

Who’s the Expert? On Multi-source Belief Change

Joseph Singleton , Richard Booth

Cardiff University

{singletonj1,boothr2}@cardiff.ac.uk

Abstract

Consider the following belief change/merging scenario. A group of information sources gives a sequence of reports about the state of the world at various instances (e.g. different points in time). The true states at these instances are unknown to us. The sources have varying levels of expertise, also unknown to us, and may be knowledgeable on some topics but not others. This may cause sources to report false statements in areas they lack expertise. What should we believe on the basis of these reports? We provide a framework in which to explore this problem, based on an extension of propositional logic with expertise formulas. This extended language allows us to express beliefs about the state of the world at each instance, as well as beliefs about the expertise of each source. We propose several postulates, provide a couple of families of concrete operators, and analyse these operators with respect to the postulates.

1 Introduction

Consider the following belief change scenario in a hospital. We observe the results of a blood test of patient 1, confirming condition X. Assuming the test is reliable, the AGM paradigm (Alchourrón, Gärdenfors, and Makinson 1985) tells us how to revise our beliefs in light of the new information. Dr. A then claims that patient 2 suffers from the same condition, but Dr. B disagrees. Given that doctors specialise in different areas and may make mistakes, who should we trust? Since the *Success* postulate ($\alpha \in K * \alpha$) assumes information is reliable, we are outside the realm of AGM revision, and must instead apply some form of *non-prioritised* revision (Hansson 1999).

Suppose it now emerges that Dr. A had earlier claimed patient 1 did *not* suffer from condition X, contrary to the test results. We now have reason to suspect Dr. A may *lack expertise* on diagnosing X, and may subsequently revise beliefs about Dr. A’s domain of expertise and the status of patient 2 (e.g. by opting to trust Dr. B instead).

While simple, this example illustrates the key features of the belief change problem we study: we consider multiple sources, whose expertise is *a priori* unknown, providing reports on various instances of a problem domain. On the basis of these reports we form beliefs both about the expertise of the sources and the state of the world in each instance.

By including a distinguished *completely reliable* source (the test results in the example) we extend AGM revision. In

some respects we also extend approaches to non-prioritised revision (e.g. selective revision (Fermé and Hansson 1999), credibility-limited revision (Hansson et al. 2001), and trust-based revision (Booth and Hunter 2018)), which assume information about the reliability of sources is known up front. The problem is also related to *belief merging* (Konieczny and Pino Pérez 2002) which deals with combining belief bases from multiple sources; a detailed comparison will be given in Section 7.

Our work is also connected to trust and belief revision, if one interprets trust as *belief in expertise*. As Yasser and Ismail (2021) note in recent work, trust and belief are inexorably linked: we should accept reports from sources we believe are trustworthy, and we should trust sources whose reports turn out to be reliable. Trust and belief should also be revised in tandem, so that we may increase or decrease trust in a source as more reports are received, and revoke or reinstate previous reports from a source as its perceived trustworthiness changes.¹

To unify the trust and belief aspects, we enrich a propositional language with *expertise statements* $E_i(\varphi)$, read as “source i has expertise on φ ”. The output of our belief change problem is then a collection of belief and knowledge sets in the extended language, describing what we *know* and *believe* about the expertise of the sources and the state of the world in each instance. For example, we should *know* reports from the reliable source are true, whereas reports from ordinary sources may only be believed.

Following recent work on logical approaches to expertise (Singleton 2021; Booth and Hunter 2018), we formally model expertise using a partition of propositional valuations for each source. Equivalently, each source has an *indistinguishability equivalence relation* over valuations. A source is an expert on a proposition φ exactly when they can distinguish every φ valuation from every $\neg\varphi$ valuation.² As in Singleton (2021), we also use *soundness formulas* $S_i(\varphi)$, which intuitively say that φ is true *up the expertise of i* . For example, if i has expertise on p but not q , then the conjunction $p \wedge q$ is sound for i whenever p holds, since we can effectively ignore q . Formally, φ is sound for i if the “actual” state of the world is indistinguishable from a φ valuation.

¹ This mutual dependence between trust and belief is also the core idea in *truth discovery* (Li et al. 2016).

ation. Note that expertise does not depend on the “actual” state, whereas soundness does. This provides a crucial link between expertise and truthfulness of information.

We then make the assumption that *sources only report sound propositions*. That is, reports are only false due to sources overstepping the bounds of their expertise. In particular, we assume sources are honest in their reports, and that experts are always right.

Note that in our introductory example, the fact that we had a report from Dr. A on patient 1 (together with reliable information on patient 1) was essential for determining the expertise of Dr. A, and subsequently the status of patient 2. While the patients are independent, reports on one can cause beliefs about the other to change, as we update our beliefs about the expertise of the sources.

In general we consider an arbitrary number of *cases*, which are seen as labels for instances of the domain. For example, a crowdsourcing worker may label multiple images, or a weather forecaster may give predictions for different locations. Each report in the input to the problem then refers to a specific case. Via these cases and the presence of the completely reliable source, we are able to model scenarios where some “ground truth” is available, listing how often sources have been correct/incorrect on a proposition (e.g. the *report histories* of Hunter (2021)). We can also generalise this scenario, e.g. by having only partial information about “previous” cases.

Throughout the paper we make the assumption that *expertise is fixed across cases*: the expertise of a source does not depend on the particular instance of the domain we look at. For instance, the expertise of Dr. A is the same for patient 1 as for patient 2. This is a simplifying assumption, and may rule out certain interpretations of the cases (e.g. if cases represent different points in time, it would be natural to let expertise evolve over time).

Contribution. Our contributions are threefold. First, we develop a logical framework for reasoning about the expertise of multiple sources and the state of the world in multiple cases. Second, we formulate a belief change problem within this framework, which allows us to explore how trust and belief should interact and evolve as reports are received from the various sources. Finally, we put forward several postulates and two concrete classes of operators – with a representation result for one class – and analyse these operators with respect to the postulates.

Paper Outline. In Section 2 we develop the formal framework. Section 3 introduces the problem and lists some core postulates. We give two constructions and specific example operators in Section 4. Section 5 introduces some further postulates concerning belief change on the basis of one new report. An analogue of selective revision (Fermé and Hansson 1999) is presented Section 6. Section 7 discusses related work, and we conclude in Section 8. Proofs are given in

² The relationship between this notion of expertise and *S5 epistemic logic* is explored in a modal logic setting in Singleton (2021), and we revisit this connection in Section 7.

the appendix of the full version of the paper (Singleton and Booth 2022).

2 The Framework

Let \mathcal{S} be a finite set of information sources. For convenience, we assume there is a *completely reliable* source in \mathcal{S} , which we denote by $*$. For example, we can treat our first-hand observations as if they are reported by $*$. Other sources besides $*$ will be termed *ordinary sources*. Let \mathcal{C} be a finite set of *cases*, which we interpret as labels for different instances of the problem domain.

Syntax. There are two levels to our formal language. To describe properties of the world in each case $c \in \mathcal{C}$, we assume a fixed finite set \mathcal{P} of propositional variables, and let \mathcal{L}_0 denote the set of propositional formulas generated from \mathcal{P} using the usual propositional connectives. We use lower case Greek letters (φ, ψ etc) for formulas in \mathcal{L}_0 . The classical logical consequence operator will be denoted by Cn_0 , and \equiv denotes equivalence of propositional formulas.

The extended *language of expertise* \mathcal{L} additionally describes the expertise of the sources, and is defined by the following grammar:

$$\Phi ::= \varphi \mid \Phi \wedge \Phi \mid \neg\Phi \mid E_i(\varphi) \mid S_i(\varphi)$$

where $i \in \mathcal{S}$ and $\varphi \in \mathcal{L}_0$. We introduce Boolean connectives $\vee, \rightarrow, \leftrightarrow$ and \perp as abbreviations. We use upper case Greek letters (Φ, Ψ etc) for formulas in \mathcal{L} . For $\Gamma \subseteq \mathcal{L}$, we write $[\Gamma] = \Gamma \cap \mathcal{L}_0$ for the propositional formulas in Γ .

The intuitive reading of $E_i(\varphi)$ is *source i has expertise on φ* , i.e., i is able to correctly identify the truth value of φ in any possible state. The intuitive reading of $S_i(\varphi)$ is that φ *sound for i to report*: that φ is true up to the expertise of i . That is, the parts of φ on which i has expertise are true. Note that both operators are restricted to propositional formulas, so we will not consider iterated formulas such as $E_i(S_j(\varphi))$.

Semantics. Let \mathcal{V} denote the set of propositional valuations over \mathcal{P} . For each $\varphi \in \mathcal{L}_0$, the set of valuations making φ true is denoted by $\text{mod}_0(\varphi)$. A *world* $W = \langle \{v_c\}_{c \in \mathcal{C}}, \{\Pi_i\}_{i \in \mathcal{S}} \rangle$ is a possible complete specification of the environment we find ourselves in:

- $v_c \in \mathcal{V}$ is the “true” valuation at case $c \in \mathcal{C}$;
- Π_i is a partition of \mathcal{V} for each $i \in \mathcal{S}$, representing the “true” expertise of source i ; and
- Π_* is the unit partition $\{\{v\} \mid v \in \mathcal{V}\}$.

Let \mathcal{W} denote the set of all worlds. Note that the partition corresponding to the distinguished source $*$ is fixed in all worlds as the finest possible partition, reflecting the fact that $*$ is completely reliable.

For any partition Π and valuation v , write $\Pi[v]$ for the unique cell in Π containing v . For a set of valuations U , write $\Pi[U] = \bigcup_{v \in U} \Pi[v]$. For brevity, we write $\Pi[\varphi]$ for $\Pi[\text{mod}_0(\varphi)]$. Then $\Pi[\varphi]$ is the set of valuations indistinguishable from a φ valuation.

For our belief change problem we will be interested in maintaining a collection of several belief sets, describing beliefs about each case $c \in \mathcal{C}$. Towards determining when a world W models such a collection, we define semantics for \mathcal{L} formulas with respect to a world and a case:

$$\begin{aligned} W, c \models \varphi &\iff v_c \in \text{mod}_0(\varphi) \\ W, c \models E_i(\varphi) &\iff \Pi_i[\varphi] = \text{mod}_0(\varphi) \\ W, c \models S_i(\varphi) &\iff v_c \in \Pi_i[\varphi] \end{aligned}$$

where $i \in \mathcal{S}$, $\varphi \in \mathcal{L}_0$, and the clauses for conjunction and negation are the expected ones. Since $\text{mod}_0(\varphi) \subseteq \Pi_i[\varphi]$ always holds, we have that $E_i(\varphi)$ holds iff there is no $\neg\varphi$ valuation which is indistinguishable from a φ valuation (c.f. Booth and Hunter (2018)). Note that since each source i has only a single partition Π_i used to interpret the expertise formulas, the truth value of $E_i(\varphi)$ does not depend on the case c . On the other hand, $S_i(\varphi)$ holds in case c iff the c -valuation of W is indistinguishable from some model of φ . That is, it is consistent with i 's expertise that φ is true.

Also note that if φ is a propositional tautology, $E_i(\varphi)$ holds for every source i . Thus, all sources are experts on *something*, even if just the tautologies.

Example 1. Let us extend the hospital example from the introduction. Let $\mathcal{S} = \{*, a, b\}$ denote the reliable source, Dr. A and Dr. B, and let $\mathcal{C} = \{c_1, c_2\}$ denote patients 1 and 2. Consider propositional variables $\mathcal{P} = \{x, y\}$, standing for condition X and Y respectively. Suppose that Dr. A has expertise on diagnosing condition Y only, whereas Dr. B only has expertise on X . For the sake of the example, suppose that patient 1 suffers from both conditions, and patient 2 suffers only from condition Y . This situation is modelled by the following world $W = \langle \{v_c\}_{c \in \{c_1, c_2\}}, \{\Pi_i\}_{i \in \{*, a, b\}} \rangle$:

$$\begin{aligned} v_{c_1} &= xy; & v_{c_2} &= \bar{x}y; \\ \Pi_a &= xy, \bar{x}y \mid x\bar{y}, \bar{x}\bar{y}; & \Pi_b &= xy, x\bar{y} \mid \bar{x}y, \bar{x}\bar{y}. \end{aligned}$$

We have $W, c \models E_a(y) \wedge E_b(x)$ for each $c \in \{c_1, c_2\}$. Also note that $W, c_1 \models x$ (patient 1 suffers from X), $W, c_1 \models S_a(\neg x)$ (it is sound for Dr. A to report otherwise; this holds since $\Pi_a[\neg x] = \{xy, \bar{x}y\} \cup \{x\bar{y}, \bar{x}\bar{y}\} \ni xy = v_{c_1}$), but $W, c_1 \not\models \neg S_b(\neg x)$ (the same formula is not sound for Dr. B; we have $\Pi_b[\neg x] = \{\bar{x}y, \bar{x}\bar{y}\} = \text{mod}_0(\neg x) \not\ni xy = v_{c_1}$).

Say Φ is valid if $W, c \models \Phi$ for all $W \in \mathcal{W}$ and $c \in \mathcal{C}$. For future reference we collect a list of validities.

Proposition 1. For any $i \in \mathcal{S}$, $c \in \mathcal{C}$ and $\varphi, \psi \in \mathcal{L}_0$, the following formulas are valid

1. $S_i(\varphi) \leftrightarrow S_i(\psi)$ and $E_i(\varphi) \leftrightarrow E_i(\psi)$, whenever $\varphi \equiv \psi$
2. $E_i(\varphi) \leftrightarrow E_i(\neg\varphi)$ and $E_i(\varphi) \wedge E_i(\psi) \rightarrow E_i(\varphi \wedge \psi)$
3. $E_i(p_1) \wedge \dots \wedge E_i(p_k) \rightarrow E_i(\varphi)$, where p_1, \dots, p_k are the propositional variables appearing in φ
4. $E_i(\varphi) \wedge S_i(\varphi) \rightarrow \varphi$, and $S_i(\varphi) \wedge \neg\varphi \rightarrow \neg E_i(\varphi)$
5. $S_i(\varphi) \wedge S_i(\neg\varphi) \rightarrow \neg E_i(\varphi)$
6. $S_*(\varphi) \leftrightarrow \varphi$ and $E_*(\varphi)$

(1) states syntax-irrelevance properties. (2) says that expertise is symmetric with respect to negation, and closed under conjunctions. Intuitively, symmetry means that i is an expert on φ if they know *whether or not* φ holds. (3) says

that expertise on each propositional variable in φ is sufficient for expertise on φ itself. (4) says that, in the presence of expertise, soundness of φ is sufficient for φ to in fact be true. (5) says that if both φ and $\neg\varphi$ are true up to the expertise of i , then i cannot have expertise on φ . Finally, (6) says that the reliable source $*$ has expertise on *all* formulas, and thus φ is sound for $*$ iff it is true.

Case-Indexed Collections. In the remainder of the paper we will be interested in forming beliefs about each case $c \in \mathcal{C}$. To do so we use collections of belief sets $G = \{\Gamma_c\}_{c \in \mathcal{C}}$, with $\Gamma_c \subseteq \mathcal{L}$, indexed by cases. Say a world W is a *model* of G iff

$$W, c \models \Phi \text{ for all } c \in \mathcal{C} \text{ and } \Phi \in \Gamma_c,$$

i.e. iff W satisfies all formulas in G in the relevant case. Let $\text{mod}(G)$ denote the models of G , and say that G is *consistent* if $\text{mod}(G) \neq \emptyset$. For $c \in \mathcal{C}$, define the *c-consequences*

$$\text{Cn}_c(G) = \{\Phi \in \mathcal{L} \mid \forall W \in \text{mod}(G), W, c \models \Phi\}.$$

We write $\text{Cn}(G)$ for the collection $\{\text{Cn}_c(G)\}_{c \in \mathcal{C}}$.

Example 2. Suppose $\mathcal{C} = \{c_1, c_2, c_3\}$, and define G by $\Gamma_{c_1} = \{S_i(p \wedge q)\}$, $\Gamma_{c_2} = \{E_i(p)\}$ and $\Gamma_{c_3} = \{E_i(q)\}$. Then, since expertise holds independently of case, any model W of G has $W, c_1 \models E_i(p) \wedge E_i(q)$. By Proposition 1 part (3), $W, c_1 \models E_i(p \wedge q)$. Since W satisfies Γ_{c_1} in case c_1 , Proposition 1 part (4) gives $W, c_1 \models p \wedge q$. Since W was an arbitrary model of G , we have $p \wedge q \in \text{Cn}_{c_1}(G)$, i.e. $p \wedge q$ is a c_1 -consequence of G . This illustrates how information about distinct cases can be brought together to have consequences for other cases.

For two collections $G = \{\Gamma_c\}_{c \in \mathcal{C}}$, $D = \{\Delta_c\}_{c \in \mathcal{C}}$, write $G \sqsubseteq D$ iff $\Gamma_c \subseteq \Delta_c$ for all c , and let $G \sqcup D$ denote the collection $\{\Gamma_c \cup \Delta_c\}_{c \in \mathcal{C}}$. With this notation, the case-indexed consequence operator satisfies analogues of the Tarskian consequence properties.³

Say a collection G is *closed* if $\text{Cn}(G) = G$. Closed collections provide an idealised representation of beliefs, which will become useful later on. For instance, when G is closed we have $E_i(\varphi) \in \Gamma_c$ iff $E_i(\varphi) \in \Gamma_d$ for all $c, d \in \mathcal{C}$; i.e. expertise statements are either present for all cases or for none. We also have $\text{Cn}_0[\Gamma_c] = [\Gamma_c]$, i.e. the propositional parts of G are (classically) closed.

In propositional logic, mod_0 is a 1-to-1 correspondence between closed sets of formulas and sets of valuations. This is not so in our setting, since some subsets of \mathcal{W} do not arise as the models of any collection. Instead, we have a 1-to-1 correspondence into a restricted collection of sets of worlds. Borrowing the terminology of Delgrande, Peppas, and Woltran (2018), say a set of worlds $S \subseteq \mathcal{W}$ is *elementary* if $S = \text{mod}(G)$ for some collection $G = \{\Gamma_c\}_{c \in \mathcal{C}}$.⁴

³ That is, (i) $G \sqsubseteq \text{Cn}(G)$, (ii) $G \sqsubseteq D$ implies $\text{Cn}(G) \sqsubseteq \text{Cn}(D)$, and (iii) $\text{Cn}(\text{Cn}(G)) = \text{Cn}(G)$.

⁴ Non-elementary sets can also exist for weaker logics (such as Horn logic (Delgrande, Peppas, and Woltran 2018)) which lack the syntactic expressivity to identify all sets of models. In our framework, \mathcal{C} -indexed collections are not expressive enough to specify *combinations of valuations*, since each Γ_c only says something about the valuation for c .

Elementariness is characterised by a certain closure condition. Say that two worlds W, W' are *partition-equivalent* if $\Pi_i^W = \Pi_i^{W'}$ for all sources i , and say W is a *valuation combination* from a set $S \subseteq \mathcal{W}$ if for all cases c there is $W_c \in S$ such that $v_c^W = v_c^{W_c}$. Then a set is elementary iff it is closed under valuation combinations of partition-equivalent worlds.

Proposition 2. $S \subseteq \mathcal{W}$ is elementary if and only if the following condition holds: for all $W \in \mathcal{W}$ and $W_1, W_2 \in S$, if W is partition-equivalent to both W_1, W_2 and W is a valuation combination from $\{W_1, W_2\}$, then $W \in S$.

3 The Problem

With the framework set out, we can formally define the problem. We seek an operator with the following behaviour:

- **Input:** A sequence of reports σ , where each report is a triple $\langle i, c, \varphi \rangle \in \mathcal{S} \times \mathcal{C} \times \mathcal{L}_0$ and $\varphi \neq \perp$. Such a report represents that *source i reports φ to hold in case c* . Note that we only allow sources to make *propositional* reports.
- **Output:** A pair $\langle B^\sigma, K^\sigma \rangle$, where $B^\sigma = \{B_c^\sigma\}_{c \in \mathcal{C}}$ is a collection of *belief sets* $B_c^\sigma \subseteq \mathcal{L}$ and $K^\sigma = \{K_c^\sigma\}_{c \in \mathcal{C}}$ is a collection of *knowledge sets* $K_c^\sigma \subseteq \mathcal{L}$.

3.1 Basic Postulates

We immediately narrow the scope of operators under consideration by introducing some basic postulates which are expected to hold. In what follows, say a sequence σ is **-consistent* if for each $c \in \mathcal{C}$ the set $\{\varphi \mid \langle *, c, \varphi \rangle \in \sigma\} \subseteq \mathcal{L}_0$ is classically consistent. Write G_{snd}^σ for the collection with $(G_{\text{snd}}^\sigma)_c = \{S_i(\varphi) \mid \langle i, c, \varphi \rangle \in \sigma\}$, i.e. the collection of soundness statements corresponding to the reports in σ .

Closure $B^\sigma = \text{Cn}(B^\sigma)$ and $K^\sigma = \text{Cn}(K^\sigma)$

Containment $K^\sigma \sqsubseteq B^\sigma$

Consistency If σ is *-consistent, B^σ and K^σ are consistent

Soundness If $\langle i, c, \varphi \rangle \in \sigma$, then $S_i(\varphi) \in K_c^\sigma$

K-bound $K^\sigma \sqsubseteq \text{Cn}(G_{\text{snd}}^\sigma \sqcup K^\emptyset)$

Prior-Extension $K^\emptyset \sqsubseteq K^\sigma$

Rearrangement If σ is a permutation of ρ , then $B^\sigma = B^\rho$ and $K^\sigma = K^\rho$

Equivalence If $\varphi \equiv \psi$ then $B^{\sigma \cdot \langle i, c, \varphi \rangle} = B^{\sigma \cdot \langle i, c, \psi \rangle}$ and $K^{\sigma \cdot \langle i, c, \varphi \rangle} = K^{\sigma \cdot \langle i, c, \psi \rangle}$

Closure says that the belief and knowledge collections are closed under logical consequence. In light of earlier remarks, this implies that the propositional belief sets $[B_c^\sigma]$ are closed under (propositional) consequence, and that $E_i(\varphi) \in B_c^\sigma$ iff $E_i(\varphi) \in B_c^\sigma$. *Containment* says that everything which is known is also believed. *Consistency* ensures the output is always consistent, provided we are not in the degenerate case where $*$ gives inconsistent reports. *Soundness* says we *know* that all reports are sound in their respective cases. This formalises our assumption that sources are *honest*, i.e. that false reports only arise due to lack of expertise. By Proposition 1 part (4) it also implies *experts are always right*: if a source has expertise on their report then it must be

true. While *Soundness* places a lower bound on knowledge, *K-bound* places an upper bound: knowledge cannot go beyond the soundness statements corresponding to the reports in σ together with the prior knowledge K^\emptyset . That is, from the point view of knowledge, a new report of $\langle i, c, \varphi \rangle$ only allows us to learn $S_i(\varphi)$ in case c (and to combine this with other reports and prior knowledge). Note that the analogous property for belief is *not* desirable: we want to be more liberal when it comes to beliefs, and allow for *defeasible inferences* going beyond the mere fact that reports are sound. *Prior-Extension* says that knowledge after a sequence σ extends the prior knowledge on the empty sequence \emptyset . *Rearrangement* says that the order in which reports are received is irrelevant. This can be justified on the basis that we are reasoning about *static worlds* for each case c , so that there is no reason to see more “recent” reports as any more or less important or truthful than earlier ones.⁵ Consequently, we can essentially view the input as a *multi-set* of belief sets – one for each source – bringing us close to the setting of belief merging. This postulate also appears as the commutativity postulate (**Com**) in the work of Schwind and Konieczny (2020). Finally, *Equivalence* says that the syntactic form of reports is irrelevant.

Taking all the basic postulates together, the knowledge component K^σ is fully determined once K^\emptyset is chosen.

Proposition 3. Suppose an operator satisfies the basic postulates. Then

1. $K^\sigma = \text{Cn}(G_{\text{snd}}^\sigma \sqcup K^\emptyset)$
2. $K^\emptyset = \text{Cn}(\emptyset)$ iff $K^\sigma = \text{Cn}(G_{\text{snd}}^\sigma)$ for all σ .

The choice of K^\emptyset depends on the scenario one wishes to model. While $\text{Cn}(\emptyset)$ is a sensible choice if the sequence σ is all we have to go on, we allow $K^\emptyset \neq \text{Cn}(\emptyset)$ in case *prior knowledge* is available (for example, the expertise of particular sources may be known ahead of time).

Another important property of knowledge, which follows from the basic postulates, says that *knowledge is monotonic*: knowledge after receiving σ and ρ together is just the case-wise union of K^σ and K^ρ .

K-conjunction $K^{\sigma \cdot \rho} = \text{Cn}(K^\sigma \sqcup K^\rho)$

K-conjunction reflects the idea that one should be cautious when it comes to knowledge: a formula should only be accepted as known if it won't be given up in light of new information.

Proposition 4. Any operator satisfying the basic postulates satisfies K-conjunction.

The postulates also imply some useful properties linking *trust* (seen as belief in expertise) and *belief/knowledge*.

Proposition 5. Suppose an operator satisfies the basic postulates. Then

1. If $\varphi \in K_c^\sigma$ and $\neg\psi \in \text{Cn}_0(\varphi)$ then $\neg E_i(\psi) \in K_c^{\sigma \cdot \langle i, c, \psi \rangle}$.
2. If $\langle i, c, \varphi \rangle \in \sigma$ and $E_i(\varphi) \in B_c^\sigma$ then $\varphi \in B_c^\sigma$.

⁵This argument is from (Delgrande, Dubois, and Lang 2006).

(1) expresses how knowledge can negatively affect trust: we should distrust sources who make reports we know to be false. (2) expresses how trust affects belief: we should believe reports from trusted sources. It can also be seen as a form of *success* for ordinary sources, and implies AGM success when $i = *$ (by Proposition 1 part (6) and *Closure*). We illustrate the basic postulates by formalising the introductory hospital example.

Example 3. Set \mathcal{S}, \mathcal{C} and \mathcal{P} as in Example 1, and consider the sequence

$$\sigma = (\langle *, c_1, x \rangle, \langle a, c_2, x \rangle, \langle b, c_2, \neg x \rangle, \langle a, c_1, \neg x \rangle).$$

What do we know on the basis of this sequence, assuming the basic postulates? First note that by *Soundness*, Proposition 1 part (6) and *Closure*, the report from $*$ gives $x \in K_c^\sigma$, i.e. reliable reports are known. *Soundness* also gives $S_a(x) \wedge S_b(\neg x) \in K_{c_2}^\sigma$. Combined with Proposition 1 parts (2), (4) and *Closure*, this yields $\neg(E_a(x) \wedge E_b(x)) \in K_c^\sigma$ for all c , formalising the intuitive idea that Drs. A and B cannot both be experts on X, since they give conflicting reports. Considering the final report from a, we get $x \wedge S_a(\neg x) \in K_{c_1}^\sigma$, and thus $\neg E_a(x) \in K_c^\sigma$ by *Closure*. So in fact Dr. A is known to be a non-expert on X.

What about beliefs? The basic postulates do not require beliefs to go beyond knowledge, so we cannot say much in general. An “optimistic” operator, however, may opt to believe that sources are experts unless we know otherwise, and thus maximise the information that can be (defeasibly) inferred from the sequence (in the next section we will introduce concrete operators obeying this principle). In this case we may believe that at least one source has expertise on x (i.e. $E_a(x) \vee E_b(x) \in B_c^\sigma$). Combined with $\neg E_a(x) \in K_c^\sigma$, *Closure* and *Containment*, we get $E_b(x) \in B_{c_2}^\sigma$. *Symmetry* of expertise together with Proposition 5 part (2) then gives $\neg x \in B_{c_2}^\sigma$, i.e. we trust Dr. B in the example and believe patient 2 does not suffer from condition X.

3.2 Model-Based Operators

While an operator is a purely syntactic object, it will be convenient to specify K^σ and B^σ in semantic terms by selecting a set of *possible* and *most plausible* worlds for each sequence σ . We call such operators *model-based*.

Definition 1. An operator is model-based if for every σ there are sets $\mathcal{X}_\sigma, \mathcal{Y}_\sigma \subseteq \mathcal{W}$ such that (i) $\mathcal{X}_\sigma \supseteq \mathcal{Y}_\sigma$; (ii) $\Phi \in K_c^\sigma$ iff $W, c \models \Phi$ for all $W \in \mathcal{X}_\sigma$; and (iii) $\Phi \in B_c^\sigma$ iff $W, c \models \Phi$ for all $W \in \mathcal{Y}_\sigma$.

In other words, K_c^σ (resp., B_c^σ) contains the formulas which hold at case c in all worlds in \mathcal{X}_σ (resp., \mathcal{Y}_σ). It follows from the relevant definitions that $\mathcal{X}_\sigma \subseteq \text{mod}(K^\sigma)$, and equality holds if and only if \mathcal{X}_σ is elementary (similarly for \mathcal{Y}^σ and B^σ). Model-based operators are characterised by our first two basic postulates.

Theorem 1. An operator satisfies *Closure* and *Containment* if and only if it is model-based.

Since we take *Closure* and *Containment* to be fundamental properties, all operators we consider from now on will be model-based. We introduce our first concrete operator.

Definition 2. Define the model-based operator weak-mb by

$$\mathcal{X}_\sigma = \mathcal{Y}_\sigma = \{W \mid W, c \models S_i(\varphi) \text{ for all } (i, c, \varphi) \in \sigma\}.$$

That is, the possible worlds \mathcal{X}_σ are exactly those satisfying the soundness constraint for each report in σ , i.e. false reports are only due to lack of expertise of the corresponding source. Syntactically, $K^\sigma = B^\sigma = \text{Cn}(G_{\text{snd}}^\sigma)$.

Clearly weak-mb satisfies *Soundness*, and one can show that it satisfies all of the basic postulates of Section 3.1.⁶ In fact, it is the *weakest* operator satisfying *Closure*, *Containment* and *Soundness*, in that for any other operator $\sigma \mapsto \langle \hat{B}^\sigma, \hat{K}^\sigma \rangle$ with these properties we have $B^\sigma \sqsubseteq \hat{B}^\sigma$ and $K^\sigma \sqsubseteq \hat{K}^\sigma$ for any σ .

Example 4. Consider weak-mb applied to the sequence $\sigma = (\langle *, c, p \rangle, \langle i, c, \neg p \wedge q \rangle)$. By *Soundness*, *Closure* and the validities from Proposition 1, we have $p \in K_c^\sigma$ and $\neg E_i(p) \in K_c^\sigma$. In fact, by *Closure*, we have $\neg E_i(p) \in K_d^\sigma$ for all cases d . However, we cannot say much about q : neither $q, \neg q, E_i(q)$ nor $\neg E_i(q)$ are in $B_c^\sigma = K_c^\sigma$.

4 Constructions

For model-based operators in Definition 1, the sets \mathcal{X}_σ and \mathcal{Y}_σ – which determine knowledge and belief – can depend on σ in a completely arbitrary manner. This lack of structure leads to very wide class of operators, and one cannot say much about them in general beyond the satisfaction of *Closure* and *Containment*. In this section we specialise model-based operators by providing two constructions.

4.1 Conditioning Operators

Intuitively, \mathcal{Y}_σ is supposed to represent the *most plausible* worlds among the possible worlds in \mathcal{X}_σ . This suggests the presence of a *plausibility ordering* on \mathcal{X}_σ , which is used to select \mathcal{Y}_σ . For our first construction we take this approach: we condition a fixed plausibility total preorder⁷ on the knowledge \mathcal{X}_σ , and obtain \mathcal{Y}_σ by selecting the minimal (i.e. most plausible) worlds.

Definition 3. An operator is a conditioning operator if there is a total preorder \leq on \mathcal{W} and a mapping $\sigma \mapsto \langle \mathcal{X}_\sigma, \mathcal{Y}_\sigma \rangle$ as in Definition 1 such that $\mathcal{Y}_\sigma = \min_{\leq} \mathcal{X}_\sigma$ for all σ .

Note that \leq is independent of σ : it is fixed before receiving any reports. All conditioning operators are model-based by definition. Clearly \mathcal{Y}_σ is determined by \mathcal{X}_σ and the plausibility order, so that to define a conditioning operator it is enough to specify \leq and the mapping $\sigma \mapsto \mathcal{X}_\sigma$. Write $W \simeq W'$ iff both $W \leq W'$ and $W' \leq W$. We now present examples of how such an ordering can be defined.

Definition 4. Define the conditioning operator var-based-cond by setting \mathcal{X}_σ in the same way as weak-mb in Definition 2, and $W \leq W'$ iff $r(W) \leq r(W')$, where

$$r(W) = - \sum_{i \in \mathcal{S}} |\{p \in \mathcal{P} \mid \Pi_i^W[p] = \text{mod}_0(p)\}|.$$

⁶ For *Consistency*, note that for any $*$ -consistent sequence σ one can form a world W such that v_c is a model of all reports from $*$ at case c , and $\Pi_i = \{\mathcal{V}\}$ for all $i \neq *$. This satisfies all the soundness constraints, so $W \in \mathcal{X}_\sigma = \mathcal{Y}_\sigma$.

⁷ A total preorder is a reflexive, transitive and total relation.

var-based-cond aims to trust each source on *as many propositional variables* as possible. One can check that var-based-cond satisfies the basic postulates.

Example 5. Revisiting the sequence $\sigma = (\langle *, c, p \rangle, \langle i, c, \neg p \wedge q \rangle)$ from Example 4 with var-based-cond, the knowledge set K_c^σ is the same as before, but we now have $q \wedge E_i(q) \in B_c^\sigma$. This reflects the “credulous” behaviour of the ranking \leq : while it is not possible to believe i is an expert on p , we should believe they are an expert on q so long as this does not conflict with soundness. For the propositional beliefs generally, we have $[B_c^\sigma] = \text{Cn}_0(p \wedge q)$. That is, var-based-cond takes the q part of the report from i (on which i is credulously trusted) while ignoring the $\neg p$ part (which is false due to report from $*$).

Definition 5. Define a conditioning operator part-based-cond with \mathcal{X}_σ as for var-based-cond, and \leq defined by the ranking function

$$r(W) = - \sum_{i \in \mathcal{S}} |\Pi_i^W|.$$

part-based-cond aims to maximise the *number of cells* in the sources’ partitions, and thereby maximise the number of propositions on which they have expertise. Unlike var-based-cond, the propositional variables play no special role. As expected, part-based-cond satisfies the basic postulates.

Example 6. Applying part-based-cond to σ from Examples 4 and 5, we no longer extract q from the report of i : $q \notin B_c^\sigma$ and $E_i(q) \notin B_c^\sigma$. Instead, we have $[B_c^\sigma] = \text{Cn}_0(p)$, and $E_i(p \vee q) \in B_c^\sigma$.

Note that both var-based-cond and part-based-cond are based on the general principle of maximising the expertise of sources, subject to the constraint that all reports are sound. This intuition is formalised by the following postulate for conditioning operators. In what follows, write $W \preceq W'$ iff Π_i^W refines $\Pi_i^{W'}$ for all $i \in \mathcal{S}$, i.e. if all sources have broadly more expertise in W than in W' .⁸

Refinement If $W \preceq W'$ then $W \leq W'$

Since \preceq is only a partial order on \mathcal{W} there are many possible total extensions; var-based-cond and part-based-cond provide two specific examples.

We now turn to an axiomatic characterisation of conditioning operators. Taken with the basic postulates from Section 3.1, conditioning operators can be characterised using an approach similar to that of Delgrande, Peppas, and Woltran (2018) in their account of *generalised AGM belief revision*.⁹ This involves a technical property Delgrande, Peppas, and Woltran call *Acyc*, which finds its roots in the *Loop* property of Kraus, Lehmann, and Magidor (1990).

Duplicate-removal If $\rho_1 = \sigma \cdot \langle i, c, \varphi \rangle$ and $\rho_2 = \rho_1 \cdot \langle i, c, \varphi \rangle$ then $B^{\rho_1} = B^{\rho_2}$ and $K^{\rho_1} = K^{\rho_2}$

Conditional-consistency If K^σ is consistent then so is B^σ

⁸ Π refines Π' if $\forall A \in \Pi, \exists B \in \Pi'$ such that $A \subseteq B$.

⁹ Note that while the result is similar, our framework is not an instance of theirs.

Inclusion-vacuity $B^{\sigma \cdot \rho} \sqsubseteq \text{Cn}(B^\sigma \sqcup K^\rho)$, with equality if $B^\sigma \sqcup K^\rho$ is consistent

Acyc If $\sigma_0, \dots, \sigma_n$ are such that $K^{\sigma_j} \sqcup B^{\sigma_{j+1}}$ is consistent for all $0 \leq j < n$ and $K^{\sigma_n} \sqcup B^{\sigma_0}$ is consistent, then $K^{\sigma_0} \sqcup B^{\sigma_n}$ is consistent

Inclusion-vacuity is so-named since it is analogous to the combination of *Inclusion* and *Vacuity* from AGM revision, if one informally views $B^{\sigma \cdot \rho}$ as the revision of B^σ by K^ρ . *Conditional-consistency* is another consistency postulate, which follows from *Consistency*, *Closure* and *Soundness*. *Acyc* is the analogue of the postulate of Delgrande, Peppas, and Woltran, which rules out cycles in the plausibility order constructed in the representation result.

As with the result of Delgrande, Peppas, and Woltran, a technical condition beyond Definition 3 is required to obtain the characterisation: say that a conditioning operator is *elementary* if for each σ the sets of worlds \mathcal{X}_σ and $\mathcal{Y}_\sigma = \min_{\leq} \mathcal{X}_\sigma$ are elementary.¹⁰

Theorem 2. Suppose an operator satisfies the basic postulates of Section 3.1.¹¹ Then it is an elementary conditioning operator if and only if it satisfies Duplicate-removal, Conditional-consistency, Inclusion-vacuity and Acyc.

The proof roughly follows the lines of Theorem 4.9 in (Delgrande, Peppas, and Woltran 2018), although some differences arise due to the form of our input as finite sequences of reports. We note that while the requirement that \mathcal{X}_σ and \mathcal{Y}_σ are elementary is a technical condition,¹² the characterisation in Proposition 2 implies a simple sufficient condition for elementariness.

Proposition 6. Suppose \leq is such that $W \simeq W'$ whenever W and W' are partition-equivalent. Then $\min_{\leq} S$ is elementary for any elementary set $S \subseteq \mathcal{W}$.

Proposition 6 implies that var-based-cond and part-based-cond are elementary. Indeed, for both operators $\mathcal{X}_\sigma = \text{mod}(G_{\text{snd}}^\sigma)$ so is elementary by definition. Since the ranking \leq for each operator only depends on the partitions of worlds, $\mathcal{Y}_\sigma = \min_{\leq} \mathcal{X}_\sigma$ is elementary also.

4.2 Score-Based Operators

The fact that the plausibility order \leq of a conditioning operator is fixed may be too limiting. For example, consider

$$\sigma = (\langle i, c, p \rangle, \langle j, c, \neg p \rangle, \langle i, d, p \rangle).$$

If one sets \mathcal{X}_σ to satisfy the soundness constraints (i.e. as in weak-mb), there is a possible world $W_1 \in \mathcal{X}_\sigma$ with $W_1, d \models \neg E_i(p) \wedge E_j(p) \wedge \neg p$ (i.e. W_1 sides with source j and p is false at d) and another world $W_2 \in \mathcal{X}_\sigma$ with $W_2, d \models E_i(p) \wedge \neg E_j(p) \wedge p$ (i.e. W_2 sides with source i). Appealing to symmetry, one may argue that neither world is *a priori* more plausible than the other, so any fixed plausibility order should have $W_1 \simeq W_2$. If these worlds are maximally plausible (e.g. if taking the “optimistic” view outlined

¹⁰ Equivalently, there is a total preorder \leq such that $\text{mod}(B^\sigma) = \min_{\leq} \text{mod}(K^\sigma)$ for all σ .

¹¹ Strictly speaking, we only need *Closure*, *Containment*, *K-conjunction*, *Equivalence* and *Rearrangement*.

¹² *Inclusion-vacuity* may fail for non-elementary conditioning.

in Example 3), conditioning gives $p \notin B_d^\sigma$ and $\neg p \notin B_d^\sigma$. However, there is an argument that W_2 should be considered more plausible than W_1 given the sequence σ , since W_2 validates the final report $\langle i, d, p \rangle$ whereas W_1 does not. Consequently, there is an argument that we should in fact have $p \in B_d^\sigma$.¹³ This shows that we need the plausibility order to be responsive to the input sequence for adequate belief change.¹⁴

As a result of this discussion, we look for operators whose plausibility ordering can depend on σ . One approach to achieve this in a controlled way is to have a ranking for each report $\langle i, c, \varphi \rangle$, and combine these to construct a ranking for each sequence σ . We represent these rankings by *scoring functions*, and call the resulting operators *score-based*.

Definition 6. An operator is *score-based* if there is a mapping $\sigma \mapsto \langle \mathcal{X}_\sigma, \mathcal{Y}_\sigma \rangle$ as in Definition 1 and functions $r_0 : \mathcal{W} \rightarrow \mathbb{N} \cup \{\infty\}$, $d : \mathcal{W} \times (\mathcal{S} \times \mathcal{C} \times \mathcal{L}_0) \rightarrow \mathbb{N} \cup \{\infty\}$ such that $\mathcal{X}_\sigma = \{W \mid r_\sigma(W) < \infty\}$ and $\mathcal{Y}_\sigma = \operatorname{argmin}_{W \in \mathcal{X}_\sigma} r_\sigma(W)$, where

$$r_\sigma(W) = r_0(W) + \sum_{\langle i, c, \varphi \rangle \in \sigma} d(W, \langle i, c, \varphi \rangle).$$

Here $r_0(W)$ is the *prior implausibility score* of W , and $d(W, \langle i, c, \varphi \rangle)$ is the *disagreement score* for world W and $\langle i, c, \varphi \rangle$. The set of most plausible worlds \mathcal{Y}_σ consists of those W which minimise the sum of the prior implausibility and the total disagreement with σ . Note that by summing the scores of each report $\langle i, c, \varphi \rangle$ with equal weight, we treat each report independently. Score-based operators generalise elementary conditioning operators with *K-conjunction*.

Proposition 7. Any elementary conditioning operator satisfying *K-conjunction* is *score-based*.

We now give a concrete example.

Definition 7. Define a *score-based operator* *excess-min* by setting $r_0(W) = 0$ and

$$d(W, \langle i, c, \varphi \rangle) = \begin{cases} |\Pi_i^W[\varphi] \setminus \operatorname{mod}_0(\varphi)|, & W, c \models S_i(\varphi) \\ \infty, & \text{otherwise.} \end{cases}$$

The set of possible worlds \mathcal{X}_σ is the same as for the earlier operators. All worlds are *a priori* equiplausible according to r_0 . The disagreement score d is defined as the number of propositional valuations in the “excess” of $\Pi_i^W[\varphi]$ which are not models of φ , i.e. the number of $\neg\varphi$ valuations which are indistinguishable from some φ valuation. The intuition here is that *sources tend to only report formulas on which they have expertise*. The minimum score 0 is attained exactly when i has expertise on φ ; other worlds are ordered by how much they deviate from this ideal.

One can verify that *excess-min* satisfies the basic postulates of Section 3.1. It can also be seen that \mathcal{X}_σ and \mathcal{Y}_σ are elementary, and *excess-min* fails *Duplicate-removal* and *Inclusion-vacuity*. It follows from Theorem 2 that *excess-min* is *not* a conditioning operator.

¹³At the very least, the case $p \in B_d^\sigma$ should not be excluded.

¹⁴In Section 5 we make this argument more precise by providing an impossibility result which shows conditioning operators with some basic properties cannot accept p in sequences such as this.

Example 7. To illustrate the differences between *excess-min* and *conditioning*, consider a more elaborate version of the example given at the start of this section:

$$\sigma = (\langle i, c, p \rightarrow q \rangle, \langle j, c, p \rightarrow \neg q \rangle, \langle *, c, p \rangle, \langle i, d, p \rangle, \langle i, d, q \rangle).$$

Here the reports of i and j in case c are consistent, but inconsistent when taken with the reliable information p from $*$. Should we believe q or $\neg q$? Both our conditioning operators *var-based-cond* and *part-based-cond* decline to decide, and have $[B_c^\sigma] = \operatorname{Cn}_0(p)$. However, since *excess-min* takes into account each report in the sequence, the fact that i reports both p and q in case d leads to $E_i(p) \wedge E_i(q) \in B_c^\sigma$. This gives $E_i(p \rightarrow q) \in B_c^\sigma$ by Proposition 1 part (3), so we can make use of the report from i in case c : we have $[B_c^\sigma] = \operatorname{Cn}_0(p \wedge q)$. This example shows that *score-based operators* can be more credulous than *conditioning operators* (e.g. we can believe $E_i(p)$ when i reports p), and can consequently hold stronger propositional beliefs.

5 One-Step Revision

The postulates of Section 3.1 only set out very basic requirements for an operator. In this section we introduce some more demanding postulates which address how beliefs should change when a sequence σ is extended by a new report $\langle i, c, \varphi \rangle$. First, we address how propositional beliefs should be affected by reliable information.

AGM-* For any σ and $c \in \mathcal{C}$ there is an AGM operator \star for $[B_c^\sigma]$ such that $[B_c^{\sigma \cdot \langle *, c, \varphi \rangle}] = [B_c^\sigma] \star \varphi$ whenever $\neg\varphi \notin K_c^\sigma$

*AGM-** says that receiving information from the reliable source $*$ acts in accordance with the well-known AGM postulates (Alchourrón, Gärdenfors, and Makinson 1985) for propositional belief revision (provided we are not in the degenerate case where the new report φ was already known to be false). Since AGM revision operators are characterised by total preorders over valuations (Grove 1988; Katsuno and Mendelzon 1991), it is no surprise that our order-based constructions are consistent with *AGM-**.

Proposition 8. *var-based-cond*, *part-based-cond* and *excess-min* satisfy *AGM-**.

Thus, we do indeed extend AGM revision in the case of reliable information. What about non-reliable information? First note that the analogue of *AGM-** for ordinary sources $i \neq *$ is *not* desirable. In particular, we should not have the *Success* postulate:

$$\varphi \in B_c^{\sigma \cdot \langle i, c, \varphi \rangle}.$$

Indeed, the sequence in Example 4 with $\varphi = \neg p \wedge q$ already shows that *Success* would conflict with the basic postulates. However, there are weaker modifications of *Success* which may be more appropriate. We consider two such postulates.

Cond-success If $E_i(\varphi) \in B_c^\sigma$ and $\neg\varphi \notin B_c^\sigma$, then $\varphi \in B_c^{\sigma \cdot \langle i, c, \varphi \rangle}$

Strong-cond-success If $\neg(E_i(\varphi) \wedge \varphi) \notin B_c^\sigma$, then $\varphi \in B_c^{\sigma \cdot \langle i, c, \varphi \rangle}$

Cond-success says that if i is deemed an expert on φ , which is consistent with current beliefs, then φ is accepted after i reports it. That is, the acceptance of φ is *conditional* on prior beliefs about the expertise of i (on φ). *Strong-cond-success* weakens the antecedent by only requiring that $E_i(\varphi)$ and φ are jointly consistent with current beliefs (i.e. i need not be considered an expert on φ). In other words, we should believe reports if there is no reason not to. It is easily shown that *Closure* and *Strong-cond-success* implies *Cond-success*. We once again revisit our examples.

Proposition 9. *var-based-cond, part-based-cond and excess-min satisfy Cond-success, and excess-min additionally satisfies Strong-cond-success.*

By omission, the reader may suppose that the conditioning operators fail *Strong-cond-success*. This is correct, and we can in fact say even more: *no* conditioning operator with a few basic properties – all of which are satisfied by var-based-cond and part-based-cond – can satisfy *Strong-cond-success*. In what follows, for a permutation $\pi : \mathcal{S} \rightarrow \mathcal{S}$ with $\pi(*) = *$, write $\pi(W)$ for the world with $v_c^{\pi(W)} = v_c^W$ and $\Pi_i^{\pi(W)} = \Pi_{\pi(i)}^W$. We have an impossibility result.

Proposition 10. *No elementary conditioning operator satisfying the basic postulates can simultaneously satisfy the following properties:*

1. $K^\emptyset = \text{Cn}(\emptyset)$
2. If π is a permutation of \mathcal{S} with $\pi(*) = *$, $W \simeq \pi(W)$
3. Refinement
4. Strong-cond-success

(1) says that before any reports are received, we only know tautologies. As remarked earlier, this is not an *essential* property, but is reasonable when no prior knowledge is available. (2) is an anonymity postulate: it says that permuting the “names” of sources does not affect the plausibility of a world, and is a desirable property in light of (1). *Refinement*, introduced in Section 4.1, says that worlds in which all sources have more expertise are preferred.

Proposition 10 highlights an important difference between conditioning and score-based operators, and hints that a fixed plausibility order may be too restrictive: we need to allow the order to be responsive to new reports in order to satisfy properties such as *Strong-cond-success*.

6 Selective Change

In the previous section we saw how a single formula φ may be accepted when it is received as an additional report. But what can we say about propositional beliefs when taking into account the *whole sequence* σ ? To investigate this we introduce an analogue of *selective revision* (Fermé and Hansson 1999), in which propositional beliefs are formed by “selecting” part of each input report (intuitively, some part consistent with the source’s expertise). In what follows, write $\sigma \upharpoonright c = \{\langle i, \varphi \rangle \mid \langle i, c, \varphi \rangle \in \sigma\}$ for the c -reports in σ .

Definition 8. *A selection scheme is a mapping f assigning to each $*$ -consistent sequence σ a function $f_\sigma : \mathcal{S} \times \mathcal{C} \times \mathcal{L}_0 \rightarrow \mathcal{L}_0$ such that $f_\sigma(i, c, \varphi) \in \text{Cn}_0(\varphi)$. An operator is*

selective if there is a selection scheme f such that for all $$ -consistent σ and $c \in \mathcal{C}$,*

$$[B_c^\sigma] = \text{Cn}_0(\{f_\sigma(i, c, \varphi) \mid \langle i, \varphi \rangle \in \sigma \upharpoonright c\}).$$

Thus, an operator is selective if its propositional beliefs in case c are formed by weakening each c -report and taking their consequences. Note that for $\sigma = \emptyset$ we get $[B_c^\sigma] = \text{Cn}_0(\emptyset)$, so selectivity already rules out non-tautological prior propositional beliefs. Also note that in the presence of *Closure*, *Containment* and *Soundness*, selectivity implies that $[B_c^\sigma] = [B_c^\rho]$, where ρ is obtained by replacing each report $\langle i, c, \varphi \rangle$ with $\langle *, c, f_\sigma(i, c, \varphi) \rangle$.

Selectivity can be characterised by a natural postulate placing an upper bound on the propositional part of B_c^σ . In what follows, let $\Gamma_c^\sigma = \{\varphi \in \mathcal{L}_0 \mid \exists i \in \mathcal{S} : \langle i, \varphi \rangle \in \sigma \upharpoonright c\}$.

Boundedness *If σ is $*$ -consistent, $[B_c^\sigma] \subseteq \text{Cn}_0(\Gamma_c^\sigma)$*

Boundedness says that the propositional beliefs in case c should not go beyond the consequences of the formulas reported in case c . In some sense this can be seen as an iterated version of *Inclusion* from AGM revision, in the case where $[B_c^\sigma] = \text{Cn}_0(\emptyset)$. We have the following characterisation.

Theorem 3. *A model-based operator is selective if and only if it satisfies Boundedness.*

This result allows us to easily analyse when conditioning and score-based operators are selective; we show in the full paper (Singleton and Booth 2022) that var-based-cond, part-based-cond and excess-min are all selective.

6.1 Case Independence

In the definition of a selection scheme, we allow $f_\sigma(i, c, \varphi)$ to depend on the case c . If one views $f_\sigma(i, c, \varphi)$ as a weakening of φ which accounts for the lack of expertise of i , this is somewhat at odds with other aspects of the framework, where expertise is independent of case. For this reason it is natural to consider *case independent* selective schemes.

Definition 9. *A selection scheme f is case independent if $f_\sigma(i, c, \varphi) \equiv f_\sigma(i, d, \varphi)$ for all $*$ -consistent σ and $i \in \mathcal{S}$, $c, d \in \mathcal{C}$ and $\varphi \in \mathcal{L}_0$.*

Say an operator is *case-independent-selective* if it is selective according to some case independent scheme. This stronger notion of selectivity can again be characterised by a postulate which bounds propositional beliefs. For any set of cases $H \subseteq \mathcal{C}$, sequence σ and $c \in \mathcal{C}$, write

$$\Gamma_c^{\sigma, H} = \{\varphi \in \mathcal{L}_0 \mid \exists i \in \mathcal{S} : \langle i, \varphi \rangle \in \sigma \upharpoonright c \text{ and } \forall d \in H : \langle i, \varphi \rangle \notin \sigma \upharpoonright d\}.$$

H-Boundedness *For any $*$ -consistent σ , $H \subseteq \mathcal{C}$ and $c \in \mathcal{C}$,*

$$[B_c^\sigma] \subseteq \text{Cn}_0 \left(\Gamma_c^{\sigma, H} \cup \bigcup_{d \in H} [B_d^\sigma] \right)$$

Note that *Boundedness* is obtained as the special case where $H = \emptyset$. We illustrate with an example.

Example 8. *Consider case c in the following sequence:*

$$\sigma = (\langle i, c, p \rangle, \langle j, c, q \rangle, \langle j, d, q \rangle, \langle k, d, r \rangle)$$

Boundedness requires that $[B_c^\sigma] \subseteq \text{Cn}_0(\{p, q\})$. However, the instance of H-Boundedness with $H = \{d\}$ makes use of the fact that j reports q in both cases c and d , and requires $[B_c^\sigma] \subseteq \text{Cn}_0(\{p\} \cup [B_d^\sigma])$. This also has an interesting implication for case d : if $\varphi \in [B_c^\sigma]$, then $p \rightarrow \varphi \in [B_d^\sigma]$. This follows since $\beta \in \text{Cn}_0(\{\alpha\} \cup \Gamma)$ iff $\alpha \rightarrow \beta \in \text{Cn}_0(\Gamma)$ for $\alpha, \beta \in \mathcal{L}_0$. Intuitively, this says that if p (from i) and q (from j) is enough to accept φ in case c , then φ is accepted in case d if p is, given that the report of q from j is repeated for d .

Theorem 4. A model-based operator is case-independent-selective if and only if it satisfies H-Boundedness.

The question of whether our concrete operators satisfy H-Boundedness (equivalently, whether they are case-independent-selective) is still open.

7 Related Work

Belief Merging. In the framework of Konieczny and Pino Pérez (2002), a merging operator Δ maps a multiset of propositional formulas $\Phi = \{\varphi_1, \dots, \varphi_n\}$ and an integrity constraint μ to a formula $\Delta_\mu(\Phi)$. This can be seen as the special case of our framework with a single case c : for Φ, μ we consider the sequence $\sigma_{\Phi, \mu}$ where $*$ reports μ and each source i reports φ_i . Any operator then gives rise to a merging operator $\Delta_\mu(\Phi) = \bigwedge [B_c^{\sigma_{\Phi, \mu}}]$.

We go beyond this setting by considering multiple cases and explicitly modelling expertise (and trust, via beliefs about expertise). While it may be possible to model expertise *implicitly* in belief merging (for example, say i is not trusted on ψ if $\Delta_\mu(\Phi) \not\vdash \psi$ when $\varphi_i \vdash \psi$), bringing expertise to the object level allows us to express more complex beliefs about expertise, such as $E_a(x) \vee E_b(x)$ in Example 3. It also facilitates postulates which refer directly to expertise, such as the weakenings of *Success* in Section 5.

However, our problem is more specialised than merging, since we focus specifically on conflicting information due to lack of expertise. Belief merging may be applied more broadly to other types of *information fusion*, e.g. subjective beliefs or goals (Grégoire and Konieczny 2006), where notions of objective expertise do not apply. While our framework *could* be applied in these settings, our postulates may no longer be desirable.

Epistemic Logic. Our notions of expertise and soundness are related to *S5 knowledge* from epistemic logic (van Ditmarsch et al. 2015). In such logics, an agent *knows* φ at a state x if φ holds at all states y “accessible” from x . Knowledge is thus determined by an *epistemic accessibility relation*, which describes the distinctions between states the agent can make. The logic of S5 arises when this relation is an equivalence relation (or equivalently, a partition).

Our previous work (Singleton 2021) – in which expertise and soundness were introduced in a modal logic framework – showed that “expertise models” are in 1-to-1 correspondence with S5 models, such that $E(\varphi)$ holds iff $A(\varphi \rightarrow K\varphi)$ holds in the S5 model, where A is the universal modality. By symmetry of expertise, we can also replace φ with its negation. Thus, expertise has a precise epistemic interpretation: it is the ability to *know whether* φ holds in any

possible state. Similarly, $S(\varphi)$ translates to $\neg K\neg\varphi$. That is, φ is sound exactly when the source does not *know* φ is false.

In the present framework, if we set $W, c \models K_i(\varphi)$ iff $\Pi_i[v_c] \subseteq \text{mod}_0(\varphi)$ and $W, c \models A\Phi$ iff $\forall v : W_{c=v}, c \models \Phi$, where $W_{c=v}$ is the world obtained from W by setting $v'_c = v$, then we have $E_i(\varphi) \equiv A(\varphi \rightarrow K_i(\varphi))$ and $S_i(\varphi) \equiv \neg K_i(\neg\varphi)$. While K_i is not quite an S5 modality (the **5** axiom requires iterating K_i , which is not possible in our framework), this shows the fundamental link between expertise, soundness and knowledge.

8 Conclusion

Summary. In this paper we studied a belief change problem – extending the classical AGM framework – in which beliefs about the state of the world in multiple cases, as well as expertise of multiple sources, must be inferred from a sequence of reports. This allowed us to take a fresh look at the interaction between trust (seen as *belief in expertise*) and belief. By inferring the expertise of the sources from the reports, we have generalised some earlier approaches to non-prioritised revision which assume expertise (or reliability, credibility, priority etc) is known up-front (e.g. (Fermé and Hansson 1999; Hansson et al. 2001; Booth and Hunter 2018; Delgrande, Dubois, and Lang 2006)). We went on to propose some concrete belief change operators, and explored their properties through examples and postulates.

We saw that conditioning operators satisfy some desirable properties, and our concrete instances make useful inferences that go beyond weak-mb. However, we have examples in which intuitively plausible inferences are blocked, and conditioning is largely incompatible with *Strong-cond-success*. Score-based operators, and in particular excess-min, offer a way around these limitations, but may come at the expense of some other seemingly reasonable postulates, such as *Duplicate-removal*.

Future Work. There are many possibilities for future work. Firstly, we have a representation result only for conditioning operators. A characterisation of score-based operators – either the class in general or the specific operator excess-min – remains to be found. This would help to further clarify the differences between conditioning and score-based operators. We have also not considered any computational issues. Determining the complexity of calculating the results of our example operators, and the complexity for conditioning and score-based operators more broadly, is left to future work. Secondly, there is scope for deeper postulate-based analysis. For example, there should be postulates governing how beliefs change in case c in response to reports in case d . We could also consider more postulates relating trust and belief, and compare these postulates with those of Yasser and Ismail (2021). Finally, our framework only deals with three levels of trust on a proposition: we can believe $E_i(\varphi)$, believe $\neg E_i(\varphi)$, or neither. Future work could investigate how to extend our semantics to talk about *graded expertise*, and thereby permit more fine-grained *degrees of trust* (Hunter 2021; Yasser and Ismail 2021; Delgrande, Dubois, and Lang 2006).

Acknowledgements

We thank the anonymous reviewers for their thoughtful comments and suggestions.

References

- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of symbolic logic* 510–530.
- Booth, R., and Hunter, A. 2018. Trust as a Precursor to Belief Revision. *J. Artif. Intell. Res.* 61:699–722.
- Delgrande, J. P.; Dubois, D.; and Lang, J. 2006. Iterated revision as prioritized merging. *KR* 6:210–220.
- Delgrande, J. P.; Peppas, P.; and Woltran, S. 2018. General belief revision. *Journal of the ACM (JACM)* 65(5):1–34.
- Fermé, E. L., and Hansson, S. O. 1999. Selective revision. *Studia Logica* 63(3):331–342.
- Grove, A. 1988. Two modellings for theory change. *Journal of philosophical logic* 157–170.
- Grégoire, E., and Konieczny, S. 2006. Logic-based approaches to information fusion. *Information Fusion* 7(1):4 – 18. Logic-based Approaches to Information Fusion.
- Hansson, S. O.; Fermé, E. L.; Cantwell, J.; and Falappa, M. A. 2001. Credibility limited revision. *Journal of Symbolic Logic* 66(4):1581–1596.
- Hansson, S. O. 1999. A survey of non-prioritized belief revision. *Erkenntnis* 50(2):413–427.
- Hunter, A. 2021. Building trust for belief revision. In Pham, D. N.; Theeramunkong, T.; Governatori, G.; and Liu, F., eds., *PRICAI 2021: Trends in Artificial Intelligence*, 543–555. Cham: Springer International Publishing.
- Katsuno, H., and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52(3):263–294.
- Konieczny, S., and Pino Pérez, R. 2002. Merging information under constraints: a logical framework. *Journal of Logic and Computation* 12(5):773–808.
- Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial intelligence* 44(1-2):167–207.
- Li, Y.; Gao, J.; Meng, C.; Li, Q.; Su, L.; Zhao, B.; Fan, W.; and Han, J. 2016. A Survey on Truth Discovery. *SIGKDD Explor. Newsl.* 17(2):1–16.
- Schwind, N., and Konieczny, S. 2020. Non-Prioritized Iterated Revision: Improvement via Incremental Belief Merging. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, 738–747.
- Singleton, J., and Booth, R. 2022. Who’s the expert? On multi-source belief change. <https://arxiv.org/abs/2205.00077>.
- Singleton, J. 2021. A logic of expertise. *ESSLLI 2021 Student Session*. <https://arxiv.org/abs/2107.10832>.
- van Ditmarsch, H.; van der Hoek, W.; Halpern, J.; and Kooi, B., eds. 2015. *Handbook of Epistemic Logic*. College Publications.
- Yasser, A., and Ismail, H. O. 2021. Trust is all you need: From belief revision to information revision. In Faber, W.; Friedrich, G.; Gebser, M.; and Morak, M., eds., *Logics in Artificial Intelligence*, 50–65. Cham: Springer International Publishing.