Social choice theory, belief merging, and strategy-proofness

Samir Chopra [a], Aditya Ghose [b], Thomas Meyer [c]

[a] *Department of Computer and Information Science*
Brooklyn College of the City University of New York
Brooklyn, NY 11210

[b] *Decision Systems Laboratory*
School of Information Technology and Computer Science
University of Wollongong
Wollongong, NSW 2522, Australia

[c] *National ICT Australia*
School of Computer Science and Engineering
University of New South Wales, 2052 Australia

Abstract

Intelligent agents have to be able to merge informational inputs received from different sources in a coherent and rational way. Several proposals have been made for information merging in which it is possible to encode the *preferences* of sources (Benferhat, Dubois, Prade, & Williams, 1999; Benferhat, Dubois, Kaci, & Prade, 2000; Lafage & Lang, 2000; Meyer, 2000, 2001; Andreka, Ryan, & Schobbens, 2001). Information merging has much in common with *social choice theory*, which aims to define operations reflecting the preferences of a society from the individual preferences of the members of the society. Given this connection, frameworks for information merging should provide satisfactory resolutions of problems raised in social choice theory. We investigate the link between the merging of *epistemic states* and some results in social choice theory. This is achieved by providing a consistent set of properties for merging akin to those used in Arrow's well-known impossibility theorem (Arrow, 1963). It is shown that in this framework Arrow's impossibility result does not hold. Similarly, by extending this to a consistent framework which includes properties corresponding to the notion of being *strategy-proof*, we show that results due to Gibbard and Satterthwaite (Gibbard, 1973; Satterthwaite, 1973, 1975) and others (Benoit, 2000; Barberá, Dutta, & Sen, 2000) do not hold in merging frameworks.

Email addresses: `schopra@sci.brooklyn.cuny.edu` (Samir Chopra), `aditya@uow.edu.au` (Aditya Ghose), `tmeyer@cse.unsw.edu.au` (Thomas Meyer).

1 Introduction

A crucial aspect of commonsense reasoning is the ability of agents to accept information from multiple sources and to put it together to come up with a composite picture that can guide further action and deliberation. Without this ability, most reasoning would quickly grind to a halt, with agents bewildered by the multiplicity and often, the mutual inconsistency of sources and the information forthcoming from them. Clearly, intelligent agents employing commonsense reasoning have to be able to merge inputs received from different sources in a coherent and rational way. Examples of this need already exist in robotics. As robotics moves from dealing with large complex industrial robots to simpler, smaller and less sophisticated— 'fast, cheap and out-of-control'— robots that rely on reduced intricacy in sensing and control, several problems have acquired significance. One of these is the problem of sensory deprivation and inconsistent sensory fusion: clearly the issue of merging of information is important here.

Several proposals for formalisms that incorporate this ability have been made for the merging of structures in which it is possible to encode the *preferences* of sources. In (Benferhat et al., 1999, 2000) *information fusion* is described in terms of possibility distributions (Dubois, Lang, & Prade, 1994; Zadeh, 1978) and the κ-framework developed in (Williams, 1995). In (Meyer, 2000, 2001), *information merging* is described in terms of epistemic states, structures in the style of (Spohn, 1988). In (Andreka et al., 2001) the *combination* of preferences is described in a framework where preferences are represented as arbitrary binary relations. In these frameworks, the notion of what would constitute 'rational' merging still remains an open question. Since it has been pointed out that information merging is similar to the operations studied in social choice theory—where the aim is to provide fair and equitable methods for aggregating the preferences of the members of a society to produce a single relation reflecting the preferences of society (Booth, 2002; Konieczny & Pino-Pérez, 1998; Maynard-Reid & Shoham, 1998; Konieczny & Pino-Pérez, 1999, 2002; Maynard-Reid & Lehmann, 2000)— it seems that one 'rationality' criterion for any proposed framework for information merging is that it deal satisfactorily with problems raised in social choice theory.

The formal similarity between the two areas is brought out best by an example. Consider an agent that is told by its two sources s_1, s_2, the following items of information:

- s_1: Stock X will rise tomorrow; stock Y will remain constant; stock Z will fall. (s_1's preferences then, are X, Y, Z)
- s_2: Stock X will fall tomorrow; stock Y will remain constant; stock Z will rise. (s_2's preferences then, are Z, Y, X)

The agent's epistemic task at this stage is to combine these two expressions of preference into an item of information that can guide its further actions i.e., which stock to purchase. Now consider a social aggregation operation that is handed the following input from the two members of society m_1, m_2 that express preferences for social goods (say, a particular taxation profile):

- m_1: Option X is highly desirable; option Y is tolerable; option Z is highly undesirable. (Source m_1's preferences then, are X, Y, Z)
- m_2: Option X is highly undesirable; option Y is tolerable; option Z is highly desirable. (Source m_2's preferences then, are Z, Y, X)

The task of the social aggregation operation is to take the inputs and to convert them into an ordering which will guide social policy i.e., the choice of which taxation profile to adopt. The prima facie similarity between the two scenarios above is obvious; unsurprisingly, so is the formal framework developed to model them.

The difficulties involved in devising a social aggregation operation are best illustrated by an example such as the Condorcet Paradox. [1]

> Suppose that Alice, Brian and Cait, are choosing between three candidates, Primus, Secunda, and Tertius, for a job. Alice prefers Primus to Secunda to Tertius. Brian prefers Secunda to Tertius to Primus. Cait prefers Tertius to Primus to Secunda. So a majority prefer Primus to Secunda, and a majority prefer Secunda to Tertius, and, paradoxically, a majority prefer Tertius to Primus. So preferences obtained by majority voting between pairs do not give a coherent ranking. Or, to put it differently, the outcome depends on the order in which the options are presented. If the first choice is between Primus and Secunda then Secunda will be eliminated and Primus will win when compared with Tertius. But if the first choice is between Primus and Tertius then Primus will be eliminated and then Secundus will win when compared with Tertius.

In this paper we extend the work in (Meyer, Ghose, & Chopra, 2001b) to investigate the link between the merging of epistemic states and some *impossibility* results in social theory: the Arrow and Gibbard-Satterthwaite theorems. Arrow showed that there is no aggregation operation satisfying certain reasonable postulates (Arrow, 1963). We show that the Arrow result does not hold in merging frameworks when preferences are represented in terms of epistemic states. Informally, epistemic states assign ranks to the valuations, or possible worlds, of the logic under consideration. We provide a list of properties to be satisfied by all rational merging operations and prove that the Arrow postulates, suitably modified to apply to this framework, can be derived from these

[1] Example taken from entry on voting paradoxes at `www.xrefer.com/entry/553842`

properties. We show that these properties are consistent.

Gibbard (Gibbard, 1973) and Satterthwaite (Satterthwaite, 1973, 1975) independently proved that under certain conditions, every reasonable method to aggregate the preferences of members of a society is vulnerable to *strategic manipulation* by the members of that society. This strategic manipulation is most relevant in voting scenarios where electors can misrepresent their actual preference profile so as to strategically elect their desired candidate. One of the Gibbard-Satterthwaite conditions on the output of the aggregation operation, single-valuedness, is quite restrictive, but similar impossibility results hold even in its absence (Benoit, 2000; Barberá et al., 2000). We extend our framework for the merging of epistemic states by adding properties which disallow various forms of manipulation. In particular, we propose properties which force merging operations to be *strategy-proof* and show that the addition of these properties results in a consistent extension of the basic framework for merging.

Why are these connections with social choice theory worth making? This is an important question since it is well-known that a move to structures richer than purely relational objects is a way to circumvent impossibility results (May, 1952; Dubois & Koning, 1991). We believe the answer to this is two-fold. At a purely formal level, there is interest in seeing how the formalisms of belief merging and social aggregation operations are really similar or dissimilar. More ambitiously, is there anything that the two areas could learn from each other? While the traffic might seem to be all one-way at this point in that social choice theory has a rich suite of impossibility theorems (Kelly, 1978) and aggregation operators which could be of value to the artificial intelligence community, it is not inconceivable that techniques –such as representational schemes and domain-specific aggregation operators– developed in the merging area would be of some use to the social choice theory community as well.

We believe the connection goes beyond the merely formal however. At a conceptual level, the provision of information by sources never takes place in a vacuum. The source's reliability, the context in which the information is provided, and most importantly, *the uses to which it can be put* are important factors in devising a method for aggregating this information. It makes sense then, to consider the connections with a formal framework developed by a community that deals with precisely this problem: a comprehensive theory for agent interactions should reflect the inputs of work aimed at formalising aspects of collective decision making. Furthermore, while social choice theory is primarily concerned with the aggregation of preferences, belief merging, from the angle we look at it, deals with the question of how to combine plausibility rankings of different sources. While these two questions are not the same, there is a sufficiently large overlap to justify a comparison of formalisms developed in the two areas.

Where does strategizing come into all of this? Voting scenarios can be seen as competitions where agents may strategize so as to get their preferences (preferred candidates) chosen by the aggregation operation. In order for the operation to be fair, it must be strategy proof. This is of value in a theory of automated agents since one can imagine a scenario in which agents have a vested interest in getting their preferences to show up higher in the composite preference profile. As an elaborate example, imagine a scenario where a small group of robots are sent on a reconnaissance mission on the surface of a remote planet to guide further exploration and mining. The information gathered in this fashion will be used by a central aggregator robot that is in charge of collecting the information collected by the discovery robots to come up with a plan for further exploration. If we assume that these robots are highly autonomous and—more ambitiously—equipped with some notion of self-interest, then the task of the central aggregator is to ensure that no robot is able to misrepresent its actual preferences so as to influence further discovery plans (perhaps the mother spacecraft is owned by a mining consortium made up of supposedly co-operating units). In this scenario, there is an assumption of autonomy on the part of the agents. Such an assumption is apt as artificial intelligence sets itself the goal of designing highly autonomous agents. Any such study will be incomplete without the investigation of scenarios in which agents are viewed as competitive and having conflicting desires and objectives. Clearly, merging operations need to deal with the possibility of the misrepresentation of beliefs so as to strategically influence the behavior of the composite system. In our robot example, we can imagine a situation in which the robots could have a vested interest in investigating a particular portion of the planet as opposed to another. The mechanisms that we describe in this paper are designed to take care of these sorts of scenarios. The connections then, between social theory and the formalisms for artificial intelligence, are timely and important at both the formal and conceptual levels. Most fundamentally, if we view the pooling or merging of epistemic resources as guiding the future actions of a group of agents then the comparison with social choice theory is especially appropriate.

The format of the paper is as follows. In Section 2 we lay the foundation for the discusssion to follow by defining epistemic states, lists and merging operations. In Section 3, we then provide a description of our framework via a presentation of some basic merging operations and their properties. In Section 4, we move on to describe the connections with social choice theory and voting procedures, pointing out the points of similarity and departure. Section 5 is devoted to a discussion of strategy-proof merging including a development of the abstract distance measures that play a crucial role in our framework. We conclude with some comments on the connections we have drawn in this study and provide some pointers to future work.

A brief note on our notation. We assume a finitely generated propositional language L closed under the usual propositional connectives and equipped with a classical model-theoretic semantics; the constants $\top, \bot$ are in L. V is the set of valuations of L and $M(\alpha)$ is the set of models of $\alpha \in L$. Classical entailment is denoted by $\models$. The set of natural numbers is denoted by $\mathbb{N}$. For $i \in \mathbb{N}$, we let $\mathcal{I}(i) = \{0, \ldots, i\}$ and $\mathcal{I}^+(i) = \{1, \ldots, i\}$. A *preorder* is a reflexive, transitive relation. A binary relation R on a set X is *connected* iff, for every $x, y \in X$, either xRy or yRx; a *total* preorder is a connected preorder. Examples involving valuations are phrased in the language with two atoms, p and q. Valuations are represented as sequences of 0s and 1s representing falsity and truth respectively: the first digit represents the truth value of p and the second one the truth value of q.

2 Epistemic states, lists and merging operations

We operate under the assumption that from the *epistemic state* of an agent the preferences of its sources can be represented as plausibility rankings of natural numbers on the valuations of L; the lower the number assigned to a valuation, the more plausible it is deemed to be. This is along the lines of work initially proposed by Spohn (Spohn, 1988) and was used in (Meyer, 2000, 2001) to define merging. Epistemic states are very similar to possibility distributions (Dubois et al., 1994) and the κ-framework (Spohn, 1988; Williams, 1995); it is relatively easy to translate between these frameworks. It is possible to use epistemic states in various ways. Our intention is for the ranks assigned to valuations to serve as markers in order to define a notion of *relative distance* between valuations, and nothing more. The reason for using it in this way is to avoid, to some extent, the problem of having to justify a particular assignment of numbers. At the same time it allows us to express the *strength* with which preferences are held; something that cannot be achieved with orderings on valuations. For example, in an epistemic state it is possible to express the information that I prefer u to v *more* than I prefer v to w.

Definition 1 *An* epistemic state Φ *is a (total) function from* V *to* $\mathbb{N}$.

It is possible to extract a consistent classical knowledge base from an epistemic state Φ by considering only those valuations with the best level of plausibility assigned to them. Let $M^i(\Phi) = \{v \in V \mid \Phi(v) = i\}$ and let $\min(\Phi) = \min\{\Phi(v) \mid v \in V\}$.

Definition 2 *A formula* $\phi \in L$ *is said to be a* knowledge base extracted from Φ *iff* $M(\phi) = M^{\min(\Phi)}(\Phi)$.

Following (Katsuno & Mendelzon, 1991), a knowledge base ϕ represents the set of wffs entailed by ϕ. Observe that the knowledge bases extracted from

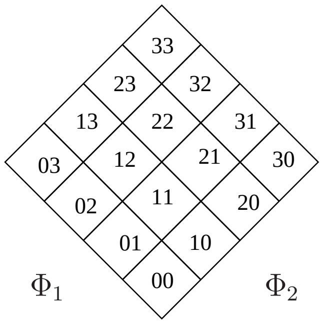

Fig. 1. A pictorial representation of an epistemic list containing two epistemic states Φ_1 and Φ_2. The sequence of two digits in each cell above indicates the natural numbers associated with valuations by the two epistemic states. A cell containing the sequence ij indicates the placement of those valuations assigned the value i by Φ_1 and assigned the value j by Φ_2.

Φ are all logically equivalent. We will often abuse notation by using $B(\Phi)$ to refer to *the* knowledge base extracted from Φ. The intention is that $B(\Phi)$ is some canonical representative of *all* the knowledge bases extracted from Φ. By extracting knowledge from an epistemic state in this way we ensure that $B(\Phi)$ will always be satisfiable, even though it may be the case that no valuation has a rank of 0 associated with it. This is in line with our informal interpretation of the natural numbers assigned to valuations; the choice of having 0 as the best plausibility rank which can be assigned to a valuation is purely a convenience.

Formally, we view merging as an operation in which the preferences of a *sequence* of sources, in the form of epistemic states, are combined to provide a new epistemic state representing the merged preferences of the sources. It is not sufficient to use finite *sets* of epistemic states, since different sources may have identical preferences, and the presence of more than one instance of an epistemic state may have a significant impact on the way in which merging is conducted.

Definition 3 *An* epistemic list E *is a finite non-empty list, or sequence, of epistemic states. We let* $|E|$ *denote the length of* E.

It is instructive to view an epistemic list pictorially as in figure 1. While such a pictorial view is only useful in representing epistemic lists containing two elements, it serves as a good foundation for understanding the principles underlying the merging of epistemic states in general.

In order for merging to be carried out at all it is crucial to make an assumption of *commensurability*; that all sources employ the same scale when they rank valuations. In practice this can be achieved by obtaining a worst level of plausibility P commonly agreed upon by all sources. Such a commitment in our formalism does not mean that any of the sources *has* to rank at least one valuation at P; it simply means that this is the worst level of plausibility that a source would ever consider attributing to any valuation. An agreement to use a particular worst level of plausibility means that all sources agree on a fixed level of *granularity*. That is, all agents are required to have the same worst level of plausibility. Note that such an approach is not without its own difficulties. It is reasonable to require that such a commonly agreed upon worst level of plausibility be no less than the worst level of plausibility initially assigned by individual sources. However, difficulties arise in deciding precisely how the diverse scales of plausibility are to be mapped on to a common scale. For example, suppose source A uses 1 as its worst level of plausibility and source B uses 2. A common level of granularity is then obtained by taking P to be at least two. Now suppose that we set $P = 5$, in order to ensure that the total number of ranks in the original two epistemic states are distributed evenly and equally to accomodate the variances in granularity. Suppose further that source B had a valuation ranked at 1, initially. In the new epistemic state, where $P = 5$, everything initially ranked at 1 by source B now have to be ranked at either rank 2 or rank 3. But which one? We can force it to choose, of course, but it might simply not have enough information available to make an informed decision.

Definition 4 *An epistemic state Φ is P-*capped*, where $P \in \mathbb{N}$, iff $\Phi(v) \leq P$ for every $v \in V$. An epistemic list E is P-*capped *iff every epistemic state in E is P-capped. The set of all P-capped epistemic lists is denoted by $\mathcal{E}^P$. The set of* all *epistemic states is denoted by $\mathcal{E}^\infty$.*

The following figure pictorially depicts a 7-capped epistemic state. Note that no valuation is assigned a rank of 0 in this example, but three are assessed as the most plausible (the ones receiving the lowest assigned rank of 1). Note also that no valuation has a rank of 7.

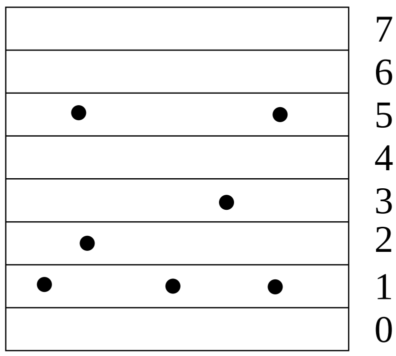

This brings us to the formal definition of merging.

Definition 5 *A P-capped merging operation Δ is a function from $\mathcal{E}^P$ to $\mathcal{E}^\infty$.*

P-capped merging does not necessarily yield P-capped epistemic states. In some cases, attempts to merge the information contained in an epistemic list may increase the granularity level of information contained in the resulting epistemic state. This is expected since an increase in granularity indicates more information available to the agent enabling a more fine-grained distinction to be made amongst different epistemic possibilities by the agent. For example, if I am assessing the chances of different teams in the basketball competition, I might rank two teams at par, but later, on receiving more information on player backgrounds, injury records and statistics from television reports and sports newspapers, rank one higher than the other. It should be noted though, that a monotonic increase of granularity is problematic since one would invariably end up with a linear ordering of valuations. A realistic account of merging should therefore also include a definition of 'contraction' like merging operations in which the level of granularity can be lowered in response to merging. This point is easily illustrated by reference to our previous example: I might have ranked one team above the other before receiving information that now forces me to revise my initial assessment and rank the two teams at the same level. In this case, the revision of plausibilities has resulted in a decrease in the granularity of my assessments of the team's chances in the competition. A development of 'contraction' like merging operations is not our present concern however, and we leave that topic aside as one for (interesting) future research.

Definition 6 *For every $n \geq 1$, every P-capped merging operation Δ is Q-*bound *for n iff for every P-capped epistemic list E s.t. $|E| = n$, $\Delta(E)$ is Q-capped and for some P-capped epistemic list F s.t. $|F| = n$ and some $v \in V$, $\Delta(F)(v) = Q$. Δ is Q-*bound *iff it is Q-*bound *for n for every $n \geq 1$. A P-capped merging operation which is Q-bound for n is referred to as (P, Q, n)-capped. Similarly, a P-capped merging operation which is Q-bound is referred to as (P, Q)-capped.*

For $n \geq 1$, every P-capped merging operation is (P, Q, n)-capped for some Q, but, as will be seen in section 3.1, need not be (P, Q)-capped for some Q. Q-boundedness should be understood as the new increased level of granularity obtained after merging has been conducted.

3 Basic merging

We are now in a position to provide some basic properties with which all P-capped merging operations ought to comply. Our claim is not that these properties *define* merging. Indeed, in section 5 we consider more desirable properties for merging which cannot be derived from (Δ0)-(Δ6) below. For the remainder of the paper we follow the convention that an epistemic list E has the form $[\Phi_1, \ldots, \Phi_{|E|}]$ and that an epistemic list F has the form $[\Psi_1, \ldots, \Psi_{|F|}]$. In the list given below v, w denote arbitrary members of V, the set of valuations.

(Δ0) $\Delta(E)(v) \geq \min\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\}$

(Δ1) If $|E| = |F|$ and if $\forall i \in \mathcal{I}^+(|E|)$ $\Phi_i(v) = \Psi_i(v)$, then $\Delta(E)(v) = \Delta(F)(v)$

(Δ2) $\forall n \geq 1$, if Δ is Q-bound for n, then $\forall q \in \mathcal{I}(Q)$ there is an $E \in \mathcal{E}^P$ and a $v \in V$ s.t. $\Delta(E)(v) = q$

(Δ3) If there is a bijection $\pi : \mathcal{I}^+(|E|) \rightarrow \mathcal{I}^+(|F|)$ such that $\Phi_i = \Psi_{\pi(i)}$ $\forall i \in \mathcal{I}^+(|E|)$, then $\Delta(E) = \Delta(F)$

(Δ4) If $\Phi_i(v) \leq \Phi_i(w)$ $\forall i \in \mathcal{I}^+(|E|)$, then $\Delta(E)(v) \leq \Delta(E)(w)$

(Δ5) If $\Delta(E)(v) \leq \Delta(E)(w)$, then $\Phi_i(v) \leq \Phi_i(w)$ for some $i \in \mathcal{I}^+(|E|)$

(Δ6) If $\Phi_i(v) = \Phi_j(v)$ $\forall i, j \in \mathcal{I}^+(|E|)$, $\Phi_i(v) \leq \Phi_i(w)$ $\forall i \in \mathcal{I}^+(|E|)$, and $\Phi_j(v) < \Phi_j(w)$ for some $j \in \mathcal{I}^+(|E|)$, then $\Delta(E)(v) < \Delta(E)(w)$

These properties need some explanation and motivation. (Δ0) requires that the rank assigned to a valuation after merging be no less (that is, no better) than the smallest (best) rank assigned to this valuation by any of the sources. This requirement can be justified by observing that there is no reason for an agent to regard a valuation as *more* plausible than suggested by any of its sources, even if all sources agree on the level of plausibility. That is, after merging a valuation should be seen as no more plausible than judged by the individual sources. (Δ1) states that the rank that Δ assigns to a valuation v is independent of the ranks assigned to any of the other valuations. This is similar in spirit to the property in social choice theory known as the *Independence of Irrelevant Alternatives* (Arrow, 1963) and is intended to capture a similar intuition. This issue will be discussed in more detail in section 4. The adoption of (Δ1) enables us to define merging in terms of aggregation functions on the natural numbers, i.e. as an operation on sequences of natural numbers. Let $seq^P = \{s \mid s = s_1, \ldots, s_n \text{ where } n \geq 1 \text{ and } s_i \in \mathcal{I}(P)\ \forall i \in \mathcal{I}^+(n)\}$. For $s \in seq^P$ we denote the length of s by $|s|$. The following proposition follows directly from (Δ1).

Proposition 1 *Let Δ be a P-capped merging operation satisfying (Δ1). Then there is a function $\delta : seq^P \rightarrow \mathbb{N}$ such that, $\forall v \in V$, $\forall E \in \mathcal{E}^P$, $\forall s \in seq^P$, if $|s| = |E|$ and $s_i = \Phi_i(v)$ $\forall i \in \mathcal{I}^+(|E|)$, then $\delta(s) = \Delta(E)(v)$.*

Merging operations on sequences thus have an indirect connection with the merging of epistemic states and it is only with the adoption of a property such as (Δ1) that this connection can be made explicit. This is in contrast with other approaches in the literature, such as that of (Benferhat et al., 1999, 2000), in which merging is defined directly in terms of sequences. We choose not to follow such an approach since the adoption of a propositional logic framework allows us to, firstly, compare our formalism for merging with other approaches such as (Benferhat et al., 1999, 2000; Lafage & Lang, 2000; Meyer, 2000, 2001; Andreka et al., 2001) and secondly, so as to make explicit the assumption ($\Delta 1$), which would otherwise have been impossible.

($\Delta 2$) is a convexity assumption. It ensures that for a merging operation bound by Q for n, no rank from 0 to Q remains unused for epistemic lists of length n. That this is a useful property can be seen by the fact that the ranks that are unused by an epistemic state are not useless and will be utilized in the epistemic state resulting from the merger of *some* epistemic list. If we did not have such a convexity assumption underlying merging, we would be in the position of having some ranks, for no special reason lying unused, thus assigning some special status to them not assigned to others. ($\Delta 3$) ensures that the order in which epistemic states occur in an epistemic list does not affect the outcome of merging. In (Meyer, 2000, 2001) this property was referred to as *commutativity* and in social choice theory it is known as *anonymity* (Kelly, 1978). ($\Delta 3$) rules out any notion of *prioritized merging* in which some sources are seen to be more important, or trustworthy, than others. This does not mean that we deem prioritized merging to be undesirable, but rather that prioritized merging depends on the existence of rational merging operations in which all sources are equally reliable. Indeed, in (Meyer, Ghose, & Chopra, 2001a) it is shown that there is a unique method of lifting non-prioritized merging operations into a prioritized setting. The adoption of ($\Delta 3$) means that it would be possible to define merging operations which receive inputs in the form of *multisets* or *bags*, instead of *lists*, of epistemic states. It is our position, however, that such assumptions should be made explicit, in the form of properties, instead of being encoded indirectly in the representational formalism. ($\Delta 4$) states that if all epistemic states in E agree that u is at least as plausible as v, then so should the resulting epistemic state. In the context of social choice theory, this is referred to as the *weak Pareto principle*. ($\Delta 5$) expects justification for regarding a valuation u as at least as plausible as v after merging has taken place: there has to be at least one epistemic state in E which regards u as at least as plausible as v. ($\Delta 5$) is a restatement of the Pareto Principle (in its contrapositive form), one of the properties used to establish Arrow's impossibility theorem in social choice theory (Arrow, 1963). The intuitions associated with ($\Delta 6$) have been discussed in (Meyer, 2000, 2001): in merging of knowledge bases, one of the basic properties is that if the knowledge bases to be merged are consistent then the merging is simply the conjunction of the knowledge bases. (This can be seen as an application of the methodological principle of minimal change of belief revision). ($\Delta 6$) is a generalisation of this principle for epistemic states. A knowledge base ϕ is a crude epistemic state in the sense that the models of ϕ are deemed to be strictly more plausible than its countermodels. Semantically, the conjunction of the knowledge bases amounts to the requirement that those valuations judged to be the most plausible by *all* knowledge bases have to *strictly* more plausible than any other valuation. In the same vein ($\Delta 6$) says that whenever all of the epistemic states in E agree on the level of plausibility of a particular interpretation u then, in the epistemic state obtained from a combination operation, u should be strictly more plausible than other valuations for which this unanimity does not hold. In this case, the lack of unanimity about a valuation v amounts to every epistemic state regarding v to be at most as plausible as u, but less plausible than u by at least one of the epistemic states. [2]

A possible additional property to consider is the following:

(Invariance) $\Delta([\Phi, \ldots, \Phi]) = \Phi$

It turns out that, although (Invariance) might be a desirable property for certain subclasses of merging operations, such as arbitration (Meyer, 2001), it is invalidated by some intuitively desirable operations, such as Δ_{Σ}, which we consider in the next section.

The following useful properties follow easily from the above properties.

Proposition 2 *Let Δ be a P-capped merging operation satisfying ($\Delta 0$).*

(1) If Δ is Q-bound for n then $Q \geq P$.
(2) If Δ satisfies ($\Delta 1$) and ($\Delta 6$), and $\exists i, j, \in \mathcal{I}^{+}(|E|)$ s.t. $\Phi_i(v) \neq \Phi_j(v)$, then $\Delta(E)(v) > \min\{\Phi_i(v) \mid i \in \mathcal{I}^{+}(|E|)\}$.

Part (1) of proposition 2 shows that the granularity level of information grows monotonically with merging, while part (2) provides lower bounds on the ranks assigned to valuations after merging. These results are consistent with the intuitions that:

- An increase in information provided by sources leads to an increase in the level of granularity of the merged epistemic state.
- The increase in plausibility of valuations is bounded in some 'rational' sense

For the latter, it is clear that merging operations which bring about unmotivated increases in the plausibility of certain valuations, such as a valuation being more plausible than judged by all of the individual sources, are unlikely to be judged as being 'rational'.

[2] Observe that dropping the first condition in the antecedent of ($\Delta 6$) is not feasible since the resulting property would be invalidated by certain reasonable merging operations, such as $\Delta_{\min 1}$.

3.1 Constructing merging operations

In this section we briefly consider some methods for constructing merging operations on epistemic states. This is not an exhaustive survey of merging operations found in the literature. The intention is merely to show that there are constructions which satisfy ($\Delta 0$)-($\Delta 6$). We consider the following merging operations (as before, in the list below v denotes an arbitrary member of V, the set of valuations):

(1) $\Delta_{\max}(E)(v) = \max\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\}$

(2) $\Delta_{\min 1}(E)(v) = \begin{cases} 2\Phi_1(v) \text{ if } \Phi_i(v) = \Phi_j(v)\ \forall i, j \in \mathcal{I}^+(|E|), \\ 2\min\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\} + 1 \text{ otherwise} \end{cases}$

(3) $\Delta_{\min 2}(E)(v) = \begin{cases} \Phi_1(v) \text{ if } \Phi_i(v) = \Phi_j(v)\ \forall i, j \in \mathcal{I}^+(|E|), \\ \min\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\} + 1 \text{ otherwise} \end{cases}$

(4) $\Delta_{\Sigma}(E)(v) = \sum_{i \in \mathcal{I}^+(|E|)} \Phi_i(v)$

These operations have been proposed and discussed in (Konieczny & Pino-Pérez, 1998; Benferhat et al., 1999, 2000; Meyer, 2000, 2001; Meyer, Ghose, & Chopra, 2001c), amongst others. Observe that $\Delta_{\max}$ and $\Delta_{\min 2}$ are (P, P)-capped, $\Delta_{\min 1}$ is $(P, 2P)$-capped, but that Δ_{Σ} is not (P, Q)-capped for any Q. We do know, however, that Δ_{Σ} is (P, nP, n)-capped for every $n \geq 1$.

$\Delta_{\min 1}$ may be thought of as a credulous merging operation; its construction may be explained as follows. Identify the valuations, if any, for which there is total agreement among all epistemic states about them being the most plausible, and take these to be the most plausible in the epistemic state resulting from the merging operation. The valuations on the next level of plausibility is obtained by considering all valuations which are deemed to be most plausible by at least one epistemic state. For the next level of plausibility we move to the valuations on which there is total agreement about them being the second most plausible set of valuations, followed by those valuations which are regarded as the second most plausible by at least one epistemic state. The process described above is repeated until all levels of plausibility for all the epistemic states have been catered for. At the heart of $\Delta_{\min 2}$ are three principles:

(1) The strong Pareto principle, which says that if all sources agree that valuation u is strictly better than v then the merging operation should rank u strictly better than v.
(2) When two sequences of natural numbers have the same minimum then the one in which there is total agreement is assigned a strictly better rank.

(3) Furthermore, it assigns to every interpretation the lowest possible value that can be assigned.

$\Delta_{\min 2}$ is the most credulous merging operation allowed by (Δ0) to (Δ6) in that it assigns valuations the best levels of plausibility allowed by the these postulates. To see why, observe that (Δ_0) requires of the value assigned to a valuation v to be no lower than the minimum value assigned to v by any of the sources. If all sources agree on the value assigned to v, then $\Delta_{\min 2}$ does indeed assign this minimum value to it. Furthermore, (Δ_6) demands that the value assigned to a valuation w on which all sources are *not* in agreement, should be *higher* than the value assigned to a valuation v for which all sources *are* in agreement, as long as all sources agree that w is no better than v. What this means, is that whenever all sources are not agreement about w, then the value assigned to w has to be *higher* than the minimum value assigned to w by any of the sources. What $\Delta_{\min 2}$ does is to assign to w the smallest value that is still bigger than this minimum.

$\Delta_{\max}$ assigns levels of plausibility by looking at the worst level of plausibility assigned to a valuation by any of the epistemic states. It may be thought of as representing a skeptical approach to merging: when confronted with differing assessments of plausibility of a particular valuation, pick the one that requires lesser commitment. Δ_{Σ} is an appropriate generalisation of an example by Lin and Mendelzon (Lin & Mendelzon, 1999) and was also independently proposed by Revesz (Revesz, 1993) as an example of weighted model fitting. The idea is simply to obtain the new plausibility ranking of a valuation by summing the plausibility rankings given by the different epistemic states. Δ_{Σ} may be considered a neutral—neither credulous nor skeptical—merging operation. Consider a 2-capped epistemic state and suppose v is ranked with a 0 by source A and a 2 by source B. Δ_{Σ} assigns 2 to v which might seem to indicate that it is as skeptical as $\Delta_{\max}$. However, while 2 is the maximum rank that $\Delta_{\max}$ can assign, 2 is about halfway between the best and the worst rank that Δ_{Σ} can assign for two sources (the worst rank being 4).

The diagrams in figure 2 represent some of the merging operations studied in this paper. The number in a cell represents the value assigned to a valuation by the merging operation. As an example, for $\Delta_{\max}$, consider the cells containing the value 3. These cells contain the valuations assigned the value 3 by Φ_1 and 0,1,2 or 3 by Φ_2 or 3 by Φ_2 and 0,1,2 or 3 by Φ_1; the merging operation assigns them all a 3. For Δ_{Σ}, the cell containing the value 6 at the top, is the cell containing the valuation assigned 3 by both epistemic states, thus explaining the assignment of the value 6. Note that in the illustration below, it is clear that while the original epistemic states considered the lowest level of plausibility to be 3, Δ_{Σ} has assigned a maximum level of 6. However, this should not be taken as indicating that the valuation assigned 6 is considered twice as implausible as the original valuations. Instead, the way to understand

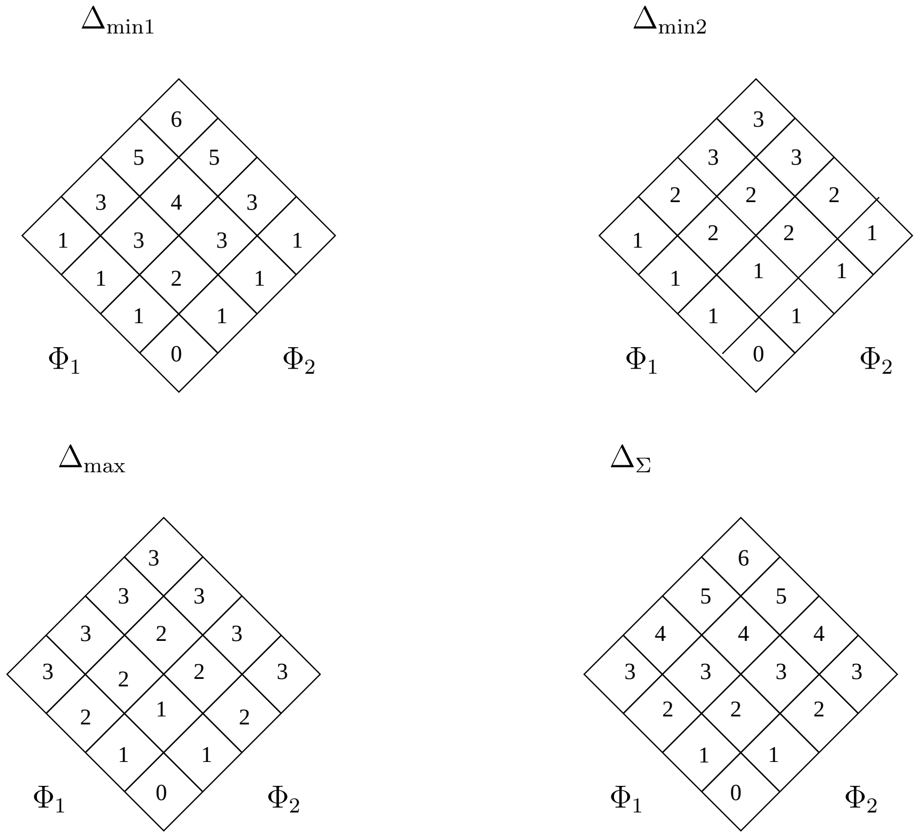

Fig. 2. A pictorial representation of four of the merging operations studied in this paper.

this appropriately is to think of the merging operation as having enabled a finer grained distinction amongst valuations. As the following result shows, the basic set of merging operations are satisfactory in a baseline sense: they satisfy the basic list of properties taken to be desirable for merging operations.

Proposition 3 $\Delta_{\max}$, Δ_{min1}, Δ_{min2} *and* Δ_{Σ} *all satisfy* ($\Delta 0$)-($\Delta 6$).

4 Social choice and merging

Social choice theory (Arrow, 1963; Sen, 1986) is a research area where the problems under scrutiny are similar to the problems encountered in merging. Social choice theory is concerned with the aggregation of preferences. There are many ways to aggregate individuals preferences to obtain societal ones: convention, established customs, religious doctrine and so on. Since most of these aren't 'fair' in being biased towards some members of society, how would a fair aggregation operation be constructed? Social choice theory abstracts the most salient requirements on the operations and then investigates the possibility of the existence of operations that could meet these requirements.

An individual's preferences are usually represented as a total preorder $\preceq$ over a (finite) set of alternatives Ω. For $x, y \in \Omega$, $x \preceq y$ means that x is at least as preferred as y. The interest then lies in the description of an aggregation operation over the preferences of n individuals, $\preceq_1, \ldots, \preceq_n$ which produces a new preference ordering over Ω. The similarities between this formalism and our framework for the merging of epistemic sets are obvious. We take Ω to be V, the set of permissible valuations, and use the total preorder on V *induced* by an epistemic state. Observe that such an induced total preorder contains less information than the epistemic state from which it was induced. For example, consider an agent that assigns a rank of 5 to w, 4 to v and 1 to u. This ranking clearly induces a preorder on valuations that tells us that an agent ranks u as more likely than v and w as more implausble than v, but it would not inform us of the fact that the agent considers v and w as *closer* to each other in plausibility than v is to u.

One of the most important results in social choice theory is Arrow's impossibility theorem (Arrow, 1963) which shows that there is no aggregation operation satisfying some intuitively desirable postulates. In this section we show that Arrow's result does not hold when recast into the framework of epistemic states. It is not our intention to show that this affects Arrow-type results in social choice scenarios. As noted in (Sen, 1970), Arrow-type results do not hold in social choice scenarions where interpersonal comparison of utilities is allowed. Social choice theorists have often used numerical scales for association with alternatives, but have attached far more significance to the numerical scales than we do. In the social choice context, such approaches are deemed problematic because of the difficulty in justifying numerical scales. By adopting the abstract distance measures discussed in Section 5, we concentrate on what properties are necessary for such a scale. Our (limited) objective is to clear the decks for merging operations and frameworks by showing that there is no need to fear the absence of 'rational' merging operations. Before presenting the postulates we note that social choice theorists fix the number of members of a society, and also assign each a unique index, while we make neither of these assumptions. This has some subtle effects on a number of postulates below, such as (IIA) and (ND).

Arrow's first postulate, dubbed *Restricted Range*, requires the result of an aggregation operation to be a total preorder on Ω. Essentially, if a list of preferences is submitted, then the output should look like one of its elements. In the context of belief revision this injunction is referred to as the *principle of categorical matching*. (Gärdenfors & Rott, 1995). In our framework this translates to the requirement that a merging operation produce an epistemic state – something which is built into our definition of merging. The second Arrow postulate, known as *Unrestricted Domain*, states that one ought to be able to apply the aggregation operation to any n-tuple of total preorders on Ω. Here the requirement is that any possible combination of preferences should be acceptable as inputs to the aggregation operation: the *a priori* rejection of certain inputs would indicate a biased aggregation operation, one that automatically rules out certain combinations as well. In our framework this translates to the requirement that merging may be applied to any (P-capped) epistemic set – again, something that is built into our definition of merging. The third Arrow postulate, known as the *Pareto Principle*, can be phrased as follows for epistemic sets:

(PP) If $\Phi_i(v) < \Phi_i(w)\ \forall i \in \mathcal{I}^+(|E|)$, then $\Delta(E)(v) < \Delta(E)(w)$

Note that (PP) is the contrapositive of ($\Delta 5$). The Pareto Principle requires that a preference that is held by all members of the input set should be respected by the aggregation operation. This is clearly intuitive: it would be strange if the aggregation operation were to produce a result that did not reflect the fact that option A was preferred to option B by all members of the group and regardless of the position of the other alternatives, this option should clearly reflect the preferences of the group.

The fourth Arrow postulate, known as the *Independence of Irrelevant Alternatives*, translates to the following postulate:

(IIA) $\forall E, F \in \mathcal{E}^P$ s.t. $|E| = |F|$, $\Phi_i(v) \leq \Phi_i(w)$ iff $\Psi_i(v) \leq \Psi_i(v)\ \forall i, j \in \mathcal{I}^+(|E|)$ implies that $\Delta(E)(v) \leq \Delta(E)(w)$ iff $\Delta(F)(v) \leq \Delta(F)(w)$

When deciding on the relative ordering of valuations v and w, (IIA) requires of us to disregard all other valuations. At least, that is the intuition. For example, if individual x prefers A to B then his ordering of these two should not be affected by the presence of the third alternative C. Why should an elector's preference for Gore over Bush be affected by the presence of Nader? The elector might vote for Nader in the final analysis but that would be because the voter preferred Nader over both alternatives and not because his ranking of Gore and Bush had changed. It is easily seen that the intuition does not hold in our more structured framework in which it is possible to define degrees of relative plausibility between valuations. Indeed our suspicion is that this intuition is not *always* properly captured in social choice contexts as well, thus providing perhaps a hint of an explanation for the impossiblity results.

Example 1 *Consider valuations* $v, w \in V$ *and let* $E = [\Phi_1, \Phi_2]$, $F = [\Psi_1, \Psi_2]$ *such that* $\Phi_1(v) = \Psi_2(w) = 0$, $\Phi_1(w) = \Phi_2(w) = \Psi_1(v) = \Psi_2(v) = 1$, *and* $\Phi_2(v) = \Psi_1(w) = 2$. *It is easily verified that* $\Phi_1(v) \leq \Phi_1(w)$ *iff* $\Psi_1(v) \leq \Psi_1(w)$ *and that* $\Phi_2(v) \leq \Phi_2(w)$ *iff* $\Psi_2(v) \leq \Psi_2(w)$. *Now consider the merging* $\Delta_{\max}$ *defined in section 3.1. It can be verified that* $\Delta_{\max}(E)(w) = \Delta_{\max}(F)(v) = 1$, *and* $\Delta_{\max}(E)(v) = \Delta_{\max}(F)(w) = 2$. *So it is not the case that* $\Delta_{\max}(E)(v) \leq \Delta_{\max}(E)(w)$ *iff* $\Delta_{\max}(F)(v) \leq \Delta_{\max}(F)(w)$. $\Delta_{\max}$ *therefore does not satisfy (IIA). Observe, however, that the ranks of* v *and* w *are obtained without reference to any of the other valuations! The rank assigned to* v *is even obtained* *without reference to w, and that of w is assigned without reference to v. In fact, it is easy to see that the rank of any valuation obtained by applying* $\Delta_{\max}$ *is independent of all other alternatives, even though* $\Delta_{\max}$ *does not satisfy (IIA).*

The example above shows that the representational framework of social choice theory is not expressive enough, and as a result, (IIA) states a stronger formal requirement than required by the intuition underlying it. Our framework enables distinctions not possible in social choice contexts via a scheme that would allow assignments of ranks to indicate strength of preferences. Take the election example above. I am sure that I do not wish to see Bush elected but am ambivalent about Gore. So I rank Bush at 10 and Gore at 5. Now Nader announces his candidacy and I rank him at 2. In doing so I did not need to refer to the ranks assigned to Bush and Gore. A voter might have been interested in voting for Gore over Bush till Nader announced his candidacy. While this voter would now rather vote for Nader, a more accurate expression of its preferences would be a scheme that would enable it to indicate that it preferred Nader to Gore, but both of these were to be far more preferable choices than Bush. In the usual social choice theory framework, where only total preorders are used, it is necessary to define independence indirectly, in terms of the ordering between two valuations. In our more structured framework this independence can be described directly, in terms of the rank assigned to a valuation. Our contention, then, is that ($\Delta 1$) is an appropriate reformulation of (IIA) since the rank of a valuation after merging has taken place is determined solely by the ranks assigned to the valuation by the various sources. The ranks assigned to other valuation do not play any role in the assignment of a rank during merging.

The last of the Arrow postulates, known as *Non-Dictatorship*, states that one source should never be able to completely dominate. In an epistemic scenario this corresponds to a situation whereby complete and total trust is not invested in a source to the extent that the source is able to bring about the disregard of conflicting information coming from other sources in any situation. For commonsense reasoning this clearly corresponds to a certain healthy skepticism. We can phrase this as follows.

(ND) For a fixed $n > 1$, there is no $i \in \mathcal{I}^+(n)$ such that, for every $E \in \mathcal{E}^P$, such that $|E| = n$, $\Phi_i(v) < \Phi_i(w)$ implies $\Delta(E)(v) < \Delta(E)(w)$ for every $v, w \in V$

Note that our version of (ND) fixes the length of the epistemic lists under consideration, something that is implicit in social choice theory scenarios.

It is easy to see that (ND) follows from ($\Delta 3$).

Proposition 4 *If a merging operation satisfies (*$\Delta 3$*) then it will also satisfy*

(ND).

From the result above and the comments on Arrow postulates which preceded it, it follows that the modified Arrow postulates all follow from (Δ0)-(Δ6). And since proposition 3 shows that there are merging operations which satisfy (Δ0)-(Δ6), we have shown that the Arrow impossibility result does not hold. It is the move from the total preorders on V to epistemic states, which have more structure than mere total preorders, which makes this possible. Note that the addition of such structure via an expression of the strength of preferences is recommended in the social choice literature as a possible method for circumventing Arrow's result. In those contexts, the selection of a suitable unit for commensuration of preferences is a question of much interest. The definition of epistemic states as functions from the set of valuations to the natural numbers—which are capped—solves this problem in merging scenarios.

5 Strategy-proof merging

Strategy-proofness is an idea that has received a great deal of attention in social choice theory, where it is frequently discussed in the context of elections. Voting theory concerns itself with the same set of problems as does social choice theory: how to define a voting procedure that most accurately and fairly represents the preferences of the population. Other concerns arise as well. Can the system be manipulated? What happens if new candidates enter or drop out of contention? Which method most accurately captures the voters' true intentions? The aim then, is to define an election procedure in which a winner is chosen in such a way that the outcome is immune to manipulation by voters, or coalitions of voters. The first impossibility result related to strategy-proofness is due to Gibbard (Gibbard, 1973) and Satterthwaite (Satterthwaite, 1973, 1975). Given some basic conditions on the number of available alternatives and the size of the electorate, and the (strong) requirement that an election procedure should produce a unique winner, their result shows that every election procedure which is non-dictatorial cannot be strategy-proof. This result is, perhaps, not particularly surprising. Consider, for example, the case in which two voters, Jack and Jill, have to choose between two candidates, Al and George. If Jack strictly prefers Al to George and Jill strictly prefers George to Al then there simply is not enough information to declare either Al or George the unambiguous winner. However, even if the requirement of producing a unique winner is relaxed it seems that Gibbard-Satterthwaite type results still hold (Benoit, 2000; Barberá et al., 2000).

Our aim in this section is to investigate notions of strategy-proofness in the context of merging. Requiring merging operations to be strategy-proof seems as necessary, and as desirable, as is the case for election procedures, or indeed, for aggregation operations in general. Constructing strategy proof procedures for merging ensures the existence of a procedure that lets the agent overcome misrepresentation of the value of the information provided by the sources and can be viewed as a form of epistemic safeguarding. The agent is able to repose a certain amount of trust in its decision making if it can be assured of the fact that the sources which present it with information are unable to systematically distort the final result to suit themselves. Before formalising the notion of strategy-proofness we consider two properties that allude to it.

Definition 7 *For $E, F \in \mathcal{E}^P$ s.t. $|E| = |F|$, and for $\mathcal{I} \subseteq \mathcal{I}^+(|E|)$, we denote by $rep(E, \mathcal{I}, F)$ the epistemic list obtained by replacing Φ_i with Ψ_i for every $i \in \mathcal{I}$.*

Intuitively $rep(E, \mathcal{I}, F)$ produces a modified version of E in which the sources mentioned in $\mathcal{I}$ have changed their preferences. That is, $rep(E, \mathcal{I}, F)$ represents the epistemic list obtained from E when Φ_i (the ith epistemic state in E) is replaced with Ψ_i (the ith epistemic state in F) for every $i \in \mathcal{I}$. For example, if $E = [\Phi_1, \Phi_2, \Phi_3]$ and $F = [\Psi_1, \Psi_2, \Psi_3]$, then $rep(E, \{2, 3\}, F) = [\Phi_1, \Psi_2, \Psi_3]$. In the properties (Mon↑) and (Mon↓) defined below, the set $\mathcal{I}$ in $rep(E, \mathcal{I}, F)$ is the singleton set $\{i\}$, from which it might seem that the notation $rep(E, \mathcal{I}, F)$ is unnecessarily clumsy. However, in some of the properties later in the section we use this notation with $\mathcal{I}$ being any subset of $\mathcal{I}^+(|E|)$.

(Mon↑) if $\Phi_i(v) \leq \Psi_i(v)$ then $\Delta(E)(v) \leq \Delta(rep(E, \{i\}, F))(v)$

(Mon↑) ensures that Δ exhibits monotonic behaviour. That is, it states that if a source worsens the rank it assigns to a valuation v, Δ will respond with a rank for v that is no better than the original. It is easy to show that (Mon↑) is equivalent to the the following property.

(Mon↓) if $\Phi_i(v) \geq \Psi_i(v)$ then $\Delta(E)(v) \geq \Delta(rep(E, \{i\}, F))(v)$

(Mon↓) states that if a source improves the rank it assigns to a valuation v, Δ will respond with a rank for v that is no worse than the original. The class of merging operations that we consider satisfies (Mon↑).[3]

Proposition 5 *If Δ satisfies (Δ4) then it also satisfies (Mon↑) (and (Mon↓)).*

This property does not guarantee strategy-proof behaviour, however, as will be shown in theorems 1 and 2. For a merging operation to be regarded as strategy-proof it must be the case that there is no incentive for any source to misrepresent its preferences. To be more precise, whenever a source provides an accurate representation of its preferences there should be a guarantee that

[3] These properties have also been considered in the context of the ELECTRE methods used in multicriteria decision-making (Roy & Bouyssou, 1993).

the result of merging will be *no less compatible* with its true preferences than if it had misrepresented its preferences.

5.1 Distance measures

Of course, the formalisation of such properties presupposes the existence of an appropriate *measure of compatibility* between epistemic states. Instead of defining one such measure, our approach will be to define a class of appropriate compatibility measures, all of which are obtained from the following abstract notion of *distance* between epistemic states. Let $\mathcal{D}$ be a set of *distance objects* and $\preceq$ a partial order on $\mathcal{D}$. We assume that $\mathcal{D}$ contains a *null element*, which we denote by $\mathbf{0}$, which is a *minimum* under $\preceq$.

Definition 8 *For any set of epistemic states $\mathcal{E}$, a function $\# : \mathcal{E} \times \mathcal{E} \to \mathcal{D}$ is a* distance measure *on $\mathcal{E}$ iff satisfies the following properties:*

(Ref) $\#(\Phi, \Psi) = \mathbf{0}$ *iff* $\Phi = \Psi$
(Sym) $\#(\Phi, \Psi) = \#(\Psi, \Phi)$

It is instructive to note that # satisfies the basic postulates for *equidistance*, used in (Borsuk & Szmielew, 1960) to define qualitative distance. There, equidistance is a four-place relation $E(a, b, c, d)$ on a set of points, which is read as "point a is just as far from point b as point c is from point d". The basic postulates for equidistance are

(E1) If $E(a, a, p, q)$ then $p = q$
(E2) $E(a, b, b, a)$
(E3) If $E(a, b, c, d)$ and $E(a, b, e, f)$ then $B(c, d, e, f)$

For a discussion of these postulates, see (Gärdenfors, 2000), pp. 15-17. For our purposes we define equidistance as a four-place relation on the set of epistemic states as follows: $E(\Phi_1, \Phi_2, \Psi_1, \Psi_2)$ iff $\#(\Phi_1, \Phi_2) = \#(\Psi_1, \Psi_2)$.

Proposition 6 *The equidistance relation defined on $\mathcal{E}^\infty$ satisfies (E1)-(E3).*

The distance measure # helps us define a measure of compatibility between epistemic states.

Definition 9 *Given a distance measure # on a set of epistemic states $\mathcal{E}$, an epistemic state Ψ is at least as compatible with Φ as Υ is, denoted by $\Psi \sqsubseteq_\Phi \Upsilon$, iff $\#(\Psi, \Phi) \preceq \#(\Upsilon, \Phi)$.*

Clearly every compatibility measure is a preorder.

At present, it is unclear how even a qualitative distance between epistemic states of differing granularities can be defined. Section 6 presents some ideas on this matter. Because of this we consider distance measures defined on P-capped epistemic states only. The restriction to distances measure on elements of $\mathcal{E}^P$ permits us to consider some more desirable properties.

(UB) If $|\Phi(v) - \Psi(v)| \leq |\Phi(v) - \Upsilon(v)|$ for every $v \in V$ then $\#(\Phi, \Psi) \preceq \#(\Phi, \Upsilon)$
(SUB) For $W \subset V$, if $|\Phi(w) - \Psi(w)| < |\Phi(w) - \Upsilon(w)|$ for every $w \in W$ and $\Phi(v) = \Psi(v) = \Upsilon(v)$ for every $v \in V \setminus W$ then $\#(\Phi, \Psi) \prec \#(\Phi, \Upsilon)$
(LB) If $\#(\Phi, \Psi) \preceq \#(\Phi, \Upsilon)$ then $\sum_{v \in V} |\Phi(v) - \Psi(v)| \leq \sum_{v \in V} |\Phi(v) - \Upsilon(v)|$
(SLB) If $\#(\Phi, \Psi) \prec \#(\Phi, \Upsilon)$ then $\sum_{v \in V} |\Phi(v) - \Psi(v)| < \sum_{v \in V} |\Phi(v) - \Upsilon(v)|$

The crucial point for our framework is that we use distances between any pair of objects that satisfy some basic properties. Recall that the numeric ranks assigned to valuations are markers with the naturals being used for convenience.

(UB) says that an epistemic state Φ is more compatible with an epistemic state Ψ than an epistemic state Υ if the difference between the ranks assigned to all valuations by Φ and Ψ respectively is less than that assigned by Φ and Υ. Note that this property requires us to ignore distances between valuations in a particular epistemic state and to concentrate on distances across epistemic states. (SUB) is a strict version of (UB); it insists that differnces between the ranks assigned to a particular set of valuations by Φ and Ψ respectively be strictly less than that assigned by Φ and Υ; for the remaining valuations, it insists that the three epstemic states agree on their rankings. (LB) is the converse of (UB); if Ψ is judged to be atleast as compatible with Φ as Υ is, then the sum of the distances between Φ and Υ is at least as great as that between Φ and Ψ. (SLB) is its corresponding strict version. These properties set limits for distance measures based on the internal structures of the epistemic states under consideration. To see why, consider the following two distance measures.

Definition 10 *(1) Consider the distance measure $\#^{\min}$ on all P-capped epistemic states for which $\mathcal{D} = \mathbb{N}$, the null element is 0, $\preceq$ is the normal linear order on natural numbers, and $\#^{\min}(\Phi, \Psi) = \sum_{v \in V} |\Phi(v) - \Psi(v)|$.*

(2) Let $\mathcal{D}$ be the set of all sequences of natural numbers of length $|V|$, and let the null element be the sequence consisting of $|V|$ zeros. For any two such sequences $s_1, \ldots, s_{|V|}$ and $t_1, \ldots, t_{|V|}$, we say that $s_1, \ldots, s_{|V|} \preceq t_1, \ldots, t_{|V|}$ iff $s_i \leq t_i$ for every $i \in \mathcal{I}^+(|V|)$. We assume that there is an enumeration $v_1, v_2, \ldots, v_{|V|}$ of the elements of V. Now let $\#^{\max}(\Phi, \Psi) = |\Phi(v_1) - \Psi(v_1)|, \ldots, \left|\Phi(v_{|V|}) - \Psi(v_{|V|})\right|$.

Observe that $\#^{\max}$ is the Pareto order and that some of the results to follow rely on its well-known properties.

Proposition 7 *$\#^{\min}$ and $\#^{\max}$ both satisfy (Ref), (Sym), (UB), (SUB), (LB) and (SLB).*

Note next that $\#^{\max}$ is the strongest of the distance measures satisfying these properties, and $\#^{\min}$ is the weakest.

Proposition 8 *Let $\#$ be a distance measure satisfying (UB) and (LB).*

(1) If $\#^{\max}(\Phi, \Psi) \preceq \#^{\max}(\Phi, \Upsilon)$ then $\#(\Phi, \Psi) \preceq \#(\Phi, \Upsilon)$.
(2) If $\#(\Phi, \Psi) \preceq \#(\Phi, \Upsilon)$ then $\#^{\min}(\Phi, \Psi) \preceq \#^{\min}(\Phi, \Upsilon)$.

Now consider the compatibility measures between P-capped epistemic states obtained from these two distance measures.

Definition 11 *$\sqsubseteq^{\min}$ is the compatibility measure obtained from $\#^{\min}$ and $\sqsubseteq^{\max}$ is the compatibility measure obtained from $\#^{\min}$. I.e. $\Psi \sqsubseteq^{\min}_{\Phi} \Upsilon$ iff $\#^{\min}(\Psi, \Phi) \preceq \#^{\min}(\Upsilon, \Phi)$ and $\Psi \sqsubseteq^{\max}_{\Phi} \Upsilon$ iff $\#^{\max}(\Psi, \Phi) \preceq \#^{\max}(\Upsilon, \Phi)$*

Observe that $\sqsubseteq^{\min}$ is a total preorder but that $\sqsubseteq^{\max}$ is just a preorder.

It turns out that $\sqsubseteq^{\max}$ is the strongest form of compatibility allowed by the properties above, and $\sqsubseteq^{\min}$ is the weakest. For ease of readibility we refer to *compatibility measures* as satisfying (Ref), (Sym), (UB) or (LB) when it is actually the distance measures used to obtain the compatibility measures which satisfy these properties.

Corollary 1 *Suppose $\sqsubseteq$ is a compatibility measure satisfying (Ref), (Sym), (UB) and (LB).*

(1) If $\Upsilon \sqsubseteq^{\max}_{\Phi} \Psi$ then $\Upsilon \sqsubseteq_{\Phi} \Psi$.
(2) If $\Upsilon \sqsubseteq_{\Phi} \Psi$ then $\Upsilon \sqsubseteq^{\min}_{\Phi} \Psi$.

The proof is immediate from proposition 8.

A possible objection to our framework may be formulated and responded to as follows. Consider three epistemic states Φ, Ψ and Υ which differ only on the values assigned to valuations v and w, and for which $\Phi(v) = 0$ and $\Phi(w) = 1$, $\Psi(v) = 3$, $\Psi(w) = 2$, $\Upsilon(v) = 3$ and $\Upsilon(w) = 4$. It is easy to check that the compatibility measures we allow will all consider Ψ to be more compatible with Φ than Υ is with Φ. But note that Φ and Υ have the same ordering of v and w, whereas in Ψ the ordering of these valuations is reversed. However, this apparently anomalous situation can be resolved by an appeal to the spirit of the property of Independence of Irrelevant Alternatives, since it compels us to ignore the distance between distinct valuations present in an epistemic state when determining the compatibility of that epistemic state with others. It is our contention that the removal of this apparent anomaly would require us to dispense with properties such as ($\Delta 1$) which try and enforce the spirit of the Independence of Irrelevant Alternatives in the merging framework.

5.2 Strategy-proof merging operators

We are now in a position to consider properties of merging operations requiring various versions of strategy-proofness. The first one we consider is intended to ensure that individual sources cannot gain anything by misrepresenting their preferences.

(ISP) $\forall E, F, G \in \mathcal{E}^P$ s.t. $|E| = |F|$ and $G = rep(E, \{i\}, F)$, $\Delta(E) \sqsubseteq_{\Phi_i} \Delta(G)$

(ISP) requires of $\Delta(E)$ to be at least as compatible with the preferences of source i than $\Delta(G)$ is, where E is the epistemic list in which i's preferences are represented accurately and G is obtained from E by i misrepresenting its preferences in some way. Given these properties a rational source will realise that is in its own interests to represent its preferences accurately.

In addition to misrepresenting its preferences, it is conceivable that a source may stand to benefit by completely *abstaining* from providing information. To consider our robot example again, take the case of a robot that simply refuses to confirm the investigations of other discovery robots. Such a lack of confirmation by abstention can affect the decision of the central aggregator (if say, it were employing some form of a majoritarian aggregation operator). Our goal is to define merging in such a way that it would be in the interest of every source to actually provide the required information, and not to omit any relevant information it may have.

Definition 12 *For $E \in \mathcal{E}^P$ and for some $\mathcal{I} \subset \mathcal{I}^+(|E|)$ we denote by $rem(E, \mathcal{I})$ the epistemic list obtained by removing Φ_i from E for every $i \in \mathcal{I}^+(|E|)$. Intuitively, $rem(E, \mathcal{I})$ produces a modified version of E in which the sources mentioned in $\mathcal{I}$ abstain from providing information.*

For example, if $E = [\Phi_1, \Phi_2, \Phi_3]$ then $rem(E, \{1, 3\}) = [\Phi_2]$.

The following property ensures that it is in the interests of a source to supply all the information it has at its disposal.

(IAP) $\forall E, G \in \mathcal{E}^P$ s.t. $G = rem(E, \{i\})$, $\Delta(E) \sqsubseteq_{\Phi_i} \Delta(G)$

(IAP) requires of $\Delta(E)$ to be at least as compatible with the preferences of source i as $\Delta(G)$, where G is obtained from E by removing the information supplied by source i.

Recall from section 3.1 that $\Delta_{\max}$ and $\Delta_{\min 2}$ are the only two (P,P)-capped merging operations we consider in this paper. As such, they are the only two merging operations which can be tested for strategy-proofness. The following result shows that $\Delta_{\max}$ is both strategy-proof and abstention-proof for any of the distance measures we regard as reasonable.

Theorem 1 *Consider any compatibility measure $\sqsubseteq$ satisfying (Ref), (Sym), (UB) and (LB).*

(1) Δ_{max} satisfies (ISP) and (IAP).
(2) Δ_{min2} satisfies (IAP) but does not satisfy (ISP).

Together, (ISP) and (IAP) thus ensure that when $\Delta_{\max}$ is applied, a source will provide its beliefs, all its beliefs, and nothing but its beliefs. Part 1 of theorem 1 shows that Gibbard-Satterthwaite style results do not hold in merging frameworks: in the context of epistemic states, it is possible to define rational merging operations which are immune to strategic manipulation or abstention by single sources.

5.3 Coalitions and coalition-proof merging

(ISP) defines strategy-proofness only relative to single sources and does not exclude the possibility of groups of sources, or coalitions, misrepresenting their preferences in such a way that the group as a whole benefits. The notion of benefit in an epistemic scenario is best explicated by reference to our example of the exploring robots. Consider a group amongst them that is interested in particular valuations being given preference by the central aggregator since this would allow a particular area of the planet surface to be explored further rather than ones they are not interested in. Such a group would be a coalition if some deliberate misrepresentation of their preferences resulted in an outcome which matches their true preferences more closely than if they had truthfully conveyed their preferences. In our example, we can imagine such a coalition forming as follows. There are three areas to be explored—X, Y, Z—by nine sources of information. A coalition of 3 members—s_2, s_4, s_5—would prefer that Z be explored first, followed by Y and X, that is their preference profile is (Z, Y, X). Assume now that the aggregator employs a majority based scheme for aggregating information. If the coalition is aware of the fact that 4 members—s_1, s_6, s_7, s_9—have the profile (X, Y, Z) and other two members—s_3, s_8—have the profile (Y, X, Z) then if the coalition were to represent its choices truthfully, it would see that its least desired item X would be the one picked. If however, it represented its profile strategically then it would make sense for it to change its profile to Y, Z, X in order to ensure that X does not get ranked first and instead that Y gets picked by the aggregator.

We denote a coalition of sources by their indices. The following *coalition-proof* property is intended to exclude the possibility of forming coalitions. It ensures that the coalition as a whole cannot benefit from misrepresenting their preferences.

(CP) $\forall \mathcal{I} \subseteq \mathcal{I}^+(|E|), \forall E, F, G \in \mathcal{E}^P$ such that $|E| = |F|$ and $G = rep(E, \mathcal{I}, F)$, it is not the case that $\Delta(G) \sqsubseteq_{\Phi_i} \Delta(E)$ for all $i \in \mathcal{I}$ and $\Delta(G) \sqsubset_{\Phi_j} \Delta(E)$ for some $j \in \mathcal{I}$

(CP) ensures that there is no group of sources which can misrepresent their preferences in such a way that no one of them is disadvantaged by it and at least one gains from it. As a result, the coalition as a whole cannot be seen to benefit from such a representation.

Observe that (CP) does not imply (ISP). If we restrict $\mathcal{I}$ in (CP) to singleton sets we obtain the following property.

(ICP) $\forall i \in \mathcal{I}^+(|E|)$, and $\forall E, F, G \in \mathcal{E}^P$ such that $|E| = |F|$ and $G = rep(E, \{i\}, F)$, $\Delta(G) \not\sqsubset_{\Phi_i} \Delta(E)$

But observe that $\sqsubseteq_{\Phi_i}$ need not be a total preorder, and so $\Delta(G) \not\sqsubset_{\Phi_i} \Delta(E)$ does not necessarily mean that $\Delta(E) \not\sqsubseteq_{\Phi_i} \Delta(G)$. Indeed, (ICP) includes the possibility that $\Delta(E) \not\sqsubseteq_{\Phi_i} \Delta(G)$ and $\Delta(G) \not\sqsubseteq_{\Phi_i} \Delta(E)$.

The option of abstention is open to coalitions of sources as well. For groups of sources the analogy with the epistemic counterpart in the robot's example is easily extended: consider a group of robots that decide together to refuse to report their findings. The following abstention-proof property is intended to exclude the possibility of the abstention of a coalition for the benefit of the coalition as a whole.

(AP) $\forall \mathcal{I} \subseteq \mathcal{I}^+(|E|), \forall E, G \in \mathcal{E}^P$ such that $G = rem(E, \mathcal{I})$, it is not the case that $\Delta(G) \sqsubseteq_{\Phi_i} \Delta(E)$ for all $i \in \mathcal{I}$ and $\Delta(G) \sqsubset_{\Phi_j} \Delta(E)$ for some $j \in \mathcal{I}$

(AP) ensures that there is no group of sources which can abstain from providing information, and by doing so ensure that no one of them is disadvantaged by it and at least one gains from it. As a result, the coalition as a whole cannot be seen to benefit from abstaining. Observe that (AP) does not imply (IAP). As is the case with (CP) and (ISP), this is because $\sqsubseteq_{\Phi_i}$ need not be a total preorder.

By providing a property which combines the requirement of being coalition-proof and abstention-proof with regard to groups of sources, we arrive at a general definition of strategy-proofness.

Definition 13 *Consider groups of sources* $\mathcal{I}, \mathcal{J} \subseteq \mathcal{I}^+(|E|)$ *such that* $\mathcal{J} \subseteq \mathcal{I}$, *and let* $\mathcal{K} = \mathcal{I} \setminus \mathcal{J}$. *For* $F \in \mathcal{E}^P$ *s.t.* $|E| = |F|$ *we let* $rr(E, \mathcal{J}, F, \mathcal{K}) = rem(rep(E, \mathcal{J}, F), \mathcal{K})$.

That is, $rr(E, \mathcal{J}, F, \mathcal{K})$ is the result obtained by first replacing Φ_j in E with Ψ_j for every $j \in \mathcal{J}$ and then removing Φ_k from the modified E for every $k \in \mathcal{K}$. For example, for $E = [\Phi_1, \Phi_2, \Phi_3, \Phi_4]$ and $F = [\Psi_1, \Psi_2, \Psi_3, \Psi_4]$, $rr(E, \{2\}, F, \{1, 4\}) = [\Psi_2, \Phi_3]$.

The property (SP) below requires of a merging operation to be *strategy-proof*: both coalition-proof and abstention-proof.

(SP) $\forall \mathcal{I}, \mathcal{J} \subseteq \mathcal{I}^+(|E|)$ such that $\mathcal{J} \subseteq \mathcal{I}$, $\forall E, F, G \in \mathcal{E}^P$ such that $|E| = |F|$ and $G = (rr(E, \mathcal{I}, F, \mathcal{I} \setminus \mathcal{J})$, it is not the case that $\Delta(G) \sqsubseteq_{\Phi_i} \Delta(E)$ for all $i \in \mathcal{I}$ and $\Delta(G) \sqsubset_{\Phi_j} \Delta(E)$ for some $j \in \mathcal{I}$

(SP) ensures that there can be no coalition of sources which misrepresent beliefs or refuse to provide information, and in doing so ensure that no member of the coalition is worse off, while at least one member j of the coalition benefits from this arrangement. (SP) ensures that there is no incentive for sources to present anything other than their true preferences, individually or in groups. Clearly (SP) implies both (CP) and (AP). It can be shown that when the distance measure $\#^{\max}$ is used, $\Delta_{\max}$ is strategy-proof in this sense.

Theorem 2 *Consider the compatibility measure* $\sqsubseteq^{\max}$.

(1) $\Delta_{\max}$ *satisfies (SP) and therefore also (CP), (AP) and (ICP).*
(2) Δ_{min2} *satisfies (AP). It does not satisfy (ICP), and therefore satisfies neither (CP) or (SP).*

These results show that when the distance measure $\#^{\max}$ is employed, $\Delta_{\max}$ is immune to manipulation and abstention from coalitions of sources as well. In addition, the fact that some strategy-proof properties are not satisfied by $\Delta_{\min 2}$ shows that these properties cannot be derived from the basic properties for merging and that their addition constitutes a strict extension of (Δ0)-(Δ6). The design of epistemic merging operations that satisfy these properties should be useful in fully illustrating the intuitions behind the postulates. That is, we expect that once specific constructions have been arrived at that satisfy these properties, the epistemic behavior they describe will further clarify the requirements placed on merging operations by the postulates.

6 Conclusion

In this paper we have drawn connections between information merging and social choice theory and shown that Arrow's impossibility result and versions

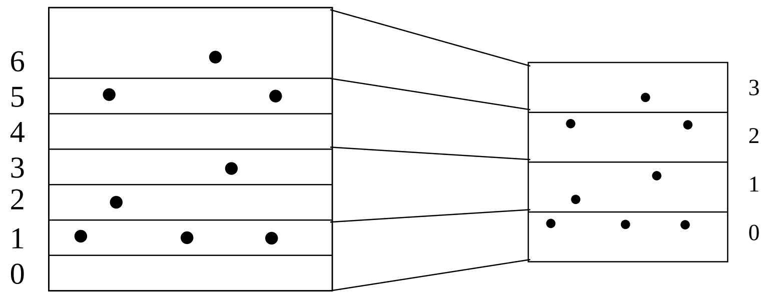

of the Gibbard-Satterthwaite theorem disappear when the representation of preferences is recast in terms of epistemic states. We started by considering a list of desirable properties for merging with no judgement made on whether these constituted any sort of canonical list. We then selected a class of merging operations and showed that they satisfied our initial list of properties. Once we reformulated the Arrow postulates for the epistemic case, we were able to show that the merging operations defined by us do enable a joint satisfaction of the postulates. Moving on to voting scenarios, we formalized concepts in the epistemic case which are the counterparts of those in voting, such as strategizing and coalition formation. We then showed that additional properties for merging would reflect the presence of these intuitions and showed that one of the merging operations defined satisfied our combined list of properties.

The results described here need to be elaborated upon, though. In section 5 the focus was on the special case of (P, P)-capped merging operations since measures of compatibility between P-capped epistemic states are then easily obtained. In the general case, however, where we are dealing with a (P, Q)-capped merging operation, we are faced with the problem of comparing epistemic states with different levels of granularity. For example, $\Delta_{\min1}$ defined in section 3.1 is a $(P, 2 \cdot P)$-capped merging operation, making it necessary to define an appropriate way of comparing P-capped epistemic states with $2 \cdot P$-capped ones. Currently it is unclear how to do so. One possible way to deal with this issue is to provide an appropriate method for mapping epistemic states with a high granularity level into epistemic states with the appropriate lower level of granularity. Such a mapping can be seen as a way to "convert" a Q-capped epistemic state to a P-capped one, thus making it possible to compare the two epistemic states. A mapping of this kind is depicted pictorially below. For example, for $\Delta_{\min1}$ we need a suitable method for mapping the elements of $\mathcal{I}(2 \cdot P)$ to $\mathcal{I}(P)$. In this case the appropriate mapping seems to be the function $\rho : \mathcal{I}(2 \cdot P) \to \mathcal{I}(P)$ such that $\rho(i) = \lceil i/2 \rceil$ (where $\lceil i/2 \rceil$ denotes the smallest integer which is no smaller than $i/2$). However, it is not clear how to determine which mappings are appropriate in the general case. At present the best we can do is to insist that a mapping ρ which converts a Q-capped epistemic state to a P-capped one should be a surjective function from $\mathcal{I}(Q)$ to $\mathcal{I}(P)$ such that $\rho(i) \leq \rho(j)$ whenever $i \leq j$.

Much work remains to be done in this area. How far can the connections between merging and social choice theory really be taken? Are there other impossibility theorems that are of relevance to the merging community and can a study of them be of use? What is the significance of epistemic sources *knowing* which merging operations are to be used on their inputs? A rich body of work has been done by the economics community in social aggregation operations and impossibility theorems; the study of these operations promises to be of much value to the merging theorist. Similarly, the rich suite of merging operations now being developed by merging theorists could conceivably be of use to the social choice theorist. The additional properties that we have suggested here should, in our opinion, be considered part of a canonical list of properties for merging operators. In that sense, our work in this study can be seen as preparatory for significant representation theorems to be proved. These would establish a class of merging operators as being rational provided they satisfy the canonical list. For the theory of autonomous agents in general, any exploration with a field that studies the possibly competitive nature of agents seeking to maximize their own interests and preferences is of value. In that sense, this connection with social choice theory, is in our opinion, timely and appropriate.

Acknowledgements

National ICT Australia is funded by the Australia Government's Department of Communications, Information and Technology and the Arts and the Australian Research Council through Backing Australia's Ability and the ICT Centre of Excellence program. It is supported by its members the Australian National University, University of NSW, ACT Government, NSW Government and affiliate partner University of Sydney. Thanks are due to the anonymous referees for some useful comments and suggestions.

A Proofs

Proposition 2. Let Δ be a P-capped merging operation satisfying ($\Delta 0$).

(1) If Δ is Q-bound for n then $Q \geq P$.

(2) If Δ satisfies ($\Delta 1$) and ($\Delta 6$), and $\exists i, j, \in \mathcal{I}^+(|E|)$ s.t. $\Phi_i(v) \neq \Phi_j(v)$, then $\Delta(E)(v) > \min\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\}$.

Proof:

(1) Follows by considering any epistemic list E such that, for some $v \in V$, $\Phi_i(v) = P$ for every $i \in \mathcal{I}^+(|E|)$.

(2) ($\Delta 0$) ensures that $\Delta(E)(v) \geq \min\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\}$. The strict inequality follows from ($\Delta 1$) and ($\Delta 6$).

□

Proposition 3. $\Delta_{\max}$, $\Delta_{\min1}$, $\Delta_{\min2}$ and Δ_{Σ} all satisfy (Δ0)-(Δ6).

Proof: We first consider $\Delta_{\max}$. (Δ0) follows by noting that $\Delta_{\max}(E)(v) = \max\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\}$. ($\Delta$1) is immediate. For ($\Delta$2), observe firstly that $\Delta_{\max}$ is P-bound. Now, pick any $p \in \mathcal{I}(P)$. We have to show that $\Delta_{\max}(E)(v) = p$ for some $v \in V$ and some $E \in \mathcal{E}^P$. To do so we let $E = [\Phi]$ and $\Phi(v) = p$ for every $v \in V$. (Δ3) is trivial. For (Δ4), suppose that $\Phi_i(v) \leq \Phi_i(w)$ for all $i \in \mathcal{I}^+(|E|)$. Then it follows that $\max\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\} \leq \max\{\Phi_i(w) \mid i \in \mathcal{I}^+(|E|)\}$. That is, $\Delta_{\max}(E)(v) \leq \Delta_{\max}(E)(w)$. For ($\Delta$5), suppose that $\Delta(E)(v) \leq \Delta(E)(w)$. That is, $\max\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\} \leq \max\{\Phi_i(w) \mid i \in \mathcal{I}^+(|E|)\}$. Suppose that j is such that $\Phi_j = \max\{\Phi_i(w) \mid i \in \mathcal{I}^+(|E|)\}$. Then it also follows that $\Phi_j(v) \leq \Phi_j(w)$. For (Δ6), suppose that $\Phi_i(v) = \Phi_j(v)\ \forall i, j \in \mathcal{I}^+(|E|)$, $\Phi_i(v) \leq \Phi_i(w)\ \forall i \in \mathcal{I}^+(|E|)$, and $\Phi_j(v) < \Phi_j(w)$ for some $j \in \mathcal{I}^+(|E|)$. Then $\Delta_{\max}((E)(v) = \Phi_i(v)\ \forall i \in \mathcal{I}^+(|E|)$. Furthermore, since there is some j such that $\Phi_j(w) > \Phi_j(v)$, it has to be the case that $\Delta(E)(w) \geq \Phi_j(w) > \Phi_i(v) = \Delta(E)(v)$.

Next we consider $\Delta_{\min2}$. (Δ0) follows by definition. (Δ1) is immediate. For (Δ2) observe that $\Delta_{\min2}$ is P-bound. Now, pick any $p \in \mathcal{I}(P)$. We have to show that $\Delta_{\min2}(E)(v) = p$ for some $v \in V$ and some $E \in \mathcal{E}^P$. To do so we let $E = [\Phi]$ and $\Phi(v) = p$ for every $v \in V$. (Δ3) is trivial. For (Δ4), suppose that $\Phi_i(v) \leq \Phi_i(w)$ for all $i \in \mathcal{I}^+(|E|)$. If $\Phi_i(v) = \Phi_j(v)$ for every $i, j \in \mathcal{I}^+(|E|)$, then $\Delta_{\min2}(E)(v) = \min\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\}$ and then $\Delta_{\min2}(E)(v) \leq \Delta_{\min2}(E)(w)$. Otherwise $\Delta_{\min2}(E)(v) = \min\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\} + 1$, and we need to consider two subcases. Subcase a: If $\Phi_i(w) = \Phi_j(w)$ for every $i, j \in \mathcal{I}^+(|E|)$, then $\Phi_i(w) \geq \Phi_j(v)$ for every $i, j \in \mathcal{I}^+(|E|)$, from which it follows that $\Delta_{\min2}(E)(w) \geq \Delta_{\min2}(E)(v)$. Subcase b: If $\Phi_i(w) \neq \Phi_j(w)$ for some $i, j \in \mathcal{I}^+(|E|)$, then $\Delta_{\min2}(E)(w) > \min\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\}$, from which it follows that $\Delta_{\min2}(E)(w) \geq \Delta_{\min2}(E)(v)$. For ($\Delta$5), suppose that $\Delta(E)_{\min2}(v) \leq \Delta(E)_{\min2}(w)$ and assume that $\Phi_i(w) < \Phi_i(v)$ for every $i \in \mathcal{I}^+(|E|)$. If $\Phi_i(w) = \Phi_j(w)$ for every $i, j \in \mathcal{I}^+(|E|)$ then it follows immediately that $\Delta_{\min2}(E)(w) < \Delta_{\min2}(E)(v)$, contradicting our supposition. So we consider the case where $\Phi_i(w) \neq \Phi_j(w)$ for some $i, j \in \mathcal{I}^+(|E|)$. If $\Phi_i(v) = \Phi_j(v)$ for every $i, j \in \mathcal{I}^+(|E|)$ then it has to be that $\min\{\Phi_i(w) \mid i \in \mathcal{I}^+(|E|)\} + 1 < \min\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\}$ and therefore $\Delta_{\min2}(E)(w) < \Delta_{\min2}(v)$, contradicting our supposition. On the other hand, if $\Phi_i(v) \neq \Phi_j(v)$ for some $i, j \in \mathcal{I}^+(|E|)$, then follows immediately that $\Delta_{\min2}(E)(w) < \Delta_{\min2}(v)$, again contradicting our supposition. So, the assumption leads to a contradiction in all case, from which it follows that $\Phi_i(v) \leq \Phi_i(w)$ for some $i \in \mathcal{I}^+(|E|)$. For (Δ6), note that, under the conditions imposed by this property, it follows that $\Delta_{\min2}(v) = \min\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\}$ and $\Delta_{\min2}(w) = \min\{\Phi_i(w) \mid i \in \mathcal{I}^+(|E|)\} + 1$, from which the result then follows immediately.

Next we consider $\Delta_{\min1}$. ($\Delta0$) follows by definition. ($\Delta1$) is immediate. For ($\Delta2$) observe that $\Delta_{\min1}$ is $(2 \cdot P)$-bound. Now, pick any $p \in \mathcal{I}(2 \cdot P)$. We have to show that $\Delta_{\min1}(E)(v) = p$ for some $v \in V$ and some $E \in \mathcal{E}^P$. To do so we let $E = [\Phi]$ and $\Phi(v) = p/2$ for every $v \in V$ if p is even, and $E = [\Phi_1, \Phi_2]$, $\Phi_1(v) = (p-1)/2$, and $\Phi_2(v) = (p+1)/2$ for every $v \in V$, if p is odd. ($\Delta3$) is trivial. For ($\Delta4$), suppose that $\Phi_i(v) \leq \Phi_i(w)$ for all $i \in \mathcal{I}^+(|E|)$. If $\Phi_i(v) = \Phi_j(v)$ for every $i, j \in \mathcal{I}^+(|E|)$, then $\Delta_{\min1}(E)(v) = 2 \cdot \min\{\Phi_i(v) \mid i \in \mathcal{I}^+(|E|)\}$ and then $\Delta_{\min1}(E)(v) \leq \Delta_{\min1}(E)(w)$. Otherwise $\Delta_{\min1}(E)(v) = 2 \cdot \min\{\Phi_i(v) \mid \ i \in \mathcal{I}^+(|E|)\} + 1$, and we need to consider two subcases. Subcase a: If $\Phi_i(w) = \Phi_j(w)$ for every $i, j \in \mathcal{I}^+(|E|)$, then $\Phi_i(w) \geq \Phi_j(v)$ for every $i, j \in \mathcal{I}^+(|E|)$, from which it follows that $\Delta_{\min1}(E)(w) \geq \Delta_{\min1}(E)(v)$. Subcase b: If $\Phi_i(w) \neq \Phi_j(w)$ for some $i, j \in \mathcal{I}^+(|E|)$, then $\Delta_{\min1}(E)(w) > 2 \cdot \min\{\Phi_i(v) \mid \ i \in \mathcal{I}^+(|E|)\}$, from which it follows that $\Delta_{\min1}(E)(w) \geq \Delta_{\min1}(E)(v)$. For ($\Delta5$), suppose that $\Delta(E)_{\min1}(v) \leq \Delta(E)_{\min1}(w)$ and assume that $\Phi_i(w) < \Phi_i(v)$ for every $i \in \mathcal{I}^+(|E|)$. If $\Phi_i(w) = \Phi_j(w)$ for every $i, j \in \mathcal{I}^+(|E|)$ then it follows immediately that $\Delta_{\min1}(E)(w) < \Delta_{\min1}(v)$, contradicting our supposition. So we consider the case where $\Phi_i(w) \neq \Phi_j(w)$ for some $i, j \in \mathcal{I}^+(|E|)$. If $\Phi_i(v) = \Phi_j(v)$ for every $i, j \in \mathcal{I}^+(|E|)$ then it has to be that $\min\{\Phi_i(w) \mid \ i \in \mathcal{I}^+(|E|)\} + 1 < \min\{\Phi_i(v) \mid \ i \in \mathcal{I}^+(|E|)\}$ and therefore $\Delta_{\min1}(E)(w) < \Delta_{\min1}(v)$, contradicting our supposition. On the other hand, if $\Phi_i(v) \neq \Phi_j(v)$ for some $i, j \in \mathcal{I}^+(|E|)$, then follows immediately that $\Delta_{\min1}(E)(w) < \Delta_{\min1}(v)$, again contradicting our supposition. So, the assumption leads to a contradiction in all case, from which it follows that $\Phi_i(v) \leq \Phi_i(w)$ for some $i \in \mathcal{I}^+(|E|)$. For ($\Delta6$), note that, under the conditions imposed by this property, it follows that $\Delta_{\min1}(v) = \min\{\Phi_i(v) \mid \ i \in \mathcal{I}^+(|E|)\}$ and $\Delta_{\min1}(w) \leq 2 \cdot \min\{\Phi_i(w) \mid \ i \in \mathcal{I}^+(|E|)\} + 1$, from which the result then follows immediately.

Finally we consider Δ_Σ. ($\Delta0$) follows by noting that $\min\{\Phi_i(v) \mid \ i \in \mathcal{I}^+(|E|)\} \leq \sum_{i \in \mathcal{I}^+(|E|)} \Phi_i(E)(v)$. ($\Delta1$) is immediate. For ($\Delta2$), observe firstly that Δ_Σ is $n \cdot P$-bound for n. Now, pick any $n > 0$ and any $p \in \mathcal{I}(n \cdot P)$. We have to show that $\Delta_\Sigma(E)(v) = p$ for some $v \in V$ and some $E \in \mathcal{E}^P$. To do so we let $E = [\Phi_1, \ldots, \Phi_n]$. Now, let $d = \lfloor p/n \rfloor$ and let $r = p - (d \cdot p)$. That is, r is the remainder when dividing p by n. We let $\Phi_i(v) = d$ for $i \in \mathcal{I}^+(n-r)$ and $\Phi_i(v) = d+1$ for $i > n-r$ and $i \leq n$. It is a matter of arithmetic to verify that $\Delta_\Sigma(E)(v) = p$. ($\Delta3$) is trivial. For ($\Delta4$), suppose that $\Phi_i(v) \leq \Phi_i(w)$ for all $i \in \mathcal{I}^+(|E|)$. It follows immediately that $\sum_{i \in \mathcal{I}^+(|E|)} \Phi_i(E)(v) \leq \sum_{i \in \mathcal{I}^+(|E|)} \Phi_i(E)(w)$. We prove the contrapositive of ($\Delta5$). Suppose that $\Phi_i(w) < \Phi_i(v)$ for all $i \in \mathcal{I}^+(|E|)$. It follows immediately that $\sum_{i \in \mathcal{I}^+(|E|)} \Phi_i(E)(w) < \sum_{i \in \mathcal{I}^+(|E|)} \Phi_i(E)(v)$. For ($\Delta6$), suppose that $\Phi_i(v) = \Phi_j(v)$ $\forall i, j \in \mathcal{I}^+(|E|)$, $\Phi_i(v) \leq \Phi_i(w)$ $\forall i \in \mathcal{I}^+(|E|)$, and $\Phi_j(v) < \Phi_j(w)$ for some $j \in \mathcal{I}^+(|E|)$. It follows immediately that $\sum_{i \in \mathcal{I}^+(|E|)} \Phi_i(E)(v) < \sum_{i \in \mathcal{I}^+(|E|)} \Phi_i(E)(w)$. $\square$

Proposition 4. If a merging operation satisfies ($\Delta 3$) then it will also satisfy (ND).

Proof: Pick an $n > 1$ and assume that for some $i \in \mathcal{I}^+(n)$, $\Phi_i(v) < \Phi_i(w)$ implies $\Delta(E)(v) < \Delta(E)(w)$ for every $v, w \in V$, and for every $E \in \mathcal{E}^P$ such that $|E| = n$. Now, if $i = 1$, let F be the epistemic list obtained from E by swapping around Φ_1 and Φ_2 in E, otherwise, let F be the epistemic list obtained from E by swapping around Φ_1 and Φ_i in E. By ($\Delta 3$), $\Delta(E) = \Delta(F)$. But this constradicts our assumption. For if $i = 1$ we can pick E in such a way that $\Phi_1(v) < \Phi_1(w)$ and $\Phi_2(w) < \Phi_2(v)$ for some $v, w \in V$, from which it follows by (ND) that $\Delta(E)(v) < \Delta(E)(w)$ and $\Delta(F)(w) < \Delta(F)(v)$. The case where $i > 1$ is similar. □

Proposition 5. If Δ satisfies ($\Delta 4$) then it also satisfies (Mon↑) (and (Mon↓)).

Proof: For (Mon↑), suppose Δ satisfies ($\Delta 4$), let $\Phi_i(v) \leq \Psi_i(v)$, and let $G = rep(E, \{i\}, F) = [\Phi_1, \ldots, \Psi_i, \ldots, \Phi_n]$. Observe that $\Phi_j(v) \leq \Phi_j(v)$ for every $j \in \mathcal{I}^+(|E|)$ such that $j \neq i$, and that $\Phi_i \leq \Psi_i$. So, by ($\Delta 4$), $\Delta(E)(v) \leq \Delta(G)(v)$. The proof for (Mon↓) is similar. □

Proposition 6. The equidistance relation defined on $\mathcal{E}^\infty$ satisfies (E1)-(E3).

Proof: For (E1), observe that that $\#(\Phi_1, \Phi_1) = \mathbf{0}$ by (Ref) and that if $\#(\Psi_1, \Psi_2) = \mathbf{0}$ then $\Psi_1 = \Psi_2$ by (Ref). (E2) follows from (Sym). For (E3), observe that if $\#(\Phi_1, \Phi_2) = \#(\Psi_1, \Psi_2)$ and $\#(\Phi_1, \Phi_2) = \#(\Upsilon_1, \Upsilon_2)$ then $\#(\Phi_1, \Phi_2) = \#(\Upsilon_1, \Upsilon_2)$. □

Proposition 7. $\#^{\min}$ and $\#^{\max}$ both satisfy (Ref), (Sym), (UB), (SUB), (LB) and (SLB).

Proof: (Ref) and (Sym) are trivial. For the remaining properties we first consider $\#^{\min}$. (LB) and (SLB) follow directly from the definition of $\#^{\min}$. For (UB), suppose $|\Phi(u) - \Psi(u)| \leq |\Phi(u) - \Upsilon(u)|$ for every $u \in V$. Then $\sum_{v \in V} |\Phi(v) - \Psi(v)| \leq \sum_{v \in V} |\Phi(v) - \Upsilon(v)|$ and therefore it is the case that $\#^{\min}(\Phi, \Psi) \preceq \#^{\min}(\Phi, \Upsilon)$. For (SUB), pick a $W \subset V$. Now suppose that $|\Phi(w) - \Psi(w)| < |\Phi(w) - \Upsilon(w)|$ for every $w \in W$ and that $\Phi(v) = \Psi(v) = \Upsilon(v)$ for every $v \in V \setminus W$. Then $\sum_{w \in W} |\Phi(w) - \Psi(w)| < \sum_{w \in W} |\Phi(w) - \Upsilon(w)|$ and therefore it is the case that $\#^{\min}(\Phi, \Psi) \prec \#^{\min}(\Phi, \Upsilon)$.

Next we consider $\#^{\max}$. (UB) and (SUB) follow directly from the definition of $\#^{\max}$. For (LB), suppose that $\#^{\max}(\Phi, \Psi) \preceq \#^{\max}(\Phi, \Upsilon)$. That is, $|\Phi(u) - \Psi(u)| \leq |\Phi(u) - \Upsilon(u)|$ for every $u \in V$. Then $\sum_{v \in V} |\Phi(v) - \Psi(v)| \leq \sum_{v \in V} |\Phi(v) - \Upsilon(v)|$. For (SLB), suppose that $\#^{\max}(\Phi, \Psi) \prec \#^{\max}(\Phi, \Upsilon)$. That is, $|\Phi(u) - \Psi(u)| \leq |\Phi(u) - \Upsilon(u)|$ for every $u \in V$, and for some $v \in V$, $|\Phi(u) - \Psi(u)| < |\Phi(u) - \Upsilon(u)|$. Then it is the case that $\sum_{v \in V} |\Phi(v) - \Psi(v)| < \sum_{v \in V} |\Phi(v) - \Upsilon(v)|$. $\square$

Proposition 8. Let $\#$ be a distance measure satisfying (UB) and (LB).

(1) If $\#^{\max}(\Phi, \Psi) \preceq \#^{\max}(\Phi, \Upsilon)$ then $\#(\Phi, \Psi) \preceq \#(\Phi, \Upsilon)$.

(2) If $\#(\Phi, \Psi) \preceq \#(\Phi, \Upsilon)$ then $\#^{\min}(\Phi, \Psi) \preceq \#^{\min}(\Phi, \Upsilon)$.

Proof:

(1) Suppose that $\#^{\max}(\Phi, \Psi) \preceq \#^{\max}(\Phi, \Upsilon)$. That means that for every $u \in V$, $|\Phi(u) - \Psi(u)| \leq |\Phi(u) - \Upsilon(u)|$. By (UB) it then follows that $\#(\Phi, \Psi) \preceq \#(\Phi, \Upsilon)$.

(2) Suppose $\#(\Phi, \Psi) \preceq \#(\Phi, \Upsilon)$. Then it follows that $\sum_{v \in V} |\Phi(v) - \Psi(v)| \leq \sum_{v \in V} |\Phi(v) - \Upsilon(v)|$ by (LB). And this means that $\#^{\min}(\Phi, \Psi) \preceq \#^{\min}(\Phi, \Upsilon)$.

$\square$

Theorem 1. Consider any compatibility measure $\sqsubseteq$ satisfying (Ref), (Sym), (UB) and (LB).

(1) Δ_{max} satisfies (ISP) and (IAP).

(2) $\Delta_{\min 2}$ satisfies (IAP) but does not satisfy (ISP).

Proof:

(1) For (ISP), let $E = [\Phi_1, \ldots, \Phi_n]$, let $F = [\Psi_1, \ldots, \Psi_n]$ and let $G = rep(E, \{i\}, F)$ for some $i \in \mathcal{I}^+(n)$. If we can show that for every $v \in V$, $|\Delta_{\max}(E)(v) - \Phi_i(v)| \leq |\Delta_{\max}(G)(v) - \Phi_i(v)|$, then the result follows from (UB). So pick a $v \in V$. If $\Delta_{\max}(G)(v) \geq \Delta_{\max}(E)(v)$ then, since $\Phi_i(v) \leq \Delta_{\max}(E)(v)$, we have that $|\Delta_{\max}(E)(v) - \Phi_i(v)| \leq |\Delta_{\max}(G)(v) - \Phi_i(v)|$. On the other hand, if $\Delta_{\max}(G)(v) < \Delta_{\max}(E)(v)$, it must have been the case that $\Phi_i(v) = \Delta_{\max}(E)(v)$, and so it follows that $|\Delta_{\max}(E)(v) - \Phi_i(v)| < |\Delta_{\max}(G)(v) - \Phi_i(v)|$.

For (IAP), let $E = [\Phi_1, \ldots, \Phi_n]$ and let $G = rem(E, \{i\})$ for some $i \in \{1, \ldots, n\}$. If we can show that for every $v \in V$, $|\Delta_{\max}(E)(v) - \Phi_i(v)| \leq |\Delta_{\max}(G)(v) - \Phi_i(v)|$, then the result follows from (UB). So pick a $v \in V$. Observe that it has to be the case that $\Delta_{\max}(G)(v) \leq \Delta_{\max}(E)(v)$. If $\Delta_{\max}(G)(v) = \Delta_{\max}(E)(v)$ then $|\Delta_{\max}(E)(v) - \Phi_i(v)| = |\Delta_{\max}(G)(v) - \Phi_i(v)|$ and we are done. On the other hand, if $\Delta_{\max}(G)(v) < \Delta_{\max}(E)(v)$, it must have been the case that $\Phi_i(v) = \Delta_{\max}(E)(v)$, and so $|\Delta_{\max}(E)(v) - \Phi_i(v)| < |\Delta_{\max}(G)(v) - \Phi_i(v)|$.

(2) For (ISP), let $E = [\Phi_1, \Phi_2]$, $F = [\Psi_1, \Psi_2]$, $\Phi_1(v) = 1$ for every $v \in V$, $\Phi_2(v) = 2$, for every $v \in V$, $\Psi_1(v) = 0$ for every $v \in V$, and $G = rep(E, \{1\}, F)$. That is, to obtain G from E, Φ_1 is replaced by

Ψ_1. It is easily verifiable that $\Delta_{\min 2}(G) \sqsubset_{\Phi_1} \Delta_{\min 2}(E)$, and so (ISP) is not satisfied. For (IAP), let $E = [\Phi_1, \ldots, \Phi_n]$, and $G = rem(E, \{i\})$ for some $i \in \mathcal{I}^+(n)$. We need to show that $|\Delta_{\min 2}(E)(v) - \Phi_i(v)| \leq |\Delta_{\min 2}(G)(v) - \Phi_i(v)|$. So pick a $v \in V$. If $\Phi_i(v) = \Phi_j(v)$ for every $i, j \in \mathcal{I}^+(n)$, then $\Delta_{\min 2}(E)(v) = \Delta_{\min 2}(G)(v)$ and so the result follows immediately. So we suppose that this is not the case and consider three cases. For the first case, if $\Delta_{\min 2}(G)(v) = \Delta_{\min 2}(E)(v)$, the result follows immediately. For the second case, suppose that $\Delta_{\min 2}(G)(v) > \Delta_{\min 2}(E)(v)$. Then, by the definition of $\Delta_{\min 2}$, it has to be the case that $\Phi_i(v) = \min\{\Phi_j(v) \mid 1 \leq j \leq n\}$, which means that $\Phi_i(v) \leq \Delta_{\min 2}(E)(v)$, from which the result follows. And for the third case, suppose that $\Delta_{\min 2}(G)(v) < \Delta_{\min 2}(E)(v)$. Then it has to be the case that $\Phi_i(v) \geq \Delta_{\min 2}(E)(v)$, from which the result follows immediately.

□

Theorem 2. Consider the compatibility measure $\sqsubseteq^{\max}$.

(1) $\Delta_{\max}$ satisfies (SP) and therefore also (CP), (AP) and (ICP).

(2) $\Delta_{\min 2}$ satisfies (AP). It does not satisfy (ICP), and therefore satisfies neither (CP) or (SP).

Proof:

(1) Let $E = [\Phi_1, \ldots, \Phi_n]$, let $F = [\Psi_1, \ldots, \Psi_n]$, and let $G = rr(E, \mathcal{J}, F, \mathcal{K})$ for some $\mathcal{I}, \mathcal{J}, \mathcal{K} \subseteq \{1, \ldots, n\}$ such that $\mathcal{K} = \mathcal{I} \setminus \mathcal{J}$. Assume that $\Delta_{\max}(G) \sqsubseteq^{\max}_{\Phi_i} \Delta_{\max}(E)$ for every $i \in \mathcal{I}$ and that $\Delta_{\max}(G) \sqsubset^{\max}_{\Phi_j} \Delta_{\max}(E)$ for some $j \in \mathcal{I}$. That is, $|\Delta_{\max}(G)(v) - \Phi_i(v)| \leq |\Delta_{\max}(E)(v) - \Phi_i(v)|$ $\forall v \in V$, $\forall i \in \mathcal{I}$ and $|\Delta_{\max}(G)(w) - \Phi_i(w)| < |\Delta_{\max}(E)(w) - \Phi_j(w)|$ for some $w \in V$ and some $j \in \mathcal{I}$. Observe firstly that, because $|\Delta_{\max}(G)(w) - \Phi_j(w)| < |\Delta_{\max}(E)(w) - \Phi_j(w)|$, it has to be the case that $\Delta_{\max}(G)(w) < \Delta_{\max}(E)(w)$. But there is some $k \in \mathcal{I}$ such that $\Delta_{\max}(E)(w) = \Phi_k(w)$, which means that $|\Delta_{\max}(G)(w) - \Phi_k(w)| > |\Delta_{\max}(E)(w) - \Phi_k(w)|$. It has to be the case that $k \in \mathcal{I}$ since $\Delta_{\max}(G)(w) < \Delta_{\max}(E)(w)$. But this contradicts the supposition that $|\Delta_{\max}(G)(v) - \Phi_i(v)| \leq |\Delta_{\max}(E)(v) - \Phi_i(v)|$ $\forall v \in V$, $\forall i \in \mathcal{I}$.

(2) For (ICP), consider the counterexample in theorem 1 which shows that $\Delta_{\min 2}$ does not satisfy (ISP). For (AP), let $E = [\Phi_1, \ldots, \Phi_n]$ and let $G = rem(E, \mathcal{I})$ for some $\mathcal{I} \subseteq \mathcal{I}^+(n)$. Assume that $\Delta_{\min 2}(G) \sqsubseteq^{\max}_{\Phi_i} \Delta_{\min 2}(E)$ for every $i \in \mathcal{I}$ and that $\Delta_{\min 2}(G) \sqsubset^{\max}_{\Phi_j} \Delta_{\min 2}(E)$ for some $j \in \mathcal{I}$. That is, $|\Delta_{\min 2}(G)(v) - \Phi_i(v)| \leq |\Delta_{\min 2}(E)(v) - \Phi_i(v)|$ $\forall v \in V$, $\forall i \in \mathcal{I}$ and $|\Delta_{\min 2}(G)(w) - \Phi_j(w)| < |\Delta_{\min 2}(E)(w) - \Phi_j(w)|$ for some $w \in V$ and some $j \in \mathcal{I}$. Observe firstly that this means that $\Phi_j(w) \neq \Delta_{\min 2}(E)(w)$. Now we consider the remaining two cases. For the first case, suppose that $\Phi_j(w) = \Delta_{\min 2}(E)(w) - 1$. Then it has to be the

case that $\Phi_j(w) = \Delta_{\min2}(G)(w)$ and that there is some $k \in \mathcal{I}$ such that $\Phi_k(w) \geq \Delta_{\min2}(E)(w)$. But this means that $|\Delta_{\min2}(E)(w) - \Phi_k(w)| < |\Delta_{\min2}(G)(w) - \Phi_k(w)|$, contradicting the supposition that $|\Delta_{\min2}(G)(v) - \Phi_i(v)| \leq |\Delta_{\min2}(E)(v) - \Phi_i(v)|\ \forall v \in V, \forall i \in \mathcal{I}$. For the remaining case, suppose that $\Phi_j(w) > \Delta_{\min2}(E)(w)$. Then it has to be the case that $\Delta_{\min2}(G)(w) > \Delta_{\min2}(E)(w)$. Now, by definition of $\Delta_{\min2}$ there is a $k \in \mathcal{I}$ such that $\Phi_k(w) \leq \Delta_{\min2}(E)(w)$. And from this it follows that $|\Delta_{\min2}(E)(w) - \Phi_k(w)| < |\Delta_{\min2}(G)(w) - \Phi_k(w)|$, again contradicting the supposition that $|\Delta_{\min2}(G)(v) - \Phi_i(v)| \leq |\Delta_{\min2}(E)(v) - \Phi_i(v)|$ $\forall v \in V, \forall i \in \mathcal{I}$.

□

References

Andreka, H., Ryan, M. D., & Schobbens, P.-Y. (2001). Operators and laws for combining preference relations. *Journal of Logic and Computation*, *12*(1), 13–53.

Arrow, K. J. (1963). *Social choice and individual values (2nd edition)*. Wiley, New York.

Barberá, S., Dutta, B., & Sen, A. (2000). Strategyproof Social Choice Correspondences. *Journal of Economic Theory, forthcoming*.

Benferhat, S., Dubois, D., Kaci, S., & Prade, H. (2000). Encoding information in possibilistic logic: A general framework for rational syntactic merging. In Horn, W. (Ed.), *ECAI 2000. 14th European Conference on Artificial Intelligence*, Amsterdam. IOS Press.

Benferhat, S., Dubois, D., Prade, H., & Williams, M.-A. (1999). A Practical Approach to Fusing Prioritized Knowledge Bases. In *Progress in Artificial Intelligence*, Vol. 1695 of *Lecture Notes In Artificial Intelligence*, pp. 222–236, Berlin. Springer-Verlag.

Benoit, J.-P. (2000). Strategyproofness for Lotteries and When Ties are Permitted. Unpublished manuscript, Department of Economics, New York University.

Booth, R. (2002). Social contraction and belief negotiation. In *Proceedings of KR'02*, pp. 375–384.

Borsuk, K., & Szmielew, W. (1960). *Foundations of Geometry*. North-Holland, Amsterdam.

Dubois, D., & Koning, J. (1991). Social choice axioms for fuzzy set aggregation. *Fuzzy sets and systems*, *43*, 257–274.

Dubois, D., Lang, J., & Prade, H. (1994). Possibilistic logic. In *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol. 3, pp. 439–513.

Gärdenfors, P. (2000). *Conceptial Spaces: The Geometry of Thought*. The MIT Press, Cambridge, Massachusetts.

Gärdenfors, P., & Rott, H. (1995). Belief revision. In Gabbay, D. M., Hogger, C. J., & Robinson, J. A. (Eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming: Epistemic and Temporal Reasoning*, Vol. 4, pp. 35–132. Clarendon Press, Oxford.

Gibbard, A. (1973). Manipulation of Voting Schemes: A General Result. *Econometrica*, *41*(4), 587–601.

Katsuno, H., & Mendelzon, A. (1991). Propositional knowledge base revision and minimal change. *Artificial Intelligence*, *52*, 263–294.

Kelly, J. S. (1978). *Arrow impossibility theorems*. Series in economic theory and mathematical economics. Academic Press, New York.

Konieczny, S., & Pino-Pérez, R. (1998). On the logic of merging. In Cohn, A. G., Schubert, L., & Shapiro, S. C. (Eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixth International Conference (KR '98)*, pp. 488–498, San Francisco, California. Morgan Kaufmann.

Konieczny, S., & Pino-Pérez, R. (1999). Merging with Integrity Constraints. In *Proceedings of the Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, ECSQARU'99. Lecture Notes in Artificial Intelligence*, Vol. 1638, pp. 233–2448.

Konieczny, S., & Pino-Pérez, R. (2002). Merging information under constraints: a qualitative framework. *Journal of Logic and Computation*, *12*(5), 773–808.

Lafage, C., & Lang, J. (2000). Logical representation of preferences for group decision making. In *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning (KR 2000), Breckenridge, Colorado, USA, 12-15 April 2000*, San Mateo, CA. Morgan Kaufmann.

Lin, J., & Mendelzon, A. O. (1999). Knowledge base merging by majority. In Pareschi, R., & Fronhoefer, B. (Eds.), *Dynamic Worlds: From the Frame Problem to Knowledge Management*. Kluwer Academic Publishers.

May, K. (1952). A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica 20*, *20*, 680–684.

Maynard-Reid, P., & Lehmann, D. (2000). Representing and Aggregating Conflicting Beliefs. In *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning (KR 2000), Breckenridge, Colorado, USA, 12-15 April 2000*, pp. 153–164, San Mateo, CA. Morgan Kaufmann.

Maynard-Reid, P., & Shoham, Y. (1998). From Belief Revision to Belief Fusion. In *Proceedings of the Third Conference on Logic and the Foundations of Game and Decision Theory (LOFT3), ICER, Torino,* *Italy,December 17-20, 1998*.

Meyer, T. (2000). Merging Epistemic States. In Mizoguchi, R., & Slaney, J. (Eds.), *PRICAI 2000: Topics in Artificial Intelligence*, Vol. 1886 of *Lecture Notes In Artificial Intelligence*, pp. 286–296, Berlin. Springer-Verlag.

Meyer, T. (2001). On the semantics of combination operations. *Journal of Applied Non-Classical Logics*, *11*(1-2), 59–84.

Meyer, T., Ghose, A., & Chopra, S. (2001a). Context-based merging. In *Common Sense 2001: Fifth symposium on logical formalizations of commonsense reasoning*.

Meyer, T., Ghose, A., & Chopra, S. (2001b). Social choice, merging and elections. In *Proceedings of the Sixth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty: ECSQARU 2001, LNAI 2143*, pp. 466–477.

Meyer, T., Ghose, A., & Chopra, S. (2001c). Syntactic representations of semantic merging operations. In *IJCAI'01 Workshop on Inconsistency in Data and Knowledge*.

Revesz, P. Z. (1993). On the Semantics of Theory Change: Arbitration between Old and New Information. In *Proceedings PODS '93, 12th ACM SIGACT SIGMOD SIGART Symposium on the Principles of Database Systems*, pp. 71–82.

Roy, B., & Bouyssou, D. (1993). Aide multicritere a la decision..

Satterthwaite, M. (1973). *The Existence of a Strategy Proof Voting Procedure: a Topic in Social Choice Theory*. Ph.D. thesis, University of Wisconsin.

Satterthwaite, M. (1975). Strategy-proofness and Arrow's Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions. *Journal of Economic Theory*, *10*(2), 187–217.

Sen, A. (1970). *Collective Choice and Social Welfare*. Holden Day, San Francisco.

Sen, A. (1986). Social choice theory. In Arrow, K. J., & Intriligator, M. D. (Eds.), *Handbook of Mathematical Economics*, Vol. III, chap. 22, pp. 1073–1181. Elsevier Science Publishers.

Spohn, W. (1988). Ordinal conditional functions: A dynamic theory of epistemic states. In Harper, W. L., & Skyrms, B. (Eds.), *Causation in Decision: Belief, Change and Statistics: Proceedings of the Irvine Conference on Probability and Causation: Volume II*, Vol. 42 of *The University of Western Ontario Series in Philosophy of Science*, pp. 105–134, Dordrecht. Kluwer Academic Publishers.

Williams, M.-A. (1995). Iterated theory base change: a computational model. In *IJCAI-95. Proceedings of the 14th International Joint Conference on Artificial Intelligence held in Montreal, PQ, August 20–25, 1995*, Vol. 2, pp. 1541–1547, San Francisco, CA. International Joint Conference on Artificial Intelligence, Morgan Kaufmann.

Zadeh, L. A. (1978). Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, *1*, 1–28.