

Decidable Reasoning in a First-Order Logic of Limited Conditional Belief

Christoph Schwering¹ and Gerhard Lakemeyer²

Abstract. In a series of papers, Liu, Lakemeyer, and Levesque address the problem of decidable reasoning in expressive first-order knowledge bases. Here, we extend their ideas to accommodate *conditional beliefs*, as in “if she is Australian, then she presumably eats Kangaroo meat.” Perhaps the most prevalent semantics of a conditional belief is to evaluate the consequent in the most-plausible worlds consistent with the premise. In this paper, we devise a technique to approximate this notion of plausibility, and complement it with Liu, Lakemeyer, and Levesque’s weak inference. Based on these ideas, we develop a logic of *limited* conditional belief, and provide soundness, decidability, and (for the propositional case) tractability results.

1 INTRODUCTION

Taking into account different contingencies is elementary for humans. For example, when we expect a guest for dinner, we might believe her to be *not* Australian, but at the same time believe that if she *is* Australian, then she presumably eats Kangaroo meat – and we would plan the menu based on these beliefs. It is no surprise that the concept of *conditional belief* is a subject of KR research, for example in belief revision [13].

The goal of this paper is to devise a reasoning service that soundly decides whether a given conditional knowledge base entails some conditional belief. As with any knowledge-based system, this requires to trade off expressivity and/or reasoning power against computational feasibility. We will address this question by beginning with a fully fledged first-order logic of conditional belief, and then weaken its inference mechanism in order to achieve decidability, while retaining the expressivity of a first-order language as well as soundness wrt the fully fledged logic.

Our approach is based on Liu, Lakemeyer, and Levesque’s (henceforth LLL) work on *limited reasoning* [26, 19, 20, 21], where they address the problem of decidable reasoning in expressive first-order knowledge bases. They stratify belief in levels: level 0 only contains the agent’s explicit beliefs; every following level k adds the beliefs that can be inferred after k case splits, that is, by branching on the possible truth assignments of k literals.

Two key features distinguish limited reasoning from other approaches to decidable first-order reasoning such as description logics. Firstly, no limits are set with regard to first-order expressivity in the query. Naturally then, soundness and decidability come at the cost of completeness. Secondly, limited reasoning solves the problem of logical omniscience [11]. Omniscience means that an agent knows

all logical consequences of his knowledge, which is obviously not realistic. LLL avoid omniscience by limiting the number of allowed case splits; intuitively, this number represents the *effort* the reasoner may spend to prove a belief.

Surprisingly perhaps, generalizing LLL’s ideas to the case of *conditional* belief is far from trivial. The following running example shall illustrate this.

Example 1 Suppose we expect a guest for dinner. We don’t know much about her preferred diet yet, but we do have some (somewhat narrow-minded) beliefs:

- Most Australians are not Italians, and vice versa.
- Australians usually eat kangaroo meat.
- We believe our guest is Italian or a vegetarian, and if she is not Italian, she presumably is Australian.
- We know that kangaroo is meat, and that vegetarians do not eat meat.

Given this conditional knowledge base, do we expect our guest to be a vegetarian in case she is not Italian? Monotonic reasoning would suggest so: our belief of her being Italian or a vegetarian yields that she must be a vegetarian if not Italian. But then she also must be Australian, hence eat kangaroo, and thus be a non-vegetarian – that is, in a monotonic logic *everything* is believed if she is not Italian, because the beliefs are inconsistent with that contingency. Conditional beliefs do not trap into inconsistency that easily. They detect that the premise “not Italian” is inconsistent with the most-plausible scenarios, and therefore go on to look for *less-plausible* scenarios where the premise holds, and check the consequent in these scenarios. In our example, we hence believe that if the guest is not Italian, she presumably is an Australian kangaroo-eater, but not a vegetarian.

It is the consistency test that makes reasoning about conditional beliefs more involved than normal beliefs. To understand why, we need some details. An agent’s epistemic state with the contingencies he considers possible can be represented as a *system of spheres* [25, 9] as in Figure 1a (page 3). The innermost sphere contains the most-plausible scenarios, and every outer sphere contains additional less-plausible scenarios. A conditional “if ϕ , then presumably ψ ” is believed iff the material implication $\phi \supset \psi$ holds in the innermost sphere that is consistent with ϕ . LLL provide us with a limited semantics of first-order logic which is suitable for checking if $\phi \supset \psi$ holds in a given sphere. On the other hand, it is insufficient for the consistency check: as the semantics is incomplete wrt classical logic, inconsistencies might go undetected, and wrong beliefs could come out true. For that reason we devise an additional *complete but unsound* semantics complementary to LLL’s: the former can be used to

¹ RWTH Aachen University, Germany and The University of New South Wales, Australia, email: schwering@kbsg.rwth-aachen.de

² RWTH Aachen University, Germany, email: gerhard@kbsg.rwth-aachen.de

soundly determine consistency, and the latter for sound inferences. Together, they will allow us to develop a logic of limited conditional belief.

The rest of the paper is structured as follows. After a survey of related work in the next section, we proceed with the technical part. To start with, we introduce the logic \mathcal{BO} for reasoning about conditional belief; it will serve as a reference point for the following results. Section 4 presents a derivative of LLL’s limited semantics of first-order logic, and introduces a novel unsound but complete semantics. The stage is then set for the paper’s main contribution in Section 5: a logic of limited conditional belief called \mathcal{BCL} , which is sound wrt \mathcal{BO} and decidable for a large class of knowledge bases and sometimes even tractable. Then we conclude and discuss future work. While we illustrate all techniques using the above example, we only sketch most proofs for space reasons. The full proofs can be found in [33].

2 RELATED WORK

There are two principal directions to tackle the problem of decidable first-order reasoning: restricting the language or restricting the inference mechanism. The earliest classes of syntactic fragments of first-order logic are prefix-vocabulary classes of formulas in prenex normal form whose symbols are taken from a certain vocabulary and whose leading quantifier prefix adheres to a specific form. Today, it is known which prefix-vocabulary classes are decidable and which are not [4]. Modern work on decidable fragments concerns bounded variable logics such as the two-variable fragment [37, 28], which connects first-order logic with propositional modal logic [8] and the prototypical description logic \mathcal{ALC} [29], which in turn has also been the basis for epistemic description logics [2].

Unlike these syntactic approaches, the direction we take here is based on restricting the inference mechanism. In other words, instead of exchanging expressivity for decidability as do syntactic fragments of first-order logic and description logics, we aim to trade off completeness against decidability. The standard tool to give semantics to knowledge and belief are Hintikka-style possible-worlds semantics [10]. Typically they imply omniscience [11], which brings along undecidability in the first-order case and intractability in the propositional case. Approaches to solve the omniscience problem include rather syntactic ones [15, 38, 6] as well as semantic approaches [22, 30, 16, 17, 5] based on tautological entailment [3, 1]. The former have the drawback of being very fine-grained and providing only little guidance as to which beliefs to include and which to leave out. The semantic approaches based on three- or four-valued semantics, by contrast, are much closer to the classical possible-worlds semantics, but miss out on many seemingly trivial inferences, such as simple unit propagation.

LLL’s work on limited belief [26, 19, 20, 21] steers a middle course between these semantic and syntactic approaches. As we shall see, their semantics has a somewhat syntactic flavour, yet it is perspicuously defined and motivated. Klassen et al. [14] recently proposed a neighbourhood semantics that avoids the syntactic flavour; however, it is restricted to the propositional case. The closest relative of our approach among LLL’s work is [20].

Although the omniscience problem and logics of limited belief have been around for over thirty years, the present paper is, as far as we know, the first to address the problem in the context of conditional belief. The underlying logic of conditional belief used as a reference for the limited logic developed in this paper is an extension of Levesque’s logic of only-knowing [23, 24] to the case of conditional belief. Its semantics is defined in terms of plausibility-ranked

possible worlds.

Lewis [25] was the first to pick up the concept of possible worlds to semantically characterize conditionals; Grove [9] adopted this view and popularized it for belief change. The idea is to sequentially nest sets of possible worlds to obtain a so-called system of spheres. Intuitively, a system of spheres induces a plausibility ranking on the worlds: the most plausible worlds are contained in the innermost set, the second-most plausible worlds in the next set, and so on. Another popular and equivalent technique are total preorders [12].

Unlike material implications (but like counterfactuals), conditional belief is non-monotonic in the sense that strengthening the antecedent is not valid. The closest relative among the non-monotonic-reasoning approaches is Pearl’s System Z [31]: it can be shown that conditional belief and only-believing subsume the non-monotonic 1-entailment in System Z [33].

As we saw already, evaluating a conditional in a system of spheres requires a consistency check. A sound consistency check, as necessitated for a sound semantics of conditionals, in turn requires a complete semantics. Relatively little work has been done in this area yet. Schaerf and Cadoli [32] address both unsound and incomplete inference, but they apply propositional techniques to restricted fragments of first-order logic. Recently, other methods of unsound reasoning have been investigated [27, 7]. However, as these approaches are not based primarily on subsumption and unit propagation, they do not fit well with LLL’s techniques. We therefore propose a new complete semantics to complement LLL’s sound inference.

3 CONDITIONAL BELIEF IN \mathcal{BO}

This section introduces the first-order logic of conditional belief \mathcal{BO} [34]. It features two modal operators to represent beliefs, whose semantics is defined through a system of spheres, that is, a sequence of nested sets of possible worlds.

The language \mathcal{BO} is defined as follows. The set of *terms* is the least set that contains all first-order variables x and all (*standard*) *names* $n \in \mathcal{N} = \{\#1, \#2, \dots\}$. The set of *formulas* is the least set of expressions such that

- $P(t_1, \dots, t_j), (t_1 = t_2), (\alpha \vee \beta), \neg\alpha, \exists x\alpha,$
- $\mathbf{B}(\phi_1 \Rightarrow \psi_1), \mathbf{O}\{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$

are formulas, where P is a predicate symbol other than $=$, t_i are terms, α and β are formulas, and ϕ_i and ψ_i are *objective* formulas, that is, ϕ_i and ψ_i mention no \mathbf{B} or \mathbf{O} operators. $P(t_1, \dots, t_j)$ and $(t_1 = t_2)$ are called (*non-*)*equality atoms*, respectively. A *literal* is an atom a or its negation $\neg a$. The complement of a literal ℓ is written $\bar{\ell}$. A formula is *ground* when it contains no variables. A *sentence* is a formula without free variables. We let $\wedge, \forall, \supset, \equiv$ be the usual abbreviations, \top stand for $\exists x(x = x)$, and \perp abbreviate $\neg\top$.

Standard names are special constants that represent all the individuals in the universe. They allow for a simpler semantics than the classical Tarskian model theory; in particular, quantification can be handled by simply substituting standard names. The formula $\mathbf{B}(\phi \Rightarrow \psi)$ intuitively expresses a conditional belief, namely the belief that if ϕ is true, then presumably ψ is also true. The formula $\mathbf{O}\{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$ goes one step further and says that the conditional beliefs $\mathbf{B}(\phi_i \Rightarrow \psi_i)$ are *all* that is believed; everything that is not a consequence of these conditional beliefs is implicitly not believed. The concept is called *only-believing*; it generalizes Levesque’s only-knowing [23, 24] to the case of conditional beliefs.

Before we give a semantics to this language, let us see how Example 1 can be expressed in \mathcal{BO} .




Figure 1. An epistemic state and two approximations. Each ellipse represents the scenarios considered possible in that sphere.

Example 2 Let A, I, V represent that the guest is Australian, Italian, a vegetarian, respectively; $E(x)$ that x is among her preferred diet; $M(x)$ that x is meat; and roo be a standard name representing kangaroo meat. Then all we believe is

- $A \Rightarrow \neg I$ and $I \Rightarrow \neg A$;
 - $A \Rightarrow E(roo)$;
 - $\top \Rightarrow I \vee V$ and $\neg I \Rightarrow A$;
 - $\neg M(roo) \Rightarrow \perp$ and $\neg \forall x (V \wedge M(x) \supset \neg E(x)) \Rightarrow \perp$.
- } = Γ

($\neg\phi \Rightarrow \perp$ expresses indefeasible knowledge of ϕ .) The question “if she is not Italian, is she presumably not a vegetarian?” then boils down to whether $\mathbf{O}\Gamma$ entails $\mathbf{B}(\neg I \Rightarrow \neg V)$.

To investigate such problems, we define a possible-worlds semantics for \mathcal{BO} .

Definition 3 A world is a set of ground non-equality atoms. An epistemic state \vec{e} is an infinite sequence of sets of worlds e_1, e_2, \dots such that $e_1 \subseteq e_2 \subseteq \dots$ and for some $p \in \{1, 2, \dots\}$, $e_p = e_{p+1} = \dots$.

Intuitively, a world is a truth assignment of the ground non-equality atoms, that is, of all $P(n_1, \dots, n_j)$ where the n_i are standard names. An epistemic state stratifies sets of such worlds with the subset relation. The intuition behind an epistemic state \vec{e} is to model a system of spheres as discussed in the introduction: e_1 contains the most-plausible worlds, e_2 adds the second-most-plausible worlds, and so on. Note that in any epistemic state \vec{e} only finitely many different sets of worlds are allowed, since the definition requires that $e_p = e_{p+1} = \dots$ for some p . (Just as well we could have defined an epistemic state to be a non-empty finite sequence of supersets; the present definition however is often easier to work with.)

Given an epistemic state \vec{e} and a world w , we can define truth of a sentence α in \mathcal{BO} , written $\vec{e}, w \models \alpha$. We let α_n^x denote the result of substituting all free occurrences of x by n . The objective part of the semantics is defined inductively as follows:

1. $\vec{e}, w \models P(n_1, \dots, n_j)$ iff $P(n_1, \dots, n_j) \in w$;
2. $\vec{e}, w \models (n_1 = n_2)$ iff n_1 and n_2 are identical names;
3. $\vec{e}, w \models (\alpha \vee \beta)$ iff $\vec{e}, w \models \alpha$ or $\vec{e}, w \models \beta$;
4. $\vec{e}, w \models \neg\alpha$ iff $\vec{e}, w \not\models \alpha$;
5. $\vec{e}, w \models \exists x \alpha$ iff $\vec{e}, w \models \alpha_n^x$ for some $n \in \mathcal{N}$.

To define the semantics of belief, it is convenient to define the *plausibility* $[\vec{e} | \phi]$ of an objective sentence ϕ in \vec{e} as the index of the first sphere consistent with ϕ :

$$[\vec{e} | \phi] = \min\{p \mid p = \infty \text{ or } \vec{e}, w \models \phi \text{ for some } w \in e_p\},$$

where $\infty \notin \{1, 2, \dots\}$ represents an “undefined” plausibility with the understanding that $p + \infty = \infty$ and $p < \infty$ for all $p \in \{1, 2, \dots\}$. Then the semantics of beliefs is as follows:

6. $\vec{e}, w \models \mathbf{B}(\phi \Rightarrow \psi)$ iff for all $p \in \{1, 2, \dots\}$, if $p \leq [\vec{e} | \phi]$ and $w' \in e_p$, then $\vec{e}, w' \models (\phi \supset \psi)$;
7. $\vec{e}, w \models \mathbf{O}\{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$ iff for all $p \in \{1, 2, \dots\}$, $w' \in e_p$ iff $\vec{e}, w' \models \bigwedge_{i: [\vec{e} | \phi_i] \geq p} (\phi_i \supset \psi_i)$.

Rules 1–5 are straightforward. Notable is perhaps that quantification can be handled substitutionally thanks to standard names. The most interesting rules are the ones for belief, of course. According to Rule 6, $\mathbf{B}(\phi \Rightarrow \psi)$ holds iff $\phi \supset \psi$ is true in the innermost sphere consistent with ϕ . Only-believing $\mathbf{O}\{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$ has the effect of $\mathbf{B}(\phi_i \Rightarrow \psi_i)$ plus maximizing every sphere: it requires the sphere e_p to contain all worlds that satisfy all $(\phi_i \supset \psi_i)$ for which $[\vec{e} | \phi_i] \geq p$.

When $\vec{e}, w \models \alpha$ for all \vec{e} (or w), we allow ourselves to omit \vec{e} (or w , respectively). As usual, we use the symbol \models also to denote entailment. In particular, $\mathbf{O}\Gamma \models \mathbf{B}(\phi \Rightarrow \psi)$ is to say that for all \vec{e} , if $\vec{e} \models \mathbf{O}\Gamma$, then $\vec{e} \models \mathbf{B}(\phi \Rightarrow \psi)$.

The fundamental property of only-believing is that it uniquely determines the epistemic state.

Theorem 4 There is a unique epistemic state \vec{e} such that $\vec{e} \models \mathbf{O}\{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$.

In fact, \vec{e} can be generated inductively: when spheres e_1, \dots, e_{p-1} are determined, we can decide whether $[\vec{e} | \phi_i] \geq p$ for every i , and thus determine the next sphere e_p using the right-hand side of Rule 7. A proof of the theorem and this construction can be found in [34]; here, we illustrate the process by continuing Example 2 instead.

Example 5 The first sphere e_1 of \vec{e} such that $\vec{e} \models \mathbf{O}\Gamma$ contains all worlds that satisfy all materialized conditionals from Γ :

$$e_1 = \{w \mid w \models (\neg A \vee \neg I) \wedge (\neg A \vee E(roo)) \wedge (I \vee V) \wedge (I \vee A) \wedge \mu\}$$

where $\mu = M(roo) \wedge \forall x (V \wedge M(x) \supset \neg E(x))$ represents our knowledge about meat and vegetarians.

For the next sphere, we need to figure out the plausibilities $[\vec{e} | \phi]$ for the conditionals $\phi \Rightarrow \psi \in \Gamma$. To begin with, we need to answer if $[\vec{e} | A] \geq 2$, that is, if e_1 is inconsistent with A . To this end, we can split on V : from V we obtain $\neg E(roo)$ (by μ) and thus $\neg A$; on the other hand, from $\neg V$ we infer I and thus $\neg A$; so indeed e_1 is inconsistent with A , that is, $[\vec{e} | A] \geq 2$. By the same argument, $[\vec{e} | \neg I] \geq 2$. It is moreover easy to see that e_1 is consistent and thus $[\vec{e} | I] = [\vec{e} | \top] = 1$. Hence the conditionals $A \Rightarrow \neg I$, $A \Rightarrow E(roo)$, $\neg I \Rightarrow A$, plus the knowledge about meat and vegetarians determine the second sphere:

$$e_2 = \{w \mid w \models (\neg A \vee \neg I) \wedge (\neg A \vee E(roo)) \wedge (I \vee A) \wedge \mu\}.$$

Again we need to check which premises are consistent with e_2 , and only the remaining conditionals determine the next sphere e_3 . It is easy to see that $[\vec{e} | A] = [\vec{e} | \neg I] = 2$, so for the third and last sphere:

$$e_3 = \{w \mid w \models \mu\}.$$

Since \vec{e} is the unique model of $\mathbf{O}\Gamma$ in \mathcal{BO} , it determines our beliefs. For example, $\mathbf{O}\Gamma \models \mathbf{B}(\neg I \Rightarrow \neg V)$ since $[\vec{e} | \neg I] = 2$ and $w \models \neg I \supset \neg V$ for all $w \in e_2$.

Now, how could a limited, decidable version of \mathcal{BO} look like? As for $\mathbf{B}(\phi \Rightarrow \psi)$, we sketched the idea in the introduction already: approximate the plausibility of ϕ from above, and then use sound inference to check whether that sphere satisfies $\phi \supset \psi$. Even if the

\models	satisfaction and entailment in \mathcal{BO}
\models^s	satisfaction in sound first-order semantics
\models^c	satisfaction in complete first-order semantics
\models^u	satisfaction and entailment in \mathcal{BOC}

Table 1. The zoo of turnstile symbols used in the paper.

plausibility approximation leads to a too-far-out sphere, this is sound because what can be inferred from an outer sphere can also be inferred from any inner sphere. But another problem is how to approximate the model of $\mathbf{O}\{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$. Example 5 shows that determining this epistemic state \vec{e} is not trivial, as reasoning is necessary to figure out which beliefs are more plausible than others, that is, which plausibilities are $\geq p$ to determine the p th sphere. Our approximation of \vec{e} will be based on a lower and an upper bound of the plausibilities of ϕ_i . As long as both bounds agree for every i on whether the plausibility of ϕ_i is $\geq p$, we can faithfully represent the p th sphere in the approximation. Once the bounds are inconsistent, though, it is not clear which conditionals shall determine the p th sphere of the approximation, and hence we skip to the final sphere, which represents at least those scenarios contained in the last sphere of \vec{e} . Two such approximated epistemic states are depicted in Figure 1: in Figure 1b the bounds are inconsistent already for the second sphere; Figure 1c faithfully represents the first two spheres, but is pessimistic about the outermost ones, that is, considers too many scenarios. It is important that the last sphere of the approximation must not be optimistic, for otherwise it might satisfy formulas that the last sphere of \vec{e} does not.

4 LIMITED OBJECTIVE REASONING

Here we introduce a sound but incomplete and another complete but unsound semantics for objective formulas. As argued before, they will lay the groundwork for \mathcal{BOC} , the limited version of \mathcal{BO} . Before we elaborate on what is meant by soundness and completeness, we define some fundamentals.

Definition 6 A *clause* is a set of literals $[\ell_1, \dots, \ell_k]$ (we use square brackets to ease readability). The empty clause is written as $[\]$. Every non-empty clause corresponds to the disjunction $(\ell_1 \vee \dots \vee \ell_k)$ (with arbitrary brackets and order). A *setup* is a set of ground clauses.

As in [26, 19, 20, 21], setups are the primitives of our limited semantics. Just like a set of possible worlds, a setup represents possibly incomplete or disjunctive information. The semantics \models^s and \models^c are sound and complete, respectively, in the following sense:

- whenever a setup s satisfies ϕ in the *sound* semantics \models^s , s classically entails ϕ , that is, every world that satisfies all $c \in s$ also satisfies ϕ in unlimited first-order logic;
- whenever a setup s classically entails ϕ , s satisfies ϕ in the *complete* semantics \models^c .

Inference in both semantics is based on unit propagation and subsumption. For example, if $[A, I]$ and $[\neg I]$ are in the setup, then $[A]$ is inferred by unit propagation, and $[A, V]$ follows by subsumption.

Definition 7 For a setup s , we write s^- to remove all subsumed clauses, s^+ to add all subsumed clauses, and $\text{UP}(s)$ to close s to-

gether with all valid equality literals under unit propagation:

$$\begin{aligned}
s^- &= \{c \in s \mid \text{for all } c' \subsetneq c, c' \notin s\}; \\
s^+ &= \{c \mid \text{for some } c' \subseteq c, c' \in s\}; \\
\text{EQ} &= \{[(n = n)], [(n \neq n')] \mid \text{distinct } n, n' \in \mathcal{N}\}; \\
\text{UP}(s) &= \text{closure of } \text{EQ} \cup s \text{ under unit propagation.}
\end{aligned}$$

We also write $\text{UP}^-(s)$ for $\text{UP}(s)^-$, and $\text{UP}^+(s)$ for $\text{UP}(s)^+$.

The following lemma states the equivalence of s , s^- , s^+ , $\text{UP}(s)$.

Lemma 8 For any world w and setup s ,

- $w \models c$ for all $c \in s$ iff
- $w \models c$ for all $c \in s^-$ iff
- $w \models c$ for all $c \in s^+$ iff
- $w \models c$ for all $c \in \text{UP}(s)$.

Proof. Showing the equivalence of (i)–(iii) is straightforward. Here we only show that (i) iff (iv). The if direction is immediate. For the converse, suppose (i) holds and let $c \in \text{UP}(s)$. The proof is by induction on the length of the derivation of c . For the base case let $c \in \text{EQ} \cup s$. If $c \in \text{EQ}$, then clearly $w \models c$. If $c \in s$, then $w \models c$ by assumption. For the induction step suppose $c \in \text{UP}(s)$ is the resolvent of $c \cup [\ell]$, $[\bar{\ell}] \in \text{UP}(s)$. By induction, $w \models c \vee \ell$ and $w \models \bar{\ell}$. Thus $w \models c$. \square

4.1 Sound but incomplete semantics

We are now ready to define the first satisfaction relation: $s, k \models^k \phi$ for a setup s , $k \in \{0, 1, 2, \dots\}$, and an objective sentence ϕ . Intuitively, k indicates how much *effort* is put into proving that ϕ is true in s . The idea is that at level $k = 0$, only obvious inferences can be made, that is, only what is subsumed by the setup s comes out true. At every higher split level $k > 0$, the reasoner may pick a literal ℓ and consider the cases where the unit clauses $[\ell]$ and $[\bar{\ell}]$ are added to the setup. When ℓ is chosen smartly, this may set off unit propagation and lead to new inferences.

We define the sound semantics as follows:

- $s, k + 1 \models^k \phi$ iff $s \cup \{[\ell]\}, k \models^k \phi$ and $s \cup \{[\bar{\ell}]\}, k \models^k \phi$ for some ground literal ℓ ;
- if c is a clause:
 $s, 0 \models^k c$ iff $c \in \text{UP}^+(s)$;
- if $(\phi \vee \psi)$ is not a clause:
 $s, 0 \models^k (\phi \vee \psi)$ iff $s, 0 \models^k \phi$ or $s, 0 \models^k \psi$;
- $s, 0 \models^k \neg(\phi \vee \psi)$ iff $s, 0 \models^k \neg\phi$ and $s, 0 \models^k \neg\psi$;
- $s, 0 \models^k \neg\neg\phi$ iff $s, 0 \models^k \phi$;
- $s, 0 \models^k \exists x \phi$ iff $s, 0 \models^k \phi_n^x$ for some $n \in \mathcal{N}$;
- $s, 0 \models^k \neg\exists x \phi$ iff $s, 0 \models^k \neg\phi_n^x$ for all $n \in \mathcal{N}$.

Let us first illustrate how the definition works by way of our running example.

Example 9 Let $s_\mu = \{[M(\text{roo})], [\neg M(n), \neg E(n), \neg V] \mid n \in \mathcal{N}\}$ and $s_1 = \{[\neg A, \neg I], [\neg A, E(\text{roo})], [I, V], [I, A]\} \cup s_\mu$. This setup corresponds to the first sphere e_1 from Example 5. There we argued that it is inconsistent with A . To obtain the same result in this limited semantics, that is, $s_1, k \models^k \neg A$, one split is needed, that is, $k \geq 1$: clearly, $[\neg A] \notin \text{UP}^+(s_1)$; but adding $[V]$ to s_1 triggers unit propagation that first yields $[\neg M(\text{roo}), \neg E(\text{roo})]$, then $[\neg E(\text{roo})]$, and

then $[\neg A]$; on the other hand, adding $[\neg V]$ yields $[I]$ and then again $[\neg A]$. Hence, $s_1, k \models \neg A$ iff $k \geq 1$. Analogously we can argue that $s_1, k \models I$ iff $k \geq 1$.

The following theorem establishes the aforementioned soundness of \models wrt classical logic. We write $s \models \phi$ to say that $w \models \phi$ for all w with $w \models c$ for all $c \in s$.

Theorem 10 *If $s, k \models \phi$, then $s \models \phi$.*

Proof. By induction on k . We show the base case $k = 0$ by subinduction on $|\phi|$. For any clause, $s, 0 \models c$ iff $c \in \text{UP}^+(s)$ only if $\text{UP}^+(s) \models c$ iff (by Lemma 8) $s \models c$. The other subinduction cases are trivial; for example, for an existential, $s, 0 \models \exists x \phi$ iff $s, 0 \models \phi_n^x$ for some $n \in \mathcal{N}$ only if (by subinduction) $s \models \phi_n^x$ for some $n \in \mathcal{N}$ only if $s \models \exists x \phi$.

For the main induction step suppose the lemma holds for k and that $s, k + 1 \models \phi$. Suppose $w \models c$ for all $c \in s$. By the Rule 1, $s \cup \{[\ell]\}, k \models \phi$ and $s \cup \{[\bar{\ell}]\}, k \models \phi$ for some ℓ . By induction, $s \cup \{[\ell]\} \models \phi$ and $s \cup \{[\bar{\ell}]\} \models \phi$. Therefore, since either $w \models \ell$ or $w \models \bar{\ell}$, we have $w \models \phi$. Hence $s \models \phi$. \square

Another interesting property is the following so-called eventual completeness for propositional formulas.

Theorem 11 *Let s be finite and ϕ be propositional.*

Then $s, k \models \phi$ for some $k \in \{0, 1, 2, \dots\}$ iff $s \models \phi$.

Proof sketch. The only-if direction follows from Theorem 10. Conversely, let k be at least the number of atoms in s and ϕ . Then we can split them all, which corresponds to testing all truth assignments for these atoms. \square

We remark that \models is a slightly restricted version of the semantics in [20] to ease the presentation. The main cost of our simplification is that we lose eventual completeness for formulas $\forall \vec{x} \phi$ where ϕ is quantifier-free.

4.2 Complete but unsound semantics

Next, we turn to the complete but unsound semantics \models . In the complete semantics, it is often more intuitive to consider the task of *disproving* that s satisfies ϕ , that is, $s, l \not\models \phi$ where $l \in \{0, 1, 2, \dots\}$. Here l specifies the reasoning effort similar to the split levels k before. In $s, k \models \phi$ one (roughly) shows that for some atoms, ϕ obviously comes out true in s under any truth assignment of these atoms (where “obvious” means after unit propagation and subsumption). By contrast, the objective for $s, l \not\models \phi$ is to show that s can be augmented with l literals so the resulting setup obviously disproves ϕ .

In particular, this requires to detect whether the setup *might* be inconsistent, because only a consistent setup can disprove ϕ . For that, we use a very simple heuristic: whenever the setup mentions some literal both positively and negatively after removing all subsumed clauses, it is deemed possibly-inconsistent. While this heuristic is of course not sophisticated, the idea is to compensate for its naivete by increasing l , that is, by more reasoning effort.

Definition 12 We write $\text{XP}(s)$ to close the set of all literals that occur in $\text{UP}^-(s)$ under unit propagation, $\text{gnd}(c)$ for the set of ground instances of c , and $s \otimes \ell$ to augment s with all ground instances of $[\ell]$ which are not obviously inconsistent with s :

$$\begin{aligned} \text{XP}(s) &= \text{UP}(\{[\ell] \mid \ell \in c \text{ for some } c \in \text{UP}^-(s)\}); \\ \text{gnd}(c) &= \{c_{\vec{n}}^{\vec{x}} \mid \vec{x} \text{ are the free variables of } c, n_i \in \mathcal{N}\}; \\ s \otimes \ell &= s \cup \{[\ell_{\vec{n}}^{\vec{x}}] \in \text{gnd}([\ell]) \mid [\bar{\ell}_{\vec{n}}^{\vec{x}}] \notin \text{UP}^+(s)\}. \end{aligned}$$

The rationale behind $s \otimes \ell$ is that often a setup may contain infinitely many instances of some clause, and we want to trigger unit propagation for all of them. For example, when $s = \{[\neg P(\#1)], [\neg P(n), Q(n)] \mid n \in \mathcal{N}\}$, then $s \otimes P(x)$ augments s with the instances $P(n)$ for all $n \neq \#1$. With unit propagation we can then infer $Q(n)$ for all $n \neq \#1$, but we avoid the empty clause.

$\text{XP}(s)$ simply serves our simple heuristic to check whether a setup might be inconsistent: it takes all literals from $\text{UP}^-(s)$ and closes them under unit resolution. The next lemma is therefore no surprise.

Lemma 13 *If $[\] \notin \text{XP}(s)$, then for some w , for all $c \in s$, $w \models c$.*

Proof. Let $[\] \notin \text{XP}(s)$ and let $w \models \ell$ iff $[\ell] \in \text{XP}(s)$, which exists as $[\ell] \notin \text{XP}(s)$ or $[\bar{\ell}] \notin \text{XP}(s)$. By subsumption, $w \models c$ for all $c \in \text{UP}^-(s)$. By Lemma 8, $w \models c$ for all $c \in s$. \square

For a setup s , effort $l \in \{0, 1, 2, \dots\}$, and an objective sentence ϕ , the complete satisfaction relation $s, l \models \phi$ is defined inductively:

1. $s, l + 1 \models \phi$ iff $s \otimes \ell, l \models \phi$ for all literals ℓ ;
2. if c is a clause:
 $s, 0 \models \neg c$ iff $[\] \in \text{XP}(s)$ or $c \notin \text{UP}^+(s)$;
3. $s, 0 \models (\phi \vee \psi)$ iff $s, 0 \models \phi$ or $s, 0 \models \psi$;
4. if $(\phi \vee \psi)$ is not a clause:
 $s, 0 \models \neg(\phi \vee \psi)$ iff $s, 0 \models \neg\phi$ and $s, 0 \models \neg\psi$;
5. $s, 0 \models \neg\neg\phi$ iff $s, 0 \models \phi$;
6. $s, 0 \models \exists x \phi$ iff $s, 0 \models \phi_n^x$ for some $n \in \mathcal{N}$;
7. $s, 0 \models \neg\exists x \phi$ iff $s, 0 \models \neg\phi_n^x$ for all $n \in \mathcal{N}$.

The main differences between \models and \models are Rules 1 and 2. It may be more intuitive to read the definition of \models from the perspective disproving. According to Rule 1 $s, l + 1 \not\models \phi$ means that we may pick some literal ℓ and show $s \otimes \ell, l \not\models \phi$. And Rule 2 says that $s, 0 \models \neg c$ when the setup is certainly consistent (since $[\] \notin \text{XP}(s)$) but satisfies c (since $c \in \text{UP}(s)$). We illustrate this with our example.

Example 14 Consider s_1 from Example 9, and let us see whether it is consistent with I , that is, $s_1, l \not\models \neg I$ for certain l . For $l = 0$, note that $\text{UP}^-(s_1)$ mentions I and $\neg I$ in clauses, so $[\] \in \text{XP}(s_1)$, and thus $s_1, 0 \models \neg I$. For $l \geq 1$, however, we are allowed to add some literal to s_1 so as to build a countermodel that clearly disproves $\neg I$. Indeed, adding, for example, $[\neg A]$ does the job: $\text{UP}^-(s_1 \otimes \neg A) = \{[\neg A], [I]\} \cup \{[M(\text{root})], [\neg E(\text{root}), \neg V], [\neg M(n), \neg E(n), \neg V] \mid n \in \mathcal{N} \setminus \{\text{root}\}\} \cup \text{EQ}$ is obviously consistent, that is, $[\] \notin \text{XP}(s_1)$, and moreover $[I] \in \text{UP}^+(s_1 \otimes \neg A)$. So by adding $\neg A$, we have shown that s_1 can falsify $\neg I$. Thus, $s_1, l \not\models \neg I$ iff $l \geq 1$.

The following theorem is the completeness result for \models .

Theorem 15 *If $s \models \phi$, then $s, l \not\models \phi$.*

Proof. By contraposition and by induction on l . For the base case let $l = 0$ and suppose $s, 0 \not\models \phi$. Then clearly $[\] \notin \text{XP}(s)$, so by Lemma 13, there is a w such that $w \models c$ for all $c \in s$ (*). We show that $w \not\models \phi$ by subinduction on $|\phi|$. For any negated clause, $s, 0 \not\models \neg c$ iff $[\] \notin \text{XP}(s)$ and $c \in \text{UP}^+(s)$ only if (by (*) and Lemma 8) $w \models c$ iff $w \not\models \neg c$. Very similarly, for any literal, $s, 0 \not\models \ell$ iff $s, 0 \not\models \neg \bar{\ell}$ only if (by the same argument as for negated clauses) $w \not\models \bar{\ell}$ iff $w \not\models \ell$. We omit the other cases; they are straightforward.

For the main induction step suppose the lemma holds for l and that $s, l + 1 \not\models \phi$. Then $s \otimes \ell, l \not\models \phi$ for some ℓ . By induction, $s \otimes \ell \not\models \phi$. By monotonicity, $s \not\models \phi$. \square

Just like \models is eventually complete for the propositional case, \models is eventually sound, that is, all invalid inferences can be detected for large enough l .

Theorem 16 Let s be finite and ϕ be propositional. Then $s, l \not\models \phi$ for some $l \in \{0, 1, 2, \dots\}$ iff $s \not\models \phi$.

Proof sketch. The only-if direction follows from Theorem 15. Conversely, when l is at least the number of atoms in s and ϕ , they can all be set to the same value as in a world that satisfies s but not ϕ . \square

5 LIMITED CONDITIONAL BELIEF IN \mathcal{BCO}

We are finally ready for \mathcal{BCO} , the logic of limited conditional belief. The language is the same as for \mathcal{BO} , except that the belief operators \mathbf{B}_k^l and \mathbf{O}_k^l are now decorated with $k, l \in \{0, 1, 2, \dots\}$ to indicate the reasoning effort, and for simplicity we disallow predicates outside of belief modalities.

Definition 17 A *limited epistemic state* \vec{s} is an infinite sequence of setups s_1, s_2, \dots such that $\text{UP}^+(s_1) \supseteq \text{UP}^+(s_2) \supseteq \dots$ and for some $p \in \{1, 2, \dots\}$, $\text{UP}^+(s_p) = \text{UP}^+(s_{p+1}) = \dots$

The idea behind limited epistemic states is the same as for unlimited epistemic states, except that in the limited case every sphere is represented as a setup instead of a set of worlds.

Recall that the plausibility of a formula is the index of the first sphere consistent with that formula. With the limited satisfaction relations from the previous section, we can approximate this notion of plausibility from below and above:

$$\begin{aligned} [\vec{s}, k \dot{\phi}] &= \min\{p \mid p = \infty \text{ or } s_p, k \not\models \neg\phi\}; \\ [\vec{s}, l \phi] &= \min\{p \mid p = \infty \text{ or } s_p, l \not\models \neg\phi\}. \end{aligned}$$

It is easy to see that increasing the effort does not impair the quality of these the approximations, as stated in the next lemma.

Lemma 18 $[\vec{s}, k \dot{\phi}] \leq [\vec{s}, k + 1 \dot{\phi}] \leq [\vec{s}, l + 1 \phi] \leq [\vec{s}, l \phi]$.

Proof. For the first inequality, it is easy to see that $s_p, k + 1 \not\models \neg\phi$ implies $s_p, k \not\models \neg\phi$. Similarly for the third inequality, $s_p, l \not\models \neg\phi$ implies $s_p, l + 1 \not\models \neg\phi$. For the remaining one, if $s_p, l + 1 \not\models \neg\phi$, then $s_p \not\models \neg\phi$ by Theorem 15, and so $s, k + 1 \not\models \neg\phi$ by Theorem 10. \square

These approximations are key to the semantics of \mathcal{BCO} . Recall that in \mathcal{BO} the semantics of conditional belief is that the plausibility of the antecedent denotes the sphere in which the material implication of antecedent and consequent should be evaluated. And for only-believing in \mathcal{BO} the p th sphere of the epistemic state is determined by those conditionals whose antecedent have a plausibility $\geq p$.

In limited reasoning, we only have the approximate plausibilities. For limited conditional belief $\mathbf{B}_k^l(\phi \Rightarrow \psi)$ the idea is therefore to approximate the plausibility of the ϕ from above. And for limited only-believing $\mathbf{O}_k^l\{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$ we shall build up the corresponding limited epistemic state only as long as the approximations from below and above are consistent: we say a limited epistemic state \vec{s} is *l_k -plausibility-consistent at $p \in \{1, 2, \dots\}$* iff for all $i \in \{1, \dots, m\}$, $[\vec{s}, k \dot{\phi}_i] \geq p$ iff $[\vec{s}, l \phi_i] \geq p$.

Moreover, in analogy to how only-believing in \mathcal{BO} maximizes every set of world of the epistemic state, we here need to minimize the setups in order to maximize the agent's non-beliefs. We say a setup s is *minimal wrt $s, k \not\models \phi$* iff $s, k \not\models \phi$ and there is no s' such that $\text{UP}^+(s') \subsetneq \text{UP}^+(s)$ and $s', k \not\models \phi$. Finally, we let $\text{NF}[v]$ stand for the prenex negation normal form of v .

Truth of a sentence α in a limited epistemic state \vec{s} , written $\vec{s} \approx \alpha$, is now defined inductively:

1. $\vec{s} \approx (\alpha \vee \beta)$ iff $\vec{s} \approx \alpha$ or $\vec{s} \approx \beta$;
2. $\vec{s} \approx \neg\alpha$ iff $\vec{s} \not\models \alpha$;

3. $\vec{s} \approx \exists x \alpha$ iff $\vec{s} \approx \alpha_n^x$ for some name n ;
4. $\vec{s} \approx \mathbf{B}_k^l(\phi \Rightarrow \psi)$ iff for all $p \in \{1, 2, \dots\}$, if $p \leq [\vec{s}, l \phi]$, then $s_p, k \not\models (\phi \supset \psi)$;
5. $\vec{s} \approx \mathbf{O}_k^l\{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$ iff for some limited epistemic state \vec{s}' , and for all $p \in \{1, 2, \dots\}$,
 - s'_p is minimal wrt $s'_p, 0 \not\models \bigwedge_{i: [\vec{s}', k \dot{\phi}_i] \geq p} \text{NF}[(\phi_i \supset \psi_i)]$;
 - $s_p = \begin{cases} s'_p & \text{if } \vec{s}' \text{ is } l_k\text{-plausibility-consistent at } 1, \dots, p; \\ s'_p \circ & \text{otherwise;} \end{cases}$ where p° is such that $\text{UP}^+(s_{p^\circ}) = \text{UP}^+(s_{p'})$ for all $p' \geq p^\circ$.

As usual and as in \mathcal{BO} , the symbol \approx is also used to denote entailment.

Rule 4 approximates the plausibility of ϕ from *above*, which avoids too-plausible spheres inconsistent with ϕ , and then applies sound inference. That way, \mathbf{B}_k^l is a conservative variant of \mathcal{BO} 's conditional belief operator.

The same spirit is behind Rule 5. The intuition is to build up the system of spheres as long as the lower and upper bound of all plausibilities are consistent. Once they are not, it is unclear how the next sphere should look like, so we skip to the “last” one, $s'_p \circ$. That last sphere is determined by conditionals which (mutually) contradict their premises, so there is no scenario where any of them could be true. The parameters k and l determine how much effort is put into checking the plausibility-consistency. Note that there may be conditionals $\phi_i \Rightarrow \psi_i$ with unsatisfiable antecedents which do not occur in the last sphere. This is because we only take those conditionals whose antecedents can be proved unsatisfiable by sound reasoning (with effort k). If we used complete instead of sound reasoning here, the outermost sphere could end up being too strong and then yield false beliefs. Figure 1 illustrates such approximations.

Before we illustrate how the semantics works, we show a unique-model property for a certain type of knowledge base.

5.1 Proper⁺ knowledge bases

The class of knowledge bases we are chiefly interested in is called proper⁺ [18]. Essentially, it requires clausal form and disallows existential quantifiers.

Definition 19 A sentence π is *proper⁺* when π is of the form $\bigwedge_j \forall \vec{x} c_j$ for clauses c_j . Then we let $\text{gnd}(\pi) = \bigcup_j \text{gnd}(c_j)$. A set of conditionals $\Gamma = \{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$ is proper⁺ when $\bigwedge_{1 \leq i \leq m} \text{NF}[(\phi_i \supset \psi_i)]$ is proper⁺.

Bringing $(\phi_i \supset \psi_i)$ into prenex negation normal form has the benefit many conditionals are proper⁺ which otherwise wouldn't. For example, $(P \wedge Q \supset R)$, which is just an abbreviation for $\neg(\neg P \vee \neg Q) \vee R$ is not proper⁺, but eliminating the double negation does the job already. Incidentally, this is also the reason why NF occurs in Rule 5.

For the remainder of this paper we let π and Γ be proper⁺. Proper⁺ knowledge bases have been shown to have attractive properties for limited belief [26, 19, 20, 21], and as we shall see these qualities also hold for conditional belief. Above all, the unique-model property from Theorem 4 carries over to limited belief.

Theorem 20 *There is a unique (modulo UP^+) limited epistemic state \vec{s} such that $\vec{s} \approx \mathbf{O}_k^l \Gamma$, that is, for all \vec{s}' such that $\vec{s}' \approx \mathbf{O}_k^l \Gamma$ and for all $p \in \{1, 2, \dots\}$, $\text{UP}^+(s_p) = \text{UP}^+(s'_p)$.*

Proof sketch. The crucial lemma is that for every proper⁺ π , s is minimal wrt $s, 0 \not\models \pi$ iff $\text{UP}^+(s) = \text{UP}^+(\text{gnd}(\pi))$, proven in [19].

With that result, the theorem can be shown by the same argument used to prove the unique-model property in \mathcal{BO} [34]. \square

Finally, here is the kangaroo example with limited belief.

Example 21 Note that Γ from Example 2 is proper⁺. Let $k = 1$ and $l = 1$, and let s_1 and s_μ be as in Example 9. Then s_1 is the first sphere of $\vec{s} \approx \mathbf{O}_k^l \Gamma$. To determine the next sphere, we first need to see whether \vec{s} is k -plausibility-consistent at 2, that is, $[\vec{s}, k \dot{\phi}] \geq 2$ iff $[\vec{s}, l \dot{\phi}] \geq 2$ for all $\phi \Rightarrow \psi \in \Gamma$. We can reuse our results from Examples 9 and 14. For example, we have shown in Example 14 that $s_1, l \not\approx \neg I$, so we have $[\vec{s}, l \dot{I}] = 1$. Similarly, in Example 9 we have shown $s_1, k \not\approx \neg A$, so $[\vec{s}, k \dot{A}] \geq 2$. That way and with Lemma 18, we obtain

- $[\vec{s}, k \dot{I}] = 1$ and $[\vec{s}, l \dot{I}] = 1$;
- $[\vec{s}, k \dot{A}] \geq 2$ and $[\vec{s}, l \dot{A}] \geq 2$;
- $[\vec{s}, k \dot{\top}] = 1$ and $[\vec{s}, l \dot{\top}] = 1$;
- $[\vec{s}, k \dot{\neg I}] \geq 2$ and $[\vec{s}, l \dot{\neg I}] \geq 2$.

The plausibilities of the last two conditionals in Example 2 are omitted, as they are vacuously ∞ . Hence, \vec{s} is k -plausibility-consistent at 2. The conditionals with plausibility ≥ 2 determine the second sphere, so we obtain

$$\text{UP}^+(s_2) = \text{UP}^+(\{[\neg A, \neg I], [\neg A, E(\text{roo})], [I, A]\} \cup s_\mu).$$

It is easy to see that $[\vec{s}, k \dot{A}] = [\vec{s}, k \dot{\neg I}] = 2$. Moreover $[\vec{s}, l \dot{A}] = [\vec{s}, l \dot{\neg I}] = 2$ can be shown by adding A to the setup. So for the final sphere s_3 we have

$$\text{UP}^+(s_3) = \text{UP}^+(s_\mu).$$

By Theorem 20, \vec{s} is the unique model of $\mathbf{O}_k^l \Gamma$, so we can now prove $\mathbf{O}_k^l \Gamma \approx \mathbf{B}_{k'}^{l'}(\neg I \Rightarrow \neg V)$ for $k' = 1, l' = 1$: since $[\vec{s}, l' \dot{\neg I}] = 2$, we only need to show $s_2, k' \not\approx I \vee \neg V$, which is easy by splitting I .

Note that for $k = 0$ or $l = 0$, the model of $\mathbf{O}_k^l \Gamma$ would have consisted of s_1 followed immediately by s_μ , because of k -plausibility-inconsistency at 2. In this case, no k' or l' would have been large enough to show $\mathbf{B}_{k'}^{l'}(\phi \Rightarrow \psi)$.

In this example, we let $k = l = k' = l' = 1$. It is easy to see that the entailment in fact holds for arbitrary $k \geq 1, l \geq 1, k' \geq 1, l' \geq 1$. It is no surprise that increasing the effort retains the beliefs in general, that is, effort behaves monotonically.

Theorem 22 Suppose $\mathbf{O}_k^l \Gamma \approx \mathbf{B}_{k'}^{l'}(\phi \Rightarrow \psi)$.

Then $\mathbf{O}_{\tilde{k}}^{\tilde{l}} \Gamma \approx \mathbf{B}_{\tilde{k}'}^{\tilde{l}'}(\phi \Rightarrow \psi)$ for all $\tilde{k} \geq k, \tilde{l} \geq l, \tilde{k}' \geq k', \tilde{l}' \geq l'$.

Proof sketch. Suppose $\vec{e} \models \mathbf{O}\Gamma$, $\vec{s} \approx \mathbf{O}_k^l \Gamma$, and $\vec{s}' \approx \mathbf{O}_{\tilde{k}}^{\tilde{l}} \Gamma$. The key argument is that \vec{s}' is at least as faithful to \vec{e} as \vec{s} is (cf. Figure 1). This is proven by inductions on \tilde{k} and \tilde{l} using Lemma 18. It is then easy to see that \vec{s}' entails at least the beliefs that \vec{s} does. Again using Lemma 18, we can then show that beliefs proved with effort k', l' can also be proved for \tilde{k}', \tilde{l}' . \square

More important perhaps is the question whether \mathcal{BOC} is sound wrt its archetype \mathcal{BO} . Indeed this is the case for belief implications with proper⁺ knowledge bases, as expressed by the following theorem.

Theorem 23 If $\mathbf{O}_k^l \Gamma \approx \mathbf{B}_{k'}^{l'}(\phi \Rightarrow \psi)$, then $\mathbf{O}\Gamma \models \mathbf{B}(\phi \Rightarrow \psi)$.

Proof sketch. By Theorem 20, there is a unique (modulo UP^+) \vec{s} such that $\vec{s} \approx \mathbf{O}_k^l \Gamma$. All spheres but its last one faithfully match the corresponding spheres of the unique \vec{e} such that $\vec{e} \models \mathbf{O}\Gamma$, and the final sphere of \vec{s} is weaker than the last sphere of \vec{e} (cf. Figure 1), so everything that can be inferred from a sphere of \vec{s} by sound inference

can also be inferred from \vec{e} . It is then easy to show that $[\vec{e}' \dot{\phi}] \leq [\vec{s}, l' \dot{\phi}]$. Since sound inference is applied to prove $(\phi \supset \psi)$, the claim follows. \square

5.2 Decidability of belief implications

Finally, we investigate computational questions of belief implications $\mathbf{O}_k^l \Gamma \approx \mathbf{B}_{k'}^{l'}(\phi \Rightarrow \psi)$. We shall see that the problem is decidable for proper⁺ knowledge bases, and in the propositional case even tractable for fixed effort.

The fundamental idea behind the decidability result is that standard names that do not occur in the knowledge base or query cannot be distinguished. Hence we only need to consider a finite number of them: those from the knowledge base and query, plus a few more (their number is bounded by the number of quantifiers and maximum arity in the knowledge base and query). We first present decision procedures S and C for $\not\approx$ and \approx , and finally the procedure B for belief implications.

Again we let π and Γ be proper⁺.

Definition 24 We let $N(\pi, \phi, j)$ contain all names that occur in the formulas π or ϕ plus $(j + 1) \cdot \max\{|\pi|_w, |\phi|_w\}$ additional names, where $|v|_w$ is the maximum of the largest number of free variables in any subformula of v and the highest arity in v . For any set of names N , we let $\text{gnd}_N(\phi)$ and $s \otimes_N \ell$ be as $\text{gnd}(\phi)$ and $s \otimes \ell$ except that the grounding is restricted to the names in N .

As sketched above, to decide $\text{gnd}(\pi), k \not\approx \phi$ it suffices to consider only names from $N(\pi, \phi, k)$ for grounding, quantification, and splitting, and it is moreover easy to see that only literals whose symbols occur in π or ϕ need to be considered. These ideas lead to the procedure $S[N, s, k, \phi] \in \{0, 1\}$ with the following inductive definition:

- $S[N, s, k + 1, \phi] = 1$ iff $S[N, s \cup \{\ell\}, k, \phi] = S[N, s \cup \{\tilde{\ell}\}, k, \phi] = 1$ for some ground literal ℓ whose symbol occurs in s or ϕ and whose names are from N ;
- if c is a clause: $S[N, s, 0, c] = 1$ iff $c \in \text{UP}^+(s)$;
- if $(\phi \vee \psi)$ is not a clause: $S[N, s, 0, (\phi \vee \psi)] = \max\{S[N, s, 0, \phi], S[N, s, 0, \psi]\}$;
- $S[N, s, 0, \neg(\phi \vee \psi)] = \min\{S[N, s, 0, \neg\phi], S[N, s, 0, \neg\psi]\}$;
- $S[N, s, 0, \neg\neg\phi] = S[N, s, 0, \phi]$;
- $S[N, s, 0, \exists x \phi] = \max\{S[N, s, 0, \phi_n^x] \mid n \in N\}$;
- $S[N, s, 0, \neg\exists x \phi] = \min\{S[N, s, 0, \neg\phi_n^x] \mid n \in N\}$.

The following theorem says that S is a decision procedure for $\text{gnd}(\pi), k \not\approx \phi$.

Theorem 25 $\text{gnd}(\pi), k \not\approx \phi$ iff $S[N, \text{gnd}_N(\pi), k, \phi] = 1$ where $N = N(\pi, \phi, k)$.

Proof sketch. The key tool to show the theorem are bijections between standard names that leave the names from π and ϕ unchanged but possibly swap any other names. For any clause c that mentions no more than $\max\{|\pi|_w, |\phi|_w\}$ names that do not occur in π or ϕ , and a bijection $*$ that swaps these names with the additional names in N , it can be shown that $c \in \text{UP}^+(\text{gnd}(\pi))$ iff $c^* \in \text{UP}^+(\text{gnd}_N(\pi))$; this basically allows us to restrict grounding of π to N . Similarly, it can be shown that quantification and splitting can be restricted to names from N . Finally it is intuitively immediate that splitting only literals whose symbols occur in π or ϕ can generate new inferences. \square

Corollary 26 $\text{gnd}(\pi), k \models^{\otimes} \phi$ is decidable. In the propositional case, the time complexity is $\mathcal{O}((|\pi| + k)^{k+1} \cdot |\phi|^{k+1} \cdot 2^k)$.

In a very similar fashion we can design a decision procedure for $\text{gnd}(\pi), l \models^{\otimes} \phi$. For analogous reasons as in \models^{\otimes} , it suffices to consider only names from $N(\pi, \phi, l)$ and to augment the setup only with literals whose symbols occur in π or ϕ . The resulting procedure $C[N, s, l, \phi] \in \{0, 1\}$ is defined inductively as follows:

- $C[N, s, l+1, \phi] = 1$ iff $C[N, s \otimes_N \ell, l, \phi] = 1$ for all (including non-ground) literals ℓ whose symbols occur in s or ϕ and whose names are from N ;
- if ℓ is a positive literal:
 $C[N, s, 0, \ell] = C[N, s, 0, \neg\bar{\ell}]$;
- if c is a clause:
 $C[N, s, 0, \neg c] = 1$ iff $[\] \in \text{XP}(s)$ or $c \notin \text{UP}^+(s)$;
- $C[N, s, 0, (\phi \vee \psi)] = \max\{C[N, s, 0, \phi], C[N, s, 0, \psi]\}$;
- if $(\phi \vee \psi)$ is not a clause:
 $C[N, s, 0, \neg(\phi \vee \psi)] = \min\{C[N, s, 0, \neg\phi], C[N, s, 0, \neg\psi]\}$;
- if $\neg\phi$ is not a clause:
 $C[N, s, 0, \neg\neg\phi] = C[N, s, 0, \phi]$;
- $C[N, s, 0, \exists x \phi] = \max\{C[N, s, 0, \phi_n^x] \mid n \in N\}$;
- $C[N, s, 0, \neg\exists x \phi] = \min\{C[N, s, 0, \neg\phi_n^x] \mid n \in N\}$.

For similar reasons as for S and \models^{\otimes} , we can show that C is a decision procedure for $\text{gnd}(\pi), l \models^{\otimes} \phi$.

Theorem 27 $\text{gnd}(\pi), l \models^{\otimes} \phi$ iff $C[N, \text{gnd}_N(\pi), l, \phi] = 1$ where $N = N(\pi, \phi, l)$.

Corollary 28 $\text{gnd}(\pi), l \models^{\otimes} \phi$ are decidable. In the propositional case, the time complexity is $\mathcal{O}((|\pi| + l)^{l+1} \cdot |\phi|^{l+1})$.

So far we have established that reasoning in \models^{\otimes} and \models^{\otimes} is decidable for proper^+ knowledge bases, and that it is tractable for given effort in the propositional case.

With these decision procedures for the objective limited semantics, it is easy to translate BOC 's semantic rules for conditional belief and only-believing to a decision procedure for limited belief implications $\mathbf{O}_k^l \Gamma \approx \mathbf{B}_{k'}^l(\phi \Rightarrow \psi)$ for $\text{proper}^+ \Gamma = \{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$. The steps of the procedure $\mathbf{B}[N, k, l, k', l', \Gamma, \phi, \psi] \in \{0, 1\}$ are as follows:

- let s'_1, \dots, s'_{m+1} be such that
 $s'_p = \text{gnd}_N(\bigwedge_{i:p=1 \text{ or } S[N, s'_{p-1}, k, \neg\phi_i]=1} \text{NF}[(\phi_i \supset \psi_i)])$;
- let $p^* = \max\{p \in \{1, \dots, m\} \mid \text{for some } i \in \{1, \dots, m\}, \max\{S[N, s'_{p'}, k, \neg\phi_i] \mid p' < p\} = \max\{C[N, s'_{p'}, l, \neg\phi_i] \mid p' < p\}\}$;
- let $(s_1, \dots, s_{p^*}, s_{p^*+1}) = (s'_1, \dots, s'_{p^*}, s'_{m+1})$;
- let $p^* = \min\{p \mid C[N, l', s_p, \neg\phi] = 0 \text{ or } p = p^* + 1\}$;
- return $S[N, k', s_{p^*}, (\phi \supset \psi)]$.

Theorem 29 $\mathbf{O}_k^l \Gamma \approx \mathbf{B}_{k'}^l(\phi \Rightarrow \psi)$ iff $\mathbf{B}[N, k, l, k', l', \Gamma, \phi, \psi] = 1$ where $N = N(\bigwedge_i \text{NF}[(\phi_i \supset \psi_i)], (\phi \supset \psi), \max\{k, l, k', l'\})$.

Proof sketch. The setups s'_1, \dots, s'_{m+1} in \mathbf{B} correspond to \vec{s}' in Rule 5 of BOC 's semantics, for it is sufficient to consider only $m+1$ many setups, since the number of different setups in \vec{s}' can be shown to be bounded by $m+1$. Then s'_{p^*} is the last sphere that is s'_k -plausibility-consistent wrt Γ , and therefore s_1, \dots, s_{p^*+1} in \mathbf{B} correspond to \vec{s} in Rule 5. Finally p^* is denotes the approximated plausibility of ϕ , and the last line evaluates $(\phi \supset \psi)$ in that sphere. \square

Corollary 30 $\mathbf{O}_k^l \Gamma \approx \mathbf{B}_{k'}^l(\phi \Rightarrow \psi)$ is decidable. For propositional Γ and ϕ , the time complexity is $\mathcal{O}(m^2 \cdot (\|\Gamma\| + j)^{2 \cdot (j+1)} + m \cdot (\|\Gamma\| + j')^{j+1} \cdot |(\phi \supset \psi)|^{j'+1})$, where $j = \max\{k, l\}$, $j' = \max\{k', l'\}$, and $\|\Gamma\| = |\bigwedge_i (\phi_i \supset \psi_i)|$.

Note that the time complexity is exponential only in the effort parameters. For fixed effort, propositional limited belief is hence tractable.

6 CONCLUSION

This paper introduces a logic of limited conditional belief. It is shown that reasoning in proper^+ knowledge bases is decidable, and even tractable in the propositional case. This is achieved by limiting the effort to be spend on the reasoning task, thereby avoiding logical omniscience while retaining the first-order expressivity in the query. Generalizing LLL's framework of limited reasoning to conditional belief turned out to be surprisingly complicated. This is chiefly due to the prominent role of plausibilities in conditional belief.

Semantically a conditional knowledge base can be uniquely represented by a system of spheres; inference then boils down to model-checking. However, as it is undecidable in general which system of spheres corresponds to the knowledge base, in limited reasoning the best we can do is work with an approximation of the system of spheres. By approximating the plausibilities of formulas from below and above, we came up with an approximative system whose first spheres faithfully represent the unlimited spheres, and whose last sphere conservatively approximates the outermost sphere of the unlimited system. Given such an approximative system of spheres for a conditional knowledge base, a limited conditional belief is evaluated by approximating the plausibility of its antecedent from above to select a sphere, and then applying sound inference in that sphere.

We see several interesting avenues of future work. First and foremost, we are currently working on an implementation of the reasoning service described here, enriched with functions in the spirit of [21]. The treatment of functions from [21] seamlessly carries over to our logic; we skipped it here for simplicity. What makes functions very attractive is, among other things, that they allow to express existentials in the knowledge base by way of Skolemization. With that implementation, we plan to explore the practical utility of limited reasoning (for example, for high-level control in robots) in general and conditional belief in particular. Then the question will arise which effort parameters to choose. A simple approach may be to iteratively increase the effort with a situation-dependent timeout after which the search procedure is aborted. Even when the expected result could not be proved within the timeout, the tried effort parameters will give the system designer insight about why his program or robot behaved the way it did.

We also plan to investigate notions of limited belief revision. The problem with many belief revision operators is that they bring along exponential growth in the number of iterated revisions. We hope to alleviate this with a limited revision operator where the revised epistemic state is an approximative structure, similar as in limited only-believing. Such a notion of limited revision could be used to devise a limited variant of the situation calculus in the style of [20] but extended to defeasible beliefs and imperfect sensing [35, 36].

ACKNOWLEDGEMENTS

This work was supported by the German National Science Foundation (DFG) research unit FOR 1513.

References

- [1] Alan R. Anderson and Nuel D. Belnap, *Entailment: The Logic of Relevance and Necessity*, Princeton University Press, 1975.
- [2] Franz Baader, Ralf Küsters, and Frank Wolter, 'Extensions to description logics', in *The Description Logic Handbook: Theory, Implementation, and Applications*, eds., Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter Patel-Schneider, pp. 219–261, (2003).
- [3] Nuel D. Belnap, 'A useful four-valued logic', in *Modern uses of multiple-valued logic*, 5–37, Springer, (1977).
- [4] Egon Börger, Erich Grädel, and Yuri Gurevich, *The Classical Decision Problem*, Springer, 1997.
- [5] James P. Delgrande, 'A framework for logics of explicit belief', *Computational Intelligence*, **11**(1), 47–88, (1995).
- [6] Ronald Fagin and Joseph Y. Halpern, 'Belief, awareness, and limited reasoning', *Artificial Intelligence*, **34**(1), 39–76, (1987).
- [7] Marcelo Finger and Renata Wassermann, 'Anytime approximations of classical logic from above', *Journal of Logic and Computation*, **17**(1), 53–82, (2007).
- [8] Erich Grädel, Phokion G. Kolaitis, and Moshe Y. Vardi, 'On the decision problem for two-variable first-order logic', *Bulletin of symbolic logic*, **3**(01), 53–69, (1997).
- [9] Adam Grove, 'Two modellings for theory change', *Journal of Philosophical Logic*, **17**(2), 157–170, (1988).
- [10] Jaakko Hintikka, *Knowledge and Belief: An Introduction to the Logic of the Two Notions*, Cornell University Press, 1962.
- [11] Jaakko Hintikka, 'Impossible possible worlds vindicated', *Journal of Philosophical Logic*, **4**(4), 475–484, (1975).
- [12] Hirofumi Katsuno and Alberto O. Mendelzon, 'Propositional knowledge base revision and minimal change', *Artificial Intelligence*, **52**(3), 263–294, (1991).
- [13] Gabriele Kern-Isberner, *Conditionals in Nonmonotonic Reasoning and Belief Revision*, Springer, 2001.
- [14] Toryn Q. Klassen, Sheila A. McIlraith, and Hector J. Levesque, 'Towards tractable inference for resource-bounded agents', in *Proceedings of the Twelfth International Symposium on Logical Formalizations of Commonsense Reasoning (Commonsense)*, (2015).
- [15] Kurt Konolige, 'A deduction model of belief', *Research notes in Artificial Intelligence*, (1986).
- [16] Gerhard Lakemeyer, 'Limited reasoning in first-order knowledge bases', *Artificial Intelligence*, **71**(2), 213–255, (1994).
- [17] Gerhard Lakemeyer, 'Limited reasoning in first-order knowledge bases with full introspection', *Artificial Intelligence*, **84**(1), 209–255, (1996).
- [18] Gerhard Lakemeyer and Hector J. Levesque, 'Evaluation-based reasoning with disjunctive information in first-order knowledge bases', in *Proceedings of the Eighth Conference on Principles of Knowledge Representation and Reasoning (KR)*, pp. 73–81, (2002).
- [19] Gerhard Lakemeyer and Hector J. Levesque, 'Decidable reasoning in a logic of limited belief with introspection and unknown individuals', in *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI)*, (2013).
- [20] Gerhard Lakemeyer and Hector J. Levesque, 'Decidable reasoning in a fragment of the epistemic situation calculus', in *Proceedings of the Fourteenth International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pp. 468–477, (2014).
- [21] Gerhard Lakemeyer and Hector J. Levesque, 'Decidable reasoning in a logic of limited belief with function symbols', in *Proceedings of the Fifteenth International Conference on Principles of Knowledge Representation and Reasoning (KR)*, (2016). To appear.
- [22] Hector J. Levesque, 'A logic of implicit and explicit belief', in *Proceedings of the Fourth National Conference on Artificial Intelligence (AAAI)*, pp. 198–202, (1984).
- [23] Hector J. Levesque, 'All I know: a study in autoepistemic logic', *Artificial Intelligence*, **42**(2), 263–309, (1990).
- [24] Hector J. Levesque and Gerhard Lakemeyer, *The Logic of Knowledge Bases*, MIT Press, 2001.
- [25] David Lewis, *Counterfactuals*, John Wiley & Sons, 1973.
- [26] Yongmei Liu, Gerhard Lakemeyer, and Hector J. Levesque, 'A logic of limited belief for reasoning with disjunctive information', in *Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR)*, (2004).
- [27] Christopher Lynch, 'Unsound theorem proving', in *Proceedings of the Workshop on Disproving: Non-Theorems, Non-Validity, Non-Provability*, pp. 1–12, (2004).
- [28] Michael Mortimer, 'On languages with two variables', *Mathematical Logic Quarterly*, **21**(1), 135–140, (1975).
- [29] Daniele Nardi and Ronald J. Brachman, 'An introduction to description logics', in *The Description Logic Handbook: Theory, Implementation, and Applications*, eds., Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter Patel-Schneider, pp. 5–44, (2003).
- [30] Peter F. Patel-Schneider, 'A decidable first-order logic for knowledge representation', *Journal of automated reasoning*, **6**(4), 361–388, (1990).
- [31] Judea Pearl, 'System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning', in *Proceedings of the Third Conference on Theoretical Aspects of Reasoning about Knowledge (TARK)*, pp. 121–135, (1990).
- [32] Marco Schaerf and Marco Cadoli, 'Tractable reasoning via approximation', *Artificial Intelligence*, **74**(2), 249–310, (1995).
- [33] Christoph Schwering, *Conditional Beliefs in Action*, Ph.D. dissertation, RWTH Aachen University, 2016.
- [34] Christoph Schwering and Gerhard Lakemeyer, 'A semantic account of iterated belief revision in the situation calculus', in *Proceedings of the Twenty-First European Conference on Artificial Intelligence (ECAI)*, pp. 801–806, (2014).
- [35] Christoph Schwering, Gerhard Lakemeyer, and Maurice Pagnucco, 'Belief revision and progression of knowledge bases in the epistemic situation calculus', in *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 3214–3220, (2015).
- [36] Christoph Schwering, Gerhard Lakemeyer, and Maurice Pagnucco, 'Belief revision and projection in the epistemic situation calculus'. Forthcoming, 2016.
- [37] Dana Scott, 'A decision method for validity of sentences in two variables', *Journal of Symbolic Logic*, **27**(377), 74, (1962).
- [38] Moshe Y. Vardi, 'On epistemic logic and logical omniscience', in *Proceedings of the First Conference on Theoretical Aspects of Reasoning about Knowledge (TARK)*, pp. 293–305, (1986).