_i = 1$ and $\bar{b}^j = \sum_i w_i b_i^j$. Hence $\bar{b}^j = \sum_i w_i (a_i^j + \beta) = \sum_i w_i a_i^j + \beta = a^j + \beta$. So $a^j = \sum_i w_i a_i^j$. ■

2.2 Probabilities

Let us begin by noting a few consequences of these theorems. The application of Theorem 4 to the problem of aggregating probability judgements across an event partition is immediate since these judgements must be non-negative and sum to one. In this context ZU expresses a unanimity condition for judgements of zero probability, naturally interpreted as the condition that if all individuals are sure that something is false then this should be accepted collectively. IIA expresses the requirement that the aggregate probability for any event is independent of the individual probability judgements regarding any other event in the partition (though, of course, the *individuals'* probabilities for the various elements across the partition are interdependent). Wagner's theorem then establishes that if the partition consists of at least three events then any probability aggregation function satisfying IIA and ZU must assign to these events probability values that are a weighted average of the values assigned to them by the individuals.

Similar results have been established for aggregation of probability functions on σ -fields of events in which the complexity of the event space allows for weaker conditions than IIA. McConway [25], for instance, shows that the linear pooling function is implied by ZU and the requirement that the aggregation function satisfies what he calls the marginalization property: that it should commute with the process of reducing the σ -field to a sub- σ -field. Equally it has been noted that linear averaging fails to satisfy some intuitively appealing conditions; notably the condition that if two events are probabilistically independent for all individuals, then they should be for the aggregate too (Laddaga's [16] independence preservation principle) and the condition that the aggregate of the individual conditional probabilities given some event E should be the conditional aggregate probability given E (the externally Bayesian condition - see Genest and Zidek [11]). The latter, but not the former, condition is in fact satisfied by geometric averages which could be a strong argument in its favour. We defer consideration of these issues till later, however.

2.3 Truth Values

Theorem 4 is equally applicable to problem of aggregating rational bivalent judgements, canonically of truth or falsity, on a set of mutually exclusive and exhaustive propositions. In this context ZU says that the unanimous judgement that some proposition is false must be accepted at the aggregate level, while IIA says that whether a proposition is judged true or false on the aggregate depends only on the truth values individuals assign to it. Once again, this is implicitly to assume that the interdependence of truth-values across the set of propositions is accommodated at the level of individual judgements. Given these assumptions, Wagner's theorem implies that aggregate truth-value of any proposition is

an average of the truth-values of individual judgements. Because truth-value judgements are bivalent, this conclusion imposes a very severe restriction on aggregation: if a proposition is judged false on the aggregate, for instance, then every individual judging it true must be zero-weighted in the aggregation function. Just how severe this is is shown by Corollary 5, which establishes that the only allocation functions on rational bivalent judgements satisfying ZU and IIA are dictatorial ones i.e. functions whose values are determined by those of a single individual.

Social Choice theory is replete with dictatorship results, but Corollary 5 is most closely connected to Dietrich and List's [7] results for the aggregation of bivalent judgements on sets of propositions that are logically connected in various possible ways. In particular their Theorem 2 establishes that an aggregation function on what they call a 'strongly connected' set of propositions that satisfies a universal domain, rationality, unanimity and independence condition must be dictatorial in form. Since a set of mutually exclusive and exhaustive propositions is strongly connected in the required sense, Corollary 5 is in all likelihood a consequence of their theorem. What is illuminating about the former, however, is that the path to dictatorship goes via linear averaging.

2.4 Utilities

The application of Theorem 6 to the aggregation of value judgements defined on a partition of possibilities is somewhat less direct because the fulfillment of the conditions for an allocation problem is not guaranteed in this case. Much depends on background assumptions about the measurability and comparability of utility judgements. Suppose that utilities are measured on a cardinal scale so that the values of the utility judgements are unique up to positive affine transformation. In this case fulfillment of the condition that matrix rows sum to a constant can be assured by the following transformation of the individual utilities. Suppose we are given utility measures u_1, u_2, \dots, u_n for our n individual perspectives which determine row sums S_1, S_2, \dots, S_n for a matrix associated with some partition X^1, \dots, X^m . Then taking the utilities of individual i^* as the point of reference, we define a zero normalisation by setting:

$$u'_i = \frac{S_{i^*}}{S_i} u_i - S_{i^*}$$

so that the sum of the $u'_i(x^j)$ is the same for each individual i ; namely 0.³ Clearly something along these lines is also possible if the numerical expressions of utility judgements are even less determinate. On the other hand, when utilities differences are comparable in some relevant sense, as they may when the utilities represent individual welfare, such a co-scaling would be tantamount to changing the description of the aggregation problem.

Even when co-scaling is possible, there is a further question as to whether it is defensible. Often it can be defended on normative grounds, though the details would depend on exactly what utility is supposed to be a measurement of and hence the nature of the judgements under consideration. The relevant case here is the one in which utility represents a subjective value judgement: canonically a value or preference ordering across a partition of prospects. In this case the utilities represent the relative strength of the individuals' preferences for each prospect *relative* to the alternatives. Individuals may differ in the relative strengths of their preferences, but there is no sense to the idea that they could systematically, across the whole partition, differ in preferences in a uniform direction. One cannot relatively prefer *all* the alternatives more than someone else.⁴ Consider, for instance, a case in which the differences of opinion concern the merits of various candidates for a job. One may argue against the proposed co-scaling of opinion on the grounds that one person may find all the candidates to be good and another that they are all bad. But then we may introduce an additional 'slack' variable into the

³This normalisation amounts to a choice of origin for the individual utilities. Nothing hangs on this choice however.

⁴Hence there can be anything like an expensive tastes argument across an entire partition of possibilities.

aggregation problem to refer to the class of potential applicants who did not in fact apply and co-scale judgements over this alternative partition.

Suppose that utilities are indeed co-scaled in some manner so that we have a utility allocation problem. In this context too, axioms ZU and U are most naturally interpreted as unanimity conditions, but given that the utility values on display may be non-unique numerical representations of comparative preferences, some caution must be attached to this interpretation. When the zero has no particular significance, then the role of ZU is really that of a scaling convention for the aggregate utility. The same applies to the unanimity condition U, although the co-scaling it requires is more demanding. IIA says that the aggregate utility of any prospect depends only on the individual utilities for that prospect. Theorem 6 then says that any utility allocation function respecting IIA and U must determine the aggregate utility by taking a weighted sum of the individual utilities.

Theorem 6 has some obvious kinship with the various characterisation theorems for Utilitarian-like aggregation functions found in the social choice literature, but the relationship is actually quite complicated. (As the comparison is of only marginal interest to our main discussion, the various technical results referred to below are proved in the appendix). In addition to the superficial differences mentioned at the beginning of the section, there is the more significant fact that our result applies only to allocation problems, not the more general aggregation problem that typically occupies social choice theorists. In fact once the assumption that the matrix row sums are equal is dispensed with, the unanimity and independence of irrelevant alternatives conditions only support a much weaker characterisation result - see Theorem 10 in the appendix. In this general case, though the aggregate judgement may still be represented as a weighted average of those of the individuals, the weights can vary with the matrix defining the particular aggregation problem. This result is roughly of the same order of strength as Harsanyi's Utilitarian theorem, but is much weaker than results obtained in other multi-profile approaches to the aggregation problem e.g. the characterisation of weighted averaging given by Gevers [12, Theorem 2].

What connects our discussion of allocation problems to the corresponding social choice literature on aggregation is the role played by 'comparability' assumptions: assumed informally here and explicitly axiomatised by Sen, D'Aspremont, Gevers and others. Crucial to the characterisation result of Gevers mentioned above, for instance, is the assumption that form of the aggregation function should be invariant with respect to the choice of origins for the individual utility measures; glossed as the assumption that utility gains, but not utility levels, are interpersonally comparable. In the context of the interpretation at hand - namely of utility as a measure of relative preference strength - the motivation for such an assumption derives from the thought that the comparison between two prospects should not depend on how we scale the measures of individual preference strength, so long as the *co*-scaling condition is met. Formally, in any case, the upshot of requiring that the transformation of an aggregation problem based on a set $\{u_i\}$ of utility measures achieved by taking for each individual i the utility $u'_i := \alpha u_i + \beta_i$ where $\alpha, \beta_i \in \mathbb{R}$ are such that the β_i but not α can vary from individual to individual, should not affect the ordering of prospects induced by the aggregate judgements, is that the weights referred to in Theorem 10 must be invariant across aggregation problems - returning us to the Gevers result. In a similar fashion requiring invariance of the ordering of prospects induced by the aggregate judgements for any transformation of the order $u'_i := \alpha_i u_i + \beta$, where the α_i can vary from individual to individual, will suffice for a dictatorship result.

The exception to this pattern of recouping existing results is the characterisation of the strict Utilitarian rule - that social utility is equal to the sum of individual utility - which in our framework does not require any assumption of interpersonal comparability: see Corollary 11 in the appendix.

2.5 The Interpretation of Weights

Two closely related questions are raised by the characterisation results for linear averaging: how are the weights to be chosen and what significance is to be attached to them? One natural answer is that the weight attached to an individual should indicate something like his or her epistemic authority on

the question at hand. As we noted before, someone's epistemic authority might reflect their access to information or their judgemental expertise or their possession of a reliable methodology. Another answer, appropriate in a different kind of context, would be that the weight attached to an individual reflects, not their epistemic authority, but their political or social authority or their moral right to have their view taken into account. The first interpretation is perhaps the most salient one when it is probability judgements that are being aggregated, the latter when it is value judgements that are under consideration. But other combinations are possible: On a cognitivist view of value judgements, for instance, the epistemic interpretation would be applicable to value judgements too, while democratic theory is full of normative arguments for equal weighting of individuals' beliefs.

On either of these two interpretations, however, there is no reason for supposing that the weight assigned to someone should be independent of the question at hand: Authority, both epistemic and political, varies with the domain of judgement. Variability in weight assignment is at odds however with the consistency of aggregate probability judgements. To see this consider a different (and more ambitious) aggregation problem; that of defining an aggregation function for any set of individual probabilities defined on a Boolean algebra of events, Ω . Assume that IIA and ZU apply to any derived matrix of individual probability judgements across a partition of Ω . Consider two different partitions of Ω and the associated $n \times m$ and $n \times m'$ matrices A and B . Let C be the $n \times m.m'$ matrix representing individual probabilities across the intersection of these two partitions with C^{jk} denoting the column of probability values for the intersection of the event whose probabilities appear in column A^j and the event whose probabilities appear in column B^k . So $a_i^j = \sum_k c_i^{jk}$ and $b_i^k = \sum_j c_i^{jk}$. Assume that for some $w_1 \dots w_n \in \mathfrak{R}$ such that $\sum_i w_i = 1$, it is the case that $c^{jk} = \sum_i w_i c_i^{jk}$. Then by the laws of probability:

$$a^j = \sum_k c^{jk} = \sum_k \sum_i w_i c_i^{jk} = \sum_i w_i \sum_k c_i^{jk} = \sum_i w_i a_i^j$$

Similarly:

$$b^k = \sum_i w_i b_i^k$$

So the weights on the individuals' judgements on the two partitions are the same independently of what kind of events they contain. But this is in stark contrast to our claim that variation in weight assignment is to be expected.

A similar conclusion follows from consideration of the corresponding wider problem of aggregating expected utility judgements defined on a Boolean algebra of prospects, though the situation is rendered considerably more complicated here by the fact that differences in individuals' utilities for prospects belonging to a coarse partition of the algebra may reflect differences in belief rather than disagreement along any non-credal dimension. The problems associated with simultaneous aggregation of utilities and probabilities are legion⁵, so let us simplify matters by assuming that individuals have identical degrees of belief, measured by probability p on Ω . Then using the matrices identified in the previous paragraph, but interpreting the values as expected utilities we can reason from the assumption that for some $w_1 \dots w_n \in \mathfrak{R}$ such that $\sum_i w_i = 1$, it is the case that $c^{jk} = \sum_i w_i c_i^{jk} p^{jk}$, and from the laws of desirability⁶ that:

$$a^j . p^j = \sum_k c^{jk} . p^{jk} = \sum_k p^{jk} \sum_i w_i c_i^{jk} = \sum_i w_i \sum_k c_i^{jk} p^{jk} = p^j \sum_i w_i a_i^j$$

Hence:

$$a^j = \sum_i w_i a_i^j$$

⁵ See for instance Mongin [26], Broome [2] and Bradley [5].

⁶ See Jeffrey [14].

Similarly:

$$b^k = \sum_i w_i b_i^k$$

So again the weights are invariant.

It is but a short step from these observations to a result that bears an obvious resemblance to the Sen's [30] 'Impossibility of a Paretian liberal' theorem and the 'No experts rights' theorem of Dietrich and List [7]. For were we to require that for some individual i the Boolean algebra Ω contains an event E_i which is such that the aggregate judgement on E_i is determined by i 's judgement on it, it would follow immediately that i was a dictator in the sense of determining the aggregate judgement on all events.

If we are to retain the idea that the pooling weights express the relative authority of the individuals then we can go one of two ways. Either we must conclude that they express, not domain specific authority, but some kind of overall or average authority with respect to all questions under consideration.⁷ Or we must allow weights to vary but give up the requirement that the aggregate judgements formed with respect to different questions be coherent. Neither is particularly enticing. The first leads directly to the problem that aggregation by linear pooling produces sub-optimal judgements, where optimality is measured with respect to the dimension of authority expressed by the pooling weights. For suppose that two individuals i and j are respectively authorities on the question of whether it is the case that A and that B . Then to form an aggregate judgement on both questions on the basis of their average competence over the two would be to produce a judgement that is inferior to one that took i 's judgement on the former and j 's on the latter. The second route, on the other hand, requires us to abandon what is often called the assumption of group or social rationality and accept that the aggregated probabilities or expected utilities are not themselves coherent probabilities or expected utilities, but simply values to be used in group decision making in a role similar to those played in individual decision making by probability and expected utility values. I shall adopt the second route here for the sake of argument and note that by dropping the group rationality condition we simultaneously evade the aforementioned criticism that linear averaging is not externally Bayesian. For this condition too presupposes that the group judgements are subject to the same rationality constraints as those of the individuals making it up.

How then are the weights to be determined? It may seem that, despite all the concessions, we have made no progress at all, for it is surely possible that the weights that should be attached to each individual's judgements may be as much an object of disagreement as the original question. Indeed, if an individual has already made the best judgement that they can in the light of the evidence that they have before them, surely each individual should attach maximum weight to his own judgement. This latter thought is mistaken, however, since it is not rationally obligatory to take our own judgements to be more reliable than those of others, even if they are the best judgements one is able to achieve. For as we noted in the previous section, others may have information or skills that one does not possess. Nonetheless the possibility of disagreement about weights remains, since we may not in fact judge authority according to the same criteria or on the basis of the same evidence. Consensus on the form of the aggregation function induced by agreement on the unanimity and independence of irrelevant alternatives conditions thus produces a second order aggregation problem over opinion on authority. And in this respect opting for geometric averaging rather than linear averaging, will leave us no better off.

Introducing considerations of epistemic and moral respect also takes us back to the central observation of the previous section: that when we respect the judgements of others, rationality requires us to revise our judgements in the light of theirs. This suggests a different path out of the problem: Instead of seeking a second order consensus on respect weights, it might be more appropriate to reshape the original aggregation problem in line with the revisions to individuals' judgements required by the

⁷I am grateful to Philippe Mongin for this suggestion.

Principle of Total Evidence. It is to a study of the implications of this kind of deliberative updating of judgements in the light of those posted by others that we now turn.

3 Deliberation

“It is natural to assume that a group sharing the same information and respecting each other’s opinions may reasonably disagree. I shall prove, on the contrary that if the members of such a group search for the truth and accept the mathematical implications of their state, then they must converge toward consensus. Disagreement is demonstrably irrational”. - Lehrer [17, p. 327]

In this section I want to explore a very simple picture of deliberation which conceives of it as a type of learning from the testimony of others. On this picture deliberation within a group takes the form of each individual reporting their opinions on a range of issues, observing the reports made by others, revising their own opinions in the light of them and then reporting their new opinions. I do not want to suggest that this, by any means, gives an exhaustive characterisation of what happens when people actually deliberate, or even (though its more plausible) that it can accommodate all the normative constraints that deliberation imposes on our belief and preference attitudes. The idea is just to give a base-line model from which we can assess potential theories of what a satisfactory deliberative resolution of differences of opinion will look like.

To flesh out a model of deliberation of this kind, the manner in which individuals should respond to the reported judgements of others needs to be spelled out. In the discussion that follows, I will focus on a model of iterative respect-driven updating separately developed by Morris DeGroot [6] and by Keith Lehrer and Carl Wagner [19]. Their model is applicable to the revision of both probability and utility judgements, but we will confine attention to the former. The model and especially the claim of Lehrer and Wagner that iterative linear pooling represents the uniquely rational way of aggregating probability judgements has been attacked from a number of different angles: See for instance Goodin [13] and the papers appearing in *Synthese* vol. 52 (1985). The discussion here picks up mainly on the themes addressed in Loewer and Laddaga [24].

3.1 A Model of Iterative Pooling

Suppose a group of individuals declare their probabilities for a set of events. Each individual’s respect for the information held by others as well as their expertise is reflected in a weighting he or she attaches to their declared judgements. This respect commits them, we have argued, to revising their probabilities in some way. In the DeGroot and Lehrer and Wagner’s model (hereafter the DLW model) this is done by each individual by adopting the respect weighted sum of the set of declared (prior) probabilities as their new (posterior) probabilities.

Consider a simple example with three individuals with prior probabilities for some particular event of 0.1, 0.2 and 0.15 respectively. Let the respect of each individual for each one of them (including themselves) be given by a 3×3 matrix with values in the interval from 0 to 1. Then the posterior probability of the event for individual i , given the posted probabilities of all, is obtained by multiplying each individual’s prior by i ’s respect for that person and taking the sum. In the case displayed below this yields consensus on a probability of 0.14 for the event in question.

$$\begin{array}{rcccl}
 & \text{Weights} & \text{Priors} & \text{Posteriors} & \\
 \text{Individual 1} & \begin{pmatrix} 0.4 & 0.2 & 0.4 \end{pmatrix} & \begin{pmatrix} 0.1 \\ 0.2 \\ 0.15 \end{pmatrix} & = & \begin{pmatrix} 0.14 \\ 0.14 \\ 0.14 \end{pmatrix} \\
 \text{Individual 2} & \begin{pmatrix} 0.6 & 0.4 & 0 \end{pmatrix} & & & \\
 \text{Individual 3} & \begin{pmatrix} 0.5 & 0.4 & 0.2 \end{pmatrix} & & &
 \end{array}$$

So long as the weights are strictly positive, revisions of this kind will produce a convergence of opinion, though not typically a consensus. Consensus *can* be induced however if this process is repeated in the following way. Suppose that each individual is apprised of everyone else's new probabilities. Then, given the earlier observation that respect for others can extend to their judgements of each other's expertise, individuals should continue to revise their opinions in the light of those of the others until they have exhausted all the information contained both in the posted judgements and in the revisions they induce on others. In general at each stage:

$$p_i^k(X) = \sum_j w_i^j \cdot p_j^{k-1}(X)$$

where:

- 1.) p_i^k is i 's probability assignment after k updates:
- 2.) w_i^j is i 's level of respect for j , and
- 3.) $\sum_j w_i^j = 1$.

Iterated revision of this kind will eventually produce a consensus in a rather broad class of cases: roughly whenever there is some individual i who respects him or herself and is such that there is a chain of strictly positive respect from each member of the group to i .

This model simplifies in one important respect: It assumes a constant respect weight at each level of the iteration. But of course I might consider that someone is a poor judge of tomorrow's weather but an impeccable judge of people's skill at judging the weather; at 'knowing who knows'. So in a more general model respect weights should be allowed to vary with level of iteration. This added realism changes nothing to the main conclusion however; namely that iterated revision of this kind will eventually produce a consensus in a rather broad class of cases (see Berger [3] and Wagner [33] for proofs of this claim).

3.2 Problems with Pooling

From our point of view the most significant thing about the DLW model is that it solves the problem that was posed in the previous section; namely that of how to determine the weights to be used when aggregating by linear averaging. It does so by harnessing a process of iterative judgement revision on the part of individuals which leads to a consensus both on the weights and the values of the variables being judged. The question remains, however: Why should individuals revise their judgements by iterative linear averaging? According to Lehrer and Wagner, this method represents the uniquely rational way of combining dissenting judgements and that the consensual judgements so arrived at are the best summaries of the information contained in the group, including information that individuals hold about one another's judgemental competences. This being so the individuals who wish to make their judgements as accurate are rationally obliged to adopt the consensual values it generates.

What is it about linear averaging that justifies this claim? Here they appeal to the characterisation results presented in the previous section, to argue that it follows from acceptance of the conditions of zero-unanimity and independence of irrelevant alternatives as constraints on the aggregation of dissenting probabilities. This argument is incomplete however. In the first place, the relation between the independence and unanimity conditions and the idea of producing an optimal summary of information held by the group is not clear. If individuals only assign probability zero to propositions that they know to be false, then presumably zero-unanimity would be justified on informational grounds (but so too would something much stronger; namely that any zero judgement be preserved by the aggregate). But what is the informational rationale for the independence condition?⁸ Secondly it does not follow

⁸In this regard it is worth observing that the idea that the weight attached to the judgement of each individual by the aggregation function should depend on individual judgements regarding the authority of one another's judgements already implies that the aggregation function does *not* satisfy the independence of irrelevant alternatives condition across a domain including judgements about the authority.

from acceptance of weighted averaging that the weights used be measures of epistemic respect or that a consensual respect weights should also be arrived at by weighted averaging. Once again, the connection between the unanimity and independence conditions and information about judgemental competences needs to be established.

I think that the best way of assessing these claims to do so in the light of whatever generally accepted principles of rational revision of judgement that we can draw on. In the case of probabilistic judgements it is natural to turn to Bayesian updating principles governing rational belief revision in the face of new evidence. Bayesian updating rules with respect to some partition of possibility space are demonstrably valid when the revision inducing experience leaves the agent's conditional probabilities, given the elements of this partition, unchanged.⁹ Identifying the correct partition is crucial, of course, and no simple matter in practice, but that is a complication that we can abstract away from here.

Suppose that our n individuals have prior probabilities p_1, \dots, p_n , such that $p_i(X) = x_i$. On the Bayesian view of things i 's revised probability for some event X , $q_i(X)$, after having observed the prior probabilities of the other individuals, should equal her conditional probability for X , given the observed probability judgements, i.e.:

$$q_i(X) = p_i(X | p_1(X) = x_1, p_2(X) = x_2, \dots, p_n(X) = x_n)$$

What is the relation between this expression and the linear pooling formula? In the simplest case of just two individuals, i and j , the dual constraints of the Bayesian and linear pooling formulas yield:

$$p_i(X | p_j(X) = x_j) = wp_j(X) + (1 - w)p_i(X)$$

where w is i 's respect weight on j 's probability judgement on X . The relation expressed here is potentially very useful. In the form given above it 'tells' i how to form her conditional probabilities in the light of her respect for j 's judgements, thereby offering a solution the hard problem of how to determine conditional probabilities given the testimony of others. On the other hand, by reorganisation we can derive the respect weights from the posterior probabilities:

$$w = \frac{p_i(X | p_j(X) = x_j) - p_i(X)}{p_j(X) - p_i(X)}$$

allowing inference of someone's respect weights from their observed belief revision - potentially useful in empirical applications of this model of deliberation.

In case i 's conditional probabilities for X given j 's expressed probability for X equals x_j , j 's prior for X , her epistemic respect for j must be at the maximum of one. And vice versa. At the other extreme, if her conditional probabilities for X given j 's judgements just equals her prior for X , her epistemic respect for j is zero. Furthermore by Bayes' Theorem :

$$p_i(X | p_j(X) = x_j) = \frac{p_i(p_j(X) = x_j | X)}{p_i(p_j(X) = x_j)} p_i(X)$$

So $p_i(X | p_j(X) = x_j) = p_i(X)$ just in case $p_i(p_j(X) = x_j | X) = p_i(p_j(X) = x_j)$, i.e. whenever j 's probability judgements are independent of the truth. So it seems that a zero respect weight on someone else's probabilities coincides with the judgement that they are probabilistically independent of the truth; a useful result.

So far so good. However a couple of problems emerge even in this simple case.

1. Suppose that i and j have the same beliefs about X at a particular point in time, but that j is subsequently able to make an additional relevant observation. If j now declares his new

⁹ See Jeffrey [15] for a demonstration and discussion of this claim.

probabilities, how should i respond? Intuitively, and provided that i does not doubt j 's powers of observation, i should simply adopt j 's new probabilities as her own. But this is tantamount to zero weighting her own judgement, which in the light of the preceding judgement would seem to be equivalent to regarding her own judgements as probabilistically independent of the truth. But i may very well regard her judgements as perfectly good, even if not as well-informed as j 's. So intuitively adopting someone else's probabilities does not commit one to the view that one's own judgements are independent of the truth.

2. Suppose that $p_i(X|p_j(X) = x_j) < p_i(X)$ because, for instance, i regards j 's judgements as systematically biased in some way. In this case i 's respect weight for j should be *negative*. But this cannot be the case in the DLW model where weights are assumed to be non-negative.

A more serious difficulty for the reconciliation of Bayesian revision and linear pooling emerges in when we consider larger groups. Suppose we have three individuals i , j and k . Then Bayesian updating requires that:

$$\begin{aligned} q_i(X) &= p_i(X|p_j(X) = x_j, p_k(X) = x_k) \\ &= \frac{p_i(p_j(X) = x_j|p_k(X) = x_k, X) \cdot p_i(p_k(X) = x_k|X)}{p_i(p_j(X) = x_j|p_k(X) = x_k) \cdot p_i(p_k(X) = x_k)} p_i(X) \end{aligned}$$

The factor of interest here is the probabilistic independence or otherwise of the judgements of individuals j and k and its significance for i . In case they are independent we obtain:

$$q_i(X) = \frac{p_i(p_j(X) = x_j|X) \cdot p_i(p_k(X) = x_k|X)}{p_i(p_j(X) = x_j) \cdot p_i(p_k(X) = x_k)} p_i(X)$$

On the other hand when k and j 's judgement on X is perfectly correlated in the sense that $p_i(p_j(X) = x_j|p_k(X) = x_k, X) = p_i(p_j(X) = x_j|p_k(X) = x_k) = 1$, we obtain:

$$q_i(X) = \frac{p_i(p_k(X) = x_k|X)}{p_i(p_k(X) = x_k)} p_i(X)$$

Clearly i 's posterior probabilities for X will generally differ in these two cases; indeed they only agree when j 's judgement on X is independent of its truth. So it cannot be that on a Bayesian account someone's posterior probabilities, given the judgements of others, depends only on these judgements and the epistemic weight that they attach to them. Crucially the method of linear pooling ignores the interdependence of the expressed judgements.

To see how this can lead us astray, compare a situation in which two scientists conduct separate experiments to try and settle some question with one in which they conduct a single experiment together. Suppose that in both cases the scientists report that as a result of their experiments they consider X to be highly probable. In the former case, we would probably want to considerably raise our own probability for X because of the convergence of expert testimony. In the latter case too we would want to raise our probability for X , but less so, because their joint testimony in favour of X is based on same information. To revise once in the light of the testimony of the first scientist and then again in the light of that of the second would in effect be to update twice on the same evidence, akin to an individual scientist conditioning twice on the same experimental result. The DLW method does not, of course, directly counsel such double-counting and respect weights could in principle be assigned with considerations of dependence in mind. Rather its weakness lies in its failure to explicitly model these considerations. In the light of this it cannot claim to give a complete account of the optimal exploitation of the information held by a group.

4 Conclusion

Our investigation shows that if we think of deliberation as a process by which we learn from the testimony of others, then Bayesian principles of belief revision are not easily reconciled with iterative linear pooling. Nor does Bayesian thinking vindicate the claim that a consensus is rationally obligatory for, with the exception of the case where the deliberators are simply pooling information and each individual's posterior for the events reported on by others is either zero or one depending on what they report, there is no guarantee that the process of conditioning on the testimony of others will lead to a convergence of opinion. Assuming here that Bayesianism provides the basic normative standard from an epistemic point of view, it follows that whatever iterative linear pooling might have going for it, it cannot be said to be rationally obligatory on epistemic grounds alone. But it is not obvious that there are any other considerations that are generally applicable that can take up the slack and vindicate the claim that it is the uniquely rational response to a diversity of opinion.

What are the wider implications of this for the accommodation of deliberation by social choice theory? We cannot of course immediately conclude that introducing deliberative considerations will lead to a whole scale reconsideration of aggregation theory: our results have been illustrative rather than demonstrative of a tension between the two. But I do not think that the difficulties are confined to the DLW-model: all attempts to reconcile linear averaging and deliberative concerns will face similar ones. Suppose, for instance, that we face an allocation problem represented in the following table.

	AB	$A\neg B$	$\neg AB$	$\neg A\neg B$
Person 1	0.3	0.7	0	0
Person 2	0.3	0	0.7	0

Assume that Person 1 has observed that A is true, that Person 2 has observed that B is true, that Person 1 knows that 2 has observed whether A is true or not, and that Person 2 knows that 1 has observed whether B is true or not. In the light of what each knows about each other, it should follow from their receipt of the reports of the other's judgement that each will adopt the following consensual judgements:

	AB	$A\neg B$	$\neg AB$	$\neg A\neg B$
Consensual	1	0	0	0

These consensual values cannot, however, be represented as a weighted average of the two perspectives we started with, however we assign weights or interpret them. So linear averaging cannot be reconciled with any plausible theory of rational deliberation. But then either the unanimity or the independence conditions (or both) underlying linear averaging will be contravened in rational deliberation.

To summarise, I claim:

1. Many problems falling under the scope of aggregation theories are either allocation problems or closely connected to ones. Moreover, the principles invoked by allocation theory are typical of those found in aggregation theories generally.
2. The constraints which the theory of allocation places on consensual judgemental values are in tension with the output of rational deliberative processes initiated by the framing of an allocation problem, when deliberation is modelled as a process of learning from others compatible with Bayesian updating principles.
3. Hence, some of the principles invoked by allocation theory (and more generally aggregation theory) need to be either qualified or abandoned.

5 Appendix: Aggregation by Averaging

We now define the aggregation problem more generally than before. Let $I = \{1, \dots, i, \dots, n\}$ be a set of n individual perspectives, $V = \{V^1, \dots, V^j, \dots, V^m\}$ be a set of m variables, with $m \geq 3$. Let $\Gamma(m, n)$ be the set of all $m \times n$ matrices with rows in \mathfrak{R}^m and $\Gamma(m, n; S)$ be the subset of $\Gamma(m, n)$ consisting of all matrices with row sums equal to S . For any $A \in \Gamma(m, n)$ let A^j be its j th column and A_i its i th row.

An aggregation function ϕ , is a mapping from $\Gamma(m, n)$ to \mathfrak{R}^m . As before we denote $\phi(A)$ by $a = \langle a^1, \dots, a^m \rangle$. Let $S_A^i := \sum_j a_i^j$, the row total for perspective i , $S_A := \sum_j a^j$, the row total for the aggregate and $\bar{S}_A := \sum_i \frac{S_A^i}{n}$, the average row total. In addition to the axioms ZU, U and IIA which serve here as constraints on aggregation functions, we will refer to the following impartiality or anonymity axiom.

Axiom 7 *Anonymity (AN):* If σ is a permutation on I and $\forall(i \in I), a_i^j = b_{\sigma(i)}^j$ then $a^j = b^j$.

Axiom 8 *Zero-dominance (ZD):* $\forall i, a_i^j > 0 \implies a^j > 0$

Theorem 9 *Let A be any matrix in $\Gamma(m, n)$ such that $a_i^j > 0$. If ϕ satisfies IIA and ZU, then $\exists w_1 \dots w_n \in \mathfrak{R}$ such that for all j , $a^j = \sum_i w_i a_i^j$. Furthermore if ϕ satisfies ZD as well, then $w_i \geq 0$.*

Proof. Note that $S_A^i > 0$. Let B be the $m \times n$ matrix defined by $b_i^j = \frac{\bar{S}_A}{S_A^i} a_i^j$. Then $\forall i, \sum_j b_i^j = \sum_j \frac{\bar{S}_A}{S_A^i} a_i^j = \bar{S}_A$ and so by definition B is an allocation matrix. Now let $\bar{\phi}$ be a function on $\Gamma(m, n; \bar{S}_A)$ such that $\bar{b}^j := \frac{\bar{S}_A}{S_A^i} a^j$, where $\bar{\phi}(B) = \langle \bar{b}^1, \dots, \bar{b}^j, \dots, \bar{b}^m \rangle$. Note that $\sum_j \bar{b}^j = \sum_j \frac{\bar{S}_A}{S_A^i} a^j = \frac{\bar{S}_A}{S_A} \sum_j a^j = \bar{S}_A$ so that $\bar{\phi}$ is an allocation function. Clearly $\bar{\phi}$ respects IIA. Furthermore if $B^j = \langle 0, 0, \dots, 0 \rangle$ then $A^j = \langle 0, 0, \dots, 0 \rangle$ and hence, since ϕ respects ZU, $a^j = 0$. So by definition $\bar{b}^j = 0$ and, hence, $\bar{\phi}$ also respects ZU. It follows by Theorem 4 that $\exists w'_1, \dots, w'_n \in \mathfrak{R}$ such that for all i , $w'_i \geq 0$, $\sum_i w'_i = 1$ and for all j , $\bar{b}^j = \sum_i w'_i b_i^j$. Hence $\bar{b}^j = \sum_i w'_i \frac{\bar{S}_A}{S_A^i} a_i^j = \bar{S}_A \sum_i \frac{w'_i}{S_A^i} a_i^j$. So $\frac{\bar{S}_A}{S_A} a^j = \bar{S}_A \sum_i \frac{w'_i}{S_A^i} a_i^j$. Hence $a^j = \sum_i \frac{w'_i \cdot S_A}{S_A^i} a_i^j$. Define $w_i = \frac{S_A}{S_A^i} w'_i$. Then as required $w_1, \dots, w_n \in \mathfrak{R}$ are such that for all j , $a^j = \sum_i w_i a_i^j$. Suppose furthermore that ϕ satisfies ZD. Then since $a_i^j > 0$, it follows both that $S_A^i > 0$ and that $a^j > 0$ and hence that $S_A > 0$. So by definition $w_i \geq 0$. ■

Theorem 10 *Assume IIA and U and ZD. Then $\exists w_1 \dots w_n \in \mathfrak{R}$ such that $a_j = \sum_i w_i a_i^j$ and, if ϕ satisfies ZD as well, $w_i \geq 0$.*

Proof. Let A be any member of $\Gamma(m, n)$ and B be an $m \times n$ matrix such that $b_i^j = a_i^j + \beta$, where $\beta > |\min(a_i^j)|$. Then $b_i^j \geq 0$. Now let $\bar{\phi}$ be the aggregation function on $\Gamma(m, n)$ defined by $\bar{b}^j := a^j + \beta$, where $\bar{\phi}(B) = \langle \bar{b}^1, \dots, \bar{b}^j, \dots, \bar{b}^m \rangle$. Since ϕ satisfies IIA, so too does $\bar{\phi}$. Furthermore, note that if $B^j = \langle 0, 0, \dots, 0 \rangle$ then $A^j = \langle -\beta, -\beta, \dots, -\beta \rangle$. Hence by Unanimity, $a^j = -\beta$ and so $\bar{b}^j = 0$. It follows that $\bar{\phi}$ satisfies ZU. Hence by Theorem 9, $\exists w_1, \dots, w_n \in \mathfrak{R}$ such that $w_i \geq 0$ and for all j , $\bar{b}^j = \sum_i w_i b_i^j$. Hence $\bar{b}^j = \sum_i w_i (a_i^j + \beta) = \sum_i w_i a_i^j + \beta = a^j + \beta$. So $a^j = \sum_i w_i a_i^j$. By Theorem 9, furthermore, if ϕ satisfies ZD as well, then $w_i \geq 0$. ■

Corollary 11 *Assume IIA, U, ZD and AN. Then $\exists \beta > 0$ such that $\forall j, a^j = \beta \sum_i a_i^j$*

Proof. We prove the theorem by showing that for any $k, l \in I$ it is the case that $w_k = w_l$. Let $D \in \Gamma(m, n)$ be identical to A save for the fact that for some variable V^j such that $a_k^j \neq a_l^j$, it is the case that $d_k^j = a_l^j$ and $d_l^j = a_k^j$. Then by AN, $a^j = \sum_i w_i a_i^j = d^j = \sum_i w_i d_i^j$, where $w_i \geq 0$. But

$$\sum_i w_i a_i^j = \sum_i w_i d_i^j$$

$$\begin{aligned} \Leftrightarrow w_k a_j^k + w_l a_j^l &= w_k d_j^k + w_l d_j^l \\ \Leftrightarrow w_k a_j^k + w_l a_j^l &= w_k a_j^l + w_l a_j^k \\ \Leftrightarrow w_k &= w_l \end{aligned}$$

Hence $\sum_i w_i a_i^j = mw \sum_i a_i^j$ where w is the common weight on the individual perspectives. But by ZD, $w > 0$. Hence $\beta := mw > 0$. ■

References

- [1] D’Aspremont, C. and Gevers, L. (1977) “Equity and the Informational Basis of Collective Choice”, *Review of Economic Studies* 44: 199-209
- [2] Broome, J. (1990) “Bolker-Jeffrey Expected Utility Theory and Axiomatic Utilitarianism”, *Review of Economic Studies* 57: 477-502
- [3] Berger, R. L. (1981) “A Necessary and Sufficient Condition for Reaching Consensus by De Groot’s Method”, *J. Amer. Statis. Assoc.* 76: 415-418
- [4] Bohman, J. (1996) *Public Deliberation: Pluralism, Complexity and Democracy*, Cambridge, Mass: MIT Press
- [5] Bradley, R. (2005) “Axiomatic Bayesian Utilitarianism and Probability Homogeneity”, *Social Choice and Welfare* 24: 221-251
- [6] De Groot, M. H. (1974) “Reaching a Consensus”, *J. Amer. Statis. Assoc.* 69: 118-121
- [7] Dietrich, F. and List, C. “A Liberal Paradox for Judgement Aggregation”, *LSE Choice Group Working Papers* 1 (6)
- [8] Dietrich, F. and List, C. “Arrow’s Theorem in Judgement Aggregation”, *LSE Choice Group Working Papers* 1 (7)
- [9] Dryzek, J. S. (1990) *Discursive Democracy*, Cambridge: Cambridge University Press
- [10] Fishkin, J. S. (1991) *Democracy and Deliberation: New Directions for Democratic Reform*, New Haven, Conn: Yale University Press
- [11] Genest, C. and J. V. Zidek (1986) “Combining Probability Distributions: A Critique and Annotated Bibliography”, *Statistical Science* 1 (1): 113-135
- [12] Gevers, L. (1979) “On Interpersonal Comparability and Social Welfare Orderings”, *Econometrica* 47: 75-89
- [13] Goodin, R. E. (2001) “Consensus Interruptus”, *The Journal of Ethics* 5: 121-131
- [14] Jeffrey, R. C. (1983) *The Logic of Decision*, 2nd ed, Chicago, University of Chicago Press
- [15] Jeffrey, R. (1992) *Probability and the Art of Judgement*, Cambridge: Cambridge University Press
- [16] Laddager, R. (1985) “Leher and the Consensus Proposal”, *Synthese* 36: 473-477
- [17] Lehrer, K. (1976) “When Rational Disagreement is Impossible”, *Noûs* 10: 327-332

- [18] Lehrer, K. (1983) "Rationality as Weighted Average", *Synthese* 57: 283-295
- [19] Lehrer, K and C. Wagner (1981) *Rational Consensus in Science and Society*, Dordrecht: Reidel
- [20] Levi, I (1980) "*The Enterprise of Knowledge : An Essay on Knowledge, Credal Probability and Chance*", Cambridge, Mass: MIT Press
- [21] Levi, I. (1985) "Consensus as Shared Agreement and Outcome of Inquiry", *Synthese* 62: 79-95
- [22] Levi, I. (1985) "Pareto-unanimity and Consensus", *Journal of Philosophy* 87
- [23] List, C. and Pettit, P. (2002) "Aggregating Sets of Judgements: An Impossibility Result", *Economics and Philosophy* 18: 89-110
- [24] Loewer, B. and Laddager, R. (1985) "Destroying the Consensus", *Synthese* 62: 3-11
- [25] McConway (1981) "Marginalization and Linear Opinion Pools", *Journal of American Statistical Association* 76: 410-14
- [26] Mongin, P. (1995) "Consistent Bayesian Aggregation." *Journal of Economic Theory* 66 (2): 313-51
- [27] Mongin, P. (1995) "The Pareto Principle and Spurious Agreement" *LSE Choice Group Working Papers* 1 (4)
- [28] Roberts, K. (1980) "Interpersonal Comparability and Social Choice Theory", *Review of Economic Studies* 47: 421-439
- [29] Seidenfeld, T., Kadane, J.B., and Schervish (1989) "On the Shared Preferences of Two Bayesian Decision Makers", *Journal of Philosophy* 86: 225-244
- [30] Sen, A. (1970) "The Impossibility of a Paretian Liberal", *Journal of Political Economy* 78: 152-7
- [31] Sen, A. (1982) *Collective Choice and Social Welfare*, Edinburgh: Oliver and Boyd
- [32] Sunstein, Cass R. (2002) "Conformity and Dissent", U Chicago Law & Economics, Olin Working Paper No. 164; and U Chicago, Public Law Research Paper No. 34
- [33] Wagner, C. (1978) "Consensus through Respect: A Model of Rational Group Decision-making", *Philosophical Studies* 34: 335-349
- [34] Wagner, C. (1982) "Allocation, Lehrer Models, and the Consensus of Probabilities", *Theory and Decision* 14: 207-220
- [35] Wagner, C. (1985) "On the Formal Properties of Weighted Averaging as a Method of Aggregation", *Synthese* 62: 97-108