Context-sensitive merging

Thomas Meyer and **Aditya Ghose**
Decision Systems Laboratory
Department of Information Systems
University of Wollongong
Wollongong, NSW 2522, Australia
e-mail: {tmeyer,aditya}@uow.edu.au

Samir Chopra
Knowledge Systems Group
Department of Artificial Intelligence
University of New South Wales
Sydney, NSW 2052, Australia
e-mail: schopra@cse.unsw.edu.au

Abstract

Intelligent agents have to be able to merge (possibly conflicting) pieces of information obtained from different sources to produce a coherent picture of the world. In this paper we propose a model for context-sensitive merging. We show that the success of the model depends, to a large extent, on the availability of suitable merging operations on epistemic states; structures in the spirit of [13]. As a result, we investigate the link between such merging operations and the aggregation operations studied in social choice theory [1]. We show that Arrow's impossibility result disappears when recast into the framework that we employ.

1 Introduction

Intelligent agents have to be able to take differing points of view into account to produce a coherent view of the world. In the artificial intelligence literature, this is *merging* [5; 10]. Research on merging usually assumes a list or multiset of pieces of information (of some kind) and provides operations producing a new piece of information — the result of merging the input. Information is usually represented as a sentence in an appropriate logic where, in the spirit of [4], the understanding is that the sentence represents the set of all sentences entailed by it. Such a sentence is referred to as a *knowledge base*. In [5] a general framework is proposed for this kind of merging by providing a set of postulates that every merging operation have to satisfy. In [10] it is argued that a more general framework is needed and that merging ought to be performed on *epistemic states* instead of knowledge bases. Epistemic states are structures in the spirit of [13]. An epistemic state contains, not only the beliefs of an agent, but also an indication of the strength of these beliefs. Merging operations of different kinds have been proposed by various authors. See, for example, [2; 7; 11; 5; 6; 3; 10; 9; 8].

Social choice theory [1; 12] is a research area which addresses problems that are similar to those in merging. Its aims are to provide an appropriate compact representation, an aggregation if you wish, of the preferences of all members in a society.

In this paper we provide a model for context-sensitive merging. The model is described from the point of view of a single agent but is structured so as to facilitate the development of a more general multi-agent model. The idea is to use the context under which information is requested to determine the reliability of sources, *of which the agent itself is one*. Information with the same reliability is combined, after which prioritised merging is applied to pieces of information corresponding to each level of reliability. We show that the success of our model depends, to a large extent, on the availability of appropriate merging operations on epistemic states. As a result, we investigate the link between merging and the aggregation operations studied in social choice theory. We show that a recasting of the Arrow postulates into our framework ensures that Arrow's celebrated impossibility result [1] disappears. Furthermore, our results indicate that some of the properties held dear in social choice theory may not be applicable in the merging context.

The structure of the paper is as follows. Section 2 is a brief review of the basic framework for merging presented in [10]. We address some shortcomings and suggest appropriate modifications. In section 3 we present our model for context-sensitive merging. Section 4 considers the relationship between social choice theory and merging. Section 5 summarises the main results in this paper and points to future research.

For the rest of this paper we assume a finitely generated propositional language L closed under the usual propositional connectives and with a classical model-theoretic semantics. U is the set of interpretations of L and $M(\alpha)$ is the set of models of $\alpha \in L$. Classical entailment is denoted by $\models$. We refer to the *range* of a function f as $ran(f)$.

2 Merging and epistemic states

In this section we provide a brief overview of the basic approach to merging presented in [10]. Meyer argues that merging ought to be conducted on the level of epistemic states and provides a plausibility ranking of the elements of U by assigning to each of them a natural number; the lower the number assigned to an interpretation, the more plausible it is deemed to be.

Definition 2.1 An *epistemic state* Φ is a function from U to $\mathbb{N}$. Given an epistemic state Φ, the *knowledge base associated* *with* Φ is some $\phi \in L$ such that $M(\phi) = \{u \mid \Phi(u) = 0\}$. An epistemic state is *consistent* iff its associated knowledge base is consistent. ■

Epistemic states are structures in the spirit of Spohn [13]. An epistemic state contains, not only the current beliefs of an agent (in the form of a knowledge base), but also information about the level of entrenchment attached to these beliefs. In particular, the current beliefs of an agent are those sentences satisfied by *all* the interpretations assigned the (best) plausibility level of 0. Syntactically, the current beliefs are represented by the knowledge base associated with an epistemic state. We sometimes abuse notation by writing Φ when we refer to the knowledge base associated with Φ. An agent is assumed to have a finite number of sources. Each source provides information about its beliefs in the form of an epistemic state. The implicit assumption is that all sources use the same language. If different sources use different languages we assume that there is some appropriate way of translating to a common language. While this is an important question in itself we adopt the position that such a translation process should precede any attempt at merging.

Another implicit assumption of Meyer's is that different sources use the same *scale* when deciding on the plausibility of interpretations. This scale is simply determined by the highest value assigned to any interpretation by any of the sources concerned. An easy way of removing this restriction is to require of each source s to specify its own personal worst level of plausibility P_s. Such a commitment does not mean that s *has* to rank at least one interpretation at P_s; it simply means that this is the worst level of plausibility that s would ever consider attributing to any interpretation. Allowing for personal levels of worst plausibility ensures that a source is able to decide on its own preferred level of granularity. A commonly agreed upon worst level of plausibility P with a granularity level acceptable to everyone is then easily obtained. We assume the existence of such a P.

Definition 2.2 An epistemic state Φ is *P-capped* iff $\Phi(u) \leq P$ for every $u \in U$. ■

Meyer defines merging as an operation that maps a finite list (or sequence) of epistemic states to an epistemic state. The principal reason for using a list instead of a set is to cater for cases in which different sources contribute identical epistemic states. But this can also be achieved more directly by tagging each epistemic state with the source from which it is obtained. We follow the latter approach.

Definition 2.3 Let S be a set of sources. We do not assume any predefined structure on sources. A *sourced* epistemic state is a pair of the form (Φ, s) where Φ is an epistemic state and $s \in S$. ■

A sourced epistemic state (Φ, s) is (in)consistent iff Φ is. Intuitively (Φ, s) is the assertion that Φ represents the beliefs of source s. A sourced epistemic state is P-capped iff Φ is.

Definition 2.4 An *epistemic set* E is a finite non-empty set of sourced epistemic states in which the same source does not occur twice. ■

E is P-capped iff all of its elements are. We let $S(E) = \{s \mid (\Phi, s) \in E\}$. That is, $S(E)$ is the set of sources occurring in E.

Definition 2.5 A P-capped merging operation Δ relative to a set of sources S is a function from the set of all P-capped epistemic sets E (such that $S(E) \subseteq S$) to the set of all epistemic states. ■

Observe that a P-capped merging operation maps a P-capped epistemic set to *any* epistemic state, not necessarily a P-capped epistemic state. The reason for this is that new information may prompt an agent to modify its preferred level of granularity (and hence its personal worst level of plausibility). As a result, the commonly agreed upon worst level of plausibility, P, might change after merging. A proper description of *how* an agent ought to modify its preferred level of granularity is beyond the scope of this paper.

2.1 Formal tools

A description of Meyer's proposed framework necessitates the provision of a range of formal tools in this section.

For a P-capped epistemic state Φ, let $\min(\Phi) = \min\{\Phi(u) \mid u \in U\}$. We denote by $seq(E)$ the set of all sequences of length $|E|$ of natural numbers ranging from 0 to P, and by $seq_{\leq}(E)$ the subset of $seq(E)$ of all sequences that are ordered non-decreasingly.

Given a set S of sources, a *reliability indicator* ri relative to S is a function from S to $\mathbb{N}$. Intuitively ri ranks the elements of S so that a lower ranking corresponds to a higher reliability. If ri is injective we say that ri is *complete*. A complete reliability indicator places a linear order of reliability on sources and therefore orders them completely. For a *complete* reliability indicator ri, an epistemic set E and $u \in U$ we let $seq^E_{ri}(u)$ be the sequence obtained by placing, for every $(\Phi, s), (\Psi, t) \in E$, $\Phi(u)$ before $\Psi(u)$ in the sequence iff $ri(s) < ri(t)$. In other words, $seq^E_{ri}(u)$ is made up of the values assigned to u by the different sources in E, and with these values ordered according to the reliability ranking as determined by ri. We let $seq^E_{\leq}(u)$ be the sequence $seq^E_{ri}(u)$ ordered non-decreasingly. Clearly $seq^E_{ri}(u) \in seq(E)$ and $seq^E_{\leq}(u) \in seq_{\leq}(E)$. Given any set seq of finite sequences of natural numbers and a total preorder $\sqsubseteq$ on seq, we define the function $\Omega^{seq}_{\sqsubseteq} : seq \to \{0, \ldots, |seq| - 1\}$ by assigning consecutive natural numbers to the elements of seq in the order imposed by $\sqsubseteq$, starting by assigning 0 to the elements lowest down in $\sqsubseteq$. We denote the *lexicographic* ordering on seq by $\sqsubseteq_{lex}$.

Definition 2.6 Given a set of P-capped epistemic states $\mathcal{E}$, the *normalisation* function η from $\mathcal{E}$ to $\mathcal{E}$ is defined as follows: $\eta(\Phi)(u) = \Phi(u) - \min(\Phi)$ for every $u \in U$. ■

Normalisation is necessary to ensure that a merging operation produces consistent epistemic states. The following example clarifies the use of normalisation.

Example 2.1 The merging operation $\Delta_{\max}$, one of the operations proposed in [10], first creates an intermediate epistemic state $\Phi^E_{\max}$ when merging the epistemic set E, by assigning to each interpretation the maximum value assigned to it by any of the sources in E. The final epistemic state is then obtained by normalising $\Phi^E_{\max}$. That is, $\Delta_{\max}(E) = \eta(\Phi^E_{\max})$.

As an example of applying $\Delta_{\max}$, let $\Phi_1(u) = \Phi_2(v) = 0$, $\Phi_1(v) = 1$, $\Phi_2(u) = 2$, and let $\Phi_1(x) = \Phi_2(x) = 3$ for all other interpretations. Now, let $E = \{(\Phi_1, s_1), (\Phi_2, s_2)\}$. It is easily verified that $\Phi^E_{\max}(v) = 1$, $\Phi^E_{\max}(u) = 2$ and $\Phi_{\max}(x) = 3$ for all other interpretations. Since the minimum value assigned to an interpretation is 1, normalisation involves the subtraction of 1 from each of the values assigned to the different interpretations, and so $\Delta_{\max}(E)(v) = 0$, $\Delta_{\max}(E)(u) = 1$ and $\Delta_{\max}(E)(x) = 2$ for all other interpretations. ∎

The merging operation $\Delta_{\max}$ used in example 2.1 is one of a whole range of available merging operations with differing intuitions. It is instructive to observe that all the merging operations proposed in [10] involve normalisation. In fact, the merging operations on knowledge bases in [5] also involve normalisation, although this is not mentioned explicitly in the paper.

2.2 A basic framework for merging

Meyer considers the following basic properties on merging operations.

(E0) $\Delta(\{(\Phi, s)\}) = \Phi$

(E1) $\Delta(E) \vDash \bot$ iff for every $(\Phi, s) \in E, \Phi \vDash \bot$

(E2) If $\Phi(u) = \Psi(u)$ for every $(\Phi, s), (\Psi, t) \in E$ and $seq^E_{\leq}(u) \sqsubset_{lex} seq^E_{\leq}(v)$, then $\Delta(E)(u) < \Delta(E)(v)$

(E3) If $\Phi(u) \leq \Phi(v)$ for every $(\Phi, s) \in E$ then $\Delta(E)(u) \leq \Delta(E)(v)$

(E4) If $\Delta(E)(u) \leq \Delta(E)(v)$ then $\Phi(u) \leq \Phi(v)$ for some $(\Phi, s) \in E$

We have modified (E0) and (E1) slightly to allow for the possibility of obtaining inconsistent epistemic states after merging has taken place. The idea is to obtain (retain, in fact) inconsistency after merging iff all sourced epistemic states to be merged are inconsistent. The intuitive justification for this is that obtaining a consistent epistemic state when merging only makes sense if at least one of the agent's sources currently has a consistent view of the world.

The next property we discuss is important for our model of merging presented in section 3. Two epistemic sets E_1 and E_2 are *epistemic-equivalent*, denoted as $E_1 \approx E_2$, iff $\{\Phi \mid \exists s \in S(E_1) \text{ s.t. } (\Phi, s) \in E_1\} = \{\Phi \mid \exists s \in S(E_2) \text{ s.t. } (\Phi, s) \in E_2\}$. Epistemic-equivalence holds whenever two epistemic sets are identical iff the sources of information are ignored.

(Comm) $E_1 \approx E_2$ implies $\Delta(E_1) = \Delta(E_2)$

(Comm) states that one should not take into account the source from which information was obtained. Such a requirement only makes sense if we assume all sources to be equally reliable. Meyer proposes the following merging operation which does not satisfy (Comm).

Definition 2.7 Let ri be a *complete* reliability indicator relative to a set of sources S (as defined in section 2.1). Let $\Phi^{(E,ri)}_{lex}(u) = \Omega^{seq(E)}_{\sqsubseteq_{lex}}(seq^E_{ri}(u))$. Then $\Delta^{ri}_{lex}(E) = \eta(\Phi^{(E,ri)}_{lex})$. ∎

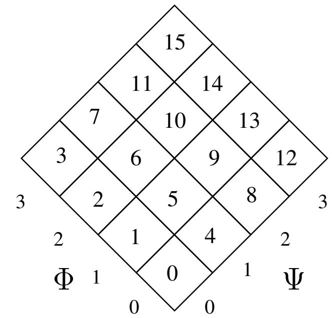

Figure 1: A representation of $\Phi^{(E,ri)}_{lex}$ which is used to obtain the merging operation Δ^{ri}_{lex}. Levels from the bottom left to top right represent the different values assigned to interpretations by Φ, starting at 0. For example, interpretations assigned a value of 1 by Φ will be placed in the cells numbered 4, 5, 6 and 7. Similarly, levels from bottom right to top left represent the different values assigned to interpretations by Ψ, starting at 0. For example, interpretations assigned the value 2 by Ψ will be placed in the cells numbered 2, 6, 10 and 14. Each particular cell therefore represents a set of interpretations based on the level of plausibility assigned to them by Φ and Ψ. For example, the cell containing a 6 represents all interpretations which are assigned a value of 1 by Φ and a value of 2 by Ψ. The number in a cell represents the value that $\Phi^{(E,ri)}_{lex}$ assigns to interpretations contained in that cell.

Δ^{ri}_{lex} is a lexicographic ordering of the interpretations based on the reliabilities on the sources imposed by ri. Figure 1 contains a graphical representation of $\Phi^{(E,ri)}_{lex}$ for an example where $E = \{(\Phi, s), (\Psi, t)\}$ and where $ri(s) < ri(t)$ (i.e. s is more reliable than t). Observe that Δ^{ri}_{lex} is obtained by normalising $\Phi^{(E,ri)}_{lex}$. It turns out that Δ^{ri}_{lex} will play an important role in our model for context-sensitive merging.

3 A model for context-sensitive merging

Our model focuses on the ability of a single agent to perform context-sensitive merging but is structured in a way that enables it to be incorporated into a model for multi-agent merging. Informally we have the following scenario in mind. An agent receives information from other agents in the form of epistemic states and *merges this with its own beliefs* (an epistemic state) to obtain a coherent view of the world. This is a generalisation of classical belief revision where an agent receives information in the form of a single sentence which takes precedence over its current beliefs. It can also be viewed as the basis for a theory of communication and learning. It differs from the setting considered in the merging literature thus far: simply merging the information obtained from a number of sources into a new knowledge base or epistemic state.

Merging is taken to be context-sensitive in two ways. Firstly, in different contexts, the reliability of a particular source may vary. If I have experts on music and political matters among my sources it makes sense to place a higher reliability on the music experts than on political experts when the context determines that information related to music is relevant, and vice versa when politics is the topic under scrutiny.

Secondly, the method of performing merging may change depending on the context. Consider the case of arbitration and majority operations; two types of merging operations studied in the literature on merging. Arbitration operations involve some aspects of give-and-take while majority operations are ones in which the opinions of the majority of sources hold sway. In some contexts it will be more appropriate to employ an arbitration operation while in others it will be more sensible to apply a majority operation. Contexts can thus be seen as *epistemic scenarios*.

To achieve context-sensitivity an agent tags each bit of information with the source from which it is obtained. One of these sources is the agent itself. Information tagged with the agent itself as the source represents the current beliefs of the agent. The agent also has available to it a whole suite of merging operations, denoted by $\mathcal{D}$, as well as a set of contexts, denoted by $\mathcal{C}$ (we do not assume any predefined structure on contexts). A particular context determines the reliability of the different sources as well as the appropriate method of merging to apply. This meta-knowledge about the appropriate merging operation and the reliability of sources is then combined in a rational way to obtain the resulting epistemic state: the agent's new beliefs about the world. Observe that in this setup an agent sees its own beliefs *as just one of many pieces of information available to it*. Its judgement of its own reliability relative to those of the other sources is based on the context. To be able to determine the reliability of sources and pick an appropriate merging operation we assume that an agent is equipped with a *context applicator*. Formally, a context applicator is a function that maps an element of $\mathcal{C}$ to a pair (Δ, ri) where $\Delta \in \mathcal{D}$ and ri is a reliability indicator (cf. section 2.1).

Definition 3.1 A merging agent A is a 7-tuple $A = (S, a, P, \Phi, \mathcal{C}, \mathcal{D}, ca)$ where S is a non-empty set of sources, $a \in S$ is the agent itself viewed as one of its own sources, P is a positive integer denoting the commonly agreed upon worst level of plausibility for the sources in S, Φ is a P-capped epistemic state denoting the beliefs of the agent, $\mathcal{C}$ is a non-empty set of contexts, $\mathcal{D}$ is a non-empty set of P-capped merging operations relative to S, and ca is a context applicator. ■

The context-sensitive merging operation $\bowtie$ that we propose is a function that maps an agent, a context and an epistemic set to an epistemic state.

Definition 3.2 Let $A = (S, a, P, \Phi, \mathcal{C}, \mathcal{D}, ca)$ be an agent, let $c \in \mathcal{C}$, and let E be a P-capped epistemic set such that $S(E) \subseteq S$. Let $ca(c) = (\Delta, ri)$ where $\Delta \in \mathcal{D}$, and $\forall i \in ran(ri)$, let $E_i = \{(\alpha, s) \mid (\alpha, s) \in E \text{ and } ri(s) = i\}$. For every $i \in ran(ri)$, let $\Phi_i = \Delta(E_i)$ and $s_i = S(E_i)$, and let $F = \bigcup_{i \in ran(ri)} \{(\Phi_i, s_i)\}$. Define the (complete) reliability indicator ri^F relative to $S(F)$ as follows: $\forall i \in ran(ri)$, $ri^F(s_i) = i$. Then $\bowtie(A, c, E) = \Delta^{ri^F}_{lex}(F)$. ■

Observe that, since ri^F is a *complete* reliability indicator, $\bowtie$ is well-defined. Observe also that $\bowtie$ does not form part of the definition of a merging agent. This is because $\bowtie$ is an operation that is fixed for all agents. A (rudimentary) multi-agent setting could thus be seen as an ordered pair $(\mathcal{A}, \bowtie)$ where $\mathcal{A}$ is a non-empty set of agents and $\bowtie$ is the context-sensitive merging operation employed by all agents in $\mathcal{A}$.

The following example clarifies the definition of $\bowtie$.

Example 3.1 Consider the agent $A = (\{a, b, c, d, e, f\}$, a,20,Φ_a,$\{c_1, c_2\}$, $\{\Delta_1, \Delta_2\}$,ca) where $ca(c_1) = (\Delta_1, ri_1)$ and where $ri_1(a) = ri_1(b) = 1$, $ri_1(c) = 2$, and $ri_1(d) = ri_1(e) = ri_1(f) = 4$. That is, agent A uses a, b, c, d, e and f as sources, where a is the agent itself viewed as one of its own sources, and these sources have all agreed on 20 as the worst level of plausibility on interpretations. Furthermore, A's own beliefs, expressed as a 20-capped epistemic state is Φ_a, A is able to distinguish between two contexts, c_1 and c_2, and is equipped with two methods of merging sourced epistemic states. When context c_1 presents itself, A regards Δ_1 as the appropriate merging operation to apply and views sources a and b as the most reliable, followed by c, and with d, e and f as the least reliable sources. Now, consider the case in which A receives information from sources b, d and e in a situation where context c_1 is applicable. That is, A receives information in the form of a 20-capped epistemic set $E = \{(\Phi_a, a), (\Phi_b, b), (\Phi_d, d), (\Phi_e, e)\}$ where (Φ_a, a) represents its own beliefs. A will apply Δ_1 to the information supplied by a (itself) and b to obtain Φ_1, and it will apply Δ_1 to the information supplied by d and e to obtain Φ_4. That is, $\Phi_1 = \Delta_1(\{(\Phi_a, a), (\Phi_b, b)\})$ and $\Phi_4 = \Delta_1(\{(\Phi_d, d), (\Phi_e, e)\})$. It then constructs the epistemic state F to look as follows: $\{(\Phi_1, \{a, b\}), (\Phi_4, \{d, e\})\}$, with $ri^F(\{a, b\}) = 1$ and $ri^F(\{d, e\}) = 4$. So F is a "super" epistemic set in which the sources are equivalence classes of initial sources *modulo* the reliability ordering obtained from ri_1. The epistemic state $\bowtie(A, c, E)$ is then obtained by applying $\Delta^{ri^F}_{lex}$ to F, which, in this case, boils down to refining Φ_1, the more reliable epistemic state according to ri^F, with Φ_4. ■

It should be clear that the choice of $\mathcal{D}$, the set of applicable merging operations on P-capped epistemic sets, plays a crucial role in the ability of an agent to perform merging in a rational fashion. At this stage we propose that every $\Delta \in \mathcal{D}$ satisfy at least (E0)-(E4) and (Comm). (Comm) is necessary because Δ will only be applied to epistemic sets for which sources are judged to be equally reliable.

4 Social choice theory and merging

As discussed in section 3 our proposed model for a context-sensitive merging agents relies heavily on the availability of merging operations on epistemic sets. Social choice theory [1; 12] is a research area where the problems under scrutiny are quite similar to the problems encountered in merging. Social choice theory is concerned with the aggregation of preferences. 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 setup and our framework for the merging of epistemic sets should be obvious. It is a matter of taking Ω to be U, the set of interpretations, and regarding the preferences as rankings of plausibility. Epistemic states, however, have more structure than the total preorders used to denote preferences, which turns out to be of crucial importance. Since epistemic states map interpretations to the set of natural numbers, they can be used to express a notion of *distance* between interpretations, *something that cannot be accomplished using total preorders*. For example, for an epistemic state Φ, $\Phi(u) = 0, \Phi(v) = 1$ and $\Phi(w) = 3$ expresses the view that the relative plausibility of u and v are closer to each other than the relative plausibility of v and w.

Arguably the most important result in social choice theory is Arrow's (in)famous impossibility theorem [1] which shows that there is no aggregation operation satisfying some intuitively desirable postulates. Our intention in this section is to provide the groundwork for comparisons between social choice theory and merging on epistemic sets. The starting point is the recasting of the original Arrow postulates into our framework.

The first postulate, dubbed *Restricted Range*, requires the result of an aggregation operation to be a total preorder on Ω. In our setup this translates to the requirement that a merging operation produce an epistemic state – something which is built into the definition of merging operations. 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. In our setup this translates to the requirement that merging may be applied to any (P-capped) epistemic set – again, something that is built into the definition of merging operations. The third Arrow postulate, known as the *weak Pareto Principle*, can be phrased as follows for epistemic sets:

(PP) If $\Phi(E)(u) < \Phi(E)(v)$ for every $(\Phi, s) \in E$ then $\Delta(E)(u) < \Delta(E)(v)$

(PP) requires that any aggregation operation should respect those preferences shared by all individuals. In our setup, if all sources agree that an interpretation u is more plausible than an interpretation v, (PP) requires of the result of any merging operation to regard u as more plausible than v as well. It is easily verified that (PP) is the contrapositive of (E4).

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

(IIA) Let E and F be P-capped epistemic sets such that $S(E) = S(F)$. If it is the case that $\Phi(u) \leq \Phi(v)$ iff $\Psi(u) \leq \Psi(v)$ for every $(\Phi, s) \in E$ and $(\Psi, s) \in F$ then it will also be the case that $\Delta(E)(u) \leq \Delta(E)(v)$ iff $\Delta(F)(u) \leq \Delta(F)(v)$

The intuition underlying (IIA) is that, when deciding on the relative ordering of interpretations u and v, (IIA) requires of us to disregard all other interpretations. This intuition can indeed be formalised appropriately as above when preferences (or plausibilities) are expressed as total preorders. However, the intuition does not carry over to our more structured setup in which it is possible to define various degrees of relative plausibility between two interpretations.

Example 4.1 Consider four epistemic states Φ_1, Φ_2, Ψ_1 and Ψ_2 such that $\Phi_1(u) = \Psi_2(v) = 0$, $\Phi_1(v) = \Phi_2(v) = \Psi_1(u) = \Psi_2(u) = 1$, and $\Phi_2(u) = \Psi_1(v) = 2$. Now, let $E = \{(\Phi_1, s), (\Phi_2, t)\}$ and $F = \{(\Psi_1, s), (\Psi_2, t)\}$. It is easily verified that $\Phi_1(u) \leq \Phi_1(v)$ iff $\Psi_1(u) \leq \Psi_1(v)$ and that $\Phi_2(u) \leq \Phi_2(v)$ iff $\Psi_2(u) \leq \Psi_2(v)$. Furthermore, it can be verified that $\Delta_{\max}(E)(v) = 0$, $\Delta_{\max}(E)(u) = 1$, $\Delta_{\max}(F)(u) = 0$, and $\Delta_{\max}(F)(v) = 1$, where $\Delta_{\max}$ is the merging operation discussed in example 2.1. So it is not the case that $\Delta_{\max}(E)(u) \leq \Delta_{\max}(E)(v)$ iff $\Delta_{\max}(F)(u) \leq \Delta_{\max}(F)(v)$. $\Delta_{\max}$ therefore does not satisfy (IIA).

Observe, however, that the ranking of u and v relative to each other are, in both cases, obtained without reference to any of the other interpretations. In fact, it can be shown that the relative ranking of any two interpretions obtained by applying $\Delta_{\max}$ is independent of irrelevant alternatives even though $\Delta_{\max}$ does not satisfy (IIA). ■

Similar results and examples can be obtained for all the merging operations proposed in [10].

What we need, therefore, is a property that adequately captures the intuition underlying (IIA) in our setup. We propose the following alternative to (IIA) as a basic postulate for merging to be added to those set out in section 2.2:

(RIIA) Let E and F be P-capped epistemic sets such that $S(E) = S(F)$ and for every $(\Phi, s) \in E$ and $(\Psi, s) \in F$, $\Phi(u) = \Psi(u)$ and $\Phi(v) = \Psi(v)$. Then $\Delta(E)(u) \leq \Delta(E)(v)$ iff $\Delta(F)(u) \leq \Delta(F)(v)$

The last of the Arrow postulates, known as *Non-Dictatorship*, states that one source should never be able to completely dominate. In our setup this can be phrased as follows for a fixed set of sources S where $|S| > 1$.

(ND) There is no $s \in S$ such that, for every P-capped epistemic set E for which $(\Phi, s) \in E$ for some Φ, it will be the case that $\Phi(u) < \Phi(v)$ implies $\Delta(E)(u) < \Delta(E)(v)$ for every $u, v \in U$.

This is a reasonable property when all sources are regarded as equally reliable, the case in which we are interested. It is easy to see that (ND) follows from (Comm).

Proposition 4.1 *If a merging operation Δ satisfies (Comm) then it will also satisy (ND).*

It turns out that most of the merging operations proposed in [10] satisfy the Arrow postulates when recast in terms of epistemic states.

Proposition 4.2 *The merging operations $\Delta_{\min}$, $\Delta_{\max}$, Δ_{Σ}, $\Delta_{R\Sigma}$, Δ_{Rmin} and Δ_{Rmax} defined in [10] satisfy Restricted Range, Unrestricted Domain, (PP), (RIIA) and (ND).*

With the exception of $\Delta_{R\Sigma}$, which is a refined version of Δ_{Σ}, all of these are generalisations of merging operations on knowledge bases found in the literature. They constitute a repertoire of merging operations with differing but equally valid intuitions associated with them.

By defining merging operations in terms of epistemic states instead of total preorders we have thus circumvented the Arrow impossibility result.

Two other properties that are widely regarded as desirable in the social choice community are of special interest in the merging context. These properties, known as Pareto's conditions, have been proposed in [5] as properties for merging of knowledge bases, and in [11] for weighted model-fitting operations. In our context they can be phrased as follows:

(P1) $\Delta(E_1) \wedge \Delta(E_2) \vDash \Delta(E_1 \cup E_2)$

(P2) If $\Delta(E_1) \wedge \Delta(E_2) \nvDash \perp$ then $\Delta(E_1 \cup E_2) \vDash \Delta(E_1) \wedge \Delta(E_2)$

It turns out that a number of merging operations proposed in [10] do not satisfy (P2). This applies to the merging operations on knowledge bases proposed in [5] as well, which can all be seen as special cases of some of the operations proposed in [10]. In the case of $\Delta_{\max}$, the merging operation discussed in example 2.1, the failure to satisfy (P2) is due to the application of normalisation (cf. definition 2.6). Consequently, the epistemic states obtained by $\Delta_{\max}$ before normalisation takes place *do* satisfy (P2).

In fact, in the case of $\Delta_{\max}$, (P1) and (P2) can be replaced with versions that are more general, and in the case of (P2), stronger as well.

Definition 4.1 For an epistemic state Φ, let $M^i(\Phi) = \{u \mid \Phi(u) \leq i\}$. ■

$M^i(\Phi)$ is thus the set of interpretations for which Φ assigns a value of at most i.

(GP1) $M^i(\Delta(E_1)) \cap M^i(\Delta(E_2)) \subseteq M^i(\Delta(E_1 \cup E_2))$

(GP2) $M^i(\Delta(E_1 \cup E_2)) \subseteq M^i(\Delta(E_1)) \cap M^i(\Delta(E_2))$

Proposition 4.3 *A merging operation Δ that satisfies (GP1) also satisfies (P1). If Δ satisfies (GP2) then it also satisfies (P2).*

When the effects of normalisation are ignored, $\Delta_{\max}$ satisfies (GP1) and (GP2).

Proposition 4.4 *The merging operation $\Delta_{\max}$ defined in [10] satisfies (GP1) and (GP2), and therefore (P1) and (P2) as well, if the normalisation step is not applied.*

Together (GP1) and (GP2) can be seen as a fairly strong *decomposition* principle which states that the merging of an epistemic set E is completely determined by the merging of any two covering subsets of E (where two subsets F and G cover E iff $F \cup G = E$). While such merging operations may form an interesting subclass, it is doubtful whether all rational merging operations are decompositional. To corroborate this conjecture we point out that the merging operation $\Delta_{R\Sigma}$ defined in [10] satisfies neither (GP1) nor (GP2); not even before normalisation.

5 Conclusion

We have provided a general framework for context-sensitive merging. The emphasis is on the viewpoint of a single agent but it was shown how to incorporate this into a multi-agent setting. Our proposal relies on the availability of suitable operations for the merging of epistemic sets. As a result, we have also investigated the link between such merging operations and the aggregation operations studied in social choice theory. In doing so we considered the postulates leading to Arrow's impossibility theorem and showed that, when recast into our setup, the impossibility result disappears. This reformulation also served to suggest an additional basic postulate to be added to the basic framework for the merging of epistemic sets presented in section 2.2. On the other hand, some properties held dear in the social choice community are inappropriate for merging.

We have provided the groundwork for merging in a multi-agent setting but a complete account of the restrictions placed on agents in the multi-agent setting still has to be developed. For example, it seems intuitively desirable to have sources divulge only part of their beliefs, depending on the context and the recipient of the information. In a similar vein, an agent might decide to select only part of the information supplied to it by a source, depending on the relevance of the information in a particular context. Furthermore, there is still disagreement about which merging operations on epistemic sets should be regarded as rational. In pursuing this question it seems useful to build on the basic comparison with social choice theory provided in section 4. Results from that section suggests that, while it is instructive to compare merging and the aggregation operations studied in social choice theory, these constructions are ultimately guided by different intuitions and the similarities should not be overemphasised.

References

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

[2] C. Baral, S. Kraus, and J. Minker. Combining multiple knowledge bases. *IEEE Transactions on Knowledge and Data Engineering*, 3(2):208–220, 1991.

[3] Salem Benferhat, Didier Dubois, Henri Prade, and Mary-Anne Williams. A Practical Approach to Fusing Prioritized Knowledge Bases. In *Progress in Artificial Intelligence*, volume 1695 of *Lecture Notes In Artificial Intelligence*, pages 222–236, Berlin, 1999. Springer-Verlag.

[4] H. Katsuno and A.O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.

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

[6] Sébastien Konieczny and Ramón Pino-Pérez. 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*, volume 1638, pages 233–2448, 1999.

[7] J. Lin. Integration of weighted knowledge bases. *Artificial Intelligence*, 83(2):363–378, 1996.

[8] Pedrito Maynard-Reid and Daniel Lehmann. 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, pages 153–164, San Mateo, CA, 2000. Morgan Kaufmann.

[9] Pedrito Maynard-Reid and Yoav Shoham. 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*, 1998.

[10] Thomas Meyer. Merging Epistemic States. In Riichiro Mizoguchi and John Slaney, editors, *PRICAI 2000: Topics in Artificial Intelligence*, volume 1886 of *Lecture Notes In Artificial Intelligence*, pages 286–296, Berlin, 2000. Springer-Verlag.

[11] P. Z. Revesz. On the semantics of arbitration. *International Journal of Algebra and Computation*, 7(2):133–160, 1997.

[12] Amartya Sen. Social choice theory. In K. J. Arrow and M. D. Intriligator, editors, *Handbook of Mathematical Economics*, volume III, chapter 22, pages 1073–1181. Elsevier Science Publishers, 1986.

[13] Wolfgang Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In William L. Harper and Brian Skyrms, editors, *Causation in Decision: Belief, Change and Statistics: Proceedings of the Irvine Conference on Probability and Causation: Volume II*, volume 42 of *The University of Western Ontario Series in Philosophy of Science*, pages 105–134, Dordrecht, 1988. Kluwer Academic Publishers.