Support + Belief = Decision Trust

Alessandro Aldini[1 [0000-0002-7250-5011]], Agata Ciabattoni[2[0000-0001-6947-8772]], Dominik Pichler[2[0009-0003-1790-2983]]✉, and Mirko Tagliaferri[1[0000-0003-3875-0512]]✉

[1] University of Urbino {alessandro.aldini,mirko.tagliaferri}@uniurb.it
[2] TU Wien {agata,dominik}@logic.at

Abstract. We present SBTrust, a logical framework designed to formalize decision trust. Our logic integrates a doxastic modality with a novel non-monotonic conditional operator that establishes a positive support relation between statements and is closely related to a known dyadic deontic modality. For SBTrust, we provide semantics, proof theory and complexity results, as well as motivating examples. Compared to existing approaches, our framework seamlessly accommodates the integration of multiple factors in the emergence of trust.

Keywords: Decision trust · non-monotonic conditional · beliefs · deontic logic

1 Introduction

Decision trust is defined as the willingness to depend on something (or somebody) with a feeling of relative security, although negative consequences are possible. This notion plays a central role in computer-mediated interactions. For instance, in e-commerce, when there is an abundance of vendors in a marketplace offering nearly identical products, customers use trust to decide whom to buy from [42]. Similarly, in the next generation Internet of Things, smart sensors, edge computing nodes, and cloud computing data centres rely on trust to share services such as data routing and analytics and to assess who to communicate with [19]. Trust is becoming increasingly important also in the interaction between AI-enhanced technologies and their users. In particular, a variety of design methodologies and policies are required to foster trust toward artificial decision-making processes [26]. In spite of their differences, in all scenarios, interactions are governed by trust evaluations that depend on various conditions, e.g., reputation scores, Quality of Service (QoS), transparency, perceived accuracy, ease of use, and the trustee's ability to behave as expected by the trustor.

Those facts drove the development of various models for assessing trust, see, e.g., [3, 12]. Yet, each existing model relies on specific conditions for the emergence of trust, thus failing in environments where multifaceted aspects play a role in establishing trust relations (see, e.g., [41] for trust in e-commerce, or [6], where it is shown that computer science students use multiple sources of information to determine their trust in AI systems). This calls for Trust models combining multiple aspects to evaluate the presence or lack of trust in the environment [43]. To address this need, we introduce SBTrust, a logic that allows reasoning about decision trust relying on varied enabling conditions.

In our logic, trusting a formula φ means that the trustor is willing to accept the formula as being true, although it might be false. This acceptance-based interpretation of trust is compatible with influential conceptual analyses of the notion of trust that show that trusting a proposition boils down to using the proposition as a premise in one's reasoning, even though the proposition might be false [20].

In SBTrust, Trust is a derived operator whose constituents are a support connective and a belief operator (hence the logic's name). Whenever it is both believed that a formula φ supports a formula ψ, and that φ is true, then ψ is φ-trusted.

The notion of support, establishing a form of positive influence between two statements, is modeled through a novel dyadic operator $\rightsquigarrow$, where $\varphi \rightsquigarrow \psi$ is read as: in the most likely φ-scenarios ψ holds. The operator $\rightsquigarrow$ yields a non-monotonic conditional sharing properties with the KLM logic $\mathbb{P}$ of preferential reasoning [29] although without validating cautious monotonicity. We characterize $\rightsquigarrow$ with semantics and proof theory; its axioms and rules turn out to axiomatize the flat (i.e., non-nested) fragment of Åqvist system $\mathbb{F}$ [7] – a foundational preference-based logic for normative reasoning. The notion of belief (what is considered to be true from a subjective standpoint) is modeled through a belief operator B, obtained through the normal modal logic $\mathbb{KD}4$. Hence, our Trust operator ($T_\varphi\psi$) is built using those ingredients - $B\varphi \wedge B(\varphi \rightsquigarrow \psi) \rightarrow T_\varphi\psi$.

In the following, we provide motivations for introducing yet a new logical framework for decision trust. The key ingredient is the support operator $\rightsquigarrow$, for which we discuss in Section 2 the (undesired and) required properties. For our logic we present syntax (Section 3), semantics (Section 3.2), and establish the connection between $\rightsquigarrow$ and Åqvist system $\mathbb{F}$. Soundness, completeness, and complexity (for the satisfiability and the model checking problem) for SBTrust are established in Section 4.

1.1 Decision trust: state of the art

Decision trust logical formalisms can be classified into the following paradigms [8]:

- **Policy-based models**: trust is obtained by implementing hard-security mechanisms based on cryptographic protocols and access control, e.g., [39]. Logical frameworks for policy-based mechanisms are defined in various papers, including [1].
- **Reputation-based models**: trust is obtained through indications of past interactions that are evaluated by gathering and manipulating performance scores for those interactions, see, e.g., [9]. Logical approaches in this setting include [2].
- **Cognitive models**: trust derives from the combination of various factors, including the agent's disposition, the importance/utility of a situation [33], and the agent's expectation and willingness [13]; several logics formalize such aspects [4, 17].

Although models within a given paradigm are employed in real-world applications (e.g., [28]), they often rely on partial features of trust or assume very specific conditions, making them limited. Policy-based models flatten trust on the use of (cryptographic) protocols and regulations that fail whenever they circularly rely on some trust conditions - *the problem of trusting the policy-makers* [27]. Reputation-based models flatten trust on scores that often represent only a proxy for trust - *the problem of the insufficiency of reputation for trust* [12]. Unlike other paradigms, cognitive models capture more nuanced notions of trust, but rely on representing agents' cognitive features, which may not always be possible and can lead to complexity issues. This creates a trade-off between effectively modeling various aspects of trust and the complexity of estimating all its constituting elements in real-world settings. The following example illustrates how the three paradigms model agent trust and highlights their shortcomings.

Example 1. [Amazon] As a leading online retailer, Amazon prioritizes consumer trust to drive transactions, enforcing protocols and vendor rules. Imagine a customer assessing whether to trust the proposition "Amazon vendor V_i is reliable" ($\boxed{GoodV_i}$). In a policy-based model of trust, the customer would trust $GoodV_i$ only if V_i meets Amazon's internal policies (e.g., V_i is a registered company). Yet, this approach has drawbacks: (*i*) requirements could be manipulated, giving the customer a false sense of security; (*ii*) trust extends beyond regulations, as customers' trust is not solely based on vendors complying with policies; (*iii*) evaluating trust would be shifted from whether a customer trusts V_i to whether it trusts Amazon (as a policy maker) and its policies - trust that would depend on other policies Amazon has to abide by, *ad infinitum* (the problem of establishing primitive trusted policy-makers). In a reputation-based model, trust in $GoodV_i$ depends solely on V_i's positive reviews. However, this has two limitations: (*i*) new vendors lack reviews, making it difficult to establish trust; (*ii*) reviews can be manipulated, (e.g., in 2017, *The Shed at Dulwich* restaurant became London's top restaurant on Tripadvisor, although serving fake food). A cognitive model of trust allows trust estimations based on agents' cognitive features (e.g., the intention of the vendor to provide a good service). However, the trust triggers are limited to the cognitive features representable within the model, restricting its application.

2 Support operator

We introduce the support operator, $\rightsquigarrow$, explaining our choice of a non-monotonic operator and its key properties with examples. Henceforth, we will shorten the reading of $\varphi \rightsquigarrow \psi$ to: given φ, then ψ is most likely. Throughout this paper, we restrict to non-nested applications of $\rightsquigarrow$ as it simplifies the technical aspects and adequately models the scenarios of interest without unnecessary complexity.

2.1 Why yet another notion of support

Various notions of support and axiomatizations as a conditional operator appear in the literature; see, e.g. [15] for three potential readings as an evidence operator. What all authors agree upon is that the operator should be non-monotonic, i.e., given $\varphi \rightsquigarrow \psi$, there is no reason why $\varphi \wedge \xi \rightsquigarrow \psi$ should be the case. This is because additional information (ξ) may undermine the previously established supporting statement. We also assume that support is a non-monotonic operator, but our rationale distinguish our view from existing ones. In particular, contraposition ($(\phi \rightsquigarrow \psi) \rightarrow (\neg\psi \rightsquigarrow \neg\phi)$), combined with right weakening (if $\phi \models \psi$, then $\chi \rightsquigarrow \phi \models \chi \rightsquigarrow \psi$), leads to monotonicity for the $\rightsquigarrow$ operator. Differently from [15], to avoid monotonicity, we give up contraposition (as motivated in Ex. 2) rather than right weakening, which is a reasonable assumption.

Example 2 (Contraposition and Modus Ponens). Assume that $GoodV_i$ supports that V_i's products are delivered fast ($\boxed{FastV_i}$), i.e., $GoodV_i \rightsquigarrow FastV_i$. This should not imply that if the delivery is slow, then it is most likely that the vendor is not a good one ($\neg FastV_i \rightsquigarrow \neg GoodV_i$), as the delay may depend on other reasons. For analogous reasons, we do not have that $GoodV_i$ and $GoodV_i \rightsquigarrow FastV_i$ imply $FastV_i$.

To proceed methodically, we draw upon the axiomatizations of non-monotonic conditionals from [29], known as KLM systems, as they serve as cornerstones for non-monotonic reasoning. We begin with an example demonstrating why the principle

$$\textbf{CM} \quad (\varphi \rightsquigarrow \psi) \wedge (\varphi \rightsquigarrow \chi) \rightarrow ((\varphi \wedge \psi) \rightsquigarrow \chi)$$

of cautious monotonicity is unfit for formalizing our notion of support (see also Rem. 5).

Example 3. Let $\boxed{DefCV_i}$ denote that customer C receives a defective item from vendor V_i, $\boxed{CompCV_i}$ that C submits a formal complaint regarding V_i, and $\boxed{RefCV_i}$ that C is refunded by V_i. It is most likely that C will complain after receiving a defective item, $DefCV_i \rightsquigarrow CompCV_i$. In the Amazon marketplace, we also have that $DefCV_i \rightsquigarrow RefCV_i$. However, having $DefCV_i \wedge RefCV_i$ does NOT mean that $CompCV_i$ is most likely, and this invalidates **CM**.

Remark 1. **CM** is tightly connected to transitivity of the underlying semantic models. In contrast to KLM logics, we do not assume transitive models, as discussed in Section 3.2.

2.2 Intuitive properties

We introduce the properties that we envision for the concept of support, illustrating their rationale by refining the scenario in Ex. 1. As will be shown in Theorem 1, many of the properties discussed below are inter-derivable, leading to a more concise axiomatization for $\rightsquigarrow$. Henceforth, by axioms, we mean axiom schemata. The naming conventions for the considered properties are taken from the KLM systems [29] and $\mathbb{F}$ [7, 34].

As the support operator $\rightsquigarrow$ applies to boolean formulas we expect all classical tautologies to be provable. Moreover, since any fact intuitively supports itself, the axiomatization of $\rightsquigarrow$ should be able to derive the following axiom

$$\textbf{ID} : \varphi \rightsquigarrow \varphi$$

The presence of this axiom highlights that $\rightsquigarrow$ does not establish a causal relation, see Remark 2. Moreover, we want our support system to not support contradictions. In essence, anything supporting a contradiction must be dismissed:

$$\textbf{ST} : (\varphi \rightsquigarrow \bot) \rightarrow \neg\varphi$$

Example 4. Let $\boxed{CompliantV_i}$ mean that vendor V_i is compliant with *"Amazon Seller Terms and Conditions"* and let V_i be a vendor with an average rating of 4.5 stars for her main product j ($\boxed{TopRatingV_{i,j}}$). Assume that $CompliantV_i$ and $TopRatingV_{i,j}$ support that V_i is a good vendor, i.e., $(CompliantV_i \wedge TopRatingV_{i,j}) \rightsquigarrow GoodV_i$. This implies that being compliant supports the connection between having good reviews and being a good vendor, as expected by Amazon and their customers, i.e., $CompliantV_i \rightsquigarrow (TopRatingV_{i,j} \rightarrow GoodV_i)$, which leads to the following axiom.

$$\textbf{SH} : ((\varphi \wedge \psi) \rightsquigarrow \chi) \rightarrow (\varphi \rightsquigarrow (\psi \rightarrow \chi))$$

This axiom expresses that deductions performed under strong assumptions may be useful even if the assumptions are not known facts.

It is quite natural to assume that if a statement supports two other statements, it supports their conjunction, as expressed by the axiom

$$\textbf{AND} : (\varphi \rightsquigarrow \psi) \wedge (\varphi \rightsquigarrow \chi) \rightarrow (\varphi \rightsquigarrow (\psi \wedge \chi))$$

Due to **ST**, this axiom can never be used to derive that a non-contradictory statement supports a contradiction. We illustrate this with Example 5, involving the lottery paradox [30], a stumbling block for default reasoning systems (see [35]).

Example 5. The paradox states that in a fair lottery, it is rational to assume that each individual ticket is likely not to win. By allowing to infer from two statements being likely that their conjunction is also likely, one concludes that two tickets are likely not to win. By iterating this reasoning, we can infer that it is likely that no ticket will win, which contradicts the fact that a winning ticket exists. The paradox does not apply to $\rightsquigarrow$. Indeed, if we assume that every ticket is most likely not to win, $\top \rightsquigarrow \neg T_i$, we can infer $\top \rightsquigarrow \bigwedge \neg T_i$ by **AND**. $\bigwedge \neg T_i$ being a contradiction, using **ST** we could derive $\neg\top$, which is impossible. This implies that our original assumption was wrong. Therefore, it cannot be assumed that $\top \rightsquigarrow \neg T_i$ for every ticket, thus stopping the paradox.

A support operator should also satisfy the **CUT** axiom, as illustrated by Example 6:

$$\textbf{CUT} : (\varphi \rightsquigarrow \psi) \wedge ((\varphi \wedge \psi) \rightsquigarrow \chi) \rightarrow (\varphi \rightsquigarrow \chi)$$

Example 6. Let $\boxed{AuthV_i}$ stand for V_i is authenticated on the Amazon marketplace. Obviously, an authenticated and compliant vendor is most likely to be a legitimate business ($\boxed{LegitV_i}$), i.e., $(AuthV_i \wedge CompliantV_i) \rightsquigarrow LegitV_i$. Moreover, due to Amazon's policies, $AuthV_i \rightsquigarrow CompliantV_i$. This implies that given $AuthV_i$ it is already most likely that V_i is legitimate, $AuthV_i \rightsquigarrow LegitV_i$.

Example 7. Let $\boxed{FairV_i}$ mean that V_i abides by the *"Acting Fairly"* policy of the "Amazon's Code of Conduct". Assume we have both $CompliantV_i \rightsquigarrow GoodV_i$ and $FairV_i \rightsquigarrow GoodV_i$. These two facts imply that it should be sufficient to satisfy $CompliantV_i$ or $FairV_i$ to be considered a good vendor, i.e., $(CompliantV_i \vee FairV_i) \rightsquigarrow GoodV_i$. This example leads to the following axiom.

$$\textbf{OR} : (\varphi \rightsquigarrow \psi) \wedge (\chi \rightsquigarrow \psi) \rightarrow ((\varphi \vee \chi) \rightsquigarrow \psi)$$

Example 8. Let $\boxed{GoodQoSV_{i,j}}$ mean that vendor V_i offers high QoS for product j. Assume that V_i sells two distinct products a and b, using the same commercial infrastructure (logistics, customer care, and so on). Hence, it would be absurd that V_i offers high QoS only for one of the two products, i.e., $\neg(GoodQoSV_{i,a} \leftrightarrow GoodQoSV_{i,b}) \rightsquigarrow \bot$. Hence, whatever $GoodQoSV_{i,a}$ supports, it should also be supported by $GoodQoSV_{i,b}$, and vice versa. This example leads to the following axiom.

$$\textbf{LL+} : (\neg(\varphi \leftrightarrow \psi) \rightsquigarrow \bot) \rightarrow ((\varphi \rightsquigarrow \chi) \leftrightarrow (\psi \rightsquigarrow \chi))$$

Example 9. Let $\boxed{GoodPriceV_i}$ and $\boxed{AmazonChoice_{i,j}}$ mean that V_i uses competitive prices, and that the product j sold by V_i is labeled as "Amazon's Choice". By Amazon's policy, $(GoodPriceV_i \wedge FastV_i \wedge TopRatingV_{i,j}) \rightarrow AmazonChoice_{i,j}$. Assume $GoodV_i \rightsquigarrow GoodPriceV_i$ (a good vendor is most likely to price well), $GoodV_i \rightsquigarrow$

$FastV_i$ (a good vendor is most likely to deliver quickly), and $GoodV_i \rightsquigarrow TopRatingV_{i,j}$ (a good vendor is most likely to have good reviews). These supports together should imply that $GoodV_i$ supports $AmazonChoice_{i,j}$. This motivates the following rule.

$$\mathbf{RCK}: \frac{\varphi_1 \wedge \cdots \wedge \varphi_n \to \varphi_{n+1}}{(\psi \rightsquigarrow \varphi_1) \wedge \cdots \wedge (\psi \rightsquigarrow \varphi_n) \to (\psi \rightsquigarrow \varphi_{n+1})}$$

of which Right Weakening (**RW**, see Remark 2), represents the particular case $n = 1$.

Example 10. Let V_i be a vendor selling product j and assume that $GoodQoSV_{i,j} \rightsquigarrow FastV_i$ and $\neg(AmazonChoice_{i,j} \rightsquigarrow DefCV_i)$. If these conditions hold, then customer C will have a good purchasing experience. Given that this implication holds, under the same hypothesis, it follows that C having a negative purchasing experience supports a contradiction. This example leads to the following "$\mathbb{S}5$-like" rule.

$$\mathbf{S5_F}: \frac{[\neg](\varphi_1 \rightsquigarrow \psi_1) \wedge \cdots \wedge [\neg](\varphi_n \rightsquigarrow \psi_n) \to \chi}{[\neg](\varphi_1 \rightsquigarrow \psi_1) \wedge \cdots \wedge [\neg](\varphi_n \rightsquigarrow \psi_n) \to (\neg\chi \rightsquigarrow \bot)}$$

where $[\neg](\varphi_i \rightsquigarrow \psi_i)$ stands for either $(\varphi_i \rightsquigarrow \psi_i)$ or its negated version $\neg(\varphi_i \rightsquigarrow \psi_i)$ and the same formulas are negated in the premise as in the conclusion of the rule. The rule is named because, when considered alongside other axioms and rules, $\mathbf{S5_F}$ grants the operator $\rightsquigarrow$ all the properties of an $\mathbb{S}5$-modality for the shallow fragment (see Theorem 3). As shown in Section 4, $\mathbf{S5_F}$ lets the operator $\rightsquigarrow$ behave locally like an absolute operator, playing a crucial role in the completeness proof.

Remark 2. (Most of) The axioms and rules discussed above are present in well-known systems. For instance, the KLM logic $\mathbb{P}$ of preferential reasoning, which interprets the dyadic operator $\varphi \mid\!\sim \psi$ as "φ *typically implies* ψ", contains the rule **RW** (see below) and axioms **ID**, **CUT**, **AND** and **OR**. I/O logics [32] and their causal versions [11], whose dyadic operator is interpreted as a dyadic obligation and a causal relation, respectively, share **RW**, **CUT**, **AND** and **OR** (but not **ID**). Note that KLM and (deontic and causal) I/O logics also contain the rule **LLE** (weaker than axiom **LL**+) below:

$$\mathbf{RW}: \frac{\varphi_1 \to \varphi_2}{(\psi \rightsquigarrow \varphi_1) \to (\psi \rightsquigarrow \varphi_2)} \qquad \mathbf{LLE}: \frac{\varphi \leftrightarrow \psi}{(\varphi \rightsquigarrow \chi) \leftrightarrow (\psi \rightsquigarrow \chi)}$$

An important difference between these logics and our $\rightsquigarrow$ operator is the direct interaction of support formulas and propositional formulas due to the rule $\mathbf{S5_F}$.

3 A logical framework for decision trust

Our logic SBTrust combines the $\rightsquigarrow$ operator with a belief operator B. For the former, we use a (subset of) the discussed axioms and rules and for the latter a $\mathbb{KD}4$ modality.[3]

[3] Other normal modal logics could be used to formalize the epistemic/doxastic notion, e.g., a $\mathbb{KD}45$ modality (see [25] for a general discussion), without major effects on the logic.

3.1 Syntax and axiomatization

The language $\mathcal{L}$ of SBTrust consists of a countable set of propositional variables or atoms (ranging over $p, q, \dots$), the connectives $\wedge$ and $\neg$ of classical logic (the other propositional connectives are defined as usual), the binary support operator $\rightsquigarrow$, and the unary belief operator B. $\mathcal{L}$ is defined by the following two layers grammar:

$$\varphi := \bot \mid p \mid \varphi \wedge \varphi \mid \neg\varphi \qquad \alpha := \varphi \mid \varphi \rightsquigarrow \varphi \mid B\alpha \mid \alpha \wedge \alpha \mid \neg\alpha$$

We use $\varphi, \psi, \chi, \delta$, and π for formulas of classical logic CL, and α and β for general formulas in $\mathcal{L}$. $\mathcal{L}_T$ and $\mathcal{L}_{CL}$ will denote the set of formulas of SBTrust and of CL, respectively. We identify theoremhood in SBTrust with derivability in its Hilbert system.

Definition 1. *SBTrust is obtained by extending any axiom system for propositional classical logic: (indicating its axioms by) (***CL***) and the Modus Ponens rule* **MP***, together with the following axioms and rules.*

For the support operator *the axiom schemata are:*

(**ID**) $\varphi \rightsquigarrow \varphi$ (**SH**) $((\psi \wedge \chi) \rightsquigarrow \varphi) \to (\psi \rightsquigarrow (\chi \to \varphi))$
(**ST**) $(\varphi \rightsquigarrow \bot) \to \neg\varphi$ (**LL+**) $(\neg(\varphi \leftrightarrow \psi) \rightsquigarrow \bot) \to ((\varphi \rightsquigarrow \chi) \leftrightarrow (\psi \rightsquigarrow \chi))$

and the corresponding inference rules are **RCK** *and* $\mathbf{S5_F}$*, with the restriction that their applications yield formulas within the language* $\mathcal{L}$*.*

For the belief operator *the axiom schemata are:*

(**KB**) $B(\alpha \to \beta) \to (B\alpha \to B\beta)$ (**DB**) $B\alpha \to \neg B\neg\alpha$ (**4B**) $B\alpha \to BB\alpha$

*and Necessitation for B (***NB***) which derives $B(\varphi)$ whenever φ has been proven.*

Trust arises as a combination of support and belief, i.e., $T_\varphi\psi := B\varphi \wedge B(\varphi \rightsquigarrow \psi)$.

Definition 2. *A derivation of φ_n is a sequence $\varphi_1, ..., \varphi_n$ where each φ_i is either an axiom instance or follows from the previous ones by applying one of the rules. We write $\Gamma \vdash \varphi$ if there is a derivation of φ or if $\psi_1 \wedge ... \wedge \psi_m \to \varphi$ for some $\psi_1, ..., \psi_m \in \Gamma$.*

Remark 3. The newly defined trust operator inherits numerous properties from the support and belief axioms and rules. While a more comprehensive study of this operator on its own is planned for future work, here we present a specific property as an initial example. By definition, $T_\varphi\psi \to B\varphi$. Using axiom **DB** and rule **NB**, we derive $\neg B\bot$ which implies $\neg T_\bot\psi$, meaning it is impossible to trust something based on a contradiction. Now, assuming $T_\varphi\bot$, we derive $B\varphi$ by definition. However, from the axioms **ST** and **KB**, we can also derive $B\neg\varphi$, leading to a contradiction by axiom **DB**. Thus, it follows that $\neg T_\varphi\psi$ is derivable whenever either φ or ψ (or both) are contradictions. This property ensures that the trust operator does not permit trusting a contradiction or trusting something based on a contradiction—a desirable feature in this framework.

We prove that all the axioms and rules stated in Section 2.2 are derivable in SBTrust.

Theorem 1. *The rules* **RW** *and* **LLE***, as well as the axioms* **AND***,* **CUT***, and* **OR** *are derivable in the system for* $\rightsquigarrow$*.*

Proof. We will show the case for **LLE**. We assume that the formula $\varphi \leftrightarrow \psi$ is provable, then we can derive $\varphi \rightsquigarrow \chi \leftrightarrow \psi \rightsquigarrow \chi$ via the following:

(1) $\neg(\varphi \leftrightarrow \psi) \to \bot$ — *Hyp.* + (**CL**)
(2) $(\neg(\varphi \leftrightarrow \psi) \rightsquigarrow \neg(\varphi \leftrightarrow \psi)) \to (\neg(\varphi \leftrightarrow \psi) \rightsquigarrow \bot)$ — (**RCK**)
(3) $\neg(\varphi \leftrightarrow \psi) \rightsquigarrow \bot$ — (**ID** + **MP**)
(4) $(\varphi \rightsquigarrow \chi) \leftrightarrow (\psi \rightsquigarrow \chi)$ — (**LL+** + **MP**)

Remark 4. From **RW** and **LLE**, we can derive the replacement of equivalent formulas within the support operator in compliance with the syntactic restrictions of the language.

Remark 5. Another reason for rejecting **CM** is that, in conjunction with **CUT** and **RW**, it permits to derive **REC** - $((\varphi \rightsquigarrow \psi) \wedge (\psi \rightsquigarrow \varphi)) \to ((\varphi \rightsquigarrow \chi) \leftrightarrow (\psi \rightsquigarrow \chi))$, which is too strong for a support operator since two statements that support each other do not necessarily support the same statements.

A strong connection holds between $\rightsquigarrow$ and $\mathbb{F}$, the dyadic deontic logic introduced in [7] and axiomatized in [34] using **MP** and the necessitation rule for $\Box$, and the axioms

CL All truth-functional tautologies
T $\Box\varphi \to \varphi$
COK $\bigcirc(\psi \to \chi/\varphi) \to (\bigcirc(\psi/\varphi) \to \bigcirc(\chi/\varphi))$
Nec $\Box\varphi \to \bigcirc(\varphi/\psi)$
ID $\bigcirc(\varphi/\varphi)$
D* $\Diamond\psi \to (\bigcirc(\varphi/\psi) \to \neg\bigcirc(\neg\varphi/\psi))$
K$_\Box$ $\Box(\varphi \to \psi) \to \Box\varphi \to \Box\psi$
5 $\Diamond\varphi \to \Box\Diamond\varphi$
Abs $\bigcirc(\varphi/\psi) \to \Box\bigcirc(\varphi/\psi)$
Ext $\Box(\varphi \leftrightarrow \psi) \to (\bigcirc(\chi/\varphi) \leftrightarrow \bigcirc(\chi/\psi))$
SH $\bigcirc(\varphi/\psi \wedge \chi) \to \bigcirc(\chi \to \varphi/\psi)$

where $\bigcirc(\psi/\varphi)$ stands for "*ψ is obligatory under the condition φ*". The flat fragment of the language of $\mathbb{F}$ –where $\Box$ and $\bigcirc$ apply only to formulas of $\mathcal{L}_{CL}$– can be translated into our language as follows:

Definition 3. *Let χ be any formula in the flat fragment of $\mathbb{F}$. The translation χ^* is*

$$\varphi^* \mapsto \varphi \qquad (\Box\varphi)^* \mapsto \neg\varphi \rightsquigarrow \bot \qquad (\bigcirc(\psi/\varphi))^* \mapsto \varphi \rightsquigarrow \psi$$

Remark 6. The dual of $\Box$ gets translated to $(\Diamond\varphi)^* := \neg(\varphi \rightsquigarrow \bot)$.

The theorem below establishes a first link between $\mathbb{F}$ and $\rightsquigarrow$.

Theorem 2. *The translation * of all axioms and rules of $\mathbb{F}$ – but* **5** *and* **Abs** *– are derivable in the axiomatization for $\rightsquigarrow$.*

Proof. The claim for **T**, **Ext**, **ID**, and **SH** follows directly from the translation. The translation of axiom **D*** is $\neg(\varphi \rightsquigarrow \bot) \to \neg((\varphi \rightsquigarrow \psi) \wedge (\varphi \rightsquigarrow \neg\psi))$. Its contraposition $((\varphi \rightsquigarrow \psi) \wedge (\varphi \rightsquigarrow \neg\psi)) \to (\varphi \rightsquigarrow \bot)$ is an instance of **AND**. The case **K**$_\Box$ follows by:

(1) $((\neg(\varphi \to \psi) \rightsquigarrow \bot) \wedge (\neg\varphi \rightsquigarrow \bot)) \to \psi$ — (**ST** + **ST** + **CL**)
(2) $((\neg(\varphi \to \psi) \rightsquigarrow \bot) \wedge (\neg\varphi \rightsquigarrow \bot)) \to (\neg\psi \rightsquigarrow \bot)$ — (**S5**$_\mathbf{F}$)

The translations of the axioms **Nec** and **COK** can be derived in a similar way and are omitted.

Remark 7. The translation of axioms **Abs** and **5** from $\mathbb{F}$ results in formulas containing nested applications of $\rightsquigarrow$. In Section 3.2, we will see that the axioms and rules for $\rightsquigarrow$ axiomatize the shallow fragment of $\mathbb{F}$. In this regard, the rule $\mathbf{S5_F}$, which does not correspond to any rule known in the literature, does not follow from the remaining axioms and rules for $\rightsquigarrow$, and it is needed to derive (some) flat formulas holding in $\mathbb{F}$.

Example 11. Assume that a vendor V_1 believes the two formulas supported by $DefCV_i$ discussed in Example 3. If V_1 sends a defective item to customer C ($DefCV_1$), from $B(DefCV_1)$ and $B(DefCV_1 \rightsquigarrow CompCV_1)$ we derive $T_{DefCV_1}(CompCV_1)$. Similarly, it also holds that $T_{DefCV_1}(RefCV_1)$. Now, assume that C does indeed receive the refund, thus $B(DefCV_1 \wedge RefCV_1)$. We show that V cannot trust that a complaint will not be submitted, $\neg T_{DefCV_1 \wedge RefCV_1}(\neg CompCV_1)$. We use the following abbreviations to write a concise derivation: Let $d := DefCV_1$, $r := RefCV_1$, and $c := CompCV_1$, and, by hypothesis, $T_d(c) \wedge T_d(r)$, i.e., $B(d) \wedge B(d \rightsquigarrow c) \wedge B(d \rightsquigarrow r)$. Hence:

$$
\begin{array}{ll}
(1)\ (d \wedge (d \rightsquigarrow c)) \rightarrow \neg(d \rightsquigarrow \neg c) & (\mathbf{ST}+\mathbf{D}^*) \\
(2)\ (d \rightsquigarrow r) \wedge \neg(d \rightsquigarrow \neg c) \rightarrow \neg((d \wedge r) \rightsquigarrow \neg c) & (\mathbf{CUT}+\mathbf{CL}) \\
(3)\ d \wedge (d \rightsquigarrow c) \wedge (d \rightsquigarrow r) \rightarrow \neg((d \wedge r) \rightsquigarrow \neg c) & (1 \wedge 2)
\end{array}
$$

Then, applying rule **NB** to (3) and using the hypothesis together with axiom **KB**, we derive $B(\neg((d \wedge r) \rightsquigarrow \neg c))$. This formula, by axiom **DB**, finally implies $\neg T_{d \wedge r}(\neg c)$.

3.2 Semantics

For evaluating a formula of the form $\varphi \rightsquigarrow \psi$, we intuitively consider only the most likely φ-scenarios and check whether ψ holds in those scenarios. This approach is inspired by preference-based logics [22], in which a conditional statement "If φ then ψ" is interpreted as among the "best" possible scenarios in which φ is true, ψ is true as well. Notice that in this approach transitivity of the relation is not assumed, in contrast to the KLM systems. The semantics for SBTrust is built on preference-based models [40] (for the support statements) and relational models (for the belief operator). Also used in KLM logics and in Åqvist system $\mathbb{F}$, preference-based models are triples $\langle S, \geq, V\rangle$, where S denotes a set of states, V a valuation function, and the preference relation $\geq \subseteq S \times S$ compares pairwise the states in S. In our context, $w \geq v$ represents a pairwise comparison stating that w is at least as likely as v ($\mathbb{F}$ uses a *better than* interpretation, while the KLM logic $\mathbb{P}$ uses a *preferred to* interpretation).[4] Our "as likely as" interpretation has to be read as comparing two *specific* elements. Like in $\mathbb{F}$, the relations apply to pairs of worlds without imposing a global ordering, and therefore, they do not necessarily have to be transitive; for an illustration, see Example 12. Although close to $\mathbb{F}$, our approach differs in two important aspects. First, in our semantics, we include a component for belief formulas. Second, instead of having a unique preference frame, in a *Trust frame*, the set of states S is partitioned into multiple preference frames $\langle S_i, \geq_i\rangle$; this allows us to consider different support systems within the same model. Note that,

[4] A preference frame consists of a set of worlds and a binary relation, exactly like a Kripke frame. The difference lies in their different usages in the model.

without partitioning, the unique given support system would necessarily have to be believed. In our interpretation, a statement $\varphi \rightsquigarrow \psi$ is true in a state $w \in S_i$ if among all the φ-states in S_i ($||\varphi||_i$) the most likely φ-states in S_i ($most(||\varphi||_i)$) satisfy ψ; $most(||\varphi||_i)$ consists of the maximal elements in the set of all φ-states of S_i according to $\succeq_i$.

In line with systems $\mathbb{F}$ and $\mathbb{P}$, our semantics include a limitedness condition: $||\varphi||_i \neq \emptyset \Rightarrow most(||\varphi||_i) \neq \emptyset$. Limitedness allows us to express when a formula φ is impossible in a partition S_i using $\varphi \rightsquigarrow \bot$. Hence, limitedness restricts SBTrust to support only non-contradictory options (with the exception of $\bot \rightsquigarrow \bot$). Consequently, an agent will never place trust in a blatant contradiction, nor will she place φ-based trust in contradicting statements, invalidating $T_\varphi\psi \wedge T_\varphi\neg\psi$. The definitions in this section characterize the frames and models on which our logic is based.

Definition 4 (Trust frame). *Let* $\mathcal{F} := \langle S, (S_i)_{i\in I}, (\succeq_i)_{i\in I}, R\rangle$ *where*
- $\langle S, R\rangle$ *is a serial and transitive Kripke frame;*
- $(S_i)_{i\in I}$ *is a partition of* S*, i.e.,* $\bigcup_{i\in I} S_i = S$ *and* $\forall i, j \in I : S_i \cap S_j = \emptyset$*;*
- *For each* $i \in I$*:* $\langle S_i, (\succeq_i)\rangle$ *is a preference frame (therefore* $\succeq_i \subseteq S_i \times S_i$*).*

Definition 5 (Trust model, truth conditions). *Let* $\mathcal{M} := \langle S, (S_i)_{i\in I}, (\succeq_i)_{i\in I}, R, V\rangle$ *where*
- $\mathcal{F} := \langle S, (S_i)_{i\in I}, (\succeq_i)_{i\in I}, R\rangle$ *is a Trust frame;*
- $V : Prop \mapsto 2^S$ *is a Valuation function;*
- *For each* $i \in I$*,* $\langle S_i, (\succeq_i), V\rangle$ *fulfills the limitedness condition: for every propositional formula* φ*, it holds that* $||\varphi||_i \neq \emptyset \Rightarrow most(||\varphi||_i) \neq \emptyset$*;*
- $\mathcal{M}, s \models p$ *iff* $s \in V(p)$*;*
- $\mathcal{M}, s \models \neg\alpha$ *iff* $\mathcal{M}, s \not\models \alpha$*;*
- $\mathcal{M}, s \models \alpha \wedge \beta$ *iff* $\mathcal{M}, s \models \alpha$ *and* $\mathcal{M}, s \models \beta$*;*
- $\mathcal{M}, s \models \varphi \rightsquigarrow \psi$ *iff* $most(||\varphi||_i) \subseteq ||\psi||_i$ *for* $s \in S_i$*;*
- $\mathcal{M}, s \models B(\varphi)$ *iff* $\forall v : (sRv \rightarrow \mathcal{M}, v \models \varphi)$*;*

$||\varphi||_i := \{v \in S_i : \mathcal{M}, v \models \varphi\}$*;* $most(||\varphi||_i) := \{s \in ||\varphi||_i : \forall v[(v \in ||\varphi||_i \wedge v \succeq_i s) \rightarrow s \succeq_i v]\}$.

Semantic consequence ($\Phi \models \alpha$) and validity ($\models \alpha$) are defined as usual.

Remark 8. Two observations are in order. Firstly, we do not assume any property on the relations $\succeq_i$ to keep the model as general as possible, and in particular, we do not require the relations to be transitive. This is related to the idea that the pairwise ordering between the states subsumes considerations about different contexts of evaluation, i.e., the sets *most*(_), which are tied to the antecedents of support formulas. The use of non-transitive models is consistent with the discussion presented in [18] about the non-transitivity of descriptive theories of preferences (see also [10] and [38] about intransitive preferences within social choice theories). Note that assuming transitivity, would closely connect $most(||\varphi \wedge \psi||_i)$ and $most(||\varphi||_i)$, forcing cautious monotonicity of the framework, which is an unwanted property for our use cases (see Examples 3 and 12). Secondly, $\mathcal{M}, w \models \varphi \rightsquigarrow \psi$ holds if w is part of a set of states in which φ supports ψ is true. An agent may or may not believe $\varphi \rightsquigarrow \psi$, independently from the fact that $\varphi \rightsquigarrow \psi$ holds or not. This gives us the possibility to capture an agent's misinformation.

The example below, inspired by [37], illustrates the reason we do not assume transitivity in the preference relation.

Example 12. Let V_1, V_2, and V_3 be three vendors with competing offers for the same product. Consider a scenario where the offers of V_1 and V_2 are similar, apart from a key attribute (e.g., delivery time), for which V_2 presents a significantly better offer. Hence, it is more likely that V_2 is preferred to V_1. For analogous reasons but for a different attribute (e.g., quality of reviews), it is possible that V_3 is preferred to V_2. However, it could be that all the attributes of V_1 and V_3 are similar enough to be incomparable. This situation is represented in the following model.

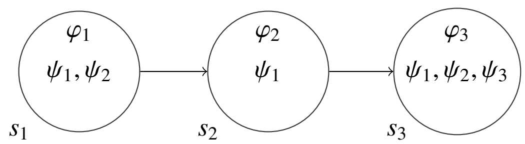

Notice that, for simplicity, we use a single-element partition $\{s_1, s_2, s_3\}$ and omit the accessibility relation R. Moreover, each formula φ_i represents that the vendor V_i is chosen (ignore, for the moment, the other formulas), and the solid arrows express the preference relations. Thus, for example, s_2 is more likely than s_1, s_3 is more likely than s_2, but it is not the case that s_3 is more likely than s_1 or vice versa. We use this example to emphasize the relation between the lack of transitivity and the rejection of **CM**. Let ψ_1 denote that V_1, V_2, and V_3 are Amazon vendors, ψ_2 denotes that Amazon offers a discount on products from long-standing vendors, and ψ_3 that Amazon offers a discount on products from new vendors. Assume that V_1 and V_3 are new vendors. Based on these assumptions, which do not affect the preference relations shown above, the following supports can be demonstrated easily: $\psi_1 \rightsquigarrow \psi_2$ and $\psi_1 \rightsquigarrow \psi_3$ (Amazon frequently supports discounts), but $(\psi_1 \wedge \psi_2) \not\rightsquigarrow \psi_3$ (it is not usual that Amazon promotes different discounts at the same time), which represent a violation of **CM**. Formally, given that $most(\|\psi_1\|) = \{s_3\}$, the support statements $\psi_1 \rightsquigarrow \psi_2$ and $\psi_1 \rightsquigarrow \psi_3$ are satisfied; however, $most(\|\psi_1 \wedge \psi_2\|) = \{s_1, s_3\}$ and s_1 violates the formula $(\psi_1 \wedge \psi_2) \rightsquigarrow \psi_3$, implying that the formula is not satisfied. Notice that if we imposed transitivity, thus relating s_1 to s_3, we would recover **CM** because in this case, $most(\|\psi_1 \wedge \psi_2\|)$ becomes $\{s_3\}$. Moreover, with **CM**, we would also get **REC** (see Remark 5), which, given that $\psi_1 \rightsquigarrow \psi_2$ and $\psi_2 \rightsquigarrow \psi_1$ hold, implies $(\psi_1 \rightsquigarrow \psi_3) \leftrightarrow (\psi_2 \rightsquigarrow \psi_3)$. Since $\psi_1 \rightsquigarrow \psi_3$ holds, we would also have $\psi_2 \rightsquigarrow \psi_3$, which is unreasonable since Amazon usually offers discounts on different types of vendors at different times.

We now examine the relation between $\mathbb{F}$ and $\rightsquigarrow$. Theorem 2 has already highlighted a syntactic connection (from $\mathbb{F}$ to $\rightsquigarrow$ via the translation * in Def. 3). Here, by using their semantics, we uncover a stronger tie. Recall that $\mathbb{F}$ is sound and complete w.r.t. all preference models $\langle S, \succeq, V\rangle$ which fulfill the limitedness condition, see [34]. We denote by $\models^{\mathbb{F}}$ the semantical consequence relation in $\mathbb{F}$.

Theorem 3. *For any set of formulas Γ and formula α in the language of $\mathbb{F}$ that do not contain nested modal operators, we have: $\Gamma \models^{\mathbb{F}} \alpha \Leftrightarrow \Gamma^* \models \alpha^*$.*

Proof. Both directions proceed by contraposition.

(⇒) We show that given a Trust model invalidating the semantical consequence for support we can find a preference model invalidating the semantical consequence for $\mathbb{F}$.

Assume that $\Gamma^* \not\models \alpha^*$. Hence, there exists a Trust model $\mathcal{M} = \langle S, (S_i)_{i\in I}, (\succeq_i)_{i\in I}, R, V\rangle$ and a state $s \in S_i$ such that $\forall\beta^* \in \Gamma^* : \mathcal{M}, s \models \beta^*$ and $\mathcal{M}, s \not\models \alpha^*$. We cut down the Trust model into a preference model as follows $\mathcal{M}' := \langle S_i, \succeq_i, V\rangle$. By definition, $\mathcal{M}'$ is a preference model fulfilling the limitedness condition.[5] Observe that no formula in $\Gamma^* \cup \{\alpha^*\}$ contains the operator B. Hence, the evaluation of the formulas in $\Gamma^* \cup \{\alpha^*\}$ at the state $s \in S_i$ coincides with the evaluation of the formulas in $\Gamma \cup \{\alpha\}$ in $\mathcal{M}$. We can therefore conclude $\forall\beta \in \Gamma : \mathcal{M}', s \models \beta$ and $\mathcal{M}', s \not\models \alpha$.

($\Leftarrow$) Given a preference model invalidating $\Gamma \models^{\mathbb{F}} \alpha$, we provide a Trust model invalidating $\Gamma^* \models \alpha^*$. Assume to have the preference model $\mathcal{M} = \langle S, \succeq, V\rangle$, fulfilling the limitedness condition and a state $s \in S$ such that $\forall\gamma \in \Gamma : \mathcal{M}, s \models \gamma$ and $\mathcal{M}, s \not\models \alpha$. We extend it into a Trust model with only one element in the partition, as follows $\mathcal{M}' := \langle S, (S_i)_{i\in I}, (\succeq_i)_{i\in I}, V, R\rangle$, with $I := \{1\}$, $S_1 := S$, $\succeq_1 := \succeq$ and $R := S \times S$. By definition, $\mathcal{M}'$ is a Trust model. Again no formula in $\Gamma^* \cup \{\alpha^*\}$ contains the operator B. Hence the evaluation of the formulas in $\Gamma^* \cup \{\alpha^*\}$ at the state $s \in S$ coincides with the evaluation of the formulas in $\Gamma \cup \{\alpha\}$ in $\mathcal{M}$. Hence $\forall\beta^* \in \Gamma^* : \mathcal{M}', s \models \beta^*$ and $\mathcal{M}', s \not\models \alpha^*$.

The Soundness and Completeness of SBTrust w.r.t. Trust models, proved in the next section, implies that the axioms and rules of $\rightsquigarrow$ axiomatize the flat fragment of $\mathbb{F}$.

4 Soundness, completeness and computational complexity

We start with the soundness of SBTrust w.r.t. Trust models.

Theorem 4 (Strong Soundness). *$\Phi \vdash \alpha \Rightarrow \Phi \models \alpha$ for any $\Phi \subseteq \mathcal{L}_T$ and $\alpha \in \mathcal{L}_T$.*

Proof. Proceed by induction on the derivation length, distinguishing cases according to the last rule applied. We show below the details for axiom **LL+** and rule $\mathbf{S5_F}$.

LL+: Given a Trust model $\mathcal{M} = \langle S, (S_i)_{i\in I}, (\succeq_i)_{i\in I}, R, V\rangle$ and a state $s \in S_i$ such that $\mathcal{M}, s \models \neg(\varphi\leftrightarrow\psi) \rightsquigarrow \bot$, then we get $most(\|\neg(\varphi\leftrightarrow\psi)\|_i) \subseteq \emptyset$. Given the limitedness assumption, this is equivalent to $\|\neg(\varphi\leftrightarrow\psi)\|_i = \emptyset$ and furthermore to $\|\varphi\leftrightarrow\psi\|_i = S_i$. Hence, φ and ψ are equivalent in every state of S_i. Therefore the sets $most(\|\varphi\|_i)$ and $most(\|\psi\|_i)$ coincide, i.e., $\mathcal{M}, s \models (\varphi \rightsquigarrow \chi)\leftrightarrow(\psi \rightsquigarrow \chi)$. $\mathbf{S5_F}$: Given a Trust model $\mathcal{M} = \langle S, (S_i)_{i\in I}, (\succeq_i)_{i\in I}, R, V\rangle$, we assume $((\neg)(\varphi_1 \rightsquigarrow \psi_1) \wedge \cdots \wedge (\neg)(\varphi_n \rightsquigarrow \psi_n)) \rightarrow \chi$ to be true in every state of $\mathcal{M}$. Given a state $s \in S_i$ such that $\mathcal{M}, s \models ((\neg)(\varphi_1 \rightsquigarrow \psi_1) \wedge \cdots \wedge (\neg)(\varphi_n \rightsquigarrow \psi_n))$ holds, it follows that $\forall w \in S_i\ \mathcal{M}, w \models ((\neg)(\varphi_1 \rightsquigarrow \psi_1) \wedge \cdots \wedge (\neg)(\varphi_n \rightsquigarrow \psi_n))$ because by virtue of the semantics of $\rightsquigarrow$, all the states of a given partition class satisfy the same (negated) support formulas. Therefore, we have that $\forall w \in S_i\ \mathcal{M}, w \models \chi$, which means $\|\neg\chi\|_i = \emptyset$ and finally $\mathcal{M}, s \models \neg\chi \rightsquigarrow \bot$.

Completeness is shown via the canonical model construction, adapted to our framework from [22]. The needed modifications are the following. First, we have to ensure that SBTrust allows us to derive all the axioms and rules required for the construction to proceed. Furthermore, unlike the models considered in [22], Trust models incorporate multiple preference frames, the belief operator B, and include the limitedness condition.

[5] In [34], limitedness is stated for every formula of $\mathbb{F}$. Since the truth sets of obligations and modalities are those of $\top$ or $\bot$, our limitedness condition is equivalent to the one given in $\mathbb{F}$.

The modifications are implemented as follows. The required axioms and rules for their proof to go through are those of $\mathbb{F}$ without $\mathbf{D}^*$ (which corresponds to limitedness). In Section 3.1 we have shown that with the exception of **5** and **Abs** the axioms of $\mathbb{F}$ are derivable in SBTrust. We prove that we can derive all the necessary properties of the canonical model even in the absence of axioms **5** and **Abs**, by relying on other rules of SBTrust, primarily on $\mathbf{S5_F}$. The multiple preference frames are handled by partitioning the maximal consistent sets used in the canonical model construction into equivalence classes containing the same support formulas. The addition of belief is easy: We equip the canonical model with the accessibility relation in the usual Kripke fashion. Incorporating the limitedness condition poses the challenge of guaranteeing that, in every preference model of our canonical model, each non-empty set $\|\varphi\|_i$ contains a maximal element according to the preference relation. We address this by using the axiom **ST**.

Definition 6. *A set $\Gamma \subseteq \mathcal{L}_T$ is called a maximal consistent set (MCS for short) if (a) $\Gamma \nvdash \bot$, and (b) For every $\alpha \in \mathcal{L}_T$ either $\alpha \in \Gamma$ or $\neg\alpha \in \Gamma$.*

Although not all the states in a model validate the same support formulas, we still need to make sure that all the states inside the same preference frame S_i do. For that reason, we partition the maximal consistent sets into equivalence classes, containing the same support formulas. We also define a set $\rightsquigarrow_\varphi(\Gamma)$ containing all formulas which are supported in a MCS Γ by a formula φ.

Definition 7. *Given $\Gamma, \Delta \subseteq \mathcal{L}_T$ and $\varphi \in \mathcal{L}_{CL}$ we define:*
- $\Gamma^{\rightsquigarrow} := \{\chi \rightsquigarrow \psi \in \Gamma\}$ *(We write $\Gamma \leftrightsquigarrow \Delta$ if $\Gamma^{\rightsquigarrow} = \Delta^{\rightsquigarrow}$);*
- $\rightsquigarrow_\varphi(\Gamma) := \{\psi : \varphi \rightsquigarrow \psi \in \Gamma\}$ *(We call Δ φ-likely for Γ if $\rightsquigarrow_\varphi(\Gamma) \subseteq \Delta$).*

Fact 1 *$\leftrightsquigarrow$ is an equivalence relation on the set of all MCSs. We write $[\Gamma]_{\leftrightsquigarrow}$ for the equivalence class containing Γ.*

Each equivalence class serves as a basis for a preference frame in our canonical model. The maximal consistent sets which are φ-likely for Γ are our candidates for the most likely φ states in the preference frame based on the equivalence class $[\Gamma]_{\leftrightsquigarrow}$, as they contain all formulas supported by φ.

Before proceeding, we need a result stating that if ψ is not supported by φ in Γ, then we can construct a MCS Δ with the same support formulas as Γ, including the negation of ψ and all the propositions supported by φ.

Lemma 1. *Given a MCS Γ and a propositional formula ψ with $\psi \notin \rightsquigarrow_\varphi(\Gamma)$, then there exists $\Delta \in [\Gamma]_{\leftrightsquigarrow}$ such that $\{\neg\psi\} \cup \rightsquigarrow_\varphi(\Gamma) \subseteq \Delta$.*

Proof. We prove the consistency of the set $A := \{\neg\psi\} \cup \rightsquigarrow_\varphi(\Gamma) \cup \Gamma^{\rightsquigarrow} \cup \{\neg(\chi \rightsquigarrow \gamma) : \chi \rightsquigarrow \gamma \notin \Gamma\}$. If consistent, we can extend it to an MCS $\Delta \subseteq [\Gamma]_{\leftrightsquigarrow}$. Assuming $A \vdash \bot$, there exist $\varphi_1, ..., \varphi_n \in \rightsquigarrow_\varphi(\Gamma)$, $\pi_1 \rightsquigarrow \psi_1, ..., \pi_m \rightsquigarrow \psi_m \in \Gamma^{\rightsquigarrow}$, and $\neg(\chi_1 \rightsquigarrow \gamma_1), ..., \neg(\chi_k \rightsquigarrow \gamma_k) \in \Gamma$ such that $\vdash (\alpha \wedge \varphi_1 \wedge ... \wedge \varphi_n \wedge \neg\psi) \rightarrow \bot$, where $\alpha := (\pi_1 \rightsquigarrow \psi_1) \wedge ... \wedge (\pi_m \rightsquigarrow \psi_m) \wedge \neg(\chi_1 \rightsquigarrow \gamma_1) \wedge ... \wedge \neg(\chi_k \rightsquigarrow \gamma_k)$. Using **CL** axioms, we obtain $\vdash \alpha \rightarrow ((\varphi_1 \wedge ... \wedge \varphi_n) \rightarrow \psi)$. From $\mathbf{S5_F}$, we derive $\vdash \alpha \rightarrow (\neg(\varphi_1 \wedge ... \wedge \varphi_n \rightarrow \psi) \rightsquigarrow \bot)$, and by **Nec**, we get $\vdash \alpha \rightarrow (\varphi \rightsquigarrow (\varphi_1 \wedge ... \wedge \varphi_n \rightarrow \psi))$. By **RCK** and the maximal consistency of the set Γ, this implies $\psi \in \rightsquigarrow_\varphi(\Gamma)$, contradicting our assumption. Thus, A is consistent.

We define the states and preference relations $\succeq_\Gamma$ for our canonical model (the index Γ is a representative of an equivalence class of the equivalence relation $\leftrightsquigarrow$). The states are (Δ, φ, i) where Δ is a MCS in the equivalence class $[\Gamma]_{\leftrightsquigarrow}$, $i \in \{0, 1, 2\}$, and φ is a propositional formula. φ and i are used to pinpoint the maximal φ-states according to the relation $\succeq_\Gamma$ and to ensure that the maximal elements of $\succeq_\Gamma$ coincide with the states satisfying the supported formulas within Γ, see Corollary 2.

For a MCS Γ we use the following notation:

– $S_\Gamma := [\Gamma]_{\leftrightsquigarrow} \times \mathcal{L}_{CL} \times \{0, 1, 2\}$ and $[\delta]_\Gamma := \{(\Delta, \varphi, i) \in S_\Gamma : \delta \in \Delta\}$.

Definition 8. *The preference relation* $\succeq_\Gamma \subseteq S_\Gamma \times S_\Gamma$ *is such that* $(\Delta, \varphi, i) \succeq_\Gamma (\Omega, \psi, j)$ *holds iff at least one of the following conditions holds:*

– Δ *is* φ*-likely for* Γ *and* $\varphi \in \Omega$*;*
– $(i = 1$ *and* $j = 0)$ *or* $(i = 2$ *and* $j = 1)$ *or* $(i = 0$ *and* $j = 2)$*.*

After having defined our preference relations we show that the maximal elements in $[\delta]_\Gamma$ are δ-likely for Γ. This is what we are aiming for as we want all the elements in $max([\delta]_\Gamma)$ to fulfill every formula supported by δ in Γ. Furthermore, if an MCS $\Delta \in S_\Gamma$ is δ-likely for Γ then (Δ, δ, i) is a maximal element in $[\delta]_\Gamma$. These results will be the core of our completeness proof since they can be used to show that a support formula is true in a state of S_Γ if and only if the formula appears in Γ. To prove this, we start proving the following technical lemma that establishes a connection between an (Δ, φ, i) appearing in $max([\delta]_\Gamma)$ and the MCS Γ.

Lemma 2. *Given* $(\Delta, \varphi, i) \in max([\delta]_\Gamma)$ *then:*
(*a*) Δ *is* φ*-likely for* Γ*;* (*b*) $\neg(\delta \to \varphi) \rightsquigarrow \bot \in \Gamma$*.*

Proof. We assume $i = 0$. The other cases are similar.

(a) Given $(\Delta, \varphi, 0) \in max([\delta]_\Gamma)$, we have $(\Delta, \varphi, 1) \succeq_\Gamma (\Delta, \varphi, 0)$ by construction of $\succeq_\Gamma$. This implies $(\Delta, \varphi, 0) \succeq_\Gamma (\Delta, \varphi, 1)$ by maximality. The latter only holds if Δ is φ-likely for Γ, and $\varphi \in \Delta$ since no other condition applies.

(b) Assume to have a MCS $\Omega \in S_\Gamma$ s.t. $\delta \in \Omega$ but $\varphi \notin \Omega$. For Ω we have $(\Delta, \varphi, 0) \not\succeq_\Gamma (\Omega, \delta, 1)$, but $(\Omega, \delta, 1) \succeq_\Gamma (\Delta, \varphi, 0)$ which contradicts $(\Delta, \varphi, 0) \in max([\delta]_\Gamma)$. Hence, it follows that such an MCS $\Omega \in S_\Gamma$ does not exist, which means $[\delta]_\Gamma \subseteq [\varphi]_\Gamma$. In other words, every MCS in S_Γ contains the formula $\delta \to \varphi$. In particular, each MCS Π with $\Pi \in [\Gamma]_{\leftrightsquigarrow}$ contains the formula $\delta \to \varphi$, which means $[\Gamma]^{\rightsquigarrow} \cup \{\neg(\chi \rightsquigarrow \gamma) : \chi \rightsquigarrow \gamma \notin \Gamma\} \cup \{\neg(\delta \to \varphi)\}$ is inconsistent. This lets us infer $[\Gamma]^{\rightsquigarrow} \cup \{\neg(\chi \rightsquigarrow \gamma) : \chi \rightsquigarrow \gamma \notin \Gamma\} \vdash \delta \to \varphi$, hence we can find finitely many support formulas in $\Gamma^{\rightsquigarrow}$ and finitely many negated support formulas in $\{\neg(\chi \rightsquigarrow \gamma) : \chi \rightsquigarrow \gamma \notin \Gamma\}$ s.t. $\vdash (\varphi_1 \rightsquigarrow \psi_1) \wedge \cdots \wedge (\varphi_n \rightsquigarrow \psi_n) \wedge \neg(\chi_1 \rightsquigarrow \gamma_1) \wedge ... \wedge \neg(\chi_k \rightsquigarrow \gamma_k) \to (\delta \to \varphi)$. By the use of $\mathbf{S5_F}$ we get $\vdash (\varphi_1 \rightsquigarrow \psi_1) \wedge \cdots \wedge (\varphi_n \rightsquigarrow \psi_n) \wedge \neg(\chi_1 \rightsquigarrow \gamma_1) \wedge ... \wedge \neg(\chi_k \rightsquigarrow \gamma_k) \to (\neg(\delta \to \varphi) \rightsquigarrow \bot)$ and finally $\neg(\delta \to \varphi) \rightsquigarrow \bot \in \Gamma$.

Corollary 1. *Given* $(\Delta, \varphi, i) \in S_\Gamma$ *then:*
(*a*) $(\Delta, \varphi, i) \in max([\delta]_\Gamma)$ *implies* Δ *is* δ*-likely for* Γ*;*
(*b*) Δ *being* δ*-likely for* Γ *implies* $(\Delta, \delta, i) \in max([\delta]_\Gamma)$*.*

Proof. (a) By $(\Delta, \varphi, i) \in max([\delta]_\Gamma)$ and Lemma 2 we derive $\neg(\delta \to \varphi) \rightsquigarrow \bot \in \Gamma$. Using the derivable rule **LLE** we get $\neg(\delta \leftrightarrow (\delta \wedge \varphi)) \rightsquigarrow \bot \in \Gamma$. Let us take an arbitrary

$\gamma \in \rightsquigarrow_\delta (\Gamma)$; this means $\delta \rightsquigarrow \gamma \in \Gamma$. Since $\neg(\delta \leftrightarrow (\delta \wedge \varphi)) \rightsquigarrow \bot \in \Gamma$ we can apply **LL+** to derive $(\delta \wedge \varphi) \rightsquigarrow \gamma \in \Gamma$. Furthermore, by applying **SH**, we get $\varphi \rightsquigarrow \delta \rightarrow \gamma \in \Gamma$. By Lemma 2 we know that Δ is φ-likely for Γ, therefore $\delta \rightarrow \gamma \in \Delta$. By assumption, we have $\delta \in \Delta$, which lets us conclude $\gamma \in \Delta$. Since γ was arbitrary we get $\rightsquigarrow_\delta (\Gamma) \subseteq \Delta$.

(b) As Δ is δ-likely for Γ, by axiom **ID** we have $\delta \in \rightsquigarrow_\delta (\Gamma) \subseteq \Delta$ and hence $\Delta \in [\delta]_\Gamma$. Taking an arbitrary MCS $\Omega \in S_\Gamma$ with $\Omega \in [\delta]_\Gamma$ and an arbitrary propositional formula π, we end up with $(\Delta, \delta, i) \succeq_\Gamma (\Omega, \pi, j)$ by Definition 8.

We can finally show that our construction works as intended, namely that every formula ψ supported by a formula δ according to a MCS Γ is contained in all the maximal δ states in the equivalence set of Γ.

Corollary 2. *Given a MCS Γ and two propositional formulas δ and ψ, it holds that $\delta \rightsquigarrow \psi \in \Gamma$ if and only if for all $(\Delta, \varphi, i) \in max([\delta]_\Gamma) : \psi \in \Delta$.*

Proof. Given $(\Delta, \varphi, i) \in max([\delta]_\Gamma)$, Corollary 1 ensures that Δ is δ-likely for Γ. Since $\psi \in \rightsquigarrow_\delta (\Gamma) \subseteq \Delta$, the claim follows. Now, by contraposition, we assume $\psi \notin \rightsquigarrow_\delta (\Gamma)$. By Lemma 1, there is an MCS $\Delta \in [\Gamma]_{\leftrightsquigarrow}$ such that $\{\neg\psi\} \cup \rightsquigarrow_\delta (\Gamma) \subseteq \Delta$. Then, Δ is δ-likely for Γ, and Corollary 1 gives $(\Delta, \varphi, i) \in max([\delta]_\Gamma)$. We conclude $\psi \notin \Delta$ by consistency.

Let us fix a set I consisting of one representative of each equivalent class of $\leftrightsquigarrow$.

Definition 9 (Canonical model). *Let $\mathcal{M}^{Can} := \langle S, (S_\Gamma)_{\Gamma \in I}, (\succeq_\Gamma)_{\Gamma \in I}, R, V\rangle$ where:*
- $S := \bigcup_{\Gamma \in I} S_\Gamma$;
- $V(p) := \{(\Delta, \varphi, i) \in S : p \in \Delta\}$;
- *the preference relation* $\succeq_\Gamma \subseteq S_\Gamma \times S_\Gamma$ *is as in Definition 8;*
- $R \subseteq S \times S$ *is defined as* $(\Delta, \varphi, i) R (\Omega, \psi, j)$ *if for all* $\alpha \in \mathcal{L}_T$: $(B(\alpha) \in \Delta \Rightarrow \alpha \in \Omega)$.

We begin the final steps of our completeness proof: the truth lemma and the proof that the canonical model $\mathcal{M}^{Can}$ (from Def. 9) is a Trust model, starting with the truth lemma. Let (Δ, π, i) be a state of the canonical model $\mathcal{M}^{Can}$, s.t. Δ is a MCS (Def. 6), and (π, i) are used to pinpoint arbitrary maximal π-states. This lemma says that a formula α is satisfied by the pointed-canonical model $\mathcal{M}^{Can}, (\Delta, \pi, i)$ iff $\alpha \in \Delta$.

Lemma 3 (Truth lemma). $\mathcal{M}^{Can}, (\Delta, \pi, i) \models \alpha$ *iff* $\alpha \in \Delta$.

Proof. Proceeds by structural induction on α. If $\alpha \in \mathcal{L}_{CL}$, the claim follows from **CL**. By i.h., we derive $||\varphi||_\Gamma = [\varphi]_\Gamma$ and $most(||\varphi||_\Gamma) = max([\varphi]_\Gamma)$ for all $\varphi \in \mathcal{L}_{CL}$.

Let α be $\varphi \rightsquigarrow \psi$. We start with the case $\alpha \in \Delta$. Take Γ with $\Delta \in [\Gamma]_{\leftrightsquigarrow}$. Given an arbitrary $(\Omega, \gamma, j) \in S_\Gamma$ with $(\Omega, \gamma, j) \in max(||\varphi||_\Gamma)$ we get $(\Omega, \gamma, j) \in max([\varphi]_\Gamma)$ by the i.h. Because $\psi \in \rightsquigarrow_\varphi (\Gamma)$ Corollary 2 implies that $\psi \in \Omega$. By the induction hypothesis, we conclude $(\Omega, \gamma, j) \models \psi$, implying $(\Delta, \pi, i) \models \varphi \rightsquigarrow \psi$. For the other direction, assume $\alpha \notin \Delta$. In this case, we have to find a maximal φ state which does not satisfy ψ. The assumption $\psi \notin \rightsquigarrow_\varphi (\Delta)$ lets us derive $\psi \notin \rightsquigarrow_\varphi (\Gamma)$. By Corollary 2, we obtain a $(\Pi, \chi, i) \in max([\varphi]_\Gamma)$ with $\psi \notin \Pi$. By induction hypothesis we get $(\Pi, \chi, i) \in max(||\varphi||_\Gamma)$ and $(\Pi, \chi, i) \models \neg\psi$. This means $(\Delta, \pi, i) \not\models \varphi \rightsquigarrow \psi$.

If α is of the form $B(\varphi)$ both directions of the claim follow directly from the induction hypothesis and the construction of R in Definition 9.

Lemma 4. *$\mathcal{M}^{Can}$ is a Trust model.*

Proof. First, we prove that for each $\Gamma \in I$ $\langle S_\Gamma, (\succeq_\Gamma), V\rangle$ fulfils the limitedness condition. Let $\varphi \in \mathcal{L}_{CL}$ and $\Gamma \in I$ with $[\varphi]_\Gamma \neq \emptyset$. Hence there exists a MCS $\Delta \in [\Gamma]_{\leftrightsquigarrow}$ with $\varphi \in \Delta$. **ST** tells us that $\neg(\varphi \rightsquigarrow \bot) \in \Delta$. This implies $\varphi \rightsquigarrow \bot \notin \Delta$ and finally $\bot \notin \rightsquigarrow_\varphi (\Gamma)$. Given that the set $\rightsquigarrow_\varphi (\Gamma)$ is closed under consequences because of **RCK**, we conclude that it is consistent. By Lemma 1, we can extend $\rightsquigarrow_\varphi (\Gamma)$ to a MCS $\Pi \in [\Gamma]_{\leftrightsquigarrow}$. By construction Π is φ-likely for Γ. By Corollary 1(b) we obtain $(\Pi, \varphi, i) \in max([\varphi]_\Gamma)$, which makes $max([\varphi]_\Gamma)$ non-empty. By Lemma 3 it follows $\|\varphi\|_\Gamma \neq \emptyset \Rightarrow most(\|\varphi\|_\Gamma) \neq \emptyset$. Since φ and Γ were arbitrary we are done. Finally, the fact that the relation R is transitive and serial, follows from the axioms for B as in standard Kripke semantics.

Theorem 5 (Strong Completeness). *$\Phi \models \alpha \Rightarrow \Phi \vdash \alpha$, for all $\Phi \subseteq \mathcal{L}_T$ and $\alpha \in \mathcal{L}_T$.*

Proof. By contraposition. From $\Phi \nvdash \alpha$ it follows that $\Phi \cup \{\neg\alpha\}$ is consistent and can therefore be extended to a MCS Δ. By Lemma 3 every formula in Δ holds in the canonical model in a state of the form (Δ, γ, i). Hence $\forall \beta \in \Phi : (\Delta, \gamma, i) \models \beta$ and $(\Delta, \gamma, i) \not\models \alpha$, which gives us $\Phi \not\models \alpha$ by Lemma 4.

Remark 9. Since the deduction theorem holds in SBTrust, Th. 5 implies compactness.

We discuss the computational complexity of SBTrust, splitting the problem into two parts: (*i*) reducing SBTrust by ignoring its support part and focusing on the Boolean and belief parts, and then (*ii*) reinstating the support part, to complete the proof.

Definition 10. *The SBTrust-reduction is obtained by reducing $\mathcal{L}_T$ to $\mathcal{L}'_T$ and transforming a model $\mathcal{M} := \langle S, \_, \_, R, V\rangle$ into a model $\mathcal{M}' := \langle S, R, V'\rangle$ as follows:*
- *φ formulas of $\mathcal{L}_T$ remain unchanged in $\mathcal{L}'_T$;*
- *All $\alpha \in \mathcal{L}_T$ of the form $\varphi \rightsquigarrow \varphi$, are mapped to fresh atoms taken from a set $Prop'$, where $Prop' \bigcap Prop = \emptyset$; all the other α formulas are adjusted accordingly;*
- *V' extends V to include in its domain $Prop'$, according to the following rule: if $\mathcal{M}, s \models \varphi \rightsquigarrow \psi$ and p' is the atom corresponding to $\varphi \rightsquigarrow \psi$, then $s \in V'(p')$.*

Theorem 6. *The decision problem for SBTrust is PSPACE-complete.*

Proof. The result for a SBTrust-reduction follows from the fact that $\mathcal{L}'_T$ is a set of $\mathbb{KD}4$ formulas and the $\mathcal{M}'$s are serial and transitive models [23]. Now, take a conjunction of the support formulas corresponding to the atoms of $Prop'$ that are mapped to true. The SAT problem for the support formulas is PSPACE: in [36], it is shown that the satisfiability problem for the logic $\mathbb{F}$ is NP-complete. The result follows by Theorem 3.

Finally, we state the complexity of the model checking problem for SBTrust.

Theorem 7. *Given $\mathcal{M} := \langle S, \_, \_, R, V\rangle$ and $\alpha \in \mathcal{L}_T$, let n be the number of states in S and r the number of pairs sRv. Let k be the number of support formulas, and k' the number of belief modalities, plus the number of atomic propositions and of connectives in α. The complexity of the model checking problem is $O(k \cdot n^2 + (k + k') \cdot (n + r))$.*

Proof. We use the splitting methodology: we first consider the modal part of the formula α and then the support part. For the first stage, apply a SBTrust-reduction to $\mathcal{M}$ and α. This takes at most $k + k'$-steps. What is left is a model-checking problem for a modal formula within a pointed Kripke model: its complexity is $O((n+r)\cdot(k+k'))$ [16]. Then, translate back the new atoms to their respective support formulas. Take the partition S_i containing the state of evaluation s. To evaluate $\varphi \rightsquigarrow \psi$, compute the two sets $most(||\varphi||_i)$ and $||\psi||_i$. The latter is straightforward. The former requires at most n^2-steps (assuming the worst case: $||\varphi||_i = S_i = S$). This must be done for all k support formulas, thus, the complexity of the whole procedure is $O(k \cdot n^2)$.

5 Conclusions and Future Works

SBTrust is a logical framework for reasoning about decision trust, built on belief and support. To formalize support, we introduced a novel non-monotonic conditional operator, axiomatizing the flat fragment of the logic $\mathbb{F}$, and based on preference semantics.

Due to the generality of these concepts, SBTrust can integrate elements from the approaches mentioned in Section 1.1 within a unified framework. More precisely, SBTrust can naturally and flexibly represent relevant factors for trust through the support conditional. For example, we can express the importance of cognitive conditions for the emergence of trust without having to give an explicit formalization of them but relying on their implicit support of trust. Following up the discussion initiated in Example 1, we illustrate how SBTrust can be used to combine different elements that contribute to establishing trust and match those to the various paradigms described in Section 1.1.

Example 13. Assume that to trust $GoodV_i$, customer C seeks to fulfill three conditions: i) a cognitive-based one; ii) a reputation-based one; iii) a policy-based one. The cognitive-based condition could be captured by the notion of occurrence Trust (denoted by formula $OccTV_i$) given in [24], which depends on multiple cognitive features of the agents involved such as the goals of C, the ability and intentions of V_i, and the effects of the actions of V_i on the goals of C. The reputation-based condition could be represented, e.g., by proposition $TopRatingV_{i,j}$, while the policy-based condition by proposition $AuthV_i$. Then, the formula $T_\Gamma(GoodV_i)$ will indicate that customer C trusts V_i as a good vendor for reason Γ, where Γ stands for $(OccTV_i \wedge TopRatingV_{i,j} \wedge AuthV_i)$.

The flexibility of SBTrust to express different conditions comes, however, at the expense of reduced deductive power, e.g., we are unable to use the internal structure of cognitive conditions to derive conclusions.

We also claim that our framework is versatile enough to encompass other notions of trust, particularly those based on supportive information, as in [21]. Notably, it also shares similarities with argumentation-based formalizations of trust [5].

Our future research will proceed in two main directions: exploring potential applications of SBTrust and the support operator, and extending the technical results presented in this paper. Application-wise, we plan to use SBTrust to describe trust dynamics in edge computing scenarios [19], by integrating various existing approaches to trust generation - policies, reputation, and cognitive features - within our unified framework. We also expect to leverage SBTrust to model the combination of the various (explainability, legal, privacy- and security-related) factors influencing the trust perception when using generative AI models [31]. From the technical point of view, we plan to study the derived operator T in isolation and identify its properties independently of support and belief. Moreover, we intend to extend SBTrust by: (*i*) using beliefs within support statements; (*ii*) providing a proof calculus, along the lines of that in [14], equipped with a prover; (*iii*) moving towards a quantitative, dynamic, and multi-agent setting.

Acknowledgements: Work partially supported by the FWF project LoDEx Grant-DOI 10.55776/I6372 and by the Italian Ministry of University and Research programme “PRIN 2022" - grant number 2022598LMZ - AsCoT-SCE - CUP: H53D23003430006.

References

1. M. Abadi. Logic in access control (tutorial notes). In Foundations of Security Analysis and Design V: FOSAD 2007/2008/2009 Tutorial Lectures, pages 145–165. Springer, 2009.
2. A. Aldini, G. Curzi, P. Graziani, and M. Tagliaferri. A probabilistic modal logic for context-aware trust based on evidence. International Journal of Approximate Reasoning, 169, 2024.
3. A. Aldini and M. Tagliaferri. A taxonomy of computational models for trust computing in decision-making procedures. In A. Jøsang, editor, European Conference on Information Warfare and Security, pages 571–578, 2018.
4. A. Aldini and M. Tagliaferri. Logics to reason formally about trust computation and manipulation. In A. Saracino and P. Mori, editors, Emerging Technologies for Authorization and Authentication, volume LNCS 11967, pages 1–15. Springer, 2020.
5. L. Amgoud and R. Demolombe. An argumentation-based approach for reasoning about trust in information sources. Argument & Computation, 5(2-3):191–215, 2014.
6. M. Amoozadeh, D. Daniels, D. Nam, A. Kumar, S. Chen, M. Hilton, R.S. Srinivasa, and M.A. Alipour. Trust in generative AI among students: An exploratory study. In 55th ACM Tech. Symp on Computer Science Education V. 1, SIGCSE 2024, page 67–73. ACM, 2024.
7. L. Åqvist. Deontic logic. In D. Gabbay and F. Guenthner, editors, Handbook of Philosophical Logic: Volume II, pages 605–714. Springer, 1984.
8. D. Artz and Y. Gil. A survey of trust in computer science and the semantic web. Journal of Web Semantics, 5(2):58–71, 2007.
9. A. Benlahbib and E.H. Nfaoui. AmazonRep: A reputation system to support Amazon’s customers purchase decision making process based on mining product reviews. In 4th Int. Conf. On Intelligent Computing in Data Sciences (ICDS), pages 1–8, 2020.
10. M.H. Birnbaum, D. Navarro-Martinez, C. Ungemach, N. Stewart, and E.G. Quispe-Torreblanca. Risky decision making: Testing for violations of transitivity predicted by an editing mechanism. Judgment and Decision Making, 11(1):75–91, 2016.
11. A. Bochman. A logic for causal reasoning. In G. Gottlob and T. Walsh, editors, Proceedings of IJCAI-03, pages 141–146. Morgan Kaufmann, 2003.
12. D. Braga, M. Niemann, B. Hellingrath, and F. Neto. Survey on computational trust and reputation models. ACM Comput. Surv., 51(5), 2018.
13. C. Castelfranchi and R. Falcone. Trust Theory: A Socio-Cognitive and Computational Model. John Wiley and Sons, 2010.
14. A. Ciabattoni and D. Rozplokhas. Streamlining input/output logics with sequent calculi. In Principles of Knowledge Representation and Reasoning, pages 146–155, 2023.
15. V. Crupi and A. Iacona. Three ways of being non-material. Studia Logica, 110:47–93, 2022.

16. R. Fagin, J.Y. Halpern, Y. Moses, and M. Vardi. Reasoning about Knowledge. The MIT Press, 2003.
17. L. Fenrong and L. Emiliano. Reasoning about belief, evidence and trust in a multi-agent setting. In Principles and Practice of Multi-Agent Systems, volume LNCS 10621, pages 71–89. Springer, 2017.
18. P. C. Fishburn. Nontransitive preferences in decision theory. Journal of Risk and Uncertainty, 4(2):113–134, 1991.
19. L. Fotia, F. Delicato, and G. Fortino. Trust in edge-based Internet of Things architectures: State of the art and research challenges. ACM Comput. Surv., 55(9), 2023.
20. K. Frost-Arnold. The cognitive attitude of rational trust. Synthese, 191(9):1957–1974, 2014.
21. D. Gambetta, editor. Trust: Making and Breaking Cooperative Relations. Blackwell, 1988.
22. D. Grossi, W. van der Hoek, and L.B. Kuijer. Reasoning about general preference relations. Artificial Intelligence, 313:103793, 2022.
23. J.Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. Artificial Intelligence, 54(3):319–379, 1992.
24. A. Herzig, E. Lorini, J.F. Hübner, and L. Vercouter. A logic of trust and reputation. Logic Journal of the IGPL, 18(1):214–244, 12 2009.
25. L. Humberstone. Philosophical Applications of Modal Logic. College Publications, 2016.
26. A. Jacovi, A. Marasović, T. Miller, and Y. Goldberg. Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in AI. In Fairness, Accountability, and Transparency, page 624–635. ACM, 2021.
27. C.D. Jensen. The importance of trust in computer security. In J. Zhou, N. Gal-Oz, J. Zhang, and E. Gudes, editors, Trust Management VIII, pages 1–12. Springer, 2014.
28. S.D. Kamvar, M.T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In 12th Int. Conf. on World Wide Web, pages 640–651, 2003.
29. S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. Artif. Intell., 44(1):167–207, 1990.
30. H.E. Kyburg. Probability and the logic of rational belief. Wesleyan University Press, 1961.
31. A. Leschanowsky, S. Rech, B. Popp, and T. Bäckström. Evaluating privacy, security, and trust perceptions in conversational AI: A systematic review. Comp. in Hum. Beh., 159, 2024.
32. D. Makinson and L. van der Torre. Input/output logics. J. of Phil. Logic, 29:383–408, 2000.
33. S. Marsh. Formalising Trust as a Computational Concept. University of Stirling, 1994.
34. X. Parent. Completeness of Åqvist's systems E and F. The Review of Symbolic Logic, 8(1):164–177, 2015.
35. D. Poole. What the lottery paradox tells us about default reasoning. In Principles of Knowledge Representation and Reasoning, pages 333–340. Morgan Kaufmann, 1989.
36. D. Rozplokhas. LEGO-like small-model constructions for Åqvist's logics. Proceedings of AIML 2024, 2024.
37. G.F. Schumm. Transitivity, preference and indifference. Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition, 52(3):435–437, 1987.
38. T. Schwartz. Cycles and Social Choice: The True and Unabridged Story of a Most Protean Paradox. Cambridge University Press, 2018.
39. K.E. Seamons, M. Winslett, L. Yu, R. Jarvis, A. Hess, J. Jacobson, B. Smith, and T. Yu. Negotiating trust on the web. IEEE Internet Computing, 6(06):30–37, nov 2002.
40. Y. Shoham. A semantical approach to non-monotonic logics. In Symposium on Logic in Computer Science (LICS 1987), pages 275–279. IEEE CS Press, 1987.
41. A. Smith. What Amazon key teaches us about trust in customer relationships. Forbes, November, 2017.
42. M. Soleimani. Buyers' trust and mistrust in e-commerce platforms: a synthesizing literature review. Inf. Syst. E-bus. Manag., 20(1):57–78, 2022.
43. M. Tagliaferri and A. Aldini. From belief to trust: A quantitative framework based on modal logic. Journal of Logic and Computation, 32(6):1017–1047, 03 2022.