

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/344070928>

On the Joint Revision of Belief and Trust

Conference Paper · August 2020

CITATIONS

0

READS

35

2 authors:



[Ammar Yasser](#)

The German University in Cairo

2 PUBLICATIONS 0 CITATIONS

SEE PROFILE



[Haythem Osman Ismail](#)

The German University in Cairo

40 PUBLICATIONS 194 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



JESICS [View project](#)



Practical Reasoning [View project](#)

On the Joint Revision of Belief and Trust

Ammar Yasser¹ and Haythem O. Ismail^{2,1}

¹ German University in Cairo, Egypt

² Cairo University, Egypt

{`ammар.abbas, haythem.ismail`}@guc.edu.eg

Abstract. Trust plays a vital role when it comes to beliefs. Deciding what to believe and what not to believe depends, to a large extent, on trust. When an information source we trust conveys a piece of information, we are more likely to believe it. On the contrary, we are more reluctant to believe information communicated by information sources we do not trust. In general, trust guides us while revising our beliefs. Despite the existence of great bodies of literature on trust and belief revision separately, a formal treatment for their intertwined relationship is lacking. Hence, in this paper, we argue that trust revision and belief revision are inseparable processes. To provide a formal treatment for the joint revision of beliefs and trust, we address issues concerning the formalization of trust in information sources and provide AGM-style postulates for the rational joint revision of the two attitudes which we refer to as information revision.

Keywords: Belief Revision · Trust Revision · Information Revision.

1 Introduction

Trust acts, even if we are not aware, as an information filter. We are willing to believe in information communicated by sources we trust, cautious about information from sources we do not trust, and suspicious about information from sources we mistrust. Trust and mistrust are constantly revised; we gain more trust in information sources the more they prove themselves to be reliable, and our trust in them erodes as they mislead us one time after the other. Such attitudes allow us to be resilient, selective and astute. If exhibited by logic-based agents, these same attitudes would make them less susceptible to holding false beliefs and, hence, less prone to excessive belief revision. Moreover, by revising trust, these agents will not forever be naively trusting nor cynically mistrusting.

Trust has been thoroughly investigated within multi-agent systems [3,9,20,19,29,22, for instance], psychology [32,8,15, for instance], and philosophy [17,14,27, for instance]. Crucially, it was also investigated in the logic-based artificial intelligence (AI) literature by several authors [4,6,24,22,16,23]. Nevertheless, we believe that there are several issues that are left unaddressed by the logical approaches. Intuitively, trust is intimately related to misleading, on one hand, and belief revision, on the other. While several logical treatments of misleading are to be found in the literature [31,7,30,18, for instance], the relation of misleading to trust erosion is often not attended to or delegated to future work. On the other hand, the extensive literature on belief revision [1,11, for example], while occasionally addressing trust-based revision of beliefs [26,28,2] does

not have much to say about the revision of trust (but see [24,26] for minimal discussions) and, as far as we know, any systematic study of jointly revising belief and trust. The goal of this paper is, hence, twofold: (i) to motivate why belief and trust revision are intertwined and should be carried out together, and (ii) to propose AGM-style postulates for the joint revision of trust and belief.

The paper is structured as follows. Section 2 describes what we mean by trust, information and information sources. It also highlights the intuitions behind joint trust and belief revision. In Section 3, we present information states, a generic structure representing information and investigating its properties. Section 4 presents a powerful notion of *relevance* which information structures give rise to. Finally, in Section 5, the formal inter-dependency of belief and trust is explored, culminating in AGM-style postulates for joint belief-trust revision.³

2 Trust and Belief

It is often noted that trust is not a dyadic relation, between the trusted and the trustee, but is a triadic relation involving an object of trust [27]. You trust your doctor with your health, your mechanic with your car, your parents to unconditionally believe you, and your mathematics professor to tell you only true statements of mathematics. Our investigation of the coupling of belief and trust lets us focus only on trust in *sources of information*. Trust in information sources comes in different forms. Among Demolombe's [5,25] different types of trust in information sources, we focus on trust in *sincerity* and *competence* since they are the two types relevant to belief revision and realistic information sources.⁴ A sincere information source is one which (if capable of forming beliefs) only conveys what it believes; a competent source is one which only conveys what is true. In this paper, we consider trust in the *reliability* of information sources, where a source is reliable if it is both sincere and competent.⁵ Note that we do not take information sources to only be cognitive agents. For example, a sensor (or perception, in general) is a possible source of information. For information sources which are not cognitive agents, reliability reduces to competence.

Rational agents constantly receive information, and are faced with the question of whether to believe or not to believe. The question is rather simple when the new information is consistent with the agent's beliefs, since no obvious risk lies in deciding either way. Things become more interesting if the new information is inconsistent with what the agent believes; if the agent decides to accept the new information, it is faced with the problem of deciding on which of its old beliefs to give up in order to maintain consistency. Principles for rationally doing this are the focus of the vast literature on belief revision [1,12, for example].

It is natural to postulate that deciding whether to believe and how to revise our beliefs—the process of belief revision—are influenced by how much we trust the source

³ For a more detailed discussion, the reader is kindly guided towards [33].

⁴ Trust in *completeness*, for example, is unrealistic since it requires that the source informs about P whenever P is true.

⁵ As suggested by [18], it is perhaps possible that breach of sincerity and competence should have different effects on belief revision; for simplicity, we do not consider this here though.

of the new piece of information. (Also see [26,28,2].) In particular, in case of a conflict with old beliefs, how much we trust in the source's reliability and how much evidence we have accumulated for competing beliefs seem to be the obvious candidates for guiding us in deciding what to do. Thus, *rational belief revision depends on trust*.

But things are more complex. For example, suppose that information source σ_1 , whom we trust very much, conveys ϕ to us. ϕ is inconsistent with our beliefs but, because we trust σ_1 , we decide to believe in ϕ and give away ψ which, together with other beliefs, implies $\neg\phi$. In this case, we say that ϕ is a *refutation* of ψ . So far, this is just belief revision, albeit one which is based on trust. But, by stopping believing in ψ , we may find it rational to revise, and decrease, our *trust* in σ_2 who earlier conveyed ψ to us. Moreover, suppose that ϕ , together with other beliefs, implies our old belief ξ . We say that ϕ is a *confirmation* of ξ . This confirmation may trigger us to revise, and increase, our trust in σ_3 who is the source of ξ . Thus, *trust revision depends on belief revision*. In fact, belief revision may be the sole factor that triggers rational trust revision in information sources.

We need not stop there though. For, by reducing our trust in σ_2 's reliability, we are perhaps obliged to stop believing (or reduce our degree of belief in) ψ' which was conveyed by σ_2 . It is crucial to note that ψ' may be totally consistent with ϕ and we, nevertheless, give it away. While we find such scenario quite plausible, classical belief revision, with its upholding of the principle of minimal change, would deem it irrational. Likewise, by increasing our trust in σ_3 we may start believing (or raise our degree of belief) in ξ' which was earlier conveyed by σ_3 . This second round of belief revision can start a second round of trust revision. It is clear that we may keep on doing this for several rounds (perhaps indefinitely) if we are really fanatic about information and its sources. Hence, we contend that belief revision and trust revision are so entangled that they need to be combined into one process of *joint belief-trust revision* or, as we shall henceforth refer to it, *information revision*.

3 Information States

In this section, we introduce formal structures for representing information in a way that would facilitate information revision.

Definition 1. An *information grading structure* \mathcal{G} is a quadruple $(\mathcal{D}_b, \mathcal{D}_t, \prec_b, \prec_t)$, where \mathcal{D}_b and \mathcal{D}_t are non-empty, countable sets; and \prec_b and \prec_t are, respectively, total orders over \mathcal{D}_b and \mathcal{D}_t .

\mathcal{D}_b and \mathcal{D}_t contain the degrees of belief and trust, respectively. They are not necessarily finite, disjoint, different or identical.⁶ Moreover, to be able to distinguish the strength by which an agent believes a proposition or trusts a source, the two sets are ordered; here, we assume them to be totally ordered.

Definition 2. An *information structure* \mathcal{I} is a quadruple $(\mathcal{L}, \mathcal{C}, \mathcal{S}, \mathcal{G})$, where

⁶ \mathcal{D}_b and \mathcal{D}_t are usually the same; however, a qualitative account of trust and belief might have different sets for grading the two attitudes.

1. \mathcal{L} is a logical language with a Tarskian consequence operator Cn ,
2. \mathcal{C} is a finite cover of \mathcal{L} whose members are referred to as topics,
3. \mathcal{S} is a non-empty finite set of information sources, and
4. \mathcal{G} is an information grading structure.

Information structures comprise our general assumptions about information. \mathcal{S} is the set of possible information sources. Possible pieces of information are statements of the language \mathcal{L} , with each piece being about one or more, but finitely many, topics as indicated by the \mathcal{L} -cover \mathcal{C} . \mathcal{L} is only required to have a Tarskian consequence operator [13]. A topic represents the scope of trust. That is, an agent trusts an information source on particular topics. Moreover, a topic is a set of statements which may be closed under all connectives, some connectives or none at all. Topics could also be disjoint or overlapping. Choosing topics to be not necessarily closed under logical connectives allows us to accommodate interesting cases. For example, \mathcal{A} may have, for the same source, a different trust value when conveying ϕ to when it conveys $\neg\phi$.

Definition 3. Let $\mathcal{I} = (\mathcal{L}, \mathcal{C}, \mathcal{S}, (\mathcal{D}_b, \mathcal{D}_t, \prec_b, \prec_t))$ be an information structure. An **information state** \mathcal{K} over \mathcal{I} is a triple $(\mathcal{B}, \mathcal{T}, \mathcal{H})$, where

1. $\mathcal{B} : \mathcal{L} \hookrightarrow \mathcal{D}_b$ is a partial function referred to as the **belief base**,
2. $\mathcal{T} : \mathcal{S} \times \mathcal{C} \hookrightarrow \mathcal{D}_t$ is a partial function referred to as the **trust base**, and
3. $\mathcal{H} \subseteq \mathcal{L} \times \mathcal{S}$, **the history**, is a finite set of pairs (ϕ, σ) where, for every $T \in \mathcal{C}$, if $\phi \in T$ then $(\sigma, T, d_t) \in \mathcal{T}$, for some $d_t \in \mathcal{D}_t$.

Trust in information sources is recorded in $\mathcal{T}(\mathcal{K})$.⁷ This is a generalization to accommodate logics with an explicit account of trust in the object language [6,23, for instance] as well as those without [22,21, for example]. $\mathcal{H}(\mathcal{K})$ acts as a formal device for recording conveyance instances. Hence, a tuple of the form (ϕ, σ) in the history denotes that formula ϕ was acquired through source σ . As with $\mathcal{T}(\mathcal{K})$, we do not require \mathcal{L} to have an explicit account for conveying.

With this setup, having trust on single propositions, as is most commonly the case in the literature [5,23, for instance], reduces to restricting all topics to be singletons. On the other hand, we may account for absolute trust in sources by having a single topic to which all propositions belong.

To put together all the pieces forming an information state, consider the following example.

Example 1. Let an information grading structure $\mathcal{G} = (\{b_1, b_2, b_3\}, \{t_1, t_2, t_3\}, \{(b_1 \prec_b b_2), (b_1 \prec_b b_3), (b_2 \prec_b b_3)\}, \{(t_1 \prec_t t_2), (t_1 \prec_t t_3), (t_2 \prec_t t_3)\})$. Hence, a belief associated with b_3 is the most preferred while a belief assigned b_1 is the least preferred. Similarly, a source associated with t_3 is the most trusted while a source that has a trust degree of t_1 is the least trusted. Given \mathcal{G} , we define information structure $\mathcal{I} = (\mathcal{L}_{\mathcal{V}}, \mathcal{C}, \{\sigma_1, \sigma_2\}, \mathcal{G})$ where language $\mathcal{L}_{\mathcal{V}}$ is a propositional language with the set $\mathcal{V} = \{P, Q, S\}$ of propositional variables and $\mathcal{C} = \{T_P, T_Q, T_S\}$ where T_P is the

⁷ It is worthy of mentioning that we use $\mathcal{B}(\mathcal{K})$, $\mathcal{T}(\mathcal{K})$, and $\mathcal{H}(\mathcal{K})$ to denote the belief base, trust base, and history of a particular information state \mathcal{K} . More importantly, we do not use such notation to denote an invocation of a function with \mathcal{K} as an argument.

set of formulas containing P as a sub-formula, likewise for T_Q and T_S with Q and S respectively. Finally, information state $\mathcal{K} = (\mathcal{B}, \mathcal{T}, \mathcal{H})$ where

- $\mathcal{B} = \{(P, b_2), (Q, b_3), (Q \rightarrow P, b_3)\}$
- $\mathcal{T} = \{(\sigma_1, T_P, t_2), (\sigma_1, T_S, t_1), (\sigma_2, T_Q, t_3), (\sigma_2, T_S, t_2)\}$
- $\mathcal{H} = \{(P, \sigma_1), (Q, \sigma_2), (Q \rightarrow P, \sigma_2), (S, \sigma_1)\}$

This example shows how information is represented in an information state. Trust attribution for information sources with different topics is recorded. Based on that, we have P, Q , and $Q \rightarrow P$ that are believed with varying degrees of belief.⁸ Moreover, it is important to note that S was conveyed by σ_1 but was not believed. There could be many reasons for not believing a conveyed proposition, one of them could be that source σ_1 has the weakest degree of trust on the topic containing S .

So far, we defined what information states are. We now define the following abbreviations of which we will later make use.

- $\sigma(\mathcal{H}(\mathcal{K})) = \{\phi \mid (\phi, \sigma) \in \mathcal{H}(\mathcal{K})\}$
- $\mathcal{S}_{\mathcal{K}} = \{\sigma \mid (\phi, \sigma) \in \mathcal{H}(\mathcal{K})\}$
- $For(\mathcal{B}(\mathcal{K})) = \{\phi \mid (\phi, d_b) \in \mathcal{B}(\mathcal{K})\}$
- $\Phi_{\mathcal{K}} = \{\phi \mid \phi \in For(\mathcal{B}(\mathcal{K})) \text{ or } (\phi, \sigma) \in \mathcal{H}(\mathcal{K}) \text{ for any } \sigma\}$

Information revision is the process of revising an information state \mathcal{K} with the conveyance of a formula ϕ by a source σ yielding a revised information state. Every information revision operator is associated with a *conveyance inclusion filter* $\mathcal{F} \subseteq \mathcal{L} \times \mathcal{S}$ which determines the conveyance instances that make it into $\mathcal{H}(\mathcal{K})$. Hence, a generic revision operator is denoted by $\times_{\mathcal{F}}$, where \mathcal{F} is the associated filter. Revising \mathcal{K} with a conveyance of ϕ by σ is denoted by $\mathcal{K} \times_{\mathcal{F}} (\phi, \sigma)$. We require all revision operators $\times_{\mathcal{F}}$ to have the same effect on the history:

$$\mathcal{H}(\mathcal{K} \times_{\mathcal{F}} (\phi, \sigma)) = \begin{cases} \mathcal{H}(\mathcal{K}) \cup \{(\phi, \sigma)\} & (\phi, \sigma) \in \mathcal{F} \\ \mathcal{H}(\mathcal{K}) & \text{otherwise} \end{cases}$$

There are three major filter types. A filter \mathcal{F} is *non-forgetful* if $\mathcal{F} = \mathcal{L} \times \mathcal{S}$; it is *forgetful* if $\emptyset \neq \mathcal{F} \subset \mathcal{S} \times \mathcal{L}$; and it is *memory-less* if $\mathcal{F} = \emptyset$. Having filters beside the non-forgetful one is to simulate realistic scenarios where an agent does not always remember every piece of information that was conveyed to it. Henceforth, the subscript \mathcal{F} will be dropped from $\times_{\mathcal{F}}$ whenever this does not result in ambiguity.

Definition 4. Let $\phi \in \mathcal{L}$ and $\sigma \in \mathcal{S}$.

1. ϕ is **more entrenched** in state \mathcal{K}_2 over state \mathcal{K}_1 , denoted $\mathcal{K}_1 \prec_{\phi} \mathcal{K}_2$, if (i) $\phi \notin Cn(For(\mathcal{B}(\mathcal{K}_1)))$ and $\phi \in Cn(For(\mathcal{B}(\mathcal{K}_2)))$; or (ii) $(\phi, b_1) \in \mathcal{B}(\mathcal{K}_1)$, $(\phi, b_2) \in \mathcal{B}(\mathcal{K}_2)$, and $b_1 \prec_b b_2$. If $\mathcal{K}_1 \not\prec_{\phi} \mathcal{K}_2$ and $\mathcal{K}_2 \not\prec_{\phi} \mathcal{K}_1$, we write $\mathcal{K}_1 \equiv_{\phi} \mathcal{K}_2$.

⁸ For this simple example, the mapping between trust degrees and belief degrees was straight forward. However, as we will see later, addressing the relationship between trust degrees and belief degrees is a complex matter.

2. σ is **more trusted on topic** T in state \mathcal{K}_2 over state \mathcal{K}_1 , denoted $\mathcal{K}_1 \prec_{\sigma,T} \mathcal{K}_2$, if $(\sigma, T, t_1) \in \mathcal{T}(\mathcal{K}_1)$, $(\sigma, T, t_2) \in \mathcal{T}(\mathcal{K}_2)$, and $t_1 \prec_t t_2$. If $\mathcal{K}_1 \not\prec_{\sigma,T} \mathcal{K}_2$ and $\mathcal{K}_2 \not\prec_{\sigma,T} \mathcal{K}_1$, we write $\mathcal{K}_1 \equiv_{\sigma,T} \mathcal{K}_2$.

Intuitively, a belief changes after revision if it is added to or removed from the belief base, or if its associated grade changes. Similarly, trust in a source regarding a topic changes after revision if the associated trust grade changes.

4 Relevant Change

As proposed earlier, the degrees of trust in sources depend on the degrees of belief in formulas conveyed by these sources and vice versa. To model such dependence, we need to keep track of which formulas and which sources are “relevant” to each other. First, we recall a piece of terminology due to [11]: $\Gamma \subset \mathcal{L}$ is a ϕ -kernel if $\Gamma \models \phi$ and, for every $\Delta \subset \Gamma$, $\Delta \not\models \phi$.

Definition 5. Let \mathcal{K} be an information state. The **support graph** $\mathfrak{G}(\mathcal{K}) = (\mathcal{S}_{\mathcal{K}} \cup \Phi_{\mathcal{K}}, \mathcal{E})$ is such that $(u, v) \in \mathcal{E}$ if and only if

1. $u \in \mathcal{S}_{\mathcal{K}}$, $v \in \Phi_{\mathcal{K}}$, and $v \in u(\mathcal{H}(\mathcal{K}))$;
2. $u \in \Phi_{\mathcal{K}}$, $v \in \Phi_{\mathcal{K}}$, $u \neq v$, and $u \in \Gamma \subseteq \Phi_{\mathcal{K}}$ where Γ is a v -kernel; or
3. $u \in \Phi_{\mathcal{K}}$, $v \in \mathcal{S}_{\mathcal{K}}$, and $(v, u) \in \mathcal{E}$.

A node u **supports** a node v if there is a simple path from u to v .

Figure 1 shows an example of the support graph for the information state presented in Example 1. Source σ_1 conveyed both P and S , thus, according to clause 1 in Definition 5, there is an edge from σ_1 to both P and S . Moreover, there is an edge from both P and S to σ_1 given the third clause in Definition 5. Similarly, there are edges from σ_2 to both Q and $Q \rightarrow P$ as well as from Q and $Q \rightarrow P$ to σ_2 . Finally, $\{Q \rightarrow P, Q\}$ is a P -kernel. Hence, we have an edge from Q to P and from $Q \rightarrow P$ to P given the second clause in the definition of the support graph.

The support graph allows us to trace back and propagate changes in trust and belief to relevant beliefs and information sources along support paths. Instances of support may be classified according to the type of relation.

Observation 1. Let \mathcal{K} be an information state.

1. $\phi \in \Phi_{\mathcal{K}}$ supports $\psi \in \Phi_{\mathcal{K}}$ if and only if $\phi \neq \psi$ and (i) $\phi \in \Gamma \subseteq \Phi_{\mathcal{K}}$ where Γ is a ψ -kernel or (ii) ϕ supports some $\sigma \in \mathcal{S}_{\mathcal{K}}$ which supports ψ .
2. $\phi \in \Phi_{\mathcal{K}}$ supports $\sigma \in \mathcal{S}_{\mathcal{K}}$ if and only if $\psi \in \sigma(\mathcal{H}(\mathcal{K}))$ and $\phi \in \Gamma \subseteq \Phi_{\mathcal{K}}$ where Γ is a ψ -kernel or ϕ supports some $\sigma' \in \mathcal{S}_{\mathcal{K}}$ which supports σ .
3. $\sigma \in \mathcal{S}_{\mathcal{K}}$ supports $\phi \in \Phi_{\mathcal{K}}$ if and only if $\psi \in \sigma(\mathcal{H}(\mathcal{K}))$ and $\psi \in \Gamma \subseteq \Phi_{\mathcal{K}}$ where Γ is a ϕ -kernel or σ supports some $\sigma' \in \mathcal{S}_{\mathcal{K}}$ which supports ϕ .
4. $\sigma \in \mathcal{S}_{\mathcal{K}}$ supports $\sigma' \in \mathcal{S}_{\mathcal{K}}$ if and only if $\sigma \neq \sigma'$ σ supports some $\phi \in \Phi_{\mathcal{K}}$ which supports σ' .

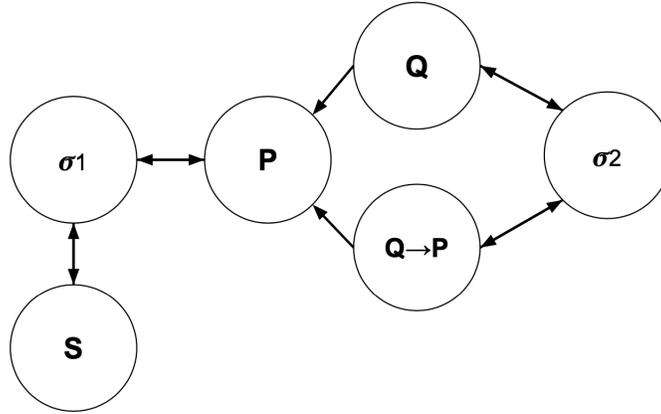


Fig. 1: The support graph where σ_2 conveyed Q and $Q \rightarrow P$ that both logically imply P which was conveyed alongside S by σ_1 .

Thus, given the first three clauses, the support relation from a formula to a formula, a formula to a source, or a source to formula may be established in two ways: (i) either purely logically via a path of only formulas or (ii) with the aid of a trust link via an intermediate source. A source can only support a source, however, by supporting a formula which supports that other source. Note that self-support is avoided by requiring support paths to be simple.

The support graph provides the basis for constructing an operator of rational information revision. Traditionally, belief revision is concerned with minimal change [10,12]. In this paper, we model minimality using relevance. However, our notion of relevance is not restricted to *logical* relevance as with classical belief revision; it also accounts for *source* relevance. When an information state \mathcal{K} is revised with formula ϕ conveyed by source σ , we would want to confine changes in belief and trust to formulas and sources relevant to ϕ , $\neg\phi$, and σ .

Definition 6. Let \mathcal{K} be an information state and u and v be nodes in $\mathfrak{G}(\mathcal{K})$. u is *v-relevant* if u supports v or v supports u . Further, if $\phi, \psi \in \mathcal{L}$ with $\Gamma_\phi \subseteq \Phi_{\mathcal{K}}$ a ϕ -kernel and $\Gamma_\psi \subseteq \Phi_{\mathcal{K}}$ a ψ -kernel where u is v -relevant for some $u \in \Gamma_\phi$ and $v \in \Gamma_\psi$, then ϕ is *ψ -relevant*.

Observation 2. If \mathcal{K} is an information state where u is v -relevant then (i) v is u -relevant; (ii) if $v \in \sigma(\mathcal{H}(\mathcal{K}))$ and $u \neq \sigma$, then u is σ -relevant; and (iii) if $v \in \mathcal{S}_{\mathcal{K}}$, $\phi \in v(\mathcal{H}(\mathcal{K}))$, and $u \neq \phi$, then, u is ϕ -relevant.

Hence, relevance is a symmetric relation. Crucially, if σ conveys ϕ , then the formulas and sources relevant to ϕ (other than σ) are exactly the formulas and sources relevant to σ (other than ϕ). For this reason, when revising with a conveyance of ϕ by σ it suffices to consider only ϕ -relevant (and $\neg\phi$ -relevant) formulas and sources.

5 Information Revision

5.1 Intuitions

Table 1 shows the possible reasonable effects on $\mathcal{B}(\mathcal{K})$ as agent \mathcal{A} revises its information state \mathcal{K} with (ϕ, σ) ; \mathcal{K}_\times is shorthand for $\mathcal{K} \times (\phi, \sigma)$ while “neither” means that neither ϕ nor $\neg\phi$ is in $Cn(For(\mathcal{B}(\mathcal{K})))$.

#	\mathcal{K}	\mathcal{K}_\times	ϕ	$\neg\phi$	Notes
B_1	neither	$(\phi, b) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K} \prec_\phi \mathcal{K}_\times$	$\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_\times$	-
B_2	neither	neither	$\mathcal{K} \equiv_\phi \mathcal{K}_\times$	$\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_\times$	-
B_3	$(\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\phi, b_2) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K} \prec_\phi \mathcal{K}_\times$	$\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_\times$	$b_1 \prec_b b_2$
B_4	$(\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\phi, b_1) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K} \equiv_\phi \mathcal{K}_\times$	$\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_\times$	-
B_5	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\phi, b_2) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K} \prec_\phi \mathcal{K}_\times$	$\mathcal{K}_\times \prec_{\neg\phi} \mathcal{K}$	-
B_6	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K})$	neither	$\mathcal{K} \equiv_\phi \mathcal{K}_\times$	$\mathcal{K}_\times \prec_{\neg\phi} \mathcal{K}$	-
B_7	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\neg\phi, b_2) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K} \equiv_\phi \mathcal{K}_\times$	$\mathcal{K}_\times \prec_{\neg\phi} \mathcal{K}$	$b_2 \prec_b b_1$
B_8	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K} \equiv_\phi \mathcal{K}_\times$	$\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_\times$	-
B_9	$\mathcal{B}(\mathcal{K}) \models \{(\neg\phi, b_1), (\phi, b_2)\}$	$(\phi, b_3) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K} \prec_\phi \mathcal{K}_\times$	$\mathcal{K}_\times \prec_{\neg\phi} \mathcal{K}$	$b_2 \prec_b b_3$
B_{10}	$\mathcal{B}(\mathcal{K}) \models \{(\neg\phi, b_1), (\phi, b_2)\}$	$(\phi, b_2) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K} \equiv_\phi \mathcal{K}_\times$	$\mathcal{K}_\times \prec_{\neg\phi} \mathcal{K}$	-
B_{11}	$\mathcal{B}(\mathcal{K}) \models \{(\neg\phi, b_1), (\phi, b_2)\}$	$(\phi, b_3) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K}_\times \prec_\phi \mathcal{K}$	$\mathcal{K}_\times \prec_{\neg\phi} \mathcal{K}$	$b_3 \prec_b b_2$
B_{12}	$\mathcal{B}(\mathcal{K}) \models \{(\neg\phi, b_1), (\phi, b_2)\}$	neither	$\mathcal{K}_\times \prec_\phi \mathcal{K}$	$\mathcal{K}_\times \prec_{\neg\phi} \mathcal{K}$	-
B_{13}	$\mathcal{B}(\mathcal{K}) \models \{(\neg\phi, b_1), (\phi, b_2)\}$	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K}_\times \prec_\phi \mathcal{K}$	$\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_\times$	-
B_{14}	$\mathcal{B}(\mathcal{K}) \models \{(\neg\phi, b_1), (\phi, b_2)\}$	$(\neg\phi, b_3) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K}_\times \prec_\phi \mathcal{K}$	$\mathcal{K}_\times \prec_{\neg\phi} \mathcal{K}$	$b_3 \prec_b b_1$

Table 1: The admissible scenarios of belief revision.

In Cases B_1 and B_2 , \mathcal{A} initially believes neither ϕ nor $\neg\phi$. After revision, B_1 shows the case where \mathcal{A} believes ϕ as it has no evidence for the contrary. However, another possible scenario would be that portrayed in Case B_2 where \mathcal{A} decides that the weight of evidence for and against ϕ is comparable and hence ϕ can not be accepted.

In the next two cases, \mathcal{A} already believes ϕ and revision with ϕ confirms what is already believed. Thus, ϕ could become more entrenched as now σ also supports ϕ (Case B_3) or ϕ 's degree may remain unchanged (Case B_4) which could occur when ϕ is believed with the maximum degree of belief, if there is such a degree, or when ϕ has only been ever conveyed by σ , who are now only confirming itself. In this latter case, \mathcal{A} might choose not to increase the degree of belief in ϕ .

\mathcal{A} believes $\neg\phi$ before revision in Cases B_5 – B_8 . Thus, revising with the conflicting piece of information ϕ could yield the following outcomes: (i) \mathcal{A} stops believing $\neg\phi$ and starts believing ϕ (Case B_5), which could occur if, for example, σ is a highly trusted source; (ii) \mathcal{A} decides that there is not enough evidence to believe either ϕ or $\neg\phi$ (Case B_6); (iii) \mathcal{A} decides not to completely give up $\neg\phi$ but there is enough evidence to make \mathcal{A} doubt $\neg\phi$ making it less entrenched (Case B_7); or (iv) \mathcal{A} decides not to change its beliefs, even when provided with ϕ (Case B_8). One scenario where this is possible is when the source is not trusted and so \mathcal{A} decides not to consider this instance of conveyance.

In general, revision may be applied to an inconsistent belief base. In such cases, what matters is retaining consistency. This is accompanied by rejecting ϕ , $\neg\phi$, or actually both. Depending on the weight of evidence and trust in σ this may result in ϕ 's becoming more entrenched (Case B_9) or not (Case B_{10}). It may also happen that ϕ becomes less entrenched (Case B_{11}), for example when some supporters of ϕ are also $\neg\phi$ -relevant. Although revision may prefer $\neg\phi$ over ϕ (Case B_{13} and Case B_{14}) this should not make $\neg\phi$ more entrenched, since revision with ϕ provides no positive evidence for its negation. Case B_{12} shows \mathcal{A} contracting both ϕ and $\neg\phi$ upon deciding that evidence for (and against) each one is comparable.

Other cases, we believe, should be forbidden for a rational operation of information revision. These cases are presented in Table 2. In Case B_{15} , \mathcal{A} is neutral. However, when provided with evidence for ϕ , surprisingly, \mathcal{A} starts believing $\neg\phi$. Cases B_{16-18} are ones where \mathcal{A} already believes ϕ and, on revising with ϕ , the following occurred: (i) ϕ becomes less entrenched (Case B_{16}); (ii) \mathcal{A} stops believing in ϕ (Case B_{17}); (iii) an extreme case where \mathcal{A} receives a confirmation for the already believed ϕ , and $\neg\phi$ ends up being believed (Case B_{18}).⁹ Cases B_{19} and B_{20} show a scenario where $\neg\phi$ becomes more entrenched. These cases are forbidden since revising with ϕ provides no support for $\neg\phi$.

#	\mathcal{K}	\mathcal{K}_\times	ϕ	$\neg\phi$	Notes
B_{15}	neither	$(\neg\phi, b) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K}_\times \equiv_\phi \mathcal{K}$	$\mathcal{K} \prec_{\neg\phi} \mathcal{K}_\times$	-
B_{16}	$(\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\phi, b_2) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K}_\times \prec_\phi \mathcal{K}$	$\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_\times$	$b_2 \prec_b b_1$
B_{17}	$(\phi, b) \in \mathcal{B}(\mathcal{K})$	neither	$\mathcal{K}_\times \prec_\phi \mathcal{K}$	$\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_\times$	-
B_{18}	$(\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\neg\phi, b_2) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K}_\times \prec_\phi \mathcal{K}$	$\mathcal{K} \prec_{\neg\phi} \mathcal{K}_\times$	-
B_{19}	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\neg\phi, b_2) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K} \equiv_\phi \mathcal{K}_\times$	$\mathcal{K} \prec_{\neg\phi} \mathcal{K}_\times$	$b_1 \prec_b b_2$
B_{20}	$\mathcal{B}(\mathcal{K}) \models \{(\neg\phi, b_1), (\phi, b_2)\}$	$(\neg\phi, b_3) \in \mathcal{B}(\mathcal{K}_\times)$	$\mathcal{K}_\times \prec_\phi \mathcal{K}$	$\mathcal{K} \prec_{\neg\phi} \mathcal{K}_\times$	$b_1 \prec_b b_3$

Table 2: The forbidden scenarios of belief revision.

Similar to our treatment of belief revision, in Table 3, we present the possible scenarios for trust change. The cases in Table 3 are based on the admissible cases of belief change presented in Table 1. In Table 3, \mathcal{K}_\times is also a shorthand for $\mathcal{K} \times (\phi, \sigma)$, θ is any ϕ -relevant and not $\neg\phi$ -relevant source, and η is any $\neg\phi$ -relevant and not ϕ -relevant source where $\sigma \neq \theta \neq \eta$.

If ϕ becomes more entrenched, then some new support was provided for ϕ . Hence, θ (or σ) might become more trusted (in light of the new support) or remain unchanged as the new evidence could be not enough to change trust. Regardless, σ and θ should not be less trusted. Also, if $\neg\phi$ becomes less entrenched as a result of revising with ϕ , it means that ϕ affected the revision even if it was not accepted. Thus σ and θ should not become less trusted. In both cases, there is no reason to for η to be more trusted. The foregoing is captured in Case T_1 . If ϕ is rejected ($\phi \notin \text{For}(\mathcal{B}(\mathcal{K}_\times))$) and at the same

⁹ These extreme reactions to revising with ϕ may perhaps be justified by complete mistrust in σ . We do not address mistrust in this paper, though.

#	Condition	Belief Revision Cases	σ	θ	η
T_1	$\mathcal{K} \prec_{\phi} \mathcal{K}_{\times}$ or $\mathcal{K}_{\times} \prec_{\neg\phi} \mathcal{K}$	$B_1, B_3, B_{5-7}, B_9, B_{10}$	$\mathcal{K}_{\times} \not\prec_{\sigma, T} \mathcal{K}$	$\mathcal{K}_{\times} \not\prec_{\theta, T} \mathcal{K}$	$\mathcal{K} \not\prec_{\eta, T} \mathcal{K}_{\times}$
T_2	ϕ rejected and $\mathcal{K}_{\times} \not\prec_{\neg\phi} \mathcal{K}$	B_2, B_8, B_{13}	$\mathcal{K} \not\prec_{\sigma, T} \mathcal{K}_{\times}$	$\mathcal{K} \not\prec_{\theta, T} \mathcal{K}_{\times}$	$\mathcal{K} \equiv_{\eta, T} \mathcal{K}_{\times}$
T_3	$\mathcal{K} \equiv_{\phi} \mathcal{K}_{\times}$ and $\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_{\times}$	B_4	$\mathcal{K} \equiv_{\sigma, T} \mathcal{K}_{\times}$	$\mathcal{K} \equiv_{\theta, T} \mathcal{K}_{\times}$	$\mathcal{K} \equiv_{\eta, T} \mathcal{K}_{\times}$
T_4	$\mathcal{K}_{\times} \prec_{\phi} \mathcal{K}$ and $\mathcal{K}_{\times} \prec_{\neg\phi} \mathcal{K}$	B_{11}, B_{12}, B_{14}	$\mathcal{K} \not\prec_{\sigma, T} \mathcal{K}_{\times}$	$\mathcal{K} \not\prec_{\theta, T} \mathcal{K}_{\times}$	$\mathcal{K} \not\prec_{\eta, T} \mathcal{K}_{\times}$

Table 3: The admissible scenarios of trust revision.

time $\neg\phi$ does not get affected, then σ and θ should not become more trusted while trust in η should stay the same as no new support was provided for/against it (Case T_2).¹⁰ If there is no change in the entrenchment of ϕ and $\neg\phi$, then there is no reason to change trust in any of σ , θ , or η (Case T_3). Finally, Case T_4 shows that if both ϕ and $\neg\phi$ become less entrenched, then there is no reason for σ , θ , or η to become more trusted. The forbidden cases of trust revision constitute all scenarios that contradict any of the aforementioned admissible cases.

Note that, in none of the cases in Table 3, do we require that trust *should* change in certain ways, only that it should not. We believe it be unwise to postulate sufficient conditions for trust change in a generic information revision operation. For example, one might be tempted to say that, if after revision with ϕ , $\neg\phi$ is no longer believed, then trust in any source supporting $\neg\phi$ *should* decrease. Things are not that straightforward, though.

Example 2. Agent \mathcal{A} 's belief base is $\{(S \rightarrow P, b_1), (Q \rightarrow \neg S, b_2)\}$. Source *Jordan*, conveys P then conveys Q . Since \mathcal{A} has no evidence against either, it believes both. Now, source *Nour*, who is more trusted than *Jordan*, conveys S . Consequently, \mathcal{A} starts believing S despite having evidence against it. To maintain consistency, \mathcal{A} also stops believing Q (because it supports $\neg S$). What should happen to \mathcal{A} 's trust in *Jordan*? We might, at first glance, think that trust in *Jordan* should decrease as he conveyed Q which is no longer believed. However, one could also argue that trust in *Jordan* should increase because he conveyed P , which is now being confirmed by *Nour*.

This example shows that setting *general* rules for how trust *must* change is almost impossible, as it depends on several factors. Whether \mathcal{A} ends up trusting *Jordan* less, more, or without change appears to depend on how the particular revision operators manipulates grades. The situation becomes more complex if the new conveyance by *Nour* supports several formulas supporting *Jordan* and refutes several formulas supported by him. In this case, how trust in *Jordan* changes (or not) would also depend on how the effects of all these support relations are aggregated. We contend that such issues should not, and cannot, be settled by general constraints on information revision.

This non-determinism about how trust changes extends to similar non-determinism about how belief changes. According to Observation 1, a formula ϕ may support another formula ψ by transitivity through an intermediate source σ . Given that, in general, the

¹⁰ In Case B_2 where \mathcal{A} is neutral and ϕ is rejected after revision, a particular operator might actually choose to decrease trust in σ . An example of this is when σ informs \mathcal{A} about ϕ and \mathcal{A} is sure that there is no way for anyone to know either ϕ or $\neg\phi$ (Schrödinger's cat).

effect of revising with ϕ on σ is non-deterministic, then so is its effect on ψ . Hence, the postulates to follow only provide *necessary* conditions for different ways belief and trust may change; the general principle being that the scope of change on revising with ϕ is limited to formulas and sources which are ϕ - and $\neg\phi$ -relevant. Postulating sufficient conditions is, we believe, ill-advised.

5.2 Postulates

In the sequel, where ϕ is a formula and σ is a source, a σ -independent ϕ -kernel is, intuitively, a ϕ -kernel that would still exist if σ did not exist. More precisely, for every $\psi \in \Gamma$, ψ is supported by some $\sigma'' \neq \sigma$, or ψ has no source. Of course, all formulas are conveyed by sources. However, given a forgetful filter, record of sources for some formulas may be missing from the history.

We believe a rational information revision operator \times should observe the following postulates on revising an information state \mathcal{K} with (ϕ, σ) and $\phi \in T$ where T is a topic. The postulates are a formalization of the intuitions outlined earlier.

- (\times_1 : **Consistency**) $Cn(For(\mathcal{B}(\mathcal{K} \times (\phi, \sigma)))) \neq \mathcal{L}$.
- (\times_2 : **Resilience**) If $Cn(\{\phi\}) = \mathcal{L}$, then $\mathcal{K} \not\prec_{\sigma, T} \mathcal{K} \times (\phi, \sigma)$.
- (\times_3 : **Supported Entrenchment**) $\mathcal{K} \times (\phi, \sigma) \prec_{\phi} \mathcal{K}$ only if $Cn(For(\mathcal{B}(\mathcal{K}))) = \mathcal{L}$.
- (\times_4 : **Opposed Entrenchment**) $\mathcal{K} \not\prec_{\neg\phi} \mathcal{K} \times (\phi, \sigma)$.
- (\times_5 : **Positive Relevance**) If $\mathcal{K} \prec_{\sigma', T} \mathcal{K} \times (\phi, \sigma)$ and $\phi \in For(\mathcal{B}(\mathcal{K} \times (\phi, \sigma)))$, then
 1. $\sigma' \neq \sigma$ is supported by ϕ ; or
 2. $\sigma' = \sigma$ and there is $\Gamma \subseteq For(\mathcal{B}(\mathcal{K}))$ where Γ is a σ -independent ϕ -kernel.
- (\times_6 : **Negative Relevance**) If $\mathcal{K} \times (\phi, \sigma) \prec_{\sigma', T} \mathcal{K}$, then
 1. $\phi \in For(\mathcal{B}(\mathcal{K} \times (\phi, \sigma)))$ and σ' is $\neg\phi$ -relevant; or
 2. $\sigma' = \sigma$, but, there is $\Gamma \subseteq For(\mathcal{B}(\mathcal{K} \times (\phi, \sigma)))$ where Γ is a $\neg\phi$ -kernel.
- (\times_7 : **Belief Confirmation**) If $\mathcal{K} \prec_{\psi} \mathcal{K} \times (\phi, \sigma)$, then, $\psi \neq \phi$ is supported by ϕ .
- (\times_8 : **Belief Refutation**) If $\mathcal{K} \times (\phi, \sigma) \prec_{\psi} \mathcal{K}$, then
 1. ψ is $\neg\phi$ -relevant and $\phi \in Cn(For(\mathcal{B}(\mathcal{K} \times (\phi, \sigma))))$ or $\mathcal{K} \times (\phi, \sigma) \prec_{\neg\phi} \mathcal{K}$; or
 2. ψ is ϕ -relevant and $\phi \notin Cn(For(\mathcal{B}(\mathcal{K} \times (\phi, \sigma))))$ or $\mathcal{K} \times (\phi, \sigma) \prec_{\phi} \mathcal{K}$.

A revised information state is consistent even if the revising formula is itself contradictory (\times_1). If ϕ is inconsistent, σ should not become more trusted¹¹ (\times_2). Following Tables 1 and 2, ϕ cannot become less entrenched unless the belief base is inconsistent (\times_3). $\neg\phi$, even if the belief base is inconsistent, should never become more entrenched (\times_4). If σ' is more trusted after revision, then (i) ϕ succeeds and (ii) either σ' is different from σ and supported by ϕ or σ' is σ and there is independent believed evidence for ϕ (\times_5). If σ' is less trusted after revision, then it must be either that ϕ succeeds and σ' (possibly identical to σ) is relevant to $\neg\phi$, or that σ' is σ and there is believed evidence for $\neg\phi$ that leads to rejecting ϕ (\times_6). ψ is more entrenched after revision only if it is supported by ϕ (\times_7). ψ is less entrenched after revision only if it is relevant to ϕ or $\neg\phi$ (or both) and the one it is relevant to is not favored by the revision (\times_8).

Given the definition of information states, support graphs, and the postulates outlined earlier, the following observations hold.

¹¹ A specific operator might choose to actually decrease trust in a source that conveys contradictions as this is a proof of its unreliability.

Observation 3. *If $\phi \in Cn(For(\mathcal{B}(\mathcal{K}))) \neq \mathcal{L}$, then $\phi \in Cn(For(\mathcal{B}(\mathcal{K} \times (\phi, \sigma))))$.*

Observation 4. *If $\neg\phi \notin Cn(For(\mathcal{B}(\mathcal{K})))$, then $\neg\phi \notin Cn(For(\mathcal{B}(\mathcal{K} \times (\phi, \sigma))))$.*

Observation 5. *If $(\psi, b_1) \in \mathcal{B}(\mathcal{K})$, $(\psi, b_2) \in \mathcal{B}(\mathcal{K} \times (\phi, \sigma))$ where $b_1 \neq b_2$, then ψ is either ϕ -relevant or $\neg\phi$ -relevant.*

Observation 6. *If $(\sigma', T, t_1) \in \mathcal{T}(\mathcal{K})$, $(\sigma', T, t_2) \in \mathcal{T}(\mathcal{K} \times (\phi, \sigma))$ and $t_1 \neq t_2$, then σ' is ϕ -relevant or $\neg\phi$ -relevant.*

Observation 7. *If $\phi \notin For(\mathcal{B}(\mathcal{K} \times (\phi, \sigma)))$, then there is no $\sigma' \in S_{\mathcal{K}}$ such that $\mathcal{K} \prec_{\sigma', T} \mathcal{K} \times (\phi, \sigma)$.*

Observation 8. *If $Cn(For(\mathcal{B}(\mathcal{K}))) \neq \mathcal{L}$, then an operator that observes \times_4 and \times_5 allows for only cases in Table 1 to occur.*

6 Conclusion and Future Work

It is our conviction that belief and trust revision are intertwined processes that should not be separated. Hence, in this paper, we argued why that is the case and provided a model for performing the joint belief-trust (information) revision with minimal assumptions on the modeling language. Then, we introduced the notion of information states that allows for the representation of information in a way that facilitates the revision process. We also introduced the support graph which is a formal structure that allows us to identify relevance relations between not only formulas, but also, information sources. Finally, we proposed the postulates that we believe any rational information revision operator should observe. Future work could go in one or more of the following directions:

1. We intend to define a representation theorem for the postulates we provided.
2. We intend to further investigate conveyance and information acquisition to further allow agents to trust/mistrust their own perception(s).
3. Lastly, we would like to add desires, intentions, and other mental attitudes to create a unified revision theory for all mental attitudes.

References

1. Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. *The journal of symbolic logic* **50**(2), 510–530 (1985)
2. Booth, R., Hunter, A.: Trust as a precursor to belief revision. *Journal of Artificial Intelligence Research* **61**, 699–722 (2018)
3. Castelfranchi, C., Falcone, R.: Principles of trust for MAS: Cognitive anatomy, social importance, and quantification. In: *Proceedings International Conference on Multi Agent Systems*. pp. 72–79. IEEE (1998)
4. Demolombe, R.: To trust information sources: a proposal for a modal logical framework. In: Castelfranchi, C., Tan, Y.H. (eds.) *Trust and Deception in Virtual Societies*, pp. 111–124. Springer Netherlands, Dordrecht (2001)

5. Demolombe, R.: Reasoning about trust: A formal logical framework. In: Jensen, C., Poslad, S., Dimitrakos, T. (eds.) *Trust Management*. pp. 291–303. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
6. Demolombe, R., Liao, C.J.: A logic of graded trust and belief fusion. In: *Proceedings of the 4th workshop on deception, fraud and trust in agent societies*. pp. 13–25 (2001)
7. van Ditmarsch, H.: Dynamics of lying. *Synthese* **191**(5), 745–777 (2014)
8. Elangovan, A., Auer-Rizzi, W., Szabo, E.: Why don't I trust you now? An attributional approach to erosion of trust. *Journal of Managerial Psychology* **22**(1), 4–24 (2007)
9. Falcone, R., Castelfranchi, C.: Social trust: A cognitive approach. In: *Trust and Deception in Virtual Societies*, pp. 55–90. Springer Netherlands (2001)
10. Gärdenfors, P., Makinson, D.: Revisions of knowledge systems using epistemic entrenchment. In: *Proceedings of the 2nd conference on Theoretical aspects of reasoning about knowledge*. pp. 83–95. Morgan Kaufmann Publishers Inc. (1988)
11. Hansson, S.O.: Kernel contraction. *The Journal of Symbolic Logic* **59**(3), 845–859 (1994)
12. Hansson, S.O.: A survey of non-prioritized belief revision. *Erkenntnis* **50**(2-3), 413–427 (1999)
13. Hansson, S.O.: *A textbook of belief dynamics - theory change and database updating*, Applied logic series, vol. 11. Kluwer (1999)
14. Hardwig, J.: The role of trust in knowledge. *The Journal of Philosophy* **88**(12), 693–708 (1991)
15. Haselhuhn, M.P., Schweitzer, M.E., Wood, A.M.: How implicit beliefs influence trust recovery. *Psychological Science* **21**(5), 645–648 (2010)
16. Herzog, A., Lorini, E., Hübner, J.F., Vercouter, L.: A logic of trust and reputation. *Logic Journal of the IGPL* **18**(1), 214–244 (2010)
17. Holton, R.: Deciding to trust, coming to believe. *Australasian Journal of Philosophy* **72**(1), 63–76 (1994)
18. Ismail, H., Attia, P.: Towards a logical analysis of misleading and trust erosion. In: Gordon, A.S., Miller, R., Turán, G. (eds.) *Proceedings of the Thirteenth International Symposium on Commonsense Reasoning, COMMONSENSE 2017, London, UK, November 6-8, 2017*. vol. 2052. CEUR-WS.org (2017)
19. Jones, A.J.: On the concept of trust. *Decision Support Systems* **33**(3), 225–232 (2002)
20. Jones, A.J., Firozabadi, B.S.: On the characterisation of a trusting agent—aspects of a formal approach. In: Castelfranchi, C., Tan, Y.H. (eds.) *Trust and Deception in Virtual Societies*, pp. 157–168. Springer Netherlands, Dordrecht (2001)
21. Jøsang, A., Ivanovska, M., Muller, T.: Trust revision for conflicting sources. In: *The 18th International Conference on Information Fusion (Fusion 2015)*. pp. 550–557. IEEE (2015)
22. Katz, Y., Golbeck, J.: Social network-based trust in prioritized default logic. In: *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI 2006)*. pp. 1345–1350 (2006)
23. Leturc, C., Bonnet, G.: A normal modal logic for trust in the sincerity. In: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*. pp. 175–183. International Foundation for Autonomous Agents and Multiagent Systems (2018)
24. Liao, C.J.: Belief, information acquisition, and trust in multi-agent systems—a modal logic formulation. *Artificial Intelligence* **149**(1), 31–60 (2003)
25. Lorini, E., Demolombe, R.: From binary trust to graded trust in information sources: A logical perspective. In: *International Workshop on Trust in Agent Societies*. pp. 205–225. Springer (2008)
26. Lorini, E., Jiang, G., Perrussel, L.: Trust-based belief change. In: T. Schaub, G. Friedrich, B.O. (ed.) *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI 2014)*. *Frontiers in Artificial Intelligence and Applications*, vol. 263, pp. 549–554. IOS Press, Amsterdam (2014)

27. McLeod, C.: Trust. In: Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2015 edn. (2015)
28. Rodenhäuser, L.B.: *A matter of trust: Dynamic attitudes in epistemic logic*. Universiteit van Amsterdam [Host] (2014)
29. Sabater, J., Sierra, C.: Review on computational trust and reputation models. *Artificial Intelligence Review* **24**(1), 33–60 (Sep 2005)
30. Sakama, C.: A formal account of deception. In: *2015 AAAI Fall Symposia*, Arlington, Virginia, USA, November 12-14, 2015. pp. 34–41. AAAI Press (2015)
31. Sakama, C., Caminada, M., Herzig, A.: A logical account of lying. In: *European Workshop on Logics in Artificial Intelligence*. pp. 286–299. Springer (2010)
32. Simpson, J.A.: Psychological foundations of trust. *Current Directions in Psychological Science* **16**(5), 264–268 (2007)
33. Yasser, A., Ismail, H.O.: Information revision: The joint revision of belief and trust. In: Varzinczak, I.J., Martinez, M.V. (eds.) *Proceedings of the 18th International Workshop on Non-Monotonic Reasoning (NMR2020)* (2020), forthcoming