CONSENSUS BY AGGREGATION AND DELIBERATION

Richard Bradley

Department of Philosophy, Logic and Scientific Method
London School of Economics
r.bradley@lse.ac.uk

ABSTRACT: On the face of it both aggregation and deliberation represent alternative ways of producing a consensus. I argue, however, that the adequacy of aggregation mechanisms should be evaluated with an eye to the effects, both possible and actual, of public deliberation. Such an evaluation is undertaken by sketching a Bayesian model of deliberation as learning from others.

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. The perspectives might be those of individuals, social classes, times or states of the world. Any situation of this kind can be represented in

Printed from: *Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz.* Ed. T. Rønnow-Rasmussen, B. Petersson, J. Josefsson & D. Egonsson, 2007. www.fil.lu.se/hommageawlodek

a table as follows, with x^j_i denoting the value of perspective i on variable X^j:

Variables

Opinions	X^1	X^2	…	X^m
I_1	X^1_1	X^2_1	…	X^m_1
I_2	X^1_2	…	…	…
…	…	…	…	…
I_n	X^1_n	…	…	X^m_n

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?

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 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 especially prevalent.

1. Unanimity 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. Universal domain conditions which require that the aggregation method be applicable over all permitted profiles of individual judgements.

Seemingly natural conditions of these kinds can have very powerful consequences. To take one example, suppose that our table represents the probability judgements of individuals over an m-fold partition of possibilities. The aggregation theorist will ask: What kinds of mappings from such sets of individual probability judgements to a 'consensual' probability are there that satisfy the three above conditions of unanimity preservation, consistency and universal domain? Surprisingly, if the mapping is such that the consensual probability for some variable x^j_i depends only on the individual probabilities for x^j_i (the independence of alternatives assumption), then it must be one that assigns to any profile of individual probabilities for x^j_i, a consensual probability that is a weighted average of the individual ones.[1]

We will return to this example later on. For the moment, note that aggregation mechanisms produce a consensus on some question of common concern, but in doing so they leave individual opinion unchanged. The consensus that is produced is thus of a derivative kind; one that

[1] See Wagner (1982) for a formal statement and proof of this result.

depends on the existence of 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.[2] 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 inquiry 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 inquiry 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 inquiry aimed a 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

[2] To 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'.

precondition for genuine involvement in the activities. (Of course, people do enter into discussions, and even inquiries, 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).

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 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 inquiry 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 inquiry 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).

In this paper, I want to examine the implications for our view of an acceptable consensus of giving explicit attention to the impact of deliberation on the opinions of individuals. In the next section, I will argue 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. In the third section, I assess the prospects for such an accommodation, looking in particular at a model developed separately by Morris DeGroot and by Keith Lehrer and Carl Wagner. 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.

2. Deliberation

Deliberation as a means of securing consensus has many advocates, especially in democratic political theory. But the mechanisms by which deliberation produces 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 (2002).

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. 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 on 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.

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. This in turn raises the question: 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.

3. Modelling Deliberation

In this section I want to assess whether the more prevalent tenets of aggregation theory are robust with respect to the implication of rational deliberation. This is difficult to do in any kind of general way, since the conclusions we can draw will depend on what we take as our instances of

theories of aggregation or deliberation. What I propose to do therefore is 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 it's 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. And to compare the implications it has for acceptable consensual judgements with those imposed by aggregation theory.

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 (1974) and by Keith Lehrer and Carl Wagner (1981). 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 (2001) and the

papers appearing in Synthese vol. 52 (1985). The discussion here picks up mainly on the themes addressed in Loewer and Laddaga (1985).

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.

	Weight 1	Weight 2	Weight 3	Prior	Posterior
Individual 1	0.4	0.2	0.4	0.1	0.14
Individual 2	0.6	0.4	0	0.2	0.14
Individual 3	0.5	0.4	0.2	0.15	0.14

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:

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

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.[3]

The DLW model is especially useful for this discussion because it provides an instance of a positive relationship between a theory of aggregation and a model of deliberation. For the process of iterative linear pooling leads to consensual probabilities that are weighted averages of individual ones. And we saw earlier on that this is precisely the form consensual probabilities must have if they formed by an aggregation mechanism that satisfies the universal domain and independence of alternatives condition and which preserves unanimous judgements. Furthermore the deliberation model tells us what weights should be used for the averaging: they are consensual judgements on the respect due to each individual's judgement.

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. I think that the best way of assessing these claims is to do so in the light of whatever generally accepted principles of rational revision of judgement that we can

[3] See Lehrer and Wagner (1981) for a proof of this claim.

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.[4] 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 n individuals have prior probabilities $p_1,...,p_n$, 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.:

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

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

[4] See Jeffrey (1992) for a demonstration and discussion of this claim.

determine conditional probabilities given the testimony of others. On the other hand, by reorganisation we can derive the respect weights from the posterior probabilities:

$$w = [(p_i(X|p_j(X) = x_j) - p_i(X)]/[p_j(X) - p_i(X)]$$

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 :

$$p_i(X|p_j(X) = x_j) = [(p_i(p_j(X) = x_j|X).p_i(X)]/p_i(p_j(X) = x_j)$$

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 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:

$$q_i(X) = p_i(X|p_j(X) = x_j, p_k(X) = x_k)$$

The factor of interest here is the probabilistic independence or otherwise of the judgements of individuals j and k and its significance for i. Compare, for instance the case in which in i's opinion, j and k's judgements on X are completely independent with the case in which they are completely dependent. If i is a Bayesian her posterior probabilities for X will agree in these two cases only 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 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¬B	¬AB	¬A¬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¬B	¬AB	¬A¬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 that the constraints which aggregation theory typically 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. Hence, some of the principles invoked by aggregation theory need to be either qualified or abandoned.

References.

De Groot, M. H. (1974) "Reaching a Consensus", J. Amer. Statis. Assoc. 69: 118-121

Genest, C. and J. V. Zidek (1986) "Combining Probability Distributions: A
Goodin, R. E. (2001) "Consensus Interruptus", The Journal of Ethics 5: 121-131

Jeffrey, R. C. (1983) The Logic of Decision, 2nd ed, Chicago, University of Chicago Press

Jeffrey, R. (1992) Probability and the Art of Judgement, Cambridge: Cambridge University Press

Laddager, R. (1985) "Leher and the Consensus Proposal", Synthese 36: 473-477

Lehrer, K. (1976) "When Rational Disagreement is Impossible", Noûs 10: 327-332

Lehrer, K. (1983) "Rationality as Weighted Average", Synthese 57: 283-295

Lehrer, K and C. Wagner (1981) Rational Consensus in Science and Society, Dordrecht: Reidel

Levi, I. (1985) "Pareto-unanimity and Consensus", Journal of Philosophy 87

Loewer, B. and Laddager, R. (1985) "Destroying the Consensus", Synthese 62: 3-11

Mongin, P. (1995) "The Pareto Principle and Spurious Agreement", LSE Choice Group Working Papers 1 (4)

Sen, A. (1982) Collective Choice and Social Welfare, Edinburgh: Oliver and Boyd

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

Wagner, C. (1982) "Allocation, Lehrer Models, and the Consensus of Probabilities", Theory and Decision 14: 207-220

Wagner, C. (1985) "On the Formal Properties of Weighted Averaging as a Method of Aggregation", Synthese 62: 97-108

26 August 2006