Universality and its Discontents   (in press, The Balkan Journal of Philosophy)

Carl G. Wagner

Department of Mathematics

The University of Tennessee

Knoxville, Tennessee 37996-1320, USA

Telephone: 865-974-4305

Fax: 865-974-6576

Email: cwagner@tennessee.edu

Universality and its Discontents

Carl G. Wagner

Abstract. In framing the concept of rational consensus, decision theorists have tended to defer to an older, established literature on social welfare theory for guidance on how to proceed. But the uncritical adoption of standards meant to regulate the reconciliation of differing *interests* has unduly burdened the development of rational methods for the synthesis of differing *judgments*. In particular, the universality conditions typically postulated in social welfare theory, which derive from fundamentally ethical considerations, preclude a sensitive treatment of special cases when carried over to the realm of judgment aggregation, especially in the case of probabilistic judgment.

Key words: consensus, probability pooling, universality axiom, weighted averaging.

1. Introduction. The thirtieth anniversary of the publication of *Rational Consensus in Science and Society* (Lehrer and Wagner 1981) offers an opportunity to survey and assess certain developments in consensus studies during the past three decades, and I am grateful to the editors of the *Balkan Journal of Philosophy* for their invitation to contribute to this special volume on communication and rational consensus. Since my own work in this area has dealt chiefly with the consensus problem for probabilistic judgments, the theorems and examples that I cite in what follows will be drawn from that area of expert judgment synthesis. But I hope to make clear that the lessons drawn from this survey admit of more general application.

In framing the concept of rational consensus, decision theorists have tended to defer to an older, established literature on social welfare theory for guidance on how to proceed. But the uncritical adoption of standards meant to regulate the reconciliation of differing *interests* has unduly burdened the development of rational methods for the synthesis of differing *judgments*. In particular, the universality conditions typically postulated in social welfare theory, which derive from fundamentally ethical considerations, preclude a sensitive treatment of special cases when carried over to the realm of judgment aggregation, especially in the case of probabilistic judgment. For requiring of a method that it produce a rational consensual probability distribution for every logically possible "profile" of individual distributions more or less forces the consensual distribution to be constructed state-by-state (or event-by-event). And, as I show, state-by-state approaches leave no room for the sort of holistic strategies required to achieve clearly desirable outcomes such as the preservation of epistemically significant cases of independence common to all individuals' probability distributions. More strikingly, suppose that the set of possible probabilities is finite (as it is in all practical cases) and the consensual probability distribution is constructed state-by-state. Then, as I have recently proved, only a dictatorship can accommodate the simple demand that if everyone assigns a state zero

probability, so does the consensual distribution. Such results show that a seemingly innocuous desire for generality of application can render the whole notion of rational consensus otiose.

2. Consensus through respect. Suppose that, after extensive discussion and the sharing of evidence, n experts have attained a state of what Lehrer has termed *dialectical equilibrium* regarding the appropriate probability distribution over some countable set Ω of possible states of the world. In this state no one sees any point to further discussion, yet the distributions $p_1, p_2, ..., p_n$ endorsed by these individuals fail to exhibit consensus. The special case of this scenario in which individuals regard each other as equals on "intelligence, perspicacity, honesty, thoroughness, and other epistemic virtues" (Gutting 1982) has, under the rubric of *the epistemic peer problem*, recently attracted the attention of a number of epistemologists and decision theorists. (Kelly 2005) has argued that individuals in this situation have no reason to further modify their probability assessments, while (Elga 2007), (Christensen 2007), and (Feldman 2007) have all advocated that individuals further revise their assessments to a single distribution that in some sense gives equal weight to each of those assessments. Naturally, this *equal weight view* requires elaboration, and may well admit of more than one reasonable interpretation.[1] As early as fifty years ago, however, (Stone 1961) suggested the weighted arithmetic mean $p = w_1 p_1 + w_2 p_2 + \cdots + w_n p_n$[2], where the weights w_i are nonnegative real numbers summing to 1, as a reasonable consensual distribution. In the case of epistemic peers one would presumably set each $w_i = 1/n$. More generally, weights would be chosen to reflect the relative expertise of the individuals.

(DeGroot 1974) and, independently, (Lehrer 1976) proposed a novel method for selecting such weights: Let each individual assign a nonnegative weight to each individual in the group, including himself, with weights assigned by an individual summing to 1. Enter these weights in an n×n matrix $W = (w_{i,j})$, where $w_{i,j}$ is the weight assigned by individual i to individual j,[3] and reflects, in some sense, i's evaluation of the expertise of j relative to other members of the group. In contributing the weights $w_{i,1}, w_{i,2}, ..., w_{i,n}$ to the i^{th} row of W, individual i commits to revising his initial distribution p_i to $p_i^{(1)} = w_{i,1} p_1 + w_{i,2} p_2 + \cdots + w_{i,n} p_n$. The revised profile $P^{(1)} = (p_1^{(1)}, p_2^{(1)}, ..., p_n^{(1)})^t$ is thus just the matrix product WP, where P denotes the $n×1$ column vector $(p_1, ..., p_n)^t$. If the revised distributions in $P^{(1)}$ fail to be identical, DeGroot and Lehrer envision the further revision of $P^{(1)}$ to $P^{(2)} = WP^{(1)} = W^2 P$ and, if necessary, of $P^{(2)}$ to $P^{(3)} = WP^{(2)} = W^3 P$, etc.

Theorems about the convergence of powers of such matrices to matrices with identical rows were invoked by DeGroot and Lehrer to identify a single sequence of *consensual weights* to employ in averaging the probability distributions $p_1, p_2, ..., p_n$. Sufficient conditions for such convergence can be cast in terms of patterns of respect communicated among the individuals.

Following Lehrer, let us say that individual i *respects* individual j if $w_{i,j} > 0$. There is a *chain of respect* from i to j if there exist individuals $i_1 = i, i_2, ..., i_k = j$ such that i_t respects i_{t+1} for $t = 1, ..., k-1$.

Theorem 2.1. Let W be a weight matrix and let E denote the set of all individuals k for which there exists a chain of respect from every other individual to k (hence, from k to k). If E is nonempty and at least one member of E respects himself, then powers of W converge to a weight matrix L with identical rows. The entries $\lambda_1, \lambda_2, ..., \lambda_n$ comprising each row of L are the unique solution to the simultaneous equations

(2.1) $$(x_1, x_2, ..., x_n)W = (x_1, x_2, ...x_n)$$

(2.2) $$x_1 + x_2 + \cdots + x_n = 1,$$

and individual i receives positive consensual weight ($\lambda_i > 0$) if and only if i belongs to E. Moreover, $\lambda_1 = \lambda_2 = \cdots = \lambda_n = 1/n$ if and only if the entries in each column of W sum to 1.

Proof. (Doob 1953). □

A more thorough discussion of the above theorem may be found in (Lehrer and Wagner 1981, Chapter 7). In particular, readers will find there an argument for adopting the unique solution to (2.1) and (2.2), whenever it exists, as the consensual weights for averaging individuals' probability distributions, even if powers of W fail to converge.

3. Weighted arithmetic means characterized. While weighted arithmetic means have the virtue of familiarity and simplicity, they constitute only one class of an extensive family of possible averaging functions. Before ruling out other possibilities, it is incumbent upon us as decision makers to know precisely what we are opting for when we opt for weighted arithmetic means, and that means identifying a set of properties that characterize such means. The first such characterization was given in (Aczél and Wagner 1980), with the following elaboration appearing in (Lehrer and Wagner 1981): Let Ω denote a countable set of possible states of the world, assumed to be mutually exclusive and exhaustive. A function $p : \Omega \rightarrow [0,1]$ is a *probability distribution* (or *probability mass function*) on Ω if and only if $\sum_{\omega \in \Omega} p(\omega) = 1$. Each probability distribution p gives rise to a *probability measure* (which, abusing notation, we also denote by p), defined for each $E \subseteq \Omega$ by $p(E) = \sum_{\omega \in E} p(\omega)$.

If n is a positive integer, we call a sequence $(p_1, ..., p_n)$ of probability distributions on Ω a *profile*. A *pooling operator T* supplies a method of reconciling the possibly differing distributions $p_1, ..., p_n$ in the form of a single, "consensual" distribution $p = T(p_1, ..., p_n)$. Of course, many pooling operators are prima facie unacceptable as a method of producing a rational

consensus, among them the *dictatorial operators* (for some fixed d, and all profiles $(p_1,..., p_n)$, $T(p_1,..., p_n) = p_d$) and the *constant operators* (which impose a fixed distribution q as the consensual distribution for every profile). There is an extensive literature on probability pooling. See (Genest and Zidek 1986) for a summary and appraisal of work done through the mid-1980s), which has tended to mimic the older literature on social welfare theory as developed, for example, by (Black 1948) and (Arrow 1951). Like social welfare theory, pooling theories lay down certain axiomatic constraints, and then try to identify the pooling operators that satisfy those constraints. Typical constraints have included:

Universal Domain (UD). The domain of the pooling operator T consists of all logically possible profiles $(p_1,..., p_n)$. That is, if Δ denotes the set of all probability distributions on Ω and Δ^n its n-fold Cartesian product, then $T : \Delta^n \to \Delta$.

State-wise Aggregation[4] (SA). For each $\omega \in \Omega$, there exists a function $f_\omega : [0,1]^n \to [0,1]$ such that for all $(p_1,..., p_n) \in \Delta^n$, $T(p_1,..., p_n)(\omega) = f_\omega(p_1(\omega),..., p_n(\omega))$.

Zero Preservation (ZP). For each $\omega \in \Omega$ and for all $(p_1,..., p_n) \in \Delta^n$, if $p_1(\omega) = \cdots = p_n(\omega) = 0$, then $T(p_1,..., p_n)(\omega) = 0$.

Theorem 3.1. If $|\Omega| \geq 3$, a pooling operator satisfies UD, SA, and ZP if and only if there exists a sequence $(w_1,..., w_n)$ of nonnegative real numbers summing to 1 such that for all $\omega \in \Omega$ and all $(p_1,..., p_n) \in \Delta^n$, $T(p_1,..., p_n)(\omega) = w_1 p_1(\omega) + \cdots + w_n p_n(\omega)$.

Proof. (Lehrer and Wagner 1981) or (Wagner 1982). ☐

Remark 3.1. When $|\Omega| = 2$, any pooling operator satisfying UD necessarily satisfies SA, and there is an extensive class of pooling operators satisfying UD, SA, and ZP (Lehrer and Wagner 1981, p. 110).

Remark 3.2. Note that while condition SA allows for the use of different functions f_ω to aggregate the individual probability assessments $p_1(\omega),..., p_n(\omega)$, this apparent flexibility turns out to be illusory when combined with UD and ZP. Indeed, all of the functions f_ω turn out to be identically equal to a single weighted arithmetic mean.

Remark 3.3. (McConway 1981) has proved a result similar in spirit to that of Theorem 3.1 involving a potentially infinite family of pooling procedures for probability measures on various sigma algebras on an arbitrary set Ω of possible states of the world. In McConway's approach, condition SA amounts to stipulating that pooling commute with marginalization (i.e., with the restriction of probability measures on a given sigma algebra to some sub-sigma algebra).

It is perhaps surprising that axioms UD, SA, and ZP so severely circumscribe the set of acceptable pooling operators. Even more striking is the fact that when the allowable values of the various probability assessments $p_i(\omega)$ and the consensual assessments $p(\omega)$ are restricted to a fixed finite subset of the real number interval $[0,1],$ then only dictatorial pooling satisfies these three axioms.

Theorem 3.2. Let V be a finite subset of the real number interval $[0,1]$ satisfying the conditions (i) $0 \in V$, (ii) $x \in V \Rightarrow 1 - x \in V$, and (iii) $x, y \in V$ and $x + y \leq 1 \Rightarrow x + y \in V$. Suppose that Ω is countable, and $|\Omega| \geq 3$. Let Δ_V denote the set of all probability distributions p on Ω such that $p(\omega) \in V$ for all $\omega \in \Omega$. A pooling operator $T : \Delta_V^n \to \Delta_V$ satisfies UD, SA, and ZP if and only if it is dictatorial.

Proof. See http://philsci-archive.pitt.edu/5196 □

4. Independence preservation. As usual, events E and $F \subseteq \Omega$ are *independent with respect to the probability measure* p if $p(E \cap F) = p(E)p(F)$. Several individuals (Laddaga 1977), (Laddaga and Loewer 1985), (Schmitt 1985) have advocated the following additional axiomatic restriction on probability pooling:

Universal Independence Preservation (UIP). For all $(p_1,...,p_n) \in \Delta^n$ and for all subsets E and F of Ω, if $p_i(E \cap F) = p_i(E)p_i(F)$ for $i = 1,...,n,$ then
$T(p_1,...,p_n)(E \cap F) = T(p_1,...,p_n)(E)T(p_1,...,p_n)(F).$

It is clear that pooling by weighted arithmetic means will often violate UIP. Indeed, only the most extreme forms of weighted arithmetic averaging satisfy UIP.

Theorem 4.1. If $|\Omega| \geq 3$, then a pooling operator satisfies UD, SA, ZP, and UIP if and only if it is dictatorial.

Proof. (Lehrer and Wagner 1983). □

Remark 4.1. If ZP is deleted from the hypotheses of the above theorem, the set of acceptable pooling operators is not enlarged in any meaningful way, consisting merely of imposed, as well as dictatorial, operators (Wagner 1984).

SA is of course a very strong condition, for it requires that the consensual probability assigned to each state ω depends only, via the function f_ω, on the probabilities assigned by individuals to that state. It is not surprising then, that it is difficult to satisfy both SA and UIP, since SA rules out the kind of holistic approach that would seem to be necessary to preserve independence of events, which commonly comprise more than one state. In addition, SA requires that $\sum_\omega f_\omega(p_1(\omega),..., p_n(\omega)) = 1,$ without any recourse to normalization. Indeed, it is this feature of

SA that is the driving force behind Theorem 3.1. It is thus natural to wonder if weakening SA to allow for subsequent normalization might accommodate UIP in some interesting ways. For this would open up the possibility of employing any of an extensive variety of weighted averaging functions to play the role of the functions f_ω.[5] In fact, this gambit does work when $|\Omega| \leq 4$, as shown in (Sundberg and Wagner 1987).[6] But when $|\Omega| \geq 5$, only dictatorial pooling accommodates UD, normalized SA, and UIP (Genest and Wagner 1987).

One response to Theorem 4.1 would be to take issue with the reasonableness of UIP, which demands of a pooling operator that it preserve *every single instance* of independence common to the distributions of the n experts. As pointed out in (Lehrer and Wagner 1983) and (Genest and Wagner 1987) there are many instances of such independence having no epistemic significance whatsoever. Suppose, for example, that you regard a certain die as fair, but I think that the die is weighted so that the probabilities of $1, 5,$ and 6 are each $1/6,$ the probabilities of 2 and 4 are each $1/12,$ and the probability of 3 is $1/3.$ On each of our probability assessments, the events 'die comes up even,' and 'die comes up a multiple of 3' are independent. But this independence is an incidental feature of our distributions to which neither of us is likely to have any theoretical commitment. Why, then, should we take care to preserve this independence in any consensual distribution that emerges from aggregating our individual distributions? On the other hand, agreed-upon cases of physical independence of certain random variables or of the independence of certain families of partitions of Ω exhibit the sort of epistemic significance that demands preservation under probability pooling. It should be noted that there are principled methods of preserving this sort of independence, both for a fixed, finite family of partitions of Ω, and for a fixed, finite set of discrete random variables (Wagner 2011). But these methods yield pooling operators that apply to profiles of probability distributions belonging to the restricted subset of Δ consisting of just those distributions with respect to which the family of partitions (respectively, set of random variables) is independent. And there is no apparent way to extend the domain of the aforementioned methods in such a way that UD is satisfied. As I argue below, however, this is no cause for concern.

5. Conclusion: Hard Cases Need Not Make Bad Law. Hard cases make bad law. This maxim, familiar to all students of the law, asserts that exceptional legal cases are not suitable as the source for generalized laws. In such cases, jurors are often left with the unpalatable alternatives of failing to do justice or creating undesirable precedents by carving out an exception to a general rule in order to accommodate morally compelling claims. The dictatorship theorems described above might seem to indicate that a similar principle operates in the realm of rational consensus theory. But any such analogy is superficial, for the law must of necessity be prepared to issue judgments in all cases, and in a way that aims for maximal consistency. But why should we subject ourselves to the straightjacket of universality conditions when devising methods of creating a rational consensus from the possibly differing opinions of a set of experts? As noted in the above section, one can do justice, so to speak, to the "hard" cases of epistemically

significant, agreed-upon independence. Yet only dictatorial (or imposed) pooling accommodates the demand for *universal* independence preservation. Similarly, demanding of a pooling method that it produce a consensual distribution for *every logically possible profile* of probability distributions (condition UD) leads almost inevitably to the adoption of pooling operators constructed in accord with state-wise aggregation (SA), or a normalized version thereof. For how can one specify a consensual distribution for every such profile, by means of a necessarily finite set of instructions, without defining that distribution state-by-state ? And we have seen how severely SA, and even normalized SA, circumscribe the set of methods available for achieving a synthesis of expert judgments, precluding the sort of holistic approach one would like to bring to bear on certain special classes of consensus problems. In the realm of rational consensus, it is not hard cases, but the demand to treat *too many cases*, that yields bad (decision-theoretic) law. That so many consensus theorists (including the present author, from whom this essay is a *mea culpa*) failed to take account of this simple observation for so long illustrates the power of faulty analogy in particularly striking form.

Footnotes

1. An analysis of an intriguing geometric interpretation of the equal weight view suggested to the author by David Jehle appears in (Shattuck and Wagner 2010).

2. That is, for each possible state of the world $\omega \in \Omega$, $p(\omega) = w_1 p_1(\omega) + w_2 p_2(\omega) + \cdots + w_n p_n(\omega)$. It is easy to check that for each ω, $\min_i \{ p_i(\omega) \} \le p(\omega) \le \max_i \{ p_i(\omega) \}$, and that $\sum_{\omega \in \Omega} p(\omega) = 1$. Stone called this way of aggregating the distributions in question *the opinion pool*.

3. Such square matrices, with nonnegative entries summing to 1 in each row, are known as *stochastic matrices*, and have been extensively studied by mathematicians, since they define a stochastic process known as a *finite Markov chain*.

4. State-wise aggregation is often termed *irrelevance of alternatives* or (confusingly) *independence*.

5. Among the possibilities are all of the so-called *weighted quasi-arithmetic means*, constructed by choosing a strictly monotone function τ, along with nonnegative weights $w_1,...,w_n$ that sum to 1, and setting $f_\omega(p_1(\omega),..., p_n(\omega)) = \tau^{-1}(w_1 \tau(p_1(\omega)) + \cdots + w_n \tau(p_n(w)))$. For example, when $\tau(x) = x^2$, then f_ω is a *weighted root-mean-square*; when $\tau(x) = x^{-1}$, then f_ω is a *weighted harmonic mean*; and when $\tau(x) = \ln x$, then f_ω is a *weighted geometric mean*.

6. In particular, normalized weighted geometric means (see note 5 *supra*) satisfy UD, normalized SA, and UIP when $|\Omega| = 4$.

References

Aczél J, Wagner C (1980) A characterization of weighted arithmetic means. SIAM J. Alg. Discr. Meth. 1, 259-260

Arrow K (1951) Social choice and individual values. Cowles Commission Monograph 12, Wiley, New York

Black D (1948) On the rationale of group decision-making. J. Polit. Econ. 56, 23-34

Christensen D (2007) Epistemology of disagreement: the good news. Philos. Rev. 116, 187-217

DeGroot M (1974) Reaching a consensus. J. Amer. Statist. Assoc. 69, 118-121

Doob J (1953) Stochastic processes. Wiley, New York

Elga A (2007) Reflection and disagreement. Noûs 41, 478-502

Feldman R (2007) Reasonable religious disagreements. In: Anthony L (ed) Philosophers without gods, Oxford University Press, Oxford. pp 194-214

Genest C, Wagner C (1987) Further evidence against independence preservation in expert judgment synthesis. Aequationes Math. 32, 74-86

Genest C, Zidek J (1986) Combining probability distributions: a critique and annotated bibliography. Statist. Sci. 1, 114-148

Gutting G (1982) Religious belief and religious skepticism. University of Notre Dame Press, Notre Dame

Kelly T (2005) Peer disagreement and higher order evidence. In: Feldman R, Warfield T (eds) Disagreement, Oxford University Press, Oxford, pp 111-174

Laddaga R (1977) Lehrer and the consensus proposal. Synthese 36, 473-477

Laddaga R, Loewer B (1985) Destroying the consensus. Synthese 62, 79-96

Lehrer K (1976) When rational agreement is impossible. Noûs 10, 327-332

Lehrer K, Wagner C (1981) Rational consensus in science and society. Reidel, Dordrecht

Lehrer K, Wagner C (1983) Probability amalgamation and the independence issue: a reply to Laddaga. Synthese 55, 339-346

McConway K (1981) Marginalisation and linear opinion pools. J. Amer. Statist. Assoc. 76, 410-414

Schmitt F (1985) Consensus, respect, and weighted averaging. Synthese 62, 25-46

Shattuck M, Wagner C (2010) Independence preservation in expert judgment synthesis. Aequationes Math. 80: 277-290

Stone M (1961) The opinion pool. Ann. Math. Statist. 32, 1339-1342

Sundberg C, Wagner C (1987) A functional equation arising in multi-agent statistical decision theory. Aequationes Math. 32: 32-37

Wagner C (1982) Allocation, Lehrer models, and the consensus of probabilities. Theory and Decision 14, 207-220

Wagner C (1984) Aggregating subjective probabilities: some limitative theorems. Notre Dame J. Formal Logic 25, 233-240

Wagner C (2010) Peer disagreement and independence preservation. Erkenntnis 74: 277-288