A Semantic Account of Iterated Belief Revision in the Situation Calculus

Christoph Schwering and **Gerhard Lakemeyer**[1]

Abstract. Recently Shapiro et al. explored the notion of iterated belief revision within Reiter's version of the situation calculus. In particular, they consider a notion of belief defined as truth in the most plausible situations. To specify what an agent is willing to believe at different levels of plausibility they make use of so-called belief conditionals, which themselves neither refer to situations or plausibilities explicitly. Reasoning about such belief conditionals turns out to be complex because there may be too many models satisfying them and negative belief conditionals are also needed to obtain the desired conclusions. In this paper we show that, by adopting a notion of only-believing, these problems can be overcome. The work is carried out within a modal variant of the situation calculus with a possible-world semantics which features levels of plausibility. Among other things, we show that only-believing a knowledge base together with belief conditionals always leads to a unique model, which allows characterizing the beliefs of an agent, after any number of revisions, in terms of entailments within the logic.

1 INTRODUCTION

Recently Shapiro, Pagnucco, Lespérance, and Levesque [16] (henceforth SPLL) explored the notion of iterated belief revision within Reiter's version of the situation calculus [13, 14]. SPLL's starting point is Scherl and Levesque's epistemic extension of the situation calculus [15], which formalizes knowledge/belief in terms of truth in all accessible situations. A drawback of this work is that it does not account for belief revision in the sense that new information which conflicts with the current beliefs (for example, through the use of sensors) would invariably lead to an epistemic state where everything is believed, as no accessible situations would be left. In order to remedy this deficiency, SPLL assign plausibilities (taken from the natural numbers) to situations and define a new notion of belief which only considers the most plausible accessible situations. In this framework, new information which conflicts with the current beliefs does not necessarily lead to inconsistency as there may well be other, less plausible situations left which are consistent with the new information. SPLL show that their approach has various desirable properties and they compare their work to the more classical approaches to belief revision and update [1, 7, 2].

In terms of knowledge representation, it seems impractical to having to specify accessible situations and plausibility levels explicitly. For one thing, in the propositional case the number of accessible situations may grow exponentially in the number of fluents, which are propositions whose truth value can be changed by actions. For another, an actual plausibility level of, say, 33 vs. 37 carries little meaning. For that reason, SPLL introduced a new operator $\Rightarrow$, inspired by conditional logic [12], where $\phi \Rightarrow \psi$ intuitively means that in all most plausible situations where ϕ holds, ψ holds as well. We will call such formulas *belief conditionals* from now on. Such statements can be viewed as constraints on the possible plausibility orderings and ideally they allow an agent to derive the right conclusions about how to revise its beliefs based on these constraints and new information acquired by sensing. In the case of SPLL, deriving consequences from such belief conditionals is complicated by the fact that there may be (infinitely) many models with different plausibility orderings satisfying these constraints. Moreover, as we will discuss in detail later, the desired consequences only obtain if negated belief conditionals are considered as well.

Problems such as these can be attributed to the fact that SPLL are not able to state that a knowledge base (KB), possibly including belief conditionals, is *all* the agent believes or that the agent *only-believes* such a KB.[2] In this paper, we propose such an approach to only-believing based on the logic $\mathcal{ES}$ [8, 10], a modal variant of the situation calculus. The approach is semantic in the sense that the language does not refer to either situations or plausibilities. Instead these are only part of the possible-world style semantics of the language. An advantage of our notion of only-believing is that it always leads to a unique model, that is, a set of facts (sentences in first-order logic) together with any number of belief conditionals always leads to a unique epistemic state. As a result of this property, the beliefs of an agent, after any number of revisions, can always be characterized in terms of entailment within the logic. We also show by way of an example taken from SPLL that we are able to obtain the same conclusions without resorting to negated belief conditionals.

The paper is organized as follows. In the next section we introduce the logic $\mathcal{ESB}$, which extends SPLL's ideas to the case of only-believing, and discuss some of its properties. In Section 3 we investigate how $\mathcal{ESB}$ handles belief revision, including a detailed example. Then we discuss related work and conclude.

2 THE LOGIC $\mathcal{ESB}$

$\mathcal{ESB}$ is a first-order modal logic which features, among others, two modal operators $\mathbf{K}$ and $\mathbf{B}$ for knowledge and belief. The $\mathbf{K}$ operator allows to express firm belief, which we simply call knowledge for the sake of distinction. The $\mathbf{B}$ operator allows to express beliefs, which can be be revised when contradicting information is obtained. In contrast to belief, knowledge can only be expanded, but cannot be revised. $\mathcal{ESB}$ is an extension of the variant of $\mathcal{ES}$ presented in [8].

[1] RWTH Aachen University, Germany,
email: (schwering,gerhard)@kbsg.rwth-aachen.de.
Supported by DAAD-Go8 project 56266625 and DFG project FOR 1513.

[2] We remark that SPLL themselves mention only-believing as an interesting open topic of future work.

While an extended version of $\mathcal{ES}$ was proposed recently in [10], we refer to the original logic because it simplifies the presentation.

2.1 The Language

The language $\mathcal{ESB}$ consists of formulas over *fluent predicates* and *rigid terms*. The set of terms is the least set which contains infinitely many variables and constant symbols and is closed under application of infinitely many function symbols. The set of well-formed formulas is the least set such that

- if P is a predicate symbol of arity $k \in \mathbb{N}$[3] and $t_1, \ldots, t_k$ are terms, then $P(t_1, \ldots, t_k)$ is an (atomic) formula;
- if t_1, t_2 are terms, then $(t_1 = t_2)$ is a formula;
- if α and α' are formulas and x is a variable, then $(\alpha \wedge \alpha')$, $\neg\alpha$, $\forall x.\alpha$ are formulas;
- if t is a term and α is a formula, $[t]\alpha$, $\Box\alpha$, and $\mathbf{P}\alpha$ are formulas;
- if α is a formula, then $\mathbf{K}\alpha$ and $\mathbf{B}\alpha$ are formulas;
- if ϕ and ψ are formulas, then $\mathbf{B}(\phi \Rightarrow \psi)$ is a formula;
- if α, ϕ_i, and ψ_i for $1 \le i \le m$, $m \in \mathbb{N}$ are formulas, then $\mathbf{O}(\alpha, \{\phi_1 \Rightarrow \psi_1, \ldots, \phi_m \Rightarrow \psi_m\})$ is a formula.

We read $[t]\alpha$ as "α holds after action t," and $\Box\alpha$ as "α holds after any sequence of actions," and $\mathbf{P}\alpha$ as "α was true before the last action." $\mathbf{K}\alpha$ ($\mathbf{B}\alpha$) is read as "the agent knows (believes) α." Knowledge, as opposed to belief, can only be expanded, but cannot be revised. The belief conditional $\mathbf{B}(\phi \Rightarrow \psi)$ is intended to express that in the most plausible scenarios where ϕ holds, ψ holds as well. For the remainder of this paper, we let Γ stand for $\{\phi_1 \Rightarrow \psi_1, \ldots, \phi_m \Rightarrow \psi_m\}$. $\mathbf{B}\Gamma$ abbreviates $\bigwedge_{i=1}^{m} \mathbf{B}(\phi_i \Rightarrow \psi_i)$. $\mathbf{O}(\alpha, \Gamma)$ captures that all the agent knows about the world is α and possibly other sentences due to the belief conditionals from Γ.

We will use $\vee$, $\exists$, $\supset$, $\equiv$, *False*, and *True* as the usual abbreviations. $\bigwedge_{i:\chi} \alpha_i$ abbreviates $\alpha_{i_1} \wedge \ldots \wedge \alpha_{i_k}$ if $i_1, \ldots, i_k$ are all the indices that satisfy χ. Instead of having different sorts for objects and actions, we lump both sorts together and allow ourselves to use any term as an action or as an object. There are two special predicates, *Poss* for the precondition and *SF* for the binary sensing result of an action.

We call a formula without free variables a *sentence*. A formula with no $\Box$, $[t]$, or $\mathbf{P}$ is called *static*. A formula with no $\mathbf{K}$, $\mathbf{B}$, or $\mathbf{O}$ is called *objective*. A formula with no fluent, $\Box$, or $[t]$ outside the scope of all $\mathbf{K}$, $\mathbf{B}$, and $\mathbf{O}$ is called *subjective*. An objective, static formula without *Poss* and *SF* is called a *fluent* formula.

To simplify the technical treatment and for the purposes of this paper, we assume that the α, ϕ, ψ, ϕ_i, and ψ_i occurring in $\mathbf{B}(\phi \Rightarrow \psi)$ and $\mathbf{O}(\alpha, \Gamma)$ are all objective.

2.2 The Semantics

The truth of an $\mathcal{ESB}$ sentence α after an action sequence z is defined with respect to two things: a world w and an epistemic state f. We write $f, w, z \models \alpha$. A world determines the truth values of all ground atoms after any sequence of actions. An epistemic state contains the possible worlds at each plausibility level.

More precisely, a world is a function from the set of ground atoms and the set of action sequences to $\{0, 1\}$. An epistemic state f is a function from $\mathbb{N}$ to the power set of the set of worlds, that is, for each plausibility level $p \in \mathbb{N}$, for any world w, $w \in f(p)$ means that w is considered possible at plausibility level p. A smaller plausibility value p indicates that the world is more plausible.

[3] As usual, $\mathbb{N}$ denotes the natural numbers including 0.

Let R denote the set of ground terms and R^* the set of sequences of ground terms, including the empty sequence $\langle\rangle$. R can be thought of as domain of discourse. This allows for quantification by substitution and equality can be simply defined to be the identity relation.

We begin with the objective part of the semantics:

1. $f, w, z \models P(r_1, \ldots, r_m)$ iff $w[P(r_1, \ldots, r_m), z] = 1$
2. $f, w, z \models (r_1 = r_2)$ iff r_1 and r_2 are identical
3. $f, w, z \models (\alpha_1 \wedge \alpha_2)$ iff $f, w, z \models \alpha_1$ and $f, w, z \models \alpha_2$
4. $f, w, z \models \neg\alpha$ iff $f, w, z \not\models \alpha$
5. $f, w, z \models \forall x.\alpha$ iff $f, w, z \models \alpha^x_r$ for all $r \in R$
6. $f, w, z \models [r]\alpha$ iff $f, w, z \cdot r \models \alpha$
7. $f, w, z \models \Box\alpha$ iff $f, w, z \cdot z' \models \alpha$ for all $z' \in R^*$
8. $f, w, z \models \mathbf{P}\alpha$ iff
   $f, w, z' \models \alpha$ where $z = z' \cdot r$ for some $z' \in R^*$, $r \in R$

To characterize what is known after an action sequence z, we define the relation $w' \simeq_z w$ for any given w (read: w' agrees with w on the sensing for z) as follows:

- $w' \simeq_{\langle\rangle} w$ for all w';
- $w' \simeq_{z \cdot r} w$ iff $w' \simeq_z w$ and $w'[SF(r), z] = w[SF(r), z]$.

$\simeq_z$ corresponds to the accessibility relations in [15] and SPLL.

9. $f, w, z \models \mathbf{K}\alpha$ iff for all $p \in \mathbb{N}$, for all $w' \simeq_z w$,
   if $w' \in f(p)$, then $f, w', z \models \alpha$

Rule 9 defines knowledge in a way similar to $\mathcal{ES}$: $\mathbf{K}\alpha$ holds if α holds in all worlds of the epistemic state (as defined by f).

To ease the presentation of the following semantic rules, it is convenient to write $f, w, z \models \mathbf{K}^p\alpha$ as shorthand for "for all $w' \simeq_z w$, if $w' \in f(p)$, then $f, w', z \models \alpha$" for any $p \in \mathbb{N}$. In other words, the macro expresses knowledge at plausibility level p. Notice that $\mathbf{K}^p$*False* holds if no world is considered possible at plausibility level p, and $\neg\mathbf{K}^p\neg\alpha$ means that there is at least one world which satisfies α at plausibility level p.
Then we have:

10. $f, w, z \models \mathbf{B}\alpha$ iff for all $p \in \mathbb{N}$,
    if $f, w, z \models \mathbf{K}^q$*False* for all $q < p$, then $f, w, z \models \mathbf{K}^p\alpha$
11. $f, w, z \models \mathbf{B}(\phi \Rightarrow \psi)$ iff for all $p \in \mathbb{N}$,
    if $f, w, z \models \mathbf{K}^q\neg\phi$ for all $q < p$, then $f, w, z \models \mathbf{K}^p(\phi \supset \psi)$

$\mathbf{B}\alpha$ and $\mathbf{B}(\phi \Rightarrow \psi)$ both emulate the behavior of their respective counterparts in SPLL: belief as truth in the most plausible worlds and belief conditionals, respectively. Recall that a smaller plausibility value p indicates that the world is more plausible. $\mathbf{B}\alpha$ holds if α holds in all of the most plausible worlds in f, that is, if α is *believed*. To understand rule 11 for $\mathbf{B}(\phi \Rightarrow \psi)$, first assume that there is some world in f which satisfies ϕ. Then $\mathbf{B}(\phi \Rightarrow \psi)$ holds iff all of the most plausible worlds which satisfy ϕ also satisfy ψ. Otherwise, if there is no world in f that satisfies ϕ, rule 11 requires $\neg\phi$ to be true in all worlds in f.

The semantics of only-believing follows:

12. $f, w, z \models \mathbf{O}(\alpha, \Gamma)$ iff for some $p_1, \ldots, p_m \in \mathbb{N} \cup \{\infty\}$,
    (a) for all $p \in \mathbb{N}$, for all $w' \simeq_z w$,
        $w' \in f(p)$ iff $f, w', z \models \alpha \wedge \bigwedge_{i:p_i \ge p}(\phi_i \supset \psi_i)$,
    (b) $f, w, z \models \mathbf{K}^p\neg\phi_i$ for all i and for all p with $p < p_i$, and
    (c) $f, w, z \models \neg\mathbf{K}^{p_i}\neg\phi_i$ for all i with $p_i \neq \infty$

Rule 12 captures the idea that α and the belief conditionals Γ are all that is known and believed, respectively. To this end, each belief conditional $\phi_i \Rightarrow \psi_i$ is assigned a plausibility $p_i \in \mathbb{N} \cup \{\infty\}$. While $f(p)$ is only defined for $p \in \mathbb{N}$, $p = \infty$ handles the case that the antecedent ϕ_i holds in no world at all. The effect of rule 12a is that it assigns a unique set of worlds to $f(p)$, namely those which are compatible with w in terms of $\simeq_z$ and which satisfy α as well as all $\phi_i \supset \psi_i$ where $p_i \geq p$. Example 12 at the end of Section 3.1 below illustrates why $p_i = p$ is not sufficient. Rule 12b requires that the conditional should not be effective at an earlier plausibility level already, and rule 12c asserts that it should indeed be effective at level p_i, that is, at plausibility level p_i, at least one world must satisfy ϕ_i (rule 12c), and in all more plausible worlds, ϕ_i must be false (rule 12b).

We write $f, w \models \alpha$ for $f, w, \langle\rangle \models \alpha$. We sometimes leave out the f (or w) in $f, w \models \alpha$ for objective (or subjective, respectively) α. A set of sentences Σ entails α (written as $\Sigma \models \alpha$) iff for every f, for every w, if $f, w \models \alpha'$ for every $\alpha' \in \Sigma$, then $f, w \models \alpha$. A sentence is valid (written as $\models \alpha$) iff $\{\} \models \alpha$.

2.3 Some Properties of $\mathcal{ESB}$

We start by showing that the logic $\mathcal{ES}$ of [8] is in fact part of $\mathcal{ESB}$. Note that $\mathcal{ES}$ only has two epistemic modalities $Know$ for knowledge and $OKnow$ for only-knowing.

For any formula α of $\mathcal{ES}$ let α^* be the translation of α to $\mathcal{ESB}$, which can be easily defined inductively on the structure of α as follows: In all cases except for $Know$ and $OKnow$, α^* is the identity function, otherwise $Know(\alpha)^* = \mathbf{K}\alpha^*$, and $OKnow(\alpha)^* = \mathbf{O}(\alpha^*, \{\})$.

The truth of an $\mathcal{ES}$ sentence after an action sequence z is defined wrt a world w and a set of possible worlds e, where a world is exactly the same as in $\mathcal{ESB}$. We write $e, w, z \models_{\mathcal{ES}} \alpha$ to denote the satisfaction relation in $\mathcal{ES}$. For space reasons and as the semantic rules of $\mathcal{ES}$ have almost identical counterparts in $\mathcal{ESB}$, we will not define them here.

Theorem 1 *Let α be a sentence of $\mathcal{ES}$. Then α is satisfiable in $\mathcal{ES}$ iff α^* is satisfiable in $\mathcal{ESB}$.*

Proof sketch. For the if direction, let ρ be a ground atom not mentioned in α and let w_i^ρ be such that $w_i^\rho[\rho, z] = i$ and $w_i^\rho[\sigma, z] = w[\sigma, z]$ for all atoms $\sigma \neq \rho$ and for all $z \in R^*$. Let $e_1 = \bigcap_{p \in \mathbb{N}} f(p)$, $e_2 = \bigcup_{p \in \mathbb{N}} f(p)$ and $e = e_1 \cup \{w_0^\rho \mid w \in e_2, w_1^\rho \notin e_1\}$. This construction guarantees that e only-knows a sentence only if all $f(p)$ do. Thus $e, w \models_{\mathcal{ES}} \alpha$ if $f, w \models \alpha^*$. For the only-if direction, $e, w \models_{\mathcal{ES}} \alpha$ implies $f, w \models \alpha^*$ for $f(p) = e$ for all $p \in \mathbb{N}$. ☐

Thus all properties of $\mathcal{ES}$ such as positive and negative introspection of knowledge immediately transfer to $\mathcal{ESB}$. Moreover, it is easy to see that full introspection holds for both $\mathbf{K}$ and $\mathbf{B}$ in all of $\mathcal{ESB}$:

Theorem 2 *Let $\mathbf{L}$ stand for either $\mathbf{K}$ or $\mathbf{B}$.*
Then $\models \Box\mathbf{L}\alpha \supset \mathbf{L}\mathbf{L}\alpha$ and $\models \Box\neg\mathbf{L}\alpha \supset \mathbf{L}\neg\mathbf{L}\alpha$.

Proof. Both statements follow immediately from the fact that for all models f, w, action sequences z, and plausibility levels p, for all $w_1, w_2 \in \{w' \mid w' \in f(p), w' \simeq_z w\}$, $w_1 \simeq_z w_2$. ☐

We now proceed with the major results of the paper: Only-believing a set of beliefs is always satisfiable, and its model is unique.

Lemma 3 $\models \Box\mathbf{O}(\alpha, \Gamma) \supset \mathbf{K}\alpha \wedge \mathbf{B}\Gamma$.

Proof. Follows immediately from the definition of rule 12. ☐

Lemma 4 *If $f \models \mathbf{O}(\alpha, \Gamma)$ for plausibility levels $p_1, \ldots, p_m$, then $\{p_1, \ldots, p_m\} \setminus \{\infty\} = \{0, \ldots, n\}$ for some $n < m$.*

Proof. Suppose $f \models \mathbf{O}(\alpha, \Gamma)$ and for some i and p, $p+1 = p_i \neq \infty$ and there is no j with $p_j = p$, that is, p is a "hole" in the plausibility ranking. By rule 12c, there is some $w' \in f(p_i)$ such that $w' \models \phi_i$. By rule 12a, for all w', $w' \in f(p)$ iff $w' \models \alpha \wedge \bigwedge_{j:p_j > p}(\phi_j \supset \psi_j)$ iff $w' \models \alpha \wedge \bigwedge_{k:p_k \geq p_i}(\phi_k \supset \psi_k)$ iff $w' \in f(p_i)$. Therefore $f(p) = f(p_i)$ and thus there is some $w' \in f(p)$ such that $w' \models \phi_i$, which contradicts rule 12b. ☐

Lemma 5 *Suppose $f \models \mathbf{O}(\alpha, \Gamma)$ and $g \models \mathbf{O}(\alpha, \Gamma)$. Let $p_1, \ldots, p_m$ be plausibility levels which satisfy rule 12 wrt f. Then for all $p \neq \infty$, there are plausibility levels $p'_1, \ldots, p'_m$ which satisfy rule 12 wrt g such that, if for all $q < p$, $\{i \mid p_i = q\} = \{i \mid p'_i = q\}$, then $\{i \mid p_i = p\} = \{i \mid p'_i = p\}$.*

Proof. We prove by contradiction. Suppose $p'_1, \ldots, p'_m$ satisfy rule 12 wrt g and for all $q < p$, $\{i \mid p_i = q\} = \{i \mid p'_i = q\}$, but $\{i \mid p_i = p\} \neq \{i \mid p'_i = p\}$.

Let $p \neq \infty$ be arbitrary and define $I = \{i \mid p_i = p\}$, $I' = \{i \mid p'_i = p\}$, $J = \{j \mid p_j > p\}$ and $J' = \{j \mid p'_j > p\}$. Note that for all $i \in I' \setminus I$, $p_i > p$, and for all $i \in I \setminus I'$, $p'_i > p$. Also note that $J' \setminus J = I \setminus I'$ and $J \setminus J' = I' \setminus I$.

Wlog assume $I' \setminus I \neq \{\}$. By rule 12a wrt g, all $w' \in g(p)$ satisfy

$$\alpha \wedge \bigwedge_{i \in I \cap I'}(\phi_i \supset \psi_i) \wedge \bigwedge_{i \in I' \setminus I}(\phi_i \supset \psi_i) \wedge \bigwedge_{j \in J \cap J'}(\phi_j \supset \psi_j) \wedge \bigwedge_{j \in J' \setminus J}(\phi_j \supset \psi_j).$$

By the above equalities, we can substitute $I' \setminus I$ in the second conjunction with $J \setminus J'$, and similarly we can replace $J' \setminus J$ in the fourth conjunction with $I \setminus I'$. Therefore, w' also satisfies the formula of rule 12a wrt f. Hence, $w' \in f(p)$. Therefore by rule 12b wrt f, for each $j \in J \setminus J'$, $w' \models \neg\phi_j$. Thus by the above equalities, for each $i \in I' \setminus I$, for all $w' \in g(p)$, $w' \models \neg\phi_i$. However, since $p \neq \infty$, by rule 12c, for each $i \in I' \setminus I$ there is some $w' \in g(p)$ with $w' \models \phi_i$. Contradiction. ☐

Theorem 6 *If $f \models \mathbf{O}(\alpha, \Gamma)$ and $g \models \mathbf{O}(\alpha, \Gamma)$, then $f = g$.*

Proof. Let $p_1, \ldots, p_m$ be plausibilities which satisfy rule 12 wrt f. Thus by rule 12a, for all p, for all w', $w' \in f(p)$ iff $w' \models \alpha \wedge \bigwedge_{i:p_i \geq p}(\phi_i \supset \psi_i)$. From Lemma 5 inductively follows that the same plausibilities satisfy rule 12 wrt g. Thus by rule 12a, for all p, for all w', $w' \in g(p)$ iff $w' \models \alpha \wedge \bigwedge_{i:p_i \geq p}(\phi_i \supset \psi_i)$. Therefore for all p, for all w', $w' \in f(p)$ iff $w' \in g(p)$. ☐

We will see in an example in the next section that this unique-model property greatly simplifies proofs of belief revision. In fact, there is even a straightforward way to generate the model:

Theorem 7 *For any α and Γ, $\mathbf{O}(\alpha, \Gamma)$ is satisfiable.*

Proof. We construct an epistemic state f such that $f \models \mathbf{O}(\alpha, \Gamma)$. Initially, set $p_1 := 0, \ldots, p_m := 0$ and $p := 0$. Then set $f(p) := \{w' \mid w' \models \alpha \wedge \bigwedge_{i:p_i \geq p}(\phi_i \supset \psi_i)\}$. Set $p_i := p + 1$ for all i which violate rule 12c, that is, there is no $w' \in f(p)$ with $w' \models \phi_i$. Then let $p := p + 1$. Repeat the loop until $p > m$, as for all i either $0 \leq p_i < m$ or $p_i = \infty$ due to Lemma 4. Finally set $f(q) := f(m)$ for all $q > m$. By construction, rules 12a, 12b, and 12c are satisfied. ☐

We remark this does not constitute an effective computation of the plausibility ranking as the method appeals to first-order entailment, an undecidable problem.

As a consequence of this theorem and Lemma 3 we obtain that $\mathbf{K}\alpha \wedge \mathbf{B}\Gamma$ is also satisfiable:

Corollary 8 *For any* α *and* Γ, $\mathbf{K}\alpha \wedge \mathbf{B}\Gamma$ *is satisfiable.*

2.4 Basic Action Theories

Lakemeyer and Levesque have shown that $\mathcal{ES}$ is able to express Reiter-style *Basic Action Theories* [10, 8]. These are intended to describe action preconditions (Σ_{pre}), action effects (Σ_{post}), sensing results (Σ_{sense}), and what holds initially (Σ_0). In $\mathcal{ESB}$ we add another component Σ_{belief}, which represents the initial beliefs of the agent in terms of belief conditionals.
More precisely, we have for a given set of predicate symbols $\mathcal{F}$:[4]

- Σ_{pre} is a singleton sentence of the form $\Box Poss(a) \equiv \pi$ for a fluent formula π;
- Σ_{post} contains for every $F \in \mathcal{F}$ a sentence $\Box[a]F(\vec{x}) \equiv \gamma_F$ where γ_F is a fluent formula;
- Σ_{sense} is a singleton sentence of the form $\Box SF(a) \equiv \varphi$ for a fluent formula φ;
- Σ_0 is a set of fluent sentences;
- Σ_{belief} is a set of belief conditionals $\phi \Rightarrow \psi$ where ϕ and ψ are fluent sentences.

The sentences in Σ_{post} are called *successor state axioms*. SSAs define how fluent truth values evolve throughout actions and incorporate Reiter's solution to the frame problem [14]. Σ_{sense} is intended to say that action a returns *true* as a sensing result if φ holds and *false* otherwise.

Let σ denote the union of Σ_{pre}, Σ_{post}, and Σ_{sense}, that is, all the non-static parts of the above sentences. A basic action theory Σ is then defined as $\Sigma_0 \cup \{\sigma, \mathbf{O}(\sigma, \Sigma_{belief})\}$.[5] In other words, we assume that the agent has correct knowledge about how actions work (σ), but its beliefs may differ from what is actually true in the world (Σ_0).

For an example basic action theory we refer to Section 3.1

3 BELIEF REVISION IN $\mathcal{ESB}$

In this section we show that $\mathcal{ESB}$ is suitable for belief revision. After showing some properties which also hold in SPLL, we give a few examples. Among other things, the examples illustrate that due to the unique-model property of only-believing (Theorem 6), proofs are much easier in our framework than in the one of SPLL.

Following SPLL, we distinguish between update actions and revision actions. Given a basic action theory Σ, an *update action* r for a formula α is a physical action that always makes true α in the real world, regardless of what was true before. Formally it is characterized by $\Sigma \models \Box[r]\alpha$ and $\Sigma \models \Box SF(r)$. A *revision action* r for a formula α, often just called a sensing action, does not have any real world effect but it may affect the agent's knowledge. In formulas, $\Sigma \models \Box[r]F(\vec{x}) \equiv F(\vec{x})$ for all $F \in \mathcal{F}$ and $\Sigma \models \Box SF(r) \equiv \alpha$. For example, to burn x is an update action for "x is burned," while checking whether x is burned is a revision action for the same formula.

Similar to SPLL we obtain the following:

Theorem 9 *For a revision action* r *for* α, $\Sigma \models \Box\alpha \supset [r]\mathbf{B}\alpha$ *and* $\Sigma \models \Box\neg\alpha \supset [r]\mathbf{B}\neg\alpha$. *For an update action* r *for* α, $\Sigma \models \Box[r]\mathbf{B}\alpha$.

[4] We assume $\Box$ has lower and $[t]$ has higher precedence than logical connectives, so that $\Box[a]F(\vec{x}) \equiv \gamma_F$ stands for $\forall a \forall \vec{x} \Box([a]F(\vec{x}). \equiv \gamma_F)$.

[5] We abuse notation and do not distinguish finite sets of sentences from conjunctions.

Theorem 10 *For any revision action* r *for* α,
$\Sigma \models \Box\mathbf{B}\neg\alpha \wedge [r]\mathbf{B}\alpha \supset [r]\mathbf{BP}(\alpha \wedge \mathbf{B}\neg\alpha)$.

In other words, if α is believed to be false and after a revision action is believed to be true, then the agent realizes that it was mistaken before.

Proof. Suppose $f, w, z \models \mathbf{B}\neg\alpha \wedge [r]\mathbf{B}\alpha$. Let $p_1 \in \mathbb{N}$ ($p_2 \in \mathbb{N}$) be maximal such that for all $p < p_1$ ($p < p_2$), there is no $w' \in f(p)$ with $w' \simeq_z w$ ($w' \simeq_{z \cdot r} w$). If the condition holds for all p_1 (p_2), let p_1 (p_2) be 0. We show $[r]\mathbf{BP}\alpha$ and $[r]\mathbf{BPB}\neg\alpha$ separately.

$f, w, z \models [r]\mathbf{BP}\alpha$ iff $f, w', z \cdot r \models \mathbf{P}\alpha$ for all $w' \in f(p_2)$ with $w' \simeq_{z \cdot r} w$ iff $f, w', z \models \alpha$ for all $w' \in f(p_2)$ with $w' \simeq_{z \cdot r} w$ iff $f, w', z \cdot r \models \alpha$ for all $w' \in f(p_2)$ with $w' \simeq_{z \cdot r} w$ (because r, as a revision action, has no physical effect on α) iff $f, w, z \models [r]\mathbf{B}\alpha$.

$f, w, z \models [r]\mathbf{BPB}\neg\alpha$ iff $f, w', z \cdot r \models \mathbf{PB}\neg\alpha$ for all $w' \in f(p_2)$ with $w' \simeq_{z \cdot r} w$ iff $f, w', z \models \mathbf{B}\neg\alpha$ for all $w' \in f(p_2)$ with $w' \simeq_{z \cdot r} w$ iff $f, w'', z \models \neg\alpha$ for all $w'' \in f(p_1)$ with $w'' \simeq_z w'$ for all $w' \in f(p_2)$ with $w' \simeq_{z \cdot r} w$ iff $f, w'', z \models \neg\alpha$ for all $w'' \in f(p_1)$ with $w'' \simeq_z w$ (because w' and w agree on the sensing throughout $z \cdot r$) iff $f, w, z \models \mathbf{B}\neg\alpha$. $\Box$

3.1 Examples

Consider the following example, which is taken from SPLL: We live in a world with two rooms, and we are always in one of them. Each of the rooms has a light which is independent from the respective other room. We can go from one room to the other, and through sensing actions we can perceive in which room we are and whether the light is on or off in the room we are currently in. We will illustrate how $\mathcal{ESB}$ handles this example, and we will particularly see why only-believing spares us the trouble of adding negative belief conditionals, $\neg\mathbf{B}(\phi \Rightarrow \psi)$, which are needed in SPLL.

We use the following symbols in the example: predicates R_1, L_1, L_2, and actions lv, $\mathsf{sR_1}$, sL. The meaning is as follows. R_1 indicates that we are in the first room, $\neg R_1$ indicates that we are in the second room. $(\neg)L_1$ and $(\neg)L_2$ represent that the light is on (off) in the first and second room, respectively. By the physical action lv we leave the current and enter the other room. The sensing action $\mathsf{sR_1}$ tells us whether or not we are currently in room one. By the sensing action sL we learn whether or not the light is on in the room we are currently in. We assume that any action is always possible. Σ_{pre}, Σ_{post}, and Σ_{sense} thus are as follows:

$$\left.\begin{array}{l} \Box Poss(a) \equiv True \\ \Box[a]R_1 \equiv \neg R_1 \wedge a = \mathsf{lv} \vee R_1 \wedge a \neq \mathsf{lv} \\ \Box[a]L_1 \equiv L_1 \\ \Box[a]L_2 \equiv L_2 \\ \Box SF(a) \equiv a = \mathsf{sL} \wedge L_1 \wedge R_1 \vee \\ \qquad\qquad a = \mathsf{sL} \wedge L_2 \wedge \neg R_1 \vee \\ \qquad\qquad a = \mathsf{lv} \vee \\ \qquad\qquad a = \mathsf{sR_1} \wedge R_1 \end{array}\right\} = \sigma$$

In reality, the light is on in both rooms and we are initially located in the second room. But we *believe* that we are in the first room and the light is off in the first room. We continue to believe that we are in room one when we learn that the light is on in the first room. Furthermore we believe that, if we are in the second room, the light there is off. Thus,

$$\Sigma_0 = \{L_1, L_2, \neg R_1\} \text{ and}$$
$$\Sigma_{belief} = \{True \Rightarrow \neg L_1 \wedge R_1, L_1 \Rightarrow R_1, \neg R_1 \Rightarrow \neg L_2\}.$$

According to our definition of a basic action theory, $\Sigma = \Sigma_0 \cup \{\sigma, \mathbf{O}(\sigma, \Sigma_{belief})\}$.

We will show that the following properties are entailed by Σ:

1. $\Sigma \models \mathbf{B}\neg L_1$
2. $\Sigma \models [\mathsf{sL}]\mathbf{B}(L_1 \wedge R_1)$
3. $\Sigma \models [\mathsf{sL}][\mathsf{sR_1}]\mathbf{B}\neg R_1$
4. $\Sigma \models [\mathsf{sL}][\mathsf{sR_1}]\mathbf{BP}(\neg R_1 \wedge \mathbf{B}R_1)$
5. $\Sigma \models \neg[\mathsf{sL}][\mathsf{sR_1}]\mathbf{B}L_1 \wedge \neg[\mathsf{sL}][\mathsf{sR_1}]\mathbf{B}\neg L_1$
6. $\Sigma \models [\mathsf{sL}][\mathsf{sR_1}][\mathsf{lv}]\mathbf{B}R_1$
7. $\Sigma \models [\mathsf{sL}][\mathsf{sR_1}][\mathsf{lv}][\mathsf{sL}]\mathbf{B}L_1$

The meaning of most properties is straightforward. Property 2 means that after sensing that the light is on in the room we are in, we believe that we are in room one. In Property 3 we learn that in fact we were in room two all along. Property 4 means that we are aware of our mistake: before sensing that we are in room two, we were in that room already, but we did not believe that. Property 5 is an example of becoming indifferent towards something: we have no opinion on the light being on or off.

Each of these properties has a corresponding counterpart in SPLL. However, SPLL need additional axioms $\neg\mathbf{B}(L_2 \wedge \neg R_1 \Rightarrow L_1)$ and $\neg\mathbf{B}(L_2 \wedge \neg R_1 \Rightarrow \neg L_1)$ to obtain the properties. We will first show that only-believing the positive belief conditionals alone entails all properties. After that we demonstrate that this is not the case for just believing the positive conditionals.

There is another slight difference between our example and SPLL's: we use a third positive belief conditional, $\neg R_1 \Rightarrow \neg L_2$. The intuitive purpose of the negative belief conditionals is to enforce possible but rather implausible worlds. However, they also have some shrouded side effects such as introducing a few new positive beliefs. Among other things, they assert in each model that $\mathbf{B}(\neg R_1 \Rightarrow \neg L_2)$ holds unless there is no world satisfying $\neg R_1$ more plausible than those created by the negative belief conditionals. In $\mathcal{ESB}$, such worlds do exist due to only-believing. Thus we believe our additional positive belief conditional is perfectly justified.

Obviously Σ is satisfiable, and the epistemic state can be generated as described in the proof of Theorem 7. The plausibility levels are $p_1 = 0$ for *True* $\Rightarrow \neg L_1 \wedge R_1$ and $p_2 = p_3 = 1$ for the other two belief conditionals, and the epistemic state is

$$f(0) = \{w' \mid w' \models \sigma \wedge \neg L_1 \wedge R_1\}$$
$$f(1) = \{w' \mid w' \models \sigma \wedge ((\neg L_1 \wedge \neg L_2) \vee R_1)\}$$
$$f(p) = \{w' \mid w' \models \sigma\} \text{ for all } p \geq 2.$$

We now show the above properties.

1. $f, w \models \mathbf{B}\neg L_1$:
   Follows because for all $w' \in f(0)$, $w' \models \neg L_1$.
2. $f, w \models [\mathsf{sL}]\mathbf{B}(L_1 \wedge R_1)$:
   For the real world w, $w[\mathit{SF}(\mathsf{sL}), \langle\rangle] = 1$, but $w'[\mathit{SF}(\mathsf{sL}), \langle\rangle] = 0$ for all $w' \in f(0)$ and thus $w' \not\simeq_{\langle \mathsf{sL}\rangle} w$. For all $w' \simeq_{\langle \mathsf{sL}\rangle} w$ either $w' \models L_1 \wedge R_1$ or $w' \models L_2 \wedge \neg R_1$. Only the former exist in $f(1)$.
3. $f, w \models [\mathsf{sL}][\mathsf{sR_1}]\mathbf{B}\neg R_1$:
   We have argued in Property 2 that there is no $w' \in f(0)$ with $w' \simeq_{\langle \mathsf{sL}\rangle} w$. Since $w, \langle \mathsf{sL}\rangle \models \neg R_1$, also $w[\mathit{SF}(\mathsf{sR_1}), \langle \mathsf{sL}\rangle] = 0$. On the other hand, as shown in Property 2, $w', \langle \mathsf{sL}\rangle \models R_1$ for all $w' \in f(1)$ with $w' \simeq_{\langle \mathsf{sL}\rangle} w$, and thus $w'[\mathit{SF}(\mathsf{sR_1}), \langle \mathsf{sL}\rangle] = 1$. Therefore, $w' \not\simeq_{\langle \mathsf{sL}, \mathsf{sR_1}\rangle} w$. Thus we arrive at plausibility level 2, and obviously there are worlds $w' \in f(2)$ with $w' \simeq_{\langle \mathsf{sL}, \mathsf{sR_1}\rangle} w$, which are precisely those with $w', \langle \mathsf{sL}, \mathsf{sR_1}\rangle \models \neg R_1$.
4. $f, w \models [\mathsf{sL}][\mathsf{sR_1}]\mathbf{BP}(\neg R_1 \wedge \mathbf{B}R_1)$:
   This property is an instance of Theorem 10, so its proof is just the proof of Theorem 10 with $z = \langle \mathsf{sL}\rangle$, $r = \mathsf{sR_1}$, $\alpha = \neg R_1$, and, according to Properties 2 and 3, $p_1 = 1$ and $p_2 = 2$.
5. $f, w \models \neg[\mathsf{sL}][\mathsf{sR_1}]\mathbf{B}L_1 \wedge \neg[\mathsf{sL}][\mathsf{sR_1}]\mathbf{B}\neg L_1$:
   As argued in Property 3, no worlds from $f(0)$ and $f(1)$ agree with w on the sensing throughout $\langle \mathsf{sL}, \mathsf{sR_1}\rangle$. According to Property 3, for all worlds $w' \in f(2)$ with $w' \simeq_{\langle \mathsf{sL}, \mathsf{sR_1}\rangle} w$, the sensing only requires $w', \langle \mathsf{sL}, \mathsf{sR_1}\rangle \models \neg R_1 \wedge L_2$, so there are some worlds with $w', \langle \mathsf{sL}, \mathsf{sR_1}\rangle \models L_1$ left and some with $w', \langle \mathsf{sL}, \mathsf{sR_1}\rangle \models \neg L_1$.
6. $f, w \models [\mathsf{sL}][\mathsf{sR_1}][\mathsf{lv}]\mathbf{B}R_1$:
   In Property 3 we have shown that $w', \langle \mathsf{sL}, \mathsf{sR_1}\rangle \models \neg R_1$ for all $w \in f(2)$ with $w' \simeq_{\langle \mathsf{sL}, \mathsf{sR_1}\rangle} w$. Since $w'[R_1, z \cdot \mathsf{lv}] = 1 - w'[R_1, z]$, we have $w', \langle \mathsf{sL}, \mathsf{sR_1}, \mathsf{lv}\rangle \models R_1$. Those worlds satisfy $w' \simeq_{\langle \mathsf{sL}, \mathsf{sR_1}, \mathsf{lv}\rangle} w$ because $\mathit{SF}(\mathsf{lv})$ is trivially *True*.
7. $f, w \models [\mathsf{sL}][\mathsf{sR_1}][\mathsf{lv}][\mathsf{sL}]\mathbf{B}L_1$:
   Due to Property 5, for some $w' \in f(2)$ with $w' \simeq_{\langle \mathsf{sL}, \mathsf{sR_1}, \mathsf{lv}\rangle} w$, we have $w', \langle \mathsf{sL}, \mathsf{sR_1}, \mathsf{lv}\rangle \models L_1$ and for others $w', \langle \mathsf{sL}, \mathsf{sR_1}, \mathsf{lv}\rangle \models \neg L_1$. The sensing is $w[\mathit{SF}(\mathsf{sL}), \langle \mathsf{sL}, \mathsf{sR_1}, \mathsf{lv}\rangle] = 1$. Since we believe to be in room one, only the worlds w' with $w', \langle \mathsf{sL}, \mathsf{sR_1}, \mathsf{lv}\rangle \models L_1$ agree on the sensing.

Note that these proofs are much simpler than proofs in SPLL. This is mainly due to the unique-model property of $\mathbf{O}$ (Theorem 6).

Finally we sketch why Properties 2 and 5 are not entailed when we just believe Σ_{belief} instead of *only*-believing it. Let Σ' denote Σ where we replace $\mathbf{O}(\sigma, \Sigma_{belief})$ with $\mathbf{K}\sigma \wedge \mathbf{B}\Sigma_{belief}$.

As for Property 5, we have seen that sensing has erased the worlds from plausibility levels 0 and 1 from our epistemic state f already. Now let g be such that $g(0) = f(0)$, $g(1) = f(1)$, and $g(p) = \{w' \mid w' \models \sigma \wedge L_1\}$ for $p \geq 2$. Then $g \models \mathbf{B}\Sigma_{belief}$ and thus $g, w \models \Sigma'$, but since $g, w \models [\mathsf{sL}][\mathsf{sR_1}]\mathbf{B}L_1$, Σ' does not entail Property 5.

For Property 2, we can exploit that $\mathbf{B}\Sigma_{belief}$ allows "holes" in the plausibility ordering. Let h be such that

$$h(0) = \{w' \mid w' \models \sigma \wedge \neg L_1 \wedge \neg L_2 \wedge R_1\}$$
$$h(1) = \{w' \mid w' \models \sigma \wedge \neg L_1 \wedge \neg L_2 \wedge \neg R_1\}$$
$$h(2) = \{w' \mid w' \models \sigma \wedge \neg L_1 \wedge L_2 \wedge \neg R_1\}$$
$$h(3) = \{w' \mid w' \models \sigma \wedge L_1 \wedge L_2 \wedge R_1\}.$$

$h \models \mathbf{B}\Sigma_{belief}$, as *True* $\Rightarrow \neg L_1 \wedge R_1$ takes effect at level 0, $L_1 \Rightarrow R_1$ takes effect at level 3, and $\neg R_1 \Rightarrow \neg L_2$ takes effect at level 1. Thus we could define whatever we want at level 2, as long as it satisfies $\neg L_1$. The sensing action sL tells us $(L_1 \wedge R_1) \vee (L_2 \wedge \neg R_1)$ holds. All $w' \in h(0)$ or $w' \in h(1)$ disagree with this sensing, but some $w' \in h(2)$ do agree. However, for all $w' \in h(2)$, $w' \not\models L_1 \wedge R_1$. Thus Σ' does not entail Property 2.

The discussed issues are also present in SPLL; as mentioned earlier they resort to negative belief conditionals to handle them. Our example shows that only-believing is an alternative, perhaps cleaner and more general way to solve this problem.

To conclude this section, we provide a few examples that illustrate the inner workings of the semantics of only-believing.

Example 11 Consider $\Gamma = \{\phi \Rightarrow \mathit{False}\}$. The effect of $\mathbf{B}\Gamma$ is the same as of $\mathbf{K}\neg\phi$ because the conditional's antecedent is unsatisfiable. Analogously, $\mathbf{O}(\mathit{True}, \Gamma)$ is equivalent to $\mathbf{O}(\neg\phi, \{\})$: The plausibility of $\phi \Rightarrow \mathit{False}$ can only be ∞, and therefore according to rule 12a, for all w' and $p \in \mathbb{N}$, $w' \in f(p)$ iff $w' \models \phi \supset \mathit{False}$ iff $w' \models \neg\phi$.

The next example shows why rule 12a requires all worlds at plausibility level p also to satisfy the implication $\phi_i \supset \psi_i$ for all conditionals with $p_i > p$.

Example 12 Consider $\Gamma = \{A \Rightarrow B, C \Rightarrow A \wedge \neg B\}$. Note that the two conditionals cannot have their antecedents true at the same plausibility level. Still, both conditionals can be effective at different levels: $f \models \mathbf{B}\Gamma$ for $f(0) = \{w' \mid w' \models (A \supset B) \wedge \neg C\}$ and $f(1) = \{w' \mid w' \models C \supset (A \wedge \neg B)\}$. The clue is to falsify the second conditional's antecedent in $f(0)$. That is precisely what rule 12a does: Suppose $g \models \mathbf{O}(\mathit{True}, \Gamma)$ where $A \Rightarrow B$ has plausibility level 0 and $C \Rightarrow A \wedge \neg B$ has plausibility level 1. By rule 12a, $w' \in g(0)$ iff $w' \models (A \supset B) \wedge (C \supset A \wedge \neg B)$ iff $w' \models (\neg A \vee B) \wedge (\neg C \vee A) \wedge (\neg C \wedge \neg B)$ iff $w' \models (\neg A \wedge \neg C) \vee (B \wedge \neg C)$ iff $w' \models (A \supset B) \wedge \neg C$. Observe that requiring $w' \models C \supset A \wedge \neg B$ for $w' \in g(0)$ in rule 12a precisely keeps out those worlds from $g(0)$ which otherwise would by rule 12c trigger the second conditional and thus make things inconsistent. Thus $f \models \mathbf{O}(\mathit{True}, \Gamma)$.

The final example demonstrates that in some cases, $\mathbf{K}\alpha \wedge \mathbf{B}\Gamma \not\models \mathbf{B}(\phi \Rightarrow \psi)$ but $\mathbf{O}(\alpha, \Gamma) \models \mathbf{B}(\phi \Rightarrow \psi)$.

Example 13 Consider $\Gamma = \{A \Rightarrow C, B \Rightarrow D\}$ and $A \wedge B \Rightarrow C \wedge D$. A model of $\mathbf{B}\Gamma$ is f such that $f(0) = \{w' \mid w' \models A \wedge \neg B \wedge C \wedge \neg D\}$, $f(1) = \{w' \mid w' \models \neg A \wedge B \wedge \neg C \wedge D\}$, $f(2) = \{w' \mid w' \models A \wedge B \wedge \neg C \wedge \neg D\}$. It is easy to see that $f \not\models \mathbf{B}(A \wedge B \Rightarrow C \wedge D)$. However, $\mathbf{O}(\mathit{True}, \Gamma) \models \mathbf{B}(A \wedge B \Rightarrow C \wedge D)$: Suppose $g \models \mathbf{O}(\mathit{True}, \Gamma)$. $w' \in g(0)$ iff $w' \models (A \supset C) \wedge (B \supset D)$, which implies $w' \models A \wedge B \supset C \wedge D$.

4 RELATED WORK

The closest relative of our work is of course SPLL [16]. SPLL builds upon the epistemic extension of Reiter's situation calculus [14] by Scherl and Levesque [15]. $\mathcal{ES}$ [8] expands the latter by only-knowing, and so is a suitable basis for our semantic characterization of SPLL plus only-believing. Like SPLL we adapt ideas from [17, 2].

Demolombe and Pozos Parra [4] define belief in terms of modal literals, which evolve according to axioms similar to Reiter's successor state axioms [14]. Unlike SPLL and us, they do not support disjunctive beliefs. The initial (non-)beliefs must be stated explicitly.

Another proposal [5] by the same authors for multi-agent belief revision uses an accessibility relation, but is able to avoid plausibilities by distinguishing between real and imaginary situations. They argue that SPLL's plausibilities are infeasible. We believe our work refutes this claim as only-believing induces unique plausibilities.

Fang and Liu's proposal [6] also supports multi-agent belief. They feature two plausibility relations, one for actions and one for situations. When an action occurs, the situations' plausibilities are updated based on the action's plausibility. Plausibilities must be assigned by hand; they do not have a $\Rightarrow$ operator like SPLL.

A more distant relative by del Val and Shoham [3] uses a circumscription policy to minimize the effects of belief update and belief revision. They also provide an operator for *believing only* which closes the initial beliefs under logical consequence.

As we mentioned in the beginning, SPLL showed how they agree and differ from the classical approaches to AGM-style belief revision [1], update [7], and iterated revision [2]. For compatibility reasons, SPLL need to assume that the actual world is always considered possible with some plausibility. In our case, this would mean that we require that the real world is an element of $f(p)$ for some p. With this it is easy to show that our approach satisfies the same postulates as SPLL. For example, they show that they satisfy all AGM postulates except $K*5$, which says that revising with α leads to an inconsistent epistemic state iff α is inconsistent. Instead, SPLL as well as we have that revising with α never leads to an inconsistent state, provided the real world is considered possible. This is because α is assumed to be the result of a sensing action, which always returns the correct value wrt the real world. We remark that SPLL can carry out their comparison with AGM and others only by assuming a particular model of a given basic action theory. In our case, since only-believing has the unique-model property, the comparison can be carried out in terms of logical entailment within $\mathcal{ESB}$.

5 CONCLUSION

The paper semantically characterizes belief change in the situation calculus in the spirit of Shapiro et al. [16]. Our logic allows to define beliefs in terms of belief conditionals and reason about how belief is updated and revised over the course of actions. In particular we have defined a novel *only-believing* operator with interesting properties: while only-believing and believing a set of belief conditionals both are always satisfiable, the former has a unique model and thus a unique plausibility ordering of beliefs, and there is a straightforward way to generate this model.

In future we plan to combine our results on only-believing with $\mathcal{ESL}$, a logic for limited reasoning about actions [11]. This work promises to allow decidable reasoning about beliefs.

Only-knowing has been shown to have a close relationship with the progression of knowledge bases after actions [9]. Thus, progression of beliefs may be another interesting application of only-believing.

REFERENCES

[1] Carlos E. Alchourron, Peter Gärdenfors, and David Makinson, 'On the logic of theory change: Partial meet contraction and revision functions', *Journal of Symbolic Logic*, **50**(2), 510–530, (1985).

[2] Adnan Darwiche and Judea Pearl, 'On the logic of iterated belief revision', *Artificial intelligence*, **89**(1), 1–29, (1997).

[3] Alvaro del Val and Yoav Shoham, 'A unified view of belief revision and update', *Journal of Logic and Computation*, **4**(5), 797–810, (1994).

[4] Robert Demolombe and Maria del Pilar Pozos Parra, 'A simple and tractable extension of situation calculus to epistemic logic', in *Foundations of Intelligent Systems*, LNCS, (2000).

[5] Robert Demolombe and Maria del Pilar Pozos Parra, 'Belief revision in the situation calculus without plausibility levels', in *Foundations of Intelligent Systems*, LNCS, (2006).

[6] Liangda Fang and Yongmei Liu, 'Multiagent knowledge and belief change in the situation calculus', in *Proc. AAAI*, (2013).

[7] Hirofumi Katsuno and Alberto O. Mendelzon, 'Propositional knowledge base revision and minimal change', *Artificial Intelligence*, **52**(3), 263–294, (1991).

[8] Gerhard Lakemeyer and Hector J. Levesque, 'Situations, si! situation terms, no!', in *Proc. KR*, (2004).

[9] Gerhard Lakemeyer and Hector J. Levesque, 'A semantical account of progression in the presence of defaults', in *Proc. IJCAI*, (2009).

[10] Gerhard Lakemeyer and Hector J. Levesque, 'A semantic characterization of a useful fragment of the situation calculus with knowledge', *Artificial Intelligence*, **175**(1), 142–164, (2011).

[11] Gerhard Lakemeyer and Hector J. Levesque, 'Decidable reasoning in a fragment of the epistemic situation calculus', in *Proc. KR*, (2014).

[12] David Lewis, *Counterfactuals*, John Wiley & Sons, 2013.

[13] John McCarthy, 'Situations, Actions, and Causal Laws', Technical Report AI Memo 2, AI Lab, Stanford University, (July 1963).

[14] Raymond Reiter, *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*, The MIT Press, 2001.

[15] Richard Scherl and Hector J. Levesque, 'Knowledge, action, and the frame problem', *Artificial Intelligence*, **144**(1–2), 1–39, (2003).

[16] Steven Shapiro, Maurice Pagnucco, Yves Lespérance, and Hector J. Levesque, 'Iterated belief change in the situation calculus', *Artificial Intelligence*, **175**(1), 165–192, (2011).

[17] Wolfgang Spohn, 'Ordinal conditional functions: A dynamic theory of epistemic states', in *Causation in Decision, Belief Change, and Statistics*, eds., William L. Harper and Brian Skyrms, (1988).