Information Revision: The Joint Revision of Belief and Trust

2 authors:

Ammar Yasser
The German University in Cairo
2 PUBLICATIONS 0 CITATIONS

Haythem Osman Ismail
The German University in Cairo
40 PUBLICATIONS 194 CITATIONS

Information Revision: The Joint Revision of Belief and Trust

Ammar Yasser[1] , **Haythem O. Ismail**[2,1]
[1]German University in Cairo, Egypt
[2]Cairo University, Egypt
ammar.abbas@guc.edu.eg, haythem.ismail@guc.edu.eg

Abstract

Most of our decisions are guided by trust, specifically decisions about what to believe and what not to believe. We accept information from sources we trust and doubt information from sources we do not trust and, in general, rely on trust in revising our beliefs. While we may have difficulty defining exactly what trust is, we can, on one hand, rather easily explain why we trust or mistrust someone and, on the other hand, occasionally revise how much we trust them. In this paper, we propose that trust revision and belief revision are inseparable processes. We address issues concerning the formalization of trust in information sources and provide AGM-style postulates for rational joint revision of the two attitudes. In doing so, we attempt to fill a number of gaps in the literature on trust, trust revision, and their relation to belief 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 (Castelfranchi and Falcone 1998; Falcone and Castelfranchi 2001; Jones and Firozabadi 2001; Jones 2002; Sabater and Sierra 2005; Katz and Golbeck 2006, for instance), psychology (Simpson 2007; Elangovan, Auer-Rizzi, and Szabo 2007; Haselhuhn, Schweitzer, and Wood 2010, for instance), and philosophy (Holton 1994; Hardwig 1991; McLeod 2015, for instance). Crucially, it was also investigated in the logic-based artificial intelligence (AI) literature by several authors (Demolombe 2001; Demolombe and Liau 2001; Liau 2003; Katz and Golbeck 2006; Herzig et al. 2010; Drawel, Bentahar, and Shakshuki 2017; Leturc and Bonnet 2018). 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 (Sakama, Caminada, and Herzig 2010; van Ditmarsch 2014; Sakama 2015; Ismail and Attia 2017, 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 (Alchourrón, Gärdenfors, and Makinson 1985; Hansson 1994; Darwiche and Pearl 1997; Van Benthem 2007, for example), while occasionally addressing trust-based revision of beliefs (Lorini, Jiang, and Perrussel 2014; Rodenhäuser 2014; Booth and Hunter 2018) does not have much to say about the revision of trust (but see (Liau 2003; Lorini, Jiang, and Perrussel 2014) 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. In Section 5, the formal inter-dependency of belief and trust is explored, culminating in AGM-style postulates for joint belief-trust revision. Finally, Section 6 presents an extended example highlighting some of the key concepts proposed in the paper.[1]

2 Trust and Belief

2.1 Trust in Information Sources

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 (McLeod 2015). You trust your doctor

[1]Because of space constraints, we were not able to provide all results in this paper. Hence, some selected proofs are available through this online appendix: proofs.

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 (Demolombe 2004; Lorini and Demolombe 2008) 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.[2] 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.[3] 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.

2.2 Joint Revision of Trust and Belief

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 (Alchourrón, Gärdenfors, and Makinson 1985; Hansson 1999a, 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 of the new piece of information. (Also see (Lorini, Jiang, and Perrussel 2014; Rodenhäuser 2014; Booth and Hunter 2018).) 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 order for an agent $\mathcal{A}$ to perform information revision, revising both its beliefs and its trust in sources, it needs to be able to recall more than just what it believes, or how much it trusts certain sources, as is most commonly the case in the literature. Hence, we introduce formal structures for representing information in a way that would facilitate information revision.

Definition 3.1. *An **information grading structure** $\mathcal{G}$ is a qunituple $(\mathcal{D}_b, \mathcal{D}_t, \prec_b, \prec_t, \delta)$, where $\mathcal{D}_b$ and $\mathcal{D}_t$ are non-empty, countable sets; $\prec_b$ and $\prec_t$ are, respectively, total orders over $\mathcal{D}_b$ and $\mathcal{D}_t$; and $\delta \in \mathcal{D}_t$.*

$\mathcal{D}_b$ and $\mathcal{D}_t$ contain the degrees of belief and trust, respectively. They are not necessarily finite, disjoint, different or identical.[4] 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. δ is interpreted as the *default* trust degree assigned to an information source with which the agent has no earlier experience.

Definition 3.2. *An **information structure** $\mathcal{I}$ is a quadruple $(\mathcal{L}, \mathcal{C}, \mathcal{S}, \mathcal{G})$, where*

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.*

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

[3] As suggested by (Ismail and Attia 2017), 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.

[4] $\mathcal{D}_b$ 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.

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 (Hansson 1999b). A topic represents the scope of trust. It 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 3.3. *Let $\mathcal{I} = (\mathcal{L}, \mathcal{C}, \mathcal{S}, (\mathcal{D}_b, \mathcal{D}_t, \prec_b, \prec_t, \delta))$ 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 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 $\mathcal{T}(\mathcal{K})$. This is a generalization to accommodate logics with an explicit account of trust in the object language (Demolombe and Liau 2001; Leturc and Bonnet 2018, for instance) as well as those without (Katz and Golbeck 2006; Jøsang, Ivanovska, and Muller 2015, for example). $\mathcal{H}(\mathcal{K})$ acts as a formal device for recording conveyance instances.[5] As with $\mathcal{T}(\mathcal{K})$, we do not require $\mathcal{L}$ to have an explicit account for conveying.[6]

With this setup, having trust on single propositions, as is most commonly the case in the literature (Demolombe 2004; Leturc and Bonnet 2018, 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.

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

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

Information revision is the process of revising an information state $\mathcal{K}$ with the conveyance of a formula ϕ by a source σ. 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 $\ltimes_{\mathcal{F}}$, where $\mathcal{F}$ is the associated filter. Revising $\mathcal{K}$ with a conveyance of ϕ by σ is denoted by $\mathcal{K} \ltimes_{\mathcal{F}} (\phi, \sigma)$. We require all revision operators $\ltimes_{\mathcal{F}}$ to have the same effect on the history:

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

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 $\varnothing \neq \mathcal{F} \subset \mathcal{S} \times \mathcal{L}$; and it is *memory-less* if $\mathcal{F} = \varnothing$. 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 $\ltimes_{\mathcal{F}}$ whenever this does not lead to ambiguity.

We now turn to what happens to the belief and trust bases of a revised information state. We start with two general definitions.

Definition 3.4. *Formula ϕ is* ***more entrenched*** *in state $\mathcal{K}_2$ over state $\mathcal{K}_1$, denoted $\mathcal{K}_1 \prec_\phi \mathcal{K}_2$ if*

1. *$\phi \notin Cn(For(\mathcal{B}(\mathcal{K}_1)))$ and $\phi \in Cn(For(\mathcal{B}(\mathcal{K}_2)))$ or*
2. *$(\phi, b_1) \in \mathcal{B}(\mathcal{K}_1)$, $(\phi, b_2) \in \mathcal{B}(\mathcal{K}_2)$, 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 $\mathcal{K}_1 \equiv_\phi \mathcal{K}_2$.

Definition 3.5. *Source 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 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. Hence, on changing the degree of belief in some formula ϕ, the degree of trust in a source σ, that previously conveyed ϕ, is likely to change. However, when the degree of trust in σ changes, the degrees of belief in formulas conveyed by σ might change as well. To model such behavior, 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 (Hansson 1994): $\Gamma \subset \mathcal{L}$ is a *ϕ-kernel* ($\phi \in \mathcal{L}$), $\Gamma \models \phi$ and, for every $\Delta \subset \Gamma$, $\Delta \not\models \phi$.

Definition 4.1. *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 following information state: Source σ_1 conveys ϕ that logically implies ψ which, in turn, is conveyed by σ_2. Hence,

[5] We chose $\mathcal{H}(\mathcal{K})$ to be a set for simplicity. However, it could be more beneficial if it was a sequence instead. An agent might need to record multiple conveyances by the same source for the same formula or distinguish more recents one without an explicit representation of time.

[6] The transfer of information is sometimes referred to using either "inform" or "communicate". In this paper, we use "convey" as a cover term for any modality of information transfer.

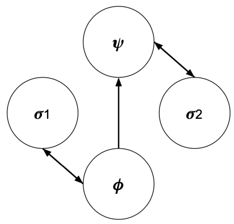

Figure 1: The support graph where σ_1 conveys ϕ which logically implies ψ which is conveyed by σ_2.

there is an edge from σ_1 to ϕ and from σ_2 to ψ given the first clause in the definition of the graph. Also, there is an edge from ϕ to ψ given the second clause. Finally, according to the last clause, there is an edge from both ϕ and ψ to σ_1 and σ_2, respectively. Note, for example, that ϕ supports ψ and that σ_1 supports σ_2. Intuitively, ϕ supports ψ directly by logically implying it; σ_1 supports σ_2 by virtue of conveying a formula (ϕ) which confirms a formula (ψ) conveyed by σ_2.

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 relata.

Observation 4.1. *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'.*

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 (Gärdenfors and Makinson 1988; Hansson 1999a). 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 want to confine changes in belief and trust to formulas and sources relevant to ϕ, $\neg\phi$, and σ.

Definition 4.2. *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.*

#	$\mathcal{K}$	$\mathcal{K}_{\ltimes}$	Notes
B_1	neither	$(\phi, b) \in \mathcal{B}(\mathcal{K}_{\ltimes})$	-
B_2	neither	neither	-
B_3	$(\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\phi, b_2) \in \mathcal{B}(\mathcal{K}_{\ltimes})$	$b_1 \prec_b b_2$
B_4	$(\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\phi, b_1) \in \mathcal{B}(\mathcal{K}_{\ltimes})$	-
B_5	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\phi, b_2) \in \mathcal{B}(\mathcal{K}_{\ltimes})$	-
B_6	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K})$	neither	-
B_7	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\neg\phi, b_2) \in \mathcal{B}(\mathcal{K}_{\ltimes})$	$b_2 \prec_b b_1$
B_8	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K}_{\ltimes})$	-

Table 1: The admissible scenarios of belief revision.

Observation 4.2. *Let $\mathcal{K}$ be an information state where u is v-relevant. The following are true.*

1. *v is u-relevant.*
2. *If $v \in \sigma(\mathcal{H}(\mathcal{K}))$ and $u \neq \sigma$, then u is σ-relevant.*
3. *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

Before formalizing the postulates of information revision, we start by presenting the intuitions of changing beliefs and trust that constitute the foundation of said formalization.

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); $\mathcal{K}_{\ltimes}$ is shorthand for $\mathcal{K} \ltimes (\phi, \sigma)$. The cases depend on whether $\phi \in Cn(For(\mathcal{B}(\mathcal{K})))$, $\neg\phi \in Cn(For(\mathcal{B}(\mathcal{K})))$, or neither ϕ or $\neg\phi$ is in $Cn(For(\mathcal{B}(\mathcal{K})))$. It is important to note that the "neither" cases are that strong only to simplify introducing the intuitions. For further simplicity, we only consider cases where $\mathcal{B}(\mathcal{K})$ (and, of course, $\mathcal{B}(\mathcal{K}_{\ltimes})$) is consistent so it is never the case that both ϕ and $\neg\phi$ are believed.

In Case B_1, $\mathcal{A}$ believes neither ϕ nor $\neg\phi$. Since $\mathcal{A}$ has no evidence to the contrary, on revising with ϕ, it is believed with some degree b, as now it is confirmed by a trusted source σ. Moreover, since $\neg\phi$ was neither refuted nor supported, it stays the same ($\mathcal{K} \prec_\phi \mathcal{K}_{\ltimes}$ and $\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_{\ltimes}$). As with Case B_1, in Case B_2, $\mathcal{A}$ is neutral about ϕ. However, on revision, $\mathcal{A}$ finds that the weight of evidence for and against ϕ are comparable so that it cannot accept ϕ ($\mathcal{K} \equiv_\phi \mathcal{K}_{\ltimes}$ and $\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_{\ltimes}$).

Unlike the previous two scenarios where $\mathcal{A}$ believed neither ϕ nor its negation, in Case B_3, $\mathcal{A}$ already believes ϕ with some degree b_1. Consequently, revision with ϕ confirms what is already believed. Since a new source σ now supports ϕ, it becomes more entrenched ($\mathcal{K} \prec_\phi \mathcal{K}_{\ltimes}$ and $\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_{\ltimes}$). On the other hand, in Case B_4, on revising with ϕ, despite $\mathcal{A}$'s already believing ϕ which is now being confirmed, ϕ does not become more entrenched ($\mathcal{K} \equiv_\phi \mathcal{K}_{\ltimes}$

#	$\mathcal{K}$	$\mathcal{K}_\ltimes$	Notes
B_9	neither	$(\neg\phi, b) \in \mathcal{B}(\mathcal{K}_\ltimes)$	-
B_{10}	$(\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\phi, b_2) \in \mathcal{B}(\mathcal{K}_\ltimes)$	$b_2 \prec_b b_1$
B_{11}	$(\phi, b) \in \mathcal{B}(\mathcal{K})$	neither	-
B_{12}	$(\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\neg\phi, b_2) \in \mathcal{B}(\mathcal{K}_\ltimes)$	-
B_{13}	$(\neg\phi, b_1) \in \mathcal{B}(\mathcal{K})$	$(\neg\phi, b_2) \in \mathcal{B}(\mathcal{K}_\ltimes)$	$b_1 \prec_b b_2$

Table 2: The forbidden scenarios of belief revision.

and $\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_\ltimes$). An example where this might occur is when ϕ is believed with the maximum degree of belief, if such degree exists, or when ϕ has only been ever conveyed by σ, who is now only confirming itself. In this latter case, $\mathcal{A}$ might choose not to increase the degree of belief in ϕ.

We now consider cases where $\mathcal{A}$ already believes $\neg\phi$. In Case B_5, revising with the conflicting piece of information ϕ coming from a highly-trusted source σ, $\mathcal{A}$ stops believing $\neg\phi$ and starts believing ϕ instead ($\mathcal{K} \prec_\phi \mathcal{K}_\ltimes$ and $\mathcal{K}_\ltimes \prec_{\neg\phi} \mathcal{K}$). Similarly, in Case B_6, $\mathcal{A}$ is presented with evidence against $\neg\phi$. After revision, $\mathcal{A}$ decides that there is not enough evidence to keep believing $\neg\phi$. However, there is also not enough evidence to believe ϕ ($\mathcal{K} \equiv_\phi \mathcal{K}_\ltimes$ and $\mathcal{K}_\ltimes \prec_{\neg\phi} \mathcal{K}$). Moreover, in Case B_7, $\mathcal{A}$ decides, on revision, that there is not enough evidence to completely give up $\neg\phi$. However, there is enough evidence to doubt $\neg\phi$ (decrease $\neg\phi$'s degree of belief) making it less entrenched ($\mathcal{K} \equiv_\phi \mathcal{K}_\ltimes$ and $\mathcal{K}_\ltimes \prec_{\neg\phi} \mathcal{K}$). On the contrary, in Case B_8, $\mathcal{A}$ decides that there is not enough evidence to change its beliefs, even when provided with ϕ, and hence $\neg\phi$ remains unchanged ($\mathcal{K} \equiv_\phi \mathcal{K}_\ltimes$ and $\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_\ltimes$). A possible scenario for this is when the source is not trusted and so $\mathcal{A}$ decides not to consider this instance of conveyance.

Other cases, we believe, should be forbidden for a rational operation of information revision. These cases are presented in Table 2.

$\mathcal{A}$ is neutral about ϕ in Case B_9. However, when provided with evidence for ϕ, $\mathcal{A}$ neither believes ϕ nor does it remain neutral. Surprisingly, $\mathcal{A}$ starts believing $\neg\phi$ ($\mathcal{K}_\ltimes \equiv_\phi \mathcal{K}$ and $\mathcal{K} \prec_{\neg\phi} \mathcal{K}_\ltimes$). ϕ is already believed in Case B_{10}. However, on getting a confirmation for ϕ, it becomes less entrenched ($\mathcal{K}_\ltimes \prec_\phi \mathcal{K}$ and $\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_\ltimes$). Similarly, in Case B_{11}, on receiving a confirmation for the already believed ϕ, $\mathcal{A}$ instead gives up believing ϕ ($\mathcal{K}_\ltimes \prec_\phi \mathcal{K}$ and $\mathcal{K} \equiv_{\neg\phi} \mathcal{K}_\ltimes$). An extreme case is that of Case B_{12} where $\mathcal{A}$, already believing ϕ, receives a confirmation thereof and upon revision, $\neg\phi$ ends up being believed ($\mathcal{K}_\ltimes \prec_\phi \mathcal{K}$ and $\mathcal{K} \prec_{\neg\phi} \mathcal{K}_\ltimes$). Finally, in Case B_{13}, $\mathcal{A}$ believes $\neg\phi$; but, when provided with evidence against it, it becomes more entrenched nevertheless ($\mathcal{K} \equiv_\phi \mathcal{K}_\ltimes$ and $\mathcal{K} \prec_{\neg\phi} \mathcal{K}_\ltimes$).

The cases in Table 2 may seem far fetched or even implausible. However, there is a line of reasoning that could accommodate such cases. Although, in this paper, we do not pursue this line or reasoning, it is at least worth a brief discussion. If agent $\mathcal{A}$ does not trust information source σ, $\mathcal{A}$ may be reluctant to believe what σ conveys, given no further supporting evidence; this much is perhaps uncontroversial. But if $\mathcal{A}$, not only does not trust σ, but strongly *mistrusts* them (given a long history of being misled by the malicious source), then $\mathcal{A}$ may reject what σ conveys and also believe its negation. To further illustrate this concept, consider a possible example.

Example 5.1. *Bob believes that there will be no classes tomorrow (ϕ) but he is not very certain about that. He meets Tim, who tells him "there will be no classes tomorrow". This is a direct confirmation of ϕ and, in normal circumstance, we should expect that Bob's belief will become more entrenched. However, Bob recalls that, time and again, Tim has viciously lied to him about cancelled classes, thereby harming his academic status. One may consider it rational in this case for Bob to lower his degree of belief in ϕ, stop believing ϕ or, in extreme cases, to opt for believing $\neg\phi$.*

Intuitions about when and how trust in some information source *should* change is very context-sensitive and 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 5.2. *Let the belief base of agent $\mathcal{A}$ be $\{(S \rightarrow P, b_1), (Q \rightarrow \neg S, b_2)\}$. Information source $Jordan$, conveys P then conveys Q. Since $\mathcal{A}$ has no evidence against either, it believes both. Now, information 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 4.1, a formula ϕ may support another formula ψ by transitivity through an intermediate source σ. Given that, in general, the 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 exit. 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 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.

($\ltimes_1$: **Closure**) $\mathcal{K} \ltimes (\phi, \sigma)$ is an information state.

($\ltimes_2$: **Default Attitude**) If $(\sigma, T, t) \notin \mathcal{T}(\mathcal{K})$, then $(\sigma, T, \delta) \in \mathcal{T}(\mathcal{K} \ltimes (\phi, \sigma))$.

($\ltimes_3$: **Consistency**) $Cn(For(\mathcal{B}(\mathcal{K} \ltimes (\phi, \sigma)))) \neq \mathcal{L}$.

($\ltimes_4$: **Resilience**) If $Cn(\{\phi\}) = \mathcal{L}$, then $\mathcal{K} \not\prec_{\sigma,T} \mathcal{K} \ltimes (\phi, \sigma)$.

($\ltimes_5$: **Supported Entrenchment**) $\mathcal{K} \ltimes (\phi, \sigma) \prec_\phi \mathcal{K}$ only if $Cn(For(\mathcal{B}(\mathcal{K}))) = \mathcal{L}$.

($\ltimes_6$: **Opposed Entrenchment**) $\mathcal{K} \not\prec_{\neg\phi} \mathcal{K} \ltimes (\phi, \sigma)$.

($\ltimes_7$: **Positive Relevance**) If $\mathcal{K} \prec_{\sigma',T} \mathcal{K} \ltimes (\phi, \sigma)$ and $\phi \in For(\mathcal{B}(\mathcal{K} \ltimes (\phi, \sigma)))$, 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.

($\ltimes_8$: **Negative Relevance**) If $\mathcal{K} \ltimes (\phi, \sigma) \prec_{\sigma',T} \mathcal{K}$, then

1. $\phi \in For(\mathcal{B}(\mathcal{K} \ltimes (\phi, \sigma)))$ and σ' is $\neg\phi$-relevant; or
2. $\sigma' = \sigma$, but, there is $\Gamma \subseteq For(\mathcal{B}(\mathcal{K} \ltimes (\phi, \sigma)))$ where Γ is a $\neg\phi$-kernel.

($\ltimes_9$: **Belief Confirmation**) If $\mathcal{K} \prec_\psi \mathcal{K} \ltimes (\phi, \sigma)$, then, $\psi \neq \phi$ is supported by ϕ.

($\ltimes_{10}$: **Belief Refutation**) If $\mathcal{K} \ltimes (\phi, \sigma) \prec_\psi \mathcal{K}$, then

1. ψ is $\neg\phi$-relevant and $\phi \in Cn(For(\mathcal{B}(\mathcal{K} \ltimes (\phi, \sigma))))$ or $\mathcal{K} \ltimes (\phi, \sigma) \prec_{\neg\phi} \mathcal{K}$; or
2. ψ is ϕ-relevant and $\phi \notin Cn(For(\mathcal{B}(\mathcal{K} \ltimes (\phi, \sigma))))$ or $\mathcal{K} \ltimes (\phi, \sigma) \prec_\phi \mathcal{K}$.

Information revision should yield an information state ($\ltimes_1$). An information source that has no prior degree of trust is associated with the default degree of trust ($\ltimes_2$). A revised information state is consistent even if the revising formula is itself contradictory ($\ltimes_3$). If ϕ is inconsistent, σ should not become more trusted[7] ($\ltimes_4$). Abiding by the admissible and forbidden cases of information revision outlined in Tables 1 and 2, ϕ cannot become less entrenched unless the belief base is inconsistent ($\ltimes_5$) while, even if the belief base is inconsistent, $\neg\phi$ should not become more entrenched ($\ltimes_6$). If an information source 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 ($\ltimes_7$). 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 ($\ltimes_7$). ψ is more entrenched after revision only if it is supported by ϕ ($\ltimes_9$). Finally, 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 ($\ltimes_{10}$).

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

5.3 Discussion

The following observations follow from the definition of information states, the support graph, and the postulates.

Observation 5.1. *Let $\mathcal{K}$ be an information state.*

1. ***Positive Entrenchment**. If $Cn(For(\mathcal{B}(\mathcal{K}))) \neq \mathcal{L}$, then, $\mathcal{K} \ltimes (\phi, \sigma) \not\prec_\phi \mathcal{K}$.*
2. ***Positive Persistence**. If $Cn(For(\mathcal{B}(\mathcal{K}))) \neq \mathcal{L}$ and $\phi \in Cn(For(\mathcal{B}(\mathcal{K})))$, then $\phi \in Cn(For(\mathcal{B}(\mathcal{K} \ltimes (\phi, \sigma))))$.*
3. ***Negative Persistence**. If $\neg\phi \notin Cn(For(\mathcal{B}(\mathcal{K})))$, then $\neg\phi \notin Cn(For(\mathcal{B}(\mathcal{K} \ltimes (\phi, \sigma))))$.*
4. ***Formula Relevance**. If $\mathcal{K} \not\equiv_\psi \mathcal{K} \ltimes (\phi, \sigma)$, then ψ is ϕ- or $\neg\phi$-relevant.*
5. ***Trust Relevance**. If $\mathcal{K} \not\equiv_{\sigma',T} \mathcal{K} \ltimes (\phi, \sigma)$, then σ' is ϕ- or $\neg\phi$-relevant.*
6. ***No Trust Increase I**. If $\phi \notin For(\mathcal{B}(\mathcal{K} \ltimes (\phi, \sigma)))$, then there is no $\sigma' \in \mathcal{S}_\mathcal{K}$ such that $\mathcal{K} \prec_{\sigma',T} \mathcal{K} \ltimes (\phi, \sigma)$.*
7. ***Rational Revision.** If $Cn(For(\mathcal{B}(\mathcal{K}))) \neq \mathcal{L}$, then an operator that observes $\ltimes_5$ and $\ltimes_6$ allows for only cases in Table 1 to occur.*

The first two clauses of Observation 5.1 follow straight away from the definition of the postulates. On the other hand, the third and fourth clauses demonstrate how the postulates managed to reflect the intuitions behind information revision that lead us to propose the support graph. As previously discussed, information revision is considered with relevant change. Thus, we achieved our goal by ensuring that if belief in a formula (or trust in a source) is revised, this formula (or source) is relevant to the formula that triggered the revision (or possibly its negation). The fifth clause highlights the fact that if the formula of revision is rejected, no extra support is provided for anyone and hence no source will be more trusted. Finally, the last clause shows how the postulates managed to capture our intuitions about belief revision highlighted in Tables 1 and 2.

Observation 5.2. *Let $\mathcal{K}_0 = \{\{\}, \{\}, \{\}\}$ be an information state. For $i > 0$, let $\mathcal{K}_i$ refer to any state resulting from the revision of $\mathcal{K}_{i-1}$ using an operator $\ltimes$, with a non-forgetful conveyance inclusion filter, which observes the postulates in Section 5.2. The following hold.*

1. ***Single Source Revision.** If $\mathcal{S} = \{\sigma\}$, then, for any information state $\mathcal{K}_i$ where $i > 0$, the maximum degree in $\{ t \mid (\sigma, T, t) \in \mathcal{T}(\mathcal{K}) \}$ is δ.*
2. ***No Trust Increase II.** If for every $\sigma \in \mathcal{S}_{\mathcal{K}_i}$, there is no source σ' that is σ-relevant, then there is no $\sigma \in \mathcal{S}_{\mathcal{K}_i}$ such that $\mathcal{K}_{i-1} \prec_{\sigma,T} \mathcal{K}_i$.*
3. ***No Trust Increase III.** If for every information state $\mathcal{K}_j$, $0 < j \leq i$, and for every source $\sigma_j \in \mathcal{S}_{\mathcal{K}_j}$ there is no*

source σ'_j that is σ_j-relevant, then the maximum degree in $\{ t \mid (\sigma, T, t) \in \mathcal{T}(\mathcal{K}_i) \}$ is δ.

The first clause in Observation 5.2 represents the case where, in fact, trust does not matter. When there is a single source, the relevance relata is reduced to logical implication between formulas. As trust can only decrease, because there can be no confirmations, information revision becomes traditional non-prioritized belief revision. The second clause draws upon the same line of reasoning. If sources are completely independent of each other, only self support is present, no source will be more trusted because, intuitively, there is no reliable independent evidence present for any of them. Last but not least, the third clause further supports our claim that if sources are not relevant to each other, relevance reduces to logical implication and, in the absence of source based support, no source will be more trusted (will not exceed the default).

An $\ltimes$ operator that observes the postulates in Section 5.2, by design, fails to observe the following AGM postulates (Alchourrón, Gärdenfors, and Makinson 1985):

- **Success**. The success postulate states that, on revising with ϕ, agent $\mathcal{A}$ should believe ϕ. The $\ltimes$ operator fails to observe AGM-success because information revision depends not only on the formula of revision but also on the source of that formula. Thus, $\mathcal{A}$ does not just accept a new piece of information.
- **Vacuity**. AGM-vacuity states that expansion with ϕ, if the belief base does not derive $\neg\phi$, is a subset of revision with ϕ. In order to draw a comparison, we have to first define expansion of information states. Let $\mathcal{K} + (\phi, \sigma)$ denote the expansion of information state $\mathcal{K}$ with ϕ conveyed by σ. Expansion just adds a formula to the belief base without checking for consistency. Obviously, $\phi \in Cn(For(\mathcal{B}(\mathcal{K} +_{\mathcal{F}} (\phi, \sigma))))$. However, since $\ltimes$ does not observe success, it is not always the case that $\phi \in Cn(For(\mathcal{B}(\mathcal{K} \ltimes (\phi, \sigma))))$, and hence expansion is not a subset of revision even if $\neg\phi \notin Cn(For(\mathcal{B}(\mathcal{K})))$. Expansion on its own does not exist for information states as, we believe, adding a piece of information from an information source should be part of revision because: i) accepting or rejecting ϕ depends, among other things, on trust in σ, ii) an expanded belief base is not necessarily a consistent one.
- **Extensionality**. Extensionality says that if $\phi \Leftrightarrow \psi$, then revision with ϕ is equivalent to revision with ψ. There is no notion of an information source in the traditional AGM approach. Again, to draw a comparison, we will consider the case where revision is taking place with ϕ and ψ both conveyed by the same information source σ. Even then, $\ltimes$ fails to observe extensionality. Since trust in a source is associated with a topic, and since topics need not be closed, it is not always the case that ϕ conveyed by σ is believed, if believed at all, to the same degree of ψ that is also conveyed by σ. Hence, revision with (ϕ, σ), in general, is not the same as revision (ψ, σ) even if $\phi \Leftrightarrow \psi$.
- **Recovery**. In our framework of information revision, there is no operation of "contraction" on its own, it has to be a part of revision. Contraction is the process of removing a formula from the consequences of a belief base. AGM-Recovery states that expansion with ϕ after contraction with ϕ should yield the original belief set before contraction (ϕ already belongs to the belief set). Thus modeling recovery in information states is enforcing that removing a formula ϕ from $\mathcal{B}(\mathcal{K})$ and then expanding with (ϕ, σ) will result in the original belief base. As with the previous cases, $\ltimes$ fails to observe recovery because if the contraction of ϕ occurred, the reintroduction of ϕ will affect the resulting degrees of belief and trust differently depending on the source of ϕ.

6 Extended Example

Let information structure $\mathcal{I} = (\mathcal{L}_{\mathcal{V}}, \mathcal{C}, \mathcal{S}, \mathcal{G})$, where

- Language $\mathcal{L}_{\mathcal{V}}$ is a propositional language with the set $\mathcal{V} = \{Arr, Inc, Doomed, Kwin, Jwin, Afather, Lymarrid, Lymother\}$ of propositional variables. The intuitive meaning of the variables is as follows. Arr means "The army of Dany will arrive". Inc means "Jon's army increased in size". $Doomed$ means "We are all doomed". $Kwin$ denotes "The Knight King wins", while $Jwin$ denotes "Jon wins". $Afather$ means that "Agon is the father of Jon". $Lymarried$ denotes that "Agon married Lyanna", and finally, $Lymother$ represents "Lyanna is the mother of Jon".
- $\mathcal{C} = \{\{\mathcal{L}_{\mathcal{V}}\}\}$.
- $\mathcal{S} = \{Tyrion, Sam, Peter, Varys, Jon\}$.
- $\mathcal{G} = (\mathbb{N}, \mathbb{N}, \prec_{\mathbb{N}}, \prec_{\mathbb{N}}, 1)$ where $\prec_{\mathbb{N}}$ is the natural order on natural numbers.

Then, we define information state $\mathcal{K}_0 = (\mathcal{B}_0, \mathcal{T}_0, \mathcal{H}_0)$ as follows:

- $\mathcal{B}_0 = \{(Arr \rightarrow Inc, 20), (Inc \rightarrow Jwin, 20),$ $(\neg Jwin \rightarrow Kwin, 20), (Kwin \rightarrow Doomed, 20),$ $(Afather, 10), (Lymarried, 6),$ $(Afather \wedge Lymarried \rightarrow Lymother, 10)\}$.
- $\mathcal{T}_0 = \{(Tyrion, 5), (Sam, 5), (Varys, 4), (Peter, 3),$ $(Jon, 10)\}$. The topic attribute was dropped from the tuples because there is only a single topic.
- $\mathcal{H}_0 = \{\}$

That is, we start revising with a consistent, non-empty, belief base, an empty history, and with an attribution of trust for all information sources.

Let $\ltimes_{\mathfrak{G}}$ be an information revision operator that will be used in this example.[8] To illustrate how it works, we need to define what we call the support degree. The support degree of a formula ϕ with respect to a source σ, is the number of believed σ-independent ϕ-kernels (other than $\{\phi\}$) and the number of sources (other than σ) that conveyed ϕ directly. Moreover, the support degree of a source σ is the sum of support degrees of all formulas it conveyed with respect to σ. The intuition is as follows. A source is supported to

[8] $\ltimes_{\mathfrak{G}}$ is just an operator created for the purpose of demonstrating interesting cases and is not a generic operator of information revision.

the extent formulas conveyed by this source are supported. However, we took into account source-independent kernels to eliminate exclusive self-support.

Given an information state $\mathcal{K}$, with a support graph $\mathfrak{G}(\mathcal{K})$, on revising with (ϕ, σ), $\ltimes_{\mathfrak{G}}$ operates as follows.

1. If ϕ is inconsistent, it will be rejected.
2. Otherwise, a degree of belief for ϕ is derived. For any formula, in this case ϕ, the degree of belief $b_\phi = Max(F, S)$. F represents the degree by which an agent believes in ϕ given all ϕ-kernels, while S represents how much an agent believes in ϕ given trust in sources that conveyed ϕ. Since a kernel is as strong as its weakest formula, let the set Γ_ϕ be the set containing, for every ϕ-kernel, the formula with the lowest degree. Then, F will be the degree of the formula with the maximum degree in Γ_ψ. Intuitively, the derived degree of belief in ϕ, given formulas, is that of its strongest support. Similarly, S will be the degree of the most trusted source among those that previously conveyed ϕ including σ.
3. Add (ϕ, b_ϕ) to the belief base. If any contradiction arises, for example $\neg\phi$ and ϕ (or any two formulas ψ and $\neg\psi$ that both belong to $Cn(For(\mathcal{B}(\mathcal{K}))))$, the derived degree of belief in $\neg\phi$ is compared to that of ϕ and the one with the lower degree is contracted. To contract any formula (for example ξ), $\ltimes_{\mathfrak{G}}$ removes, recursively, from every single ξ-kernel the formula with the lowest degree.
4. Once the beliefs are consistent, the support graph will be reconstructed. Given the new graph, trust in every ϕ- or $\neg\phi$-relevant source σ' is, possibly, revised. If $(\sigma', t_1) \in \mathcal{T}(\mathcal{K})$ and $(\sigma', t_2) \in \mathcal{T}(\mathcal{K}\ltimes_{\mathfrak{G}})$, while the support degree of σ' in $\mathcal{K}$ is d_1 and the support degree of σ' in $\mathcal{K}\ltimes_{\mathfrak{G}}$ is d_2, then the new degree of trust $t_2 = t_1 + (d_2 - d_1)$. The proposed trust update formula ensures that for any source, if the support degree increases, trust increases, and if the support degree decreases trust decreases.
5. Finally, given the new trust degrees derived in the previous step, for every formula ψ that is ϕ- or $\neg\phi$-relevant, a possible new degree of belief is derived in the same way b_ϕ was derived in step 2.

Observation 6.1. $\ltimes_{\mathfrak{G}}$ *observes* $\ltimes_{3-10}$ *of Section 5.2*

We now follow the changes to the information state of agent $\mathcal{A}$ as it observes the following conveyance instances. Every information state $\mathcal{K}_i$ is the result of revision of $\mathcal{K}_{i-1}$ using $\ltimes_{\mathfrak{G}}$, starting from $\mathcal{K}_0$. The conveyance inclusion filter is non-forgetful, hence every instance of conveyance will make it to the history.

First Instance: $Peter$ conveys Arr. Since $\mathcal{A}$ has no evidence against Arr, $\mathcal{A}$ believes Arr. As there is no evidence for Arr, its degree of belief will be equal to that of the trust in its source $Peter$. There was no confirmations nor refutations to any formulas, so the trust base remains unchanged. $\mathcal{K}_1$ is as follows: $\mathcal{B}(\mathcal{K}_1) = \mathcal{B}(\mathcal{K}_0) \cup \{(Arr, 3)\}$, $\mathcal{T}(\mathcal{K}_1) = \mathcal{T}(\mathcal{K}_0)$, and $\mathcal{H}(\mathcal{K}_1) = \{(Arr, Peter)\}$.

Second Instance: $Tyrion$ conveys $Doomed$. Similar to the previous case, $\mathcal{A}$ believes $Doomed$. $\mathcal{K}_2$ is as follows: $\mathcal{B}(\mathcal{K}_2) = \mathcal{B}(\mathcal{K}_1) \cup \{(Doomed, 5)\}$, $\mathcal{T}(\mathcal{K}_2) = \mathcal{T}(\mathcal{K}_1)$, and $\mathcal{H}(\mathcal{K}_2) = \mathcal{H}(\mathcal{K}_1) \cup \{(Doomed, Tyrion)\}$.

Third Instance: Sam conveys $Kwin$. $Kwin$ confirms $Doomed$. Since $Doomed$ had no kernels and was conveyed only by $Tyrion$, the support degree of $Tyrion$ in $\mathcal{K}_2$ (d_1) was 0. This is why $Tyrion$ was not more trusted in $\mathcal{K}_2$. However, after Sam conveyed $Kwin$, there is a $Tyrion$-independent $Doomed$-kernel $\{Kwin, Kwin \rightarrow Doomed\}$. Thus, the support degree of $Doomed$ in $\mathcal{K}_3$ with respect to $Tyrion$ is 1, which makes $Tyrion$'s new support degree (d_2) 1 as well. $Tyrion$ will be more trusted by a value equal to $d_2 - d_1 = 1$. Because $Tyrion$ is more trusted, $Tyrion$-relevant formulas could be more entrenched. In this particular case, $Doomed$ will be more entrenched. Hence, $\mathcal{K}_3$ is as follows: $\mathcal{B}(\mathcal{K}_3) = \mathcal{B}(\mathcal{K}_1) \cup \{(Doomed, 6), (Kwin, 5)\}$, $\mathcal{T}(\mathcal{K}_3) = (\mathcal{T}(\mathcal{K}_1) \setminus \{(Tyrion, 5)\}) \cup \{(Tyrion, 6)\}$, and $\mathcal{H}(\mathcal{K}_3) = \mathcal{H}(\mathcal{K}_2) \cup \{(Kwin, Sam)\}$.

Fourth Instance: $Peter$ conveys Inc. The newly conveyed Inc is supported by a kernel $\{Arr, Arr \rightarrow Inc\}$. However, since the only source supporting both Arr and Inc is $Peter$ himself, $Peter$'s support degree will not increase, but Inc will be believed as $\mathcal{A}$ has no evidence against it. $\mathcal{K}_4$ is as follows: $\mathcal{B}(\mathcal{K}_4) = \mathcal{B}(\mathcal{K}_3) \cup \{(Inc, 3)\}$, $\mathcal{T}(\mathcal{K}_4) = \mathcal{T}(\mathcal{K}_3)$, and $\mathcal{H}(\mathcal{K}_4) = \mathcal{H}(\mathcal{K}_3) \cup \{(Inc, Peter)\}$.

Fifth Instance: $Varys$ conveys $Jwin$. $\mathcal{A}$ has no evidence against $Jwin$ thus $\mathcal{A}$ believes it. Moreover, both Arr and Inc are believed propositions supporting $Jwin$. Now, there are 2 $Varys$-independent $Jwin$-kernels. Namely: $\{Inc, Inc \rightarrow Jwin\}$ and $\{Arr, Arr \rightarrow Inc, Inc \rightarrow Jwin\}$. Hence, the support degree of $Varys$ becomes 2. Thus, $Varys$ will be more trusted with a value of 2. As with the third instance, since $Varys$ is more trusted, the formulas that $Varys$ supports could be more entrenched and hence $\mathcal{K}_5$ is as follows: $\mathcal{B}(\mathcal{K}_5) = \mathcal{B}(\mathcal{K}_4) \cup \{(Jwin, 6)\}$, $\mathcal{T}(\mathcal{K}_5) = (\mathcal{T}(\mathcal{K}_4) \setminus \{(Varys, 4)\}) \cup \{(Varys, 6)\}$, and $\mathcal{H}(\mathcal{K}_5) = \mathcal{H}(\mathcal{K}_4) \cup \{(Jwin, Varys)\}$.

Sixth Instance: $Varys$ conveys $Lymother$. $\mathcal{A}$ already has evidence for $Lymother$. $Lymother$ has a single kernel $\{Afather, Lymarried, Afather \wedge Lymarried \rightarrow Lymother\}$. Since this kernel is not dependent on $Varys$, $Varys$'s support degree increases by 1 (due to $Lymother$) resulting in $Varys$ being more trusted. The weakest formula in the $Lymother$-kernel has a degree of 6. However, trust in $Varys$, a source who directly conveyed $Lymother$, is 7 and hence $Lymother$ will have a degree of belief equal to 7. As sources become more trusted, belief in formulas conveyed by these sources could increase. Hence, belief in $Jwin$ will increase and $\mathcal{K}_6$ is as follows: $\mathcal{B}(\mathcal{K}_6) = \mathcal{B}(\mathcal{K}_4) \cup \{(Jwin, 7), (Lymother, 7)\}$, $(\mathcal{T}(\mathcal{K}_5) \setminus \{(Varys, 6)\}) \cup \{(Varys, 7)\}$, and $\mathcal{H}(\mathcal{K}_6) = \mathcal{H}(\mathcal{K}_5) \cup \{(Lymother, Varys)\}$.

Seventh Instance: Jon himself after the battle conveys $\neg Jwin$. Here, $\neg Jwin$ supports $Kwin$ and $Doomed$. However, it is a direct refutation to $Jwin$ and provides evidence against Inc and Arr. This is the first time $\mathcal{A}$ has evidence against the newly conveyed formula. However, Jon has the highest degree of trust and hence $\neg Jwin$ will have a higher degree of belief than $Jwin$. Thus, $\mathcal{A}$ will choose to remove $Jwin$ as follows.

$Jwin$ has three kernels: $\Gamma_1 = \{Jwin\}$, $\Gamma_2 = \{Inc \rightarrow Jwin, Inc\}$ and $\Gamma_3 = \{Arr, Arr \rightarrow Inc, Inc \rightarrow Jwin\}$. The operator will remove the formula with the lowest degree from every kernel. Γ_1 has a single formula so it is removed and hence $\mathcal{A}$ gives up $Jwin$. Moreover, in Γ_2, Inc has a lower degree than $Inc \rightarrow Jwin$ thus $\mathcal{A}$ will give up Inc. Finally, following the same line of reasoning, Arr will be removed from Γ_3.

The support degree of $Varys$ in $\mathcal{K}_7$ is 1 as opposed to 3 in $\mathcal{K}_6$. Hence, $Varys$'s support degree decreased by 2 and, subsequently, $Varys$ becomes less trusted. Although $Peter$ is $Jwin$-relevant, according to the definition of $\ltimes_{\mathfrak{S}}$, in this particular case, trust in $Peter$ will not decrease. Both $Tyrion$ and Sam received a new confirmation and their support degrees increased by 1 resulting in them being more trusted which lead formulas supported by them to become more entrenched. Jon is not supported by any sources or formulas so trust in Jon will remain unchanged. $\mathcal{K}_7$ is as follows: $\mathcal{B}(\mathcal{K}_7) = \mathcal{B}(\mathcal{K}_0) \cup \{(Doomed, 7), (Kwin, 6), (Lymother, 6), (\neg Jwin, 10)\}$, $\mathcal{T}(\mathcal{K}_7) = \{(Tyrion, 7), (Sam, 6), (Varys, 4), (Peter, 3), (Jon, 10)\}$, and $\mathcal{H}(\mathcal{K}_7) = \mathcal{H}(\mathcal{K}_6) \cup \{(\neg Jwin, Jon)\}$.

The case of $Lymother$ is a very interesting case. $Lymother$ became less entrenched after revision with $\neg Jwin$. In traditional AGM-approaches $Lymother$ would not have been considered relevant to $\neg Jwin$ and hence it would not change according to the principle of minimality. However, as we previously argued, belief in a formula depends on trust in sources of said formula. Thus, when trust in $Varys$ decreased, irrelevant from $Lymother$, formulas conveyed by $Varys$ (including $Lymother$) were subject to revision.

7 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. Moreover, we introduced the support graph which is a novel formal structure that highlights the 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, giving rise to an explainable AI architecture.

References

Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *The journal of symbolic logic* 50(2):510–530.

Booth, R., and Hunter, A. 2018. Trust as a precursor to belief revision. *Journal of Artificial Intelligence Research* 61:699–722.

Castelfranchi, C., and Falcone, R. 1998. Principles of trust for MAS: Cognitive anatomy, social importance, and quantification. In *Proceedings International Conference on Multi Agent Systems*, 72–79. IEEE.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial intelligence* 89(1-2):1–29.

Demolombe, R., and Liau, C.-J. 2001. A logic of graded trust and belief fusion. In *Proceedings of the 4th workshop on deception, fraud and trust in agent societies*, 13–25.

Demolombe, R. 2001. To trust information sources: a proposal for a modal logical framework. In Castelfranchi, C., and Tan, Y.-H., eds., *Trust and Deception in Virtual Societies*. Dordrecht: Springer Netherlands. 111–124.

Demolombe, R. 2004. Reasoning about trust: A formal logical framework. In Jensen, C.; Poslad, S.; and Dimitrakos, T., eds., *Trust Management*, 291–303. Berlin, Heidelberg: Springer Berlin Heidelberg.

Drawel, N.; Bentahar, J.; and Shakshuki, E. 2017. Reasoning about trust and time in a system of agents. *Procedia Computer Science* 109:632–639.

Elangovan, A.; Auer-Rizzi, W.; and Szabo, E. 2007. Why don't I trust you now? An attributional approach to erosion of trust. *Journal of Managerial Psychology* 22(1):4–24.

Falcone, R., and Castelfranchi, C. 2001. Social trust: A cognitive approach. In Castelfranchi, C., and Tan, Y.-H., eds., *Trust and Deception in Virtual Societies*. Dordrecht: Springer Netherlands. 55–90.

Gärdenfors, P., and Makinson, D. 1988. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the 2nd conference on Theoretical aspects of reasoning about knowledge*, 83–95. Morgan Kaufmann Publishers Inc.

Hansson, S. O. 1994. Kernel contraction. *The Journal of Symbolic Logic* 59(3):845–859.

Hansson, S. O. 1999a. A survey of non-prioritized belief revision. *Erkenntnis* 50(2-3):413–427.

Hansson, S. O. 1999b. *A textbook of belief dynamics - theory change and database updating*, volume 11 of *Applied logic series*. Kluwer.

Hardwig, J. 1991. The role of trust in knowledge. *The Journal of Philosophy* 88(12):693–708.

Haselhuhn, M. P.; Schweitzer, M. E.; and Wood, A. M. 2010. How implicit beliefs influence trust recovery. *Psychological Science* 21(5):645–648.

Herzig, A.; Lorini, E.; Hübner, J. F.; and Vercouter, L. 2010. A logic of trust and reputation. *Logic Journal of the IGPL* 18(1):214–244.

Holton, R. 1994. Deciding to trust, coming to believe. *Australasian Journal of Philosophy* 72(1):63–76.

Ismail, H., and Attia, P. 2017. Towards a logical analysis of misleading and trust erosion. In Gordon, A. S.; Miller, R.; and Turán, G., eds., *Proceedings of the Thirteenth International Symposium on Commonsense Reasoning, COMMONSENSE 2017, London, UK, November 6-8, 2017*, volume 2052 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Jones, A. J., and Firozabadi, B. S. 2001. On the characterisation of a trusting agent—aspects of a formal approach. In Castelfranchi, C., and Tan, Y.-H., eds., *Trust and Deception in Virtual Societies*. Dordrecht: Springer Netherlands. 157–168.

Jones, A. J. 2002. On the concept of trust. *Decision Support Systems* 33(3):225–232.

Jøsang, A.; Ivanovska, M.; and Muller, T. 2015. Trust revision for conflicting sources. In *The 18th International Conference on Information Fusion (Fusion 2015)*, 550–557. IEEE.

Katz, Y., and Golbeck, J. 2006. Social network-based trust in prioritized default logic. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI 2006)*, 1345–1350.

Leturc, C., and Bonnet, G. 2018. A normal modal logic for trust in the sincerity. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems*, 175–183. International Foundation for Autonomous Agents and Multiagent Systems.

Liau, C.-J. 2003. Belief, information acquisition, and trust in multi-agent systems—a modal logic formulation. *Artificial Intelligence* 149(1):31–60.

Lorini, E., and Demolombe, R. 2008. From binary trust to graded trust in information sources: A logical perspective. In *International Workshop on Trust in Agent Societies*, 205–225. Springer.

Lorini, E.; Jiang, G.; and Perrussel, L. 2014. Trust-based belief change. In T. Schaub, G. Friedrich, B. O., ed., *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI 2014)*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, 549–554. Amsterdam: IOS Press.

McLeod, C. 2015. Trust. In Zalta, E. N., ed., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, fall 2015 edition.

Rodenhäuser, L. B. 2014. *A matter of trust: Dynamic attitudes in epistemic logic*. Universiteit van Amsterdam [Host].

Sabater, J., and Sierra, C. 2005. Review on computational trust and reputation models. *Artificial Intelligence Review* 24(1):33–60.

Sakama, C.; Caminada, M.; and Herzig, A. 2010. A logical account of lying. In *European Workshop on Logics in Artificial Intelligence*, 286–299. Springer.

Sakama, C. 2015. A formal account of deception. In *2015 AAAI Fall Symposia, Arlington, Virginia, USA, November 12-14, 2015*, 34–41. AAAI Press.

Simpson, J. A. 2007. Psychological foundations of trust. *Current Directions in Psychological Science* 16(5):264–268.

Van Benthem, J. 2007. Dynamic logic for belief revision. *Journal of applied non-classical logics* 17(2):129–155.

van Ditmarsch, H. 2014. Dynamics of lying. *Synthese* 191(5):745–777.