To appear in Artificial Intelligence
(Elsevier journal ISSN: 0004-3702, www.journals.elsevier.com/artificial-intelligence)
accepted for publication on 17 June 2020

Probabilistic Reasoning About Epistemic Action Narratives

Fabio Aurelio D'Asaro, Antonis Bikakis, Luke Dickens, Rob Miller[1]

Department of Information Studies, University College London

Abstract

We propose the action language EPEC – Epistemic Probabilistic Event Calculus – that supports probabilistic, epistemic reasoning about narratives of action occurrences and environmentally triggered events, and in particular facilitates reasoning about future belief-conditioned actions and their consequences in domains that include both perfect and imperfect sensing actions. To provide a declarative semantics for sensing and belief conditioned actions in a probabilistic, narrative setting we introduce the novel concept of an epistemic *reduct*. We then formally compare our language with two established frameworks for probabilistic reasoning about action – the action language PAL by Baral et al, and the extension of the situation calculus to reason about noisy sensors and effectors by Bacchus et al. In both cases we prove a correspondence with EPEC for a class of domains representable in both frameworks.

Keywords: Reasoning about Actions, Epistemic Reasoning, Narrative Reasoning, Probabilistic Reasoning, Conditional Actions, Imperfect Sensing

1. Introduction

The action language EPEC – Epistemic Probabilistic Event Calculus – described in this article combines probabilistic reasoning, epistemic reasoning, reasoning about the general effects of actions, and reasoning about particular narratives of action occurrences (i.e. *events* along an explicitly represented

[1]Authors other than the first author listed alphabetically.

time-line). Each of these topics has its own venerable history of AI research, and notable work has already been done in some sub-combinations. For example, the work of Moore, Scherl, Levesque, Belle, Bacchus and others concerns epistemic (and sometimes also probabilistic) reasoning about actions (see e.g. [1], [2], [3] [4]), the work of Ma et. al. [5] facilitates epistemic reasoning about narratives of events, and the work of Baral, Skarlatidis, Artikis and others focuses on probabilistic reasoning about action narratives (see e.g. [6], [7]). (See Section 4.3 for a full discussion of related work.) However, to our knowledge little or no previous research has been undertaken towards a full integration of all four of these topics, and we aim to demonstrate the benefits of such an integration in this paper. EPEC is related to PEC (Probabilistic Event Calculus) [8], an earlier probabilistic framework for narrative reasoning that did not contain any epistemic features. The utility of our EPEC framework is partly illustrated by the following (imaginary) example medical scenario.

Example 1.1 (Epectisis)**.** A doctor is 95% certain that a patient has made skin contact with her, and 80% certain that this patient is suffering from *epectisis*, a rare disease caused by a bacterial infection. With such contact, epectisis is typically passed on 75% of the time. A course of *epecillin* is known to eliminate the disease 99% of the time if taken before symptoms manifest themselves, although with a 15% risk of side effects. A blood test gives a pre-symptom indication of the disease, but with a 10% false positive and a 5% false negative error rate. The doctor decides that she will undertake the blood test, and if after this she still has a more than 50% belief that she is infected she will take a course of epecillin. She reasons that, assuming that she did not have epectitis prior to contact with her patient, in this way she will eventually be at least 93.1% sure that she does not have the disease, while giving herself only a 2.875% chance of suffering epecillin's occasional side effects. (A probability tree diagram of this domain is given in Appendix B.) [end example]

The scenario above has a number of interesting features, all of which can be represented in our EPEC framework (see Appendix C for the corresponding

EPEC domain description $\mathcal{D}_e$ and example entailments). First, it includes a *narrative*, in this case containing a single probable past event – the doctor is 95% certain that she had skin contact with a patient. In deciding a course of action, the doctor appends two future events to this narrative – performing a blood test and (conditionally) taking epecillin. This reflects an abductive view of plan specification commonplace in the context of event-calculus-like frameworks (see e.g. [9]). Second, some of the *causal information* about actions is probabilistic – in general, contact has a 75% probability of causing infection, and taking epecillin has a 15% probability of causing a side effect. Third, one of the actions mentioned – performing a blood test – is a *sensing action* in that it has an effect on the doctor's (probabilistic) knowledge, so that the doctor's plan has an epistemic dimension. Moreover, the sensing is *imperfect*, with the possibility of false positives and negatives. Fourth, the doctor's plan includes a *conditional action*, conditioned on a future belief state – *if* after performing the blood test she has a strong belief that she is infected *then* she will take the medicine. We view this as a key feature of epistemic planning, in that sensing actions and actions conditioned on (revised) beliefs resulting directly or indirectly from sensing outcomes must go hand-in-hand, or there is little point in including sesning actions within a plan. A principal advantage of EPEC is that it allows for explicit representation and probabilistic reasoning about future belief-conditioned actions and their consequences in domains that include both perfect and imperfect sensing actions. This is a key contribution of our work.

The medical domain is one of a number of application areas where we envisage EPEC having a useful role. Other domains in which A.I. and knowledge-based applications have to reason with imperfect sensor input include cognitive and mobile robotics, and physical monitoring systems. Not all of the representational features of EPEC are illustrated by the example above. Descriptions of more complex domains may for example include non-Boolean-valued properties (*fluents*), concurrent actions, conflicting concurrent noisy sensory inputs, and, importantly, events (probabilistically) *triggered* by the environment under certain conditions. This latter feature in turn allows EPEC to be used to model domains involving *decay* and analogous dynamic behaviours.

To reflect the high degree of non-determinism inherent in probabilistic models, as well as the epistemic nature of EPEC domains, EPEC's semantics uses a structure of "possible worlds", each with its own timeline and overall probability attached. The semantics is developed in two stages. First, the "non-epistemic" case is considered, where all action occurrences are considered to be executed by the environment and are independent from the agent's belief state. Second, the semantics is generalised to the epistemic case, where possible worlds are considered to have two components – a timeline of actual environmental conditions and events, together with history of the agent's sensory experience and decisions regarding its own actions. We show how the notion of an epistemic *reduct* can be used to model this second case in terms of the first, while taking into account the agent's sensory input, changing belief state and associated decision making process as time progresses.

In order to progress understanding of the space of formalisms available for probabilistic reasoning about actions, we conclude with an investigation of the relation of EPEC to two established frameworks in this area of research. These are the action language PAL developed by Baral, Tran and Tuan [6], and the extension of the situation calculus to reason about noisy sensors and effectors by Bacchus, Halpern and Levesque [3]. In both cases we provide a general translation procedure of a class of domains written in these languages into EPEC, and prove that probabilistic entailment is preserved under the translations.

In summary, the main contribution of this paper is the formulation of an action language, EPEC (Epistemic Probabilistic Event Calculus), and associated semantics that supports probabilistic, epistemic reasoning about narratives of (potentially simultaneous) action occurrences and environmentally triggered events, and in particular facilitates reasoning about future belief-conditioned actions and their consequences in domains that include both perfect and imperfect sensing actions, and potentially simultaneous and/or conflicting sensory inputs. Additional contributions are formal comparisons with two established frameworks for probabilistic reasoning about actions by provably correct translations into EPEC.

The paper is organised as follows. Section 2 gives some general background about the topic areas related to this research. Section 3 describes the syntax, semantics and key properties of EPEC. Section 4 describes the translations into EPEC of domains written in the frameworks in [6] and [3], and follows with a wider discussion of related work. In Section 5 we briefly comment on ongoing experiments implementing EPEC. Section 6 concludes the paper with a final summary and remarks about possible future directions of research.

2. Background and Overview of Approach

2.1. Reasoning About Actions and Narratives

Logic-based reasoning about actions as a field of A.I. research was arguably triggered in 1969 by McCarthy and Hayes' proposal for a *Situation Calculus* (SC) [10], with its ontology of situations, actions, and time-varying properties of the world potentially affected by actions called *fluents*. Perhaps the most well-known subsequent formulation of the SC is that of Reiter and his colleagues (see for example [11]). Reiter's *basic action theories* (BATs) support classical-logic reasoning about the consequences of (all) hypothetical sequences of actions embedded in a forward-branching tree of situations. They also provide a principled non-monotonic solution to the *frame problem* – roughly, the problem of succinctly and flexibly expressing that most actions do not affect most fluents under most circumstances (see e.g. [12] for an indepth discussion of this issue).

The basic BAT formulation and core ontology mentioned above do not however support *narrative reasoning* – it is not possible to assert that any particular sequence of actions has (or has not) actually occurred or will (or will not) actually occur, because the independent flow of time is not represented and all possible action sequences are given equal hypothetical status. The *Event Calculus* (EC), with its alternative ontology of actions, fluents and timepoints, was originally developed partly to address this early limitation of the SC[2]. It was

[2]Various extensions to the SC have been subsequently developed to enable narrative rea-

first formulated as a logic program [15], and later in classical logic (see e.g. [9]). Because the EC includes the flow of time in its core ontology, usually as an integer or real-number timeline, EC domain descriptions can incorporate simple predicate assertions that particular actions (or "events") have occurred or will occur at particular times. For example, a non-probabilistic description of Example 1.1 might include the sentence $Happens(SkinContact, -1.0)$, to assert that skin contact occurred an hour ago. Similar assertions using future time-points are used for EC-based abductive planning (see [16]). Depending on the domain or reasoning task, narratives may sometimes take into account events in the environment as well as actions performed by an agent, and this gives rise to another aspect of the frame problem – the problem of succinctly and flexibly expressing that most events or actions do not occur most of the time. This is generally solved in EC frameworks by minimisation (e.g. predicate completion or circumscription) of the *Happens* predicate or its equivalent. The inclusion of narrative capability in the EC in a simple and natural way has made it the obvious ontological starting point for development of EPEC. Indeed, we see EPEC as a natural development from previous EC-inspired work on PEC (Probabilistic Event Calculus [8]), EFEC (Epistemic Functional Event Calculus [5]) and Modular-$\mathcal{E}$ [17]. We do however envisage that near-equivalent representations could be formulated in SC-inspired frameworks if suitably augmented with narrative capabilities.

EPEC is formulated as an *action language*, with its own specialised syntax and semantics, rather than as a classical logic theory or logic program. Gelfond and Lifschitz [18, 19] were the first to propose this particular methodology for research into reasoning about actions, and it has subsequently been employed in the context of the EC in [20] and [17]. The general benefits of the action language approach are detailed in [18, 19, 20, 17], and here we use it in order to highlight the essential properties and features of EPEC in as simple a way

soning, albeit at the expense of the simplicity of the original SC formulation. See for example [13], [14] and [11] Chapter 7.

as possible, and to distinguish between the declarative specification of the class of domains that we wish to model and any particular implementation choice.

2.2. Epistemic Reasoning About Actions

Much of the work on combining formal theories of knowledge with theories of action has its roots in the work of Moore [1], who merged a modal logic of knowledge with the situation calculus. This SC-based approach has subsequently been developed by a number of researchers, and notably in publications by Scherl, Levesque and collaborators [21, 2, 22, 23, 24, 25] addressing various issues related to sensing and acting and appropriately extending Reiter's solution to the frame problem. Modal logic represents knowledge and the lack of it via an accessibility relation between possible worlds – an agent's lack of knowledge about the truth of a proposition p is modelled by making accessible both a possible world in which p holds and a possible world in which $\neg p$ holds. In the SC-based theories of knowledge and action, this accessibility relation is represented as an *epistemic fluent* that connects corresponding situations in different possible worlds, and which is amenable to change via sensing actions.

The possible worlds approach to constructing a theory of knowledge and action has also been applied to the Event Calculus. The *Epistemic Functional Event Calculus* (EFEC) [5] regards possible worlds as alternative narrative timelines (as opposed to alternative situation trees), again connected by an epistemic fluent. By using its narrative ontology, EFEC allows for reasoning about knowledge of the past, present and future, and incorporates the notion illustrated in Example 1.1 of knowledge-conditioned actions embedded along a time-line. For these reasons we have used its ontology as a starting point for the development of EPEC, although EFEC is non-probabilistic and not able to model the related notions of reasoning with uncertainty, degrees of belief or imperfect sensing.

2.3. Reasoning with Uncertainty and Probability Theory

Probability theory has been a mainstream tool for A.I. applied to domains with a degree of uncertainty since the the mid-1980s [26, 27, 28], and in particular has been employed in models of causality [29]. In the present work we follow the Bayesian view of probabilities as representing (justified) degrees of belief, and use this notion of degree of belief in place of the binary logic concept of certain knowledge. EPEC retains the semantic structure of possible worlds employed in the frameworks such as EFEC mentioned in Section 2.2, but relates possible worlds in terms of a *probability distribution* rather than in terms of an accessibility relation. Temporal changes in the agent's degree of belief in past, present and future values of fluents, as a result of (imperfect) sensing actions, are captured using *conditional probabilities*. In the remainer of this section we offer the reader intuitions about how some essential ideas from discrete probability theory, and ideas relating logic and probability, underly our approach.

Probability theory models a discrete domain as a collection of *random variables* (RVs), each of which can take on one of a finite number of values. A *joint probability distribution* over a set of RVs assigns a numerical probability between 0 and 1 to each combination of values for the RVs, in such a way that the probabilities for all combinations sum to 1. As will be seen in Section 3, in the case of EPEC domains, at one granularity the RVs can be taken to be fluent/time-point pairs (each representing the value of the fluent at that time-point), action/time-point pairs (each representing whether or not the action has occurred at that time-point), and sensing outcomes (each representing the fluent value correctly or incorrectly sensed at the point where the relevant sensing action occurs). However, application of the principles of default persistence of fluent values and the closed world assumption for action occurrences (jointly contributing to a solution to the frame problem) allows EPEC's syntax and semantics to be defined at a different granularity, in terms of a finite collection of RVs, by considering only an initial time-point, the time-points at which there is the possibility of an action occurring, and the maximal change-free intervals between such time-points. Dependencies (in the probabilistic sense) between the RVs are indirectly captured by the domain-dependent causal and narrative propositions that comprise an EPEC domain description. We make use of some standard concepts and results from probability theory, including *marginal* and *conditional probability*, the *sum* and *product rules*, and *Bayes' Theorem.*

Definitions and commentary on these can be found for example in [30].

We also make use of the concept of a *probability function* that assigns probabilities to a set of propositional formulas in a consistent way. The definition given here is a minor generalisation of that of Paris [31, Chapter 2, p.10], in that it does not require the entailment relation to which it refers to be with respect to all interpretations of the language (i.e. classical propositional entailment), but only that it is with respect to some non-empty set of interpretations.

Definition 2.1 (Probability Function, Conditional Probability for Formulas)**.** Let $\models_\bullet$ be the entailment relation defined in the standard way with respect to some non-empty set of interpretations of a propositional language L, and let $\mathbf{F}_{set}$ be a set of propositional formulas of L closed under the propositional operators. A *probability function over* $\mathbf{F}_{set}$ *w.r.t.* $\models_\bullet$ is a function $p : \mathbf{F}_{set} \mapsto [0, 1]$ such that for all $\varphi, \psi \in \mathbf{F}_{set}$:

1. if $\models_\bullet \varphi$, then $p(\varphi) = 1$, and
2. if $\varphi \models_\bullet \neg\psi$ then $p(\varphi \vee \psi) = p(\varphi) + p(\psi)$.

For $p(\psi) \neq 0$, the associated *conditional probability* of φ given ψ is defined as

$$p(\varphi \mid \psi) = \frac{p(\varphi \wedge \psi)}{p(\psi)} \tag{1}$$

[end definition]

The properties of probability functions listed in [31] follow straightforwardly:

- $p(\neg\varphi) = 1 - p(\varphi)$,
- If $\models_\bullet \varphi$ then $p(\neg\varphi) = 0$,
- If $\varphi \models_\bullet \psi$ then $p(\varphi) \leq p(\psi)$,
- If $\models_\bullet \varphi \leftrightarrow \psi$ then $p(\varphi) = p(\psi)$,
- $p(\varphi \vee \psi) = p(\varphi) + p(\psi) - p(\varphi \wedge \psi)$.
- If $p(\psi) \neq 0$, then $p(\cdot \mid \psi) : \mathbf{F}_{set} \to [0, 1]$ is a probability function.

Note that if p is first defined as a probability distribution over the interpretations on which $\models_\bullet$ is based, and the probability $p(\varphi)$ for each $\varphi \in \mathbf{F}_{set}$ is then defined as the sum of the probabilities of the interpretations in which φ evaluates to true, then p is automatically a probability function over $\mathbf{F}_{set}$.

3. Epistemic Probabilistic Event Calculus (EPEC)

3.1. Informal Overview

The key components of an EPEC domain language (defined formally in Definition 3.1 below) are *(many-valued) fluents*, *environmental actions*, *agent actions*, and *instants* (timepoints). *Literals* such as $F = V$ and $A = \mathit{true}$ assign values to fluents and truth-values to actions, and these literals are combined into *formulas* using the standard propositional conectives, and 'time-stamped' *i-formulas* using an '@' connective (Definitions 3.3 and 3.4 below). Sets of literals that mention each fluent and action exactly once are called *states* and sets that mention each fluent but no actions are called *fluent states*. Subsets of (fluent) states are called *partial (fluent) states* (Definition 3.7 below). Finally, an *outcome* O is a pair of the form $(\tilde{X}, P^+)$ where $\tilde{X}$ is a partial fluent state and $P^+ \in (0, 1]$ is a non-zero probability (Definition 3.8 below). These components are the building blocks for the six types of *propositions* that comprise an EPEC domain description (each defined formally in the next section):

- *v-propositions* of the form "F **takes-values** $\langle V_1, \ldots, V_m \rangle$" that declare what values each fluent F may take (see Definition 3.2),
- *o-propositions* of the form "A **occurs-at** I **with-prob** P^+ **if-holds** θ" indicating that environmental action A occurs at instant I with probability P^+ if formula θ holds (see Definition 3.5),
- *p-propositions* of form "A **performed-at** I **with-prob** P^+ **if-believes** $(\theta, \bar{P})$" stating that agent action A is performed at instant I with probability P^+ if formula θ is believed to hold with probability in interval $\bar{P}$ (see Definition 3.6),
- *c-propositions* of the form "θ **causes-one-of** $\{O_1, \ldots, O_m\}$" indicating that for each $O_i = (\tilde{X}_i, P_i^+)$ there is a probability of P_i^+ that formula θ will cause the changes identified in partial fluent state $\tilde{X}_i$ (see Definition 3.10),

- *i-propositions* of the form "**initially-one-of** $\{O_1, \ldots, O_m\}$" indicating that for each $O_i = (\tilde{S}_i, P_i^+)$ there is a probability of P_i^+ that the initial state of the world is the one identified in the fluent state $\tilde{S}_i$ (see Definition 3.11), and
- *s-propositions* of the form "θ **senses** X **with-accuracies M**" stating that θ holding causes the value of the fluent or action X to be sensed with accuracy given by the confusion matrix **M** (see Definition 3.12).

The utility of EPEC is that, by expressing the available narrative (including the agent's strategy) and causal information of the domain we wish to model in terms of these various proposition types, we are able to formally "entail" conclusions such as those of the doctor in Example 1.1 (modelled in full in EPEC in Appendix C). As we will see at the end of Section 3.4 below, these conclusions are in the form of *b-propositions* ('b' for 'believes') such as

at 3 **believes** $[\mathit{DocHasEpectisis}=\mathit{false}]@3$ **with-probs**
$\{ (\langle\{((\mathit{DoBloodTest}, \mathit{DocHasEpectisis}), \mathit{false})\}@1\rangle, 0.425, 0.9314),$
$(\langle\{((\mathit{DoBloodTest}, \mathit{DocHasEpectisis}), \mathit{true})\}@1, \{\mathit{TakeEpecillin}\}@2\rangle, 0.575, 0.9907)\}$

The first line of this proposition indicates the time-point or 'instant' at which the belief holds ('**at** 3'), and the temporal formula ('$[\mathit{DocHasEpectisis}=\mathit{false}]@3$') believed. The remaining lines of the proposition give the various alternative strengths of belief under the various alternative ways the represented scenario might unfold. The second line shows one possible sequence in the secenario of sensing experiences and agent actions ('$\langle\{((\mathit{DoBloodTest}, \mathit{DocHasEpectisis}), \mathit{false})\}@1\rangle$') – in this case doing a blood test at instant 1 and the result showing negative as regards having epectisis – together with the probability of that sequence occurring ('0.425'), and, if that sequence does occur, the strength of belief ('0.9314') the agent will then hold for as regards the first line. The third line is similar, describing an alternative unfolding of the scenario. So in English the whole proposition reads "at time 3 [i.e. after the blood test and possibly taking epicillin] there is a 0.425 probability that the blood test will have shown negative, in which case the doctor will be 93.14% sure she does not have the disease, and a 0.575 probability that the blood test will have shown positive, in which case the doctor will have taken epecillin and will therefore be 99.07% confident that she no longer has the disease".

We define our entailment relation, and prove some related properties, in the next three sections. In Section 3.2 we give the syntax of the various propositions that make up EPEC domain descriptions. Then in Section 3.3 we give a semantics to domain descriptions that do not contain sensing or belief conditioned actions ("non-epistemic" domains), so that the probabilities of possible worlds can be calculated in isolation. Then in Section 3.4 we use this as a building block for the semantics of full epistemic domains. For ease of reading the notation used in this section is summarised in Appendix A.

3.2. EPEC Syntax

We begin with the basic vocabulary of our framework:

Definition 3.1 (Domain Language)**.** An EPEC *domain language* is a tuple $\langle \mathcal{F}, \mathcal{A}, \mathcal{A}_e, \mathcal{A}_a, \mathcal{V}, \mathit{vals}, \mathcal{I}, \leq, \bar{0}\rangle$, where $\mathcal{F}$ is a finite non-empty set of *fluents*, $\mathcal{A}$ is a finite set of *actions*, $\mathcal{A}_e$ is a finite set of *environmental actions*, $\mathcal{A}_a$ is a finite set of *agent actions*, $\mathcal{A} = \mathcal{A}_e \cup \mathcal{A}_a$ and $\mathcal{A}_e \cap \mathcal{A}_a = \emptyset$, $\mathcal{V}$ is a finite non-empty set of *values* such that $\{\mathit{false}, \mathit{true}\} \subseteq \mathcal{V}$, *vals* is a function mapping elements in $\mathcal{F} \cup \mathcal{A}$ to tuples of elements (without repetitions) from $\mathcal{V}$ (i.e. *vals* specifies the values that each fluent and action can take), and $\mathcal{I}$ is a non-empty set of *instants* (i.e. time-points) with a minimum element $\bar{0}$ w.r.t. a total ordering $\leq$ over $\mathcal{I}$. For every $A \in \mathcal{A}$, $\mathit{vals}(A) = \langle \mathit{false}, \mathit{true}\rangle$ and for any $X \in \mathcal{F} \cup \mathcal{A}$ the expression $V \in \mathit{vals}(X)$ means that if $\mathit{vals}(X) = \langle V_1, \ldots, V_n\rangle$ then $V = V_i$ for some $1 \leq i \leq n$. [end definition]

The "epectisis" example of Section 1 can be represented using three boolean-valued fluents *PatientHasEpectisis*, *DocHasEpectisis* and *DocHasSideEffects*, an environmental action *SkinContact*, and two agent actions *DoBloodTest* and *TakeEpecillin*. For this example there are many choices for the set $\mathcal{I}$ of instants (e.g. the non-negative reals or integers) – for simplicity we use the finite set $\{-1, 0, 1, 2, 3\}$[3]. See Appendix C for the full EPEC representation $\mathcal{D}_e$ of Example 1.1, from which various example statements and expressions are taken below.

As mentioned above, EPEC domain descriptions are comprised of a finite number of propositions of various types (summarised again in Appendix A on page 81). Domain descriptions include *v-propositions* ('v' for 'value') such as

$$\mathit{PatientHasEpectisis}\ \textbf{takes-values}\ \langle \mathit{false}, \mathit{true} \rangle \tag{EP1}$$

to indicate the ordered set of values that each fluent can take. The general definition is:

Definition 3.2 (v-proposition)**.** A *v-proposition* has the form

$$F\ \textbf{takes-values}\ \langle V_1, \ldots, V_m \rangle \tag{2}$$

where $m \geq 1$ and $\langle V_1, \ldots, V_m \rangle = vals(F)$. [end definition]

EPEC uses *i-literals* such as $[\mathit{DocHasEpectisis} = \mathit{true}]@0$ to associate particular values with particular fluents, and such as $[\mathit{SkinContact} = \mathit{true}]@{-1}$ to symbolise occurrences and non-occurrences of actions, at particular instances. These can be composed into *i-formulas* such as $[\mathit{SkinContact} = \mathit{false}]@{-1} \rightarrow [\mathit{DocHasEpectisis} = \mathit{false}]@0$ as shown in the following two definitions:

Definition 3.3 (Fluent and Action Literals, i-Literals)**.** A *fluent literal* is an expression of the form $F = V$ for some $F \in \mathcal{F}$ and $V \in \mathit{vals}(F)$. A fluent is *boolean* if $\mathit{vals}(F) = \langle \mathit{false}, \mathit{true} \rangle$. An *action literal* is either $A = \mathit{false}$ or $A = \mathit{true}$ for some $A \in \mathcal{A}$. Where no ambiguity can arise, $Z = \mathit{true}$ and $Z = \mathit{false}$ are sometimes abbreviated to Z and $\neg Z$ respectively for a fluent or action Z. A *literal* is either a fluent literal or an action literal, and an *i-literal* is an expression of the form $[L]@I$ for some literal L and some $I \in \mathcal{I}$. [end definition]

[3] We use -1 as the minimum instant here simply because in the scenario the *SkinContact* occurrence happened in the past. The use of a negative integer has no semantic significance for EPEC.

Definition 3.4 (Formulas, Fluent Formulas, i-Formulas)**.** The set of *formulas*, denoted by Θ, is the closure of the set of literals under $\wedge$, $\vee$, $\neg$ and $\rightarrow$. A formula θ is said to be a *fluent formula* if it contains no action literals. The set of *i-formulas*, denoted by Φ, is the closure of the set of i-literals under $\wedge$, $\vee$, $\neg$ and $\rightarrow$. The shorthand $[\theta]@I$ stands for the i-formula formed from the formula θ and the instant I by replacing all literals L occurring in θ by $[L]@I$ (e.g. $[F\!=\!V \rightarrow F'\!=\!V']@3$ is shorthand for $[F\!=\!V]@3 \rightarrow [F'\!=\!V']@3$). The symbol $\top$ stands for an arbitrary tautological formula. [end definition]

The narrative components of EPEC domain descriptions are described with *p-* and *o-propositions* (‘p’ for ‘performed’ and ‘o’ for ‘occurred’), which optionally employ formulas within their syntax to express the conditions under which particular actions may happen at particular times. For example, the fact in Example 1.1 that a patient (who is part of the environment) may have made skin contact with the doctor (the agent) can be expressed with the (conditionless) o-proposition

$$\textit{SkinContact}\ \textbf{occurs-at}\ -1\ \textbf{with-prob}\ 0.95 \qquad \text{(EP5)}$$

and the doctor’s plan of action can be described with the two p-propositions

$$\textit{DoBloodTest}\ \textbf{performed-at}\ 1 \qquad \text{(EP10)}$$

$$\begin{array}{l}\textit{TakeEpecillin}\ \textbf{performed-at}\ 2 \\ \qquad\qquad \textbf{if-believes}\ (\textit{DocHasEpectisis}\!=\!\textit{true}, (0.5, 1]) \end{array} \qquad \text{(EP11)}$$

More generally, o- and p-propositions may optionally have both probabilities and conditions attached to them, as shown in the following two definitions:

Definition 3.5 (o-proposition)**.** An *o-proposition* has the form

$$A\ \textbf{occurs-at}\ I\ \textbf{with-prob}\ P^+\ \textbf{if-holds}\ \theta \qquad (3)$$

for some action $A \in \mathcal{A}_e$, instant I, $P^+ \in (0, 1]$ and fluent formula θ. For an o-proposition $\underline{\mathbf{o}}$ of the form (3), θ is called the *body of* $\underline{\mathbf{o}}$ or *body*($\underline{\mathbf{o}}$), and $\underline{\mathbf{o}}$ is said to *have instant* I and *principal action* A. If $P^+ = 1$ then the “ **with-prob** P^+” part of the proposition may be omitted, and if θ is $\top$ then the “ **if-holds** θ” part may be omitted. [end definition]

Definition 3.6 (p-proposition)**.** A *p-proposition* has the form

$$A \textbf{ performed-at } I \textbf{ with-prob } P^+ \textbf{ if-believes } (\theta, \bar{P}) \tag{4}$$

for some action $A \in \mathcal{A}_a$, instant I, $P^+ \in (0,1]$, fluent formula θ, and (open, half-open or closed) interval $\bar{P}$ with endpoints in $[0,1]$ (i.e. a probability range). For a p-proposition $\underline{\mathbf{p}}$ of the form (4), $(\theta, \bar{P})$ is called the *body of* $\underline{\mathbf{p}}$ or *body*($\underline{\mathbf{p}}$), and $\underline{\mathbf{p}}$ is said to *have instant* I and *principal action* A. If $P^+ = 1$ then the " **with-prob** P^+" part of the proposition may be omitted, and if θ is $\top$ and $1 \in \bar{P}$ then the " **if-believes** $(\theta, \bar{P})$" part may be omitted. [end definition]

EPEC uses *c-propositions* ('c' for 'causes') to model knowledge about the general (probabilistic) effects of actions on the agent's environment in various circumstances, independently from any particular narrative. An example from the "epectisis" domain is:

$$\begin{aligned} &\mathit{TakeEpecillin} \wedge \mathit{DocHasEpectisis} \textbf{ causes-one-of} \\ &\quad \{ \ (\{\neg \mathit{DocHasEpectisis}, \mathit{DocHasSideEffects}\}, 0.1485), \\ &\quad\quad (\{\neg \mathit{DocHasEpectisis}\}, 0.8415), \\ &\quad\quad (\{\mathit{DocHasSideEffects}\}, 0.0015), \\ &\quad\quad (\emptyset, 0.0085) \ \} \end{aligned} \tag{EP7}$$

The attached probabilities in this particular example (taken together with those in the twin c-proposition (EP8) in Appendix C) reflect that the two potential effects of *TakeEpecillin* are independent[4]. But in general the effects captured in c-propositions need not be independent.

To give the general definition of a c-proposition we need terms for its components. Sets such as $\{\neg \mathit{DocHasEpectisis}, \mathit{DocHasSideEffects}\}$ of fluent literals are *partial fluent states*, and a partial fluent state paired with a probability, such as $(\{\neg \mathit{DocHasEpectisis}, \mathit{DocHasSideEffects}\}, 0.1485)$, is called an *outcome*. The c-proposition above has four outcomes in its *head*, and, as for all c-propositions, their combined *weight* (i.e. the sum of their probabilities) is 1. The next four definitions formalise this terminology, and give some further notation that will be required in Sections 3.3 and 3.4 to describe EPEC's semantics.

[4]In the standard sense in probability theory, i.e. that random variables A and B are independent if for all values a and b, $\Pr(A = a).\Pr(B = b) = \Pr(A = a, B = b)$.

Definition 3.7 (State, Partial State, Fluent State)**.** A *state* S is a set of literals, exactly one for each $F \in \mathcal{F}$ and $A \in \mathcal{A}$. A *partial state* is a subset $X \subseteq S$ of a state S. The subset of a partial state X containing exactly the fluent literals in X is a *partial fluent state*, and is denoted by $X{\restriction}\mathcal{F}$. For S a state, $S{\restriction}\mathcal{F}$ is called a *fluent state*. The subset of X containing exactly the action literals in X is denoted by $X{\restriction}\mathcal{A}$. The set of all states is denoted by $\mathcal{S}$, and the set of all partial states is denoted by $\mathcal{X}$. Finally, the sets $\{S{\restriction}\mathcal{F} \mid S \in \mathcal{S}\}$ and $\{X{\restriction}\mathcal{F} \mid X \in \mathcal{X}\}$ are denoted by $\tilde{\mathcal{S}}$ and $\tilde{\mathcal{X}}$ respectively. [end definition]

Definition 3.8 (Outcome, Projection Functions)**.** An *outcome* is a pair of the form $(\tilde{X}, P^+)$ where $\tilde{X}$ is a partial fluent state and $P^+ \in (0, 1]$ (i.e. P^+ is a non-zero probability). The *projection functions* χ and π are such that for any outcome $O = (\tilde{X}, P^+)$, $\chi(O) = \tilde{X}$ and $\pi(O) = P^+$. The set of all outcomes $\tilde{\mathcal{X}} \times (0, 1]$ is denoted by $\mathcal{O}$. [end definition]

Definition 3.9 (Weight of a Set of Outcomes)**.** Given a finite set of outcomes $B = \{O_1, O_2, \ldots, O_m\}$ the *weight of* B is defined as

$$\pi(B) = \sum_{i=1}^{m} \pi(O_i).$$

(i.e. the sum of the probabilities of its elements.) [end definition]

We can now give the general definition of a c-proposition. In it, and for the remainder of the paper, literals of the EPEC domain language are interpreted as atomic propositions when using the standard propositional entailment symbol $\models$, so that for example $(F = \mathit{true} \wedge F' = \mathit{false}) \models F = \mathit{true}$.

Definition 3.10 (c-proposition)**.** A *c-proposition* has the form

$$\theta \textbf{ causes-one-of } \{O_1, O_2, \ldots, O_m\} \tag{5}$$

where, for $i = 1, \ldots, m$, $O_i \in \mathcal{O}$, $\chi(O_i) \neq \chi(O_j)$ when $i \neq j$, θ is a formula such that $\theta \models (A = \mathit{true})$ for at least one $A \in \mathcal{A}$, and $\pi(\{O_1, \ldots, O_m\}) = 1$. For a c-proposition $\underline{\mathbf{c}}$ of the form (5), the formula $\mathit{body}(\underline{\mathbf{c}}) = \theta$ and the set $\mathit{head}(\underline{\mathbf{c}}) = \{O_1, \ldots, O_m\}$ are the *body* and *head* of $\underline{\mathbf{c}}$, respectively. Outcome O_i is often omitted from $\mathit{head}(\underline{\mathbf{c}})$ if $\chi(O_i) = \emptyset$ (leaving $\pi(O_i)$ implicit since $\pi(\{O_1, \ldots, O_m\}) = 1$). [end definition]

To unambiguously calculate the probabilities of fluents taking particular values after a particular series of actions, we also need to assign probabilities to their possible initial values. EPEC domain descriptions include a unique *i-proposition* (‘i’ for ‘initially’) which lists the fluent states that have a non-zero initial probability. In Example 1.1 the doctor is reasoning on the assumption that she did not initially have the disease, therefore the i-proposition for that domain is

$$\begin{array}{l} \textbf{initially-one-of} \\ \{\ (\{\mathit{PatientHasEpectisis}, \neg \mathit{DocHasEpectisis}, \neg \mathit{DocHasSideEffects}\}, 0.8), \\ \quad (\{\neg \mathit{PatientHasEpectisis}, \neg \mathit{DocHasEpectisis}, \neg \mathit{DocHasSideEffects}\}, 0.2)\ \} \end{array} \tag{EP4}$$

The general definition of the form of an i-proposition is as follows:

Definition 3.11 (i-proposition)**.** An *i-proposition* has the form

$$\textbf{initially-one-of } \{O_1, O_2, \ldots, O_m\} \tag{6}$$

where, for $i = 1, \ldots, m$, $O_i \in \mathcal{O}$, $\pi(\{O_1, \ldots, O_m\}) = 1$, $\chi(O_i) \in \tilde{\mathcal{S}}$, and $\chi(O_i) \neq \chi(O_j)$ when $i \neq j$ (i.e. $\chi(O_i)$ and $\chi(O_j)$ are mutually exclusive fluent states). [end definition]

Finally, we introduce *s-propositions* (‘s’ for ‘senses’) to capture the general effects of sensing actions. The effectiveness of the blood test in Example 1.1 is represented as:

$$\mathit{DoBloodTest} \textbf{ senses } \mathit{DocHasEpectisis} \textbf{ with-accuracies } \begin{pmatrix} 0.9 & 0.1 \\ 0.05 & 0.95 \end{pmatrix} \tag{EP9}$$

The leading diagonal in the *accuracy* (or *confusion*) *matrix* gives the probabilities that the test correctly indicates the doctor’s condition (e.g. if the doctor has the disease there is a 95% probability that the test will correctly show this). A 2×2 matrix is used here because *DocHasEpectisis* is a 2-valued fluent. The other two entries indicate the probabilities of false positives and false negatives, with the order of the rows and columns as indicated in the v-proposition (EP2) for *DocHasEpectisis*. The general definition for an s-proposition is:

Definition 3.12 (s-proposition)**.** Let $X \in \mathcal{F} \cup \mathcal{A}_e$ and $vals(X) = \langle V_1, \ldots, V_m \rangle$. An *s-proposition* has the form

$$\theta \textbf{ senses } X \textbf{ with-accuracies } \mathbf{M} \tag{7}$$

where $\theta \models (A = true)$ for some $A \in \mathcal{A}_a$, and $\mathbf{M}$ is an $m \times m$ matrix with all elements in $[0, 1]$. For an s-proposition $\underline{\mathbf{s}}$ of the form (7), θ is called the *body of* $\underline{\mathbf{s}}$, or *body*($\underline{\mathbf{s}}$), and X is called the *object of* $\underline{\mathbf{s}}$, or *object*($\underline{\mathbf{s}}$). The pair (θ, X) is called the *signature of* $\underline{\mathbf{s}}$ and denoted by *sig*($\underline{\mathbf{s}}$). The element $\mathbf{M}_{i,j}$ in $\mathbf{M}$ represents the probability that, given that V_i is the actual value of X when θ occurs, the value V_j is sensed. Hence $\mathbf{M}$ is subject to the condition that each row adds to 1:

$$\forall 1 \leq i \leq m, \quad \sum_{j=1}^{m} \mathbf{M}_{i,j} = 1 \tag{8}$$

The s-proposition "θ **senses** X **with-accuracies** $\mathbf{I}_M$" (where $\mathbf{I}_M$ is the $m \times m$ identity matrix, representing perfect sensing) is sometimes abbreviated to "θ **senses** X". [end definition]

An EPEC domain description is a finite collection of the types of proposition mentioned above (and summarised in Appendix A on page 81), but with some restrictions to ensure that the different propositions do not contradict each other. For example, we cannot include two o-propositions that both state that a particular action occurs at a particular time, but with different probabilities. Neither can we include two c-propositions that state that the same action(s) in the same circumstances have two different sets of outcomes. To describe these restrictions precisely, we employ the following definition of *(in)compatibility* of formulas.

Definition 3.13 (Compatibility of Formulas)**.** Given a partial state X and a formula θ, we sometimes write $X \models \theta$ to indicate that $\bigwedge_{L \in X} L \models \theta$. Two formulas θ_1 and θ_2 are *compatible* if there is a state S such that $S \models \theta_1 \wedge \theta_2$, and *incompatible* otherwise. [end definition]

We now give the definition of an EPEC domain description. As an example, the full domain description of Example 1.1 is given in Appendix C.

Definition 3.14 (Domain Description)**.** A *domain description* is a finite set $\mathcal{D}$ of v-propositions, c-propositions, p-propositions, o-propositions, i-propositions and s-propositions such that:

(i) $\mathcal{D}$ contains exactly one v-proposition for each $F \in \mathcal{F}$ [see Def. 3.2],

(ii) $\mathcal{D}$ contains exactly one i-proposition [see Def. 3.11],

(iii) for any two distinct c-propositions in $\mathcal{D}$ with bodies θ_1 and θ_2 [see Def. 3.10], θ_1 and θ_2 are incompatible [see Def. 3.13],

(iv) for any given $A \in \mathcal{A}_e$ and $I \in \mathcal{I}$, if $\mathcal{D}$ contains a pair of o-propositions "A **occurs-at** I **with-prob** P_1^+ **if-holds** θ_1" and "A **occurs-at** I **with-prob** P_2^+ **if-holds** θ_2", then θ_1 and θ_2 are incompatible [see Defs 3.13, 3.5],

(v) for any given $A \in \mathcal{A}_a$ and $I \in \mathcal{I}$, if $\mathcal{D}$ contains a pair of p-propositions "A **performed-at** I **with-prob** P_1^+ **if-believes** $(\theta_1, \bar{P}_1)$" and "A **performed-at** I **with-prob** P_2^+ **if-believes** $(\theta_2, \bar{P}_2)$" then $\theta_1 = \theta_2$ and $\bar{P}_1 \cap \bar{P}_2 = \emptyset$ [see Def. 3.6],

(vi) no two s-propositions in $\mathcal{D}$ have the same signature [see Def. 3.12].

[end definition]

In the remainder of the paper, since by Definition 3.14 each s-proposition within a domain description $\mathcal{D}$ has a unique signature (θ, X), we will sometimes refer to its accuracy matrix as $\mathbf{M}_{\mathcal{D}}(\theta, X)$.

3.3. Semantic Entailment and Results for Non-epistemic Domains

In this section we give a semantics to *non-epistemic* domain descriptions that do not contain sensing actions or belief-conditioned action occurrences:

Definition 3.15 (Ne-domain Description)**.** An *ne-domain description* (*non-epistemic domain description*) is a domain description that contains no s- or p-propositions. [end definition]

In what follows, $\mathcal{L}$, $\mathcal{D}$ and $\mathcal{N}$ will generally signify an arbitrary domain language, domain description and ne-domain description respectively. As an example ne-domain, consider an 80% probable single roll of a dice which has been painted so that once face shows 'a', two faces show 'b' and three faces show 'c'. We model this scenario with three instants 0, 1 and 2, and supposing that, before the probable roll, the dice is initially showing either an 'a' or a 'b' on its uppermost face, with equal probabilities. The ne-domain description $\mathcal{N}_d$ is:

Example 3.1 (Lettered Dice Roll)**.** $\mathcal{N}_d$ consists of the following propositions:

Face **takes-values** $\langle a, b, c \rangle$ (D1)

initially-one-of $\{(\{Face = a\}, 0.5), (\{Face = b\}, 0.5)\}$ (D2)

Roll **causes-one-of** (D3)

$$\{\ (\{Face = a\}, 0.17), (\{Face = b\}, 0.33), (\{Face = c\}, 0.5)\ \}$$

Roll **occurs-at** 1 **with-prob** 0.8 if-holds $\top$ (D4)

[end example]

For ne-domain descriptions, the key semantic structure is a *world* (analogous to a modal logic possible world). Worlds are effectively timelines labelled with all fluent values and action occurrences/non-occurrences at each instant:

Definition 3.16 (World)**.** A *world* is a function $W : \mathcal{I} \rightarrow \mathcal{S}$. The set of all worlds is denoted by $\mathcal{W}$. [end definition]

In Example 3.1 there is one 3-valued fluent and one (binary-valued) action, and hence 6 possible states in the set $\mathcal{S}$. As there are 3 instants, this gives a total of $6^3 = 216$ (mostly nonsensical) worlds, four examples of which are:

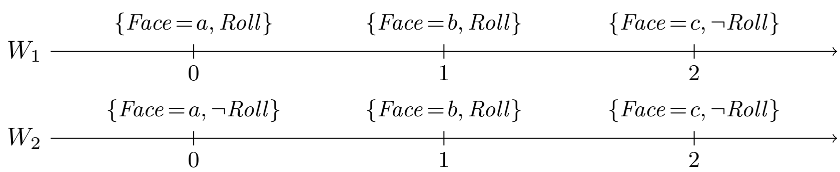

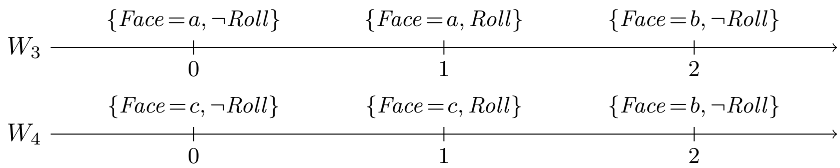

Definitions 3.17 to 3.24 which follow allow us to distinguish between worlds that are compatible with the ne-domain description in question – we call these *well-behaved* worlds – and those that are not. (Example 3.1 has 8 out of 216 well-behaved worlds.) We start with a notion of *satifaction* of an i-formula:

Definition 3.17 (Satisfaction of an i-formula, Logical Consequence for i-formulas)**.** Given a world W and a literal L, W *satisfies an i-formula* $[L]@I$, written $W \Vdash [L]@I$, iff $L \in W(I)$. Otherwise, $W \not\Vdash [L]@I$. The definition of $\Vdash$ is recursively extended for arbitrary i-formulas as follows: if φ and ψ are i-formulas, $W \Vdash \varphi \wedge \psi$ iff $W \Vdash \varphi$ and $W \Vdash \psi$, and $W \Vdash \neg\varphi$ iff $W \not\Vdash \varphi$. The i-formulas $\varphi \vee \psi$ and $\varphi \rightarrow \psi$ are interpreted as shorthand for $\neg(\neg\varphi \wedge \neg\psi)$ and $\neg(\varphi \wedge \neg\psi)$ respectively. Given a (possibly empty) set Δ of i-formulas, $W \Vdash \Delta$ iff $W \Vdash \psi$ for all $\psi \in \Delta$. Given an i-formula φ and a set Δ of i-formulas $\Delta \Vdash \varphi$ if for all $W \in \mathcal{W}$ such that $W \Vdash \Delta$, $W \Vdash \varphi$ also holds. For two i-formulas φ and ψ, $\psi \Vdash \varphi$ is shorthand for $\{\psi\} \Vdash \varphi$, and $\Vdash \varphi$ is shorthand for $\emptyset \Vdash \varphi$. [end definition]

Examples of i-formula satisfaction for the lettered dice domain (Example 3.1) are $W_1 \Vdash ([Roll]@0 \wedge [Roll]@1)$, $W_2 \Vdash [Face{=}c]@2$ and $W_3 \Vdash [Roll]@1$.

The first criterion that a world must satisfy to be well-behaved with respect to an ne-domain description $\mathcal{N}$ is that the action occurrences it identifies along its timeline exactly match those represented as o-propositions in $\mathcal{N}$:

Definition 3.18 (Closed World Assumption for Actions)**.** A world W satisfies the *closed world assumption for actions* (CWA) *w.r.t. an ne-domain description* $\mathcal{N}$ if for all $A \in \mathcal{A}$ and $I \in \mathcal{I}$:

(i) if $W \mid\models [A]@I$ then there exists some $P^+ \in (0, 1]$ and fluent formula θ such that $W \mid\models [\theta]@I$ and "A **occurs-at** I **with-prob** P^+ **if-holds** θ" is in $\mathcal{N}$, and

(ii) for any θ such that $W \mid\models [\theta]@I$ and the o-proposition "A **occurs-at** I **with-prob** 1 **if-holds** θ" is in $\mathcal{N}$, $W \mid\models [A]@I$.

[end definition]

Of the four worlds pictured for Example 3.1, W_2, W_3 and W_4 all satisfy the CWA w.r.t. $\mathcal{N}_d$, but W_1 fails condition (i) above since $W_1 \mid\models [Roll]@0$.

The second criterion that a world must satisfy to be well-behaved with respect to an ne-domain description $\mathcal{N}$ is that it is compatible with one of the alternative initial conditions identified in the domain description's i-proposition. (Note that the next few definitions apply to an arbitrary domain description $\mathcal{D}$, and thus also to any ne-domain description $\mathcal{N}$.)

Definition 3.19 (Initial Choice and Initial Consistency)**.** Let $\mathcal{D}$ be a domain description with (unique) i-proposition "**initially-one-of** $\{O_1, O_2, \ldots, O_m\}$". Each $O_1, O_2, \ldots, O_m$ is called an *initial choice of* $\mathcal{D}$. A world W is said to *satisfy the initial condition of* $\mathcal{D}$ if there exists an initial choice O_i of $\mathcal{D}$ such that $W(\bar{0}) \restriction \mathcal{F} = \chi(O_i)$ [See Defs 3.1, 3.7, 3.8]. In this case it is said that W *and* O_i *are initially consistent with each other w.r.t.* $\mathcal{D}$. [end definition]

For Example 3.1, W_1, W_2 and W_3 all satisfy the initial condition of $\mathcal{N}_d$ since $W_1(0)\restriction\mathcal{F} = W_2(0)\restriction\mathcal{F} = W_3(0)\restriction\mathcal{F} = \{Face{=}a\} = \chi(\{Face{=}a\}, 0.5)$, but W_4 does not since $W_4(0)\restriction\mathcal{F} = \{Face{=}c\}$ which does not appear in (D2).

The last criterion that a world must satisfy to be well-behaved is that the changes in fluent values along its time-line are explainable in terms of the domain description's c-propositions. To formalise this, we need to be able to identify which c-propositions are activated at which time-points in that world:

Definition 3.20 (Cause Occurrence)**.** Let θ be the body of a c-proposition $\underline{\mathbf{c}}$ in a domain description $\mathcal{D}$ and $I \in \mathcal{I}$. If $W \Vdash [\theta]@I$ then it is said that *a cause occurs at instant I in W w.r.t. to $\mathcal{D}$*, and that *$\underline{\mathbf{c}}$ is activated at I in W w.r.t. $\mathcal{D}$*. The set $occ_{\mathcal{D}}(W)$ is the set $\{I \in \mathcal{I} \mid \text{a cause occurs at } I \text{ in } W\}$. The function $cprop_{\mathcal{D}}$ with domain $\{(W, I) \mid W \in \mathcal{W}, I \in occ_{\mathcal{D}}(W)\}$ is defined as $cprop_{\mathcal{D}}(W, I) = \underline{\mathbf{c}}$ where $\underline{\mathbf{c}}$ is the (unique) c-proposition activated at I in W. [end definition]

In Example 3.1 the single c-proposition (D3) has body *Roll*. Since $W_1 \Vdash [Roll]@0 \wedge [Roll]@1$, (D3) is activated at 0 and at 1 in W_1, $occ_{\mathcal{N}_d}(W_1) = \{0, 1\}$, and $cprop_{\mathcal{N}_d}(W_1, 0) = cprop_{\mathcal{N}_d}(W_1, 1) = \text{(D3)}$.

We next define an *effect choice* for a given world. This selects one particular outcome from each c-proposition at each point that it is activated:

Definition 3.21 (Effect Choice)**.** Let W be a world and $\mathcal{D}$ a domain description. An *effect choice for W w.r.t. $\mathcal{D}$* is a function $ec : occ_{\mathcal{D}}(W) \to \mathcal{O}$ such that for all instants $I \in occ_{\mathcal{D}}(W)$, $ec(I) \in head(cprop_{\mathcal{D}}(W, I))$. [end definition]

In Example 3.1, one of the nine possible effect choices for W_1 w.r.t. $\mathcal{N}_d$ is the effect choice ec_1^d, specifed as $ec_1^d(0) = (\{Face{=}b\}, 0.33)$, $ec_1^d(1) = (\{Face{=}c\}, 0.5)$.

To describe the expected effect of one of a c-proposition's outcomes on the state in which it is activated, we define the notion of a *fluent state update*:

Definition 3.22 (Fluent State Update)**.** Given a fluent state $\tilde{S}$ and a partial fluent state $\tilde{X}$, the *update of $\tilde{S}$ w.r.t. $\tilde{X}$*, written $\tilde{S} \oplus \tilde{X}$, is the fluent state $(\tilde{S} \ominus \tilde{X}) \cup \tilde{X}$, where $\tilde{S} \ominus \tilde{X}$ is the partial fluent state formed by removing all fluent literals from $\tilde{S}$ of the form $F{=}V$ for some F and V' such that $F{=}V' \in \tilde{X}$. The operator $\oplus$ is left-associative, so e.g. $\tilde{S} \oplus \tilde{X} \oplus \tilde{X}'$ is understood as $((\tilde{S} \oplus \tilde{X}) \oplus \tilde{X}')$. [end definition]

For example, $\{Dice{=}a\} \oplus \{Dice{=}b\} = \{Dice{=}b\}$, and $\{PatientHasEpectisis, \neg DocHasEpectisis, \neg DocHasSideEffects\} \oplus \{DocHasEpectisis\} = \{PatientHasEpectisis, DocHasEpectisis, \neg DocHasSideEffects\}$.

We can now state our final criterion for a world to be well-behaved – it must satisfy the *justified change condition*. Definition 3.23 below states that changes in fluents' values along a world's time-line occur only immediately after instants where a c-proposition is activated, and that all the fluents' new values must appear in a single outcome of that c-proposition. The definition thus encapsulates a solution to the frame problem, as it ensures that fluents not explicitly affected by the activation of a c-proposition have a default persistence. Like all EPEC's definitions, Definition 3.23 is stated in a manner that allows it to be applied to continuous as well as discrete time-lines. To aid understanding, the reader may find it helpful to first note that the effect choice ec_1^d described after Definition 3.21 provides a justification for the fluent changes along world W_1, since, for example

$$\begin{aligned} W_1(2){\restriction}\mathcal{F} = \{Face{=}c\} &= \{Face{=}a\} \oplus \{Face{=}b\} \oplus \{Face{=}c\} \\ &= (W_1(0){\restriction}\mathcal{F}) \oplus \chi(\{Face{=}b\}, 0.33) \oplus \chi(\{Face{=}c\}, 0.5) \\ &= (W_1(0){\restriction}\mathcal{F}) \oplus \chi(ec_1^d(0)) \oplus \chi(ec_1^d(1)) \end{aligned}$$

Definition 3.23 (Justified Change)**.** A world W satisfies the *justified change condition w.r.t.* $\mathcal{D}$ if and only if there exists an effect choice ec w.r.t. $\mathcal{D}$ such that for all instants I and I' with $I < I'$, ec maps the instants in $occ_{\mathcal{D}}(W) \cap [I, I') = \{I_1, \ldots, I_n\}$ to $O_1, O_2, \ldots, O_n$ respectively, where $I_1, \ldots, I_n$ are ordered w.r.t. $\leq$, and

$$W(I'){\restriction}\mathcal{F} = (W(I){\restriction}\mathcal{F}) \oplus \chi(O_1) \oplus \chi(O_2) \oplus \cdots \oplus \chi(O_n) \tag{9}$$

If a world W satisfies the justified change condition for some effect choice ec, W and ec are said to be *consistent* with each other *w.r.t.* $\mathcal{D}$. [end definition]

In fact W_1, W_3 and W_4 all satisfy the justified change condition w.r.t $\mathcal{N}_d$, whereas W_2 has an unjustified change from instant 0 to instant 1.

Definition 3.24 summarises the conditions for a world to be well-behaved:

Definition 3.24 (Well-behaved World)**.** A world is *well-behaved w.r.t. an ne-domain description* $\mathcal{N}$ if it satisfies the closed world assumption for actions w.r.t. $\mathcal{N}$ [see Def. 3.18], the initial condition of $\mathcal{N}$ [see Def. 3.19] and the justified change condition w.r.t. $\mathcal{N}$ [see Def. 3.23]. The set of well-behaved worlds w.r.t. $\mathcal{N}$ is denoted by $\mathcal{W}_{\mathcal{N}}$. [end definition]

Therefore, by Definitions 3.18 to 3.24, $W_3 \in \mathcal{W}_{\mathcal{N}_d}$ but $W_1, W_2, W_4 \notin \mathcal{W}_{\mathcal{N}_d}$. Note that, since the sets $\mathcal{F}$, $\mathcal{A}$ and $\mathcal{V}$ of fluents, actions and values are all finite, and a domain description contains only a finite number of propositions, the CWA (Definition 3.18) and justified change condition (Definition 3.23) together ensure that the set $\mathcal{W}_{\mathcal{N}}$ of well-behaved worlds is always finite, for any ne-domain description $\mathcal{N}$.

We can use the notion of a well-behaved world to adapt the notion of entailment (symbolised by $|\models$) given in Definition 3.17 to a specific ne-domain $\mathcal{N}$. $\mathcal{N}$-*entailment* is symbolised by $|\models_{\mathcal{N}}$ in the following definition:

Definition 3.25 ($\mathcal{N}$-entailment)**.** Given an i-formula φ, a set Δ of i-formulas and an ne-domain description $\mathcal{N}$, Δ $\mathcal{N}$-*entails* φ, written $\Delta \mathrel{|\models_{\mathcal{N}}} \varphi$, if for all well-behaved worlds $W \in \mathcal{W}_{\mathcal{N}}$ such that $W \mathrel{|\models} \Delta$, $W \mathrel{|\models} \varphi$ also holds [see Def. 3.17]. For two i-formulas ψ and φ, $\psi \mathrel{|\models_{\mathcal{N}}} \varphi$ is shorthand for $\{\psi\} \mathrel{|\models_{\mathcal{N}}} \varphi$, and $|\models_{\mathcal{N}} \varphi$ is shorthand for $\emptyset \mathrel{|\models_{\mathcal{N}}} \varphi$. **[end definition]**

For example, $[\mathit{Face}=a]@0 \mathrel{|\models_{\mathcal{N}_d}} [\mathit{Face}=a]@1$.

To complete the semantics of an ne-domain description, it remains to define a probability distribution over the set of well-behaved worlds that properly accounts for the various probabilities embedded in the propositions. We begin with the notion of a *trace* of a well-behaved world, which is an initial choice coupled with an effect choice that can account for the values of fluents along the world's time-line. We can straightforwardly evaluate the probability of a given trace, conditional on the action occurrences in the associated world actually taking place, by taking the product of the probabilities of the individual outcomes selected by the trace:

Definition 3.26 (Trace)**.** Let W be a well-behaved world w.r.t. $\mathcal{N}$. A *trace of* W *w.r.t.* $\mathcal{N}$ is a pair $\langle ic, ec\rangle$ where ic is an initial choice [see Def. 3.19] consistent with W and ec is an effect choice [see Def. 3.21] consistent with W. In this case $\langle ic, ec\rangle$ is also a *trace of* $\mathcal{N}$. The set $traces(W, \mathcal{N})$ is the set of all traces of W w.r.t. $\mathcal{N}$. For an arbitrary trace $tr = \langle ic, ec\rangle$ of $\mathcal{N}$, we identify the (unique) corresponding well-behaved world as W_{tr}, and define the *evaluation of* tr, written $\epsilon(tr)$, as:

$$\epsilon(tr) \;=\; \pi(ic)\;.\prod_{I\in occ_{\mathcal{D}}(W_{tr})} \pi(ec(I)) \tag{10}$$

(Where $\pi : \mathcal{O} \mapsto (0, 1]$ is as defined in Definition 3.8.) [end definition]

For Example 3.1, $\langle ic_2^d, ec_2^d\rangle$ is a trace for the well-behaved world W_3, where $ic_2^d = (\{Face{=}a\}, 0.5)$ and $ec_2^d(1) = (\{Face{=}b\}, 0.33)$, so that $\epsilon(\langle ic_2^d, ec_2^d\rangle) = 0.165$. Note that although in this example all eight well behaved worlds have a unique trace, in the general case a well-behaved world might have more that one trace. To see this, consider the following variation of the c-proposition (D3):

Roll **causes-one-of** (D3′)
$\{\ (\{Face{=}a\}, 0.16), (\{Face{=}b\}, 0.31), (\{Face{=}c\}, 0.47), (\emptyset, 0.06)\ \}$

indicating that very occasionally *Roll* fails to have any effect on the dice, and consider the well-behaved world W_5:

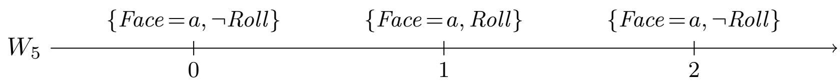

If (D3′) replaces (D3) in $\mathcal{N}_d$ there are two effect choices ec' and ec'' for W_5, where $ec'(1) = (\{Face{=}a\}, 0.16)$ and $ec''(1) = (\emptyset, 0.06)$. Hence W_5 has two traces.

To calculate the probability that the exact sequence of action occurrences/non-occurrences entailed by a particular well-behaved world W actually takes place, we can take the product of the probabilities/complement probabilities indicated in the corresponding o-propositions. We call this product the *narrative evaluation* of $\mathcal{N}$ w.r.t. W:

Definition 3.27 (Narrative Evaluation)**.** Given an o-proposition $\underline{\mathbf{o}}$ of the form "A **occurs-at** I **with-prob** P^+ **if-holds** θ" and a world W, the *narative evaluation of* $\underline{\mathbf{o}}$ *w.r.t.* W is defined as

$$\epsilon(\underline{\mathbf{o}}, W) = \begin{cases} 1 & \text{if } W \mid\not\models [\theta]@I \\ P^+ & \text{if } W \mid\models [\theta]@I \text{ and } W \mid\models [A]@I \\ 1 - P^+ & \text{if } W \mid\models [\theta]@I \text{ and } W \mid\models [\neg A]@I \end{cases} \tag{11}$$

For an ne-domain description $\mathcal{N}$ the above definition is extended to:

$$\epsilon(\mathcal{N}, W) = \prod_{\underline{\mathbf{o}} \in \mathcal{N}} \epsilon(\underline{\mathbf{o}}, W). \tag{12}$$

If $\mathcal{N}$ contains no o-propositions then $\epsilon(\mathcal{N}, W) = 1$. [end definition]

For Example 3.1, $\epsilon(\mathcal{N}_d, W_3) = \epsilon((\text{D4}), W_3) = 0.8$.

We now define our probability distribution, the *ne-model-function*, over the well-behaved worlds of an ne-domain description $\mathcal{N}$, symbolised as $M^{ne}_{\mathcal{N}}$, and prove that it is indeed a probability distribution in Proposition 3.1 that follows. The probability of a well-behaved world W is the product of the probability of its narrative component (action occurrences) being true and the probability of the particular chosen causal effects of those actions having taken place.

Definition 3.28 (Ne-model-function)**.** The *ne-model-function* of an ne-domain description $\mathcal{N}$ is the function $M^{ne}_{\mathcal{N}} : \mathcal{W}_{\mathcal{N}} \mapsto [0, 1]$ defined for each well-behaved world W as:

$$M^{ne}_{\mathcal{N}}(W) = \epsilon(\mathcal{N}, W) \quad \cdot \sum_{tr \in traces(W, \mathcal{N})} \epsilon(tr) \tag{13}$$

The ne-model-function $M^{ne}_{\mathcal{N}}$ is extended to a function $M^{ne}_{\mathcal{N}} : \Phi \mapsto [0, 1]$ over i-formulas in the following way:

$$M^{ne}_{\mathcal{N}}(\varphi) = \sum_{W \mid\models \varphi} M^{ne}_{\mathcal{N}}(W) \tag{14}$$

and if ψ is such that $M^{ne}_{\mathcal{N}}(\psi) \neq 0$, then the function $M^{ne}_{\mathcal{N}}(\cdot \mid \psi) : \Phi \mapsto [0, 1]$ is defined as:

$$M^{ne}_{\mathcal{N}}(\varphi \mid \psi) = \frac{M^{ne}_{\mathcal{N}}(\varphi \wedge \psi)}{M^{ne}_{\mathcal{N}}(\psi)} \tag{15}$$

[end definition]

Using the above definition in the context of Example 3.1, we can see for example that $M^{ne}_{\mathcal{N}_d}(W_3) = 0.8 \times (0.5 \times 0.33) = 0.132$, and since W_3 (labelled as $W_{[aRb]}$ in the diagram below) is the only well-behaved world that satifies both $[Face\!=\!a]@0$ and $[Face\!=\!b]@2$, equation (14) gives $M^{ne}_{\mathcal{N}_d}([Face\!=\!a]@0 \wedge [Face\!=\!b]@2) = 0.132$. Inspection of the eight well-behaved worlds for $\mathcal{N}_d$ pictured below shows that two of them, $W_{[a\neg Ra]}$ and $W_{[b\neg Rb]}$, satisfy $[\neg Roll]@1$, both with probability $0.2 \times 0.5 = 0.1$, and the first of these worlds also satisfies $[Face\!=\!a]@2$. So by equation (15), $M^{ne}_{\mathcal{N}_d}(\ [Face\!=\!a]@2 \ \mid\ [\neg Roll]@1\ ) = \frac{0.1}{(0.1+0.1)} = 0.5$.

	0	1	2	
$W_{[aRa]}$	$\{Face\!=\!a, \neg Roll\}$	$\{Face\!=\!a, Roll\}$	$\{Face\!=\!a, \neg Roll\}$	$0.8 \times (0.5 \times 0.17)$
$W_{[aRb]}$	$\{Face\!=\!a, \neg Roll\}$	$\{Face\!=\!a, Roll\}$	$\{Face\!=\!b, \neg Roll\}$	$0.8 \times (0.5 \times 0.33)$
$W_{[aRc]}$	$\{Face\!=\!a, \neg Roll\}$	$\{Face\!=\!a, Roll\}$	$\{Face\!=\!c, \neg Roll\}$	$0.8 \times (0.5 \times 0.5)$
$W_{[bRa]}$	$\{Face\!=\!b, \neg Roll\}$	$\{Face\!=\!b, Roll\}$	$\{Face\!=\!a, \neg Roll\}$	$0.8 \times (0.5 \times 0.17)$
$W_{[bRb]}$	$\{Face\!=\!b, \neg Roll\}$	$\{Face\!=\!b, Roll\}$	$\{Face\!=\!b, \neg Roll\}$	$0.8 \times (0.5 \times 0.33)$
$W_{[bRc]}$	$\{Face\!=\!b, \neg Roll\}$	$\{Face\!=\!b, Roll\}$	$\{Face\!=\!c, \neg Roll\}$	$0.8 \times (0.5 \times 0.5)$
$W_{[a\neg Ra]}$	$\{Face\!=\!a, \neg Roll\}$	$\{Face\!=\!a, \neg Roll\}$	$\{Face\!=\!a, \neg Roll\}$	0.2×0.5
$W_{[b\neg Rb]}$	$\{Face\!=\!b, \neg Roll\}$	$\{Face\!=\!b, \neg Roll\}$	$\{Face\!=\!b, \neg Roll\}$	0.2×0.5

The numbers in the right of the diagram above sum to 1, and so by definition form a probability distribution over the well-behaved worlds of $\mathcal{N}_d$. The following proposition confirms that this is true in the general case.

Proposition 3.1. Given an ne-domain description $\mathcal{N}$ [see Def. 3.15], the ne-model $M^{ne}_{\mathcal{N}} : \mathcal{W}_{\mathcal{N}} \mapsto [0, 1]$ [see Def. 3.28] is a probability distribution over $\mathcal{W}_{\mathcal{N}}$.

Proof: See Appendix D.1. ■

Corollary 3.1. Given an ne-domain description $\mathcal{N}$, the ne-model $M^{ne}_{\mathcal{N}} : \Phi \mapsto [0, 1]$ [see Def. 3.28] is a probability function over Φ w.r.t. $|\!\models_{\mathcal{N}}$ [see Defs 2.1, 3.4, 3.25].

Proof: The corollary follows directly from Proposition 3.1, Definitions 2.1, 3.17 and 3.25 of a probability function, satisfaction of an i-formula and $\mathcal{N}$-entailment respectively, equation (14) of Definition 3.28, and the remarks in the last paragraph of Section 2.3, taking i-literals as atomic propositional symbols and well-behaved worlds as interpretations. ■

Specific evaluations of $M^{ne}_{\mathcal{N}}(\varphi)$ and $M^{ne}_{\mathcal{N}}(\varphi \mid \psi)$ using equations (14) and (15) of Definition 3.28 can be written in action-language-style syntax using *h-propositions*. We write "$\mathcal{N} \mid\models \varphi$ **holds-with-prob** P^+" to mean that $M^{ne}_{\mathcal{N}}(\varphi) = P^+$, and we write $M^{ne}_{\mathcal{N}}(\varphi \mid \psi) = P^+$ as "$(\mathcal{N} \mid \psi) \mid\models \varphi$ **holds-with-prob** P^+".

3.4. Semantic Entailment and Results for Epistemic Domains

In this section we return to consideration of full domain descriptions that may include p-propositions (describing belief-conditioned agent action occurrences) and s-propositions (describing the effects of sensing actions). We also return to using the Epectitis domain (Example 1.1, listed in full in Appendix C) as a running example. For convenience here, we will use the abbreviated fluent and action names *PtHaEp* (*PatientHasEpectisis*), *DrHaEp* (*DocHasEpectisis*), *DrHaSiEf* (*DocHasSideEffects*), *SkCo* (*SkinContact*), *DoBlTe* (*DoBloodTest*) and *TkEc* (*TakeEpecillin*). We represent the result of a sensing action as follows:

Definition 3.29 (Sensing Outcome)**.** A *sensing outcome* is a pair of the form $((\theta, X), V)$ for some signature (θ, X) of an s-proposition [see Def. 3.12] and some value $V \in vals(X)$. [end definition]

For example, the blood-test sensing outcome $((DoBlTe, DrHaEp), true)$ represents an outcome of an activation of the s-proposition (EP9) (see Appendix C), identified by its unique signature $(DoBlTe, DrHaEp)$, in which the blood-test (either correctly or incorrectly) shows positive.

Sensing outcomes and performances of agent actions, represent the information potentially available to the agent, and are collected together in a *sensing and acting history*. Each sensing and acting history records a possible sequence of experiences and behaviour (inputs and outputs) that the agent might have as it progresses through the time-line. Note that, this reflects our assumption that the agent is aware of its own actions:

Definition 3.30 (Sensing and Acting History, Indistinguishable Histories)**.** A *sensing and acting history* (or *history*) is a function H from $\mathcal{I}$ to the power set of sensing outcomes and agent actions. The set of all histories is denoted by $\mathcal{H}$. For a given instant I, two histories H and H' are *indistinguishable up to (and excluding)* I if $H(I') = H'(I')$ for all $I' < I$. For a given history H and instant I, the equivalence class of histories indistinguishable from H up to I is denoted by $[H]_{<I}$. [end definition]

Equivalence classes of the form $[H]_{<I}$ represent agent experiences which are identical up to I, and will be used later to evaluate whether the agent will execute belief-conditioned actions at the instant I. An example of a history (compatible with Example 1.1) is $H_1^e(-1) = H_1^e(0) = H_1^e(3) = \emptyset$, $H_1^e(1) = \{\mathit{DoBlTe}, ((\mathit{DoBlTe}, \mathit{DrHaEp}), \mathit{true})\}$ and $H_1^e(2) = \{\mathit{TkEc}\}$. In words, the experience that H_1^e encodes is that the doctor performs the blood-test (sensing action) at 1 which gives a positive result. In response, the doctor takes epecillin at 2. No other actions or sensing outcomes are experienced. An alternative history (incompatible with the example) is $H_2^e(-1) = H_2^e(0) = H_2^e(1) = H_2^e(2) = H_2^e(3) = \{\mathit{TkEc}\}$ (the doctor senses nothing, but takes epecillin at every instant).

Whereas for ne-domain descriptions the key semantic structure is a world (see Definition 3.16), in the case of epistemic domains the key semantic structure is a world paired with a history. We call such a pairing an *h-world*:

Definition 3.31 (h-world)**.** An *h-world* is a pair (W, H) for a world $W \in \mathcal{W}$ and a sensing history $H \in \mathcal{H}$. [See Defs 3.16 and 3.30] [end definition]

EPEC's semantics identifies the set of h-worlds that contain a history and world pair that are compatible both with each other and with the domain description in question – these are called *well-behaved h-worlds* – and defines a probability distribution over this set. Some complexity arises because at all points in the narrative of any particular h-world where the agent makes decisions about performing belief-conditioned actions, the strengths of those beliefs (represented as probabilities) have to be evaluated by consideration of all h-worlds and their respective probabilities. However, the as strengths of those evaluated beliefs influence whether the belief-conditioned actions fire, they thus influence the probability distribution over h-worlds on which the belief was evaluated. For example, in the (well-behaved) h-world (W_1^e, H_1^e) pictured below, the doctor's decision to take epecillin at instant 2 is based on her greater than 50% belief that she has epectitis due to the (false) positive blood test at instant 1, and this in turn results in her suffering from epecillin's side effects at instant 3. Note that it is only the history up to instant 1 that is relevant to the doctor's decision to take epecillin, and we will see that in the general case it is equivalence classes of indistinguishable histories (see Definition 3.30) up to agent decision points such as this that are key for defining EPEC's semantics.

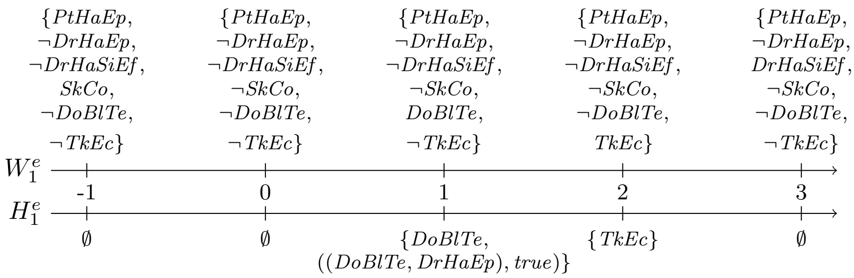

The next three definitions allow us to assign a probability to a history, and to an equivalence class of histories (indistinguishable up to a given instant), in both cases conditioned on a particular world. First, we define a *sensing occurrence* in a particular world, and the associated mapping $socc_{\mathcal{D}}$. For a given h-world, $socc_{\mathcal{D}}$ indicates which s-propositions (if any) are activated (occur) at each instant, and what both the true and the sensed values of the associated fluents are for these activations:

Definition 3.32 (Sensing Occurrence)**.** Let $\mathcal{D}$ be a domain description, $\underline{\mathbf{s}}$ be an s-proposition in $\mathcal{D}$ with body θ, and $I \in \mathcal{I}$. If $W \mid\models [\theta]@I$ then it is said that that *a sensing action occurs at instant* I *in* W *w.r.t. to* $\mathcal{D}$, and that $\underline{\mathbf{s}}$ *is activated at* I *in* W *w.r.t.* $\mathcal{D}$. Let (W, H) be an h-world. For any instant $I \in \mathcal{I}$, $socc_{\mathcal{D}}((W, H), I)$ is defined as

$$\{((\theta, X), V, V') \mid \underline{\mathbf{s}} \text{ is an s-proposition in } \mathcal{D},\ sig(\underline{\mathbf{s}}) = (\theta, X), \\ W \mid\models [\theta \wedge X = V]@I,\ ((\theta, X), V') \in H(I)\}$$

[end definition]

For example, the blood-test s-proposition (EP9) is activated at 1 in W_1^e w.r.t. $\mathcal{D}_e$, and, since $W_1^e \mid\models [DoBlTe \wedge DrHaEp = false]@1$ and $((DoBlTe, DrHaEp), true) \in H_1^e(1)$, then $socc_{\mathcal{D}_e}((W_1^e, H_1^e), 1) = \{((DoBlTe, DrHaEp), false, true)\}$.

We can now define the following well-behavedness property for h-worlds. For an h-world to be well-behaved, it must satisfy the *closed world assumption for sensing and acting* (*CWSA*). This ensures that the history gives exactly one sensed value for each of the s-propositions activated in the world (condition (i)), no sensed values where no s-proposition is activated (condition (ii)), and also that the history properly describes the agent's awareness of its own actions in the world (condition (iii)):

Definition 3.33 (CWSA)**.** Given a domain description $\mathcal{D}$, an h-world (W, H) satisfies the *closed world assumption for sensing and acting (CWSA) w.r.t.* $\mathcal{D}$ if for each $I \in \mathcal{I}$:

(i) for each s-proposition $\underline{\mathbf{s}}$ activated at I in W, $(sig(\underline{\mathbf{s}}), V) \in H(I)$ for exactly one value V,

(ii) for each $((\theta, X), V') \in H(I)$ there is an s-proposition in $\mathcal{D}$ with signature (θ, X) activated at I in W, and

(iii) for each $A \in \mathcal{A}_a$, $A \in H(I)$ if and only if $W \mid\models [A = true]@I$.

[end definition]

Inspection of the h-world (W_1^e, H_1^e) illustrated on page 31 shows that it meets the three conditions above and therefore satifies the CWSA w.r.t. $\mathcal{D}_e$.

We next address the issue of assigning an overall probability to an entire h-world (W, H). We start by assigning a probability to H conditioned on W, which we call an *evaluation* of H w.r.t. W. Definition 3.34 below says that the evaluation of H given W is the product of the accuracy matrix entries for each of the sensing results contained in H. To aid reading of this definition, recall (from Definition 3.12) that $\mathbf{M}(\theta, X)$ refers to the accuracy matrix in the s-proposition with unique signature (θ, X) (where θ is the condition of the s-proposition and X is the fluent being sensed). The element $\mathbf{M}(\theta, X)_{i,j}$ in $\mathbf{M}(\theta, X)$ represents the probability that, given that V_i is the actual value of X when θ occurs, the value V_j is sensed. Recall also (from Definition 3.32) that $socc_{\mathcal{D}}((W, H), I)$ gives a sensed value of a fluent according to H together with the actual value according to W, which in turn identify this specific element in $\mathbf{M}(\theta, X)$.

Definition 3.34 (History Evaluation)**.** Let $\mathcal{D}$ be a domain-description, (W, H) an h-world, and $\mathcal{J}_H^W = \{I \mid I \in \mathcal{I}, socc_{\mathcal{D}}((W, H), I) \neq \emptyset\}$ (i.e. $\mathcal{J}_H^W$ is the set of instants at which some sensing occurs). For all $I \in \mathcal{J}_H^W$ let

$$\mathbf{Mprod}_H^W(I) = \prod_{((\theta, X), V_i, V_j) \in socc_{\mathcal{D}}((W, H), I)} \mathbf{M}(\theta, X)_{i,j}$$

where for each expression in this product V_i and V_j are the ith and jth elements of $vals(X)$ respectively. The *evaluation of H given W w.r.t. $\mathcal{D}$* is defined as:

$$\epsilon_{\mathcal{D}}(H \mid W) = \begin{cases} 0 & \text{if } (W, H) \text{ does not satisfy the CWSA} \\ 1 & \text{if } (W, H) \text{ satisfies CWSA and } \mathcal{J}_H^W = \emptyset \\ \prod_{I \in \mathcal{J}_H^W} \mathbf{Mprod}_H^W(I) & \text{otherwise} \end{cases}$$

For a class $[H]_{<I}$ of indistinguishable sensing histories up to I, $\epsilon_{\mathcal{D}}([H]_{<I} \mid W)$ denotes the sum:

$$\sum_{H' \in [H]_{<I}} \epsilon_{\mathcal{D}}(H' \mid W)$$

[end definition]

Note that $\epsilon_{\mathcal{D}}(H \mid W)$ is equal to 1 when (W, H) satisfies the CWSA and there are no sensing occurrences, because there is no uncertainty that is coming from sensing noise. When there are sensing occurrences, the first (inner) product in the definition allows for concurrent (and potentially conflicting) sensing at any particular instant, whereas the second (outer) product allows for different instants at which sensing occurs along the time-line. The use of products reflects the fact that we consider sensing acts to be probabilistically independent (no sensing act interferes with any other, even when they are concurrent). For the h-world (W_1^e, H_1^e) illustrated on page 31, $\mathcal{J}_{H_1^e}^{W_1^e} = \{1\}$, $socc_{\mathcal{D}_e}((W_1^e, H_1^e), 1) = \{((\mathit{DoBlTe}, \mathit{DrHaEp}), \mathit{false}, \mathit{true})\}$, and so $\mathbf{M}_{\mathbf{prod}}{}_{H_1^e}^{W_1^e}(1) = \mathbf{M}(\mathit{DoBlTe}, \mathit{DrHaEp})_{0,1} = 0.1$ (see (EP9) in Appendix C). Therefore $\epsilon_{\mathcal{D}_e}(H_1^e \mid W_1^e) = 0.1$, and since in this particular case H_1^e is the only history that can be paired with W_1^e so that (W_1^e, H_1^e) satisfies the CWSA, $\epsilon_{\mathcal{D}_e}([H_1^e]_{<2} \mid W_1^e) = 0.1$.

Two key ideas underly the remaining definitions for EPEC's semantics. The first is the intuition that, from the point of view of an agent residing in world W after experiencing all the events in a particular history H, the domain description $\mathcal{D}$ can effectively be "reduced" to a corresponding ne-domain description. This is because the belief-conditioned decisions embedded in the p-propositions will have already been made, making those p-propositions no different in retrospect from o-propositions with either a *true* or a *false* condition, and the relevant probabilities from matrix entries in the s-propositions can be accounted for by the overall probability of the history H given the world W. This idea motivates the notion of a *reduct* in Definitions 3.37 and 3.38 below. The second idea is that, using the product rule $p(X, Y) = p(Y \mid X)p(X)$ for random variables X and Y from probability theory (notation here as in [30]), we can express a probability distribution over h-worlds in terms of a marginal probability $p(W)$ for each world and a conditional probability $p(H \mid W)$ for each history given a particular world. We already have two such candidate probability distributions, $M_{\mathcal{N}}^{ne}$ and $\epsilon_{\mathcal{D}}$, from Definitions 3.28 and 3.34 respectively. Definition 3.36 below utilises these to provide a general joint probability distribution called the *pre-model-function*, written $\tilde{M}_{\mathcal{D}}^{\mathcal{N}}$, with respect to any domain description $\mathcal{D}$ and ne-domain description $\mathcal{N}$. However, the subsequent focus will then be on ne-domain descriptions $\mathcal{N}$ that are reducts of $\mathcal{D}$ in the sense described informally above.

In order to re-express p-propositions as o-propositions in a reduct ne-domain description, we need to re-classify agent actions as environmental actions in the underlying domain language:

Definition 3.35 (Ne-transform of a Domain Language)**.** The *ne-transform* of a domain language $\mathcal{L} = \langle \mathcal{F}, \mathcal{A}, \mathcal{A}_e, \mathcal{A}_a, \mathcal{V}, vals, \mathcal{I}, \leq, \bar{0} \rangle$, written $\mathcal{L}^{ne}$, is the domain language $\langle \mathcal{F}, \mathcal{A}, \mathcal{A}_e \cup \mathcal{A}_a, \emptyset, \mathcal{V}, vals, \mathcal{I}, \leq, \bar{0} \rangle$. (i.e. Agent actions are re-classified as environmental actions.) [end definition]

We now give the definition of a pre-model function $\tilde{M}^{\mathcal{N}}_{\mathcal{D}}$. Intuitively, this is the probability distribution over h-worlds induced by $\mathcal{D}$ and $\mathcal{N}$ which describes the likelihood for a given world W and history H of the agent receiving the sensory input in H and evaluating it according to the s-propositions in $\mathcal{D}$, but acting and causing change according to W and evaluating this according to $\mathcal{N}$. As indicated above, for generality this abstract definition is for an arbitrary $\mathcal{D}$ and $\mathcal{N}$. But, for a given $\mathcal{D}$, our objective in subsequent definitions is to identify the unique $\mathcal{N}$ (the "reduct") that mirrors the rational behaviour of the agent according to its full specification in $\mathcal{D}$. This provides a stepping-stone to enable the full epistemic semantics of EPEC to be couched in terms of its non-epistemic part. Note that including o-propositions in $\mathcal{N}$ that mirror some of the p-propositions in $\mathcal{D}$ is the key to identifying this unique $\mathcal{N}$.

Definition 3.36 (Pre-model-function)**.** Given a domain description $\mathcal{D}$ written in language $\mathcal{L}$ and an ne-domain description $\mathcal{N}$ written in $\mathcal{L}^{ne}$, the *pre-model-function of* $\mathcal{D}$ *w.r.t.* $\mathcal{N}$ is the function $\tilde{M}^{\mathcal{N}}_{\mathcal{D}} : \mathcal{W} \times \mathcal{H} \mapsto [0, 1]$ (with $\mathcal{H}$ defined in $\mathcal{L}$) defined as:

$$\tilde{M}^{\mathcal{N}}_{\mathcal{D}}(W, H) = \begin{cases} \epsilon_{\mathcal{D}}(H \,|\, W) \cdot M^{ne}_{\mathcal{N}}(W) & \text{if } W \in \mathcal{W}_{\mathcal{N}} \\ 0 & \text{otherwise} \end{cases}$$

The domain of $\tilde{M}^{\mathcal{N}}_{\mathcal{D}}$ is extended to equivalence classes of histories and to i-formulas in the usual way:

$$\tilde{M}^{\mathcal{N}}_{\mathcal{D}}(W, [H]_{<I}) = \sum_{H' \in [H]_{<I}} \tilde{M}^{\mathcal{N}}_{\mathcal{D}}(W, H')$$

$$\tilde{M}^{\mathcal{N}}_{\mathcal{D}}(\varphi, H) = \sum_{W \| \models \varphi} \tilde{M}^{\mathcal{N}}_{\mathcal{D}}(W, H) \quad \text{and} \quad \tilde{M}^{\mathcal{N}}_{\mathcal{D}}(\varphi, [H]_{<I}) = \sum_{W \| \models \varphi} \tilde{M}^{\mathcal{N}}_{\mathcal{D}}(W, [H]_{<I})$$

[end definition]

The following proposition establishes that $\tilde{M}^{\mathcal{N}}_{\mathcal{D}}$ is indeed a probability distribution over the set of h-worlds.

Proposition 3.2. Let $\mathcal{D}$ and $\mathcal{N}$ be domain and ne-domain descriptions in languages $\mathcal{L}$ and $\mathcal{L}^{ne}$ respectively. The pre-model-function $\tilde{M}_{\mathcal{D}}^{\mathcal{N}}$ of $\mathcal{D}$ w.r.t. $\mathcal{N}$ is a probability distribution over $\mathcal{W} \times \mathcal{H}$, i.e. $\sum_{(W,H)\in\mathcal{W}\times\mathcal{H}} \tilde{M}_{\mathcal{D}}^{\mathcal{N}}(W,H) = 1$.

Proof: See Appendix D.2. ■

As an example pre-model-function, consider $\tilde{M}_{\mathcal{D}_e}^{\mathcal{N}_e}$, where $\mathcal{D}_e$ is the domain description (EP1)–(EP11) listed in Appendix C, and $\mathcal{N}_e$ is the ne-domain description that consists of propositions (EP1)–(EP8) together with the o-propositions

DoBloodTest **occurs-at** 1 (EP12)

TakeEpecillin **occurs-at** 2 (EP13)

Note that the world W_1^e illustrated on page 31 is a well-behaved world with respect to $\mathcal{N}_e$ (i.e. $W_1^e \in \mathcal{W}_{\mathcal{N}_e}$). Inspection of the probabilities embedded in (EP4), (EP5), (EP6) and (EP8) of Appendix C shows that $M_{\mathcal{N}_e}^{ne}(W_1^e) = 0.8 \times 0.95 \times 0.25 \times 0.15 = 0.0285$. The calculations on page 34 give $\epsilon_{\mathcal{D}_e}(H_1^e \mid W_1^e) = 0.1$, so $\tilde{M}_{\mathcal{D}_e}^{\mathcal{N}_e}(W_1^e, H_1^e) = 0.1 \times 0.0285 = 0.00285$.

The next definition takes us a step closer to formalising the notion of a reduct of a domain description (discussed informally on page 34) by specifying how to either eliminate a p-proposition or convert it to an o-proposition with respect to a probability distribution over worlds:

Definition 3.37 (Covers, Reduct of a p-proposition)**.** Let $\underline{\mathbf{p}}$ be the p-proposition "A **performed-at** I **with-prob** P^+ **if-believes** $(\theta, \bar{P})$" in the language $\mathcal{L}$, $\underline{\mathbf{o}}$ be the o-proposition "A **occurs-at** I **with-prob** P^+" in the language $\mathcal{L}^{ne}$, and d be a probability distribution over worlds. Then d *covers* $\underline{\mathbf{p}}$ if $\sum_{W\,|\!\models[\theta]@I} d(W) \in \bar{P}$. In this case $\underline{\mathbf{p}}$ *reduces to* the singleton set $\{\underline{\mathbf{o}}\}$ and $\{\underline{\mathbf{o}}\}$ *is the reduct of* $\underline{\mathbf{p}}$ w.r.t. d. Otherwise $\underline{\mathbf{p}}$ *reduces to* $\emptyset$ and and $\emptyset$ *is the reduct of* $\underline{\mathbf{p}}$ w.r.t. d. [end definition]

Clearly, for any d (EP10) reduces to {(EP12)}, and (EP11) reduces to either {(EP13)} or $\emptyset$ depending on whether $\sum_{W\,|\!\models[DrHaEp]@2} d(W) > 0.5$ or not.

Definition 3.38 below states that the reduct of a domain description $\mathcal{D}$ w.r.t. a history H is $\mathcal{D}$ itself but with all s-propositions removed, and each p-proposition replaced with its reduct. For each p-proposition, the probability distribution used in its reduction is conditional on the probability of the history H up to the instant of the p-proposition (expressed as an equivalence class of indistinguishable histories). This is because it is the agent's sensing and acting history up to that instant that influences its belief in the condition of the p-proposition. Definition 3.38 is stated in general terms for any probability distribution d over the set of h-worlds, but in subsequent definitions the probability distribution used in the reduction will be the pre-model-function $\tilde{M}^{\mathcal{N}}_{\mathcal{D}}$ [see Definition 3.36]. (Note that if d gives a zero probability for a particular equivalence class of indistinguishable histories then any associated conditional probability is undefined, so, for mathematical completeness, in this case we use the unconditional marginal probability d over worlds to reduce the p-proposition, although histories with zero probability do not contribute to EPEC's semantics in subsequent definitions.)

Definition 3.38 (Reduct of a Domain Description)**.** Let $\mathcal{D}$ be domain description, d be a probability distribution over $\mathcal{W} \times \mathcal{H}$ and H be a history. Then *the reduct of* $\mathcal{D}$ *w.r.t.* d *and* H is the ne-domain description obtained by first removing all the s-propositions and p-propositions from $\mathcal{D}$, and then for each removed p-proposition $\underline{\mathbf{p}}$ with instant I, conjoining the ne-domain description with the reduct of $\underline{\mathbf{p}}$ w.r.t. the probability distribution $d(\,\cdot \mid [H]_{<I})$ over $\mathcal{W}$, defined as follows:

$$d(\,\cdot \mid [H]_{<I}) = \begin{cases} \dfrac{\sum_{H' \in [H]_{<I}} d(\,\cdot\,, H')}{\sum_{H' \in [H]_{<I}} d(H')} & \text{if } \sum_{H' \in [H]_{<I}} d(H') \neq 0 \\ d(\,\cdot\,) & \text{if } \sum_{H' \in [H]_{<I}} d(H') = 0 \end{cases}$$

[end definition]

Definition 3.39 and Proposition 3.3 below show that for a given history H there is at most one "reasonable" reduct of a domain description $\mathcal{D}$, which we denote $\mathcal{R}^{\mathcal{D}}_{H}$. Definition 3.39 also implicitly gives a theoretical procedure for finding $\mathcal{R}^{\mathcal{D}}_{H}$, as follows: (i) Arbitrarily generate an ne-domain description $\mathcal{N}$ from $\mathcal{D}$ by deleting all the s-propositions, replacing some of the p-propositions by equivalent but conditionless o-propositions, and deleting the rest of the p-propositions (at most 2^n such $\mathcal{N}$s may be generated in this way, where n is the number of p-propositions in $\mathcal{D}$). (ii) Form the pre-model-function $\tilde{M}^{\mathcal{N}}_{\mathcal{D}}$ with the arbitrarily chosen $\mathcal{N}$. (iii) Reduce $\mathcal{D}$ w.r.t. $\tilde{M}^{\mathcal{N}}_{\mathcal{D}}$ and H. (iv) If step (iii) results in $\mathcal{N}$ itself, then $\mathcal{N}$ is in fact $\mathcal{R}^{\mathcal{D}}_{H}$. If it does not, go back to step (i).

Definition 3.39 (Reduct Set)**.** Let $\mathcal{D}$ be a domain description and H a history. Then the *reduct set of* $\mathcal{D}$ *w.r.t.* H, denoted $R(\mathcal{D}, H)$, is the set of ne-domain descriptions such that $\mathcal{N} \in R(\mathcal{D}, H)$ if and only if the reduct of $\mathcal{D}$ w.r.t. the pre-model-function $\tilde{M}^{\mathcal{N}}_{\mathcal{D}}$ and history H is $\mathcal{N}$ itself. **[end definition]**

Proposition 3.3. Let $\mathcal{D}$ be a domain description and H a history. Then the reduct set $R(\mathcal{D}, H)$ of $\mathcal{D}$ w.r.t. H contains at most one element. If $R(\mathcal{D}, H) \neq \emptyset$ this unique element is denoted $\mathcal{R}^{\mathcal{D}}_{H}$.

Proof: See Appendix D.3. ■

The intuition behind Proposition 3.3 is that for a completely specified sensor input stream (as defined by H) the agent's belief state is uniquely and fully determined at every instant, and therefore, since the agent must exactly follow the prescription of what to perform provided by the p-propositions in $\mathcal{D}$, this will result in a unique sequence of agent action performances (re-expressed as additional o-propositions in the reduct). The p-propositions in $\mathcal{D}$ not re-instated as o-propositions in $\mathcal{R}^{\mathcal{D}}_{H}$ are exactly those whose belief-preconditions are not met by the agent's belief state w.r.t. H.

For the "epectisis" domain $\mathcal{D}_e$, it turns out that $\mathcal{N}_e = \mathcal{R}^{\mathcal{D}_e}_{H^1_e}$ [see pages 31 and 36]. This is because plugging in the condition of (EP11) into Definitions 3.37 and 3.38 and referring to the various probabilities embedded in $\mathcal{D}_e$'s propositions gives $\sum_{W \,|\!\models [DrHaEp]@2} \tilde{M}^{\mathcal{N}_e}_{\mathcal{D}_e}(W \mid [H^e_1]_{<2}) = 0.9264$, and this value falls within $(0.5, 1]$ so that (EP11) reduces to (EP13).

We can now formally define when an h-world is well-behaved, giving a second criterion by applying Definition 3.24 of a well-behaved world w.r.t an ne-domain description using $\mathcal{R}^{\mathcal{D}}_{H}$. (Recall that the first criterion is that the h-world satisfy the CWSA as discussed on page 32):

Definition 3.40 (Well-behaved h-world)**.** Let $\mathcal{D}$ be a domain description. The h-world (W, H) is *well-behaved w.r.t.* $\mathcal{D}$ if (i) it satisfies the CWSA, and (ii) $R(\mathcal{D}, H)$ is non-empty and W is well-behaved w.r.t. $\mathcal{R}_H^{\mathcal{D}}$. [end definition]

Since by inspection W_1^e is well-behaved w.r.t. $\mathcal{N}_e = \mathcal{R}^{\mathcal{D}_e}_{H_e^1}$, then (W_1^e, H_1^e) is well-behaved w.r.t. $\mathcal{D}_e$.

We can now state the main overarching definition for the semantics of EPEC. Definition 3.41 defines the *model-function* mapping from h-worlds to $[0, 1]$ (analogous to the ne-model function for ne-domains – see Definition 3.28), and Proposition 3.4 immediately after confirms that the model-function is a probability distribution over h-worlds:

Definition 3.41 (Model Function)**.** The *model-function of a domain description* $\mathcal{D}$ is the function $M_{\mathcal{D}} : \mathcal{W} \times \mathcal{H} \mapsto [0, 1]$ defined as follows:

$$M_{\mathcal{D}}(W, H) = \begin{cases} \tilde{M}_{\mathcal{D}}^{\mathcal{R}_H^{\mathcal{D}}}(W, H) & \text{if } (W, H) \text{ is well-behaved w.r.t. } \mathcal{D} \\ 0 & \text{otherwise} \end{cases} \tag{16}$$

[end definition]

Proposition 3.4. Let $\mathcal{D}$ be a domain description. Then the model-function $M_{\mathcal{D}}$ of $\mathcal{D}$ is a probability distribution over $\mathcal{W} \times \mathcal{H}$, i.e. $\sum_{(W,H) \in \mathcal{W} \times \mathcal{H}} M_{\mathcal{D}}(W, H) = 1$.

Proof: See Appendix D.4. ■

Since $M_{\mathcal{D}}$ is a joint probability distribution over worlds and histories, we can refer to marginal probabilities $M_{\mathcal{D}}(W)$, $M_{\mathcal{D}}(H)$, $M_{\mathcal{D}}([H]_{<I})$, etc. and conditional probabilities $M_{\mathcal{D}}(W \mid H)$, $M_{\mathcal{D}}(H \mid W)$, $M_{\mathcal{D}}(W \mid [H]_{<I})$ etc. using the standard definitions. Like $M_{\mathcal{N}}^{ne}$, $M_{\mathcal{D}}$ can also be straightforwardly extended to provide probabilities for i-formulas: $M_{\mathcal{D}}(\varphi, H) = \sum_{W \models \varphi} M_{\mathcal{D}}(W, H)$.

At the beginning of Section 3 on page 10 we informally described how domain descriptions entail *b-propositions*. We can now formally define this entailment, starting with the definition of the general form of a b-proposition:

Definition 3.42 (b-proposition)**.** A *b-proposition* has the form

$$\textbf{at } I \textbf{ believes } \varphi \textbf{ with-probs } \{([H_1], B_1, P_1), \ldots, ([H_m], B_m, P_m)\}$$

for some instant I, i-formula φ, and real numbers $B_1, \ldots, B_m \in (0, 1]$ and $P_1, \ldots, P_m \in [0, 1]$ such that $\sum_{i=1}^{m} B_i = 1$. Each $[H_i]$ is an equivalence class $[H_i]_{<I}$ of histories represented as "$\langle H_i(I_1)@I_1, \ldots, H_i(I_n)@I_n \rangle$" for instants $I_1, \ldots, I_n$ such that, for $1 \leq j \leq n$, $I_j < I$ and $H_i(I_j) \neq \emptyset$, and with repetitions of actions entailed by a sensing outcome in $H_i(I_j)$ removed[5]. [end definition]

We use the model-function $M_{\mathcal{D}}$ to ascertain when a b-proposition is entailed by a domain descripion $\mathcal{D}$:

Definition 3.43 (Entailment for Domain Descriptions)**.** The b-proposition "$\textbf{at } I \textbf{ believes } \varphi \textbf{ with-probs } \{([H_1], B_1, P_1), \ldots, ([H_m], B_m, P_m)\}$" *is entailed by the domain description* $\mathcal{D}$ iff, for $1 \leq i \leq m$, $M_{\mathcal{D}}(\varphi \mid [H_i]_{<I}) = P_i$ and $M_{\mathcal{D}}([H_i]_{<I}) = B_i$. [end definition]

Intuitively, if a domain description $\mathcal{D}$ entails a b-proposition

$$\textbf{at } I \textbf{ believes } \varphi \textbf{ with-probs } \{([H_1], B_1, P_1), \ldots, ([H_m], B_m, P_m)\}$$

this means that at instant I the agent will believe that the i-formula φ holds with one of the probabilities $P_1, \ldots, P_m$ depending on the input/output from its sensors/effectors, which is "recorded" in $[H_i]$ and has an associated probability B_i of actually being experienced/executed. One example of a b-proposition entailed by $\mathcal{D}_e$ is that on page 11, repeated here for convenience:

$$\begin{array}{l} \textbf{at } 3 \textbf{ believes } [\neg \mathit{DrHaEp}]@3 \textbf{ with-probs} \\ \quad \{\ (\langle \{((\mathit{DoBlTe}, \mathit{DrHaEp}), \mathit{false})\}@1 \rangle, 0.425, 0.9314), \\ \quad\ \ (\langle \{((\mathit{DoBlTe}, \mathit{DrHaEp}), \mathit{true})\}@1, \{\mathit{TkEc}\}@2 \rangle, 0.575, 0.9907)\ \} \end{array}$$

(see page 11 for a natural language interpretation of this). Other examples are:

$$\textbf{at } 0 \textbf{ believes } [\mathit{DocHasEpectisis} = \mathit{true}]@0 \textbf{ with-probs } \{(\langle\rangle, 1, 0.57)\}$$

which reads "the doctor currently (instant 0) has a 57% belief that she has epectisis", and

[5] For example, $\{((\mathit{DoBlTe}, \mathit{DrHaEp}), \mathit{false}), \mathit{DoBlTe}\}@1$ has the repeated DoBlTe removed so that it is written simply as $\{((\mathit{DoBlTe}, \mathit{DrHaEp}), \mathit{false})\}@1$.

$$\textbf{at } 0 \textbf{ believes } [\mathit{DocHasSideEffects} = \mathit{true}]@3 \textbf{ with-probs } \{(\langle\rangle, 1, 0.087675)\}$$

which reads "the doctor currently (instant 0) has a 8.77% belief that she will have side effects after executing her plan (instant 3)".

4. Comparison with Existing Formalisations

In this section we examine the relationship of our approach to existing formalisations for representing dynamic probabilistic domains. In sub-sections 4.1 and 4.2 we demonstrate a formal equivalence between EPEC and two established frameworks for probabilistic reasoning about actions, in cases where the classes of representable domain features coincide. In both cases we do this by showing that probabilistic entailment is preserved under a general translation procedure into EPEC. The two formalisations that we translate are PAL ("Probabilistic Action Language") [6], and the extended Situation Calculus framework proposed by Bacchus, Halpern and Levesque [3] which we will refer to here as "BHL" after its authors. In sub-section 4.3 we present a *feature chart* to summarise and contrast the types of domain phenomena that these and various other frameworks are able to represent and reason about.

4.1. PAL and EPEC

PAL is a widely cited, probabilistic version of the language $\mathcal{A}$ [19], and as such inherits a basic ontology of actions and fluents, thus supporting reasoning about the effects of (all) hypothetical action sequences arranged in a situation-calculus-like tree structure, as opposed to the likelyhood and likely consequences of actual action occurrences embedded in an independent time line (i.e. probabilistic narrative reasoning). Its development was motivated by a desire for an elaboration tolerant representation of Markov decision processes, and was influenced by Pearl's work on *functional causal models* [29]. The authors remark that "*our formulation can be considered as a generalization of Pearl's formulation of causality to a dynamic setting with a more elaboration tolerant representation*". PAL does not support epistemic features such as sensing and belief-conditioned actions, and so is compared here with non-epistemic EFEC.

4.1.1. PAL Syntax and Semantics

For readability and to avoid ambiguity, in the remainder of Section 4.1, PAL terms appear with an overhead dot ($\dot{}$) wherever they might otherwise be confused with a similar EPEC term.

A PAL language signature is divided into four sorts: $\dot{\mathbf{F}}$ for *fluents*, $\dot{\mathbf{U}}_I$ for *inertial unknown variables*, $\dot{\mathbf{U}}_N$ for *non-inertial unknown variables*, and $\dot{\mathbf{A}}$ for *actions*. A PAL *theory* is a set of *propositions*, each of which belongs to one of three categories: the *domain description language* PAL_D, the probability description language PAL_P, or the *observation language* PAL_O. The semantics of PAL defines an entailment relation $\models_{PAL}$ between PAL theories and *queries* defined in a fourth category called the *query language* PAL_Q.

We next describe the general forms of propositions in each of these four PAL categories, using the same conventions as in [6] – that $\dot{\psi}$ is a fluent formula, $\dot{\varphi}$ is a formula of fluents and unknown variables (both inertial and non-intertial), $\dot{\theta}$ is a formula of fluent and inertial variables, $\dot{u}$ is an unknown variable, $\dot{a}$ is an action, and $\dot{n}$ is a real value in $[0, 1]$. PAL_D consists of *dynamic causal laws* of the form "$\dot{a}$ **causes** $\dot{\psi}$ **if** $\dot{\varphi}$", *static causal laws* of the form "$\dot{\theta}$ **causes** $\dot{\psi}$", and *executability conditions* of the form "**impossible** $\dot{a}$ **if** $\dot{\varphi}$". PAL_P consists of propositions of the form "**probability of** $\dot{u}$ **is** $\dot{n}$". PAL_O consists of propositions of the form "$\dot{\psi}$ **obs_after** $\dot{a_1}, \ldots, \dot{a_n}$", with "**initially** $\dot{\psi}$" being a shorthand for "$\dot{\psi}$ **obs_after** $\dot{a_1}, \ldots, \dot{a_n}$" when $\dot{n} = 0$. Finally, PAL_Q consists of *queries* of the form "**probability of** $[\dot{\varphi}$ **after** $\dot{a_1}, \ldots, \dot{a_n}]$ **is** $\dot{n}$".

For full formal details of the semantics of PAL the reader is referred to [6], but we summarise them here, highlighting the parts and notation needed to describe the relationship of PAL to EPEC. Fluents and unknown variables are Boolean-valued, and a *state* $\dot{s}$ is an interpretation of fluents and unknown variables that obeys the constraints imposed by the static causal laws in PAL_D, regarded as classical implications. The *u-state* $\dot{s}_u$ is $\dot{s}$ but excluding the interpretation of fluents, the set of states $\dot{I}(\dot{s}_u)$ is defined as $\{\dot{s}' \mid \dot{s}'_u = \dot{s}_u\}$, and $\dot{s}_N$ signifies the interpretation of only the non-inertial unknown variables in $\dot{s}$.

Following action language convention, the propositions in PAL_D characterise a unique transition function $\dot{\Phi}$ such that $\dot{\Phi}(\dot{a}, \dot{s})$ is the set of states that may be reached after executing action $\dot{a}$ in state $\dot{s}$. $\dot{\Phi}$ is defined, and in particular takes into account static causal laws, in much the same way as in earlier work by McCain and Turner [32], and is recursively extended to sequences of actions in a standard way. However, $\dot{\Phi}$ imposes no constraints on the way non-inertial unknown variables change during a transition, thus allowing these variables to be used to introduce random factors into actions' effects, with frequencies controlled by the probability distribution defined via the propositions in PAL_P. Inertial unknown variables play a similar randomising role, but their values cannot change at all during a sequence of state transitions so that they effectively introduce an random element only to the intial state. Finally, the hypothetical observations in PAL_O indirectly filter out some of the otherwise possible combinations of initial values for fluents and unknown variables, so that the entailed probabilities in queries are conditional probabilities, conditioned the probability of initially being in a state not excluded by PAL_O propositions.

Only a limited number of the EPEC propositions (EP1)-(EP11) listed in Appendix C (modelling the Epectisis Example 1.1) are translatable into PAL. The v-propositions (EC1)-(EC3) are implicit in PAL since all fluents and unknown variables are Boolean. The i-proposition (EC4) is not translatable into PAL, because PAL does not provide a mechanism to assign differing probabilities to intial values of fluents.[6] The narrative proposition (EP5), stating that the doctor probably had skin contact with the patient, can also not be represented in PAL, since, although [6] includes some brief proposals about extending PAL with narrative capabilities, these do not include the possibility of attaching probabilities to action occurrences. And since PAL does not model epistemic features, (EP9)–(EP11) are also not translatable. However, PAL can encapsulate (EP6)–(EP8), modelling the basic causal information in this domain, in a

[6]One way to approximate this functionality in PAL might be to introduce an "initialising" action into the language, with probabilistic effects, and insist that each query included this as the first action in the query's action sequence.

succinct fashion as follows:

$$\dot{SkinContact}\ \textbf{causes}\ \dot{DocHasEpectisis}\ \textbf{if}\ (\dot{PatientHasEpectisis} \wedge \dot{Inertial}_1) \quad \text{(PAL-EP1)}$$

$$\textbf{probability of}\ \dot{Inertial}_1\ \textbf{is}\ 0.75 \quad \text{(PAL-EP2)}$$

$$\dot{TakeEpecillin}\ \textbf{causes}\ \neg\dot{DocHasEpectisis}\ \textbf{if}\ (\dot{DocHasEpectisis} \wedge \dot{Inertial}_2) \quad \text{(PAL-EP3)}$$

$$\textbf{probability of}\ \dot{Inertial}_2\ \textbf{is}\ 0.99 \quad \text{(PAL-EP4)}$$

$$\dot{TakeEpecillin}\ \textbf{causes}\ \dot{DocHasSideEffects}\ \textbf{if}\ \dot{Inertial}_3 \quad \text{(PAL-EP5)}$$

$$\textbf{probability of}\ \dot{Inertial}_3\ \textbf{is}\ 0.15 \quad \text{(PAL-EP6)}$$

We can also model the assumption that the doctor does not initially have epectisis with the hypothetical observation

$$\textbf{initially}\ \neg\dot{DocHasEpectisis} \quad \text{(PAL-EP7)}$$

(PAL-EP1)–(PAL-EP7) together entail queries such as

$$\textbf{probability of}\ [\dot{DocHasEpectisis}\ \textbf{after}\ \dot{SkinContact}, \dot{TakeEpecillin}]\ \textbf{is}\ (0.5 \times 0.75 \times 0.01)$$

The 0.5 in the above proposition reflects the fact that PAL's semantics automatically assigns a 0.5 initial unconditional probability to $\dot{PatientHasEpectisis}$, as it does to all fluents in the absence of any static causal laws. This example also illustrates a limitation to the sense in which PAL is elaboration tolerant: the two possible effects of taking epecillin – curing epectisis and giving side-effects – are able to be represented in two separate dynamic causal laws, (PAL-EP3) and (PAL-EP5). But this separation relies on the fact that the two effects are probabilistically independent and so can be modelled using two (probabilistically independent) inertial unknown variables $\dot{Inertial}_2$ and $\dot{Inertial}_3$.

To summarise, unlike EPEC, PAL does not support the modelling of domains that include probabilistic information about initial fluent values, probabilistic information about action occurrences, probabilistic events in the environment, non-boolean fluents, concurrent actions, or epistemic features such as sensing and belief-conditioned actions. On the other hand, EPEC does not support counterfactual probabilistic reasoning (enabled in PAL via hypothetical observations). However, we show in the next sub-sections, culminating in Proposition 4.1, that PAL domains where all hypothetical observations concern the initial state can be encoded as equivalent EPEC domains, with the action sequences in PAL queries interpreted as simple EPEC narratives, and that in these cases the numerical probabilities in PAL and EPEC entailments coincide.

For the remainder of Section 4.1, we assume a finite PAL language signature $\langle \dot{\mathbf{F}}, \dot{\mathbf{U}}_I, \dot{\mathbf{U}}_N, \dot{\mathbf{A}} \rangle$, a PAL theory $\mathcal{P}$ consisting of a finite number of PAL_D and PAL_P propositions and a single PAL_O proposition "**initially** $\dot{\psi}_0$", and a PAL query $\mathcal{Q}$ ="**probability of** $[\dot{\varphi}_q$ **after** $\dot{a}_1, \ldots, \dot{a}_k]$ **is** n". We define a corresponding EPEC domain language (see Definition 3.1) with respect to $\mathcal{P}$ and $\mathcal{Q}$ as $\mathcal{L}(\mathcal{P}, \mathcal{Q}) = \langle \dot{\mathbf{F}} \cup \dot{\mathbf{U}}_I \cup \dot{\mathbf{U}}_N,\ \dot{\mathbf{A}},\ \dot{\mathbf{A}},\ \emptyset,\ \langle \mathit{false}, \mathit{true} \rangle,\ \{0, \ldots, k\},\ \leq,\ 0 \rangle$. In other words, we regard all PAL fluents and unknown variables as EPEC fluents, all PAL actions as EPEC environmental actions (because PAL is non-epistemic), and form the set of EPEC instants from the number of actions in the PAL query $\mathcal{Q}$ under consideration plus an initial instant 0. Without loss of generality we also assume that for any two distinct dynamic causal laws "$\dot{a}$ **causes** $\dot{\psi}$ **if** $\dot{\varphi}$" and "$\dot{a}$ **causes** $\dot{\psi}'$ **if** $\dot{\varphi}'$" in PAL_D it is the case that $\dot{\varphi} \models \neg\dot{\varphi}'$, and that for each action $\dot{a}$ there is exactly one executability condition "**impossible** $\dot{a}$ **if** $\dot{\varphi}_a$" (if $\dot{a}$ is always executable then $\dot{\varphi}_a = \bot$). [7] We will refer to a PAL theory that meets all of these requirements as being *normalised*, and futhermore assume that $\mathcal{P}$ is *consistent* in the sense that PAL_D permits at least one state $\dot{s}$ such that $\dot{s} \models \dot{\psi}_0$ and PAL_P contains exactly one proposition for each unknown variable

[7] Any pair of PAL dynamic causal laws "$\dot{a}$ **causes** $\dot{\psi}$ **if** $\dot{\varphi}$" and "$\dot{a}$ **causes** $\dot{\psi}'$ **if** $\dot{\varphi}'$" can be re-written as the three laws "$\dot{a}$ **causes** $\dot{\psi} \wedge \dot{\psi}'$ **if** $\dot{\varphi} \wedge \dot{\varphi}'$", "$\dot{a}$ **causes** $\dot{\psi}$ **if** $\dot{\varphi} \wedge \neg\dot{\varphi}'$" and "$\dot{a}$ **causes** $\dot{\psi}'$ **if** $\neg\dot{\varphi} \wedge \dot{\varphi}'$". Any pair of executability conditions "**impossible** $\dot{a}$ **if** $\dot{\varphi}_a$" and "**impossible** $\dot{a}$ **if** $\dot{\varphi}_a{}'$" can be re-written as "**impossible** $\dot{a}$ **if** $\dot{\varphi}_a \vee \dot{\varphi}_a{}'$".

that specifies a unique probability in the range $(0, 1)$.

4.1.2. PAL to EPEC Domain Encoding

We form an EPEC domain description $\mathcal{N}(\mathcal{P}, \mathcal{Q})$ from the PAL domain theory $\mathcal{P}$ and query $\mathcal{Q}$ as follows. The unique i-proposition in $\mathcal{N}(\mathcal{P}, \mathcal{Q})$ is

$$\textbf{initially-one-of } \{(\dot{s_1}, \dot{P}(\dot{s_1})), \ldots, (\dot{s_n}, \dot{P}(\dot{s_n}))\}$$

where $\{\dot{s_1}, \ldots, \dot{s_n}\}$ is the set of PAL states (defined via the static causal laws in PAL_D), and $\dot{P}(\dot{s_i})$ is the unconditional probability of state $\dot{s_i}$ as defined in equation (0.7) of [6] (i.e. $\dot{P}(\dot{s}) = \frac{\dot{P}(\dot{s}_u)}{|I(\dot{s}_u)|}$, where $\dot{P}(\dot{s}_u)$ is a simple product of probabilities/complement-probabilities taken from PAL_P propositions). For each action $\dot{a_i}$ appearing in $\mathcal{Q}$ = "**probability of** $[\dot{\varphi}_q$ **after** $\dot{a_1}, \ldots, \dot{a_k}]$ **is** n" we include the o-proposition "$\dot{a_i}$ **occurs-at** $i - 1$ **if-holds** $\neg\dot{\varphi_{a_i}}$".

To translate each PAL dynamic causal law into a set of EPEC c-propositions, we first define the formula $only(\dot{a})$ for each $\dot{a} \in \dot{\mathbf{A}}$ as follows:

$$only(\dot{a}) = \dot{a} \wedge \left(\bigwedge_{\dot{a}' \in \dot{\mathbf{A}}, \dot{a}' \neq \dot{a}} \neg \dot{a}' \right)$$

Given the dynamic causal law "$\dot{a}$ **causes** $\dot{\psi}$ **if** $\dot{\varphi}$" in PAL_D, then for each state $\dot{s}$ such that $\dot{s} \models \dot{\varphi} \wedge \neg\dot{\varphi}_a$ and $\dot{\Phi}(\dot{a}, \dot{s}) = \{\dot{s_1}, \ldots, \dot{s_m}\}$, $\mathcal{N}(\mathcal{P}, \mathcal{Q})$ includes the c-proposition

$$only(\dot{a}) \wedge \dot{s} \textbf{ causes-one-of } \{(\dot{s_1}, \dot{P}_{\dot{a}}(\dot{s_1} \,|\, \dot{s})), \ldots, (\dot{s_m}, \dot{P}_{\dot{a}}(\dot{s_m} \,|\, \dot{s}))\}$$

where $\dot{P}_{\dot{a}}(\dot{s}' \,|\, \dot{s})$, the conditional probability of transitioning to $\dot{s}'$ due to $\dot{a}$ given $\dot{s}$, is as defined in equation (0.9) of [6] (i.e. $\dot{P}_{\dot{a}}(\dot{s}' \,|\, \dot{s}) = \frac{2^{|\dot{\mathrm{U}}_N|}}{|\dot{\Phi}(\dot{a}, \dot{s})|} \cdot \dot{P}(\dot{s}'_N)$).[8]

[8]There is a typographical error in equation (0.9) of [6], where $\frac{2^{|\dot{\mathrm{U}}_N|}}{|\dot{\Phi}(\dot{a}, \dot{s})|}$ is incorrectly written as $\frac{|\dot{\Phi}(\dot{a}, \dot{s})|}{2^{|\dot{\mathrm{U}}_N|}}$. This is corrected by the authors in [33].

4.1.3. Correspondence between entailments

Given the above encoding from PAL domain theories to EPEC ne-domains, we can now state an equivalence between PAL and EPEC entailments as follows:

Proposition 4.1 (PAL/EPEC Correspondence)**.** Let $\mathcal{P}$ be the normalised, consistent PAL theory $PAL_D \cup PAL_P \cup \{$"**initially** $\dot{\psi}_0$"$\}$, let $\mathcal{Q}$ be the PAL query "**probability of** $[\dot{\varphi}_q$ **after** $\dot{a}_1, \ldots, \dot{a}_k]$ **is** n", let $\mathcal{N}(\mathcal{P}, \mathcal{Q})$ be the EPEC encoding of $\mathcal{P}$ and $\mathcal{Q}$, and let α be the EPEC i-formula $[\dot{a}_1]@0 \wedge \ldots \wedge [\dot{a}_k]@k-1$. Then $\mathcal{P} \models_{PAL} \mathcal{Q}$ if and only if

$$\Big(\mathcal{N}(\mathcal{P}, \mathcal{Q}) \Big| [\dot{\psi}_0]@0\Big) \;\;|\!\!\models\; \alpha \wedge [\dot{\varphi}_q]@k \textbf{ holds-with-prob } n$$

Proof: See Appendix D.5. ■

Note that in the case where one or more of $\dot{a}_1, \ldots, \dot{a}_k$ are non-executable, both EPEC and PAL will calculate n as equal to 0. This is ensured in the EPEC entailment by the inclusion of α as a conjunct in the h-proposition. Conversely, when no actions have executability conditions α is true in all EPEC well-behaved worlds and so has a probability of 1.

Proposition 4.1 shows a general correspondence between PAL and EPEC even when PAL domains include static causal laws (often referred to elsewhere as *ramifications*). For the sake of generality, the particular EPEC encodings of PAL domains used in Proposition 4.1 are very much less compact than the PAL representations. However, in practice, for many domains (especially those without static causal laws) these EPEC representations can, by inspection, be reduced to equivalent EPEC domains that are very much more compact. For example, [6] includes the following probabilistic version $\mathcal{P}_{ys}$ of the Yale Shooting Problem as a PAL domain:

$\dot{Shoot}$ **causes** $\neg \dot{Alive}$ **if** $\dot{Loaded} \wedge \dot{u}_1$ (PAL-YS1)

$\dot{Load}$ **causes** $\dot{Loaded}$ **if** $\dot{u}_2$ (PAL-YS2)

probability of $\dot{u}_1$ **is** p_1 (PAL-YS3)

$$\textbf{probability of } \dot{u}_2 \textbf{ is } p_2 \tag{PAL-YS4}$$

$$\textbf{initially } \dot{\mathit{Alive}} \wedge \neg \dot{\mathit{Loaded}} \tag{PAL-YS5}$$

and shows that $\mathcal{P}_{ys} \models_{PAL}$ **probability of** $[\dot{\mathit{Alive}}$ **after** $\dot{\mathit{Load}}, \dot{\mathit{Shoot}}]$ **is** $1 - p_1.p_2$. The EPEC encoding procedure described above encodes (PAL-YS1) and (PAL-YS2) as a total of 12 separate c-propositions, and includes each of the 16 PAL states as an initial outcome in the EPEC i-proposition. However, clearly in this case the domain can be reduced to an equivalent EPEC ne-domain $\mathcal{N}_{ys}$ that dispenses with the unknown variables, as follows:

$$\mathit{Shoot} \wedge \neg \mathit{Load} \wedge \mathit{Loaded} \textbf{ causes-one-of } \{(\{\neg \mathit{Alive}\}, p_1), (\emptyset, 1 - p_1)\} \tag{YS1}$$

$$\mathit{Load} \wedge \neg \mathit{Shoot} \textbf{ causes-one-of } \{(\{\mathit{Loaded}\}, p_2), (\emptyset, 1 - p_2)\} \tag{YS2}$$

$$\textbf{initially-one-of } \{(\{\mathit{Alive}, \neg \mathit{Loaded}\}, 1)\} \tag{YS3}$$

$$\mathit{Load} \textbf{ occurs-at } 0 \tag{YS4}$$

$$\mathit{Shoot} \textbf{ occurs-at } 1 \tag{YS5}$$

It is straightforward to show that $\mathcal{N}_{ys} \mathrel{|\!\!\models} [\mathit{Alive}]@2$ **holds-with-prob** $1 - p_1 \cdot p_2$.

4.2. BHL and EPEC

We now turn our attention to the "BHL" extension of the Situation Calculus (SC) described by Bacchus, Halpern and Levesque in [3] to enable probabilistic reasoning about noisy sensors and effectors. BHL builds upon the SC-based solution to the frame problem employed in Reiter's *basic action theories* (BATs) (see e.g. [11]). We assume here that the reader is somewhat familiar with this standard variant of the situation calculus, and so restrict ourselves to a few summarising remarks[9]. BATs are sorted classical logic theories with sorts for *actions* and *situations*. Boolean-valued fluents are represented by predicates and non-Boolean fluents by functions, and in both cases these are parameterised by

[9] Additionally, Bacchus et al. offer a concise four page summary of BATs in Section 2 of [3].

a situation argument thus allowing their (truth-)value to change from situation to situation. For example, the formula[10] $\mathit{Holding}(x, s) \land \neg \mathit{Broken}(x, s)$ might represent that the agent is holding unbroken object x in situation s. The term $do(a, s)$ represents the situation arising from performing action a in situation s, so that $do(\mathsf{drop}(x), s)$ might signify the situation resulting from dropping object x in situation s. Foundational axioms ensure that all situations are either *initial*[11] or equal to a nested *do* term rooted at an initial situation, that there is an *actual initial situation* S_0 and that situations rooted at different initial situations or arising from different action sequences are distinct. A unique *precondition axiom* for each action A of the form $\textsc{Poss}(\mathsf{A}, s) \equiv \phi(\mathsf{A}, s)$ for some formula ϕ specifies the necessary and sufficient characteristics of A and s for A to be executable in s, e.g. $\textsc{Poss}(\mathsf{drop}(x), s) \equiv \mathit{Holding}(x, s)$. The various ways in which actions cause changes in the world are captured by *effect axioms*, such as $\textsc{Poss}(\mathsf{drop}(x), s) \rightarrow \mathit{Broken}(x, do(\mathsf{drop}(x), s))$. Reiter's solution to the frame problem involves mechanically transforming the collection of effect axioms into a set of *successor-state* axioms, one for each fluent, that completely specify the state of each fluent after any particular executable action is performed in any particular situation. For example, the successor state axiom for *Broken* might be $\textsc{Poss}(a, s) \rightarrow [\mathit{Broken}(x, do(a, s)) \equiv (a = \mathsf{drop}(x) \lor (\mathit{Broken}(x, s) \land a \neq \mathsf{mend}(x)))]$. BATs also include uniqueness-of-names axioms for actions and, optionally, axioms (partially) describing the initial situation.

In addition to building on BATs, BHL utilises notation developed to represent "complex", possibly non-deterministic actions in the SC-based GOLOG programming language [34]. Complex action abbreviations are defined recursively using the primitive actions in BATs as a base case. The abbreviation $Do(\delta, s, s^+)$ signifies that s^+ is a possible situation arising from executing the complex action δ in situation s. For a primitive (deterministic) action a, $Do(a, s, s^+)$ is

[10] In all formulas, free variables are taken to be implicitly universally quantified with maximum scope.

[11] This is a relaxation of Reiter's original constraint that there be a unique initial situation S_0, necessary to model epistemic domains.

an abbreviation for $\textsc{Poss}(a, s) \wedge s^{+} = do(a, s)$. GOLOG, and hence BHL, represents a non-deterministic action as an arbitrary choice among a collection of deterministic actions. In particular, if δ_x is an action term containing the free variable x then the complex action term $\pi x.\delta_x$ represents the more specific action formed by arbitrarily choosing a particular value for x. Hence $Do(\pi x.\delta_x, s, s^{+})$ is an abbreviation for $\exists x.Do(\delta_x, s, s^{+})$. Other complex action abbreviations are defined via the Do predicate in a similar way: $Do(\delta_1, s, s^{+}) \vee Do(\delta_2, s, s^{+})$ abbreviates to $Do(\delta_1 \mid \delta_2, s, s^{+})$ to represent non-deterministic binary choice, and $Do([\delta_1; \delta_2], s, s^{+})$ is an abbreviation for $\exists s'.(Do(\delta_1, s, s') \wedge Do(\delta_1, s', s^{+}))$, to represent sequencing.

Bacchus et al. illustrate BHL's approach to representing noisy sensors and effectors with a running example of a robot moving along an unbounded one-dimensional surface, supposing that the robot has only imperfect control of its movement via a "noisy-advance" action, and only imperfect ability to ascertain its current position via a "noisy-sense-position" sensing action. We employ this same example to summarise BHL's features here. The fact that in attempting to move x units the robot might in fact move y units is modelled with the definitional axiom $\mathsf{noisy\text{-}advance}(x) \overset{\text{def}}{=} \pi y.\mathsf{advance}(x, y)$, where the primitive action $\mathsf{advance}(x, y)$ means "actually move y units while attempting to move x units". A bound b on the numerical difference between x and y is achieved with the $\textsc{Poss}$ predicate: $\textsc{Poss}(\mathsf{advance}(x, y), s) \equiv |x - y| \leq b$. The robot can only perform $\mathsf{advance}(x, y)$ indirectly via $\mathsf{noisy\text{-}advance}(x)$. A similar mechanism is employed to model the robot's limited ability to sense its current position. The primitive sensing action $\mathsf{sense\text{-}position}(x, y)$ means "sense that the position is x when the actual position is y". In this case the robot has no direct control over either x or y, since y is defined by the current situation and x by the degree of reliability of its sensor. Hence the robot can only directly perform a parameter-free noisy-sense-position action defined by the expression $\mathsf{noisy\text{-}sense\text{-}position} \overset{\text{def}}{=} \pi x, y.\mathsf{sense\text{-}position}(x, y)$, with an associated precondition axiom $\textsc{Poss}(\mathsf{sense\text{-}position}(x, y), s) \;\equiv\; y = \mathit{position}(s) \;\wedge\; |x - y| \leq c$. We refer here to actions such as $\mathsf{noisy\text{-}advance}(x)$ and noisy-sense-position that directly correspond to activations of the robot's affectors and sensors as *agent-performable*. We follow BHL terminology in referring to the first x argument in $\mathsf{advance}(x, y)$ and $\mathsf{sense\text{-}position}(x, y)$ as the *nominal* value of these primitive non-agent-performable actions, and the second y argument as the *actual* value.

The epistemic features of BHL build upon the work of Moore [1] and Scherl and Levesque [21]. A binary fluent K is included in the language as a dynamic accessibility relation between situations. An agent is deemed to know in situation s that a situation-parameterised formula ϕ is true precisely if it is true of all situations s' such that $K(s', s)$. The formula $\forall s'.K(s', s) \rightarrow \phi[s']$ expressing this is abbreviated in BHL to $\textsc{Know}(\phi[s_{know}], s)$, where s_{know} is a special "placeholder" situation term. The successor-state axiom for K is a modification of the successor-state axiom in [21] that takes into account the agent's awareness of the imprecision of its effectors (ordinary actions) and sensors (sensing actions) using a notion of *observational indistinguishability* (OI), explained below.

Observational indistinguishability is modelled using a ternary predicate $\textsc{Oi}$. The expression $\textsc{Oi}(a, a', s)$ signifies that in situation s the agent cannot distinguish between actions a and a'. A BHL theory is assumed to include a collection of *observation-indistinguishability-axioms* (OIAs) of the form $\textsc{Oi}(a, a', s) \equiv \phi(a, a', s)$, one for each primitive action a. An agent-performable complex action that includes the possibility of executing a must therefore also include the alternative possibilities of executing any action observation-indistinguishable from a. In the robot example, the OIA for $\mathsf{advance}$ is $\textsc{Oi}(\mathsf{advance}(x, y), a', s) \equiv \exists y'.a' = \mathsf{advance}(x, y')$, and the OIA for $\mathsf{sense\text{-}position}$ is $\textsc{Oi}(\mathsf{sense\text{-}position}(x, y), a', s) \equiv \exists y'.a' = \mathsf{sense\text{-}position}(x, y')$, reflecting that the robot is only aware of the nominal values, and not the actual values, of these two actions. The successor-state axiom for K, which takes observational indistinguishability into account and reflects the fact that noisy actions decrease the precision of knowledge, is:

$$\textsc{Poss}(a, s) \rightarrow \\ [K(s'^{+}, do(a, s)) \equiv \exists a', s'.(\textsc{Poss}(a', s') \wedge s'^{+} = do(a', s') \wedge \textsc{Oi}(a, a', s) \wedge K(s', s))]$$

To allow the addition of probabilistic information to a domain description, BHL incorporates two real-valued functions, p and ℓ. The functional fluent $p(s', s)$ is analogous to $K(s', s)$ and represents the relative degree of belief or *weight* the agent assigns situation s' from the standpoint of s. The initial relationship between p and K is given by the extra foundational axiom $\forall s.p(s, S_0) \geq 0 \wedge (\neg K(s, S_0) \rightarrow p(s, S_0) = 0)$, ensuring that initial situations known to be impossible are given a zero weighting. The additional logical apparatus involving the function ℓ described below ensures that analogous constraints continue to hold in all subsequent situations. The agent's degree of belief in a formula ϕ when in situation s is written $\text{BEL}(\phi[s_{know}], s)$, which is an abbreviation for the expression:

$$\frac{\sum_{\{s':\phi[s']\}} p(s', s)}{\sum_{s'} p(s', s)}$$

Although it is not part of the notation in [3], in the discussion below we will refer to the denominator $\sum_{s'} p(s', s)$ in this expression as $\text{PROB}(s)$, noting that $\text{BEL}(\phi[s_{know}], s)$ is undefined when $\text{PROB}(s) = 0$. For a complex action term δ we will write $\text{PROB}(\delta, s) = r$ as an abbreviation for $Do(\delta, s, s') \rightarrow \text{PROB}(s') = r$.

The expression $\ell(a, s)$ ("ℓ" for "likelihood") is the probability of primitive action a being selected for execution among all its OI alternatives when an agent-performable action encapsulating these alternatives is executed. ℓ is constrained to be a probability distribution over each OI set of actions in every situation, via a set of domain-dependent *action-likelihood* axioms of the form $\ell(a, s) = z \equiv \phi(a, s)$. For example, in the robot domain the action-likelihood axiom for $\mathsf{advance}$ might be $\ell(\mathsf{advance}(x, y), s) = Normal((y - x)/\sigma)$, where *Normal* is a discrete appoximation of a Gaussian probability distribution. The successor-state axiom for p is defined in terms of POSS, OI and ℓ as follows:

$$\begin{aligned} &\text{POSS}(a, s) \rightarrow p(s'^{+}, (a, s)) = \\ &\qquad\qquad \textbf{if } \exists a', s'.(s'^{+} = (a', s') \ \wedge\ \text{POSS}(a', s') \ \wedge\ \text{OI}(a, a', s)) \\ &\qquad\qquad\quad \textbf{then } \ p(s', s) \times \ell(a', s) \\ &\qquad\qquad\quad \textbf{else } \ 0 \end{aligned}$$

This axiom combined with the constraints put on ℓ ensures that, for any situation S_α rooted at S_0 and such that $\text{PROB}(S_\alpha, S_0) > 0$, $\text{BEL}(\cdot\,, S_\alpha)$ is a probability function (in the sense of Definition 2.1) with respect to the set of situations K-related to S_α (regarded as interpretations of the fluents in the domain).

To summarise, a BHL *extended action theory* extends a standard BAT with: (i) axioms for describing the initial situation in terms of the fluents K and p, (ii) successor state axioms for K and p, (iii) one observation-indistinguishability axiom for each action, (iv) one action-likelihood axiom for each action, (v) axioms to ensure that $\textsc{Bel}$ is a probability function, and (vi) abbreviations for $\textsc{Know}$, $\textsc{Bel}$ and Do and for agent-performable complex actions.

4.2.1. An Example BHL Domain and its EPEC Counterpart

Before demonstrating a formal correspondence between BHL and EPEC, we give a preliminary intuition about the relationship between BHL and EPEC representations by formulating an example domain in both frameworks. This is a finite-domain version of the one-dimensional robot example in [3], elements of which have already been discussed above. In what follows, we assume that the arguments of advance are of sort $\{-20, \ldots, -1, 1, \ldots, 20\}$ and that the function *position* and the arguments of sense-position are of sort $\{0, \ldots, 20\}$, i.e. that the robot can move backwards and forwards among positions 0 to 20. We make appropriate adjustments to the precondition and action-likelihood axioms accordingly. The domain-dependent part of the BHL theory is:

Initial beliefs:

$\textsc{Bel}(\mathit{position}(s_{\mathit{know}}) = 10, S_0) = 4/13$
$\textsc{Bel}(\mathit{position}(s_{\mathit{know}}) = 11, S_0) = 8/13$
$\textsc{Bel}(\mathit{position}(s_{\mathit{know}}) = 12, S_0) = 1/13$

Precondition axioms:

$\textsc{Poss}(\mathsf{advance}(x, y), s) \equiv (0 \leq \mathit{position}(s) + y \leq 20)$
$\textsc{Poss}(\mathsf{sense\text{-}position}(x, y), s) \equiv y = \mathit{position}(s)$

Successor-state axiom:

$\textsc{Poss}(a, s) \rightarrow \mathit{position}(do(a, s)) =$
$\quad\quad \textbf{if } \exists x, y.a = \mathsf{advance}(x, y) \textbf{ then } \mathit{position}(s) + y \textbf{ else } \mathit{position}(s)$

Observation-indistinguishability axioms:

$\textsc{Oi}(\mathsf{advance}(x, y), a', s) \equiv \exists y'.a' = \mathsf{advance}(x, y')$
$\textsc{Oi}(\mathsf{sense\text{-}position}(x, y), a', s) \equiv \exists y'.a' = \mathsf{sense\text{-}position}(x, y')$

Action-likelihood axioms:

$\ell(\mathsf{advance}(x,y),s) =$
  if $(0 < x{+}\mathit{position}(s) < 20)$ **then**
    $\{$ **if** $y{=}x$ **then** 0.5 **else if** $|y-x|{=}1$ **then** 0.25 **else** 0 $\}$
  else if $x{+}\mathit{position}(s){=}0$ **then**
    $\{$ **if** $y{=}x$ **then** 0.75 **else if** $y{=}x{+}1$ **then** 0.25 **else** 0 $\}$
  else if $x{+}\mathit{position}(s){=}20$ **then**
    $\{$ **if** $y{=}x$ **then** 0.75 **else if** $y{=}x{-}1$ **then** 0.25 **else** 0 $\}$
  else if $x{+}\mathit{position}(s){<}0$ **then**
    $\{$ **if** $y{+}\mathit{position}(s){=}0$ **then** 1 **else** 0 $\}$
  else if $x{+}\mathit{position}(s){>}20$ **then**
    $\{$ **if** $y{+}\mathit{position}(s){=}20$ **then** 1 **else** 0 $\}$

$\ell(\mathsf{sense\text{-}position}(x,y),s) =$
  if $(0 < x < 20)$ **then**
    $\{$ **if** $y{=}x$ **then** 0.5 **else if** $|y-x|{=}1$ **then** 0.25 **else** 0 $\}$
  else if $x{=}0$ **then**
    $\{$ **if** $y{=}0$ **then** 0.75 **else if** $y{=}1$ **then** 0.25 **else** 0 $\}$
  else if $x{=}20$ **then**
    $\{$ **if** $y{=}20$ **then** 0.75 **else if** $y{=}19$ **then** 0.25 **else** 0 $\}$

Agent-performable complex actions:

$\mathsf{noisy\text{-}advance}(x) \stackrel{\text{def}}{=} \pi y.\mathsf{advance}(x,y)$
$\mathsf{noisy\text{-}sense\text{-}position} \stackrel{\text{def}}{=} \pi x,y.\mathsf{sense\text{-}position}(x,y)$

We will call the BHL extended action theory that includes the axioms above $\mathcal{T}_R$. To translate $\mathcal{T}_R$ into an EPEC domain description $\mathcal{D}_R$, we take the set $\mathcal{A}_a$ of agent actions to be $\{\mathit{NoisySensePosition}, \mathit{NoisyAdvance}(-20), \ldots, \mathit{NoisyAdvance}(-1), \mathit{NoisyAdvance}(1), \ldots, \mathit{NoisyAdvance}(20)\}$ corresponding to all ground instances of agent-performable actions in $\mathcal{T}_R$. As with PAL, and since EPEC allows for concurrent action execution, we need notation to express that only a single action is executed in a particular context. Once again we use $\mathit{only}(a)$ as an abbreviation for the formula $a \wedge \bigwedge_{a' \in \mathcal{A}_a, a' \neq a} \neg a'$. For brevity the c-propostions specified below are expressed in terms of schemas, taking x and z as meta-variables appropriately instantiated in individual propositions, and regarding arithmetic expressions as evaluated to specific integers. $\mathcal{D}_R$ is:

v-proposition:

$\mathit{Position}$ **takes-values** $\langle 0, \ldots, 20 \rangle$

i-proposition:

initially-one-of

$\{(\{Position\!=\!10\}, 4/13), (\{Position\!=\!11\}, 8/13), (\{Position\!=\!12\}, 1/13)\}$

c-proposition schemas:

for $0 < x + z < 20$:

$only(NoisyAdvance(x)) \wedge Position\!=\!z$ **causes-one-of** $\{(\{Position\!=\!z{+}x\}, 0.5), (\{Position\!=\!z{+}x{-}1\}, 0.25), (\{Position\!=\!z{+}x{+}1\}, 0.25)\}$

for $x + z = 0$:

$only(NoisyAdvance(x)) \wedge Position\!=\!z$ **causes-one-of** $\{(\{Position\!=\!0\}, 0.75), (\{Position\!=\!1\}, 0.25)\}$

for $x + z = 20$:

$only(NoisyAdvance(x)) \wedge Position\!=\!z$ **causes-one-of** $\{(\{Position\!=\!20\}, 0.75), (\{Position\!=\!19\}, 0.25)\}$

for $x + z < 0$:

$only(NoisyAdvance(x)) \wedge Position\!=\!z$ **causes-one-of** $\{(\{Position\!=\!0\}, 1)\}$

for $x + z > 20$:

$only(NoisyAdvance(x)) \wedge Position\!=\!z$ **causes-one-of** $\{(\{Position\!=\!20\}, 1)\}$

s-proposition:

$only(NoisySensePosition)$ **senses** $Position$ **with-accuracies** M

where matrix M represents the distribution given by $\ell(\textsf{sense-position}(x, y), s)$, i.e.:

$$M = \begin{pmatrix} 0.75 & 0.25 & 0 & \dots & 0 & 0 & 0 \\ 0.25 & 0.5 & 0.25 & 0 & \dots & 0 & 0 \\ 0 & 0.25 & 0.5 & 0.25 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & 0 & 0.25 & 0.5 & 0.25 & 0 \\ 0 & \dots & \dots & 0 & 0.25 & 0.5 & 0.25 \\ 0 & \dots & \dots & \dots & 0 & 0.25 & 0.75 \end{pmatrix}$$

To describe the relationship between BEL formulas entailed by $\mathcal{T}_R$ and b-propostions entailed by $\mathcal{D}_R$, it is convenient to define a new BHL complex action abbreviation $\textsf{noisy-sense-position}(x) \stackrel{\text{def}}{=} \pi y.\textsf{sense-position}(x, y)$. Although x-instantiations of $\textsf{noisy-sense-position}(x)$ are not agent-performable (since the agent cannot choose the values that its sensors detect), the agent-performable noisy-sense-position action can be "unfolded" into the set of all x-instantiations of noisy-sense-position(x) such that $\textsc{Prob}(\textsf{noisy-sense-position}(x)) > 0$, each of which corresponds to a sensing component of a particular EPEC sensing and acting history. To illustrate, we consider the agent-performable action sequence [noisy-advance(1); noisy-advance(-1); noisy-sense-position], which for convenience we further abbreviate to $\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}$. The abbreviation $\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}(x)$ is similar but with noisy-sense-position action replaced by noisy-sense-position(x).

We model $\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}$ in EPEC by augmenting $\mathcal{D}_R$ with a narrative embedded in a four-instant timeline $\{0, 1, 2, 3\}$ as follows:

$NoisyAdvance(1)$ **performed-at** 0

$NoisyAdvance(-1)$ **performed-at** 1

$NoisySensePosition$ **performed-at** 2

Let us suppose that we are interested in discovering the extent to which the agent believes it is at position 11 after executing $\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}$. Unfolding this action sequence into all possible x-instantiations of $\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}(x)$ such that $\textsc{Prob}(\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}(x), S_0) > 0$ gives us the following 9 BHL entailments:

$\mathcal{T}_R \models Do(\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}(7), S_0, s) \rightarrow \textsc{Bel}(position(s_{know}) = 11, s) = 0$

$\mathcal{T}_R \models Do(\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}(8), S_0, s) \rightarrow \textsc{Bel}(position(s_{know}) = 11, s) = 0$

$\mathcal{T}_R \models Do(\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}(9), S_0, s) \rightarrow \textsc{Bel}(position(s_{know}) = 11, s) = 0$

$\mathcal{T}_R \models Do(\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}(10), S_0, s) \rightarrow \textsc{Bel}(position(s_{know}) = 11, s) = 34/103$

$\mathcal{T}_R \models Do(\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}(11), S_0, s) \rightarrow \textsc{Bel}(position(s_{know}) = 11, s) = 136/235$

$\mathcal{T}_R \models Do(\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}(12), S_0, s) \rightarrow \textsc{Bel}(position(s_{know}) = 11, s) = 17/41$

$\mathcal{T}_R \models Do(\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}(13), S_0, s) \rightarrow \textsc{Bel}(position(s_{know}) = 11, s) = 0$

$\mathcal{T}_R \models Do(\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}(14), S_0, s) \rightarrow \textsc{Bel}(position(s_{know}) = 11, s) = 0$

$\mathcal{T}_R \models Do(\textsf{na}(1)\textsf{na}(-1)\textsf{nsp}(15), S_0, s) \rightarrow \textsc{Bel}(position(s_{know}) = 11, s) = 0$

That is to say, for example, that after attempting to advance one unit forwards followed by one unit backwards and then sensing its position to be 12 the robot will believe that the probability that it is in position 11 to be 17/41. The collection of BHL entailments above corresponds to the single EPEC entailment:

$$
\begin{aligned}
&\mathcal{D}_R \mathrel{|}\!\models \textbf{at } 3 \textbf{ believes } [\mathit{Position}\!=\!11]@3 \textbf{ with-probs}\\
&\qquad\{(\langle\{\mathit{NoisyAdvance}(1)\}@0, \{\mathit{NoisyAdvance}(-1)\}@1,\\
&\qquad\quad \{((\mathit{NoisySensePosition}, \mathit{Position}), 7)\}@2\rangle, 1/208, 0),\\
&\qquad(\langle\{\mathit{NoisyAdvance}(1)\}@0, \{\mathit{NoisyAdvance}(-1)\}@1,\\
&\qquad\quad \{((\mathit{NoisySensePosition}, \mathit{Position}), 8)\}@2\rangle, 1/26, 0),\\
&\qquad(\langle\{\mathit{NoisyAdvance}(1)\}@0, \{\mathit{NoisyAdvance}(-1)\}@1,\\
&\qquad\quad \{((\mathit{NoisySensePosition}, \mathit{Position}), 9)\}@2\rangle, 109/832, 0),\\
&\qquad(\langle\{\mathit{NoisyAdvance}(1)\}@0, \{\mathit{NoisyAdvance}(-1)\}@1,\\
&\qquad\quad \{((\mathit{NoisySensePosition}, \mathit{Position}), 10)\}@2\rangle, 103/416, 34/103),\\
&\qquad(\langle\{\mathit{NoisyAdvance}(1)\}@0, \{\mathit{NoisyAdvance}(-1)\}@1,\\
&\qquad\quad \{((\mathit{NoisySensePosition}, \mathit{Position}), 11)\}@2\rangle, 235/832, 136/235),\\
&\qquad(\langle\{\mathit{NoisyAdvance}(1)\}@0, \{\mathit{NoisyAdvance}(-1)\}@1,\\
&\qquad\quad \{((\mathit{NoisySensePosition}, \mathit{Position}), 12)\}@2\rangle, 41/208, 17/41),\\
&\qquad(\langle\{\mathit{NoisyAdvance}(1)\}@0, \{\mathit{NoisyAdvance}(-1)\}@1,\\
&\qquad\quad \{((\mathit{NoisySensePosition}, \mathit{Position}), 13)\}@2\rangle, 67/832, 0),\\
&\qquad(\langle\{\mathit{NoisyAdvance}(1)\}@0, \{\mathit{NoisyAdvance}(-1)\}@1,\\
&\qquad\quad \{((\mathit{NoisySensePosition}, \mathit{Position}), 14)\}@2\rangle, 7/416, 0),\\
&\qquad(\langle\{\mathit{NoisyAdvance}(1)\}@0, \{\mathit{NoisyAdvance}(-1)\}@1,\\
&\qquad\quad \{((\mathit{NoisySensePosition}, \mathit{Position}), 15)\}@2\rangle, 1/832, 0)\}
\end{aligned}
$$

Note that, unlike EPEC's b-proposition, the BHL entailments provide only the probabilistic belief in a formula given a particular sequence of sensed values, and not the associated probability of experiencing those particular sensed values. However, using our additional "Prob" notation we can also generate entailments such as $\mathcal{T}_R \models \textsc{Prob}(\mathsf{na}(1)\mathsf{na}(-1)\mathsf{nsp}(12), S_0) = 41/208$ to confirm that these associated probabilities correspond in both BHL and EPEC.

4.2.2. A General Correspondence Between BHL and EPEC

The translation of the one-dimentional robot domain from $\mathcal{T}_R$ to $\mathcal{D}_R$ generalises to a wide class of finite BHL representations, as we show next.

To guarantee translatability, we impose the following restricitions on a BHL extended action theory $\mathcal{T}$. We assume the $\mathcal{T}$ is written in a sorted predicate calculus where all sorts except the sort of situations are finite, and in particular that there are a finite number of actions and fluents, with all fluents being functional and taking values over finite sets. We assume $\mathcal{T}$ distinguishes between physical actions and sensing actions, with dedicated sorts for each. Furthermore, for each agent-performable action (e.g. noisy-advance and

$\mathsf{noisy\text{-}sense\text{-}position}$) we assume there exist auxiliary actions (e.g. $\mathsf{advance}$ and $\mathsf{sense\text{-}position}$) whose purpose is to specify the error profile of the corresponding agent-performable action. In more detail, we assume that physical actions have the form $\mathsf{noisy\text{-}a}(x) \stackrel{\text{def}}{=} \pi x.\mathsf{a}(x, y)$ where x is the nominal (intended) value and y is the actual value. For these actions we assume an observation-indistinguishability axiom of the form $\textsc{Oi}(\mathsf{a}(x, y), a', s) \equiv \exists y'.a' = \mathsf{a}(x, y')$. Similarly, we assume that sensing actions have the form $\mathsf{noisy\text{-}sense\text{-}f} \stackrel{\text{def}}{=} \pi x, y.\mathsf{sense\text{-}f}(x, y)$ where where x is the nominal value read on the sensor and y is the actual value, and that there is an associated observation-indistinguishability axiom of the form $\textsc{Oi}(\mathsf{sense\text{-}f}(x, y), a', s) \equiv \exists y'.a' = \mathsf{sense\text{-}f}(x, y')$. The agent-performable action $\mathsf{noisy\text{-}sense\text{-}f}$ implicitly represents the sensing of fluent f. We assume that the action-likelihood axioms in $\mathcal{T}$ have been written in such away that non-executable actions are given zero likelihood, i.e. that $\mathcal{T} \models \neg\textsc{Poss}(a, s) \rightarrow \ell(a, s) = 0$. Finally, we assume that the initial situation is fully specified in terms of $\textsc{Bel}$ in the sense that $\mathcal{T}$ includes a collection of axioms of the form $\textsc{Bel}(\phi_i(s_{know}), S_0) = p_i$ for $1 \leq i \leq n$, such that $\sum_{i=1}^{n} p_i = 1$ and each ϕ_i is a maximally consistent conjunction of fluent literals (corresponding to an EPEC state). We will refer to an extended action theory with these characteristics as a *normalised* theory.

The translation of $\mathcal{T}$ into EPEC is with respect to a given arbitrary sequence of agent-performable actions α, and is written $\mathcal{D}(\mathcal{T}, \alpha)$. For each functional fluent f of sort $\{v_1, \ldots, v_n\}$, $\mathcal{D}(\mathcal{T}, \alpha)$ includes the v-proposition

$$F \textbf{ takes-values } \langle v_1, \ldots, v_n \rangle$$

and $\mathcal{D}(\mathcal{T}, \alpha)$ includes the i-proposition

$$\textbf{initially-one-of } \{(\phi_1, p_1), \ldots, (\phi_n, p_n)\}$$

corresponding to the collection of $\textsc{Bel}$ axioms regarding S_0 mentioned above.

To specify the c-propositions in $\mathcal{D}(\mathcal{T}, \alpha)$ we use the effect axioms relating to physical actions from which the successor-state axioms of $\mathcal{T}$ have been generated. For a given agent-performable action $\mathsf{noisy\text{-}a}(x)$ and associated primitive actions $\mathsf{a}(x, y)$ we assume these are of the general form:

$$\mathrm{Poss}(\mathsf{a}(x,y),s) \wedge \psi_1(s) \rightarrow F_1(do(\mathsf{a}(x,y),s)) = v_1(y)$$
$$\vdots$$
$$\mathrm{Poss}(\mathsf{a}(x,y),s) \wedge \psi_m(s) \rightarrow F_m(do(\mathsf{a}(x,y),s)) = v_m(y)$$

and that the action-likelihood axiom for $\mathsf{a}(x,y)$ is of the general form

$$\begin{aligned}\ell(\mathsf{a}(x,y),s) = \ & \textbf{if } \phi_1(x,s) \textbf{ then } p_1(x,y)\\ & \textbf{else if } \phi_2(x,s) \textbf{ then } p_2(x,y)\\ & \quad\vdots\\ & \textbf{else if } \phi_n(x,s) \textbf{ then } p_n(x,y)\\ & \textbf{else } p_{n+1}(x,y)\end{aligned}$$

Without loss of generality, we further assume that $\{\psi_1(s),\ldots,\psi_m(s)\}$ can be partitioned as $\{\psi_1(s),\ldots,\psi_{m_1}(s)\} \cup \ldots \cup \{\psi_{m_n+1}(s),\ldots,\psi_m(s)\}$, in such a way that $\{\psi_1(s),\ldots,\psi_{m_1}(s)\}$ is maximally consistent with $\phi_1(x,s)$, $\{\psi_{m_1+1}(s),\ldots,\psi_{m_2}(s)\}$ is maximally consistent with $\neg\phi_1(x,s) \wedge \phi_2(x,s)$, and so on, ending with $\{\psi_{m_n+1}(s),\ldots,\psi_m(s)\}$ being maximally consistent with $\neg\phi_1(x,s) \wedge \cdots \wedge \neg\phi_n(x,s)$. This ensures that the preconditions of the c-propositions in the translation are mutually exclusive. Then the following c-proposition schemas are included in $\mathcal{D}(\mathcal{T},\alpha)$, where groundings of x range over all values in its sort:

$$\mathit{only}(\mathit{NoisyA}(x)) \wedge \phi_1(x) \wedge \psi_1 \wedge \cdots \wedge \psi_{m_1} \textbf{ causes-one-of}$$
$$\{(\{F_1 = v_1(x,g),\ldots,F_{m_1} = v_{m_1}(x,g)\}, p_1(x,g)) \mid g \text{ is a grounding of } y \text{ in } \mathsf{a}(x,y) \text{ s.t. } p_1(x,g) > 0\}$$

$$\mathit{only}(\mathit{NoisyA}(x)) \wedge \neg\phi_1(x) \wedge \phi_2(x) \wedge \psi_{m_1+1} \wedge \cdots \wedge \psi_{m_2} \textbf{ causes-one-of}$$
$$\{(\{F_{m_1+1} = v_{m_1+1}(x,g),\ldots,F_{m_2} = v_{m_2}(x,g)\}, p_2(x,g)) \mid g \text{ is a grounding of } y \text{ in } \mathsf{a}(x,y) \text{ s.t. } p_2(x,g) > 0\}$$

$$\vdots$$

$$\mathit{only}(\mathit{NoisyA}(x)) \wedge \neg\phi_1(x) \wedge \cdots \wedge \neg\phi_{n-1}(x) \wedge \phi_n(x) \wedge \psi_{m_{n-1}+1} \wedge \cdots \wedge \psi_{m_n} \textbf{ causes-one-of}$$
$$\{(\{F_{m_{n-1}+1} = v_{m_{n-1}+1}(x,g),\ldots,F_{m_n} = v_{m_n}(x,g)\}, p_n(x,g)) \mid g \text{ is a grounding of } y \text{ in } \mathsf{a}(x,y) \text{ s.t. } p_n(x,g) > 0\}$$

$$\mathit{only}(\mathit{NoisyA}(x)) \wedge \neg\phi_1(x) \wedge \cdots \wedge \neg\phi_n(x) \wedge \psi_{m_n+1} \wedge \cdots \wedge \psi_m \textbf{ causes-one-of}$$
$$\{(\{F_{m_n+1} = v_{m_n+1}(x,g),\ldots,F_m = v_m(x,g)\}, p_{n+1}(x,g)) \mid g \text{ is a grounding of } y \text{ in } \mathsf{a}(x,y) \text{ s.t. } p_{n+1}(x,g) > 0\}$$

To generate s-propositions for the sensing action noisy-sense-f, we assume an action-likelihood axiom for sense-f of the following general form:

$$\ell(\textsf{sense-f}(x,y),s) = \begin{array}{l} \textbf{if } y \neq f(s) \textbf{ then } 0 \\ \textbf{else if } \phi_1(s) \textbf{ then } p_1(x,y) \\ \quad\quad\vdots \\ \textbf{else if } \phi_n(s) \textbf{ then } p_n(x,y) \\ \textbf{else } p_{n+1}(\textsf{sense-f}(x,y)) \end{array}$$

The following s-propositions are included in $\mathcal{D}(\mathcal{T},\alpha)$:

$only(NoisySenseF) \wedge \phi_1$ **senses** F **with-accuracies** M_1

$only(NoisySenseF) \wedge \neg\phi_1 \wedge \phi_2$ **senses** F **with-accuracies** M_2

$\vdots$

$only(NoisySenseF) \wedge \neg\phi_1 \wedge \cdots \wedge \neg\phi_{n-1} \wedge \phi_n$ **senses** F **with-accuracies** M_n

$only(NoisySenseF) \wedge \neg\phi_1 \wedge \cdots \wedge \neg\phi_n$ **senses** F **with-accuracies** M_{n+1}

where M_i is defined for each $i \in \{1,\ldots,n\}$ as:

$$M_i = \begin{pmatrix} p_i(v_1,v_1) & \ldots & p_i(v_m,v_1) \\ \vdots & \vdots & \vdots \\ p_i(v_1,v_m) & \ldots & p_i(v_m,v_m) \end{pmatrix}$$

and $v_1,\ldots,v_m$ are the values fluent F can take, ordered as in the corresponding v-proposition.

The agent-performable action sequence α translates straightforwardly to an EPEC narrative. If $\alpha = [\textsf{noisy-act}_1;\ldots;\textsf{noisy-act}_k]$ then we assume a timeline of instants $\{0,1,\ldots,k\}$ and include the following p-propositions in $\mathcal{D}(\mathcal{T},\alpha)$:

$NoisyAct_1$ **performed-at** 0

$\vdots$

$NoisyAct_k$ **performed-at** $k-1$

Proposition 4.2 below shows how entailments from $\mathcal{T}$ and $\mathcal{D}(\mathcal{T},\alpha)$ correspond. In the statement of the theorem, the notation $T([H])$, where $[H]$ is a history component within a b-proposition, signifies the BHL action sequence extracted from $[H]$ in which nominal (sensed) values of sensing actions have been instantiated to match the corresponding sensed values in $[H]$. (For example, if $[H]' = \langle\{NoisyAdvance(4)\}@0,\ \{((NoisySensePosition, Position),6)\}@1\rangle$ then

$T([H]')$ is [noisy-advance(4); noisy-sense-position(6)]).

Proposition 4.2 (BHL/EPEC Correspondence)**.** Let $\mathcal{T}$ be a normalised extended action theory and let α be an executable sequence of n agent-performable actions. Let $\Delta = \{\alpha_1, \ldots, \alpha_m\}$ be the maximal set of distinct executable action sequences such that for each $\alpha_i \in \Delta$, $\textsc{Prob}(\alpha_i, S_0) > 0$ and α_i has been obtained from α by replacing each noisy sensing action, say $\pi x, y.\mathsf{sense\text{-}f}(x, y)$, with $\pi y.\mathsf{sense\text{-}f}(v, y)$ for some sensed value v for the nominal variable x. Then

$$\text{for all } 1 \leq i \leq m,\ \mathcal{T} \models Do(\alpha_i, S_0, s) \rightarrow \textsc{Bel}(\phi[s_{know}], s) = P_i$$
$$\text{if and only if}$$
$$\mathcal{D}(\mathcal{T}, \alpha) \mid\models \textbf{at } n \textbf{ believes } [\phi]@n \textbf{ with-probs } P$$

where $P = \{([H_1], B_1, P_1), \ldots, ([H_m], B_m, P_m)\}$ is such that for each $1 \leq i \leq m$ $T([H_i]) = \alpha_i$ and $B_i = \textsc{Prob}(\alpha_i, S_0)$.

Proof: See Appendix D.6. ■

4.2.3. Representing the Epectitis Example in BHL

We conclude our discussion of BHL by examining the extent to which Example 1.1 (epectisis) can be represented as a BHL extended action theory, in order to give some additional insight into the relationship between BHL and EPEC. As is the case for PAL, not all of the EPEC propositions (EP1)-(EP11) for this domain (listed in Appendix C) are translatable into BHL. BHL does not include mechanisms to model arrative information, i.e. information about what actions actually occur at different timepoints, such as that contained in (EP5), (EP10) and (EP11). Moreover, the occurrence of *SkinContact* is probabilistic in this narrative (with a probability of 0.95 that it occurred), and the occurrence of *TakeEpecillin* is probabilistic-belief-conditioned. The capacity of an SC-based framework to model these additional types of action characteristics has, to our knowledge, not yet been explored.

We therefore focus on representing the effects (modelled in (EP6)–(EP8)) of actions *SkinContact* and *TakeEpecillin*, and of sensing action *DoBloodTest* (described in (EP9)). Both of these actions are non-deterministic, and so in BHL need to be represented by a collection of obesrvationally-indistinguishable, deterministic primitive actions, one for each possible non-deterministic outcome.

The action *TakeEpecillin* may or may not cure epectisis and may or not cause side effects, so four BHL primitive action versions are needed to cover the possible combinations of these eventualities. We therefore define the agent-performable action take-epecillin as:

$$\mathsf{take\text{-}epecillin}[\,] \mid \mathsf{take\text{-}epecillin}[\neg e] \mid \mathsf{take\text{-}epecillin}[s] \mid \mathsf{take\text{-}epecillin}[\neg e,s]$$

where the symbols $\neg e$ and s in [ ]'s indicate the eventualities of curing epectisis and causing side effects respectively. Similarly, we have:

$$\mathsf{skin\text{-}contact} \stackrel{\text{def}}{=} \mathsf{skin\text{-}contact}[\,] \mid \mathsf{skin\text{-}contact}[e]$$

For this domain we can assume all primitive actions are always executable, i.e. $\forall a, s.\textsc{Poss}(a, s)$. So the successor-state axioms for the fluents *docHasEpectisis* and *docHasSideEffects* are:

$$\begin{aligned}&\mathit{docHasEpectisis}(\mathit{do}(a, s)) \equiv (a = \mathsf{skin\text{-}contact}[e] \vee \\ &\qquad (\mathit{docHasEpectisis}(s) \wedge a \neq \mathsf{take\text{-}epecillin}[\neg e] \wedge a \neq \mathsf{take\text{-}epecillin}[\neg e,s]))\end{aligned}$$

$$\begin{aligned}&\mathit{docHasSideEffects}(\mathit{do}(a, s)) \equiv \\ &\qquad (\mathit{docHasSideEffects}(s) \vee a = \mathsf{take\text{-}epecillin}[s] \vee a = \mathsf{take\text{-}epecillin}[\neg e,s])\end{aligned}$$

Observation-indistinguishability (OI) axioms are need for each primitive action, all of the same general form. For example, the OI axiom for take-epecillin[ ] is:

$$\begin{aligned}&\textsc{Oi}(\mathsf{take\text{-}epecillin}[\,], a', s) \equiv (a' = \mathsf{take\text{-}epecillin}[\,] \vee \\ &\qquad a' = \mathsf{take\text{-}epecillin}[\neg e] \vee a' = \mathsf{take\text{-}epecillin}[s] \vee a' = \mathsf{take\text{-}epecillin}[\neg e,s])\end{aligned}$$

The probabilistic information contained in (EP6)–(EP8) is encapsulated in six action-likelihood axioms:

$$\ell(\mathsf{skin\text{-}contact}[e], s) = \textbf{ if } \mathit{patientHasEpectitis}(s) \textbf{ then } 0.75 \textbf{ else } 0$$

$$\ell(\mathsf{skin\text{-}contact}[\,], s) = \textbf{ if } \mathit{patientHasEpectitis}(s) \textbf{ then } 0.25 \textbf{ else } 1$$

$$\ell(\mathsf{take\text{-}epecillin}[\neg e,s], s) = \textbf{ if } \mathit{docHasEpectitis}(s) \textbf{ then } 0.1485 \textbf{ else } 0$$

$$\ell(\mathsf{take\text{-}epecillin}[\neg e], s) = \textbf{ if } \mathit{docHasEpectitis}(s) \textbf{ then } 0.8415 \textbf{ else } 0$$

$$\ell(\mathsf{take\text{-}epecillin}[s], s) = \textbf{ if } \mathit{docHasEpectitis}(s) \textbf{ then } 0.0015 \textbf{ else } 0.15$$

$$\ell(\mathsf{take\text{-}epecillin}[\,], s) = \textbf{ if } \mathit{docHasEpectitis}(s) \textbf{ then } 0.0085 \textbf{ else } 0.85$$

To represent (EP9) in BHL, we need two versions of *DoBloodTest*, so we define $\mathsf{do\text{-}blood\text{-}test} \stackrel{\text{def}}{=} \mathsf{do\text{-}blood\text{-}test}[-] \mid \mathsf{do\text{-}blood\text{-}test}[+]$ where $\mathsf{do\text{-}blood\text{-}test}$ is agent-performable, and the primitive actions on the righthand side of the definition represent the blood test actions that return negative and positive results respectively. Their observation-indistinguishability axioms are:

$$\textsc{Oi}(\mathsf{do\text{-}blood\text{-}test}[-], a', s) \equiv a' = \mathsf{do\text{-}blood\text{-}test}[-]$$
$$\textsc{Oi}(\mathsf{do\text{-}blood\text{-}test}[+], a', s) \equiv a' = \mathsf{do\text{-}blood\text{-}test}[+]$$

and the likelihood axioms for each are:

$$\ell(\mathsf{do\text{-}blood\text{-}test}[-], s) = \textbf{ if } \neg docHasEpectitis(s) \textbf{ then } 0.9 \textbf{ else } 0.05$$
$$\ell(\mathsf{do\text{-}blood\text{-}test}[+], s) = \textbf{ if } \neg docHasEpectitis(s) \textbf{ then } 0.1 \textbf{ else } 0.95$$

Finally, we can translate the initial conditions in (EP4) to:

$$\textsc{Bel}(patientHasEpectitis(s_{know}) \wedge \neg docHasEpectitis(s_{know}) \\ \wedge \neg docHasSideEffects(s_{know}), S_0) = 0.8$$

$$\textsc{Bel}(\neg patientHasEpectitis(s_{know}) \wedge \neg docHasEpectitis(s_{know}) \\ \wedge \neg docHasSideEffects(s_{know}), S_0) = 0.2$$

It is straightforward to confirm that for this domain EPEC and BHL entailments also coincide, for those parts of the narrative that can be approximated with a BHL Do term. For example, if we label the above BHL encoding as $\mathcal{T}_e$ then we have:

$$\mathcal{T}_e \models Do([\mathsf{skin\text{-}contact}; \mathsf{do\text{-}blood\text{-}test}[-]], S_0, s) \rightarrow \\ \textsc{Bel}(docHasEpectisis(s_{know}), s) = 1/13$$

$$\mathcal{T}_e \models \textsc{Prob}([\mathsf{skin\text{-}contact}; \mathsf{do\text{-}blood\text{-}test}[-]], S_0) = 39/100$$

$$\mathcal{T}_e \models Do([\mathsf{skin\text{-}contact}; \mathsf{do\text{-}blood\text{-}test}[+]], S_0, s) \rightarrow \\ \textsc{Bel}(docHasEpectisis(s_{know}), s) = 57/61$$

$$\mathcal{T}_e \models \textsc{Prob}([\mathsf{skin\text{-}contact}; \mathsf{do\text{-}blood\text{-}test}[+]], S_0) = 61/100$$

These entailments coincide with:

$$(\mathcal{D}_e \mid [SkinContact = true]@{-1}) \;\; |\!\models \\ \textbf{at } 2 \textbf{ believes } [DocHasEpectisis = true]@2 \textbf{ with-probs} \\ \{(\langle\{((DoBloodTest, DocHasEpectisis), false)\}@1\rangle, 39/100, 1/13), \\ (\langle\{((DoBloodTest, DocHasEpectisis), true)\}@1\rangle, 61/100, 57/61)\}$$

	PEC [8]	EPEC	PAL [6]	PSC [35]	BHL [3]	Modular-$\mathcal{E}$ [17]	EFEC [5]	Prob-EC [7]	MLN-EC [36]	Language $\mathcal{E}$+ [37]
Classical logic formalism				✓	✓		✓			
Action language	✓	✓	✓			✓				✓
Supports narratives	✓	✓	✓			✓	✓	✓	✓	
Functional fluents	✓	✓		✓	✓		✓			
Elaboration tolerant			✓			✓*	✓	✓	✓	
Supports ramifications			✓			✓				✓
Supports qualifications						✓				✓
Probabilistic reasoning	✓	✓	✓	✓	✓			✓	✓	✓
Epistemic reasoning		✓		✓	✓		✓			✓
Concurrent actions	✓	✓				✓	✓	✓	✓	
Triggered actions	✓	✓					✓			
Imperfect sensing		✓		✓	✓					
Reasoning about knowledge of past, present and future		✓					✓			
Belief conditioned actions		✓✓					✓			✓
Supports continuous probability distributions				✓	✓✓✓					
Supports only knowing					✓					
Mixes probabilities and non-determinism				✓	✓					✓

Table 1: A 'feature chart' of frameworks for Reasoning About Actions and some of the domain features they support. The double checkmark (✓✓) indicates that EPEC can model actions whose epistemic precondition incorporates strengths of belief. The triple checkmark (✓✓✓) indicates that BHL can deal seamlessly with discrete and continuous probability distributions. The checkmark with star (✓*) indicates that Modular-E is both 'modular' and elaboration tolerant and argues the inseparability of these properties.

4.3. Other Related Work and a Feature Chart

To conclude this section, we provide a brief survey of other languages for Reasoning About Actions that also support some form of uncertain reasoning. In particular we compare these in Table 1 in terms of the domain features they support. Although our focus is on probabilistic reasoning, Table 1 also lists two non-probabilistic languages that have influenced our work, Modular-$\mathcal{E}$ [17] and EFEC [5], since they include features such as the ability to reason about knowledge of past, present and future, and support for resolving conflicts between concurrent actions.

BHL [3], already discussed in detail in Section 4.2, is a cornerstone of early work integrating probabilistic degrees of belief with logical aparatus for reasoning about actions, and has served as an inspiration for similar approaches (e.g. [35]). Although it does not support narrative reasoning and belief-conditioned actions, it has been extended and enriched with several other features not provided for by EPEC. For example, [4, 38] add support for continuous (as well as mixed) probability distributions, and [39] applies BHL and these extensions to the problem of localisation, i.e. to the case where an agent moves in a (multi-dimensional) world and can sense its position. The problem of extending BHL with a modality known as 'only knowing', which allows for a precise specification of what is and what is not known within a logical theory of actions, has also been tackled in [40]. The Probabilistic Situation Calculus [35] ('PSC' in Table 1) is similar to BHL in that it is based on Reiter's Situation Calculus, and supports continuous probability distributions and observations. Its semantics is given in terms of randomly reactive automata and implemented in Mathematica.

PAL, already discussed in detail in Section 4.1, provides an action-language representation for Markov Decision Processes. Although it is ontologically close to the Situation Calculus, it allows for a limited class of non-probabilistic narratives. Additionally, it incorporates a degree of elaboration tolerance – the end-user can define additional random variables alongside an existing theory provided these variables are probabilistically independent. PAL does not support sensing or epistemic reasoning, and so in these repects the planning-oriented action language $\mathcal{E}+$ [37] is an advance. $\mathcal{E}+$ supports both sensing and belief-conditioned actions, and the authors of [37] provide algorithms for the efficient computation of plans, which makes it an advanced tool for epistemic planning. However, sensing actions are always assumed to be perfect and so the association of a confusion matrix (see EPEC's Definition 3.12) with its semantics is not supported.

The Event Calculus has not been enhanced with any form of probabilistic reasoning until more recently. To our knowledge, the first attempts to do so are MLN-EC [36] and the closely related Prob-EC [7]. These two languages give a probabilistic semantics to the Event Calculus, respectively using Markov Logic Networks [41] and a recent probabilistic dialect of Prolog called ProbLog [42]. Both are based on a discrete-time reworking of the Event Calculus and applied to the task of event recognition, where a set of complex activities must be detected when a set of time-stamped short-term activities is received as input. They provide separate support for causal rules (MLN-EC) and probabilistic events (Prob-EC), which in EPEC are integrated together. These two frameworks have opened the way for other probabilistic extensions of the Event Calculus. For example, [43] extends Prob-EC to deal with uncertain event observations, uncertain effects, and uncertain composite event definitions.

Outside of the immediate sub-field of Reasoning about Actions, the area of Contingent Planning deals with the problem of (efficiently) generating plans in non-deterministic environments, sometimes under uncertainty [44, 45, 46]. As in the domains EPEC aims to model, this leads to the necessity of including sensing actions in plans, and generates plans that branch according to each possible outcome of these sensing actions. Although the main focus of this area of research is on procedures for planning rather than on knowledge representation, we envisage that the problem of synthesizing plans in EPEC could fruitfully exploit this body of work.

5. Implementation of EPEC

Although full technical descriptions and evaluations of implementations of EPEC are beyond the scope of this paper, in this section we briefly outline our implementation experiments to date[12]. Prior work in [8] describes a provably exact implementation of PEC (a subset of the non-epistemic fragment of EPEC), using the Answer Set Programming grounder and solver Clingo [47] to generate well-behaved worlds and determine their probability mass in order to answer temporal projection queries. However, as discussed elsewhere (see e.g. [36]), similar implementations do not scale well for a wide set of domains. For this reason, other implementations have been developed and/or are currently under development with a view to answering queries efficiently. A runtime ASP version of the implementation proposed in [8], augmented with epistemic capabilities, is currently being developed within the context of the AVATEA project [48], where it is being used for recognising and reacting to high-level activities detected from data-streams annotated with probabilities. In this runtime version, only queries about the present state of the world are allowed. To adjust for this constraint, new events are recompiled in a new i-proposition as they are received. This makes the reasoning task lighter and allows for efficient computation of exact probabilities, at the expense of the ability to reason about the past. An ASP implementation of the full EPEC is also available, which exploits Clingo's integration with Python to implement the reduct mechanism. Preliminary experiments indicate that it suffers from scalibility issues similar to those of [8], but we envisage that it might be used effectively for applications where the system is not required to produce immediate answers, e.g. for story understanding.

In order to provide an implementation for application areas that require

[12]Some of the existing implementations are publicly available: see `https://github.com/dasaro/pec` for the ASP implementation of the non-epistemic fragment of EPEC, `https://github.com/dasaro/pec-anglican` for the approximate Anglican implementation of the non-epistemic fragment of EPEC, `https://gitlab.com/fdasaro/pec-runtime` for the runtime ASP version, or `https://gitlab.com/fdasaro/epec-vanilla` for an ASP implementation of full EPEC.

scalability, we have also experimented with approximate inference methods. Our experiments have involved transforming EPEC domains into corresponding partially observable Markov decision processes (POMDPs) [49] (or MDPs in the case of non-epistemic domains) and then sampling from them. We have used Anglican [50], a probabilistic programming language, for this, since it allows one to describe a model and then use in-built sampling algorithms to perform probabilistic inference.

6. Summary, Discussion and Future Work

This paper presents and describes the action language EPEC – Epistemic Probabilistic Event Calculus – that combines epistemic, probabilistic, causal and narrative reasoning within a natural and intuitive syntax. A key feature of EPEC is that it supports the representation of, and reasoning about, uncertain (i.e. probabilistic) information concerning narratives which can contain two different types of events: environmentally triggered action occurrences, and belief-conditioned agent action executions. In particular, agent actions may simultaneously trigger changes in the environment and allow fluent values to be perfectly or imperfectly sensed by the agent. Both environmental and agent actions may occur concurrently, and the underlying time structure can be either discrete or continuous. Concurrent sensing conflicts are appropriately resolved. The semantics of EPEC takes advantage of both a modal logic possible-worlds approach to reasoning about knowledge, and a Bayesian view of probability that generalises the notion of knowledge (true belief) to that of justified belief. The resultant formalism allows us to model and reason about an important class of domains that to our knowledge are not easily or adequately representable by existing frameworks. For example, a large proportion of medical and other technical information is both probabilistic and causal in nature, and it is common to take a scenario- or narrative-based approach to reasoning about such domains. Our future work plans include experimenting with EPEC as a decision support tool in such contexts, for example by developing appropriate front-end interfaces that allow it be be used directly by patients choosing between alternative courses of treatment for medical conditions.

The syntax and associated semantics of EPEC could in future be expanded in a number of ways. In particular, the *reduct* mechanism employed within the semantics in Section 3.4 relates an agent's sensing history to the actual environmental history in a global, non-temporally directed way, and this opens the possibility of extending EPEC in order to model agent actions conditioned by beliefs about future as well as past states. For example, an expanded syntax might include p-propositions containing i-formulas rather than formulas, such as

$$A\ \textbf{performed-at}\ I\ \textbf{with-prob}\ P^{+}\ \textbf{if-believes}\ (\textbf{do-results-in}\ \varphi, \bar{P})$$

$$A\ \textbf{performed-at}\ I\ \textbf{with-prob}\ P^{+}\ \textbf{if-believes}\ (\textbf{dont-do-results-in}\ \varphi, \bar{P})$$

where the i-formula φ potentially refers to instants both before and after I.

Although the focus of this paper has been in providing a logical formalism for the declarative specification of a class of domains, we are mindful of the need to engineer implementable frameworks. [8] describes an Answer Set Programming (ASP) approach to implementing an earlier (slightly more restrictive) class of ne-domain descriptions. We have already made progress in expanding this implementation to epistemic domains (see the footnotes in Section 5 for availability of code), and aim to empirically evaluate the efficiency of this implementation using a variety of domains and automated tasks. We have also experimented with computing approximate EPEC entailments using the probabilistic programming language Anglican. Emprical comparisons of these two types of implementation will allow us to investigate the trade-off, if any, between inefficiency and approximation.

Finally, we would like to explore various modes of reasoning using EPEC as an undelying semantic foundation. As mentioned in Section 5 we have developed a preliminary translation of EPEC domains into partially observable Markov decision processes (POMDPs) [49], which opens the possibility of developing an EPEC-based approach to *epistemic planning* building upon the growing body of existing work on probabilistic planning using POMDP models, probabilistic programming and related techniques, e.g. [51, 52, 53]. Moreover, as manually encoding domains to represent real world problems raises concerns around imprecision and cost, we also aim to explore methods for approximating real world domains from observed narratives. Most straightforwardly, we envisage combining a pre-defined EPEC domain description template with associated observed narratives to perform probabilistic inference (e.g. see [30]) on free-parameters within the template, such as the probabilities of a c-proposition[13]. More ambitiously, we would like to explore how domain descriptions might be automatically or semi-automatically learned at a deeper structural level.

Acknowledgements

We would like to thank the three anonymous referees for their careful, helpful and detailed reviews of earlier drafts of this paper. Their and the editor's observations and suggestions have lead to considerable improvements in both the presentation and the technical content of our work.

References

[1] R. C. Moore, A Formal Theory of Knowledge and Action, in: J. R. Hobbs, R. C. Moore (Eds.), Formal Theories of the Commonsense World, Ablex, Norwood, NJ., 1985, pp. 319–358.

[2] R. B. Scherl, H. J. Levesque, Knowledge, Action, and the Frame Problem, Artificial Intelligence 144 (1–2) (2003) 1–39.

[3] F. Bacchus, J. Y. Halpern, H. J. Levesque, Reasoning about Noisy Sensors and Effectors in the Situation Calculus, Artificial Intelligence 111 (1-2) (1999) 171–208.

[13] A probabilistic programming language like Anglican provides a readily available set of tools for this kind of inference

[4] V. Belle, H. Levesque, Reasoning about Continuous Uncertainty in the Situation Calculus, in: Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13, AAAI Press, 2013, pp. 732–738.

[5] J. Ma, R. Miller, L. Morgenstern, T. Patkos, An Epistemic Event Calculus for ASP-based Reasoning About Knowledge of the Past, Present and Future, in: Proceedings of the 19th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, LPAR-19, 2013.

[6] C. Baral, N. Tran, L.-C. Tuan, Reasoning About Actions in a Probabilistic Setting, in: Eighteenth National Conference on Artificial Intelligence, AAAI/IAAI, 2002, pp. 507–512.

[7] A. Skarlatidis, A. Artikis, J. Filippou, G. Paliouras, A Probabilistic Logic Programming Event Calculus, Theory and Practice of Logic Programming 15 (2015) 213–245.

[8] F. A. D'Asaro, A. Bikakis, L. Dickens, R. Miller, Foundations for a Probabilistic Event Calculus, in: Logic Programming and Nonmonotonic Reasoning, LPNMR'17, 2017, pp. 57–63.

[9] R. Miller, M. Shanahan, Some Alternative Formulations of the Event Calculus, in: Computational Logic: Logic Programming And Beyond, Springer, 2002, pp. 452–490.

[10] J. McCarthy, P. Hayes, Some Philosophical Problems from the Standpoint of Artificial Intelligence, in: B. Meltzer, D. Michie (Eds.), Machine Intelligence 4, Edinburgh University Press, 1969, pp. 463–502.

[11] R. Reiter, Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems, MIT Press, 2001.

[12] M. Shanahan, Solving the Frame Problem - a Mathematical Investigation of the Common Sense Law of Inertia, MIT Press, 1997.

[13] R. Miller, M. Shanahan, Narratives in the Situation Calculus, Journal of Logic and Computation 4 (5) (1994) 513–530.

[14] J. Pinto, Temporal Reasoning in the Situation Calculus, PhD Thesis, University of Toronto, Department of Computer Science, 1994.

[15] R. Kowalski, M. Sergot, A Logic-based Calculus of Events, New Generation Computing 4 (1) (1986) 67–95.

[16] K. Eshghi, Abductive Planning with Event Calculus, in: Logic Programming, Proceedings of the Fifth International Conference and Symposium, Seattle, Washington, USA, August 15-19, 1988 (2 Volumes), 1988, pp. 562–579.

[17] A. Kakas, L. Michael, R. Miller, $\mathcal{M}$odular-$\mathcal{E}$ and the Role of Elaboration Tolerance in Solving the Qualification Problem, Artificial Intelligence 175 (1) (2011) 49–78.

[18] M. Gelfond, V. Lifschitz, Representing Actions in Extended Logic Programming, in: K. Apt (Ed.), Proceedings of the Joint International Conference and Symposium on Logic Programming (JICSLP'92), MIT Press, 1992, pp. 559–573.

[19] M. Gelfond, V. Lifschitz, Representing Action and Change by Logic Programs, Journal of Logic Programming 17 (2/3–4) (1993) 301–321.

[20] A. C. Kakas, R. Miller, A Simple Declarative Language for Describing Narratives With Actions, Journal of Logic Programming 31 (1-3) (1997) 157–200.

[21] R. B. Scherl, H. J. Levesque, The Frame Problem and Knowledge-Producing Actions, in: In Proceedings of the Eleventh National Conference on Artificial Intelligence, AAAI Press/The MIT Press, 1993, pp. 689–695.

[22] R. B. Scherl, Reasoning about the Interaction of Knowledge, Time and Concurrent Actions in the Situation Calculus, in: Proceedings of the 18th

International Joint Conference on Artificial Intelligence, IJCAI'03, Morgan Kaufmann Publishers Inc., 2003, pp. 1091–1096.

[23] S. Zimmerbaum, R. Scherl, Sensing Actions, Time, and Concurrency in the Situation Calculus, in: Proceedings of Intelligent Agents VII Agent Theories Architectures and Languages 2000, Lecture Notes in Computer Science, vol 1986, Springer, Berlin, Heidelberg, 2001, pp. 31–45.

[24] R. Reiter, On Knowledge-Based Programming with Sensing in the Situation Calculus, ACM Transactions on Computational Logic 2 (2000) 433–457.

[25] Y. Lesperance, H. J. Levesque, F. Lin, R. B. Scherl, Ability and Knowing How in the Situation Calculus, Studia Logica 66 (1995) 2000.

[26] J. Pearl, Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1988.

[27] P. Cheeseman, In Defense of Probability, in: Proceedings of the 9th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'85, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1985, pp. 1002–1009.

[28] F. V. Jensen, T. D. Nielsen, Bayesian Networks and Decision Graphs, 2nd Edition, Springer Publishing Company, Incorporated, 2007.

[29] J. Pearl, Causality: Models, Reasoning and Inference, 2nd Edition, Cambridge University Press, New York, NY, USA, 2009.

[30] C. M. Bishop, Pattern Recognition and Machine Learning, Springer-Verlag, Berlin, Heidelberg, 2006.

[31] J. Paris, The Uncertain Reasoner's Companion: a Mathematical Perspective, Vol. 39, Cambridge University Press, 2006.

[32] N. McCain, H. Turner, A Causal Theory of Ramifications and Qualifications, in: Proceedings of the Fourteenth International Joint Conference on

Artificial Intelligence, IJCAI 95, Montréal Québec, Canada, August 20-25 1995, 2 Volumes, 1995, pp. 1978–1984.

[33] N. Tran, C. Baral, Encoding Probabilistic Causal Model in Probabilistic Action Language, in: Proceedings of the 19th National Conference on Artifical Intelligence, AAAI'04, AAAI Press, 2004, pp. 305–310.

[34] H. J. Levesque, R. Reiter, Y. Lespérance, F. Lin, R. B. Scherl, GOLOG: A Logic Programming Language for Dynamic Domains, Journal of Logic Programming 31 (1-3) (1997) 59–83.

[35] P. Mateus, A. Pacheco, J. Pinto, A. Sernadas, C. Sernadas, Probabilistic Situation Calculus, Annals of Mathematics and Artificial Intelligence 32 (1/4) (2001) 393–431.

[36] A. Skarlatidis, G. Paliouras, A. Artikis, G. A. Vouros, Probabilistic Event Calculus for Event Recognition, ACM Transactions on Computational Logic (TOCL) 16 (2) (2015) 11.

[37] L. Iocchi, T. Lukasiewicz, D. Nardi, R. Rosati, Reasoning about Actions with Sensing under Qualitative and Probabilistic Uncertainty, ACM Transactions on Computational Logic (TOCL) 10 (1).

[38] V. Belle, H. J. Levesque, Reasoning about Discrete and Continuous Noisy Sensors and Effectors in Dynamical Systems, Artifial Intelligence 262 (2018) 189–221.

[39] V. Belle, H. J. Levesque, A logical theory of localization, Studia Logica 104 (4) (2016) 741–772.

[40] V. Belle, G. Lakemeyer, H. J. Levesque, A First-Order Logic of Probability and Only Knowing in Unbounded Domains, in: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, 2016, pp. 893–899.

[41] M. Richardson, P. Domingos, Markov logic networks, Machine Learning 62 (1) (2006) 107–136.

[42] L. De Raedt, A. Kimmig, H. Toivonen, ProbLog: A Probabilistic Prolog and Its Application in Link Discovery, in: Proceedings of the 20th International Joint Conference on Artifical Intelligence, Vol. 7 of IJCAI'07, 2007, pp. 2462–2467.

[43] K. McAreavey, K. Bauters, W. Liu, J. Hong, The Event Calculus in Probabilistic Logic Programming with Annotated Disjunctions, in: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS '17, 2017, pp. 105–113.

[44] S. M. Majercik, M. L. Littman, Contingent Planning Under Uncertainty via Stochastic Satisfiability, in: Proceedings of the Sixteenth National Conference on Artificial Intelligence and the Eleventh Innovative Applications of Artificial Intelligence Conference Innovative Applications of Artificial Intelligence, AAAI '99/IAAI '99, 1999, pp. 549–556.

[45] R. Brafman, G. Shani, Online belief tracking using regression for contingent planning, Artificial Intelligence 241 (2016) 131–152.

[46] A. Albore, H. Palacios, H. Geffner, A Translation-based Approach to Contingent Planning, in: Proceedings of the 21st International Jont Conference on Artifical Intelligence, IJCAI'09, 2009, pp. 1623–1628.

[47] M. Gebser, R. Kaminski, B. Kaufmann, T. Schaub, Clingo = ASP + control: Preliminary report, CoRR Archive abs/1405.3694.

[48] F. A. D'Asaro, A. Origlia, S. Rossi, Towards a Logic-Based Approach for Multi-Modal Fusion and Decision Making during Motor Rehabilitation Sessions, in: Proceedings of the 20th Workshop "From Objects to Agents" (WOA), 2019.

[49] L. P. Kaelbling, M. L. Littman, A. R. Cassandra, Planning and Acting in Partially Observable Stochastic Domains, Artificial Intelligence 101 (1-2) (1998) 99–134.

[50] D. Tolpin, J. W. van de Meent, H. Yang, F. Wood, Design and Implementation of Probabilistic Programming Language Anglican, arXiv preprint arXiv:1608.05263.

[51] K. Bauters, K. McAreavey, J. Hong, Y. Chen, W. Liu, L. Godo, C. Sierra, Probabilistic planning in agentspeak using the pomdp framework, in: I. Hatzilygeroudis, V. Palade, J. Prentzas (Eds.), Combinations of Intelligent Methods and Applications, Springer International Publishing, 2016, pp. 19–37.

[52] V. Belle, Probabilistic Planning by Probabilistic Programming, CoRR Archive abs/1801.08365.

[53] S. Ross, J. Pineau, S. Paquet, B. Chaib-draa, Online planning algorithms for pomdps, CoRR abs/1401.3436.

Appendices

A. Notation and Notational Conventions Used in Section 3

Symbol(s)	Used For	See Def(s)
$\mathcal{I}$	the set of all instants	Def. 3.1 (p. 12)
$I, I', I_1, I'', I_2, \ldots$	elements of $\mathcal{I}$ (instants)	Def. 3.1 (p. 12)
$\mathcal{A}$, $\mathcal{A}_e$, $\mathcal{A}_a$	the sets of all actions, environmental actions and agent actions	Def. 3.1 (p. 12)
$A, A', A_1, A'', A_2, \ldots$	elements of $\mathcal{A}$ (actions)	Def. 3.1 (p. 12)
$\mathcal{F}$	the set of all fluents	Def. 3.1 (p. 12)
$F, F', F_1, F'', F_2, \ldots$	elements of $\mathcal{F}$ (fluents)	Def. 3.1 (p. 12)
$\mathcal{V}$	the set of all values	Def. 3.1 (p. 12)
$V, V', V_1, V'', V_2, \ldots$	elements of $\mathcal{V}$ (values)	Def. 3.1 (p. 12)
Θ	the set of all formulas	Def. 3.4 (p. 14)
$\theta, \theta', \theta_1, \theta'', \theta_2, \ldots$	formulas	Def. 3.4 (p. 14)
Φ	the set of all i-formulas	Def. 3.4 (p. 14)
$\varphi, \varphi', \varphi_1, \varphi'', \varphi_2, \ldots$	i-formulas	Def. 3.4 (p. 14)
$P, P', P_1, P'', P_2, \ldots$	real values in $[0, 1]$ (probabilities)	
$P^+, P^+_1, P^+_2, \ldots$	real values in $(0, 1]$ (non-zero probabilities)	
$\mathcal{S}$	the set of all states	Def. 3.7 (p. 16)
$S, S', S_1, S'', S_2, \ldots$	elements of $\mathcal{S}$ (states)	Def. 3.7 (p. 16)
$\mathcal{X}$	the set of all partial states	Def. 3.7 (p. 16)
$X, X', X_1, X'', X_2, \ldots$	elements of $\mathcal{X}$ (partial states)	Def. 3.7 (p. 16)
$\tilde{\mathcal{S}}$	the set of all fluent states	Def. 3.7 (p. 16)

table continued on next page

Symbol(s)	Used For	See Def(s)
$\tilde{S}, \tilde{S}', \tilde{S}_1, \tilde{S}'', \tilde{S}_2, \ldots$	elements of $\tilde{\mathcal{S}}$ (fluent states)	Def. 3.7 (p. 16)
$\tilde{\mathcal{X}}$	the set of all partial fluent states	Def. 3.7 (p. 16)
$\tilde{X}, \tilde{X}', \tilde{X}_1, \tilde{X}'', \tilde{X}_2, \ldots$	elements of $\tilde{\mathcal{X}}$ (partial fluent states)	Def. 3.7 (p. 16)
$\restriction\mathcal{F}$, $\restriction\mathcal{A}$	restriction operators for states and partial states	Def. 3.7 (p. 16)
$\mathcal{O}$	the set $\tilde{\mathcal{X}} \times (0, 1]$ of all outcomes	Def. 3.8 (p. 16)
$O, O', O_1, O'', O_2, \ldots$	elements of $\mathcal{O}$ (outcomes)	Def. 3.8 (p. 16)
χ, π	projection functions for outcomes, weight for a set of outcomes	Defs 3.8, 3.9 (p. 16)
$\mathcal{D}$, $\mathcal{D}'$, $\mathcal{D}_x$, etc.	domain descriptions	Def. 3.14 (p. 19)
$\mathcal{N}$, $\mathcal{N}'$, $\mathcal{N}_x$, etc.	ne-domain descriptions	Def. 3.15 (p. 20)
$\mathcal{W}$	the set of all worlds	Def. 3.16 (p. 20)
$W, W', W_1, \ldots$	elements of $\mathcal{W}$ (worlds)	Def. 3.16 (p. 20)
$\Vert\!\models$	satisfaction/entailment for worlds/i-formulas	Def. 3.17 (p. 21)
CWA	Closed World Assumption for Actions	Def. 3.18 (p. 22)
ic, ic', etc.	initial choice	Def. 3.19 (p. 22)
$occ_{\mathcal{D}}(W)$	the set of instances for which a cause occurs in W	Def. 3.20 (p. 23)
$cprop_{\mathcal{D}}(W, I)$	the c-proposition activated at I in W	Def. 3.20 (p. 23)
ec, ec', etc.	effect choice	Def. 3.21 (p. 23)
$\tilde{S} \oplus \tilde{X}$	update of $\tilde{S}$ w.r.t. $\tilde{X}$	Def. 3.22 (p. 23)
$\mathcal{W}_{\mathcal{N}}$	the set of well-behaved worlds w.r.t. $\mathcal{N}$	Def. 3.24 (p. 24)

table continued on next page

Symbol(s)	Used For	See Def(s)
$\|\models_{\mathcal{N}}$	$\mathcal{N}$-entailment	Def. 3.25 (p. 25)
tr, tr', ...	trace	Def. 3.26 (p. 26)
$\epsilon(tr)$	evaluation of a trace	Def. 3.26 (p. 26)
$\epsilon(\mathcal{N}, W)$	evaluation of a narrative	Def. 3.27 (p. 27)
$M^{ne}_{\mathcal{N}}$	ne-model-function	Def. 3.28 (p. 27)
$((\theta, X), V)$	sensing outcome	Def. 3.29 (p. 29)
$\mathcal{H}$	the set of all sensing and acting histories (histories)	Def. 3.30 (p. 30)
$H, H', H_1, H'', H_2, \ldots$	elements of $\mathcal{H}$ (histories)	Def. 3.30 (p. 30)
$[H]_{<I}$	class of histories indistinguishable from H up to I	Def. 3.30 (p. 30)
(W, H)	h-world	Def. 3.31 (p. 30)
$((\theta, X), V, V')$	sensing occurrence	Def. 3.32 (p. 32)
$socc_{\mathcal{D}}((W, H), I)$	the set of sensing occurrences in (W, H) at I	Def. 3.32 (p. 32)
CWSA	closed world assumption for sensing and acting	Def. 3.33 (p. 32)
$\epsilon_{\mathcal{D}}(H \mid W)$	history evaluation (of H given W w.r.t. $\mathcal{D}$)	Def. 3.34 (p. 33)
$\tilde{M}^{\mathcal{N}}_{\mathcal{D}}$	pre-model-function of $\mathcal{D}$ w.r.t. $\mathcal{N}$	Def. 3.36 (p. 35)
$R(\mathcal{D}, H)$	reduct set of $\mathcal{D}$ w.r.t. H	Def. 3.39 (p. 38)
$\mathcal{R}^{\mathcal{D}}_{H}$	reduct of $\mathcal{D}$ w.r.t. H	Prop. 3.3 (p. 38)
$M_{\mathcal{D}}$	model-function	Def. 3.41 (p. 39)

table continued on next page

Proposition	Proposition Form	Intuitive meaning
v-proposition	F **takes-values** $\langle V_1, \ldots, V_m \rangle$	Declares what values fluent F may take
o-proposition	A **occurs-at** I **with-prob** P^+ **if-holds** θ	Environmental action A occurs at instant I with probability P^+ if formula θ holds
p-proposition	A **performed-at** I **with-prob** P^+ **if-believes** $(\theta, \bar{P})$	Agent action A is performed at instant I with probability P^+ if formula θ is believed to hold with probability in the interval $\bar{P}$
c-proposition	θ **causes-one-of** $\{O_1, \ldots, O_m\}$	For each $O_i = (\tilde{X}_i, P_i^+)$ there is a probability of P_i^+ that formula θ will cause the changes identified in partial fluent state $\tilde{X}_i$
i-proposition	**initially-one-of** $\{O_1, \ldots, O_m\}$	For each $O_i = (\tilde{S}_i, P_i^+)$ there is a probability of P_i^+ that the initial state of the world is the one identified in the fluent state $\tilde{S}_i$
s-proposition	θ **senses** X **with-accuracies** $\mathbf{M}$	θ holding causes the value of the fluent or action X to be sensed with accuracy given by the confusion matrix $\mathbf{M}$

B. Probability Tree Diagram for the Epectisis Example

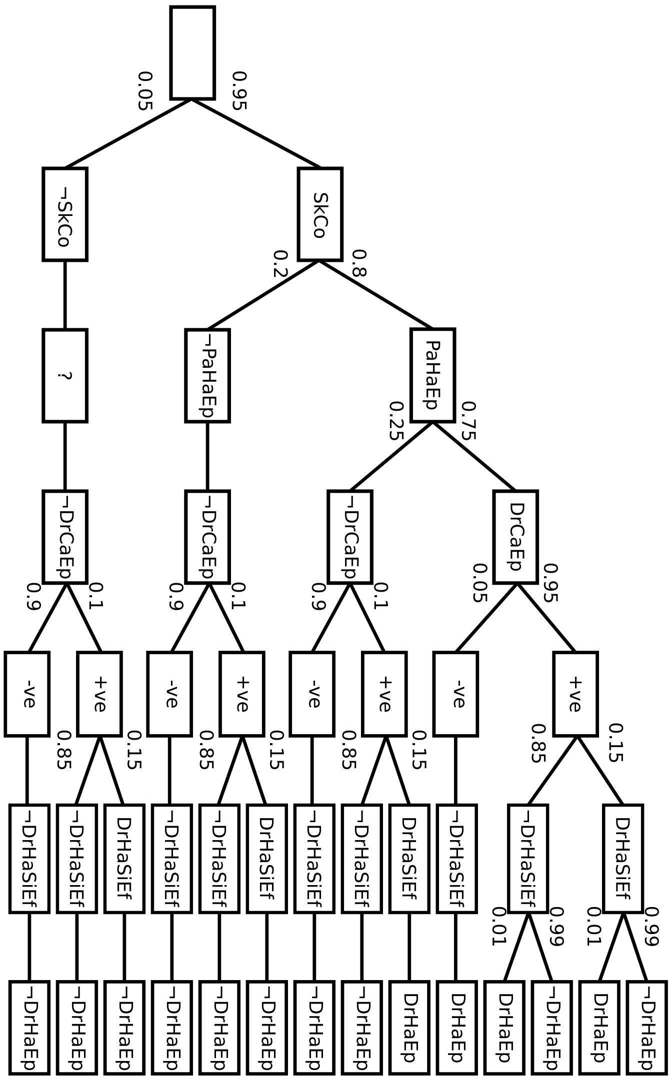

Figure 1: Tree diagram of potential events in the epectitis Example 1.1. Columns refer to key aspects on which the narrative may differ, and forks indicate narrative branches with branching probabilities given. Nodes labelled: ? indicate that the infection status of the patient can take either value; *DrCaEp* (¬*DrCaEp*) indicate that the Doctor does (does not) catch Epectitis from the patient; +ve (-ve) indicate a positive (negative) blood test; and *DrHaEp* (¬*DrHaEp*) indicate that the doctor has (does not have) epectitis at the final instant. Other labels reuse shorthand introduced in Section 3.4.

C. Domain Language and Description for Example 1.1

This is the EPEC domain language and domain description $\mathcal{D}_e$ for the Epectisis example, for simplicity taking the set of instants to be a finite set of integers.

Domain language: $\langle \mathcal{F}^{doc}, \mathcal{A}^{doc}, \mathcal{A}_e^{doc}, \mathcal{A}_a^{doc}, \mathcal{V}^{doc}, vals^{doc}, \mathcal{I}^{doc}, \leq, -1\rangle$, where

- $\mathcal{F}^{doc} = \{PatientHasEpectisis, DocHasEpectisis, DocHasSideEffects\}$,
- $\mathcal{A}_e^{doc} = \{SkinContact\}$,
- $\mathcal{A}_a^{doc} = \{DoBloodTest, TakeEpecillin\}$,
- $\mathcal{V}^{doc} = \{true, false\}$,
- $\mathcal{I}^{doc} = \{-1, 0, 1, 2, 3\}$.

Domain description $\mathcal{D}_e$:

$PatientHasEpectisis$ **takes-values** $\langle false, true\rangle$ (EP1)

$DocHasEpectisis$ **takes-values** $\langle false, true\rangle$ (EP2)

$DocHasSideEffects$ **takes-values** $\langle false, true\rangle$ (EP3)

initially-one-of (EP4)
$\{\ (\{PatientHasEpectisis = true, DocHasEpectisis = false, DocHasSideEffects = false\}, 0.8),$
$(\{PatientHasEpectisis = false, DocHasEpectisis = false, DocHasSideEffects = false\}, 0.2)\ \}$

$SkinContact$ **occurs-at** -1 **with-prob** 0.95 (EP5)

$(SkinContact = true \wedge PatientHasEpectisis = true)$ (EP6)
causes-one-of $\{\ (\{DocHasEpectisis = true\}, 0.75), (\emptyset, 0.25)\ \}$

$TakeEpecillin = true \wedge DocHasEpectisis = true$ **causes-one-of** (EP7)
$\{\ (\{DocHasEpectisis = false, DocHasSideEffects = true\}, 0.1485),$
$(\{DocHasEpectisis = false\}, 0.8415),$
$(\{DocHasSideEffects = true\}, 0.0015),$
$(\emptyset, 0.0085)\ \}$

$TakeEpecillin = true \wedge DocHasEpectisis = false$ **causes-one-of** (EP8)
$\{\ (\{DocHasSideEffects = true\}, 0.15),$
$(\emptyset, 0.85)\ \}$

$$\textit{DoBloodTest} = \textit{true}\ \textbf{senses}\ \textit{DocHasEpectisis}\ \textbf{with-accuracies} \begin{pmatrix} 0.9 & 0.1 \\ 0.05 & 0.95 \end{pmatrix} \tag{EP9}$$

$$\textit{DoBloodTest}\ \textbf{performed-at}\ 1 \tag{EP10}$$

$$\textit{TakeEpecillin}\ \textbf{performed-at}\ 2\ \textbf{if-believes}\ (\textit{DocHasEpectisis} = \textit{true}, (0.5, 1]) \tag{EP11}$$

Example EPEC entailments:

"The doctor currently has a 57% belief that she has epectisis":
$\mathcal{D}_e \mathrel{|}\joinrel\models$ **at** 0 **believes** $[\textit{DocHasEpectisis} = \textit{true}]@0$ **with-probs** $\{(\langle\rangle, 1, 0.57)\}$

"At time 3 [i.e. after the blood test and possibly taking epicillin] there is a 0.425 probability that the blood test will have shown negative, in which case the doctor will be 93.14% sure she does not have the disease, and a 0.575 probability that the blood test will have shown positive, in which case the doctor will have taken epecillin and will therefore be 99.07% confident that she no longer has the disease":
$\mathcal{D}_e \mathrel{|}\joinrel\models$ **at** 3 **believes** $[\textit{DocHasEpectisis} = \textit{false}]@3$ **with-probs**
$\{\ (\langle\{((\textit{DoBloodTest}, \textit{DocHasEpectisis}), \textit{false})\}@1\rangle, 0.425, 0.9314),$
$(\langle\{((\textit{DoBloodTest}, \textit{DocHasEpectisis}), \textit{true})\}@1, \{\textit{TakeEpecillin}\}@2\rangle, 0.575, 0.9907)\}$

"The doctor currently has a 8.77% belief that she will have side effects after executing her plan":
$\mathcal{D}_e \mathrel{|}\joinrel\models$ **at** 0 **believes** $[\textit{DocHasSideEffects} = \textit{true}]@3$ **with-probs** $\{(\langle\rangle, 1, 0.087675)\}$

D. Proofs of Propositions

D.1. Proof of Proposition 3.1

The statement of the proposition is:

Proposition 3.1. Given an ne-domain description $\mathcal{N}$ [see Def. 3.15], the ne-model $M^{ne}_{\mathcal{N}} : \mathcal{W}_{\mathcal{N}} \mapsto [0,1]$ [see Def. 3.28] is a probability distribution over $\mathcal{W}_{\mathcal{N}}$.

We prove this using a series of auxilliary definitions, notation, lemmas and corollaries as follows.

Definition D.1 (Restricted Domain Description)**.** Let $\mathcal{N}$ be an ne-domain description. $\mathcal{N}_{\leq I}$ denotes the ne-domain description obtained from $\mathcal{N}$ by removing all o-propositions with instants $> I$, $\mathcal{N}_{<I}$ denotes the ne-domain description obtained by removing all o-propositions with instants $\geq I$, and $\mathcal{N}_{\emptyset}$ denotes the ne-domain description obtained by removing all o-propositions, i.e. $\mathcal{N}_{\emptyset} = \mathcal{N}_{<\bar{0}}$. [end definition]

Definition D.2 (Fluent-indistinguishable, Indistinguishable)**.** Let W and W' be worlds and I be an instant. W is *fluent-indistinguishable from* W' *up to* I iff $W(I')\restriction\mathcal{F} = W'(I')\restriction\mathcal{F}$ for all $I' \leq I$. W is *indistinguishable from* W' *up to* I iff it is fluent indistinguishable from W' up to I and $W(I') = W'(I')$ for all $I' < I$. [end definition]

Definition D.3 (Occurence Narrative)**.** An *occurrence narrative* (usually denoted as $\mathbf{N}$ with appropriate subscript/superscript) is a finite set of o-propositions. For $\mathcal{N}$ an ne-domain description, the occurrence narrative *narr*($\mathcal{N}$) is the set of o-propositions in $\mathcal{N}$. As an extension of Definition 3.27, the *evaluation of an occurrence narrative* $\mathbf{N}$ *w.r.t. a world* W, denoted $\epsilon(\mathbf{N}, W)$, is defined as:

$$\epsilon(\mathbf{N}, W) = \prod_{\underline{\mathbf{o}} \in \mathbf{N}} \epsilon(\underline{\mathbf{o}}, W).$$

If $\mathbf{N}$ contains no o-propositions then $\epsilon(\mathbf{N}, W) = 1$. [end definition]

Definition D.4 (InstantOf)**.** For an o-proposition $\underline{\mathbf{o}}$, *instantOf*($\underline{\mathbf{o}}$) signifies the instant $\underline{\mathbf{o}}$ has, and for an occurrence narrative $\mathbf{N}$ this is extended to:

$$\mathit{occInstants}(\mathbf{N}) = \{\mathit{instantOf}(\underline{\mathbf{o}}) \mid \underline{\mathbf{o}} \in \mathbf{N}\}$$ [end definition]

Notation D.1 (Expanded Notation for a Trace)**.** [see Def 3.26] We will sometimes write the trace $\langle ic, ec\rangle$ of W w.r.t. $\mathcal{N}$ as the tuple $\langle ic, ec(I_1), \ldots, ec(I_n)\rangle$ of outcomes, where $\{I_1, \ldots, I_n\} = occ_{\mathcal{D}}(W)$ and $I_j < I_{j+1}$ for all $1 \leq j < n$.

The transition function in the following definition gives the probability of moving from state S to the fluent state $\tilde{S}'$ within $\mathcal{N}$, independently of its particular narrative:

Definition D.5 (Transition Set, Transition Function)**.** Given an ne-domain description $\mathcal{N}$, a state S and a fluent state $\tilde{S}'$, the *transition set* $tset_{\mathcal{N}}(S, \tilde{S}')$ is defined as follows:

- if $\mathcal{N}$ contains a (unique) c-proposition $\underline{\mathbf{c}}$ such that $S \models body(\underline{\mathbf{c}})$, then $tset_{\mathcal{N}}(S, \tilde{S}') = \{O \in head(\underline{\mathbf{c}}) \mid (S{\restriction}\mathcal{F}) \oplus \chi(O) = \tilde{S}'\}$,
- if there is no such c-proposition and $S{\restriction}\mathcal{F} = \tilde{S}'$ then $tset_{\mathcal{N}}(S, \tilde{S}') = \{(\emptyset, 1)\}$,
- otherwise, $tset_{\mathcal{N}}(S, \tilde{S}') = \emptyset$.

The *transition function* for $\mathcal{N}$ is the function $t_{\mathcal{N}} : \mathcal{S} \times \tilde{\mathcal{S}} \rightarrow [0, 1]$ defined by $t_{\mathcal{N}}(S, \tilde{S}') = \pi(tset_{\mathcal{N}}(S, \tilde{S}'))$ [with π as in Definition 3.9]. [end definition]

The transition function is used in the following lemma to express $M^{ne}_{\mathcal{N}}(W)$ in terms of the model of a well-behaved world w.r.t. an appropriately restricted domain description:

Lemma D.1. Let $\mathcal{N}$ be an ne-domain description and W a world such that $occ_{\mathcal{N}}(W) = \{I_1, \ldots, I_n\} \neq \emptyset$ where $I_1, \ldots, I_n$ are ordered w.r.t. $\leq$, and let $\underline{\mathbf{c}}$ be the c-proposition activated in W at I_n w.r.t. $\mathcal{N}$. Then W is well-behaved w.r.t. $\mathcal{N}$ if and only if (i) there exists a unique world W' well-behaved w.r.t. $\mathcal{N}_{<I_n}$ which is indistinguishable from W up to I_n, (ii) for all $I > I_n$, $W(I){\restriction}\mathcal{F} = \tilde{S}^{W}_{>I_n}$ where $\tilde{S}^{W}_{>I_n} = (W(I_n){\restriction}\mathcal{F}) \oplus \chi(O)$ for some outcome $O \in head(\underline{\mathbf{c}})$, and (iii) W satisfies the CWA for actions w.r.t. $\mathcal{N}$.

Furthermore, if W is well-behaved w.r.t. $\mathcal{N}$, and putting $\mathcal{N}' = \mathcal{N}_{<I_n}$, then

$$M^{ne}_{\mathcal{N}}(W) = \frac{\epsilon(\mathcal{N}, W)}{\epsilon(\mathcal{N}', W')} \cdot M^{ne}_{\mathcal{N}'}(W') \cdot t_{\mathcal{N}}(W(I_n), \tilde{S}^{W}_{>I_n}) \tag{17}$$

Proof:

"Only if" subproof:

Let W be well-behaved w.r.t. $\mathcal{N}$. Let $tr = \langle ic, ec(I_1), \ldots, ec(I_n)\rangle$ be an arbitrary trace of W w.r.t. $\mathcal{N}$ and consider the tuple $tr' = \langle ic, ec(I_1), \ldots, ec(I_{n-1})\rangle$. Since W is well-behaved w.r.t. $\mathcal{N}$, since $\mathcal{N}$ and $\mathcal{N}_{<I_n}$ differ only by one or more o-propositions occurring at I_n, and since tr' does not mention any instant strictly greater than I_{n-1}, it is possible to construct a world W' well-behaved w.r.t. $\mathcal{N}_{<I_n}$ which has trace tr' w.r.t. $\mathcal{N}_{<I_n}$ and which is fluent-indistinguishable from W up to instant I_n, by putting $W'(I) = W(I)$ for all $I < I_n$, $W'(I)\restriction\mathcal{F} = W(I_n)\restriction\mathcal{F}$ for all $I \geq I_n$, and $\neg A \in W'(I)$ for all $I \geq I_n$ and $A \in (A)$. W' is well-behaved by construction, and unique because $\mathcal{N}_{<I_n}$ does not contain o-propositions with instants greater than I_{n-1} so that no other assignment of states to $W'(I)$ for $I \geq I_n$ will satisfy the CWA and justified change condition w.r.t. $\mathcal{N}_{<I_n}$. Condition (ii) is satisfied since W satisfies the justified change condition w.r.t. $\mathcal{N}$, and condition (iii) is satisfied by definition since W is well-behaved.

"If" subproof:

Let W' be well-behaved w.r.t. $\mathcal{N}_{<I_n}$ and let $occ_{\mathcal{N}_{<I_n}}(W') = \{I_1, \ldots, I_{n-1}\}$. Let $tr' = \langle ic, ec(I_1), \ldots, ec(I_{n-1})\rangle$ be a trace of W' w.r.t. $\mathcal{N}_{<I_n}$ and construct the tuple $tr = \langle ic, ec(I_1), \ldots, ec(I_{n-1}), O\rangle$ for the outcome $O \in head(\underline{\mathbf{c}})$ such that $(W(I)\restriction\mathcal{F}) = (W(I_n)\restriction\mathcal{F}) \oplus \chi(O)$ for all $I > I_n$. Since W' is well-behaved w.r.t. $\mathcal{N}_{<I_n}$ and indistinguishable from W up to I_n by hypothesis (i), then tr is a trace of W w.r.t. $\mathcal{N}$. Since $\mathcal{N}$ and $\mathcal{N}_{<I_n}$ share the same i-proposition then W satisfies the initial condition of $\mathcal{N}$. By hypothesis (ii) W satisfies the justified change condition w.r.t. $\mathcal{N}$, and so by hypothesis (iii) W is well-behaved.

"Furthermore" subproof:

Let $\tilde{S}^{W}_{>I_n}$ and $\underline{\mathbf{c}}$ be as in the statement of the proposition. The above subproofs show that any trace tr of W w.r.t. $\mathcal{N}$ can be constructed from a trace tr' of an appropriate W' by appending some $O \in head(\underline{\mathbf{c}})$ such that $\tilde{S}^{W}_{>I_n} = (W(I_n)\restriction\mathcal{F}) \oplus \chi(O)$ for the mapping at I_n, i.e. some $O \in tset_{\mathcal{N}}(W(I_n), \tilde{S}^{W}_{>I_n})$.

Definition 3.28 now implies

$$\begin{aligned}
M^{ne}_{\mathcal{N}}(W) &= \epsilon(\mathcal{N}, W) \quad \cdot \sum_{tr \in traces(W,\mathcal{N})} \epsilon(tr) \\
&= \epsilon(\mathcal{N}', W') \cdot \frac{\epsilon(\mathcal{N}, W)}{\epsilon(\mathcal{N}', W')} \cdot \pi(tset_{\mathcal{N}}(W(I_n), \tilde{S}^W_{>I_n})) \quad \cdot \sum_{tr' \in traces(W',\mathcal{N}')} \epsilon(tr') \\
&= \frac{\epsilon(\mathcal{N}, W)}{\epsilon(\mathcal{N}', W')} \cdot t_{\mathcal{N}}(W(I_n), \tilde{S}^W_{>I_n}) \cdot \left(\epsilon(\mathcal{N}', W') \quad \cdot \sum_{tr' \in traces(W',\mathcal{N}')} \epsilon(tr') \right) \\
&= \frac{\epsilon(\mathcal{N}, W)}{\epsilon(\mathcal{N}', W')} \cdot t_{\mathcal{N}}(W(I_n), \tilde{S}^W_{>I_n}) \cdot M^{ne}_{\mathcal{N}'}(W')
\end{aligned}$$

This is well defined since $\epsilon(\mathcal{N}', W') > 0$ for any ne-domain description $\mathcal{N}'$ and well-behaved world W'. ■

Corollary D.1. Let $\mathcal{N}$ be an ne-domain description and let I be an instant. Then W is well-behaved w.r.t. $\mathcal{N}_{\leq I}$ if and only if (i) there exists a unique world W' well-behaved w.r.t. $\mathcal{N}_{<I}$ which is indistinguishable from W up to I, (ii) for all $I' > I$ and, if $I \in occ_{\mathcal{N}_{\leq I}}(W)$, some $O \in head(cprop_{\mathcal{N}_{\leq I}}(W, I))$, then $W(I') \restriction \mathcal{F} = \tilde{S}^W_{>I}$ where

$$\tilde{S}^W_{>I} = \begin{cases} (W(I) \restriction \mathcal{F}) \oplus \chi(O) & \text{if } I \in occ_{\mathcal{N}_{\leq I}}(W) \\ W(I) \restriction \mathcal{F} & \text{otherwise} \end{cases}$$

and (iii) W satisfies the CWA for actions w.r.t. $\mathcal{N}_{\leq I}$.
Furthermore, when W is well-behaved w.r.t. $\mathcal{N}_{\leq I}$:

$$M^{ne}_{\mathcal{N}_{\leq \mathcal{I}}}(W) = \frac{\epsilon(\mathcal{N}_{\leq I}, W)}{\epsilon(\mathcal{N}_{<I}, W')} \cdot M^{ne}_{\mathcal{N}_{<\mathcal{I}}}(W') \cdot t_{\mathcal{N}}(W(I), \tilde{S}^W_{>I}) \tag{18}$$

Proof:
If $I \in occ_{\mathcal{N}_{\leq I}}(W)$ then the corollary follows directly from Lemma D.1 since the domain description $\mathcal{D}_{\leq I}$ satisfies all of its hypotheses. If $I \notin occ_{\mathcal{N}_{\leq I}}(W)$ then the if and only if parts of the corollary follow directly by setting $W = W'$, and as no c-proposition is activated in W at I and therefore W has the same traces w.r.t. both $\mathcal{N}_{\leq I}$ and $\mathcal{N}_{<I}$, equation (18) follows since $\epsilon(\mathcal{N}_{\leq I}, W) = \epsilon(\mathcal{N}_{<I}, W)$, $M^{ne}_{\mathcal{N}_{\leq \mathcal{I}}}(W) = M^{ne}_{\mathcal{N}_{<\mathcal{I}}}(W)$, and $t_{\mathcal{N}}(W(I), \tilde{S}^W_{>I}) = 1$. ■

Lemma D.2 (Transition Function Normalisation)**.** For any ne-domain description $\mathcal{N}$ and any state S, $\sum_{\tilde{S}' \in \tilde{\mathcal{S}}} t_{\mathcal{N}}(S, \tilde{S}') = 1$.

Proof:

Proof is by cases:

Case 1. If there is no c-proposition $\underline{\mathbf{c}}$ such that S entails $body(\underline{\mathbf{c}})$, then it follows from Definition D.5 that

$$\sum_{\tilde{S}' \in \tilde{\mathcal{S}}} t_{\mathcal{N}}(S, \tilde{S}') = t_{\mathcal{N}}(S, S \restriction \mathcal{F}) = \pi((\emptyset, 1)) = 1$$

Case 2. Let $\underline{\mathbf{c}}$ be the unique c-proposition such that $S \models body(\underline{\mathbf{c}})$. Then applying the definition of $t_{\mathcal{N}}$ from Definition D.5 gives

$$\sum_{\tilde{S}' \in \tilde{\mathcal{S}}} t_{\mathcal{N}}(S, \tilde{S}') = \sum_{\tilde{S}' \in \tilde{\mathcal{S}}} \pi(tset_{\mathcal{N}}(S, \tilde{S}')) \tag{19}$$

Notice that for a fixed outcome O, it is impossible to have $O \in tset_{\mathcal{N}}(S, \tilde{S}')$ and $O \in tset_{\mathcal{N}}(S, \tilde{S}'')$ for two distinct fluent states $\tilde{S}', \tilde{S}''$ as this would imply $\tilde{S}' = (S \restriction \mathcal{F}) \oplus \chi(O) = \tilde{S}''$. Hence it is sufficient to show that $\{O \in tset_{\mathcal{N}}(S, \tilde{S}') \mid \tilde{S}' \in \tilde{\mathcal{S}}\} = head(\underline{\mathbf{c}})$, as this implies that the sum (19) equals 1, since $\pi(head(\underline{\mathbf{c}})) = 1$ by definition of a c-proposition.

By definition of a transition set, $\{O \in tset_{\mathcal{N}}(S, \tilde{S}') \mid \tilde{S}' \in \tilde{\mathcal{S}}\} \subseteq head(\underline{\mathbf{c}})$. Conversely, for any $O \in head(\underline{\mathbf{c}})$, $O \in tset_{\mathcal{N}}(S, \tilde{S}')$ for $\tilde{S}' = (S \restriction \mathcal{F}) \oplus \chi(O)$, hence $head(\underline{\mathbf{c}}) \subseteq \{O \in tset_{\mathcal{N}}(S, \tilde{S}') \mid \tilde{S}' \in \tilde{\mathcal{S}}\}$. ■

Lemma D.3 (Occurrence Narrative Normalisation)**.** Let $\mathcal{N}$ be an ne-domain description, I be an instant and $\mathbf{N}_I$ be the (possibly empty) occurrence narrative that contains exactly those o-propositions in $\mathcal{N}$ that have instant I. Let $\{W_1, \ldots, W_m\}$ be a maximal set of well-behaved worlds w.r.t. $\mathcal{N}$ such that $W_i(I) \restriction \mathcal{F} = W_j(I) \restriction \mathcal{F}$ for all $1 \leq i, j \leq m$ and $W_i(I) \neq W_j(I)$ when $i \neq j$. Then

$$\sum_{j=1}^{m} \epsilon(\mathbf{N}_I, W_j) = 1$$

Proof:

Fix a maximal set $\{W_1, \ldots, W_m\}$ as in the hypothesis and let $\tilde{S}$ be the fluent state such that $\tilde{S} = W_i(I)\lceil\mathcal{F}$ for all $1 \leq i \leq m$. Let $\{\underline{\mathbf{o}}_1, \ldots, \underline{\mathbf{o}}_k\}$ be a maximal set of o-propositions in $\mathbf{N}_I$ of the form

$$A_i \textbf{ occurs-at } I \textbf{ with-prob } P_i^+ \textbf{ if-holds } \theta_i$$

such that $\tilde{S} \models \theta_i$ and $P_i^+ < 1$, for all $1 \leq i \leq k$.

From maximality of $\{W_1, \ldots, W_m\}$ and the CWA for actions it follows that $m = 2^k$, with every world in this set having a different assignment of actions $A_1, \ldots, A_k$ to truth values at instant I. Let $\bar{x} = \langle x_1, \ldots, x_k \rangle$ be a k-dimensional vector representing a specific assignment of $A_1, \ldots, A_k$ to truth values, where each $x_i \in \{0, 1\}$. Therefore the sum $\sum_{j=1}^{m} \epsilon(\mathbf{N}_I, W_j)$ evaluates to:

$$\sum_{\bar{x}} \left(\prod_{i=1}^{k} (P_i^+)^{x_i} \left(1 - P_i^+\right)^{1-x_i} \right) = \prod_{i=1}^{k} \left(P_i^+ + (1 - P_i^+)\right) = 1$$

■

Lemma D.4 (Causality)**.** Let $\mathcal{N}$ be an ne-domain description and $\mathcal{N}'$ be an ne-domain description such that $\mathcal{N}' = \mathcal{N} \cup \mathbf{N}_{I'}$ where $\mathbf{N}_{I'}$ is an occurrence narrative such that all o-propositions in $\mathbf{N}_{I'}$ have instant I', and all o-propositions in $\mathcal{N}$ have instants less than I'. Then, for any $I < I'$ and any formula θ,

$$M^{ne}_{\mathcal{N}'}([\theta]@I) = M^{ne}_{\mathcal{N}}([\theta]@I)$$

Proof:

In the summations in the equations below, W ranges over the set w.r.t. well-behaved worlds w.r.t. $\mathcal{N}$, W' ranges over the set of well-behaved worlds w.r.t. $\mathcal{N}'$, $[W]^{I'}_{\mathcal{N}'}$ signifies the set of well-behaved worlds w.r.t. $\mathcal{N}'$ that are indistinguishable from W up to I', and $\tilde{S}^{W'}_{>I'}$ is as defined in Corollary D.1. Corollary D.1 and Lemma D.2 yield:

$$\begin{aligned}
M^{ne}_{\mathcal{N}'}([\theta]@I) &= \sum_{W' \models [\theta]@I} M^{ne}_{\mathcal{N}'}(W') \\
&\overset{\text{Cor.D.1}}{=} \sum_{W \models [\theta]@I} \; \sum_{W' \in [W]^{I'}_{\mathcal{N}'}} \frac{\epsilon(\mathcal{N}', W')}{\epsilon(\mathcal{N}, W)} \cdot M^{ne}_{\mathcal{N}}(W) \cdot t_{\mathcal{N}'}(W'(I'), \tilde{S}^{W'}_{>I'}) \\
&= \sum_{W \models [\theta]@I} M^{ne}_{\mathcal{N}}(W) \sum_{W' \in [W]^{I'}_{\mathcal{N}'}} \frac{\epsilon(\mathcal{N}', W')}{\epsilon(\mathcal{N}, W)} \cdot t_{\mathcal{N}'}(W'(I'), \tilde{S}^{W'}_{>I'})
\end{aligned}$$

According to Corollary D.1 every world in $[W]^{I'}_{\mathcal{N}'}$ can be reconstructed from its state at instant I' and the single persisting state that it takes at instants strictly greater than I'. Let $\{W'_1, \ldots, W'_m\}$ be a maximal set of well-behaved worlds in $[W]^{I'}_{\mathcal{N}'}$ such that $W'_i(I') \neq W'_j(I')$ when $i \neq j$ for $1 \leq i, j \leq m$. Then the above chain of equalities continues as follows:

$$
\begin{aligned}
&= \sum_{W|\models[\theta]@I} M^{ne}_{\mathcal{N}}(W) \sum_{j=1}^{m} \sum_{\tilde{S}' \in \tilde{\mathcal{S}}} \frac{\epsilon(\mathcal{N}', W'_j)}{\epsilon(\mathcal{N}, W)} \cdot t_{\mathcal{N}'}(W'_j(I'), \tilde{S}') \\
&= \sum_{W|\models[\theta]@I} M^{ne}_{\mathcal{N}}(W) \sum_{j=1}^{m} \frac{\epsilon(\mathcal{N}', W'_j)}{\epsilon(\mathcal{N}, W)} \cdot \sum_{\tilde{S}' \in \tilde{\mathcal{S}}} t_{\mathcal{N}'}(W'_j(I'), \tilde{S}') \\
&\overset{\text{Lem.D.2}}{=} \sum_{W|\models[\theta]@I} M^{ne}_{\mathcal{N}}(W) \sum_{j=1}^{m} \frac{\epsilon(\mathcal{N}', W'_j)}{\epsilon(\mathcal{N}, W)} \\
&\overset{\text{Lem.D.3}}{=} \sum_{W|\models[\theta]@I} M^{ne}_{\mathcal{N}}(W) = M^{ne}_{\mathcal{N}}([\theta]@I)
\end{aligned}
$$

and notice that we can apply Lemma D.3 since $\frac{\epsilon(\mathcal{N}', W_j)}{\epsilon(\mathcal{N}, W)} = \epsilon(\mathbf{N}_{I'}, W_j)$. ■

We can now prove Proposition 3.1, re-stated again here:

Proposition 3.1. Given an ne-domain description $\mathcal{N}$ [see Def. 3.15], the ne-model $M^{ne}_{\mathcal{N}} : \mathcal{W}_{\mathcal{N}} \mapsto [0, 1]$ [see Def. 3.28] is a probability distribution over $\mathcal{W}_{\mathcal{N}}$.

Proof:

We need to show that

$$\sum_{W \in \mathcal{W}_{\mathcal{N}}} M^{ne}_{\mathcal{N}}(W) = 1 \qquad (20)$$

This is proved by induction on the size of *occInstants*$(narr(\mathcal{N}))$:

Base Case: Suppose *occInstants*$(narr(\mathcal{N})) = \emptyset$, i.e. $\mathcal{N}$ contains no o-propositions. Since each $W \in \mathcal{W}_{\mathcal{N}}$ satifies the CWA w.r.t. $\mathcal{N}$, $occ_{\mathcal{N}}(W) = \emptyset$. So each $W \in \mathcal{W}_{\mathcal{N}}$ has a unique trace of the form $\langle ic \rangle$ for some initial choice of $\mathcal{N}$, and can be written as $W_{\langle ic \rangle}$. by Definition 3.27, $\epsilon(\mathcal{N}, W) = 1$ for every world W, so that by Equation (13) of Definition 3.28:

$$\sum_{W \in \mathcal{W}_{\mathcal{N}}} M^{ne}_{\mathcal{N}}(W) \;=\; \sum_{W_{\langle ic \rangle} \in \mathcal{W}_{\mathcal{N}}} \epsilon(\langle ic \rangle)$$

By Definition 3.19 each initial choice ic is an O_i taken from the unique i-proposition "**initially-one-of** $\{O_1, \ldots, O_m\}$" of $\mathcal{N}$, so by Definition 3.11 and

Definition 3.26:

$$\sum_{W_{\langle ic \rangle} \in \mathcal{W}_{\mathcal{N}}} \epsilon(\langle ic \rangle) \;=\; \sum_{i=1}^{m} \pi(O_i) \;=\; 1$$

Inductive step: Suppose that $\mathit{occInstants}(\mathit{narr}(\mathcal{N})) = \{I_1, \ldots, I_n\}$ where $I_1, \ldots, I_n$ are ordered w.r.t. $\leq$, so that $\mathcal{N} = \mathcal{N}_{\leq I_n}$. Let $\mathcal{N}' = \mathcal{N}_{<I_n}$ so that by the inductive hypothesis $\sum_{W \in \mathcal{W}_{\mathcal{N}'}} M^{ne}_{\mathcal{N}'}(W) = 1$. Since the i-formula $[\top]@\bar{0}$ is satisfied by every world W [see Definitions 3.4 and 3.17] then by Lemma D.4:

$$\sum_{W \in \mathcal{W}_{\mathcal{N}}} M^{ne}_{\mathcal{N}}(W) \;=\; M^{ne}_{\mathcal{N}}([\top]@\bar{0}) \;=\; M^{ne}_{\mathcal{N}'}([\top]@\bar{0}) \;=\; \sum_{W \in \mathcal{W}_{\mathcal{N}'}} M^{ne}_{\mathcal{N}'}(W) \;=\; 1 \qquad \blacksquare$$

D.2. Proof of Proposition 3.2

The statement of the proposition is:

Proposition 3.2. Let $\mathcal{D}$ and $\mathcal{N}$ be domain and ne-domain descriptions in languages $\mathcal{L}$ and $\mathcal{L}^{ne}$ respectively. The pre-model-function $\tilde{M}^{\mathcal{N}}_{\mathcal{D}}$ of $\mathcal{D}$ w.r.t. $\mathcal{N}$ is a probability distribution over $\mathcal{W} \times \mathcal{H}$, i.e. $\sum_{(W,H) \in \mathcal{W} \times \mathcal{H}} \tilde{M}^{\mathcal{N}}_{\mathcal{D}}(W, H) = 1$.

and we prove this via the following lemma and corollary.

Lemma D.5. Let $\mathcal{D}$ be a domain description, (W, H) be an h-world satisfying the CWSA w.r.t. $\mathcal{D}$, $I \in \mathcal{I}$, and let $\mathcal{J}^W_H$ be as defined in Definition 3.34. Then,

$$\epsilon_{\mathcal{D}}([H]_{<I} \mid W) = \prod_{I' \in \mathcal{J}^W_H, I' < I} \left(\prod_{((\theta, X), V_i, V_j) \in \mathit{socc}_{\mathcal{D}}((W,H), I')} \mathbf{M}(\theta, X)_{i,j} \right)$$

Proof:

In all of the following equations I' ranges over (the finite set) $\mathcal{J}^W_H$. By Definition 3.34,

$$\begin{aligned} \epsilon_{\mathcal{D}}([H]_{<I} \mid W) &= \sum_{H' \in [H]_{<I}} \epsilon_{\mathcal{D}}(H' \mid W) \\ &= \sum_{H' \in [H]_{<I}} \left(\prod_{I' \in \mathcal{J}^W_H} \left(\prod_{((\theta, X), V_i, V_j) \in \mathit{socc}_{\mathcal{D}}((W,H'), I')} \mathbf{M}(\theta, X)_{i,j} \right) \right) \end{aligned}$$

Recall that $[H]_{<I}$ is the equivalence class such that if $H', H'' \in [H]_{<I}$ then $H'(I') = H''(I')$ for all $I' < I$. So, for any $X \in \mathcal{F} \cup \mathcal{A}$, $V_i \in vals(X)$ and $I' \geq I$, if $((\theta, X), V_i) \in H(I')$ then for all $V_j \neq V_i$ such that $V_j \in vals(X)$ there exists some other $H' \in [H]_{<I}$ such that (W, H') also satisfies CWSA w.r.t. $\mathcal{D}$ and $((\theta, X), V_j) \in H'(I')$. Therefore the above sum can be rewritten as:

$$\underbrace{\prod_{I'<I}\left(\prod_{((\theta,X),V_i,V_j)\in socc_{\mathcal{D}}((W,H),I')} \mathbf{M}(\theta,X)_{i,j}\right)}_{\text{common factor to all } H \in [H]_{<I}} \prod_{I'\geq I}\left(\sum_{((\theta,X),V_i)\in H(I')}\left(\sum_{j=1}^{m}\mathbf{M}(\theta,X)_{i,j}\right)\right)$$

and since (by Equation (8) of Definition 3.12) $\sum_{j=1}^{m} \mathbf{M}(\theta, X)_{i,j} = 1$ the lemma is proved. ■

Corollary D.2. Let $\mathcal{D}$ be a domain description and W a world. Then

$$\sum_{H\in\mathcal{H}} \epsilon_{\mathcal{D}}(H \mid W) = 1.$$

Proof:

The Corollary follows directly from Lemma D.5 by considering the equivalence class $[H]_{<\bar{0}} = \mathcal{H}$. ■

We can now prove Proposition 3.2, re-stated again here:

Proposition 3.2. Let $\mathcal{D}$ and $\mathcal{N}$ be domain and ne-domain descriptions in languages $\mathcal{L}$ and $\mathcal{L}^{ne}$ respectively. The pre-model-function $\tilde{M}_{\mathcal{D}}^{\mathcal{N}}$ of $\mathcal{D}$ w.r.t. $\mathcal{N}$ is a probability distribution over $\mathcal{W} \times \mathcal{H}$, i.e. $\sum_{(W,H)\in\mathcal{W}\times\mathcal{H}} \tilde{M}_{\mathcal{D}}^{\mathcal{N}}(W, H) = 1$.

Proof:

The proof follows directly from the product rule and Definition 3.36 of $\tilde{M}_{\mathcal{D}}^{\mathcal{N}}$, since $M_{\mathcal{N}}^{ne}$ is a probability distribution over $\mathcal{W}$ [Proposition 3.2] and, for each $W \in \mathcal{W}$, $\epsilon_{\mathcal{D}}(\cdot \mid W)$ is a probability distribution over $\mathcal{H}$ [Corollary D.2]. ■

D.3. Proof of Proposition 3.3

The statement of the proposition is:

Proposition 3.3. Let $\mathcal{D}$ be a domain description and H a history. Then the reduct set $R(\mathcal{D}, H)$ of $\mathcal{D}$ w.r.t. H contains at most one element. If $R(\mathcal{D}, H) \neq \emptyset$ this unique element is denoted $\mathcal{R}_H^{\mathcal{D}}$.

Proof: Proof is by contradiction. Let $\mathcal{N}$ and $\mathcal{N}'$ be two distinct ne-domain descriptions in $R(\mathcal{D}, H)$. Since $\mathcal{N}$ and $\mathcal{N}'$ can differ only in their o-propositions, without loss of generality we can assume that there is an o-proposition $\underline{\mathbf{o}}$ with instant I such that $\underline{\mathbf{o}} \in \mathcal{N}$ and $\underline{\mathbf{o}} \notin \mathcal{N}'$, and that there is no o-proposition with instant less than I on which $\mathcal{N}$ and $\mathcal{N}'$ differ. Furthermore $\underline{\mathbf{o}}$ must be the reduct of some p-proposition $\underline{\mathbf{p}} \in \mathcal{D}$. Let $\underline{\mathbf{p}}$ have body $(\theta, \bar{P})$. Since $\underline{\mathbf{o}} \in \mathcal{N}$ and $\underline{\mathbf{o}} \notin \mathcal{N}'$, $\tilde{M}^{\mathcal{N}}_{\mathcal{D}}(\theta \mid [H]_{<I}) \in \bar{P}$ and $\tilde{M}^{\mathcal{N}'}_{\mathcal{D}}(\theta \mid [H]_{<I}) \notin \bar{P}$ implying

$$\frac{\tilde{M}^{\mathcal{N}}_{\mathcal{D}}([\theta]@I, [H]_{<I})}{\tilde{M}^{\mathcal{N}}_{\mathcal{D}}([H]_{<I})} \neq \frac{\tilde{M}^{\mathcal{N}'}_{\mathcal{D}}([\theta]@I, [H]_{<I})}{\tilde{M}^{\mathcal{N}'}_{\mathcal{D}}([H]_{<I})}$$

To prove that this is not the case, it suffices to show that $\tilde{M}^{\mathcal{N}}_{\mathcal{D}}([\theta]@I, [H]_{<I}) = \tilde{M}^{\mathcal{N}'}_{\mathcal{D}}([\theta]@I, [H]_{<I})$ for an arbitrary θ, as this also accounts for $\tilde{M}^{\mathcal{N}}_{\mathcal{D}}([H]_{<I}) = \tilde{M}^{\mathcal{N}'}_{\mathcal{D}}([H]_{<I})$ by setting $\theta = \top$. Let $\{W_1, \ldots, W_m\}$ be a maximal set of distinct representatives of classes of indistinguishable worlds up to I that are well-behaved w.r.t. $\mathcal{N}$, and let $[W_1]^I_{\mathcal{N}}, \ldots, [W_m]^I_{\mathcal{N}}$ be the classes they represent. For each $[W_i]^I_{\mathcal{N}}$ let $\varphi_{[W_i]^I_{\mathcal{N}}}$ be an i-formula mentioning only instants less than I that uniquely characterises $[W_i]^I_{\mathcal{N}}$ in the sense that $W \mid\models \varphi_{[W_i]^I_{\mathcal{N}}}$ if and only if $W \in [W_i]^I_{\mathcal{N}}$ (such formulas exist because in the limit whole states can be characterised as conjunctions of literals, and well-behaved worlds can be uniquely identified by finite conjunctions of states at action points and states that persist between action points). Furthermore let $\varphi_i = [\theta]@I \wedge \varphi_{[W_i]^I_{\mathcal{N}}}$. Then,

$$\begin{aligned}\tilde{M}^{\mathcal{N}}_{\mathcal{D}}([\theta]@I, [H]_{<I}) &= \sum_{W \mid\models [\theta]@I} \epsilon_{\mathcal{D}}([H]_{<I} \mid W) \cdot M^{ne}_{\mathcal{N}}(W) \\ &= \sum_{i=1}^{m} \left(\sum_{W \mid\models \varphi_i} \epsilon_{\mathcal{D}}([H]_{<I} \mid W) \cdot M^{ne}_{\mathcal{N}}(W) \right)\end{aligned}$$

Lemma D.5 shows that $\epsilon_{\mathcal{D}}([H]_{<I} \mid W)$ depends only on instants up to I. Therefore, for each φ_i, $\epsilon_{\mathcal{D}}([H]_{<I} \mid W)$ has the same value for all W such that $W \mid\models \varphi_i$, so we can denote this value by by $\epsilon_{\mathcal{D}}([H]_{<I} \mid \varphi_i)$. The above sequence of equalities therefore continues as follows:

$$= \sum_{i=1}^{m} \epsilon_{\mathcal{D}}([H]_{<I} \mid \varphi_i) \left(\sum_{W \mid\models \varphi_i} M^{ne}_{\mathcal{N}}(W) \right) = \sum_{i=1}^{m} \epsilon_{\mathcal{D}}([H]_{<I} \mid \varphi_i) \cdot M^{ne}_{\mathcal{N}}(\varphi_i)$$

Since the two ne-domain descriptions $\mathcal{N}$ and $\mathcal{N}'$ disagree only on o-propositions having instant $\geq I$, $M^{ne}_{\mathcal{N}}(\varphi_i) = M^{ne}_{\mathcal{N}'}(\varphi_i)$ for all $1 \leq i \leq m$ and therefore the above sequence of equalities continues:

$$= \sum_{i=1}^{m} \epsilon_{\mathcal{D}}([H]_{<I} \mid \varphi_i) \cdot M^{ne}_{\mathcal{N}'}(\varphi_i) = \tilde{M}^{\mathcal{N}'}_{\mathcal{D}}(\theta, [H]_{<I})$$

which contradicts the assumption that $\mathcal{N}$ and $\mathcal{N}'$ are distinct and therefore completes the proof by contradiction. ■

D.4. Proof of Proposition 3.4

The statement of the proposition is:

Proposition 3.4. Let $\mathcal{D}$ be a domain description. Then the model-function $M_{\mathcal{D}}$ of $\mathcal{D}$ is a probability distribution over $\mathcal{W} \times \mathcal{H}$, i.e. $\sum_{(W,H) \in \mathcal{W} \times \mathcal{H}} M_{\mathcal{D}}(W, H) = 1$.

To prove this, we first adapt Definitions D.3 and D.4 to p-propositions:

Definition D.6 (Performance Narrative)**.** A *performance narrative* (usually denoted as $\mathbf{P}$ with appropriate subscript/superscript) is a finite set of p-propositions. For $\mathcal{D}$ an EPEC domain description, the performance narrative *pnarr*($\mathcal{D}$) is the set of p-propositions in $\mathcal{D}$. [end definition]

Definition D.7 (InstantOf for p-propositions)**.** For a p-proposition $\underline{\mathbf{p}}$, *instantOf*($\underline{\mathbf{p}}$) signifies the instant $\underline{\mathbf{p}}$ has, and for a performance narrative $\mathbf{P}$ this is extended to:

$$\mathit{perfInstants}(\mathbf{P}) = \{\mathit{instantOf}(\underline{\mathbf{p}}) \mid \underline{\mathbf{p}} \in \mathbf{P}\}$$ [end definition]

We will also need the following definitions and lemma:

Definition D.8 (Reducts of Domain Description)**.** For any EPEC domain description $\mathcal{D}$, the set of *reducts w.r.t.* $\mathcal{D}$, written $R(\mathcal{D})$, is the set $\{\mathcal{R}^{\mathcal{D}}_H \in R(\mathcal{D}, H) \mid H \in \mathcal{H}, R(\mathcal{D}, H) \neq \emptyset\}$. [end definition]

Definition D.9 (History Class)**.** Let $\mathcal{D}$ be an EPEC domain description. Let $\{H_1, \ldots, H_n\}$ be a maximal set of histories such that $R(\mathcal{D}, H_i) \neq R(\mathcal{D}, H_j)$ when $i \neq j$ and $R(\mathcal{D}, H_i) \neq \emptyset$ for all $1 \leq i \leq n$. The associated equivalence classes $[H_1]_{\mathcal{D}}, \ldots, [H_n]_{\mathcal{D}}$ are called *history classes of* $\mathcal{D}$. [end definition]

Lemma D.6. Let $\mathcal{D}$ be an EPEC domain description and let $[H_1]_{\mathcal{D}}, \ldots, [H_n]_{\mathcal{D}}$ be its history classes. Then,

$$\sum_{(W,H)} M_{\mathcal{D}}(W,H) = \sum_{W} \sum_{1 \le i \le n} M^{ne}_{\mathcal{R}_{H_i}}(W) \epsilon_{\mathcal{D}}([H_i]_{\mathcal{D}} \mid W)$$

Proof:

In the following, the h-worlds considered in the sums range over the set of h-worlds that are well-behaved w.r.t. $\mathcal{D}$ (since those that are not do not contribute to the sum):

$$\sum_{(W,H)} M_{\mathcal{D}}(W,H) = \sum_{(W,H)} \tilde{M}^{\mathcal{R}_H}_{\mathcal{D}}(W,H) = \sum_{(W,H)} M^{ne}_{\mathcal{R}_H}(W) \cdot \epsilon_{\mathcal{D}}(H \mid W)$$

$$= \sum_{W} \sum_{1 \le i \le n} \sum_{H \in [H_i]_{\mathcal{D}}} M^{ne}_{\mathcal{R}_{H_i}}(W) \cdot \epsilon_{\mathcal{D}}(H \mid W)$$

$$= \sum_{W} \sum_{1 \le i \le n} M^{ne}_{\mathcal{R}_{H_i}}(W) \sum_{H \in [H_i]_{\mathcal{D}}} \epsilon_{\mathcal{D}}(H \mid W) = \sum_{W} \sum_{1 \le i \le n} M^{ne}_{\mathcal{R}_{H_i}}(W) \epsilon_{\mathcal{D}}([H_i]_{\mathcal{D}} \mid W)$$

■

We can now prove the principal proposition, restated here for convenience:

Proposition 3.4. Let $\mathcal{D}$ be a domain description. Then the model-function $M_{\mathcal{D}}$ of $\mathcal{D}$ is a probability distribution over $\mathcal{W} \times \mathcal{H}$, i.e. $\sum_{(W,H) \in \mathcal{W} \times \mathcal{H}} M_{\mathcal{D}}(W,H) = 1$.

Proof:

By induction on $\mathit{perfInstants}(\mathit{pnarr}(\mathcal{D}))$.

Base case: Let $\mathit{perfInstants}(\mathit{pnarr}(\mathcal{D})) = \emptyset$. Therefore, there are no p-propositions in $\mathcal{D}$, and for all $H \in \mathcal{H}$, $R(\mathcal{D}, H) = \{\mathcal{N}\}$ for an ne-domain description $\mathcal{N}$ which is equal to $\mathcal{D}$ with all the s-propositions removed. Then,

$$\sum_{(W,H)} M_{\mathcal{D}}(W,H) = \sum_{(W,H)} \tilde{M}^{\mathcal{N}}_{\mathcal{D}}(W,H) = \sum_{(W,H)} M^{ne}_{\mathcal{N}}(W) \cdot \epsilon_{\mathcal{D}}(H \mid W)$$

$$= \sum_{W} M^{ne}_{\mathcal{N}}(W) \sum_{H} \epsilon_{\mathcal{D}}(H \mid W) = \sum_{W} M^{ne}_{\mathcal{N}}(W) = 1$$

Inductive case: Let $\mathit{perfInstants}(\mathit{pnarr}(\mathcal{D})) = \{I_1, \ldots, I_n\}$ and let $\mathcal{D}'$ be the EPEC domain description $\mathcal{D}$ with all the p-propositions having instants $\geq I_n$ removed.

The inductive hypothesis is:

$$\sum_{(W,H)} M_{\mathcal{D}'}(W, H) = 1$$

Let $\mathcal{D} \backslash \mathcal{D}' = \{\underline{\mathbf{p}}_1, \ldots, \underline{\mathbf{p}}_k\}$ (possibly empty) and let $\underline{\mathbf{o}}_1, \ldots, \underline{\mathbf{o}}_k$ be the corresponding converted o-propositions. Then, for $1 \leq i \leq n$, the reduct $\mathcal{R}^{\mathcal{D}}_{H_i}$ is obtained from one of the reducts in $R(\mathcal{D}')$ by adding zero or more o-propositions from $\{\underline{\mathbf{o}}_1, \ldots, \underline{\mathbf{o}}_k\}$.

Let $R(\mathcal{D}) = \{\mathcal{R}^{\mathcal{D}}_{H_1}, \ldots, \mathcal{R}^{\mathcal{D}}_{H_n}\}$ and $R(\mathcal{D}') = \{\mathcal{R}^{\mathcal{D}'}_{H'_1}, \ldots, \mathcal{R}^{\mathcal{D}'}_{H'_m}\}$, with $n \geq m$. Let $\mathcal{R}^{\mathcal{D}}_{H_i} = \mathcal{R}^{\mathcal{D}'}_{H'_{r_i}} \cup \{\underline{\mathbf{o}}_{i_1}, \ldots, \underline{\mathbf{o}}_{r_i}\}$ for each $1 \leq i \leq n$. Viceversa, each $\mathcal{R}^{\mathcal{D}'}_{H'_k} \in R(\mathcal{D}')$ is a member of multiple reducts $\mathcal{R}^{\mathcal{D}}_{H_{k_1}}, \ldots, \mathcal{R}^{\mathcal{D}}_{H_{k_{t_k}}} \in R(\mathcal{D})$ (in the sense that $\mathcal{R}^{\mathcal{D}'}_{H'_k} \subseteq \mathcal{R}^{\mathcal{D}}_{H_{k_1}}, \ldots, \mathcal{R}^{\mathcal{D}'}_{H'_k} \subseteq \mathcal{R}^{\mathcal{D}}_{H_{k_t}}$). Then,

$$\sum_{(W,H)} M_{\mathcal{D}}(W, H) = \sum_{W} \sum_{1 \leq i \leq n} M^{ne}_{\mathcal{R}^{\mathcal{D}}_{H_i}}(W) \epsilon_{\mathcal{D}}([H_i]_{\mathcal{D}} \mid W)$$

$$= \sum_{W} \sum_{1 \leq k \leq m} M^{ne}_{\mathcal{R}^{\mathcal{D}'}_{H'_k}}(W'_W) \epsilon_{\mathcal{D}'}([H'_k]_{\mathcal{D}'} \mid W'_W) \left(\sum_{1 \leq i \leq t_k} \frac{M^{ne}_{\mathcal{R}^{\mathcal{D}}_{H_{k_i}}}(W)}{M^{ne}_{\mathcal{R}^{\mathcal{D}'}_{H'_k}}(W'_W)} \cdot \frac{\epsilon_{\mathcal{D}}([H'_{k_i}]_{\mathcal{D}} \mid W)}{\epsilon_{\mathcal{D}'}([H'_k]_{\mathcal{D}'} \mid W'_W)} \right)$$

where W'_W is the unique world well-behaved w.r.t. $R^{\mathcal{D}'}_{H'_k}$ which is indistinguishable from W up to I_n (see Corollary D.1).

Let

$$\tilde{t}_{k,i}(W, I) = \frac{\epsilon(R^{\mathcal{D}}_{H_i}, W)}{\epsilon(R^{\mathcal{D}'}_{H_{k_i}}, W'_W)} \cdot t_{R^{\mathcal{D}}_{H_i}}(W'_W(I), \tilde{S}^W_{>I})$$

where $\tilde{S}^W_{>I_n}$ is the fluent state taken by W at instants $> I$ (see Corollary D.1).

Then from Corollary D.1 we have that:

$$\frac{M^{ne}_{\mathcal{R}^{\mathcal{D}}_{H_{k_i}}}(W)}{M^{ne}_{\mathcal{R}^{\mathcal{D}'}_{H'_k}}(W'_W)} = \frac{\epsilon(R^{\mathcal{D}}_{H_i}, W)}{\epsilon(R^{\mathcal{D}'}_{H_{k_i}}, W'_W)} \cdot t_{R^{\mathcal{D}}_{H_i}}(W'_W(I_n), \tilde{S}^W_{>I_n}) = \tilde{t}_{k,i}(W, I_n)$$

In the following, for a fluent state $\tilde{S}$ we write $W^{W'}_{\tilde{S},I}$ for the world that is equal to W' up to I and has persisting fluent state $\tilde{S}$ at instants $> I$. Then note that applying Lemmas D.2 and D.3 gives:

$$\sum_{\tilde{S}} \tilde{t}_{k,i}(W^{W'}_{\tilde{S},I}, I) = 1$$

Then the above chain of equalities continues:

$$= \sum_{W' \in \mathcal{W}} \sum_{1 \leq k \leq m} M^{ne}_{\mathcal{R}^{\mathcal{D}'}_{H'_k}}(W') \epsilon_{\mathcal{D}'}([H'_k]_{\mathcal{D}'} \mid W') \left(\sum_{1 \leq i \leq t_k} \sum_{\tilde{S}} \tilde{t}_{k,i}(W^{W'}_{\tilde{S},I_n}, I_n) \cdot \frac{\epsilon_{\mathcal{D}}([H'_{k_i}]_{\mathcal{D}} \mid W^{W'}_{\tilde{S},I_n}, I_n)}{\epsilon_{\mathcal{D}'}([H'_k]_{\mathcal{D}'} \mid W')} \right)$$

and note that for all $1 \leq i \leq t_k$

$$\epsilon_{\mathcal{D}}([H_{k_i}]_{\mathcal{D}} \mid W^{W'}_{\tilde{S},I_n}, I_n) = \epsilon_{\mathcal{D}'}([H'_k]_{\mathcal{D}'} \mid W') \cdot c_{k,i}$$

for some appropriate constant $c_{k,i}$ (according to the outcomes being possibly sensed in H_{k_i} at instants I such that $I_{n-1} \leq I < I_n$), and for any fixed k, note that $\sum_{1 \leq i \leq t_k} c_{k,i} = 1$ as $W^{W'}_{\tilde{S},I_n}$ is sensed in every possible way.

However, the constant $c_{k,i}$ does not depend on the fluent state $\tilde{S}$, since the information according to which the two reducts $\mathcal{R}^{\mathcal{D}}_{H_{k_i}}$ and $\mathcal{R}^{\mathcal{D}}_{H_{k_j}}$ (for $i \neq j$) are obtained must have been sensed at some instants $\geq I_{n-1}$ and $< I_n$. Therefore,

$$= \sum_{W'} \sum_{1 \leq k \leq m} M^{ne}_{\mathcal{R}^{\mathcal{D}'}_{H'_k}}(W') \epsilon_{\mathcal{D}'}([H'_k] \mid W') \left(\sum_{1 \leq i \leq t_k} c_{k,i} \sum_{\tilde{S}} \tilde{t}_{k,i}(W^{W'}_{\tilde{S},I_n}, I_n) \right)$$

$$= \sum_{W'} \sum_{1 \leq k \leq m} M^{ne}_{\mathcal{R}^{\mathcal{D}'}_{H'_k}}(W') \epsilon_{\mathcal{D}'}([H'_k]_{\mathcal{D}'} \mid W') \left(\sum_{1 \leq i \leq t_k} c_{k,i} \right)$$

and finally

$$= \sum_{W'} \sum_{1 \leq k \leq m} M^{ne}_{\mathcal{R}^{\mathcal{D}'}_{H'_k}}(W') \epsilon_{\mathcal{D}'}([H'_k]_{\mathcal{D}'} \mid W') = \sum_{(W,H)} M_{\mathcal{D}'}(W, H) = 1$$

via the inductive hypothesis. ■

D.5. Proof of Proposition 4.1

The statement of the proposition is:

> **Proposition 4.1** (PAL/EPEC Correspondence)**.** Let $\mathcal{P}$ be the normalised, consistent PAL theory $PAL_D \cup PAL_P \cup \{$"**initially** $\dot{\psi}_0$"$\}$, let $\mathcal{Q}$ be the PAL query "**probability of** $[\dot{\varphi}_q$ **after** $\dot{a}_1, \ldots, \dot{a}_k]$ **is** n", let $\mathcal{N}(\mathcal{P},\mathcal{Q})$ be the EPEC encoding of $\mathcal{P}$ and $\mathcal{Q}$, and let α be the EPEC i-formula $[\dot{a}_1]@0 \wedge \ldots \wedge [\dot{a}_k]@k-1$. Then $\mathcal{P} \models_{PAL} \mathcal{Q}$ if and only if
>
> $$\left(\mathcal{N}(\mathcal{P},\mathcal{Q}) \,\middle|\, [\dot{\psi}_0]@0\right) \;|\!\!\models\; \alpha \wedge [\dot{\varphi}_q]@k \;\textbf{holds-with-prob}\; n$$

Proof:

Let $\mathcal{Q}^p =$ "$\dot{\varphi}_q$ **after** $\dot{a}_1, \ldots, \dot{a}_k$" be the multi-action plan contained in the query $\mathcal{Q}$, let $\mathcal{O} =$ "**initially** $\dot{\psi}_0$" be the single hypothetical observation in $\mathcal{P}$, and let $\mathcal{N} = \mathcal{N}(\mathcal{P},\mathcal{Q})$. Then by Definition 3 in [6] and Definition 3.28, it is sufficient to show that $\dot{P}(\mathcal{Q}^p \mid \mathcal{O}) = M^{ne}_{\mathcal{N}}(\alpha \wedge [\dot{\varphi}_q]@k \mid [\dot{\psi}_0]@0)$. Using definitions 1–3 and equations (0.7)–(0.16) in [6] to expand $\dot{P}(\mathcal{Q}^p \mid \mathcal{O})$ gives:

$$
\begin{aligned}
\dot{P}(\mathcal{Q}^p \mid \mathcal{O}) &= \sum_{\dot{s}} \dot{P}(\dot{s} \mid \mathcal{O})\dot{P}(\mathcal{Q}^p \mid \dot{s}) && \text{[by (0.16)]}\\
&= \sum_{\dot{s}} \dot{P}(\dot{s})\dot{P}(\mathcal{O} \mid \dot{s})\dot{P}(\mathcal{Q}^p \mid \dot{s}) \Big/ \sum_{\dot{s}} \dot{P}(\dot{s})\dot{P}(\mathcal{O} \mid \dot{s}) && \text{[by (0.15)]}\\
&= \sum_{\substack{\dot{s} \models \dot{\psi}_0 \\ \dot{s}^{(i)} \in \dot{\Phi}(\dot{a}_i, \dot{s}^{(i-1)}) \\ \dot{s}^{(k)} \models \dot{\varphi}_q}} \dot{P}(\dot{s})\dot{P}_{\dot{a}_1}(\dot{s}^{(1)} \mid \dot{s}) \cdots \dot{P}_{\dot{a}_k}(\dot{s}^{(k)} \mid \dot{s}^{(k-1)}) \Big/ \sum_{\dot{s} \models \dot{\psi}_0} \dot{P}(\dot{s}) && \text{[by (0.13)]}\\
&= \sum_{\substack{\dot{s} \models \dot{\psi}_0 \\ \dot{s}^{(i)} \in \dot{\Phi}(\dot{a}_i, \dot{s}^{(i-1)}) \\ \dot{s}^{(k)} \models \dot{\varphi}_q}} \frac{\dot{P}(\dot{s}_u)}{|\dot{I}(\dot{s}_u)|} \left(\prod_{i=1}^{k} \frac{2^{|\dot{\mathrm{U}}_N|}}{\left|\dot{\Phi}(\dot{a}_i, \dot{s}^{(i-1)})\right|} \cdot \dot{P}\left(\dot{s}^{(i)}_N\right) \right) \Big/ \sum_{\dot{s} \models \dot{\psi}_0} \frac{\dot{P}(\dot{s}_u)}{|\dot{I}(\dot{s}_u)|} && \text{[by (0.7),(0.9)]} \;\; \textbf{(A)}
\end{aligned}
$$

Expanding $M^{ne}_{\mathcal{N}}(\alpha \wedge [\dot{\varphi}_q]@k \mid [\dot{\psi}_0]@0)$ using the definitions in Section 3.3 and

Appendix D Section D.1 gives:

$$
\begin{aligned}
M^{ne}_{\mathcal{N}}\left(\alpha \wedge [\dot{\varphi}]@k \,\middle|\, [\dot{\psi}]@0\right) &= \frac{M^{ne}_{\mathcal{N}}(\alpha \wedge [\dot{\varphi}]@k \wedge [\dot{\psi}]@0)}{M^{ne}_{\mathcal{N}}([\dot{\psi}]@0)} \\
&= \frac{M^{ne}_{\mathcal{N}}(\alpha \wedge [\dot{\varphi}]@k \wedge [\dot{\psi}]@0)}{M^{ne}_{\mathcal{N}_{<0}}([\dot{\psi}]@0)} = \frac{\sum_{W \mid\models \alpha \wedge [\dot{\varphi}]@k \wedge [\dot{\psi}]@0} M^{ne}_{\mathcal{N}}(W)}{\sum_{W \mid\models [\dot{\psi}]@0} M^{ne}_{\mathcal{N}_{<0}}(W)} \\
&= \frac{\sum_{W \mid\models \alpha \wedge [\dot{\varphi}]@k \wedge [\dot{\psi}]@0} M^{ne}_{\mathcal{N}_{<0}}(W^{(1)}) \left(\prod_{i=1}^{k} t_{\mathcal{N}}\left(W(i-1), \tilde{S}^{W^{(i+1)}}_{>i-1}\right)\right)}{\sum_{W \mid\models [\psi]@0} M^{ne}_{\mathcal{N}_{<0}}(W)} \qquad \textbf{(B)}
\end{aligned}
$$

where for a given W, for each i, $W^{(i+1)}$ is the unique world that is well-behaved w.r.t. $\mathcal{N}_{<i}$ and indistinguishable from W up to instant i.

That expressions **(A)** and **(B)** are equal follows because:

- $M^{ne}_{\mathcal{N}_{<0}}(W^{(1)})$ evaluates to $\frac{\dot{P}(\dot{s}_u)}{|\dot{I}(\dot{s}_u)|}$ when the fluent state $\tilde{W}(0) = \dot{s}$.
- The product of transition functions $\left(\prod_{i=1}^{k} t_{\mathcal{N}}\left(W(i-1), \tilde{S}^{W^{(i+1)}}_{>i-1}\right)\right)$ evaluates to

$$
\left(\prod_{i=1}^{k} \frac{2^{|\dot{\mathrm{U}}_N|}}{\left|\dot{\Phi}(\dot{a}_i, \dot{s}^{(i-1)})\right|} \cdot \dot{P}\left(\dot{s}^{(i)}_N\right)\right)
$$

  when the fluent state and action in $W(i-1)$ are set to $\dot{s}^{(i-1)}$ and $\dot{a}_i$ respectively, and the fluent state $\tilde{S}^{W^{(i+1)}}_{>i-1} = \dot{s}^{(i)}$.

- The sum in the numerator of **(B)** evaluates to 0 when any one of $\dot{a}_1, \ldots, \dot{a}_k$ is non-executable, as satisfying the corresponding non-executability precondition would render the CWA unsatisfied, so that the set of well-behaved worlds such that $W \mid\models \alpha \wedge [\dot{\phi}]@k \wedge [\dot{\psi}]@0$ would be empty.

■

D.6. Proof of Proposition 4.2

The statement of the proposition is:

Proposition 4.2 (BHL/EPEC Correspondence)**.** Let $\mathcal{T}$ be a normalised extended action theory and let α be an executable sequence of n agent-performable actions. Let $\Delta = \{\alpha_1, \ldots, \alpha_m\}$ be the maximal set of distinct executable action sequences such that for each $\alpha_i \in \Delta$, $\textsc{Prob}(\alpha_i, S_0) > 0$ and α_i has been obtained from α by replacing each noisy sensing action, say $\pi x, y.\textsf{sense-f}(x, y)$, with $\pi y.\textsf{sense-f}(v, y)$ for some sensed value v for the nominal variable x. Then

$$\text{for all } 1 \leq i \leq m,\ \mathcal{T} \models Do(\alpha_i, S_0, s) \rightarrow \textsc{Bel}(\phi[s_{know}], s) = P_i$$

$$\text{if and only if}$$

$$\mathcal{D}(\mathcal{T}, \alpha) \mid\models \textbf{at } n \textbf{ believes } [\phi]@n \textbf{ with-probs } P$$

where $P = \{([H_1], B_1, P_1), \ldots, ([H_m], B_m, P_m)\}$ is such that for each $1 \leq i \leq m$ $T([H_i]) = \alpha_i$ and $B_i = \textsc{Prob}(\alpha_i, S_0)$.

Proof:

For brevity, $\mathcal{D}$ in this proof is shorthand for $\mathcal{D}(\mathcal{T}, \alpha)$. We prove the proposition by showing the equivalence for an arbitrary i, $1 \leq i \leq m$, between $\textsc{Bel}(\phi[s_{know}], s)$ and $M_{\mathcal{D}}([\phi]@n \mid H_i)$, for an arbitrary s such that $Do(\alpha_i, S_0, s)$.

Let $s = do([a_1; \ldots; a_n], s'_0)$ for primitive actions $a_1, \ldots, a_n$ and some initial situation s'_0 such that $p(s'_0, S_0) > 0$. By definition, $\textsc{Bel}(\phi[s_{know}], s)$ is shorthand for

$$\frac{\sum_{s':\phi[s_{know}/s']} p(s', s)}{\sum_{s'} p(s', s)}$$

which, in turn, using the successor-state axiom for p is equal to

$$\frac{\sum_{s'=do(a'_n, s'_{n-1}):\phi[s_{know}/s']} p(s'_0, S_0) \cdot \prod_{i=1}^{n} \ell(a'_i, s'_{i-1})}{\sum_{s'=do(a'_n, s'_{n-1})} p(s'_0, S_0) \cdot \prod_{i=1}^{n} \ell(a'_i, s'_{i-1})} \tag{21}$$

where the $s'_1, \ldots, s'_{n-1}, a'_1, \ldots, a'_n$ are situations and actions satisfying the conditions in the **if** part of the successor-state axiom for p, i.e. $s'_1 = do(a'_1, s'_0)$, $\textsc{Oi}(a_1, a'_1, s'_0)$, $s'_2 = do(a'_2, s'_1)$, $\textsc{Oi}(a_2, a'_2, s'_1)$ and so on up to n.

Notice that $\mathcal{D}$ does not contain any p-propositions with epistemic preconditions. Therefore it only has one reduct independently of any specific sensing history, say $\mathcal{R}$. By definition, the model function (for well-behaved h-worlds) expands to:

$$M_{\mathcal{D}}(W, H_i) = \tilde{M}_{\mathcal{D}}^{\mathcal{R}_{\mathcal{D}}}(W, H_i) = \epsilon_{\mathcal{D}}(H_i \mid W) M_{\mathcal{R}}^{ne}(W)$$

We are interested in $M_{\mathcal{D}}([\phi]@n \mid H_i)$. This evaluates to:

$$M_{\mathcal{D}}([\phi]@n \mid H_i) = \frac{\sum_{W \mid \models [\phi]@n} \epsilon_{\mathcal{D}}(H_i \mid W) M_{\mathcal{R}}^{ne}(W)}{\sum_{W \in \mathcal{W}} \epsilon_{\mathcal{D}}(H_i \mid W) M_{\mathcal{R}}^{ne}(W)} \tag{22}$$

Recall from the proof of Proposition 4.1 in Appendix D.5 that $M_{\mathcal{R}}^{ne}(W)$ can be written as

$$M_{\mathcal{R}_{<0}}^{ne}(W^{(1)}) \left(\prod_{i=1}^{k} t_{\mathcal{R}} \left(W(i-1), \tilde{S}_{>i-1}^{W^{(i+1)}} \right) \right)$$

So (22) evaluates to:

$$\frac{\sum_{W \mid \models [\phi]@n} \epsilon_{\mathcal{D}}(H_i \mid W) M_{\mathcal{R}_{<0}}^{ne}(W^{(1)}) \left(\prod_{i=1}^{k} t_{\mathcal{R}} \left(W(i-1), \tilde{S}_{>i-1}^{W^{(i+1)}} \right) \right)}{\sum_{W \in \mathcal{W}} \epsilon_{\mathcal{D}}(H_i \mid W) M_{\mathcal{R}_{<0}}^{ne}(W^{(1)}) \left(\prod_{i=1}^{k} t_{\mathcal{R}} \left(W(i-1), \tilde{S}_{>i-1}^{W^{(i+1)}} \right) \right)} \tag{23}$$

That (21) and (23) are equal then follows since:

- The value of $p(s'_0, S_0)$ equals one of $p_1, \ldots, p_n$ as they appear in the translation to an EPEC i-proposition on page 58, depending on what is true in the possible situation s'_0. By definition of an i-proposition there exists some W such that this also equals $M_{\mathcal{R}_{<0}}^{ne}(W^{(1)})$. (Again using the notation from Appendix D.5.)

- If a_i is not a sensing action, then $\ell(a'_i, s'_{i-1})$ is obtained by applying one of the functions $p_1(\cdot,\cdot), \ldots, p_{n+1}(\cdot,\cdot)$ as they appear in the action-likelihood axiom on page 59 to a'_i. a_i has effects as specified in its effect axioms, whose preconditions are relative to s'_{i-1}. By translation there is a W such that (i) $W(i-1)$ contains the agent-performable action corresponding to a_i together with the state representation of s'_{i-1}, and (ii) the result of a_i is captured in $\tilde{S}_{>i-1}^{W^{(i+1)}}$. This is reflected in the term $t_{\mathcal{R}} \left(W(i-1), \tilde{S}_{>i-1}^{W^{(i+1)}} \right)$

in (23).

- If a_i is a sensing action, then $\ell(a'_i, s'_{i-1})$ is taken into account within $\epsilon_{\mathcal{D}}(H_i \mid W)$, by unfolding Definition 3.34 and considering the entry corresponding to a_i in the matrix resulting from its translation.

■