Evaluation and development of formal consensus methods; ACERA Project No 0706

Mark Colyvan; University of Sydney

Progress Report; July 2007

Acknowledgements

This report is a product of the Australian Centre of Excellence for Risk Analysis (ACERA). In preparing this report, the authors acknowledge the financial and other support provided by the Department of Agriculture, Fisheries and Forestry (DAFF), the University of Melbourne, Australian Mathematical Sciences Institute (AMSI) and Australian Research Centre for Urban Ecology (ARCUE).

Disclaimer

This report has been prepared by consultants for the Australian Centre of Excellence for Risk Analysis (ACERA) and the views expressed do not necessarily reflect those of ACERA. ACERA cannot guarantee the accuracy of the report, and does not accept liability for any loss or damage incurred as a result of relying on its accuracy.

Table of contents

1. Activities to date

Conferences and Workshops:

- NCEAS Working group (including Mark Burgman, Helen Regan and Mark Colyvan) Jan–Feb 2007 in Santa Barbara
- AEDA Conservation Planning Workshop (Katie Steele attended), 8–11 May in Brisbane
- ACERA Robust Multi-Criteria Decision Analysis Workshop (Helen Regan, Mark Burgman and Katie Steele attending), 12-17 Aug 2007 in Hobart.
- ACERA Workshop on Formal Consensus Methods planned for Aug 21–24 in Sydney (project members plus a couple of invited experts on formal consensus), including an afternoon session 2–5pm on 24th August for selected DAFF invitees.

Relevant Publications:

- Colyvan, M. 'Relative Expectation Theory', Journal of Philosophy, forthcoming.
- Colyvan, M. 'Population Ecology', chapter in S. Sarkar and A. Plutynski (eds.), A Companion to the Philosophy of Biology, Blackwell, 2007.
- Colyvan, M. 'No Expectations', Mind, Vol. 115, No. 459, (July 2006), pp. 695–702.
- Colyvan, M. and Regan, H.M. 'Legal Decisions and the Reference-Class Problem', International Journal of Evidence and Proof, Volume 11, 2007.
- Regan. H.M., Colyvan, M. and Markovchick-Nicholls, L. 'A Formal Model for Consensus and Negotiation in Environmental Management', Journal of Environmental Management, Vol. 80, No. 2 (July 2006), pp. 167–76.
- Steele, K. 'Distinguishing Indeterminate Belief from "Risk-averse" Preferences', Synthese, forthcoming.
- Steele, K. 'The Precautionary Principle: A New Approach to Public Decision-Making?', Law, Probability and Risk, forthcoming.
- Steele, K. Regan, H.M. and Colyvan, M. 'Right Decisions or Happy Decision Makers?', Social Epistemology, forthcoming.
- Regan, H.M., Davis, F.W., Andelman, S.J., Widyanata, A. and Freese, M. 'Comprehensive criteria for biodiversity evaluation in conservation planning', Biodiversity and Conservation, forthcoming.

Work in Progress:

- Colyvan, M. 'What's Maths Got to Do with It?'
- Colyvan, M., Cox, D. and Steele, K. 'Modelling the Moral Dimension of Decisions', under consideration at Nous.
- Colyvan, M. Lubell, M., Regan, H.M. and Skyrms, B. 'The Conservation Game: Survey of Game-theoretic Methods in Conservation-management Applications'.
- Colyvan, M. and Steele, K. 'Ecology and Environmental Decisions', invited contribution to B. Brown, K. de Laplante, and K. Peacock (eds.), Handbook of the Philosophy of Science, Volume 11: Philosophy of Ecology, North Holland/Elsevier, forthcoming.
- Linquist, S., Colyvan, M., Grey, W., Griffiths, P. Odenbaugh, J. and Possingham, H. 'Current Trends and Future Directions in the Philosophy of Ecology'.
- Steele, K., Colyvan, M., Regan, H.M. and Burgman, M. 'A Survey of Formal Consensus Methods for Conservation Management'.

- Steele, K. review of Husain Sarkar's 'Group Rationality in Scientific Research', invited by the Notre Dame Philosophical Reviews.
- Steele, K. 'Limitations of the decision-theoretic model of desire', under consideration at Erkenntnis.
- Steele, K. 'Pareto-optimality and the distinction between deliberation and aggregation in group choice'.

Mark Colyvan's Relevant Presentations:

- 7 May 2007: 'Formal Models of Consensus', presented in the workshop on "Methodological Problems in the Social Sciences" in the Tilburg Center for Logic and Philosophy of Science, Tilburg University, Tilburg, Netherlands.
- 3 May 2007: 'Probability and Law', presented to the Departments of Theoretical and Practical Philosophy at Lund University, Lund, Sweden.
- 23 April 2007: 'Relative Expectation Theory', presented to the Department of Philosophy at the Ludwig-Maximilians University, Munich, Germany.
- 16 December 2006: 'The Use of Mathematics to Describe Biological Systems', presented at the 3rd Queensland Biohumanities Conference, University of Queensland, Brisbane, Australia.
- 9 October 2006: 'Do the Laws of Ecology Lie?' presented in the History and Philosophy of Science seminar series at the University of Sydney, Sydney, Australia.
- 25 August 2006: 'Probability in Law' presented to the TC Beirne School of Law at the University of Queensland, Brisbane, Australia.
- 18 July 2006: 'Risk and the Limitations of (Classical) Probability Theory' presented at the Australian Society for Risk Analysis Conference at the University of Melbourne, Melbourne, Australia.
- 2 July 2006: 'What's Maths Got to Do with It?', Presidential Address at the 2006 Australasian Association of Philosophy Conference at the Australian National University, Canberra, Australia.
- 29 June 2006: 'Mathematical Models in Ecology and Conservation Biology', presented at the 2nd Queensland Biohumanities Conference, University of Queensland, Brisbane, Australia.

Helen Regan's Relevant Presentations:

- March 2006: 'A formal model for consensus and negotiation in land-use planning', presented at the 21st Annual Symposium of the United States Regional Chapter of the International Association for Landscape Ecology, San Diego, USA.

Katie Steele's Relevant Presentations:

- 9 May 2007: 'Group decision models: Balancing truth and fairness', presented at the AEDA Conservation Planning Workshop, Bardon Conference Centre, Brisbane, Australia.
- 22 November 2006: Comments on Philip Pettit's 'Rationality, reasoning and regulation: the case of group agents', Centre for Time "Minds, Mobs and Memories" Conference, University of Sydney.
- 26 May 2006: 'What is it Rational to Value?', presented at the Formal Epistemology Workshop, UC Berkeley.
- 15 April 2006: 'Modelling the Moral Dimension of Decisions', 8th Annual CMU/Pitt Graduate Student Philosophy Conference, Pittsburgh.

2. Summary of main findings

The project members have co-authored a paper on formal consensus methods entitled "Right Decisions or Happy Decision Makers?" (Steele *et al.* (forthcoming)), listed above. Another paper connected to this project is Regan *et al.* (2006), also listed above. Both papers focus on the merits of the Lehrer-Wagner method for achieving consensus.[1] This method is intended to produce a group opinion, given the independent opinions of members, where both the group and member opinions are expressed as arrays of (positive) numerical estimates that sum to some value (we concentrate on probability distributions). A brief description of the Lehrer-Wagner method is as follows: after considerable group discussion permitting group members to share evidence and state reasons for their opinions, they each nominate a probability distribution and also provide weightings of respect for all other members of the group. If these initial probability distributions differ, each member computes a revised distribution that is the weighted average (based on their personal weightings of respect) of the distributions of all members in the group. If disagreement remains after the new probability distributions have been calculated, each person computes a new weighted average based again on their personal weightings of respect compounded with the revised distributions of all group members. The process continues until all group members converge upon a common probability distribution. Convergence is guaranteed under some reasonable assumptions about the respect weightings. In addition to the convergence result, the algorithm arguably accommodates all the relevant information available (i.e. the initial probability distributions of all group members, as well as their weightings of respect for other group members).

In our analysis of the Lehrer-Wagner formal method of consensus, we concentrated on the distinction between consensus and compromise in group decision-making. This distinction is not well attended to in the literature on formal consensus methods. Indeed, some algorithms discussed in this literature aim for consensus, while others are best interpreted as aiming for a mere group compromise. The Lehrer-Wagner method is most persuasive as a model of how a group comprised of rational individuals should ideally reach *consensus*, given a sufficient amount of "respect" amongst group members. The idea of a consensus is that the final group opinion just equates to the opinion of all individual members. In other words, the members have each updated their own beliefs so that they now hold the same opinion; this will be the group's opinion. While such a group result will clearly be maximally satisfactory to all group members, we cannot expect this kind of agreement in all group decision situations. In certain cases, there will not be the kind of group "respect" that facilitates consensus, and not just because the group has been ill-chosen or because the members are not sufficiently open-minded. Some issues upon which a group must decide involve a mix of fact- and value-components. For instance, the problem of deciding which of a number of species should receive priority in a conservation project may involve a mixture of fact and value judgment. And individuals' value opinions are arguably not the sort of opinions that they are inclined to update based on the opinions of others. It is thus important to recognise that a formal method intended to achieve consensus, such as the Lehrer-Wagner method, will not be appropriate in all circumstances. Sometimes a group must settle for a mere compromise, and in such cases we may want to appeal to an alternative model of group deliberation.

While it is important to explore alternative decision models that facilitate compromise with respect to opinions expressed as arrays of numerical values, the Lehrer-Wagner method can also be given this weaker interpretation. Typically, the assessment of methods for achieving group compromise proceeds by considering the mathematical properties, or axioms, that the

[1] For a thorough description of the method, see Lehrer, K. and Wagner, C. (1981)'Rational Consensus in Science and Society'. Dordrecht: D. Reidel Publishing Company.

methods uphold. This is an approach that was championed by Arrow (1963)[2]: first we consider what axioms a reasonable group choice algorithm should satisfy, and then we consider what algorithms do in fact have these properties. If no algorithms satisfy one's chosen axioms (as per Arrow's famous impossibility result) then these axioms may need to be weakened. Wagner (1982) characterises the Lehrer-Wagner method amongst the class of "weighted arithmetic mean functions" that are intended to solve the problem of determining a group array of (positive) numerical values that sum to a specified value, given individual inputs of a similar nature.[3] Weighted arithmetic mean functions can be shown to satisfy a number of axioms, including the "independence of irrelevant alternatives" condition and "zero unanimity". The former axiom holds that the value of a particular position in the group array should depend only on the corresponding values in the individual members' arrays. The latter axiom holds that if all group members have a zero-value in a particular position of their individual arrays, then the group should assign a zero-value in the same position. These are arguably necessary properties for any reasonable group compromise algorithm. In any case, alternatives to the Lehrer-Wagner model should be assessed in terms of the axioms that they uphold.

[2] See Arrow, K.J. (1963) 'Social choice and individual values'. New York: John Wiley.
[3] See Wagner, C. (1982) 'Allocation, Lehrer Models and the Consensus of Probabilities', Theory and Decision 14: 207–220.

3. Impact (Highlights)

1. Mark Colyvan (CI) was invited to talk about the preliminary results of the project at the conference "Methodological Issues in the Social Sciences" at the Tilburg Center for Logic and Philosophy of Science, Tilburg, The Netherlands, 7 May 2007.

2. Katie Steele (Postdoc) has been invited to visit the Tilburg Center for Logic and Philosophy of Science (April–May 2008) to collaborate with Professor Stephan Hartmann on Social Choice Theory and matters related to the consensus project.

3. Helen Regan (CI) was successful in gaining funding on a series of NCEAS Working Groups in Santa Barbara, California, to work on topics related to the current project. (These working groups also included other CIs Mark Burgman and Mark Colyvan.)

4. Summary plan for future work

We have begun work on several projects that extend our previous work on the Lehrer-Wagner formal consensus method. The first is a survey of the range of formal methods in the literature that are competitors to the Lehrer-Wagner method and potentially appropriate to environmental decision-making. This will appear as a paper (mentioned above): Steele, K., Colyvan, M., Regan, H.M. and Burgman, M. 'A Survey of Formal Consensus Methods for Conservation Management'. The Lehrer-Wagner method is an example of a "weighted arithmetic mean function", but there are also "weighted geometric mean functions", which involve the multiplication rather than the addition of individual estimates. In addition, there is a whole class of "Bayesian methods" for determining a group probability distribution, which involve a prior probability distribution being updated by the opinions of each group member. The axiomatic approach will be important for comparing the merits of these alternative group methods. We need to consider what are the properties that the various "consensus" algorithms have, and thus what are the strengths and drawbacks of particular algorithms, given various types of group judgments that must be determined in a range of contexts.

When assessing the axioms that underpin the various consensus or aggregation methods, it is important to consider what makes adherence to a given axiom important. For instance, why might an axiom like "independence of irrelevant alternatives" be important? We hold that there are two major ways in which such axioms might be justified: on grounds of truth or on grounds of procedural fairness. The truth-tracking aspect is that group judgment aggregation methods are often thought to yield a result that has more chance of being true to the facts than any individual member's judgment. So some axioms might be defended because they contribute to this result. Or else, given certain extra assumptions about the propensity of individual members to make the "right" judgment, it might be shown that a particular group aggregation method is a better truth-tracker than another.[4] Alternatively, we might prefer to assess the axioms upheld by various aggregation methods in terms of how these axioms contribute to procedural fairness. For instance, this is surely the reasoning behind the "neutrality" condition, which holds that the group algorithm should not be sensitive to the particular labels (or names) attached to individuals. We think the truth-tracking versus fairness distinction is important and should be explored. In particular, the distinction seems very relevant to the difference between group methods that achieve consensus and those that achieve compromise. Indeed, One of the work-in-progress papers mentioned above is intended to pursue this issue—Steele, K. 'Pareto-optimality and the distinction between deliberation and aggregation in group choice'.

Another issue associated with the Lehrer-Wagner method that deserves further investigation is the representation of the uncertainty surrounding the final group result. As mentioned, the Lehrer-Wagner method guarantees convergence on a group array of judgments (e.g. a group probability distribution) under some reasonable assumptions about the respect weightings. It must be noted, however, that it is long-run (i.e., after possibly infinitely many iterations) convergence we're talking about here. Investigation into the mathematical properties of the convergence (i.e. rate of convergence, whether it is monotone, and questions about stability) under different circumstances would clearly be of great practical importance. The nature of the convergence to a group result for a given application might serve as an indication of the robustness or the level of uncertainty associated with the group result in that instance. Alternatively, one might prefer a group aggregation method that makes uncertainty explicit in the form of confidence intervals for the values in the group array. We might require the individual inputs to be an array of interval estimates rather than an array of point estimates,

[4] Condorcet's jury theorem, for instance, amounts to the result that majority rule yields a very high chance (which increases as the number of group members increases) of getting the right choice between two alternatives if each group member has over 0.5 chance of making the right choice.

or we might determine some way of producing a group array of intervals from precise individual arrays (by analysing rate of convergence, for instance). In any case, none of the formal consensus methods outlined in the literature explicitly consider the representation of uncertainty associated with group judgment. This would be a very useful thing to investigate.

The other major factor that must be considered in investigating group aggregation methods is their appropriateness and ease of implementation in environmental decision-making contexts. For instance, a potential drawback of the Lehrer-Wagner method is its dependence on group members assigning weightings of "respect" to each other. While efforts could be made to ensure that any such respect weightings are kept confidential, experts might still be reluctant to make explicit their assessments of other experts within the group. In fact, it is this problem that Regan *et al*. (2006) seek to address—they suggest a method for generating respect measures for use in the Lehrer-Wagner model that is based on the numerical opinions that members provide. (They model respect as the "distance" from an individual's opinion to another member's opinion.) This particular issue is specific to the Lehrer-Wagner method and other related weighted arithmetic mean functions. There are other more general questions to consider, however, when considering how effective a group aggregation method will be in practice. For instance, we want any such method to have a rationale that is reasonably transparent—it should be possible to explain the rationale to an expert group in a small amount of time, and the logic of the group aggregation method should, of course, be convincing. It is preferable that we do not ask experts to provide precise numerical estimates on issues (like "respect") that are inherently vague or complex. On the other hand, decision-making does, in the end, require some unavoidably difficult probability/value estimations, and this fact must be recognised.

A primary assumption underlying the aggregation methods discussed thus far is that experts will represent their opinions honestly. Indeed, the group result, whether it is genuine consensus or a mere compromise, is supposed to represent the group as a whole. The idea is that the group comes together to ultimately act as a team. In a significant number of environmental decision-making contexts, however, it is likely that an expert group will not have this amount of cohesiveness. The values/goals of members might be so diverse that it is better to model the group in terms of a more adversarial model, such as a game model, as opposed to a "social choice" model. The distinction between these different kinds of group scenarios is not well marked out in the literature. There are some synergies: for instance, game theory has been introduced into social-choice models in order to predict when group members will be likely to share the relevant information that they have, and when they will be disposed, rather, to engage in "strategic voting". Group members may well be inclined to pervert a group process if this will further their preferences, and such preferences need not be selfishly motivated. For instance, if a group member ultimately thinks that a particular conservation action should be undertaken, they may be inclined to overstate the extinction risk of a relevant species. The interface between social choice and game theoretic models is something that we would like to pursue in future research because we think it has great relevance for environmental decision-making.

5. Timeline for future work

12–17 Aug 2007: ACERA Robust Multi-Criteria Decision Analysis Workshop (Helen Regan, Mark Burgman and Katie Steele attending), in Hobart.

21–24 August 2007: ACERA Workshop on Formal Consensus Methods in Sydney (project members plus a couple of invited experts on formal consensus), including an afternoon session 2–5pm on 24th August for selected DAFF invitees.

July–December 2007: work on papers identified as “work in progress”, i.e. survey of formal consensus methods and investigation of truth-tracking versus procedural fairness justifications of group aggregation methods. In addition, work on extra issues identified in the August workshops.

January–June 2008: Continue work on papers and presentations.

April–May 2008: Katie Steele to visit the Tilburg Center for Logic and Philosophy of Science, Tilburg University, the Netherlands, on a special research fellowship.

June 2008: Paper presentation at the Formal Epistemology Workshop in Austin.

July–December 2008: Concentrate on identified future research—interface between game theoretic and social choice models. Plan and draft relevant papers.

October 2008: Possible 2nd meeting of project group to review results and discuss further research ideas.

November 2008: Paper presentation at the Philosophy of Science Association meeting in Pittsburgh.

January–June 2009: Focus on implementation of consensus/game theoretic model(s) that has(have) been identified to have particular strengths and ease of implementation. This may require putting together a computer program.

July–December 2009: Complete papers and computer program. Concentrate on policy recommendations that have arisen from the research.