

Don't Plan for the Unexpected: Planning Based on Plausibility Models

Mikkel Birkegaard Andersen Thomas Bolander
Martin Holm Jensen

October 10, 2013

Abstract

We present a framework for automated planning based on plausibility models, as well as algorithms for computing plans in this framework. Our plausibility models include postconditions, as ontic effects are essential for most planning purposes. The framework presented extends a previously developed framework based on dynamic epistemic logic (DEL), without plausibilities/beliefs. In the pure epistemic framework, one can distinguish between strong and weak epistemic plans for achieving some, possibly epistemic, goal. By taking all possible outcomes of actions into account, a strong plan guarantees that the agent achieves this goal. Conversely, a weak plan promises only the possibility of leading to the goal. In real-life planning scenarios where the planning agent is faced with a high degree of uncertainty and an almost endless number of possible exogenous events, strong epistemic planning is not computationally feasible. Weak epistemic planning is not satisfactory either, as there is no way to qualify which of two weak plans is more likely to lead to the goal. This seriously limits the practical uses of weak planning, as the planning agent might for instance always choose a plan that relies on serendipity. In the present paper we introduce a planning framework with the potential of overcoming the problems of both weak and strong epistemic planning. This framework is based on plausibility models, allowing us to define different types of plausibility planning. The simplest type of plausibility plan is one in which the goal will be achieved when all actions in the plan turn out to have the outcomes found most plausible by the agent. This covers many cases of everyday planning by human agents, where we—to limit our computational efforts—only plan for the most plausible outcomes of our actions.

1 Introduction

Whenever an agent deliberates about the future with the purpose of achieving a goal, she is engaging in the act of planning. Automated Planning is a widely studied area of AI dealing with such issues under many different assumptions and restrictions. In this paper we consider *planning under uncertainty* [11] (nondeterminism and partial observability), where the agent has knowledge and beliefs about the environment and how her actions affect it. We formulate scenarios using plausibility models obtained by merging the frameworks in [5, 21].

Example 1 (The Basement). An agent is standing at the top of an unlit stairwell leading into her basement. If she walks down the steps in the dark, it's likely that she




Figure 1: Three plausibility models.

will trip. On the other hand, if the lights are on, she is certain to descend unharmed. There is a light switch just next to her, though she doesn't know whether the bulb is broken.

She wishes to find a plan that gets her safely to the bottom of the stairs. Planning in this scenario is contingent on the situation; e.g. is the bulb broken? Will she trip when attempting her descent? In planning terminology a plan that *might* achieve the goal is a *weak solution*, whereas one that *guarantees* it is a *strong solution*.

In this case, a weak solution is to simply descend the stairs in the dark, risking life and limb for a trip to the basement. On the other hand, there is no strong solution as the bulb might be broken (assuming it cannot be replaced). Intuitively, the *best* plan is to flick the switch (expecting the bulb to work) and then descend unharmed, something neither weak nor strong planning captures.

Extending the approach in [2] to a logical framework incorporating beliefs via a plausibility ordering, we formalise plans which an agent considers most likely to achieve her goals. This notion is incorporated into algorithms developed for the framework in [2], allowing us to synthesise plans like the *best* one in Example 1.

In the following section we present the logical framework we consider throughout the paper. Section 3 formalises planning in this framework, and introduces the novel concept of plausibility solutions to planning problems. As planning is concerned with representing possible ways in which the future can unfold, it turns out we need a belief modality corresponding to a globally connected plausibility ordering, raising some technical challenges. Section 4 introduces an algorithm for plan synthesis (i.e. generation of plans). Further we show that the algorithm is terminating, sound and complete. To prove termination, we must define bisimulations and bisimulation contractions.

2 Dynamic Logic of Doxastic Ontic Actions

The framework we need for planning is based on a dynamic logic of doxastic ontic actions. Actions can be epistemic (changing knowledge), doxastic (changing beliefs), ontic (changing facts) or any combination. The following formalisation builds on the *dynamic logic of doxastic actions* [5], adding postconditions to event models as in [21]. We consider only the single-agent case. Before the formal definitions are given, we present some intuition behind the framework in the following example, which requires some familiarity with epistemic logic.

Example 2. Consider an agent and a coin biased towards heads, with the coin lying on a table showing heads (h). She contemplates tossing the coin and realizes that it can land either face up, but (due to nature of the coin) believes it will land heads up. In either case, after the toss she knows exactly which face is showing.

The initial situation is represented by the *plausibility model* (defined later) \mathcal{M} and the contemplation by \mathcal{M}'' (see Figure 1). The two worlds u_1 , u_2 are epistemically indistinguishable ($u_1 \not\sim u_2$) and represent the observable non-deterministic outcome of

the toss. The *dashed* directed edge signifies a (*global*) *plausibility relation*, where the direction indicates that she finds u_2 more plausible than u_1 (we overline proposition symbols that are false).

Example 3. Consider again the agent and biased coin. She now reasons about shuffling the coin under a dice cup, leaving the dice cup on top to conceal the coin. She cannot observe which face is up, but due to the bias of the coin believes it to be heads. She then reasons further about lifting the dice cup in this situation, and realises that she will observe which face is showing. Due to her beliefs about the shuffle she finds it most plausible that heads is observed.

The initial situation is again \mathcal{M} . Consider the model \mathcal{M}' , where the *solid* directed edge indicates a *local plausibility relation*, and the direction that v_2 is believed over v_1 . By *local* we mean that the two worlds v_1, v_2 are (*epistemically*) *indistinguishable* ($v_1 \sim v_2$), implying that she is ignorant about whether h or $\neg h$ is the case.¹ Together this represents the concealed, biased coin. Her contemplations on lifting the cup is represented by the model \mathcal{M}'' as in the previous example.

In Example 2 the agent reasons about a non-deterministic action whose outcomes are *distinguishable* but not equally plausible, which is different from the initial contemplation in Example 3 where the outcomes are *not* distinguishable (due to the dice cup). In Example 3 she subsequently reasons about the observations made after a sensing action. In both examples she reasons about the future, and in both cases the final result is the model \mathcal{M}'' . In Example 8 we formally elaborate on the actions used here.

It is the nature of the agent’s ignorance that make \mathcal{M}' and \mathcal{M}'' two inherently different situations. Whereas in the former she is ignorant about h due to the coin being concealed, her ignorance in the latter stems from not having lifted the cup yet. In general we can model ignorance either as a consequence of epistemic indistinguishability, or as a result of not yet having acted. Neither type subsumes the other and both are necessary for reasoning about actions. We capture this distinction by defining both local and global plausibility relations. The end result is that local plausibility talks about belief in a particular epistemic equivalence class, and global plausibility talks about belief in the entire model. We now remedy the informality we allowed ourselves so far by introducing the necessary definitions for a more formal treatment.

Definition 4 (Dynamic Language). Let a countable set of propositional symbols P be given. The language $L(P)$ is given by the following BNF:

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid K\phi \mid B^\psi\phi \mid X\phi \mid [\mathcal{E}, e]\phi$$

where $p \in P$, \mathcal{E} is an *event model* on $L(P)$ as (simultaneously) defined below, and $e \in D(\mathcal{E})$. K is the *local* knowledge modality, B^ψ the *global* conditional belief modality, X is a (non-standard) *localisation* modality (explained later) and $[\mathcal{E}, e]$ the dynamic modality.

We use the usual abbreviations for the other boolean connectives, as well as for the dual dynamic modality $\langle \mathcal{E}, e \rangle \phi := \neg [\mathcal{E}, e] \neg \phi$ and unconditional (or absolute) global belief $B\phi := B^\top \phi$. The duals of K and B^ψ are denoted \widehat{K} and \widehat{B}^ψ .

$K\phi$ reads as “the (planning) agent knows ϕ ”, $B^\psi\phi$ as “conditional on ψ , the (planning) agent believes ϕ ”, and $[\mathcal{E}, e]\phi$ as “after all possible executions of (\mathcal{E}, e) , ϕ holds”. $X\phi$ reads as “locally ϕ ”.

¹In the remainder, we use (in)distinguishability without qualification to refer to epistemic (in)distinguishability.

Definition 5 (Plausibility Models). A *plausibility model* on a set of propositions P is a tuple $\mathcal{M} = (W, \sim, \leq, V)$, where

- W is a set of *worlds*,
- $\sim \subseteq W \times W$ is an equivalence relation called the *epistemic relation*,
- $\leq \subseteq W \times W$ is a connected well-preorder called the *plausibility relation*,²
- $V : P \rightarrow 2^W$ is a *valuation*.

$D(\mathcal{M}) = W$ denotes the *domain* of \mathcal{M} . For $w \in W$ we name (\mathcal{M}, w) a *pointed plausibility model*, and refer to w as the *actual world* of (\mathcal{M}, w) . $<$ denotes the *strict plausibility relation*, that is $w < w'$ iff $w \leq w'$ and $w' \not\leq w$. \simeq denotes *equiplausibility*, that is $w \simeq w'$ iff $w \leq w'$ and $w' \leq w$.

In our model illustrations a directed edge from w to w' indicates $w' \leq w$. By extension, strict plausibility is implied by unidirected edges and equiplausibility by bidirected edges. For the models in Figure 1, we have $v_1 \sim v_2$, $v_2 < v_1$ in \mathcal{M}' and $u_1 \not\sim u_2$, $u_2 < u_1$ in \mathcal{M}'' . The difference between these two models is in the epistemic relation, and is what gives rise to local (solid edges) and global (dashed edges) plausibility. In [5] the local plausibility relation is defined as $\trianglelefteq := \sim \cap \leq$; i.e. $w \trianglelefteq w'$ iff $w \sim w'$ and $w \leq w'$. \trianglelefteq is a *locally well-preordered relation*, meaning that it is a union of *mutually disjoint well-preorders*. Given a plausibility model, the domain of each element in this union corresponds to an \sim -equivalence class.

Our distinction between local and global is not unprecedented in the literature, but it can be a source of confusion. In [5], \leq was indeed connected (i.e. global), but in later versions of the framework [6] this was no longer required. The iterative development in [19] also discuss the distinction between local and global plausibility (named preference by the author). Relating the notions to the wording in [5], \leq captures *a priori* beliefs about *virtual* situations, *before* obtaining any direct information about the actual situation. On the other hand, \trianglelefteq captures *a posteriori* beliefs about an *actual* situation, that is, the agent's beliefs *after* she obtains (or assumes) information about the actual world.

\mathcal{M}'' represents two distinguishable situations (v_1 and v_2) that are a result of reasoning about the future, with v_2 being considered more plausible than v_1 . These situations are identified by restricting \mathcal{M}'' to its \sim -equivalence classes; i.e. $\mathcal{M}'' \upharpoonright \{v_1\}$ and $\mathcal{M}'' \upharpoonright \{v_2\}$. Formally, given an epistemic model \mathcal{M} , the *information cells* in \mathcal{M} are the submodels of the form $\mathcal{M} \upharpoonright [w]_{\sim}$ where $w \in D(\mathcal{M})$. We overload the term and name any \sim -connected plausibility model on P an *information cell*. This use is slightly different from the notion in [6], where an information cell is an \sim -equivalence class rather than a restricted model. An immediate property of information cells is that $\leq = \trianglelefteq$; i.e. the local and global plausibility relations are identical. A partition of a plausibility model into its information cells corresponds to a *localisation* of the plausibility model, where each information cell represents a local situation. The (later defined) semantics of X enables reasoning about such localisations using formulas in the dynamic language.

Definition 6 (Event Models). An *event model* on the language $L(P)$ is a tuple $\mathcal{E} = (E, \sim, \leq, pre, post)$, where

- E is a finite set of (*basic*) *events*,

²A well-preorder is a reflexive, transitive binary relation s.t. every non-empty subset has minimal elements [6].




Figure 2: Three event models.

- $\sim \subseteq E \times E$ is an equivalence relation called the *epistemic relation*,
- $\leq \subseteq E \times E$ is a connected well-preorder called the *plausibility relation*,
- $pre : E \rightarrow L(P)$ assigns to each event a *precondition*,
- $post : E \rightarrow (P \rightarrow L(P))$ assigns to each event a *postcondition* for each proposition. Each $post(e)$ is required to be only finitely different from the identity.

$D(\mathcal{E}) = E$ denotes the *domain* of \mathcal{E} . For $e \in E$ we name (\mathcal{E}, e) a *pointed event model*, and refer to e as the *actual event* of (\mathcal{E}, e) . We use the same conventions for accessibility relations as in the case of plausibility models.

Definition 7 (Product Update). Let $\mathcal{M} = (W, \sim, \leq, V)$ and $\mathcal{E} = (E, \sim', \leq', pre, post)$ be a plausibility model on P resp. event model on $L(P)$. The *product update* of \mathcal{M} with \mathcal{E} is the plausibility model denoted $\mathcal{M} \otimes \mathcal{E} = (W', \sim'', \leq'', V')$, where

- $W' = \{(w, e) \in W \times E \mid \mathcal{M}, w \models pre(e)\}$,
- $\sim'' = \{((w, e), (v, f)) \in W' \times W' \mid w \sim v \text{ and } e \sim' f\}$,
- $\leq'' = \{((w, e), (v, f)) \in W' \times W' \mid e <' f \text{ or } (e \simeq' f \text{ and } w \leq v)\}$,
- $V'(p) = \{(w, e) \in W' \mid \mathcal{M}, w \models post(e)(p)\}$ for each $p \in P$.

The reader may consult [4, 5, 6, 21] for thorough motivations and explanations of the product update. Note that the event model's plausibilities take priority over those of the plausibility model (action-priority update).

Example 8. Consider Figure 2, where the event model \mathcal{E} represents the biased non-deterministic coin toss of Example 2, \mathcal{E}' shuffling the coin under a dice cup, and \mathcal{E}'' lifting the dice cup of Example 3. We indicate \sim and \leq with edges as in our illustrations of plausibility models. Further we use the convention of labelling basic events e by $\langle pre(e), post(e) \rangle$. We write $post(e)$ on the form $\{p_1 \mapsto \phi_1, \dots, p_n \mapsto \phi_n\}$, meaning that $post(e)(p_i) = \phi_i$ for all i , and $post(e)(q) = q$ for $q \notin \{p_1, \dots, p_n\}$.

Returning to Example 2 we see that $\mathcal{M} \otimes \mathcal{E} = \mathcal{M}''$ where $u_1 = (w, e_1), u_2 = (w, e_2)$. In \mathcal{E} we have that $e_2 < e_1$, which encodes the bias of the coin, and $e_1 \not\sim e_2$ encoding the observability, which leads to u_1 and u_2 being distinguishable.

Regarding Example 3 we have that $\mathcal{M} \otimes \mathcal{E}' = \mathcal{M}'$ (modulo renaming). In contrast to \mathcal{E} , we have that $f_1 \sim f_2$, representing the inability to see the face of the coin due to the dice cup. For the sensing action \mathcal{E}'' , we have $\mathcal{M} \otimes \mathcal{E}' \otimes \mathcal{E}'' = \mathcal{M}''$, illustrating how, when events are equiplausible ($g_1 \simeq g_2$), the plausibilities of \mathcal{M}' carry over to \mathcal{M}'' .

We've shown examples of how the interplay between plausibility model and event model can encode changes in belief, and further how to model both ontic change and sensing. In [1] there is a more general treatment of action types, but here such a

classification is not our objective. Instead we simply encode actions as required for our exposition and leave these considerations as future work.

Among the possible worlds, \leq gives an ordering defining what is believed. Given a plausibility model $\mathcal{M} = (W, \sim, \leq, V)$, any non-empty subset of W will have one or more minimal worlds with respect to \leq , since \leq is a well-preorder. For $S \subseteq W$, the set of \leq -minimal worlds, denoted $Min_{\leq} S$, is defined as:

$$Min_{\leq} S = \{s \in S \mid \forall s' \in S : s \leq s'\}.$$

The worlds in $Min_{\leq} S$ are called the *most plausible* worlds in S . The worlds of $Min_{\leq} D(\mathcal{M})$ are referred to as the *most plausible* of \mathcal{M} . With belief defined via minimal worlds (see the definition below), the agent has the same beliefs for any $w \in D(\mathcal{M})$. Analogous to most plausible worlds, an information cell \mathcal{M}' of \mathcal{M} is called *most plausible* if $D(\mathcal{M}') \cap Min_{\leq} D(\mathcal{M}) \neq \emptyset$ (\mathcal{M}' contains at least one of the most plausible worlds of \mathcal{M}).

Definition 9 (Satisfaction Relation). Let a plausibility model $\mathcal{M} = (W, \sim, \leq, V)$ on P be given. The satisfaction relation is given by, for all $w \in W$:

$$\begin{array}{ll} \mathcal{M}, w \models p & \text{iff } w \in V(p) \\ \mathcal{M}, w \models \neg\phi & \text{iff not } \mathcal{M}, w \models \phi \\ \mathcal{M}, w \models \phi \wedge \psi & \text{iff } \mathcal{M}, w \models \phi \text{ and } \mathcal{M}, w \models \psi \\ \mathcal{M}, w \models K\phi & \text{iff } \mathcal{M}, v \models \phi \text{ for all } w \sim v \\ \mathcal{M}, w \models B^{\psi}\phi & \text{iff } \mathcal{M}, v \models \phi \text{ for all } v \in Min_{\leq}\{u \in W \mid \mathcal{M}, u \models \psi\} \\ \mathcal{M}, w \models X\phi & \text{iff } \mathcal{M} \upharpoonright [w]_{\sim}, w \models \phi \\ \mathcal{M}, w \models [\mathcal{E}, e]\phi & \text{iff } \mathcal{M}, w \models pre(e) \text{ implies } \mathcal{M} \otimes \mathcal{E}, (w, e) \models \phi \end{array}$$

where $\phi, \psi \in L(P)$ and (\mathcal{E}, e) is a pointed event model. We write $\mathcal{M} \models \phi$ to mean $\mathcal{M}, w \models \phi$ for all $w \in D(\mathcal{M})$. Satisfaction of the dynamic modality for non-pointed event models \mathcal{E} is introduced by abbreviation, viz. $[\mathcal{E}]\phi := \bigwedge_{e \in \mathcal{D}(\mathcal{E})} [\mathcal{E}, e]\phi$. Furthermore, $\langle \mathcal{E} \rangle \phi := \neg [\mathcal{E}]\neg\phi$.³

The reader may notice that the semantic clause for $\mathcal{M}, w \models X\phi$ is equivalent to the clause for $\mathcal{M}, w \models [\mathcal{E}, e]\phi$ when $[\mathcal{E}, e]$ is a public announcement of a *characteristic formula* [17] being true exactly at the worlds in $[w]_{\sim}$ (and any other world modally equivalent to one of these). In this sense, the X operator can be thought of as a public announcement operator, but a special one that always announces the current information cell. In the special case where \mathcal{M} is an information cell, we have for all $w \in D(\mathcal{M})$ that $\mathcal{M}, w \models X\phi$ iff $\mathcal{M}, w \models \phi$.

3 Plausibility Planning

The previous covered a framework for dealing with knowledge and belief in a dynamic setting. In the following, we will detail how a rational agent would adapt these concepts to model her own reasoning about how her actions affect the future. Specifically, we will show how an agent can predict whether or not a particular plan leads to a desired goal. This requires reasoning about the *conceivable* consequences of actions without *actually* performing them.

³Hence, $\mathcal{M}, w \models \langle \mathcal{E} \rangle \phi \Leftrightarrow \mathcal{M}, w \models \neg [\mathcal{E}]\neg\phi \Leftrightarrow \mathcal{M}, w \models \neg(\bigwedge_{e \in \mathcal{D}(\mathcal{E})} [\mathcal{E}, e]\neg\phi) \Leftrightarrow \mathcal{M}, w \models \bigvee_{e \in \mathcal{D}(\mathcal{E})} \neg [\mathcal{E}, e]\neg\phi \Leftrightarrow \mathcal{M}, w \models \bigvee_{e \in \mathcal{D}(\mathcal{E})} \langle \mathcal{E}, e \rangle \phi$.




Figure 3: The situation before and after attempting to pay with a debit card, plus the event model depicting the attempt. This illustrates that the most plausible information cell can contain the least plausible world.

Two main concepts are required for our formulation of planning, both of which build on notions from the logic introduced in the previous section. One is that of states, a representation of the planning agent’s view of the world at a particular time. Our states are plausibility models. The other concept is that of actions. These represent the agent’s view of everything that can happen when she does something. Actions are event models, changing states into other states via product update.

In our case, the agent has knowledge and beliefs about the initial situation, knowledge and beliefs about actions, and therefore also knowledge and beliefs about the result of actions.

3.1 Reasoning About Actions

Example 10 (Friday Beer). Nearing the end of the month, an agent is going to have an end-of-week beer with her coworkers. Wanting to save the cash she has on hand for the bus fare, she would like to buy the beer using her debit card. Though she isn’t certain, she believes that there’s no money (\bar{m}) on the associated account. Figure 3 shows this initial situation as \mathcal{M} , where \bar{t} signifies that the transaction hasn’t been completed. In this small example her goal is to make t true.

When attempting to complete the transaction (using a normal debit card reader), a number of different things can happen, captured by \mathcal{E} in Figure 3. If there is money on the account, the transaction will go through (e_2), and if there isn’t, it won’t (e_1). This is how the card reader operates most of the time and why e_1 and e_2 are the most plausible events. Less plausible, but still possible, is that the reader malfunctions for some other reason (e_3). The only feedback the agent will receive is whether the transaction was completed, not the reasons why it did or didn’t ($e_1 \sim e_3 \not\sim e_2$). That the agent finds out whether the transaction was successful is why we do not collapse e_1 and e_2 to one event e' with $pre(e') = \top$ and $post(e')(t) = m$.

$\mathcal{M} \otimes \mathcal{E}$ expresses the agent’s view on the possible outcomes of attempting the transaction. The model \mathcal{M}' is the bisimulation contraction of $\mathcal{M} \otimes \mathcal{E}$, according to the definition in Section 4.1 (the world (w_1, e_3) having been removed, as it is bisimilar to (w_1, e_1)).

\mathcal{M}' consists of two information cells, corresponding to whether or not the transaction was successful. What she believes will happen is given by the global plausibility relation. When actually attempting the transaction the result will be one of the information cells of \mathcal{M}' , namely $\mathcal{M}_{\bar{t}} = \mathcal{M}' \upharpoonright \{(w_1, e_1), (w_2, e_3)\}$ or $\mathcal{M}_t = \mathcal{M}' \upharpoonright \{(w_2, e_2)\}$,

in which she will know $\neg t$ and t respectively. As (w_1, e_1) is the most plausible, we can say that she *expects* to end up in (w_1, e_1) , and, by extension, in the information cell $\mathcal{M}_{\bar{t}}$: She expects to end up in a situation where she knows $\neg t$, but is ignorant concerning m . If, unexpectedly, the transaction *is* successful, she will know that the balance is sufficient (m). The most plausible information cell(s) in a model are those the agent expects. That (w_2, e_3) is in the expected information cell, when the globally more plausible world (w_2, e_2) is not, might seem odd. It isn't. The partitioning of \mathcal{M} into the information cells $\mathcal{M}_{\bar{t}}$ and \mathcal{M}_t suggests that she will sense the value of t ($\neg t$ holds everywhere in the former, t everywhere in the latter). As she expects to find out that t does not to hold, she expects to be able to rule out all the worlds in which t *does* hold. Therefore, she expects to be able to rule out (w_2, e_2) and *not* (w_2, e_3) (or w_1, e_1). This gives $\mathcal{M}' \models BX(K\neg t \wedge B\neg m \wedge \tilde{K}m)$: She expects to come to know that the transaction has failed and that she will believe there's no money on the account (though she does consider it possible that there is).

Under the definition of planning that is to follow in Section 3.2, an agent has a number of actions available to construct plans. She needs a notion of which actions can be considered at different stages of the planning process. As in the planning literature, we call this notion *applicability*.

Definition 11 (Applicability). An event model \mathcal{E} is said to be *applicable* in a plausibility model \mathcal{M} if $\mathcal{M} \models \langle \mathcal{E} \rangle \top$.

Unfolding the definition of $\langle \mathcal{E} \rangle$, we see what applicability means:

$$\begin{aligned} \mathcal{M} \models \langle \mathcal{E} \rangle \top &\Leftrightarrow \forall w \in D(\mathcal{M}) : \mathcal{M}, w \models \langle \mathcal{E} \rangle \top \Leftrightarrow \\ &\forall w \in D(\mathcal{M}) : \mathcal{M}, w \models \bigvee_{e \in D(\mathcal{E})} \langle \mathcal{E}, e \rangle \top \Leftrightarrow \\ &\forall w \in D(\mathcal{M}), \exists e \in D(\mathcal{E}) : \mathcal{M}, w \models \langle \mathcal{E}, e \rangle \top \Leftrightarrow \\ &\forall w \in D(\mathcal{M}), \exists e \in D(\mathcal{E}) : \mathcal{M}, w \models \text{pre}(e) \text{ and } \mathcal{M} \otimes \mathcal{E}, (w, e) \models \top \Leftrightarrow \\ &\forall w \in D(\mathcal{M}), \exists e \in D(\mathcal{E}) : \mathcal{M}, w \models \text{pre}(e). \end{aligned}$$

This says that no matter which is the actual world (it must be one of those considered possible), the action defines an outcome. This concept of applicability is equivalent to the one in [1]. The discussion in [9, sect. 6.6] also notes this aspect, insisting that actions must be *meaningful*. The same sentiment is expressed by our notion of applicability.

Proposition 12. *Given a plausibility model \mathcal{M} and an applicable event model \mathcal{E} , we have $D(\mathcal{M} \otimes \mathcal{E}) \neq \emptyset$.*

The product update $\mathcal{M} \otimes \mathcal{E}$ expresses the outcome(s) of doing \mathcal{E} in the situation \mathcal{M} , in the planning literature called *applying* \mathcal{E} in \mathcal{M} . The dynamic modality $[\mathcal{E}]$ expresses reasoning about what holds after applying \mathcal{E} .

Lemma 13. *Let \mathcal{M} be a plausibility model and \mathcal{E} an event model. Then $\mathcal{M} \models [\mathcal{E}]\phi$ iff $\mathcal{M} \otimes \mathcal{E} \models \phi$.*

Proof. $\mathcal{M} \models [\mathcal{E}]\phi \Leftrightarrow \forall w \in \mathcal{D}(\mathcal{M}) : \mathcal{M}, w \models [\mathcal{E}]\phi \Leftrightarrow$
 $\forall w \in \mathcal{D}(\mathcal{M}) : \mathcal{M}, w \models \bigwedge_{e \in \mathcal{D}(\mathcal{E})} [\mathcal{E}, e]\phi \Leftrightarrow \forall (w, e) \in \mathcal{D}(\mathcal{M}) \times \mathcal{D}(\mathcal{E}) : \mathcal{M}, w \models [\mathcal{E}, e]\phi \Leftrightarrow$
 $\forall (w, e) \in \mathcal{D}(\mathcal{M}) \times \mathcal{D}(\mathcal{E}) : \mathcal{M}, w \models \text{pre}(e) \text{ implies } \mathcal{M} \otimes \mathcal{E}, (w, e) \models \phi \Leftrightarrow$
 $\forall (w, e) \in D(\mathcal{M} \otimes \mathcal{E}) : \mathcal{M} \otimes \mathcal{E}, (w, e) \models \phi \Leftrightarrow \mathcal{M} \otimes \mathcal{E} \models \phi. \quad \square$




Figure 4: An information cell, \mathcal{M}_0 , and two event models, flick and desc.




Figure 5: The models resulting from applying the actions flick and desc in \mathcal{M}_0 . Reflexive edges are not shown and the transitive closure is left implicit.

Here we are looking at *global satisfaction*, by evaluating $[\mathcal{E}]\phi$ in all of \mathcal{M} , rather than a specific world. The reason is that evaluation in planning must happen from the perspective of the planning agent and its “information state”. Though one of the worlds of \mathcal{M} is the actual world, the planning agent is ignorant about which it is. Whatever plan it comes up with, it must work in all of the worlds which are indistinguishable to the agent, that is, in the entire model. A similar point, and a similar solution, is found in [13].

Example 14. We now return to the agent from Example 1. Her view of the initial situation (\mathcal{M}_0) and her available actions (flick and desc) are seen in Figure 4. The propositional letters mean t : “top of stairs”, l : “light on”, b : “bulb working”, s : “switch on” and u : “unharmed”. Initially, in \mathcal{M}_0 , she believes that the bulb is working, and knows that she is at the top of the stairs, unharmed and that the switch and light is off: $\mathcal{M}_0 \models Bb \wedge K(t \wedge u \wedge \neg l \wedge \neg s)$.

flick and desc represent flicking the light switch and trying to descend the stairs, respectively. Both require being at the top of the stairs (t). f_1 of flick expresses that if the bulb is working, turning on the switch will turn on the light, and f_2 that if the bulb is broken or the switch is currently on, the light will be off. The events are epistemically distinguishable, as the agent will be able to tell whether the light is on or off. desc describes descending the stairs, with or without the light on. e_1 covers the agent descending the stairs unharmed, and can happen regardless of there being light or not. The more plausible event e_2 represents the agent stumbling, though this can only happen in the dark. If the light is on, she will descend safely. Definition 11 and Lemma 13 let us express the action sequences possible in this scenario.

- $\mathcal{M}_0 \models \langle \text{flick} \rangle \top \wedge \langle \text{desc} \rangle \top$. The agent can initially do either flick or desc.
- $\mathcal{M}_0 \models [\text{flick}] \langle \text{desc} \rangle \top$. After doing flick, she can do desc.
- $\mathcal{M}_0 \models [\text{desc}] (\neg \langle \text{flick} \rangle \top \wedge \neg \langle \text{desc} \rangle \top)$. Nothing can be done after desc.

Figure 5 shows the plausibility models arising from doing flick and desc in \mathcal{M}_0 . Via Lemma 13 she can now conclude:

- $\mathcal{M}_0 \models [\text{flick}] (Kb \vee K\neg b)$: Flicking the light switch gives knowledge of whether the bulb works or not.
- $\mathcal{M}_0 \models [\text{flick}] BKb$. She expects to come to know that it works.
- $\mathcal{M}_0 \models [\text{desc}] (K\neg t \wedge B\neg u)$. Descending the stairs in the dark will definitely get her to the bottom, though she believes she will end up hurting herself.

3.2 Planning

We now turn to formalising planning and then proceed to answer two questions of particular interest: How do we verify that a given plan achieves a goal? And can we compute such plans? This section deals with the first question, plan verification, while the second, plan synthesis, is detailed in Section 4.

Definition 15 (Plan Language). Given a finite set \mathbf{A} of event models on $L(P)$, the *plan language* $\mathcal{L}(P, \mathbf{A})$ is given by:

$$\pi ::= \mathcal{E} \mid \text{skip} \mid \text{if } \phi \text{ then } \pi \text{ else } \pi \mid \pi; \pi$$

where $\mathcal{E} \in \mathbf{A}$ and $\phi \in L(P)$. We name members π of this language *plans*, and use $\text{if } \phi \text{ then } \pi$ as shorthand for $\text{if } \phi \text{ then } \pi \text{ else skip}$.

The reading of the plan constructs are “do \mathcal{E} ”, “do nothing”, “if ϕ then π , else π' ”, and “first π then π' ” respectively. In the translations provided in Definition 16, the condition of the if-then-else construct becomes a K -formula, ensuring that branching depends only on worlds which are distinguishable to the agent. The idea is similar to the *meaningful plans* of [9], where branching is allowed on *epistemically interpretable formulas* only.

Definition 16 (Translation). Let α be one of s , w , sp or wp . We define an α -*translation* as a function $[\cdot]_\alpha : \mathcal{L}(P, \mathbf{A}) \rightarrow (L(P) \rightarrow L(P))$:

$$[\mathcal{E}]_\alpha \phi := \langle \mathcal{E} \rangle \top \wedge \begin{cases} [\mathcal{E}] XK\phi & \text{if } \alpha = s \\ \widehat{K} \langle \mathcal{E} \rangle XK\phi & \text{if } \alpha = w \\ [\mathcal{E}] BXK\phi & \text{if } \alpha = sp \\ [\mathcal{E}] \widehat{B}XK\phi & \text{if } \alpha = wp \end{cases}$$

$$[\text{skip}]_\alpha \phi := \phi$$

$$[\text{if } \phi' \text{ then } \pi \text{ else } \pi']_\alpha \phi := (K\phi' \rightarrow [\pi]_\alpha \phi) \wedge (\neg K\phi' \rightarrow [\pi']_\alpha \phi)$$

$$[\pi; \pi']_\alpha \phi := [\pi]_\alpha ([\pi']_\alpha \phi)$$

We call $[\cdot]_s$ the *strong translation*, $[\cdot]_w$ the *weak translation*, $[\cdot]_{sp}$ the *strong plausibility translation* and $[\cdot]_{wp}$ the *weak plausibility translation*.

The translations are constructed specifically to make the following lemma hold, providing a semantic interpretation of plans (leaving out **skip** and $\pi_1; \pi_2$).

Lemma 17. *Let \mathcal{M} be an information cell, \mathcal{E} an event model and ϕ a formula of $L(P)$. Then:*

1. $\mathcal{M} \models [\mathcal{E}]_s \phi$ iff $\mathcal{M} \models \langle \mathcal{E} \rangle \top$ and for each information cell \mathcal{M}' of $\mathcal{M} \otimes \mathcal{E}$: $\mathcal{M}' \models \phi$.
2. $\mathcal{M} \models [\mathcal{E}]_w \phi$ iff $\mathcal{M} \models \langle \mathcal{E} \rangle \top$ and for some information cell \mathcal{M}' of $\mathcal{M} \otimes \mathcal{E}$: $\mathcal{M}' \models \phi$.
3. $\mathcal{M} \models [\mathcal{E}]_{sp} \phi$ iff $\mathcal{M} \models \langle \mathcal{E} \rangle \top$ and for each most plausible information cell \mathcal{M}' of $\mathcal{M} \otimes \mathcal{E}$: $\mathcal{M}' \models \phi$.

4. $\mathcal{M} \models [\mathcal{E}]_{wp}\phi$ iff $\mathcal{M} \models \langle \mathcal{E} \rangle \top$ and for some most plausible information cell \mathcal{M}' of $\mathcal{M} \otimes \mathcal{E} : \mathcal{M}' \models \phi$.
5. $\mathcal{M} \models [\text{if } \phi' \text{ then } \pi \text{ else } \pi']_{\alpha}\phi$ iff
 $(\mathcal{M} \models \phi' \text{ implies } \mathcal{M} \models [\pi]_{\alpha}\phi)$ and $(\mathcal{M} \not\models \phi' \text{ implies } \mathcal{M} \models [\pi']_{\alpha}\phi)$.

Proof. We only prove 4 and 5, as 1–4 are very similar. For 4 we have:

$$\begin{aligned}
\mathcal{M} \models [\mathcal{E}]_{wp}\phi &\Leftrightarrow \mathcal{M} \models \langle \mathcal{E} \rangle \top \wedge [\mathcal{E}] \widehat{B}XK\phi \Leftrightarrow^{\text{Lemma 13}} \\
\mathcal{M} \models \langle \mathcal{E} \rangle \top \text{ and } \mathcal{M} \otimes \mathcal{E} \models \widehat{B}XK\phi &\Leftrightarrow \\
\mathcal{M} \models \langle \mathcal{E} \rangle \top \text{ and } \forall (w, e) \in D(\mathcal{M} \otimes \mathcal{E}) : \mathcal{M} \otimes \mathcal{E}, (w, e) \models \widehat{B}XK\phi &\Leftrightarrow^{\text{Prop. 12}} \\
\mathcal{M} \models \langle \mathcal{E} \rangle \top \text{ and } \exists (w, e) \in \text{Min}_{\leq} D(\mathcal{M} \otimes \mathcal{E}) : \mathcal{M} \otimes \mathcal{E}, (w, e) \models XK\phi &\Leftrightarrow \\
\mathcal{M} \models \langle \mathcal{E} \rangle \top \text{ and } \exists (w, e) \in \text{Min}_{\leq} D(\mathcal{M} \otimes \mathcal{E}) : \mathcal{M} \otimes \mathcal{E} \upharpoonright [(w, e)]_{\sim}, (w, e) \models K\phi &\Leftrightarrow \\
\mathcal{M} \models \langle \mathcal{E} \rangle \top \text{ and } \exists (w, e) \in \text{Min}_{\leq} D(\mathcal{M} \otimes \mathcal{E}) : \mathcal{M} \otimes \mathcal{E} \upharpoonright [(w, e)]_{\sim} \models \phi &\Leftrightarrow \\
\mathcal{M} \models \langle \mathcal{E} \rangle \top \text{ and in some most plausible information cell } \mathcal{M}' \text{ of } \mathcal{M} \otimes \mathcal{E}, \mathcal{M}' \models \phi. &
\end{aligned}$$

For if-then-else, first note that:

$$\begin{aligned}
\mathcal{M} \models \neg K\phi' \rightarrow [\pi]_{\alpha}\phi &\Leftrightarrow \forall w \in D(\mathcal{M}) : \mathcal{M}, w \models \neg K\phi' \rightarrow [\pi]_{\alpha}\phi \Leftrightarrow \\
\forall w \in D(\mathcal{M}) : \mathcal{M}, w \models \neg K\phi' \text{ implies } \mathcal{M}, w \models [\pi]_{\alpha}\phi &\Leftrightarrow^{\mathcal{M} \text{ is an info. cell}} \\
\forall w \in D(\mathcal{M}) : \text{if } \mathcal{M}, v \models \neg\phi' \text{ for some } v \in D(\mathcal{M}) \text{ then } \mathcal{M}, w \models [\pi]_{\alpha}\phi &\Leftrightarrow \\
\text{if } \mathcal{M}, v \models \neg\phi' \text{ for some } v \in D(\mathcal{M}) \text{ then } \forall w \in D(\mathcal{M}) : \mathcal{M}, w \models [\pi]_{\alpha}\phi &\Leftrightarrow \\
\mathcal{M} \not\models \phi' \text{ implies } \mathcal{M} \models [\pi']_{\alpha}\phi. &
\end{aligned}$$

Similarly, we can prove:

$$\mathcal{M} \models K\phi' \rightarrow [\pi]_{\alpha}\phi \Leftrightarrow \mathcal{M} \models K\phi' \text{ implies } \mathcal{M} \models [\pi']_{\alpha}\phi.$$

Using these facts, we get:

$$\begin{aligned}
\mathcal{M} \models [\text{if } \phi' \text{ then } \pi \text{ else } \pi']_{\alpha}\phi &\Leftrightarrow \mathcal{M} \models (K\phi' \rightarrow [\pi]_{\alpha}\phi) \wedge (\neg K\phi' \rightarrow [\pi']_{\alpha}\phi) \Leftrightarrow \\
\mathcal{M} \models K\phi' \rightarrow [\pi]_{\alpha}\phi \text{ and } \mathcal{M} \models \neg K\phi' \rightarrow [\pi']_{\alpha}\phi &\Leftrightarrow \\
(\mathcal{M} \models \phi' \text{ implies } \mathcal{M} \models [\pi]_{\alpha}\phi) \text{ and } (\mathcal{M} \not\models \phi' \text{ implies } \mathcal{M} \models [\pi']_{\alpha}\phi). &
\end{aligned}$$

□

Using XK (as is done in all translations) means that reasoning after an action is relative to a particular information cell (as $\mathcal{M}, w \models XK\phi \Leftrightarrow \mathcal{M} \upharpoonright [w]_{\sim}, w \models K\phi \Leftrightarrow \mathcal{M} \upharpoonright [w]_{\sim} \models \phi$).

Definition 18 (Planning Problems and Solutions). Let P be a finite set of propositional symbols. A planning problem on P is a triple $\mathcal{P} = (\mathcal{M}_0, \mathbf{A}, \phi_g)$ where

- \mathcal{M}_0 is a finite information cell on P called the *initial state*.
- \mathbf{A} is a finite set of event models on $L(P)$ called the *action library*.
- $\phi_g \in L(P)$ is the *goal (formula)*.

A plan $\pi \in \mathcal{L}(P, \mathbf{A})$ is an α -*solution* to \mathcal{P} if $\mathcal{M}_0 \models [\pi]_{\alpha}\phi_g$. For a specific choice of $\alpha = s/w/sp/wp$, we will call π a *strong/weak/strong plausibility/weak plausibility-solution* respectively.




Figure 6: Event model for replacing a broken bulb.

Given a π , we wish to check whether π is an α -solution (for some particular α) to \mathcal{P} . This can be done via model checking the dynamic formula given by the translation $[\pi]_\alpha \phi_g$ in the initial state of \mathcal{P} .

A strong solution π is one that guarantees that ϕ_g will hold after executing it (“ π achieves ϕ_g ”). If π is a weak solution, it achieves ϕ_g for at least one particular sequence of outcomes. Strong and weak plausibility-solutions are as strong- and weak-solutions, except that they need only achieve ϕ_g for *all of/some of* the most plausible outcomes.

Example 19. The basement scenario (Example 1) can be formalised as the planning problem $\mathcal{P}_B = (\mathcal{M}_0, \{\text{flick}, \text{desc}\}, \phi_g)$ with \mathcal{M}_0 , flick and desc being defined in Figure 4 and $\phi_g = \neg t \wedge u$. Let $\pi_1 = \text{desc}$. We then have that:

$$\begin{aligned} \mathcal{M}_0 \models [\text{desc}]_w (\neg t \wedge u) &\Leftrightarrow \mathcal{M}_0 \models \langle \text{desc} \rangle \top \wedge \widehat{K} \langle \text{desc} \rangle XK(\neg t \wedge u) \Leftrightarrow^{\text{desc is applicable}} \\ \mathcal{M}_0 \models \widehat{K} \langle \text{desc} \rangle XK(\neg t \wedge u) &\Leftrightarrow \exists w \in D(\mathcal{M}_0) : \mathcal{M}_0, w \models \langle \text{desc} \rangle XK(\neg t \wedge u). \end{aligned}$$

Picking w_1 , we have

$$\begin{aligned} \mathcal{M}_0, w_1 \models \langle \text{desc} \rangle XK(\neg t \wedge u) &\Leftrightarrow \mathcal{M}_0 \otimes \text{desc}, (w_1, e_1) \models XK(\neg t \wedge u) \Leftrightarrow \\ \mathcal{M}_0 \otimes \text{desc} \upharpoonright [(w_1, e_1)]_\sim &\models (\neg t \wedge u) \end{aligned}$$

which holds as seen in Figure 5. Thus, π_1 is a weak solution. Further, Lemma 17 tells us that π_1 is not a *sp* solution, as u does not hold in the (most plausible) information cell $\mathcal{M} \otimes \text{desc} \upharpoonright \{(w_1, e_2), (w_2, e_2)\}$.

The plan $\pi_2 = \text{flick}; \text{desc}$ is a strong plausibility solution, as can be verified by $\mathcal{M}_0 \models [\pi_2]_{sp} (\neg t \wedge u)$. Without an action for replacing the lightbulb, there are no strong solutions. Let **replace** be the action in Figure 6, where $\text{post}(r_1)(u) = \neg s$ signifies that if the power is on, the agent will hurt herself, and define a new problem $\mathcal{P}'_B = \{\mathcal{M}_0, \{\text{flick}, \text{desc}, \text{replace}\}, \phi_g\}$. Then $\pi_3 = \text{flick}; (\text{if } \neg l \text{ then flick; replace; flick}); \text{desc}$ is a strong solution (we leave verification to the reader): If the light comes on after flicking the switch (as expected) she can safely walk down the stairs. If it does not, she turns off the power, replaces the broken bulb, turns the power on again (this time knowing that the light will come on), and then proceeds as before.

Besides being an *sp*-solution, π_2 is also a *w*- and a *wp*-solution, indicating a hierarchy of strengths of solutions. This should come as no surprise, given both the formal and intuitive meaning of planning and actions presented so far. In fact, this hierarchy exists for any planning problem, as shown by the following result which is a consequence of Lemma 17 (stated without proof).

Lemma 20. *Let $\mathcal{P} = (\mathcal{M}_0, A, \phi_g)$ be a planning problem. Then:*

- *Any strong solution to \mathcal{P} is also a strong plausibility solution:*
 $\mathcal{M}_0 \models [\pi]_s \phi_g \Rightarrow \mathcal{M}_0 \models [\pi]_{sp} \phi_g.$
- *Any strong plausibility solution to \mathcal{P} is also a weak plausibility solution:*
 $\mathcal{M}_0 \models [\pi]_{sp} \phi_g \Rightarrow \mathcal{M}_0 \models [\pi]_{wp} \phi_g.$
- *Any weak plausibility solution to \mathcal{P} is also a weak solution:*
 $\mathcal{M}_0 \models [\pi]_{wp} \phi_g \Rightarrow \mathcal{M}_0 \models [\pi]_w \phi_g.$

4 Plan Synthesis

In this section we show how to synthesise conditional plans for solving planning problems. Before we can give the concrete algorithms, we establish some technical results which are stepping stones to proving termination of our planning algorithm, and hence decidability of plan existence in our framework.

4.1 Bisimulations, contractions and modal equivalence

We now define bisimulations on plausibility models. For our purpose it is sufficient to define bisimulations on \sim -connected models, that is, on information cells.⁴ First we define a *normal plausibility relation* which will form the basis of our bisimulation definition.

Definition 21 (Normality). Given is an information cell $\mathcal{M} = (W, \sim, \leq, V)$ on P . By slight abuse of language, two worlds $w, w' \in W$ are said to *have the same valuation* if for all $p \in P$: $w \in V(p) \Leftrightarrow w' \in V(p)$. Define an equivalence relation on W : $w \approx w'$ iff w and w' has the same valuation. Now define $w \preceq w'$ iff $\text{Min}_{\leq}([w]_{\approx}) \leq \text{Min}_{\leq}([w']_{\approx})$. This defines the *normal plausibility relation*. \mathcal{M} is called *normal* if $\preceq = \leq$. The *normalisation* of $\mathcal{M} = (W, \sim, \leq, V)$ is $\mathcal{M}' = (W, \sim, \preceq, V)$.

Definition 22 (Bisimulation). Let $\mathcal{M} = (W, \sim, \leq, V)$ and $\mathcal{M}' = (W', \sim', \leq', V')$ be information cells on P . A non-empty relation $\mathcal{R} \subseteq W \times W'$ is a *bisimulation* between \mathcal{M} and \mathcal{M}' (and $\mathcal{M}, \mathcal{M}'$ are called *bisimilar*) if for all $(w, w') \in \mathcal{R}$:

[atom] For all $p \in P$: $w \in V(p)$ iff $w' \in V'(p)$.

[forth] If $v \in W$ and $v \preceq w$ then there is a $v' \in W'$ s.t. $v' \preceq' w'$ and $(v, v') \in \mathcal{R}$.

[back] If $v' \in W'$ and $v' \preceq' w'$ then there is a $v \in W$ s.t. $v \preceq w$ and $(v, v') \in \mathcal{R}$.

If \mathcal{R} has domain W and codomain W' , it is called *total*. If $\mathcal{M} = \mathcal{M}'$, it is called an *autobisimulation* (on \mathcal{M}). Two worlds w and w' of an information cell $\mathcal{M} = (W, \sim, \leq, V)$ are called *bisimilar* if there exists an autobisimulation \mathcal{R} on \mathcal{M} with $(w, w') \in \mathcal{R}$.

We are here only interested in total bisimulations, so, unless otherwise stated, we assume this in the following. Note that our definition of bisimulation immediately implies that there exists a (total) bisimulation between any information cell and its normalisation. Note also that for normal models, the bisimulation definition becomes the standard modal logic one.⁵

Lemma 23. *If two worlds of an information cell have the same valuation they are bisimilar.*

Proof. Assume worlds w and w' of an information cell $\mathcal{M} = (W, \sim, \leq, V)$ have the same valuation. Let \mathcal{R} be the relation that relates each world of \mathcal{M} to itself and additionally relates w to w' . We want to show that \mathcal{R} is a bisimulation. This amounts to showing [atom], [forth] and [back] for the pair $(w, w') \in \mathcal{R}$. [atom] holds trivially since $w \approx w'$.

⁴The proper notion of bisimulation for plausibility structures is explored in more detail by Andersen, Bolander, van Ditmarsch and Jensen in ongoing research. A similar notion for slightly different types of plausibility structures is given in [20]. Surprisingly, Demey does not consider our notion of bisimulation in his thorough survey [10] on different notions of bisimulation for plausibility structures.

⁵We didn't include a condition for the epistemic relation, \sim , in [back] and [forth], simply because we are here only concerned with \sim -connected models.

For [forth], assume $v \in W$ and $v \preceq w$. We need to find a $v' \in W$ s.t. $v' \preceq w'$ and $(v, v') \in \mathcal{R}$. Letting $v' = v$, it suffices to prove $v \preceq w'$. Since $w \approx w'$ this is immediate: $v \preceq w \Leftrightarrow \text{Min}_{\leq}([v]_{\approx}) \leq \text{Min}_{\leq}([w]_{\approx}) \stackrel{w \approx w'}{\Leftrightarrow} \text{Min}_{\leq}([v]_{\approx}) \leq \text{Min}_{\leq}([w']_{\approx}) \Leftrightarrow v \preceq w'$. [back] is proved similarly. \square

Unions of autobisimulations are autobisimulations. We can then in the standard way define the (*bisimulation*) *contraction* of a normal information cell as its quotient with respect to the union of all autobisimulations [8].⁶ The contraction of a non-normal model is taken to be the contraction of its normalisation. In a contracted model, no two worlds are bisimilar, by construction. Hence, by Lemma 23, no two worlds have the same valuation. Thus, the contraction of an information cell on a finite set of proposition symbols P contains at most $2^{|P|}$ worlds. Since any information cell is bisimilar to its contraction [8], this shows that there can only exist finitely many non-bisimilar information cells on any given finite set P .

Two information cells \mathcal{M} and \mathcal{M}' are called *modally equivalent*, written $\mathcal{M} \equiv \mathcal{M}'$, if for all formulas ϕ in $L(P)$: $\mathcal{M} \models \phi \Leftrightarrow \mathcal{M}' \models \phi$. Otherwise, they are called *modally inequivalent*. We now have the following standard result (the result is standard for standard modal languages and bisimulations, but it is not trivial that it also holds here).

Theorem 24. *If two information cells are (totally) bisimilar they are modally equivalent.*

Proof. We need to show that if \mathcal{R} is a total bisimulation between information cells \mathcal{M} and \mathcal{M}' , then for all formulas ϕ of $L(P)$: $\mathcal{M} \models \phi \Leftrightarrow \mathcal{M}' \models \phi$. First we show that we only have to consider formulas ϕ of the static sublanguage of $L(P)$, that is, the language without the $[\mathcal{E}, e]$ modalities. In [5], reduction axioms from the dynamic to the static language are given for a language similar to $L(P)$. The differences in language are our addition of postconditions and the fact that our belief modality is defined from the global plausibility relation rather than being localised to epistemic equivalence classes. The latter difference is irrelevant when only considering information cells as we do here. The former difference of course means that the reduction axioms presented in [5] will not suffice for our purpose. [21] shows that adding postconditions to the language without the doxastic modalities only requires changing the reduction axiom for $[\mathcal{E}, e]p$, where p is a propositional symbol. Thus, if we take the reduction axioms of [5] and replace the reduction axiom for $[\mathcal{E}, e]p$ by the one in [21], we get reduction axioms for our framework. We leave out the details.

We now need to show that if \mathcal{R} is a total bisimulation between information cells \mathcal{M} and \mathcal{M}' , then for all $[\mathcal{E}, e]$ -free formulas ϕ of $L(P)$: $\mathcal{M} \models \phi \Leftrightarrow \mathcal{M}' \models \phi$. Since \mathcal{R} is total, it is sufficient to prove that for all $[\mathcal{E}, e]$ -free formulas ϕ of $L(P)$ and all $(w, w') \in \mathcal{R}$: $\mathcal{M}, w \models \phi \Leftrightarrow \mathcal{M}', w' \models \phi$. The proof is by induction on ϕ . In the induction step we are going to need the induction hypothesis for several different choices of \mathcal{R}, w and w' , so what we will actually prove by induction on ϕ is this: For all formulas ϕ of $L(P)$, if \mathcal{R} is a total bisimulation between information cells \mathcal{M} and \mathcal{M}' on P and $(w, w') \in \mathcal{R}$, then $\mathcal{M}, w \models \phi \Leftrightarrow \mathcal{M}', w' \models \phi$.

The base case is when ϕ is propositional. Then the required follows immediately from [atom], using that $(w, w') \in \mathcal{R}$. For the induction step, we have the following

⁶More precisely, let \mathcal{M} be a normal information cell and let \mathcal{R} be the union of all autobisimulations on \mathcal{M} . Then the contraction $\mathcal{M}' = (W', \sim', \leq', V')$ of \mathcal{M} has as worlds the equivalence classes $[w]_{\mathcal{R}} = \{w' \mid (w, w') \in \mathcal{R}\}$ and has $[w]_{\mathcal{R}} \leq' [w']_{\mathcal{R}}$ iff $v \leq v'$ for some $v \in [w]_{\mathcal{R}}$ and $v' \in [w']_{\mathcal{R}}$.

cases of ϕ : $\neg\psi, \psi \wedge \gamma, X\psi, K\psi, B^\gamma\psi$. The first two cases are trivial. So is $X\psi$, as $X\psi \leftrightarrow \psi$ holds on any information cell. For $K\psi$ we reason as follows. Let \mathcal{R} be a total bisimulation between information cells \mathcal{M} and \mathcal{M}' with $(w, w') \in \mathcal{R}$. Using that \mathcal{R} is total and that \mathcal{M} and \mathcal{M}' are both \sim -connected we get: $\mathcal{M}, w \models K\psi \Leftrightarrow \forall v \in W: \mathcal{M}, v \models \psi \stackrel{\text{i.h.}}{\Leftrightarrow} \forall v' \in W': \mathcal{M}', v' \models \psi \Leftrightarrow \mathcal{M}', w' \models K\psi$.

The case of $B^\gamma\psi$ is more involved. Let $\mathcal{M}, \mathcal{M}', \mathcal{R}, w$ and w' be as above. By symmetry, it suffices to prove $\mathcal{M}, w \models B^\gamma\psi \Rightarrow \mathcal{M}', w' \models B^\gamma\psi$. So assume $\mathcal{M}, w \models B^\gamma\psi$, that is, $\mathcal{M}, v \models \psi$ for all $v \in \text{Min}_{\leq}\{u \in W \mid \mathcal{M}, u \models \gamma\}$. We need to prove $\mathcal{M}', v' \models \psi$ for all $v' \in \text{Min}_{\leq'}\{u' \in W' \mid \mathcal{M}', u' \models \gamma\}$. So let $v' \in \text{Min}_{\leq'}\{u' \in W' \mid \mathcal{M}', u' \models \gamma\}$. By definition of $\text{Min}_{\leq'}$ this means that:

$$\text{for all } u' \in W', \text{ if } \mathcal{M}', u' \models \gamma \text{ then } v' \leq' u'. \quad (1)$$

Choose an $x \in \text{Min}_{\leq}\{u \in W \mid u \approx u' \text{ and } (u', v') \in \mathcal{R}\}$. We want to use (1) to show that the following holds:

$$\text{for all } u \in W, \text{ if } \mathcal{M}, u \models \gamma \text{ then } x \leq u. \quad (2)$$

To prove (2), let $u \in W$ with $\mathcal{M}, u \models \gamma$. Choose u' with $(u, u') \in \mathcal{R}$. The induction hypothesis implies $\mathcal{M}', u' \models \gamma$. We now prove that $v' \leq' \text{Min}_{\leq'}([u']_{\approx})$. To this end, let $u'' \in [u']_{\approx}$. We need to prove $v' \leq' u''$. Since $u'' \approx u'$, Lemma 23 implies that u' and u'' are bisimilar. By induction hypothesis we then get $\mathcal{M}', u'' \models \gamma$.⁷ Using (1) we now get $v' \leq' u''$, as required. This shows $v' \leq' \text{Min}_{\leq'}([u']_{\approx})$. We now have $\text{Min}_{\leq'}([v']_{\approx}) \leq' v' \leq' \text{Min}_{\leq'}([u']_{\approx})$, and hence $v' \preceq u'$. By [back] there is then a v s.t. $(v, v') \in \mathcal{R}$ and $v \preceq u$. By choice of x , $x \leq \text{Min}_{\leq}([v]_{\approx})$. Using $v \preceq u$, we now finally get: $x \leq \text{Min}_{\leq}([v]_{\approx}) \leq \text{Min}_{\leq}([u]_{\approx}) \leq u$. This shows that (2) holds.

From (2) we can now conclude $x \in \text{Min}_{\leq}\{u \in W \mid \mathcal{M}, u \models \gamma\}$ and hence, by original assumption, $\mathcal{M}, x \models \psi$. By choice of x there is an $x' \approx x$ with $(x', v') \in \mathcal{R}$. Since $\mathcal{M}, x \models \psi$ and $x' \approx x$, we can again use Lemma 23 and the induction hypothesis to conclude $\mathcal{M}, x' \models \psi$. Since $(x', v') \in \mathcal{R}$, another instance of the induction hypothesis gives us $\mathcal{M}', v' \models \psi$, and we are done. \square

Previously we proved that there can only be finitely many non-bisimilar information cells on any finite set P . Since we have now shown that bisimilarity implies modal equivalence, we immediately get the following result, which will be essential to our proof of termination of our planning algorithms.

Corollary 25. *Given any finite set P , there are only finitely many modally inequivalent information cells on P .*

4.2 Planning Trees

When synthesising plans, we explicitly construct the search space of the problem as a labelled AND-OR tree, a familiar model for planning under uncertainty [11]. Our AND-OR trees are called *planning trees*.

Definition 26 (Planning Tree). A *planning tree* is a finite, labelled AND-OR tree in which each node n is labelled by a plausibility model $\mathcal{M}(n)$, and each edge (n, m) leaving an OR-node is labelled by an event model $\mathcal{E}(n, m)$.

⁷Note that we here use the induction hypothesis for the autobisimulation on \mathcal{M}' linking u' and u'' , not the bisimulation \mathcal{R} between \mathcal{M} and \mathcal{M}' .

Planning trees for planning problems $\mathcal{P} = (\mathcal{M}_0, \mathbf{A}, \phi_g)$ are constructed as follows: Let the initial planning tree T_0 consist of just one OR-node $root(T_0)$ with $\mathcal{M}(root(T_0)) = \mathcal{M}_0$ (the root labels the initial state). A planning tree for \mathcal{P} is then any tree that can be constructed from T_0 by repeated applications of the following non-deterministic tree expansion rule.

Definition 27 (Tree Expansion Rule). Let T be a planning tree for a planning problem $\mathcal{P} = (\mathcal{M}_0, \mathbf{A}, \phi_g)$. The tree expansion rule is defined as follows. Pick an OR-node n in T and an event model $\mathcal{E} \in \mathbf{A}$ applicable in $\mathcal{M}(n)$ with the proviso that \mathcal{E} does not label any existing outgoing edges from n . Then:

1. Add a new AND-node m to T with $\mathcal{M}(m) = \mathcal{M}(n) \otimes \mathcal{E}$, and add an edge (n, m) with $\mathcal{E}(n, m) = \mathcal{E}$.
2. For each information cell \mathcal{M}' in $\mathcal{M}(m)$, add an OR-node m' with $\mathcal{M}(m') = \mathcal{M}'$ and add the edge (m, m') .

The tree expansion rule is similar in structure to—and inspired by—the expansion rules used in tableau calculi, e.g. for modal and description logics [12]. Note that the expansion rule applies only to OR-nodes, and that an applicable event model can only be used once at each node.

Considering single-agent planning a two-player game, a useful analogy for planning trees are game trees. At an OR-node n , the agent gets to pick any applicable action \mathcal{E} it pleases, winning if it ever reaches an information model in which the goal formula holds (see the definition of solved nodes further below). At an AND-node m , the environment responds by picking one of the information cells of $\mathcal{M}(m)$ —which of the distinguishable outcomes is realised when performing the action.

Without restrictions on the tree expansion rule, even very simple planning problems might be infinitely expanded (e.g. by repeatedly choosing a no-op action). Finiteness of trees (and therefore termination) is ensured by the following blocking condition.

B The tree expansion rule may not be applied to an OR-node n for which there exists an ancestor OR-node m with $\mathcal{M}(m) \equiv \mathcal{M}(n)$.⁸

Lemma 28 (Termination). *Any planning tree built by repeated application of the tree expansion rule under condition **B** is finite.*

Proof. Planning trees built by repeated application of the tree expansion rule are finitely branching: the action library is finite, and every plausibility model has only finitely many information cells (the initial state and all event models in the action library are assumed to be finite, and taking the product update of a finite information cell with a finite event model always produces a finite result). Furthermore, condition **B** ensures that no branch has infinite length: there only exists finitely many modally inequivalent information cells over any language $L(P)$ with finite P (Corollary 25). König’s Lemma now implies finiteness of the planning tree. \square

Example 29. Let’s consider a planning tree in relation to our basement scenario (cf. Example 19). Here the planning problem is $\mathcal{P}_B = (\mathcal{M}_0, \{\text{flick}, \text{desc}\}, \phi_g)$ with \mathcal{M}_0 , flick and desc being defined in Figure 4 and $\phi_g = \neg t \wedge u$. We have illustrated the planning tree T in Figure 7. The root n_0 is an OR-node (representing the initial state \mathcal{M}_0), to

⁸Modal equivalence between information cells can be decided by taking their respective bisimulation contractions and then compare for isomorphism, cf. Section 4.1.

- If $\alpha = w$ then at least one child of n is α -solved.
- If $\alpha = sp$ then all most plausible children of n are α -solved.
- If $\alpha = wp$ then at least one of the most plausible children of n is α -solved.

Let T denote any planning tree for a planning problem $\mathcal{P} = (\mathcal{M}_0, \mathbf{A}, \phi_g)$. Below we show that when an OR-node n of T is α -solved, it is possible to construct an α -solution to the planning problem $(\mathcal{M}(n), \mathbf{A}, \phi_g)$. In particular, if the root node is α -solved, an α -solution to \mathcal{P} can be constructed. As it is never necessary to expand an α -solved node, nor any of its descendants, we can augment the blocking condition \mathcal{B} in the following way (parameterised by α where α is one of s , w , sp or wp).

\mathcal{B}_α The tree expansion rule may not be applied to an OR-node n if one of the following holds: 1) n is α -solved; 2) n has an α -solved ancestor; 3) n has an ancestor OR-node m with $\mathcal{M}(m) \equiv \mathcal{M}(n)$.

A planning tree that has been built according to \mathcal{B}_α is called an α -*planning tree*. Since \mathcal{B}_α is more strict than \mathcal{B} , Lemma 28 immediately gives finiteness of α -planning trees—and hence termination of any algorithm building such trees by repeated application of the tree expansion rule. Note that a consequence of \mathcal{B}_α is that in any α -planning tree an α -solved OR-node is either a leaf or has exactly one α -solved child. We make use of this in the following definition.

Definition 31 (Plans for Solved Nodes). Let T be any α -planning tree for $\mathcal{P} = (\mathcal{M}_0, \mathbf{A}, \phi_g)$. For each α -solved node n in T , a plan $\pi(n)$ is defined recursively by:

- if $\mathcal{M}(n) \models \phi_g$, then $\pi(n) = \text{skip}$.
- if n is an OR-node and m its α -solved child, then $\pi(n) = \mathcal{E}(n, m); \pi(m)$.
- if n is an AND-node and m_1, \dots, m_k its α -solved children, then
 - If $k = 1$ then $\pi(n) = \pi(m_1)$.
 - If $k > 1$ then for all $i = 1, \dots, k$ let δ_{m_i} denote a formula true in $\mathcal{M}(m_i)$ but not in any of the $\mathcal{M}(m_j) \not\equiv \mathcal{M}(m_i)$ and let $\pi(n) =$
if δ_{m_1} then $\pi(m_1)$ else if δ_{m_2} then $\pi(m_2)$ else \dots if δ_{m_k} then $\pi(m_k)$.

Note that the plan $\pi(n)$ of a α -solved node n is only uniquely defined up to the choice of δ -formulas in the if-then-else construct. This ambiguity in the definition of $\pi(n)$ will not cause any troubles in what follows, as it only depends on formulas satisfying the stated property. We need, however, to be sure that such formulas always exist and can be computed. To prove this, assume n is an AND-node and m_1, \dots, m_k its α -solved children. Choose $i \in \{1, \dots, k\}$, and let m_{n_1}, \dots, m_{n_l} denote the subsequence of m_1, \dots, m_k for which $\mathcal{M}(m_{n_j}) \not\equiv \mathcal{M}(m_i)$. We need to prove the existence of a formula δ_{m_i} such that $\mathcal{M}(m_i) \models \delta_{m_i}$ but $\mathcal{M}(m_{n_j}) \not\models \delta_{m_i}$ for all $j = 1, \dots, l$. Since $\mathcal{M}(m_{n_j}) \not\equiv \mathcal{M}(m_i)$ for all $j = 1, \dots, l$, there exists formulas δ_j such that $\mathcal{M}(m_i) \models \delta_j$ but $\mathcal{M}(m_{n_j}) \not\models \delta_j$. We then get that $\delta_1 \wedge \delta_2 \wedge \dots \wedge \delta_l$ is true in $\mathcal{M}(m_i)$ but none of the $\mathcal{M}(m_{n_j})$. Such formulas can definitely be computed, either by brute force search through all formulas ordered by length or more efficiently and systematically by using characterising formulas as in [2] (however, characterising formulas for the present formalism are considerably more complex than in the purely epistemic framework of the cited paper).

Let n be a node of a planning tree T . We say that n is *solved* if it is α -solved for some α . If n is s -solved then it is also sp -solved, if sp -solved then wp -solved, and if

wp -solved then w -solved. This gives a natural ordering $s > sp > wp > w$. Note the relation to Lemma 20. We say that a solved node n has *strength* α , if it is α -solved but not β -solved for any $\beta > \alpha$, using the aforementioned ordering.

Example 32. Consider again the planning tree T in Figure 7 for the planning problem $\mathcal{P}_B = (\mathcal{M}_0, \{\text{flick}, \text{desc}\}, \phi_g)$ with $\phi_g = \neg t \wedge u$. Each solved node has been labelled by its strength. The reader is encouraged to check that each node has been labelled correctly according to Definition 30. The leafs satisfying the goal formula ϕ_g have strength s , by definition. The strength of the root node is sp , as its uppermost child has strength sp . The reason this child has strength sp is that *its* most plausible child has strength s .

We see that T is an sp -planning tree, as it is possible to achieve T from n_0 by applying tree expansions in an order that respects \mathcal{B}_{sp} . However, it is not the smallest sp -planning tree for the problem, as e.g. the lower subtree is not required for n_0 to be sp -solved. Moreover, T is *not* a w -planning tree, as \mathcal{B}_w would have blocked further expansion once either of the three solved leafs were expanded.

In our soundness result below, we show that plans of α -solved roots are always α -solutions to their corresponding planning problems. Applying Definition 31 to the sp -planning tree T gives an sp -solution to the basement planning problem, viz. $\pi(n_0) = \text{flick}; \text{desc}; \text{skip}$. This is the solution we referred to as the *best* in Example 1: Assuming all actions result in their most plausible outcomes, the best plan is to flick the switch and then descend. After having executed the first action of the plan, flick, the agent will know whether the bulb is broken or not. This is signified by the two distinct information cells resulting from the flick action, see Figure 7. An agent capable of replanning could thus choose to revise her plan and/or goal if the bulb turns out to be broken.

Theorem 33 (Soundness). *Let α be one of s, w, sp or wp . Let T be an α -planning tree for a problem $\mathcal{P} = (\mathcal{M}_0, \mathbf{A}, \phi_g)$ such that $\text{root}(T)$ is α -solved. Then $\pi(\text{root}(T))$ is an α -solution to \mathcal{P} .*

Proof. We need to prove that $\pi(\text{root}(T))$ is an α -solution to \mathcal{P} , that is, $\mathcal{M}_0 \models [\pi(\text{root}(T))]_{\alpha} \phi_g$. Since \mathcal{M}_0 is the label of the root, this can be restated as $\mathcal{M}(\text{root}(T)) \models [\pi(\text{root}(T))]_{\alpha} \phi_g$. To prove this fact, we will prove the following stronger claim:

For each α -solved OR-node n in T , $\mathcal{M}(n) \models [\pi(n)]_{\alpha} \phi_g$.

We prove this by induction on the height of n . The base case is when n is a leaf (height 0). Since n is α -solved, we must have $\mathcal{M}(n) \models \phi_g$. In this case $\pi(n) = \text{skip}$. From $\mathcal{M}(n) \models \phi_g$ we can conclude $\mathcal{M}(n) \models [\text{skip}]_{\alpha} \phi_g$, that is, $\mathcal{M}(n) \models [\pi(n)]_{\alpha} \phi_g$. This covers the base case. For the induction step, let n be an arbitrary α -solved OR-node n of height $h > 0$. Let m denote the α -solved child of n , and m_1, \dots, m_l denote the children of m . Let m_{n_1}, \dots, m_{n_k} denote the subsequence of m_1, \dots, m_l consisting of the α -solved children of m . Then, by Definition 31,

- If $k = 1$ then $\pi(n) = \mathcal{E}(n, m); \pi(m_{n_1})$.
- If $k > 1$ then $\pi(n) = \mathcal{E}(n, m); \pi(m)$ where $\pi(m) =$
if $\delta_{m_{n_1}}$ then $\pi(m_{n_1})$ else if $\delta_{m_{n_2}}$ then $\pi(m_{n_2})$ else \dots if $\delta_{m_{n_k}}$ then $\pi(m_{n_k})$.

We here consider only the (more complex) case $k > 1$. Our goal is to prove $\mathcal{M}(n) \models [\pi(n)]_{\alpha} \phi_g$, that is, $\mathcal{M}(n) \models [\mathcal{E}(n, m); \pi(m)]_{\alpha} \phi_g$. By the induction hypothesis we have $\mathcal{M}(m_{n_i}) \models [\pi(m_{n_i})]_{\alpha} \phi_g$ for all $i = 1, \dots, k$ (the m_{n_i} are of lower height than n).

Claim 1. $\mathcal{M}(m_{n_i}) \models [\pi(m)]_\alpha \phi_g$ for all $i = 1, \dots, k$.

Proof of claim. Let i be given. We need to prove

$$\mathcal{M}(m_{n_i}) \models \left[\text{if } \delta_{m_{n_1}} \text{ then } \pi(m_{n_1}) \text{ else } \dots \text{ if } \delta_{m_{n_k}} \text{ then } \pi(m_{n_k}) \right]_\alpha \phi_g.$$

Note that by using item 5 of Lemma 17 it suffices to prove that for all $j = 1, \dots, k$,

$$\mathcal{M}(m_{n_i}) \models \delta_{m_{n_j}} \text{ implies } \mathcal{M}(m_{n_i}) \models [\pi(m_{n_j})]_\alpha \phi_g. \quad (3)$$

Let $j \in \{1, \dots, k\}$ be chosen arbitrarily. Assume first $j = i$. By induction hypothesis we have $\mathcal{M}(m_{n_j}) \models [\pi(m_{n_j})]_\alpha \phi_g$, and hence $\mathcal{M}(m_{n_i}) \models [\pi(m_{n_j})]_\alpha \phi_g$. From this (3) immediately follows. Assume now $j \neq i$. By the construction of the δ -formulas, either $\mathcal{M}(m_{n_j}) \equiv \mathcal{M}(m_{n_i})$ or $\mathcal{M}(m_{n_i}) \not\models \delta_{m_{n_j}}$. In the latter case, (3) holds trivially.

In case of $\mathcal{M}(m_{n_j}) \equiv \mathcal{M}(m_{n_i})$ we immediately get $\mathcal{M}(m_{n_i}) \models [\pi(m_{n_j})]_\alpha \phi_g$, since by induction hypothesis we have $\mathcal{M}(m_{n_j}) \models [\pi(m_{n_j})]_\alpha \phi_g$. This concludes the proof of the claim.

Note that by definition of the tree expansion rule (Definition 27), $\mathcal{M}(m_1), \dots, \mathcal{M}(m_l)$ are the information cells in $\mathcal{M}(m)$.

Claim 2. *The following holds:*

- If $\alpha = s$ (w), then for every (some) information cell \mathcal{M}' in $\mathcal{M}(m)$: $\mathcal{M}' \models [\pi(m)]_\alpha \phi_g$.
- If $\alpha = sp$ (wp), then for every (some) most plausible information cell \mathcal{M}' in $\mathcal{M}(m)$: $\mathcal{M}' \models [\pi(m)]_\alpha \phi_g$.

Proof of claim. We only consider the most complex cases, $\alpha = sp$ and $\alpha = wp$. First consider $\alpha = sp$. Let \mathcal{M}' be a most plausible information cell in $\mathcal{M}(m)$. We need to prove $\mathcal{M}' \models [\pi(m)]_\alpha \phi_g$. Since, as noted above, $\mathcal{M}(m_1), \dots, \mathcal{M}(m_l)$ are the information cells in $\mathcal{M}(m)$, we must have $\mathcal{M}' = \mathcal{M}(m_i)$ for some $i \in \{1, \dots, l\}$. Furthermore, as \mathcal{M}' is among the most plausible information cells in $\mathcal{M}(m)$, m_i must by definition be a most plausible child of m . Definition 30 then gives us that m_i is α -solved. Thus $m_i = m_{n_j}$ for some $j \in \{1, \dots, k\}$. By Claim 1 we have $\mathcal{M}(m_{n_j}) \models [\pi(m)]_\alpha \phi_g$, and since $\mathcal{M}' = \mathcal{M}(m_i) = \mathcal{M}(m_{n_j})$ this gives the desired conclusion. Now consider the case $\alpha = wp$. Definition 30 gives us that at least one of the most plausible children of m are α -solved. By definition, this must be one of the m_{n_i} , $i \in \{1, \dots, k\}$. Claim 1 gives $\mathcal{M}(m_{n_i}) \models [\pi(m)]_\alpha \phi_g$. Since m_{n_i} is a most plausible child of m , we must have that $\mathcal{M}(m_{n_i})$ is among the most plausible information cells in $\mathcal{M}(m)$. Hence we have proven that $[\pi(m)]_\alpha \phi_g$ holds in a most plausible information cell of $\mathcal{M}(m)$.

By definition of the tree expansion rule (Definition 27), $\mathcal{M}(m) = \mathcal{M}(n) \otimes \mathcal{E}(n, m)$. Thus we can replace $\mathcal{M}(m)$ by $\mathcal{M}(n) \otimes \mathcal{E}(n, m)$ in Claim 2 above. Using items 1–4 of Lemma 17, we immediately get from Claim 2 that independently of α the following holds: $\mathcal{M}(n) \models [\mathcal{E}(n, m)]_\alpha [\pi(m)]_\alpha \phi_g$ (the condition $\mathcal{M}(n) \models \langle \mathcal{E}(n, m) \rangle \top$ holds trivially by the tree expansion rule). From this we can then finally conclude $\mathcal{M}(n) \models [\mathcal{E}(n, m); \pi(m)]_\alpha \phi_g$, as required. \square

Theorem 34 (Completeness). *Let α be one of s , w , sp or wp . If there is an α -solution to the planning problem $\mathcal{P} = (\mathcal{M}_0, A, \phi_g)$, then an α -planning tree T for \mathcal{P} can be constructed, such that $\text{root}(T)$ is α -solved.*

Proof. First note that we have $[\text{skip}; \pi]_\alpha \phi_g = [\text{skip}]_\alpha ([\pi]_\alpha \phi_g) = [\pi]_\alpha \phi_g$. Thus, we can without loss of generality assume that no plan contains a subexpression of the form $\text{skip}; \pi$. The length of a plan π , denoted $|\pi|$, is defined recursively by: $|\text{skip}| = 1$; $|\mathcal{E}| = 1$; $|\text{if } \phi \text{ then } \pi_1 \text{ else } \pi_2| = |\pi_1| + |\pi_2|$; $|\pi_1; \pi_2| = |\pi_1| + |\pi_2|$.

Claim 1. *Let π be an α -solution to $\mathcal{P} = (\mathcal{M}_0, \mathbf{A}, \phi_g)$ with $|\pi| \geq 2$. Then there exists an α -solution of the form $\mathcal{E}; \pi'$ with $|\mathcal{E}; \pi'| \leq |\pi|$.*

Proof of claim. Proof by induction on $|\pi|$. The base case is $|\pi| = 2$. We have two cases, $\pi = \text{if } \phi \text{ then } \pi_1 \text{ else } \pi_2$ and $\pi = \pi_1; \pi_2$, both with $|\pi_1| = |\pi_2| = 1$. If π is the latter, it already has desired the form. If $\pi = \text{if } \phi \text{ then } \pi_1 \text{ else } \pi_2$ then, by assumption on π , $\mathcal{M}_0 \models [\text{if } \phi \text{ then } \pi_1 \text{ else } \pi_2]_\alpha \phi_g$. Item 5 of Lemma 17 now gives that $\mathcal{M}_0 \models \phi$ implies $\mathcal{M}_0 \models [\pi_1]_\alpha \phi_g$ and $\mathcal{M}_0 \not\models \phi$ implies $\mathcal{M}_0 \models [\pi_2]_\alpha \phi_g$. Thus we must either have $\mathcal{M}_0 \models [\pi_1]_\alpha \phi_g$ or $\mathcal{M}_0 \models [\pi_2]_\alpha \phi_g$, that is, either π_1 or π_2 is an α -solution to \mathcal{P} . Thus either $\pi_1; \text{skip}$ or $\pi_2; \text{skip}$ is an α -solution to \mathcal{P} , and both of these have length $|\pi|$. This completes the base case. For the induction step, consider a plan π of length $l > 2$ which is an α -solution to \mathcal{P} . We again have two cases to consider, $\pi = \text{if } \phi \text{ then } \pi_1 \text{ else } \pi_2$ and $\pi = \pi_1; \pi_2$. If $\pi = \pi_1; \pi_2$ is an α -solution to \mathcal{P} , then π_1 is an α -solution to the planning problem $\mathcal{P}' = (\mathcal{M}_0, \mathbf{A}, [\pi_2]_\alpha \phi_g)$, as $\mathcal{M}_0 \models [\pi_1; \pi_2]_\alpha \phi_g \Leftrightarrow \mathcal{M}_0 \models [\pi_1]_\alpha [\pi_2]_\alpha \phi_g$. Clearly $|\pi_1| < l$, so the induction hypothesis gives that there is an α -solution $(\mathcal{E}; \pi'_1)$ to \mathcal{P}' , with $|\mathcal{E}; \pi'_1| \leq |\pi_1|$. Then, $\mathcal{E}; \pi'_1; \pi_2$ is an α -solution to \mathcal{P} and we have $|\mathcal{E}; \pi'_1; \pi_2| = |\mathcal{E}; \pi'_1| + |\pi_2| \leq |\pi_1| + |\pi_2| = |\pi|$. If $\pi = \text{if } \phi \text{ then } \pi_1 \text{ else } \pi_2$ is an α -solution to \mathcal{P} , then we can as above conclude that either π_1 or π_2 is an α -solution to \mathcal{P} . With both $|\pi_1| < l$ and $|\pi_2| < l$, the induction hypothesis gives the existence an α -solution $\mathcal{E}; \pi'$, with $|\mathcal{E}; \pi'| \leq |\pi|$. This completes the proof of the claim.

We now prove the theorem by induction on $|\pi|$, where π is an α -solution to $\mathcal{P} = (\mathcal{M}_0, \mathbf{A}, \phi_g)$. We need to prove that there exists an α -planning tree for \mathcal{P} in which the root is α -solved. Let T_0 denote the planning tree for \mathcal{P} only consisting of its root node with label \mathcal{M}_0 . The base case is when $|\pi| = 1$. Here, we have two cases, $\pi = \text{skip}$ and $\pi = \mathcal{E}$. In the first case, the planning tree T_0 already has its root α -solved, since $\mathcal{M}_0 \models [\text{skip}]_\alpha \phi_g \Leftrightarrow \mathcal{M}_0 \models \phi_g$. In the second case, $\pi = \mathcal{E}$, we have $\mathcal{M}_0 \models [\mathcal{E}]_\alpha \phi_g$ as $\pi = \mathcal{E}$ is an α -solution to \mathcal{P} . By definition, this means that \mathcal{E} is applicable in \mathcal{M}_0 , and we can apply the tree expansion rule to T_0 , which will produce:

1. A child m of the root node with $\mathcal{M}(m) = \mathcal{M}_0 \otimes \mathcal{E}$.
2. Children m_1, \dots, m_l of m , where $\mathcal{M}(m_1), \dots, \mathcal{M}(m_l)$ are the information cells of $\mathcal{M}(m)$.

Call the expanded tree T_1 . Since $\mathcal{M}_0 \models [\mathcal{E}]_\alpha \phi_g$, Lemma 17 implies that for every/some/every most plausible/some most plausible information cell \mathcal{M}' in $\mathcal{M}_0 \otimes \mathcal{E}$, $\mathcal{M}' \models \phi_g$ (where $\alpha = s/w/sp/wp$). Since $\mathcal{M}(m_1), \dots, \mathcal{M}(m_l)$ are the information cells of $\mathcal{M}_0 \otimes \mathcal{E}$, we can conclude that every/some/every most plausible/some most plausible child of m is α -solved. Hence also m and thus n are α -solved. The base is hereby completed.

For the induction step, let π be an α -solution to \mathcal{P} with length $l > 1$. Let T_0 denote the planning tree for \mathcal{P} consisting only of its root node with label \mathcal{M}_0 . By Claim 1, there exists an α -solution to \mathcal{P} of the form $\mathcal{E}; \pi'$ with $|\mathcal{E}; \pi'| \leq |\pi|$. As $\mathcal{M}_0 \models [\mathcal{E}; \pi']_\alpha \phi_g \Leftrightarrow \mathcal{M}_0 \models [\mathcal{E}]_\alpha [\pi']_\alpha \phi_g$, \mathcal{E} is applicable in \mathcal{M}_0 . Thus, as in the base case, we can apply the tree expansion rule to T_0 which will produce nodes as in 1 and 2 above. Call the expanded tree T_1 . Since $\mathcal{M}_0 \models [\mathcal{E}]_\alpha [\pi']_\alpha \phi_g$, items 1–4 of Lemma 17

implies that for every/some/every most plausible/some most plausible information cell in $\mathcal{M}_0 \otimes \mathcal{E}$, $[\pi']_\alpha \phi_g$ holds. Hence, for every/some/every most plausible/some most plausible child m_i of m , $\mathcal{M}(m_i) \models [\pi']_\alpha \phi_g$. Let m_{n_1}, \dots, m_{n_k} denote the subsequence of m_1, \dots, m_l consisting of the children of m for which $\mathcal{M}(m_{n_i}) \models [\pi']_\alpha \phi_g$. Then, by definition, π' is an α -solution to each of the planning problem $\mathcal{P}_i = (\mathcal{M}(m_{n_i}), \mathbf{A}, \phi_g)$, $i = 1, \dots, k$. As $|\pi'| < |\mathcal{E}; \pi'| \leq l$, the induction hypothesis gives that α -planning trees T'_i with α -solved roots can be constructed for each \mathcal{P}_i . Let T_2 denote T_1 expanded by adding each planning tree T'_i as the subtree rooted at \mathcal{M}_{n_i} . Then each of the nodes m_{n_i} are α -solved in T , and in turn both m and $\text{root}(T_2)$ are α -solved. The final thing we need to check is that T_2 has been correctly constructed according to the tree expansion rule, more precisely, that condition \mathcal{B}_α has not been violated. Since each T'_i has in itself been correctly constructed in accordance with \mathcal{B}_α , the condition can only have been violated if for one of the non-leaf OR-nodes m' in one of the T'_i s, $\mathcal{M}(m') \equiv \mathcal{M}(\text{root}(T_2))$. We can then replace the entire planning tree T_2 by a (node-wise modally equivalent) copy of the subtree rooted at m' , and we would again have an α -planning tree with an α -solved root. \square

4.3 Planning Algorithm

In the following, let \mathcal{P} denote any planning problem, and α be one of s , w , sp or wp . With all the previous in place, we now have an algorithm for synthesising an α -solution to \mathcal{P} , given as follows.

PLAN(α, \mathcal{P})

- 1 Let T be the α -planning tree only consisting of $\text{root}(T)$ labelled by the initial state of \mathcal{P} .
- 2 Repeatedly apply the tree expansion rule of \mathcal{P} to T until no more rules apply satisfying condition \mathcal{B}_α .
- 3 If $\text{root}(T)$ is α -solved, return $\pi(\text{root}(T))$, otherwise return FAIL.

Theorem 35. PLAN(α, \mathcal{P}) is a terminating, sound and complete algorithm for producing α -solutions to planning problems \mathcal{P} . Soundness means that if PLAN(α, \mathcal{P}) returns a plan, it is an α -solution to \mathcal{P} . Completeness means that if \mathcal{P} has an α -solution, PLAN(α, \mathcal{P}) will return one.

Proof. Termination comes from Lemma 28 (with \mathcal{B} replaced by the stronger condition \mathcal{B}_α), soundness from Theorem 33 and completeness from Theorem 34 (given any two \mathcal{B}_α -saturated α -planning trees T_1 and T_2 for the same planning problem, the root node of T_1 is α -solved iff the root node of T_2 is). \square

With PLAN(α, \mathcal{P}) we have given an algorithm for solving α -parametrised planning problems. The α parameter determines the strength of the synthesised plan π , cf. Lemma 20. Whereas the cases of weak ($\alpha = w$) and strong ($\alpha = s$) plans have been the subject of much research, the generation of weak plausibility ($\alpha = wp$) and strong plausibility ($\alpha = sp$) plans based on pre-encoded beliefs is a novelty of this paper. Plans taking plausibility into consideration have several advantages. Conceptually, the basement scenario as formalised by \mathcal{P}_B (cf. Example 19) allowed for several weak solutions (with the shortest one being hazardous to the agent) and *no* strong solutions. In this case, the synthesised strong plausibility solution corresponds to the course of action a rational agent (mindful of her beliefs) should take. There are also computational advantages. An invocation of PLAN(sp, \mathcal{P}) will expand at most as many nodes

as an invocation of $\text{PLAN}(s, \mathcal{P})$ before returning a result (assuming the same order of tree expansions). As plausibility plans only consider the most plausible information cells, we can prune non-minimal information cells during plan search.

We also envision using this technique in the context of an agent framework where planning, acting and execution monitoring are interleaved.⁹ Let us consider the case of strong plausibility planning ($\alpha = sp$). From some initial situation an *sp*-plan is synthesised which the agent starts executing. If reaching a situation that is not covered by the plan, she restarts the process from this point; i.e. she *replans*. Note that the information cell to replan from is present in the tree as a sibling of the most plausible information cell(s) expected from executing the last action. Such replanning mechanisms allow for the *repetition* of actions necessary in some planning problems with cyclic solutions.

We return one last time to the basement problem and consider a modified *replace* action such that the replacement light bulb might, though it is unlikely, be broken. This means that there is no strong solution. Executing the *sp*-solution *flick; desc*, she would replan after *flick* if that action didn't have the effect of turning on the light. A strong plausibility solution from this point would then be *flick; replace; flick; desc*.

5 Related and Future Work

In this paper we have presented α -solutions to planning problems incorporating ontic, epistemic and doxastic notions. The cases of $\alpha = sp/sw$ are, insofar as we are aware, novel concepts not found elsewhere in the literature. Our previous paper [2] concerns the cases $\alpha = s/w$, so that framework deals only with *epistemic* planning problems without a doxastic component. Whereas we characterise solutions as formulas, [1] takes a semantic approach to strong solutions for epistemic planning problems. In their work plans are sequences of actions, requiring conditional choice of actions at different states to be encoded in the action structure itself. By using the $\mathcal{L}(P, A)$ we represent this choice explicitly.

The meaningful plans of [9, chap. 2] are reminiscent of the work in this paper. Therein, plan verification is cast as validity of an EDL-consequence in a given system description. Like us, they consider single-agent scenarios, conditional plans, applicability and incomplete knowledge in the initial state. Unlike us, they consider only deterministic epistemic actions (without plausibility). In the multi-agent treatment [9, chap. 4], action laws are translated to a fragment of DEL with only public announcements and public assignments, making actions singleton event models. This means foregoing nondeterminism and therefore sensing actions.

Epistemic planning problems in [15] are solved by producing a sequence of pointed epistemic event models where an external variant of applicability (called *possible at*) is used. Using such a formulation means outcomes of actions are fully determined, making conditional plans and weak solutions superfluous. As noted by the authors, and unlike our framework, their approach does not consider factual change. We stress that [1, 15, 9] all consider the multi-agent setting which we have not treated here.

In our work so far, we haven't treated the problem of where domain formulations come from, assuming just that they are given. Standardised description languages are vital if modal logic-based planning is to gain wide acceptance in the planning community. Recent work worth noting in this area includes [7], which presents a specification language for the multi-agent belief case.

⁹Covering even more mechanisms of agency is *situated planning* [11].

As suggested by our construction of planning trees, there are several connections between our approach for $\alpha = s$ and two-player imperfect information games. First, product updates imply perfect recall [18]. Second, when the game is at a node belonging to an information set, the agent knows a proposition only if it holds throughout the information set. Finally, the strong solutions we synthesise are very similar to mixed strategies. A strong solution caters to any information cell (contingency) it may bring about, by selecting exactly one sub-plan for each [3].

Our work relates to [11], where the notions of strong and weak solutions are found, but without plausibilities. Their belief states are sets of states which may be partitioned by observation variables. The framework in [16] describes strong conditional planning (prompted by nondeterministic actions) with partial observability modelled using a fixed set of observable state variables. Our partition of plausibility models into information cells follows straight from the definition of product update. A clear advantage in our approach is that *actions* readily encode both nondeterminism and partial observability. [14] shows that the *strong plan existence problem* for the framework in [2] is 2-EXP-complete. In our formulation, $\text{PLAN}(s, \mathcal{P})$ answers the same question for \mathcal{P} (it gives a strong solution if one exists), though with a richer modal language.

We would like to do plan verification and synthesis in the multi-agent setting. We believe that generalising the notions introduced in this paper to multi-pointed plausibility and event models are key. Plan synthesis in the multi-agent setting is undecidable [1], but considering restricted classes of actions as is done in [15] seems a viable route for achieving decidable multi-agent planning. Other ideas for future work include replanning algorithms and learning algorithms where plausibilities of actions can be updated when these turn out to have different outcomes than expected.

Acknowledgements

For valuable comments at various stages of the work presented in this article, we would like to extend our gratitude to the following persons: Patrick Allo, Alexandru Baltag, Johan van Benthem, Hans van Ditmarsch, Jens Ulrik Hansen, Sonja Smets and the anonymous reviewers.

References

- [1] ANDERSEN, M. B., AND BOLANDER, T. Epistemic planning for single- and multi-agent systems. *Journal of Applied Non-Classical Logics* 21, 1 (2011), 9–34.
- [2] ANDERSEN, M. B., BOLANDER, T., AND JENSEN, M. H. Conditional epistemic planning. In *JELIA (2012)*, L. F. del Cerro, A. Herzig, and J. Mengin, Eds., Lecture Notes in Computer Science, Springer.
- [3] AUMANN, R. J., AND HART, S., Eds. *Handbook of Game Theory with Economic Applications*. Elsevier, 1992.
- [4] BALTAG, A., AND MOSS, L. S. Logics for Epistemic Programs. *Synthese* 139 (2004), 165–224.
- [5] BALTAG, A., AND SMETS, S. Dynamic belief revision over multi-agent plausibility models. In *G. Bonanno, W. van der Hoek, M. Woolridge (eds.), Proceedings of*

- the 7th Conference on Logic and the Foundations of Game and Decision (LOFT 2006)* (2006), pp. 11–24.
- [6] BALTAG, A., AND SMETS, S. A qualitative theory of dynamic interactive belief revision. In *Logic and the Foundation of Game and Decision Theory (LOFT7)* (2008), G. Bonanno, W. van der Hoek, and M. Wooldridge, Eds., vol. 3 of *Texts in Logic and Games*, Amsterdam University Press, pp. 13–60.
- [7] BARAL, C., GELFOND, G., PONTELLI, E., AND SON, T. C. An action language for reasoning about beliefs in multi-agent domains. In *Proceedings of the 14th International Workshop on Non-Monotonic Reasoning* (2012).
- [8] BLACKBURN, P., AND VAN BENTHEM, J. Modal logic: A semantic perspective. In *Handbook of Modal Logic* (2006), Elsevier.
- [9] DE LIMA, T. *Optimal Methods for Reasoning about Actions and Plans in Multi-Agents Systems*. PhD thesis, IRIT, University of Toulouse 3, France, 2007.
- [10] DEMEY, L. Some remarks on the model theory of epistemic plausibility models. *Journal of Applied Non-Classical Logics* 21, 3-4 (2011), 375–395.
- [11] GHALLAB, M., NAU, D. S., AND TRAVERSO, P. *Automated Planning: Theory and Practice*. Morgan Kaufmann, 2004.
- [12] HORROCKS, I., HUSTADT, U., SATTLER, U., AND SCHMIDT, R. Computational modal logic. In *Handbook of Modal Logic* (2006), Elsevier.
- [13] JAMROGA, W., AND ÅGOTNES, T. Constructive knowledge: what agents can achieve under imperfect information. *Journal of Applied Non-Classical Logics* 17, 4 (2007), 423–475.
- [14] JENSEN, M. H. The computational complexity of single agent epistemic planning (manuscript).
- [15] LÖWE, B., PACUIT, E., AND WITZEL, A. DEL planning and some tractable cases. *Lecture Notes in Computer Science* 6953 (2011).
- [16] RINTANEN, J. Complexity of planning with partial observability. In *ICAPS* (2004), S. Zilberstein, J. Koehler, and S. Koenig, Eds., AAAI, pp. 345–354.
- [17] VAN BENTHEM, J. Dynamic odds and ends. Technical Report ML-1998-08, University of Amsterdam, 1998.
- [18] VAN BENTHEM, J. Games in dynamic-epistemic logic. *Bulletin of Economic Research* 53, 4 (2001), 219–48.
- [19] VAN DITMARSCH, H. Prolegomena to dynamic logic for belief revision. *Synthese* 147 (2005), 229–275.
- [20] VAN DITMARSCH, H. Revocable belief revision. *Studia Logica* (2013, to appear).
- [21] VAN DITMARSCH, H., AND KOOI, B. Semantic results for ontic and epistemic change. In *LOFT 7* (2008), Amsterdam University Press, pp. 87–117.