


This article was downloaded by:[Colyvan, Mark]
On: 6 December 2007
Access Details: [subscription number 788079378]
Publisher: Routledge
Informa Ltd Registered in England and Wales Registered Number: 1072954
Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



Social Epistemology

A Journal of Knowledge, Culture and Policy

Publication details, including instructions for authors and subscription information:
<http://www.informaworld.com/smpp/title~content=t713765921>

Right Decisions or Happy Decision-makers?

Katie Steele; Helen M. Regan; Mark Colyvan; Mark A. Burgman

Online Publication Date: 01 October 2007

To cite this Article: Steele, Katie, Regan, Helen M., Colyvan, Mark and Burgman, Mark A. (2007) 'Right Decisions or Happy Decision-makers?', *Social Epistemology*, 21:4, 349 - 368

To link to this article: DOI: 10.1080/02691720601159711

URL: <http://dx.doi.org/10.1080/02691720601159711>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article maybe used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Right Decisions or Happy Decision-makers?

Katie Steele, Helen M. Regan, Mark Colyvan and Mark A. Burgman

Group decisions raise a number of substantial philosophical and methodological issues. We focus on the goal of the group decision exercise itself. We ask: What should be counted as a good group decision-making result? The right decision might not be accessible to, or please, any of the group members. Conversely, a popular decision can fail to be the correct decision. In this paper we discuss what it means for a decision to be “right” and what components are required in a decision process to produce happy decision-makers. Importantly, we discuss how “right” decisions can produce happy decision-makers, or rather, the conditions under which happy decision-makers and right decisions coincide. In a large range of contexts, we argue for the adoption of formal consensus models to assist in the group decision-making process. In particular, we advocate the formal consensus convergence model of Lehrer and Wagner (1981), because a strong case can be made as to why the underlying algorithm produces a result that should make each of the experts in a group happy. Arguably, this model facilitates true consensus, where the group choice is effectively each person’s individual choice. We analyse Lehrer and Wagner’s algorithm for reaching consensus on group probabilities/utilities in the context of complex decision-making for conservation biology. While many conservation decisions are driven by a search for objective utility/probability distributions (regarding extinction risks of species and the like), other components of conservation management primarily concern the interests of stakeholders. We conclude with cautionary notes on mandating consensus in decision scenarios for which no fact of the matter exists. For such decision settings alternative types of social choice methods are more appropriate.

Katie Steele is a postdoctoral fellow in the Department of Philosophy at the University of Sydney, Australia. Correspondence to: Katie Steele, Department of Philosophy, University of Sydney, Sydney, NSW, 2006, Australia. Email: katiesiobhansteele@yahoo.com

Helen M. Regan is an Assistant Professor of Ecology in the Department of Biology at San Diego State University, USA.

Mark Colyvan is Professor of Philosophy and Director of the Sydney Centre for the Foundations of Science, University of Sydney, Australia.

Mark A. Burgman is Professor of Environmental Science and Director of the Australian Centre of Excellence for Risk Analysis, University of Melbourne, Australia.

Keywords: Formal Consensus; Group Decision; Social Choice; Conservation Decision-making

Introduction

Group decisions raise a number of substantial philosophical and methodological issues. Here we wish to focus on the goal of the group decision exercise itself: is there an objectively correct answer to the question of what the group ought to do? After all, the right decision (with respect to empirical facts) might not be obvious to anyone, and thus might not please any of the group members. And conversely, a popular decision can fail to be the correct decision. Here we explore how the approach taken to social choices should depend on whether there really is a fact of the matter regarding what is the best action, and whether such facts are accessible to group members. A group in the business of determining the right course of action may need a certain methodology for group decision-making. On the other hand, if the exercise is the more modest one of ensuring that the decision-makers are happy, a different approach to the group decision process may be warranted.

We thus consider the relationship between right decisions and happy decision-makers. It is worth noting upfront, however, that our understanding of “happy” decision-makers is a rather qualified one. Here we will put aside the extreme case of where the group decision exercise is intended to merely make the group members happy in some psychological sense. Rather, we use the term “happy” in a normative sense; our interest is in the sort of group decision process that ought to make individual contributors happy upon rational reflection. To this end we consider the nature of the group process and the features that should make for happy decision-makers, given a specific context.

The issues we address have implications in many practical settings. We will highlight the practical significance of the issues by focusing on various conservation management decisions. Such decisions are particularly interesting in this regard. First, the decisions in question are not driven by pure matters of taste—there is often significant scientific input into the decision process—so it would seem that for the most part the search is for the right decision, as determined by the facts. The experts involved in conservation management decisions, however, usually display a diverse range of expertise and knowledge, from social and political sciences to ecology and natural resource management. As we shall see, these issues complicate the question of whether the objectively correct decision is accessible to individual group members, and what kind of decision process ought to make the members of a group happy. Indeed, it is widely held that diverse input, when synthesized in an appropriate way, will also lead to the right decision. Secondly, while not pure matters of taste, many conservation decisions ultimately affect stakeholders that have differing agendas. Moreover, in many applied conservation contexts, there is a mandate for diverse stakeholder groups to come to agreement on a course of action (Brower, Reedy, and Yelin-Kefer 2001; Gregory, Daniels, and Fields 2001). In such contexts it is believed that the act will have a greater chance of success if all parties are happy with it, irrespective of the rightness of the

decision. In this paper, we explore the intersection of these two ideals: happiness and rightness.

Before we start on the main topic we need to say a little about the coherence of the notion of a group decision. As in decisions elsewhere, group decisions depend on utilities (or at least preferences) and probabilities (most commonly, subjective probabilities). But both of these notions seem to be essentially tied to an individual. The very ideas of group subjective probability and group preferences might seem mistaken. Indeed, some researchers caution that if experts do not have similar views then it is meaningless to aggregate values of individuals into a group preference (Bolloju 2001; O'Leary 1993). This paper addresses this very issue. Our notion of happy decision-makers is tied to whether the selected group aggregation process is meaningful and warranted in the context in which it is used. We take a moment, however, to survey prior foundational work on this issue, starting with group preferences.

There is a rich and interesting literature on the problem of determining group preferences from individual preferences (e.g. Arrow 1963; Harsanyi 1976; Sen 1970). Of course the existence of this literature, in itself, doesn't guarantee the coherence of the notion of group preferences. Indeed, one of the lessons one might take from Arrow's (1963) famous impossibility result is that group preferences are quite unlike individual preferences. But still we must, and we do, determine group preferences from individual preferences—in many situations a decision must be made by the group. An election, for instance, is nothing other than an attempt to determine group preferences from the individual preferences of the voters. The approach we explore in this paper is different from, but is in the tradition of, Arrow (1963), Harsanyi (1976) and Sen (1970) of determining group preferences from individual preferences.

Making sense of group subjective probabilities is a little more difficult. But there must be some sense to be made of group beliefs. For example, society at large has a strong (but not universally high) degree of belief in the theory of evolution, and this belief has increased since 1859. Pettit (2004) goes so far as to suggest that at least some community groups with an explicit agenda can be considered institutional persons, in that they are required to exhibit consistency in their judgements over time. Moreover, such group judgements can be understood as supervening on, as opposed to being mere aggregations of individual judgements, on account of what is known as the "discursive dilemma" (Pettit 2004; Pettit and List 2004).¹ Here we are talking about determining group probabilities over propositions, rather than binary acceptance of propositions, but similar issues regarding the group aggregation process arise. We might consider impossibility results such as the "discursive dilemma" to cast doubt on the coherence of group judgement, or else, like Pettit, we might think that such results demonstrate that groups have a "life of their own" and are not mere aggregations of individuals. In any case, making sense of group beliefs is again something that needs to be done independently of present purposes. So having addressed, if not dispensed with, concerns about the core notions employed in group choice, let us move on to the issue of the objectivity of probabilities and utilities. This lies at the heart of the notion of right decisions.

Section 1: “Right” Decisions and Why They are Hard to Come By

Let us consider a “right” decision to be one that results in the objectively correct choice of action, given the information available. Clearly there is still a substantial question about what this amounts to. Within standard decision theory the goal is always the maximization of expected utility. In this framework, once we have decided on the “right” probability and utility functions, there is a fact of the matter about the ranking of two actions; the best action at our disposal (though sometimes there will not be a unique such action) will be the one with greatest expected utility. This final calculation is the least contentious aspect of decision-making. Of more interest is whether we can in the first place determine “right” probability and utility functions. In this discussion we will use a hard-line definition of “rightness”—we will consider “right” probability and utility functions to be those that are objectively correct, given the available evidence and decision context. If we define “right” probabilities and utilities in this way, then there is no distinction between what is “right” for a single decision-maker and what is “right” for a group (so long as the same evidence is available to all group members). Either the probabilities and utilities are correct when compared with the way the world is, or they are not. The question is whether this kind of “rightness” is generally accessible to decision-makers, or even makes sense.

Consider probabilities first. It is not too implausible that for all states of the world there is a (perhaps unknown) fact of the matter about what the objective probability is.² For instance, it is generally accepted that an ordinary coin has an objective probability of 0.5 for landing heads, provided the context or experimental set-up is one in which the coin is tossed fairly. Where there is much (relevant) frequency data available, talk of objective probabilities seems warranted. Indeed, there are some well known convergence theorems in the Bayesian literature (see Earman 1992) showing that any initial probabilistic belief function over events of the right type will eventually converge on a singular (objective) probability distribution, given sufficient frequency data. So in principle at least, if there are objective chances and every individual in the group aligns their degrees of belief with these objective chances, then we could arrive at group probabilities on objective grounds. This is all well and good, however, for simple kinds of events that lend themselves to endless repetition, like coin tosses. More complex events, such as those concerning ecosystem functioning that are likely to arise in conservation contexts, make talk of objective probabilities seem out of place. Objective chances for such events may theoretically exist, but they are a matter of speculation rather than fact.

Next consider utilities. Here is where the most serious difficulties lie. It seems implausible that even individual utility functions have an objective basis, let alone that group utility functions are objective. This particular issue arises in conservation management regularly when individuals disagree on the utility of alternative outcomes resulting from an environmental management decision. Consider restoration of an anthropogenically disturbed natural area (e.g. damming of a river that drowns an existing lake and creates an artificial lake, such as Lake Pedder in Tasmania, Australia). Suppose one places very high value on the resultant artificial lake and its ecosystem and

believes that any perceived utility in the restoration of the lake to its original state (by decommissioning the dam and reintroducing lost species and substrate) is substantially outweighed by the costs of restoration and the uncertainty that such a state could be achieved (Crowley 1999). Another individual might place little to no value on the lake due to its unnatural creation and perceive that any costs incurred in restoration would be substantially outweighed by the ecological and aesthetic value gained in shifting the system closer to its natural state prior to the building of the dam. These types of debates regarding the value of restoration occur regularly in conservation biology (Hildebrand, Watts, and Randle 2005). Who is right here? To what might we appeal to settle such a disagreement over values? Even if there is a “God’s-eye point of view” from which there is an objectively correct answer, that vantage point is not available to human decision-makers.

It is important to recognize, however, that subjectivity doesn’t mean that anything goes with respect to utilities. There are formal constraints on admissible utility functions (the von Neumann-Morgenstern axioms)³ and it might also be argued that there are some further ethical constraints (see Colyvan, Cox, and Steele forthcoming). Moreover, the decision context might be suitably narrow so as to make a particular choice of utility function the “right” one. For example, suppose that saving a particular species from extinction is the agreed goal of all members in the group. Then the utility function should be such that, with a suitable *ceteris paribus* clause, outcomes that yield the highest chance of persistence for the species have greatest utility (Maguire, Seal, and Brussard 1987; Regan et al. 2005). (If saving a particular species was the primary, but not the only, goal, then we might choose utility functions that reflect some minimum requirement for risk of extinction, and which then seek to optimize other social preferences.) Of course, due to problems of complexity, the objective chance of persistence of a species will not be obvious to decision-makers, nevertheless the utilities in a decision problem, like the probabilities, may concern matters of fact.

Even in the limited decision contexts where it is agreed that utilities should reflect facts about the world, decision-makers will not have direct access to these objective facts. Likewise, while in some contexts the relevant objective probabilities may be obvious to decision-makers, our comments above indicate that these kinds of decision situations are few and far between. Admittedly, even when objective probabilities and utilities are not known exactly, the right course of action might still be obvious to a group, because one option might have greatest expected utility according to all plausible probability and utility pairs. This kind of testing is known as “sensitivity analysis”; the aim is to determine how robust a particular choice is, given error estimates for the probability and/or utility distributions. Just as objective facts pertinent to a decision are rarely clear-cut, it will also be rare for one option to be the best act under any conceivable probability/utility assignment. The decision-makers must make use of best-estimate probability and utility distributions. Deciding on appropriate error margins and determining how such uncertainty should affect choice are very important further considerations (see Halpern et al. 2006). We do not address this kind of uncertainty (uncertainty about probability distributions and utility assignments) in

this paper; we focus rather on group agreement about best-estimate probabilities and utilities. Making links with decision models that accommodate uncertainty would be a useful and important extension of this work. In the next section we pursue the issue of why it is legitimate for experts to endorse different best-estimate probability distributions.

Section 2: Why Diversity of Opinion is Desirable

We might accept that there is an objectively correct probability distribution relevant to a particular decision context. But if the desired objective probability is unknown, a decision-maker can hardly appeal to the authority of the objective probability. In practice, we can only work with whatever information is available, and adjust the result according to the pertinent background information. (Of course, one option might be to gather more evidence, however, when there is an imperative to make a decision, as there is in many conservation applications, this is often not possible.) For instance, in order to estimate the chance of survival of a captive-bred Sumatran rhino reintroduced into the wild, it would be advantageous if the wildlife manager had access to frequency information pertaining to a large group of reintroduced Sumatran rhinos that are identical to this rhino in all respects causally relevant to the particular scenario under consideration. Then we might say that the wildlife manager's probability estimate was objectively correct. In practice, of course, the wildlife manager can only use the imperfect information available (based on a small number of prior reintroductions of Sumatran rhinos or similar species and perhaps other non-frequency-based information), and tailor the probability estimate according to the relevant peculiarities of the captive-bred rhino and its proposed new habitat. (For a discussion of the consequences of using imperfect frequency information in captive breeding decisions for the Sumatran rhino, see Rabinowitz 1995).

We want to emphasize here that good probability judgements almost always require subjective expert judgement. In nearly all contexts sufficient frequency data is simply not available,⁴ and so probability estimates must be fashioned according to the particularities of the case in question. But the expertise that a single individual brings to bear on a probability estimate will depend on past experience and training (leaving aside personal bias for the moment), the particular aspects of the system that he/she is most familiar with, the particular model favoured, and so on. For complex multi-faceted systems, such as those likely to be involved in conservation decisions, any one expert will only understand a portion of all the relevant details. This highlights the importance of consulting a range of experts—they will ideally have differing and complementary insights about the structure of the system and the processes involved. This is division of labour in research and decision-making. It is clearly required in decision contexts such as those that are the focus of this paper. Diversity of opinion amongst experts is not just something that we must deal with because it is inevitable, it is also a desirable state of affairs resulting from the fact that experts have differing specialties and motivational biases and thus see things from differing perspectives.

Section 3: Benefits of a Structured Group Decision Process

We argue that if members of a group recognize and value the importance of disagreement among experts, they will see group decision-making as an exercise in negotiating different opinions, rather than a debate about who is right. For the moment, we focus on decisions that require the input of scientific experts rather than the full spectrum of stakeholders. We have more to say about cases involving conflicting stakeholder interests later. Our interest is in how an expert group comes to accept a particular probability distribution in the course of decision-making. (We are also interested in how an expert group comes to accept a utility distribution; for many kinds of decisions we are considering, utilities can be treated similarly to probabilities, provided they are normalized distributions. We have more to say about utilities later.) Given that we want experts to take their role in a group decision process seriously, it is important that they be happy with how their opinion is taken into account. Furthermore, it is often in the interest of the scientific community at large to project a unified opinion to the public, as the handling of issues such as global climate change will attest to. We want to know what kind of group process, then, will make for happy expert decision-makers. There is a great deal in the literature describing the kinds of processes that result in group member satisfaction (Clemen 1996; Engelhardt and Caplan 1986; Kaplan 1992; Valverde 2001). Group member satisfaction is more likely to occur when each members' views are acknowledged and incorporated in the decision-making process. Group member satisfaction is less likely in situations where marginalization and bullying occur, or when members are forced to defer to the view of a dominant member against their will. Group members generally expect that the negotiation process will come to some resolution with a group decision at the end of it.

A well-functioning group may engage in constructive group debate that concludes with the acceptance of a group probability function, but anyone who has sat on a committee, or been a member of a family, cooperative or social group will agree that this is by no means guaranteed. In most situations, if there are not measures to ensure that all group members have a say, then less assertive members will not be heard and thus will not contribute to the final result. Just assuming that informal discussion will lead to group agreement is a dangerous ideal that will generally lead to the wishes of the dominant members being forced upon the rest (Peterson, Peterson, and Peterson 2005). Some of the problems that can occur in group decision-making exercises for which there are no consensus processes are: a dominant group member can manipulate group members to reach a position these other members do not hold (Hamilton 2003; Steinel and De Dreu 2004); the formation of social cliques within the group can isolate and alienate other group members that have unique expertise (Thomas-Hunt, Ogden, and Neale 2003); and idiosyncronies of group size and group member status can lead to deference to a single group member irrespective of that member's depth of knowledge (Ohtsubo and Masuchi 2004). Deference to a dominant group member, due to demonstrated self-assurance, or straight out intimidation, irrespective of that member's level of expertise, is a problem with serious consequences in group decision-making for environmental problems, especially in light of recent studies revealing the

complex (and often inverse) relationship between self-assurance and predictive accuracy (and indeed expertise and predictive accuracy, see Tetlock 2005).

A great deal has also been written on how groups should systematically determine a shared probability distribution, given the differing views of members. The methods are typically categorized as either behavioural or formal (Burgman 2005). The distinction concerns the extent to which numerical methods are employed to ensure that individual opinions converge upon a group result. Formal methods dictate mathematically how individual estimates should be combined; behavioural methods advocate a process of discussion and reflection, but ultimately rely on group members to update their individual estimates in whatever manner they each see fit. The two categories are not necessarily exclusive. Almost all group decisions are likely to be enhanced by an initial structured discussion allowing members to share evidence, state their position, question others' conclusions and update their own views.⁵ Behavioural methods address this aspect of the decision process (Burgman 2005). For example, discussion might be structured such that each group member is able to articulate the reasoning behind their probability estimate. Reasons for differences might then be debated and mere confusions addressed. Some methods combine behavioural and formal processes to reach a final group result. One such method that is commonly used in conservation and environmental decision-making contexts is the Delphi method (this method is outlined in Burgman 2005, along with other methods that combine behavioural and formal techniques). The Delphi method involves a number of stages whereby group members submit their probability estimates (or views) via questionnaires and are then provided with group statistics such as the modal probability estimate and the inter-quartile range; this information may then be used by members to reassess their own probability estimate (using whatever method each member sees fit), and the process is repeated again for the revised values, and so on until the group *hopefully* comes to some resolution.

While we value the communicative aspect of decision-making, and recognize the importance of theoretical work in devising formats for successful group discussion, we nonetheless draw attention to what we consider to be important ingredients of a happy decision: that the process is repeatable and it terminates. A decision process is repeatable if the same inputs (at a particular stage of the process) would subsequently yield the same result. The process is said to terminate if it guarantees a single group probability distribution, or failing this, the process has an identifiable point of completion, even if the group is left with a set of probability distributions amongst which no further distinctions can be made. The problem with behavioural aggregation processes is that they do not assure either of these characteristics. The Delphi technique is more explicit than others with respect to what information should guide revision of probability estimates; group statistics are provided at every stage, with the expectation that group members will reassess their own probability estimates in the light of others' contributions. But the procedure nonetheless lacks guidance as to how this individual updating should be done. Furthermore, there is no obvious motive for individuals to refine their probability estimates on the basis of group statistics. It is simply expected that they will each defer to the group as a whole, but the manner in which each does so will not be

identical, and any one member need not employ a consistent updating rationale. That all members will eventually converge on one probability distribution is merely wishful thinking.

It is important to bear in mind the distinction between consensus and compromise. In discussions of both behavioural and formal probability aggregation methods, it is generally assumed that the end result of the group procedure is a consensus probability distribution, rather than a mere compromise. But “consensus” tends to be used loosely. As we understand the distinction, at least, *consensus* is the state where all agents come to agree on the matter in question; compromise is where the agents do not agree about the matter in question but they agree, in the spirit of co-operation, to settle for something other than what they believe to be best. So, for example, a committee charged with determining how much money to devote to purchasing land for a wildlife reserve, may start out disagreeing on the appropriate amount but after discussion all committee members come to agree on an appropriate amount. This would be consensus. On the other hand, the discussion might not lead to anyone changing their view about the appropriate amount to be devoted to the wildlife reserve but the group may decide, for example, to average the amounts recommended by the committee members. This is a compromise. In general, formal (mathematical) methods such as averaging achieve compromise. Here, we outline a formal method for achieving consensus. It is clear that while compromises have their virtues, all other things being equal, consensus is preferable. By its very nature, consensus ensures that every member of the group gets what they want—everyone is happy. Compromise, on the other hand, does not ensure this. Indeed, compromise can often result in nobody getting what they want—the average, for example, does not necessarily coincide with the preferences or beliefs of any group member.

Section 4: Why We Should be Happier with Formal Consensus

Formal consensus methods prescribe particular methods for combining individual probability estimates. (The particular method we will focus on aims for genuine “consensus”, or so we will argue.) We have stressed that structured group discussion remains an initial phase allowing members to discuss the evidence pertaining to an unknown consequence or forecast. But at some point the discussion reaches its limit (i.e. when all evidence has been identified and no one in the group is further moved by others’ arguments regarding the significance of the evidence). Lehrer and Wagner (1981) refer to this state of collective stand-off as *reflective equilibrium* (i.e. each individual has reached their own private conclusions based on the information at hand). At this point in the decision process, a formal method for aggregating the probability estimates of individual members is most appropriate. As stated, a formal method ensures repeatability and the termination of the decision process, given a set of inputs from individual group members. It is important to employ the formal method at the point of reflective equilibrium—we want members to have refined their probability distributions on the basis of evidence and arguments that have been brought to their attention, and not purely out of deference to others. Admittedly it may be difficult for

experts to distinguish clearly between instances in which they are swayed by argument and instances in which they are swayed by another's reputation, but it is at least plausible that people are conscious of this distinction to some extent.⁶

Some object to any sort of formal group decision method, claiming that it effectively takes a decision out of the hands of the group or committee. But formal methods are not intended to alienate human decision-makers; the formal method merely models the behaviour of a well-functioning group. This is an important point. While formal methods are normative and employ precise mathematical algorithms, justification for these methods depends on what are considered characteristics of a well-functioning group. Even if there are no observed cases of groups spontaneously functioning in such a structured manner, models can nonetheless act as idealized descriptions of the world. In any case, formal consensus methods are justified only insofar as they appeal to our reflected intuitions about what makes for a good consensus decision. They are intended to model the process by which an ideal agent updates his/her opinions based on the opinions of others in the group. Such formal methods are no more alienating than, say, a single decision-maker setting up a decision table to work out his/her expected utilities and thus what course of action he/she should pursue. If the alienation objection stands, it stands as an objection to all formal models used in *all aspects* of decision-making. But this is taking things way too far. There is nothing alienating about formal models *per se*. We should not think that a population model, for instance, used to make decisions about conserving endangered species, makes predictions that have no connection to our observations of the world, and nor should we think that a consensus model gives decision results that have no relation to "real life" decision-making.

There is the further question of what justifies formal consensus methods. It might be argued that individual persons are the best judges of their own beliefs. Why then do we need a formal algorithm to help us update our beliefs in response to the differing opinions of our colleagues? Answered simply, formal consensus methods are supposed to reflect how ideal agents reason, and we do not all act as ideal agents all of the time. Often (arguably always) our reasoning is fallible, and updating beliefs in the context of a group decision is no exception. Kahneman and Tversky (e.g. 1982) have brought considerable attention to the common mistakes we tend to make in individual choice situations. If we employ a formal consensus method at the point of reflective equilibrium for a group, then we can avoid any mistakes in reasoning that group members may happen to make. A standard formal decision method is likely to be more reliable than the reasoning process of group members. Moreover, such methods are carefully designed so that they obey various principles or axioms that seem intuitively desirable. We would be wise to make use of this specialized work in judgement aggregation. A group that is convened to decide upon the best management plan for protecting a particular species will consist of experts in pertinent areas of ecology, botany and the like, but we do not expect these same persons to be experts in methods of consensus. Once an assembled group of experts have benefited from one another's knowledge about the relevant facts underpinning a probability estimate, there is arguably nothing further for them to decide upon that cannot be done better by a well-planned formal consensus method.

As can be expected, there is controversy among decision theorists as to what model best describes how a well-functioning group ideally behaves. Argument revolves around how well the models appear to be motivated, and also, in the tradition of Arrow, what axiomatic structure these models have. (Note that the group choice scenarios we have been considering are different from that featured in Arrow's impossibility theorem because we are assuming that for many expert decisions interpersonal comparison of utility is valid.⁷) Reviews of the competing formal consensus methods have been done elsewhere (see Cooke 1991; French 1985; Winkler 1968). Here, we outline one method in particular—the consensus convergence model of Lehrer and Wagner (1981)—because a single example is sufficient for our purposes. Moreover, there is considerable agreement about what kinds of properties a good consensus model should have. For instance, a straightforward desirable property is “unanimity”, which holds that “if all members agree on a probability, then the combined probability must also agree” (Clemen and Winkler 1999, 189). The Lehrer–Wagner model satisfies unanimity, as well as a number of other conditions (Lehrer and Wagner 1981). Moreover, the model is well justified; Lehrer and Wagner give a convincing account of how their method for rational consensus models the way an agent should reason when his/her views differ from others whose opinion he/she respects.

Before continuing let us outline the Lehrer–Wagner method.⁸ The algorithm is as follows: at the point of reflective equilibrium group members each nominate a probability estimate and also provide weightings of respect for all other members of the group. If these initial probability estimates differ, each member computes a revised estimate that is the weighted average (based on their personal weightings of respect) of the estimates of all members in the group. If disagreement remains after the new probability estimates have been calculated, each person computes a new weighted average based again on their personal weightings of respect compounded with the revised estimates of all group members. The process continues until all group members converge upon a common probability estimate. Convergence is guaranteed under some reasonable assumptions about the respect weightings.⁹ In addition to the convergence result, the algorithm accommodates all the relevant information available (i.e. the probability estimates of all group members, as well as their weightings of respect for other group members). The way in which members are required to revise their probability estimates at each stage also seems intuitively correct.¹⁰ If one group member has a certain level of respect for the opinions of another on a certain topic, then surely the first group member is rationally committed to revising her estimate to incorporate this information. We give this last idea more attention in the next section; we support the claim but think that it also suggests limitations for the use of formal consensus methods.

Just how members should determine their measures of respect for other group members is perhaps the most difficult aspect of the model to operationalize. Lehrer and Wagner (1981) think that “disinterested” respect weightings are most justifiable; these are based on assessments of competence, rather than similarity of opinion, with respect to the particular problem under consideration. Regan, Colyvan, and Markovchick-Nicholls (2006), however, outline a case for using respect weightings based on similarity of opinion (which Lehrer and Wagner refer to as “egoistic” respect weightings). The

chief benefit of the latter approach is that it allows the computation of respect weightings based solely on the individuals' probability estimates—group member A_i 's measure of respect for group member A_j is assigned based on the distance A_j 's view is from A_i 's. This may be distinctly advantageous in practical situations where group members are unwilling or unable to quantify their respect for the competence of others in the group—it is one thing to acknowledge a relative ordering of respect for other individuals in a group, it is an entirely different matter to place an abstract numerical value on the levels of respect a person has for the views or expertise of other members of the group. While there is evidence that group members do assign levels of respect for other members (although perhaps non-numerical) based on similarity to their own position (Yaniv 2004), this approach does not allow for the incorporation of other factors that might influence an individual's level of respect for another group member (e.g. past behaviour, motivation behind probability assignment, perceived honesty of disclosure).

Section 5: Stakeholders Versus Experts and the Limitations of Formal Consensus

Given the limited number of circumstances in which an individual expert is likely to know the “right” probability distribution, we think that there will be a great many cases for which consensus is appropriate. Indeed, as mentioned above, consensus is mandated in many conservation contexts to ensure all stakeholders are satisfied with the management decision, the idea being that decisions reached through consensus are more likely to be successfully implemented and less likely to end up in legal battles (although for problems with such mandates see Brower, Reedy, and Yelin-Kefer 2001; Gregory, Daniels, and Fields 2001; Peterson, Peterson, and Peterson 2005). Note that the experts in conservation decisions are sometimes also stakeholders, perhaps because they represent agencies with particular mandated goals, or because their academic work is motivated by a special interest. Such experts may be prepared to legally challenge a decision if they are not satisfied with the group process. We think that formal consensus methods such as the Lehrer–Wagner model have attractive features (such as the convergence result), and have much greater rational justification than any *ad hoc* process for combining probability estimates that might arise in the course of informal group discussion. Having acknowledged their significance, however, we want to consider what might be the limitations of using such formal consensus methods. Consensus may not always be possible (e.g. when stakeholders are in genuine conflict), or even desirable, in which case other group decision-making methods may need to be explored.

The constraints on the respect function that Lehrer and Wagner stipulate in order to ensure consensus provide a good starting point for testing the limitations of the model. Convergence is guaranteed only if there is “communication of respect” amongst group members, which amounts to each individual being connected to every other through a chain of respect. (For instance, there is a chain of respect between one group member and a second member if the first gives positive respect to someone who gives positive respect to someone else who gives positive respect to the second group member (Lehrer

and Wagner 1981, 27).) At least one person in the group also has to give positive respect to him/herself. If this “communication of respect” condition is not met, then convergence upon a common probability distribution is not guaranteed for the group. And in such cases, the usefulness of the consensus model is questionable. An obvious conclusion that we can draw here is that the make-up of the group is very important when we are trying to achieve consensus. But the choice of group members is of course prior to any decision-making process, and so we might say that the consensus model cannot be blamed for any shortcomings in group selection.¹¹ Surely if we are talking about a group of experts assembled to decide upon an issue within the scope of their expertise, then we would hope that there is at least “communication of respect” within the group. If this is not the case, then we would recommend reviewing the choice of group, rather than reviewing the choice of decision process.

Potential problems with respect weightings highlight a deeper issue that we need to address. Recall our distinction between “consensus” and “compromise”. When we talk about “consensus” we mean that the group comes to some kind of genuine agreement. The probability distribution that all group members converge upon should be rationally upheld by all; each member should be genuinely committed to this group estimate at the end of the revision process. In order to explain this point, we might liken the consensus algorithm to the Bayesian method of updating probabilities. A Bayesian’s best estimate of the probability of an event is supposedly based on his/her entire set of knowledge; and where new relevant data become available the prior probability is updated in the appropriate way to incorporate this data. We might contend that the Lehrer–Wagner consensus method simply extends this rationale that a probability estimate should be based on *all* the available data. Here our data is not just the scientific knowledge of the agent; it includes the estimates of other agents whose opinion we respect. The respect weightings an agent assigns to other group members should be such that he/she considers the revised weighted probability to be his/her new best estimate.

The Lehrer–Wagner model is elegant in that it aims to illustrate a consensus process that produces a result all group members are rationally committed to. While many group aggregation algorithms exist (Xu and Da 2003; Yeh et al. 1999), not all formal consensus methods intend the final estimate to have such strong rational justification. What is different about the Lehrer–Wagner model is the special meaning given to respect weightings and the way they are contributed and used by all group members. Of course, we acknowledge that respect of this kind is difficult to operationalize. For this reason, heuristic respect metrics (such as proposed by Regan, Colyvan, and Markovchick-Nicholls (2006) make for a useful modification of the original Lehrer–Wagner model. Nevertheless, it is instructive to consider the idealized version of the model in order to better understand the types of decisions best suited to consensus. They should be the types of decisions that involve genuine shifting of opinion according to what fellow group members think about the issue at hand.

So what sort of questions seem suited to such shifts in opinion? Surely it is only issues for which there is a fact of the matter that the group is trying to approximate. Indeed it seems misguided to talk of belief at all unless there is also belief in an underlying fact of

the matter. To do otherwise would be to commit oneself to a kind of Moore paradox (I believe P but there is no fact of the matter about P). An example may help here. Consider a group of experts assembled to determine the probability of persistence for a number of species under different management plans. Each recognizes that this is a complex problem, and that the different members of the group have different types of expertise about the ecosystem in question. All members produce estimates of the probability of persistence that accord with the evidence as they see it. While each member cannot trace how the other members have determined their respective probabilities (because each person will use different conceptual models and have a different appreciation of the evidence), there is reason to think that the other group members contribute important alternative perspectives that, when taken together and weighted appropriately, *ideally* converge upon the truth.¹² Even if convergence is guaranteed, the truth is not. All experts may be biased, even in judgements about objective probabilities (e.g. Hynes and Vanmarche 1977).

Let us contrast a decision scenario such as that described above with ones that involve substantial value judgement, where there is no obvious fact of the matter that the group is trying to approximate. Many examples spring to mind if we are considering reaching consensus on a utility function.¹³ For instance, suppose there are limited funds available to purchase land for wildlife reserves for the protection of threatened species. One option is to secure a site that will be protective of a flagship or charismatic threatened species, as this might promote conservation to the general public and garner more support (such as for tigers, rhinoceroses, gorillas, birds, pandas). Another option is to purchase a site that is protective of lesser known threatened species, such as many plant and insect species, which are afforded less protection in environmental policy and generally receive less public attention (e.g. virtually no insects are protected in most jurisdictions, mainly because of a lack of scientific and public interest in their welfare). Let's further suppose that the umbrella or focal species concepts cannot be applied here, that is one species cannot serve as a surrogate for other species (Lambeck 1997), and which species you'd prefer to protect with the purchase of land is nothing more than a question of taste. Here it seems that no objectivity grounds the choice of utility function. In situations like this, it is questionable whether formal consensus is the appropriate social choice mechanism—if there is no fact of the matter, why would group members be moved to revise their respective utility estimates in the light of their colleagues' opinions? Without such genuine belief revision, the consensus procedure of converging upon a shared group distribution lacks the kind of strong justification that it otherwise enjoys.

Admittedly the cases where convergence upon a group distribution seems unreasonable will tend to be decisions about utilities rather than probabilities. In this discussion we have focused on the application of consensus methods to probability distributions. Almost all interpretations of probability presume a connection between probabilities and the way things are in the world—while the Bayesian recognizes the subjectivity in probability estimates, objective Bayesians, at least, maintain that degrees of belief are supposed to respond to the facts (frequency data, our best scientific theories and the like). So it seems that whenever a group is trying to agree upon a

probability distribution, there will be some objective fact that underwrites the exercise and that the final probability (hopefully) responds to the facts in an appropriate way. Thus, provided experts recognize their responsibilities, and assuming the group is constructed so there is “communication of respect”, convergence upon a shared group probability distribution will generally be meaningful. While at first glance utilities look like an altogether different and more problematic matter, in many cases formal consensus methods will be highly applicable here too. Utilities need not always involve insurmountable value differences. Many decisions will have sufficiently narrow scope such that utilities can be tied to facts about the world. The example above about the probability of persistence of a particular species under different management options can be considered a utility assignment. The required estimates are probabilities, but they act as utilities in this particular decision (see Regan et al. 2005 and references therein for examples).

Our point is not that utility distributions are unsuited to formal consensus methods, but there will be some cases where the process of converging upon a shared group utility distribution makes less sense. These are cases when there is no fact of the matter that the group is trying to approximate. It might be a question of negotiating genuinely conflicting stakeholder interests, rather than coming to agreement about some feature of the world. It is not that the group members do not respect one another in these circumstances. Of course it is hoped that they each respect one another in all the morally commendable ways. But respect has a special meaning in formal consensus models. For the Lehrer–Wagner model, in particular, respecting another’s opinion means being willing to revise one’s own estimate to incorporate (with appropriate weighting) this alternative opinion. To use a simple example: while an agent might respect the fact that another agent rates vanilla ice-cream higher than chocolate, this is not the kind of respect that moves the first agent to alter her own ice-cream preferences.

Section 6: Concluding Remarks

Is the group at a loss in situations where members are not inclined to alter their own estimates based on the opinions of others? Of course not. Here, we have been talking about the process whereby a group comes to agree on a particular probability or utility distribution. If the decision situation is such that genuine agreement on a distribution would be forced or contrived, the group need not call it a day; there are other approaches to group decision-making that aim for compromise, if not consensus. Rather than trying to forge prior agreement on both the probability and utility functions relevant to a decision, it might be more appropriate to let individual members determine their own separate orderings for the available options, based on their own expected utility calculations. This type of group decision will be more like an election than consensus; it is the type of problem dealt with in traditional social choice theory.¹⁴

For example, Fox et al. (2004) reported on planning for forests in North-Eastern Tasmania in which the viability of 11 species was explored. No single planning option maximized prospects for all species. Any management action may act optimally for some (or one) species and suboptimally for others. In such circumstances, actions

cannot be defended on the grounds of objectivity. Individual utility functions are merely a matter of taste; different members simply have different preferences due to their particular biological interests, their aesthetic appreciation of nature, or due to some other socially mediated value. Alternatively, the group may allow options that satisfy, rather than optimize, conservation utilities and then maximize other values, such as timber or water. To determine a final ranking of the management options, the group may well combine their utility distributions according to a formal consensus method, such as the Lehrer–Wagner method outlined above. This might give a reasonable and fair result; the method does, after all, accommodate the utility assignments of all members. But for cases like this where the idea of a shared group utility distribution lacks rationale, we think it more appropriate to conceive of the decision process as a negotiation or compromise, rather than unanimous agreement, amongst members. For instance, some kind of voting method¹⁵ for deciding which species to prioritize might be the most transparent way to proceed.

In conclusion, decision procedures that should yield happy decision-makers will depend on the type of problem under discussion. If the problem is one where there are experts in the group who clearly know the relevant objective probability or utility distribution, then there is good reason for all group members to defer to this subgroup. We have noted, however, that most predictions regarding the impacts of our actions and their relative worth, particularly for complex decisions in conservation biology, are not within the knowledge base of any one individual. Moreover, it may never be worth assuming that one or a few individuals are vastly superior judges, because the overconfidence of experts makes such assignments of superiority error-prone and susceptible to biases resulting from psychology and context. In most cases, the group should be happy about diversity of opinion—it demonstrates (providing the group is well chosen) the varied expertise of members. But happy resolution of diverse opinion involves all views being taken into account in a consistent and justified manner. This is the beauty of formal consensus methods. The particular model that we've focused on—the Lehrer–Wagner model—clearly demonstrates why group members should be happy with a formal consensus. On this model, all members supposedly adopt (in the sense that they have reason to believe) the group probability/utility distribution at the end of the process. Thus the group choice is effectively each person's individual choice. As we have pointed out, however, this is all well and good when the issue is such that the combined expert estimates ideally converge upon the truth. Striving for the happiness of consensus where there is none to be found may lead to frustration. Hence, the requirement of consensus in many group decision contexts in conservation biology may be pure folly. Where we have a simple question of taste, the group decision may be better conceived as a compromise; in such cases we recommend an alternative type of social choice algorithm.

Acknowledgements

The authors would like to thank the Australian Research Council Centre of Excellence for Mathematics and Statistics of Complex Systems for sponsoring the Working Group

on “Decision Making for Complex Problems in Conservation” at the University of Melbourne, Melbourne, Australia in February and March 2005. This paper arose out of that working group and the authors are indebted to all the participants of the working group, but particularly to Yakov Ben-Haim, for ongoing discussion. We are also indebted to two anonymous referees for many very useful suggestions.

Notes

- [1] Pettit and List (2004) show that “proposition-wise majority voting over the judgments held by multiple individuals on some interconnected propositions can lead to inconsistent collective judgments on these propositions”. They call this the “discursive dilemma”. (Indeed, Arrow’s (1963) result is shown to be a special case of the “discursive dilemma”.) According to Pettit (2004), the group has to determine a strategy to avoid such collective inconsistencies of judgment, and in so doing becomes more than just an aggregate of individuals.
- [2] We make use of the idea of objective chance, but we also acknowledge that the nature of objective chance continues to puzzle philosophers (Hájek 2003b).
- [3] Recall that we are interested in *normative* decision theory, so while real agents may have irrational or inconsistent preference orderings, we do not allow such orderings here.
- [4] And even if there were exhaustive frequency data, there would likely still be disagreement regarding the appropriate reference class upon which to base the probability estimate. For a general discussion of the “reference class problem” see Hájek (2003a).
- [5] There may be exceptions. For instance, Clemen and Winkler (1999, 197) cite sources who claim empirical results show that in some cases “interaction of any kind amongst experts led to increased over-confidence and hence poorer calibration of group probability judgments”.
- [6] Of course, being conscious of the distinction and being able to demarcate reputation from argument reliably in practice is another matter. But all we are suggesting here is that as an idealization of the formal method, this assumption is not without some intuitive plausibility.
- [7] See, for instance Mackay (1980) for a discussion of Arrow’s theorem, and the significance of particular constraints such as the assumption that interpersonal comparison of utility is not valid.
- [8] For a thorough description of the method see Lehrer and Wagner (1981).
- [9] Although, 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.
- [10] Lehrer and Wagner (1981) are very persuasive on this point.
- [11] Having said that, the Lehrer–Wagner method actually does allow some room to manoeuvre with respect to group membership—the assignment of weights of respect means that not all members of the group have the same voice in the decision process. Every member has the same power to assign weights as the others but once the weights are assigned the end result may be that some members are marginalized while others are given greater say. Of course, such weightings do not change the group membership but they can, in effect, exclude some members from having any serious impact on the deliberations.
- [12] This example raises another, related issue, about how we should treat vagueness. Disciplines like conservation biology and conservation management rely heavily on vague terms such as “vulnerable” and the like (Regan, Colyvan, and Burgman 2000, 2002). The predominant view about vagueness is that while there are cases where vague terms determinately apply and cases where they determinately do not apply, with the borderline cases there is no fact of the matter about whether the term applies or not. Be that as it may, agents can still have beliefs about, and change their beliefs about, matters involving vagueness, though perhaps not for borderline cases. In any case, we set such complications aside for present purposes.

- [13] We are putting aside further problems with interpersonal comparisons of utility for the moment.
- [14] For an account of the broad issues of social choice theory, and the differing proposed group decision algorithms, see Arrow (1963), Harsanyi (1976) and Sen (1970).
- [15] See previous note for references. There are many different voting methods, each with different properties. An obvious method is the well-known “majority rules” (with respect to which of a pair of options is preferred), but this algorithm suffers from what is called the “voting paradox” (Arrow 1963, 2).

References

- Arrow, K. J. 1963. *Social choice and individual values*. 2nd ed. New York: John Wiley.
- Bolloju, N. 2001. Aggregation of analytic hierarchy process models based on similarities in decision makers' preferences. *European Journal of Operational Research* 128: 499–508.
- Brower, A., C. Reedy, and J. Yelin-Kefer. 2001. Consensus versus conservation in the Upper Colorado River Basin Recovery Implementation Program. *Conservation Biology* 15 (4): 1001–07.
- Burgman, M. A. 2005. *Risks and decisions for conservation and environmental management*. Cambridge: Cambridge University Press.
- Clemen, R. T. 1996. *Making hard decisions: An introduction to decision analysis*. 2nd ed. Belmont, CA: Duxbury Press.
- Clemen, R. T., and R. L. Winkler. 1999. Combining probability distributions from experts in risk analysis. *Risk Analysis* 19 (2): 187–203.
- Colyvan, M., D. Cox, and K. Steele. Forthcoming. Modelling the moral dimension of decisions.
- Cooke, R. M. 1991. *Experts in uncertainty: Opinion and subjective probability in science*. New York: Oxford University Press.
- Crowley, K. 1999. Lake Pedder's loss and failed restoration: Ecological politics meets liberal democracy in Tasmania. *Australian Journal of Political Science* 34 (3): 409–24.
- Earman, J. 1992. *Bayes or bust?: A critical examination of Bayesian Confirmation Theory*. Cambridge, MA and London: MIT Press.
- Engelhardt, H., and H. Caplan. 1986. Patterns of controversy and closure: The interplay of knowledge, values, and political forces. In *Scientific controversies: Case studies in the resolution and closure of disputes in science and technology*, edited by H. Engelhardt and H. Caplan. New York: Cambridge University Press.
- Fox, J. C., T. J. Regan, S. A. Bekessy, B. A. Wintle, M. J. Brown, J. M. Meggs, K. Bonham, R. Mesibov, M. A. McCarthy, S. A. Munks, P. Wells, R. Brereton, K. Graham, J. Hickey, P. Turner, M. Jones, W. E. Brown, N. Mooney, S. Grove, K. Yamada, and M. A. Burgman. 2004. Linking landscape ecology and management to population viability analysis. Report 2: Population viability analyses for eleven forest dependent species. Unpublished report to Forestry Tasmania. University of Melbourne.
- French, S. 1985. Group consensus probability distributions: A critical survey. In *Bayesian statistics*, edited by J. M. Bernardo, M. H. De Groot, D. V. Lindley and A. F. M. Smith. Amsterdam: North-Holland.
- Gregory, R., T. Daniels, and D. Fields. 2001. Decision aiding, not dispute resolution: Creating insights through structured environmental decisions. *Journal of Policy Analysis and Management* 20 (3): 415–32.
- Hájek, A. 2003a. Conditional probability is the very guide of life. In *Probability is the very guide of life: The philosophical uses of chance*, edited by H. E. Kyburg, Jr. and M. Thalos. Chicago and La Salle, IL: Open Court.
- Hájek, A. 2003b. Interpretations of probability. In *The Stanford encyclopedia of philosophy (Summer 2003 edition)*, edited by E. N. Zalta [cited 30 May 2006]. Available from <http://plato.stanford.edu/archives/sum2003/entries/probability-interpret/>; INTERNET.

- Halpern, B. S., H. M. Regan, H. P. Possingham, and M. A. McCarthy. 2006. Accounting for uncertainty in marine reserve design. *Ecology Letters* 9: 2–11.
- Hamilton, R. W. 2003. Why do people suggest what they do not want? Using context effects to influence others' choices. *Journal of Consumer Research* 29: 492–506.
- Harsanyi, J. C. 1976. *Essays on ethics, social behavior, and scientific explanation*. Dordrecht: Reidel.
- Hildebrand, R. H., A. C. Watts, and A. M. Randle. 2005. The myths of restoration ecology. *Ecology and Society* 10 (1) [cited 30 May 2006]: 19. Available from <http://www.ecologyandsociety.org/vol10/iss1/art19/>; INTERNET.
- Hynes, M. E., and E. H. Vanmarche. 1977. Reliability of embankment performance predictions. In *Mechanics of engineering*, edited by R. N. Dubey and N. C. Lind. Toronto: University of Waterloo Press.
- Kahneman, D., and A. Tversky. 1982. On the study of statistical intuitions. *Cognition* 11 (2): 123–41.
- Kaplan, S. 1992. “Expert opinion” versus “expert opinions”. Another approach to the problem of eliciting/combining/using expert opinion in PRA. *Reliability Engineering and System Safety* 35: 61–72.
- Lambeck, R. J. 1997. Focal species: A multi-species umbrella for nature conservation. *Conservation Biology* 11 (4): 849–56.
- Lehrer, K., and C. Wagner. 1981. *Rational consensus in science and society*. Dordrecht: Reidel.
- MacKay, A. F. 1980. *Arrow's Theorem: The paradox of social choice*. New Haven and London: Yale University Press.
- Maguire, L. A., U. S. Seal, and P. F. Brussard. 1987. Managing critically endangered species: A case study of the Sumatran rhino. In *Viable populations for conservation*, edited by M. E. Soulé. Cambridge: Cambridge University Press.
- Ohtsubo, Y., and A. Masuchi. 2004. Effects of status difference and group size in group decision making. *Group Processes & Intergroup Relations* 7 (2): 161–72.
- O'Leary, D. E. 1993. Determining differences in expert judgment: Implications for knowledge acquisition and validation. *Decision Sciences* 24 (2): 395–407.
- Peterson, M. N., M. J. Peterson, and T. R. Peterson. 2005. Conservation and the myth of consensus. *Conservation Biology* 19 (3): 762–67.
- Pettit, P. 2004. Groups with minds of their own. In *Socializing metaphysics: The nature of social reality*, edited by F. F. Schmitt. New York: Rowman and Littlefield.
- Pettit, P., and C. List. 2004. Aggregating sets of judgments: Two impossibility results compared. *Synthese* 140 (1–2): 207–35.
- Rabinowitz, A. 1995. Helping a species go extinct—The Sumatran rhino in Borneo. *Conservation Biology* 9 (3): 482–88.
- Regan, H. M., Y. Ben-Haim, B. Langford, W. G. Wilson, P. Lundberg, S. J. Andelman, and M. A. Burgman. 2005. Robust decision making under severe uncertainty for conservation management. *Ecological Applications* 15 (4): 1471–77.
- Regan, H. M., M. Colyvan, and M. A. Burgman. 2000. A proposal for fuzzy IUCN categories and criteria. *Biological Conservation* 92 (1): 101–08.
- Regan, H. M., M. Colyvan, and L. Markovchick-Nicholls. 2002. A taxonomy and treatment of uncertainty for ecology and conservation biology. *Ecological Applications* 12 (2): 618–28.
- Regan, H. M., M. Colyvan, and L. Markovchick-Nicholls. 2006. A formal model for consensus and negotiation in environmental management. *Journal of Environmental Management* 80 (2): 167–76.
- Sen, A. 1970. *Collective choice and social welfare*. San Francisco: Holden Day.
- Steinel, W., and C. K. W. De Dreu. 2004. Social motives and strategic misrepresentation in social decision making. *Journal of Personality and Social Psychology* 86 (3): 419–34.
- Tetlock, P. E. 2005. *Expert political judgment: How good is it? How can we know?* Princeton, NJ: Princeton University Press.
- Thomas-Hunt, N. C., T. Y. Ogdan, and M. A. Neale. 2003. Who's really sharing? Effects of social and expert status on knowledge exchange within groups. *Management Science* 49 (4): 464–77.

- Valverde, L. J. 2001. Expert judgment resolution in technically-intensive policy disputes. In *Assessment and management of environmental risks*, edited by I. Linkov and J. Palma-Oliveira. Dordrecht: Kluwer Academic.
- Winkler, R. L. 1968. The consensus of subjective probability distributions. *Management Science* 15: 361–75.
- Xu, Z. S., and Q. L. Da. 2003. An overview of operators for aggregating information. *International Journal of Intelligent Systems* 18 (9): 953–69.
- Yaniv, I. 2004. Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes* 93 (1): 1–13.
- Yeh, J.-M., C. Lin, B. Kreng, and J.-Y. Gee. 1999. A modified procedure for synthesising ratio judgements in the analytic hierarchy process. *Journal of the Operational Research Society* 50 (8): 867–73.