arXiv:cs/9910016v1 [cs.AI] 21 Oct 1999

Probabilistic Agent Programs*

Jürgen Dix†
University of Koblenz, Dept. of Computer Science
D-56075 Koblenz, Germany

Mirco Nanni
University of Pisa, Dept. of Computer Science
I-56125 Pisa, Italy

V.S. Subrahmanian‡
Dept. of CS, University of Maryland
College Park, MD 20752, USA

20th June 2018

Abstract

Agents are small programs that autonomously take actions based on changes in their environment or "state." Over the last few years, there have been an increasing number of efforts to build agents that can interact and/or collaborate with other agents. In one of these efforts, Eiter, Subrahmanian, and Pick (1999) have shown how agents may be built on top of legacy code. However, their framework assumes that agent states are completely determined, and there is no uncertainty in an agent's state. Thus, their framework allows an agent developer to specify how his agents will react when the agent is 100% sure about what is true/false in the world state. In this paper, we propose the concept of a *probabilistic agent program* and show how, given an arbitrary program written in any imperative language, we may build a declarative "probabilistic" agent program on top of it which supports decision making in the presence of uncertainty. We provide two alternative semantics for probabilistic agent programs. We show that the second semantics, though more epistemically appealing, is more complex to compute. We provide sound and complete algorithms to compute the semantics of *positive* agent programs.

*Most proofs are contained in the appendix.

†This work was carried out when the author was visiting the University of Maryland from January-October 1999.

‡This work was supported by the Army Research Office under Grants DAAH-04-95-10174, DAAH-04-96-10297, DAAG-55-97-10047 and DAAH04-96-1-0398, by the Army Research Laboratory under contract number DAAL01-97-K0135 and by an NSF Young Investigator award IRI-93-57756.

1 Introduction

Over the last few years, there has been increasing interest in the area of software agents. Such agents provide a wide variety of services including identification of interesting newspaper articles, software robots that perform tasks (and plan) on a user's behalf, content based routers, agent based telecommunication applications, and solutions to logistics applications. *IMPACT* (see |http://www.cs.umd.edu/projects/impact/—) is a multinational project whose aim is to define a formal theory of software agents, implement (appropriate fragments of) the theory efficiently, and develop an appropriate suite of applications on top of this implementation. An *IMPACT* agent manages a set of data types/structures (including a message box) through a set of application program interface (API) function calls. The state of the agent at a given point in time is a set of objects belonging to these data types. Each agent has a set of integrity constraints that its state must always satisfy. When an agent's state changes (due to external events such as receipt of a message), the agent tries to modify its state so that the integrity constraints are satisfied. To do this, it has a suite of actions, and an *agent program* that specifies the operating principles (what is permitted, what is forbidden, what is obligatory, etc., and under what conditions?). (Eiter, Subrahmanian, and Pick 1999; Eiter and Subrahmanian 1999) provides a detailed study of the semantics and complexity of such agents, (Eiter, Subrahmanian, and Rogers 1999) contains compile-time and run-time algorithms, while (Arisha, Ozcan, Ross, Subrahmanian, Eiter, and Kraus 1999) focuses on system architecture.

Past work on *IMPACT* assumes that all agents reason with a complete and certain view of the world. However, in many real world applications, agents have only a partial, uncertain view of what is true in the world. Though an agent may need to reason about uncertainty for many reasons, in this paper, we will assume that the main cause of uncertainty in an agent is due to its state being uncertain. For example, when an image processing agent is asked to identify an enemy vehicle, it might return the fact that vehicle v_1 is a T72 tank (with 60–70% probability) and a T-80 tank (with 20-45% probability). However, this raises several problems, the first of which is that as an action can only be executed if its precondition is true in the current state, if the agent doesn't know what the state is, then it cannot determine which of its actions are executable, and which are not. Second, even if an action is executable, the state that results may not be precisely determinable either. One consequence of all this is that the semantics of agent programs change significantly when such uncertainties arise.

The main contributions (and organization) of this paper may now be summed up as follows.

1. In Section 2, we present a brief overview of agents (without any uncertainty involved) as described in (Eiter, Subrahmanian, and Pick 1999).

2. Then, in Section 3, we define the concept of a probabilistic code call, which is the basic syntactic construct through which uncertainty in abstract data

types manifests itself.

3. In Section 4, we define the syntax of probabilistic agent programs. Specifically, we show that probabilistic agent programs allow an agent developer to specify the permissions, obligations, forbidden actions, etc. associated with an agent depending not only on the probabilities that certain conditions hold in the agent's state, but also on the developer's assumptions about the relationship between these conditions (e.g. the probability that a conjunction holds in a given state depends not only on the probabilities of the conjuncts involved, but also on the dependencies if any between the conjuncts).

4. In Section 5, we develop three formal semantics for probabilistic agent programs which extend each other as well as the semantics for (ordinary, non probabilistic) agent programs defined by Eiter, Subrahmanian, and Pick (1999). We also provide results relating these diverse semantics.

5. Then, in Section 6, we develop a sound and complete algorithm to compute the semantics defined when only positive agent programs are considered. We also show that the classical agent programs of Eiter, Subrahmanian, and Pick (1999) are a special case of our probabilistic programs.

6. In Section 7, we provide an alternative, Kripke style semantics for agent programs. In contrast to the previous "family" of semantics which assume that an agent's precondition must be true with 100% probability for the agent to execute it, this semantics also allows an agent to execute it when it is not sure (with 100% probability) that the action's precondition is true. We extend all three semantics of agent programs defined earlier in Section 5 to handle these intuitions. Unfortunately, as we show in this section, this desire for a "more sophisticated" sematics comes at a high computational price.

2 Preliminaries

In *IMPACT*, each agent $\mathfrak{a}$ is built on top of a body of software code (built in any programming language) that supports a well defined application programmer interface (either part of the code itself, or developed to augment the code). Hence, associated with each agent $\mathfrak{a}$ is a body of software code $\mathcal{S}^{\mathfrak{a}}$ defined as follows.

Definition 2.1 (Software Code) *We may characterize the code on top of which an agent is built as a triple* $\mathcal{S} =_{def} (\mathcal{T}_{\mathcal{S}}, \mathcal{F}_{\mathcal{S}}, \mathcal{C}_{\mathcal{S}})$ *where:*

1. $\mathcal{T}_{\mathcal{S}}$ *is the set of all data types managed by* $\mathcal{S}$*,*

2. $\mathcal{F}_{\mathcal{S}}$ *is a set of predefined functions which makes access to the data objects managed by the agent available to external processes, and*

3. *$\mathcal{C}_{\mathcal{S}}$ is a set of type composition operations. A type composition operator is a partial n-ary function c which takes types $\tau_1, \ldots, \tau_n$ as input, and yields a type $c(\tau_1, \ldots, \tau_n)$ as output. As c is a partial function, c may only be defined for certain arguments $\tau_1, \ldots, \tau_n$, i.e., c is not necessarily applicable on arbitrary types.*

When $\mathfrak{a}$ is clear from context, we will often drop the superscript $\mathfrak{a}$. Intuitively, $\mathcal{T}_{\mathcal{S}}$ is the set of all data types managed by $\mathfrak{a}$, $\mathcal{F}_{\mathcal{S}}$ is the set of all function calls supported by $\mathcal{S}$'s application programmer interface (*API*). $\mathcal{C}_{\mathcal{S}}$ is the set of ways of creating new data types from existing data types. This characterization of a piece of software code is widely used (cf. the *Object Data Management Group*'s *ODMG* standard (Cattell, R. G. G., et al. 1997) and the *CORBA* framework (Siegal 1996)). Each agent also has a message box having a well defined set of associated code calls that can be invoked by external programs.

Example 2.2 [Surveillance Example] Consider a surveillance application where there are hundreds of (identical) surveillance agents, and a geographic agent. The data types associated with the surveillance and geographic agent include the standard `int`,`bool`,`real`,`string`,`file` data types, plus those shown below:

Surveillance Agent	**Geographic Agent**
image:record of imageid:file; day:date; time:int; location:string imagedb: *setof* image;	map:↑ quadtree; **quadtree:record of** place:string; xcoord:int; ycoord:int; pop:int nw,ne,sw,se:↑ quadtree

A third agent may well merge information from these two agents, tracking a sequence of surveillance events.

The surv agent may support a function surv : *identify*() which takes as input, an image, and returns as output, the set of all identified vehicles in it. It may also support a function called surv : *turret*() that takes as input, a vehicle id, and returns as output, the type of gun-turret it has. Likewise, the geo agent may support a function geo : *getplnode*() which takes as input a map and the name of a place and returns the set of all nodes with that name as the place-field, a function geo : *getxynode*() which takes as input a map and the coordinates of a place and returns the set of all nodes with that coordinate as the node, a function called geo : *range*() that takes as input a map, an x, y coordinate pair, and a distance r and returns as output, the set of all nodes in the map (quadtree) that are within r units of location (x, y).

Throughout this paper, we will expand on this simple example and use it to illustrate and motivate the various definitions in the paper.

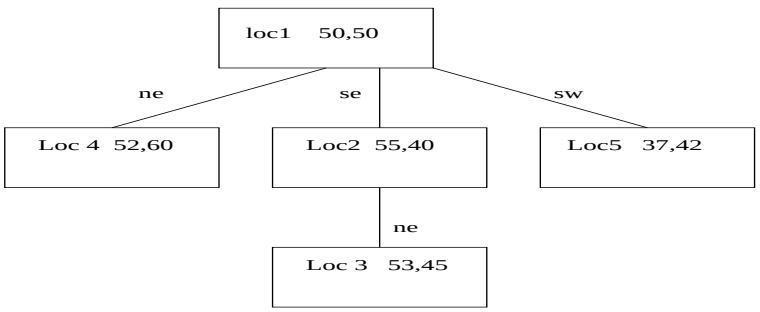

Figure 1: Example quadtree for Surveillance Application.

Definition 2.3 (State of an Agent) *The state of an agent at any given point t in time, denoted $\mathcal{O}_{\mathcal{S}}(t)$, consists of the set of all instantiated data objects of types contained in $\mathcal{T}_{\mathcal{S}}^{\mathtt{a}}$.*

An agent's state may change because it took an action, or because it received a message. Throughout this paper we will assume that except for appending messages to an agent $\mathtt{a}$'s mailbox, another agent $\mathtt{b}$ cannot directly change $\mathtt{a}$'s state. However, it might do so indirectly by shipping the other agent a message requesting a change.

Example 2.4 For example, the state of the Geographic Agent may consist of two quadtrees (one of which, map1, is shown in Figure 1), and the type "map" may contain two objects, map1, and map2, pointing to these two quadtrees, respectively. (The figure doesn't show population values explicitly. Assume the population values are 20,000 for Loc1, 28,000 for Loc2, 15,000 for Loc3, and 40,000 for Loc4, and 8000 for Loc5.)

Queries and/or conditions may be evaluated w.r.t. an agent state using the notion of a code call atom and a code call condition defined below.

Definition 2.5 (Code Call/Code Call Atom) *If $\mathcal{S}$ is the name of a software package, f is a function defined in this package, and $(d_1, \ldots, d_n)$ is a tuple of arguments of the input type of f, then $\mathcal{S}:f(\mathtt{d_1}, \ldots, \mathtt{d_n})$ is called a* code call.

If `cc` *is a code call, and* `X` *is either a variable symbol or an object of the output type of* `cc`*, then* $\mathbf{in}(\mathtt{X}, \mathtt{cc})$ *is called a* code call atom.

For instance, in the Surveillance example, $\mathtt{geo}:\mathit{getplnode}(\mathtt{map1}, \mathtt{"Loc1"})$ returns the set containing just the single node referring to Loc1 in Figure 1. Likewise, the code call $\mathtt{geo}:\mathit{range}(\mathtt{map1}, 55, 50, 11)$ returns the set containing the nodes labeled Loc1 and Loc2.

Definition 2.6 (Code Call Condition) *A* code call condition χ *is defined as follows:*

1. *Every code call atom is a code call condition.*
2. *If* $\mathtt{s}, \mathtt{t}$ *are either variables or objects, then* $\mathtt{s} = \mathtt{t}$ *is a code call condition.*
3. *If* $\mathtt{s}, \mathtt{t}$ *are either integers/real valued objects, or are variables over the integers/reals, then* $\mathtt{s} < \mathtt{t}, \mathtt{s} > \mathtt{t}, \mathtt{s} \geq \mathtt{t}, \mathtt{s} \leq \mathtt{t}$ *are code call conditions.*
4. *If* χ_1, χ_2 *are code call conditions, then* $\chi_1 \,\&\, \chi_2$ *is a code call condition.*

A code call condition satisfying any of the first three criteria above is an atomic *code call condition.*

An example code call condition is shown below.

Example 2.7 $\mathbf{in}(\mathtt{X}, \mathtt{geo} : \mathit{range}(\mathtt{map1}, \mathtt{55}, \mathtt{50}, \mathtt{11})) \,\&\, X.pop > 25,000$ is a code call condition that is satisfied by only one node in map1, viz. the Loc2 node.

Each agent has an associated set of **integrity constraints**—only states that satisfy these constraints are considered to be *valid* or *legal* states. An integrity constraint is an implication whose consequent is a code call atom, and whose antecedent is a code call condition. Appendix A contains a detailed definition.

Each agent has an action-base describing various actions that the agent is capable of executing. Actions change the state of the agent and perhaps the state of other agents' `msgboxes`. As in classical AI, all actions have an associated precondition (a code call condition that the agent state must satisfy for the action to be executable) and an add/delete list. Appendix A contains detailed definitions from (Eiter, Subrahmanian, and Pick 1999).

For instance, the `geo` agent may have an *insert* action that adds a node to the map. Likewise, the `surv` agent may also have an *insert* action which inserts a new image into the image database. Both these agents also have an action that sends a message.

Each agent has an associated "notion of concurrency," **conc**, which a set of actions and an agent state as input, and produces as output, a single action that reflects the combination of all the input actions. (Eiter, Subrahmanian, and Pick 1999) provides examples of three different notions of concurrency. We will sometimes abuse notation write $\mathbf{conc}(S, \mathcal{O})$ to denote the new state obtained by concurrently executing the actions in S in state $\mathcal{O}$.

Each agent has an associated set of *action constraints* that define the circumstances under which certain actions may be concurrently executed. As at any given point t in time, many sets of actions may be concurrently executable, each agent has an *Agent Program* that determines what actions the agent can take, what actions the agent cannot take, and what actions the agent must take. Agent programs are defined in terms of status atoms defined below.

Definition 2.8 (Status Atom/Status Set) *If $\alpha(\vec{t})$ is an action, and $Op \in \{\mathbf{P}, \mathbf{F}, \mathbf{W}, \mathbf{Do}, \mathbf{O}\}$, then $Op\alpha(\vec{t})$ is called a* status atom. *If A is an action status atom, then $A, \neg A$ are called* status literals. *A* status set *is a finite set of ground status atoms.*

Intuitively, $\mathbf{P}\alpha$ means α is permitted, $\mathbf{F}\alpha$ means α is forbidden, $\mathbf{Do}\,\alpha$ means α is actually done, and $\mathbf{W}\alpha$ means that the obligation to perform α is waived.

Definition 2.9 (Agent Program) *An agent program $\mathcal{P}$ is a finite set of rules of the form*

$$A \quad \leftarrow \quad \chi \,\&\, L_1 \,\&\, \ldots \,\&\, L_n$$

where χ is a code call condition and $L_1, \ldots, L_n$ are status literals.

The semantics of agent programs are well described in (Eiter, Subrahmanian, and Pick 1999; Eiter and Subrahmanian 1999)—due to space reasons, we do not explicitly recapitulate them here, though Appendix A contains a brief overview of the semantics.

3 Probabilistic Code Calls

Consider a code call of the form $\mathtt{d}:f(\mathtt{args})$. This code call returns a set of objects. If an object o is returned by such a code call, then this means that o is *definitely* in the result of evaluating $\mathtt{d}:f(\mathtt{args})$. However, there are many cases, particularly in applications involving reasoning about knowledge, where a code call may need to return an "uncertain" answer. In our our surveillance example, $\mathtt{surv}:\mathit{identify}(\mathtt{image1})$ tries to identify all objects in a given image—however, it is well known that image identification is an uncertain task. Some objects may be identified with 100% certainty, while in other cases, it may only be possible to say it is either a T72 tank with 40–50% probability, or a T80 tank with (50-60%) probability.

Definition 3.1 (Random Variable of Type τ) *A* random variable *of type τ is a finite set* **RV** *of objects of type τ, together with a probability distribution $\wp$ that assigns real numbers in the unit interval $[0, 1]$ to members of* **RV** *such that $\Sigma_{o \in \mathbf{RV}} \wp(o) \leq 1$.*

It is important to note that in classical probability theory (Ross 1997), random variables satisfy the stronger requirement that $\Sigma_{o \in \mathbf{RV}} \wp(o) = 1$. However, in many real life situations, a probability distribution may have missing pieces, which explains why we have chosen a weaker definition. However, the classical probability case when $\Sigma_{o \in \mathbf{RV}} \wp(o) = 1$ is an instance of our more general definition.

Definition 3.2 (Probabilistic Code Call $\mathtt{a}:_{\mathbf{RV}} f(\mathtt{d}_1, \ldots, \mathtt{d}_n)$)
Suppose $\mathtt{a}:f(\mathtt{d}_1, \ldots, \mathtt{d}_n)$ is a code call whose output type is τ. The probabilistic

code call *associated with* $\mathtt{a}:f(\mathtt{d}_1,\ldots,\mathtt{d}_n)$*, denoted* $\mathtt{a}:_{\mathbf{RV}} f(\mathtt{d}_1,\ldots,\mathtt{d}_n)$*, returns a set of random variables of type* τ *when executed on state* $\mathcal{O}$*.*

The following example illustrates the use of probabilistic code calls.

Example 3.3 Consider the code call $\mathsf{surv}:\mathit{identify}(\mathtt{image1})$. This code call may return the following two random variables.

$$\langle\{t72, t80\}, \{\langle t72, 0.5\rangle, \langle t80, 0.4\rangle\}\rangle \text{ and } \langle\{t60, t84\}, \{\langle t60, 0.3\rangle, \langle t84, 0.7\rangle\}\rangle$$

This says that the image processing algorithm has identified two objects in image1. The first object is either a T72 or a T80 tank with 50% and 40% probability, respectively, while the second object is either a T60 or a T84 tank with 30% and 70% probability respectively.

Probabilistic code calls and code call conditions look exactly like ordinary code calls and code call conditions — however, as a probabilistic code call returns a set of *random variables*, probabilistic code call atoms are true or false with some probability.

Example 3.4 Consider the probabilistic code call condition

$$\mathbf{in}(\mathtt{X}, \mathsf{surv}:_{\mathbf{RV}} \mathit{identify}(\mathtt{image1})) \,\&\, \mathbf{in}(\mathtt{A1}, \mathsf{surv}:_{\mathbf{RV}} \mathit{turret}(\mathtt{X})).$$

This code call condition attempts to find all vehicles in "image1" with a gun turret of type A1. Let us suppose that the first code call returns just one random variable specifying that image1 contains one vehicle which is either a T72 (probability 50%) or a T80 tank (probability 40%). When this random variable (X) is passed to the second code call, it returns one random variable with two values—A1 with probability 30% and A2 with probability 65%. What is the probability that the code call condition above is satisfied by a particular assignment to X?

The answer to this question depends very much upon the knowledge we have (if any) about the dependencies between the identification of a tank as a T-72 or a T-80, and the type of gun turret on these. For instance, if we know that all T72's have A2 type turrets, then the probability of the conjunct being true when X is a T72 tank is 0. On the other hand, it may be that the turret identification and the vehicle identification are independent for T80s—hence, when X is set to T80, the probability of the conjunct being true is $0.4 \times 0.3 = 0.12$.

Unfortunately, this is not the only problem. Other problems also arise, as shown in the following example.

Example 3.5 Suppose we consider a code call χ returning the following two random variables.

$$\begin{aligned} \mathbf{RV}_1 &= \langle\{a, b\}, \wp_1\rangle \\ \mathbf{RV}_2 &= \langle\{b, c\}, \wp_2\rangle \end{aligned}$$

Suppose $\wp_1(a) = 0.9, \wp_1(b) = 0.1, \wp_2(b) = 0.8, \wp_2(c) = 0.1$. What is the probability that b is in the result of the code call χ?

Answering this question is problematic. The reason is that we are told that there are at most two objects returned by χ. One of these objects is either a or b, and the other is either b or c. This leads to four possibilities, depending on which of these is true. The situation is further complicated because in some cases, knowing that the first object is b may preclude the second object from being b—this would occur, for instance, if χ examines photographs each containing two different people and provides identifications for each. a, b and c may be potential id's of such people returned by the image processing program. In such cases, the same person can never be pictured with himself or herself.

Of course, in other cases, there may be no reason to believe that knowing the value of one of two objects tells us anything about the value of the second object. For example if we replace people with colored cubes (with a denoting amber cubes, b black, and c cyan), there is no reason to believe that two identical black cubes cannot be pictured next to each other.

One could argue, however, that the above reasoning is incorrect because if two objects are completely identical, then they must be the same. This means that if we have two distinct black cubes, then these two black cubes must be *distinguishable* from one another via some property such as their location in the photo, or their *Id*s. This is Leibniz's well known *extensionality principle*. Hence, we will require the results of a probabilistic code call to be *coherent* in the following sense.

Definition 3.6 (Coherent Probabilistic Code Call) *A probabilistic code call is* coherent *iff for all distinct* $\langle X_1, \wp_1\rangle, \langle X_2, \wp_2\rangle$, $X_1 \cap X_2 = \emptyset$.

Throughout this paper, only coherent probabilistic code calls are considered. Thus, the expression "probabilistic code call" assumes coherence.

Definition 3.7 (Satisfying a Code Call Atom) *Suppose* $\mathtt{a}:_{\mathbf{RV}} f(\mathtt{d_1}, \ldots, \mathtt{d_n})$ *is a ground probabilistic code call and* o *is an object of the output type of this code call w.r.t. agent state* $\mathcal{O}$. *Suppose* $[\ell, u]$ *is a closed subinterval of the unit interval* $[0, 1]$.

- $o \models_{\mathcal{O}}^{[\ell,u]} \mathbf{in}(\mathtt{X}, \mathtt{a}:_{\mathbf{RV}} f(\mathtt{d_1}, \ldots, \mathtt{d_n}))$
  if there is a $(Y, \wp)$ *in the answer returned by evaluating* $\mathtt{a}:_{\mathbf{RV}} f(\mathtt{d_1}, \ldots, \mathtt{d_n})$ *w.r.t.* $\mathcal{O}$ *such that* $o \in Y$ *and* $\ell \leq \wp(o) \leq u$.

- $o \models_{\mathcal{O}}^{[\ell,u]} \mathbf{not_in}(\mathtt{X}, \mathtt{a}:_{\mathbf{RV}} f(\mathtt{d_1}, \ldots, \mathtt{d_n}))$
  if for all random variables $(Y, \wp)$ *returned by evaluating* $\mathtt{a}:_{\mathbf{RV}} f(\mathtt{d_1}, \ldots, \mathtt{d_n})$ *w.r.t.* $\mathcal{O}$, *either* $o \notin Y$ *or* $\wp(o) \notin [\ell, u]$.

Probabilistic code call conditions are defined in exactly the same way as code call conditions. However, extending the above definition of "satisfaction" to probabilistic code call conditions is highly problematic because (as shown in

Examples 3.4 and 3.5), the probability that a conjunction is true depends not only on the probabilities of the individual conjuncts, but also on the dependencies between the events denoted by these conjuncts. The notion of a *probabilistic conjunction strategy* defined below captures these different ways of computing probabilities via an abstract definition.

Definition 3.8 (Probabilistic Conjunction Strategy $\otimes$)
A probabilistic conjunction strategy *is a mapping $\otimes$ which maps a pair of probability intervals to a single probability interval satisfying the following axioms:*

1. ***Bottomline:*** $[L_1, U_1] \otimes [L_2, U_2] \leq [\min(L_1, L_2), \min(U_1, U_2)]$ *where* $[x, y] \leq [x', y']$ *if* $x \leq x'$ *and* $y \leq y'$.
2. ***Ignorance:*** $[L_1, U_1] \otimes [L_2, U_2] \subseteq [\max(0, L_1 + L_2 - 1), \min(U_1, U_2)]$.
3. ***Identity:*** *When* $(e_1 \wedge e_2)$ *is consistent and* $[L_2, U_2] = [1, 1]$, $[L_1, U_1] \otimes [L_2, U_2] = [L_1, U_1]$.
4. ***Annihilator:*** $[L_1, U_1] \otimes [0, 0] = [0, 0]$.
5. Commutativity: $[L_1, U_1] \otimes [L_2, U_2] = [L_2, U_2] \otimes [L_1, U_1]$.
6. ***Associativity:*** $([L_1, U_1] \otimes [L_2, U_2]) \otimes [L_3, U_3] = [L_1, U_1] \otimes ([L_2, U_2] \otimes [L_3, U_3])$.
7. ***Monotonicity:*** $[L_1, U_1] \otimes [L_2, U_2] \leq [L_1, U_1] \otimes [L_3, U_3]$ *if* $[L_2, U_2] \leq [L_3, U_3]$.

Intuitively, in the above definition, $[L_1, U_1], [L_2, U_2]$ are intervals in which the probability of events e_1, e_2 are known to lie, and $[L_1, U_1] \otimes [L_2, U_2]$ returns a probability range for the co-occurrence of both these events. The Bottomline axiom says that the probability of the conjunct is smaller than the probabilities of the individual events. When we know nothing about the relationship between the events e_1, e_2, Boole (1854) has shown that the probability of the conjunction must lie in the interval $[\max(0, L_1 + L_2 - 1), \min(U_1, U_2)]$. This is what is stated in the Ignorance axiom. The identity and annihilator axioms specify what happens when one of the events is deterministic (i.e. not probabilistic). The axioms of commutativity and associativity are self explanatory. The monotonicity axiom says that if we sharpen the probability range of one of the two events, then the probability range of the conjunctive event is also sharpened.

The concept of a conjunction strategy is very general, and has as special cases, the following well known ways of combining probabilities.

1. When we do not know the dependencies between e_1, e_2, we may use the conjunction strategy $\otimes_{ig}$ defined as $([L_1, U_1] \otimes_{ig} [L_2, U_2]) \equiv [\max(0, L_1 + L_2 - 1), \min(U_1, U_2)]$.
2. When e_1, e_2 have maximal overlap, use the positive correlation conjunctive strategy $\otimes_{pc}$ defined as $([L_1, U_1] \otimes_{pc} [L_2, U_2]) \equiv [\min(L_1, L_2), \min(U_1, U_2)]$.

3. When e_1, e_2 have minimal overlap, use the negative correlation conjunctive strategy $\otimes_{nc}$ defined as $([L_1, U_1] \otimes_{nc} [L_2, U_2]) \equiv [\max(0, L_1 + L_2 - 1), \max(0, U_1 + U_2 - 1)]$.

4. When the two events occur independently, use the independence conjunction strategy $([L_1, U_1] \otimes_{in} [L_2, U_2]) = [L_1 \cdot L_2, U_1 \cdot U_2]$.

4 Probabilistic Agent Programs: Syntax

We are now ready to define the syntax of a probabilistic agent program (pap for short). This syntax builds upon the well studied *annotated logic paradigm* proposed by(Subrahmanian 1987), and later studied extensively (Kifer and Subrahmanian 1992; Ng and Subrahmanian 1993b; Ng and Subrahmanian 1993a).

4.1 Annotation Syntax

We assume the existence of an *annotation language* $\mathbf{L}^{ann}$—the constant symbols of $\mathbf{L}^{ann}$ are the real numbers in the unit interval $[0, 1]$. In addition, $\mathbf{L}^{ann}$ contains a finite set of function symbols, each with an associated arity, and a (possibly infinite) set of variable symbols, ranging over the unit interval $[0, 1]$. All function symbols are *pre-interpreted* in the sense that associated with each function symbol f of arity k is a fixed function from $[0, 1]^k$ to $[0, 1]$.

Definition 4.1 (Annotation Item) *We define annotation items inductively as follows:*

- *Every constant and every variable of* $\mathbf{L}^{ann}$ *is an annotation item.*
- *If* f *is an annotation function of arity* n *and* $\mathsf{ai}_1, \ldots, \mathsf{ai}_n$ *are annotation items, then* $f(\mathsf{ai}_1, \ldots, \mathsf{ai}_n)$ *is an annotation item.*

An annotation item is ground *if no annotation variables occur in it.*

For instance, $0, 0.9, (V+0.9), (V+0.9)^2$ are all annotation items if V is a variable in $\mathbf{L}^{ann}$ and "+", "^{2}" are annotation functions of arity 2.

Definition 4.2 (Annotation $[\mathsf{ai}_1, \mathsf{ai}_2]$) *If* $\mathsf{ai}_1, \mathsf{ai}_2$ *are annotation items, then* $[\mathsf{ai}_1, \mathsf{ai}_2]$ *is an* annotation. *If* $\mathsf{ai}_1, \mathsf{ai}_2$ *are both ground, then* $[\mathsf{ai}_1, \mathsf{ai}_2]$ *is a* ground annotation.

For instance, $[0, 0.4], [0.7, 0.9], [0.1, \frac{V}{2}], [\frac{V}{4}, \frac{V}{2}]$ are all annotations. The annotation $[0.1, \frac{V}{2}]$ denotes an interval only when a value in $[0, 1]$ is assigned to the variable V.

Definition 4.3 (Annotated Code Call Condition $\chi : \langle [\mathsf{ai}_1, \mathsf{ai}_2], \otimes \rangle$) *If* χ *is a probabilistic code call condition,* $\otimes$ *is a conjunction strategy, and* $[\mathsf{ai}_1, \mathsf{ai}_2]$ *is an annotation, then* $\chi : \langle [\mathsf{ai}_1, \mathsf{ai}_2], \otimes \rangle$ *is an* annotated code call condition. $\chi : \langle [\mathsf{ai}_1, \mathsf{ai}_2], \otimes \rangle$ *is* ground *if there are no variables in either* χ *or in* $[\mathsf{ai}_1, \mathsf{ai}_2]$.

Intuitively, the ground annotated code call condition $\chi : \langle[\mathsf{ai}_1, \mathsf{ai}_2], \otimes\rangle$ says that the probability of χ being true (under conjunction strategy $\otimes$) lies in the interval $[\mathsf{ai}_1, \mathsf{ai}_2]$. For example, when $X, A1$ are ground,

$$\mathbf{in}(\mathtt{X}, \mathrm{surv} :_{\mathbf{RV}} \mathit{identify}(\mathtt{image1})) \,\&\, \mathbf{in}(\mathtt{A1}, \mathrm{surv} :_{\mathbf{RV}} \mathit{turret}(\mathtt{X})) : \langle[0.3, 0.5], \otimes_{ig}\rangle$$

is true *if and only if* the probability that X is identified by the surv agent and that the turret is identified as being of type A1 lies between 30 and 50% assuming that nothing is known about the dependencies between turret identifications and identifications of objects by surv.

We are now ready to define the concept of a probabilistic agent program.

Definition 4.4 (Probabilistic Agent Programs $\mathcal{PP}$) *Suppose Γ is a conjunction of annotated code calls, and $A, L_1, \ldots, L_n$ are action status atoms. Then*

$$A \leftarrow \Gamma, L_1, \ldots, L_n \tag{1}$$

is a probabilistic action rule. *For such a rule r, we use $B^+_{as}(r)$ to denote the positive action status atoms in $\{L_1, \ldots, L_n\}$, and $B^-_{as}(r)$ to denote the set of negative action status liters in $\{L_1, \ldots, L_n\}$.*

A probabilistic agent program *is a finite set of probabilistic action rules.*

A simple example of a probabilistic agent program is given below.

Example 4.5 [Probabilistic Agent Program] Consider an intelligent sensor agent that is performing surveillance tasks. The following rules specify a small pap that such an agent might use.

$$\begin{array}{rcl}
\mathbf{Do}\, \mathit{send_warn}(X) & \leftarrow & \mathbf{in}(\mathtt{F}, \mathrm{surv} : \mathit{file}(\mathtt{imagedb})) \,\& \\
& & \mathbf{in}(\mathtt{X}, \mathrm{surv} :_{\mathbf{RV}} \mathit{identify}(\mathtt{F})) \,\& \\
& & \mathbf{in}(\mathtt{A1}, \mathrm{surv} :_{\mathbf{RV}} \mathit{turret}(\mathtt{X}))) : \langle[0.7, 1.0], \otimes_{ig}\rangle \\
& & \neg\mathbf{F}\mathit{send_warn}(X). \\
\mathbf{F}\mathit{send_warn}(X) & \leftarrow & \mathbf{in}(\mathtt{X}, \mathrm{surv} :_{\mathbf{RV}} \mathit{identify}(\mathtt{F})) \,\& \\
& & \mathbf{in}(\mathtt{L}, \mathrm{geo} :_{\mathbf{RV}} \mathit{getplnode}(\mathtt{X.location})) \,\& \\
& & \mathbf{in}(\mathtt{L}, \mathrm{geo} :_{\mathbf{RV}} \mathit{range}(\mathtt{100}, \mathtt{100}, \mathtt{20})).
\end{array}$$

This agent operates according to two very simple rules. The first rule says that it sends a warning whenever it identifies an enemy vehicle as having a gun turret of type A1 with over 70% probability, as long as sending such a warning is not forbidden. The second rule says that sending a warning is forbidden if the enemy vehicle is within 20 units of distance from location (100,100).

5 Probabilistic Agent Programs: Semantics

We are now ready to define the semantics of paps. The semantics of paps will be defined via the concept of a *probabilistic status set* (defined below).

Definition 5.1 (Probabilistic Status Set $\mathcal{PS}$) *A* probabilistic status set *is any set $\mathcal{PS}$ of ground action status atoms over $\mathcal{S}$. For any operator $\mathsf{Op} \in \{\mathbf{P}, \mathbf{Do}, \mathbf{F}, \mathbf{O}, \mathbf{W}\}$, we denote by $\mathsf{Op}(\mathcal{PS})$ the set $\{\alpha \mid \mathsf{Op}(\alpha) \in \mathcal{PS}\}$. Similarly, we use $\neg\mathcal{PS}$ to denote the set $\{\neg A \mid A \in \mathcal{PS}\}$.*

It will turn out that given any probabilistic agent program, and an (uncertain) agent state evaluated using probabilistic code calls, the meaning of the pap w.r.t. the state may be defined via a set of probabilistic status sets that have some desirable properties. These properties fall into three broad categories:

1. the probabilistic status set must be "closed" under the rules in the pap;
2. the probabilistic status set must be deontically consistent (e.g. it cannot require something to be both permitted and forbidden) and it must not violate the action constraints;
3. the probabilistic status set must not lead to a new state that violates the integrity constraints associated with the agent;

5.1 Satisfaction of Annotated Formulae

In this section, we define what it means for an agent state to satisfy an annotated code call condition.

Definition 5.2 (Satisfying an Annotated Code Call Condition) *Suppose $\mathcal{O}$ is an agent state, and $\chi : \langle[\mathsf{ai}_1, \mathsf{ai}_2], \otimes\rangle$ is a ground annotated code call condition. $\mathcal{O}$ is said to* satisfy *$\chi : \langle[\mathsf{ai}_1, \mathsf{ai}_2], \otimes\rangle$, denoted $\mathcal{O} \models^{[\mathsf{ai}_1,\mathsf{ai}_2]} \chi : \langle[\mathsf{ai}_1, \mathsf{ai}_2], \otimes\rangle$ iff:*

- *χ is of the form $o = o$ (where o is an object), or*
- *χ is of the form $r_1 < r_2$, where r_1, r_2 are real numbers (or integers) such that r_1 is less than r_2, or*
- *χ is of the form $\mathbf{in}(\mathtt{X}, \mathtt{a} :_{\mathbf{RV}} f(\mathtt{d}_1, \ldots, \mathtt{d}_n))$ and $o \models_{\mathcal{O}}^{[\mathsf{ai}_1,\mathsf{ai}_2]} \mathbf{in}(\mathtt{X}, \mathtt{a} :_{\mathbf{RV}} f(\mathtt{d}_1, \ldots, \mathtt{d}_n))$, or*
- *χ is of the form $\mathbf{not_in}(\mathtt{o}, \mathtt{a} :_{\mathbf{RV}} f(\mathtt{d}_1, \ldots, \mathtt{d}_n))$ and the following holds $o \models_{\mathcal{O}}^{[\mathsf{ai}_1,\mathsf{ai}_2]} \mathbf{not_in}(\mathtt{X}, \mathtt{a} :_{\mathbf{RV}} f(\mathtt{d}_1, \ldots, \mathtt{d}_n))$, or*
- *χ is of the form $\chi_1 \wedge \chi_2$ and $[\ell_1, u_1], [\ell_2, u_2]$ are the tightest intervals such that $\mathcal{O} \models^{[\ell_1,u_1]} \chi_1$ and $\mathcal{O} \models^{[\ell_2,u_2]} \chi_2$ and $[\mathsf{ai}_1, \mathsf{ai}_2] \supseteq [\ell_1, u_1] \otimes [\ell_2, u_2]$.*

$\mathcal{O}$ is said to satisfy *a non-ground annotated code call $\chi : \langle[\mathsf{ai}_1, \mathsf{ai}_2], \otimes\rangle$ iff $\mathcal{O}$ satisfies all ground instances of $\chi : \langle[\mathsf{ai}_1, \mathsf{ai}_2], \otimes\rangle$.*

5.2 Closure and $\mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS})$

We may associate with any pap $\mathcal{PP}$, an operator $\mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS})$ which maps probabilistic status sets to probabilistic status sets.

Definition 5.3 (Operator $\mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS})$) *Suppose $\mathcal{PP}$ is a probabilistic agent program, $\mathcal{O}_\mathcal{S}$ is an agent state, and $\mathcal{PS}$ is a probabilistic status set. Then*

$$\boldsymbol{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}) = \{\mathsf{Op}\,\alpha \mid \mathsf{Op}\,\alpha \text{ is the head of a ground instance } r \text{ of a rule in } \mathcal{PP} \text{ satisfying the 4 conditions below}\}$$

1. $B^+_{as}(r) \subseteq \mathcal{PS}$ *and* $\neg.B^-_{as}(r) \cap \mathcal{PS} = \emptyset$, *and*
2. *For every annotated code call condition* $\chi : \langle[\mathsf{ai}_1, \mathsf{ai}_2], \otimes\rangle$ *in the body of* r, *it is the case that* $\mathcal{O}_\mathcal{S} \models^{[\mathsf{ai}_1,\mathsf{ai}_2]} \chi : \langle[\mathsf{ai}_1, \mathsf{ai}_2], \otimes\rangle$ *and*
3. *if* $\mathsf{Op} \in \{\mathbf{P}, \mathbf{O}, \mathbf{Do}\}$, *then* $\mathcal{O}_\mathcal{S} \models^{[1,1]} Pre(\alpha)$ *and*
4. *for every action status atom of the form* $\mathsf{Op}\,\beta$ *in* $B^+_{as}(r)$ *such that* $\mathsf{Op} \in \{\mathbf{P}, \mathbf{O}, \mathbf{Do}\}$, $\mathcal{O}_\mathcal{S} \models^{[1,1]} Pre(\beta)$.

The first part of this definition says that for a rule to fire, the action status atoms in its body must be "true" w.r.t. $\mathcal{PS}$. The second condition says that annotated code call conditions in a rule body must be satisfied in the current object state for the rule to fire. The third part is more tricky. It says that if $\mathbf{O}\alpha$ or $\mathbf{Do}\,\alpha$ or $\mathbf{P}\alpha$ is in the head of a rule, then for the rule to fire, the precondition of the action must be true with 100% probability. The final condition is similar w.r.t. to positive action status atoms in the body. *Thus, for now, we are assuming that for an agent to perform an action (or even be permitted to perform an action), it must be 100% sure that the action's precondition is true* (later in Section 8, we will provide an alternate, more complex semantics that does not require this).

Definition 5.4 (Closure under Program Rules) *$\mathcal{PS}$ is said to be closed under the rules of* pap *$\mathcal{PP}$ in state $\mathcal{O}^p$ iff $\boldsymbol{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}) \subseteq \mathcal{PS}$.*

5.3 Deontic/Action Consistency/Closure

The concept of deontic/action consistency requires that probabilistic status sets satisfy the agent's action constraints and commonsense axioms about deontic modalities.

Definition 5.5 (Deontic and Action Consistency) *A probabilistic status set $\mathcal{PS}$ is* deontically consistent *with respect to an agent state $\mathcal{O}$ iff it satisfies the following rules for any ground action α:*

- *If* $\mathbf{O}\alpha \in \mathcal{PS}$, *then* $\mathbf{W}\alpha \notin \mathcal{PS}$.
- *If* $\mathbf{P}\alpha \in \mathcal{PS}$, *then* $\mathbf{F}\alpha \notin \mathcal{PS}$.
- *If* $\mathbf{P}\alpha \in \mathcal{PS}$, *then* $\mathcal{O} \models^{[1,1]} Pre(\alpha)$.

A probabilistic status set $\mathcal{PS}$ is action consistent *w.r.t. $\mathcal{O}$ iff for every action constraint of the form*

$$\{\alpha_1(\vec{X}_1), \ldots, \alpha_k(\vec{X}_k)\} \hookleftarrow \chi \tag{2}$$

either $\mathcal{O} \not\models^{[1,1]} \chi$ or $\{\alpha_1(\vec{X}_1), \ldots, \alpha_k(\vec{X}_k)\} \not\subseteq \mathbf{Do}\,(\mathcal{PS})$.

The following example illustrate the concept of deontic and action consistency.

Example 5.6 Suppose we have a resource allocation agent having two actions — *send_A*() and *send_B*() — each of which sends a unit of the resource respectively to agents A B. To execute either of them, we need to have at least one unit of resource, and to execute them together we need at least 2 units:

$$\begin{array}{l} Pre(send_A()) = \mathbf{in}(\mathtt{X}, \mathtt{allocator}: avail_rsc()) \ \&\ X > 0. \\ Pre(send_B()) = \mathbf{in}(\mathtt{X}, \mathtt{allocator}: avail_rsc()) \ \&\ X > 0. \\ \{send_to_A(), send_to_B()\} \hookleftarrow \mathbf{in}(\mathtt{X}, \mathtt{allocator}: avail_rsc()) \ \&\ X < 2 \end{array}$$

Suppose the agent's current state $\mathcal{O}$ is one in which *avail_rsc*() returns 1. Then

$$\begin{array}{rl} \mathcal{PS} = & \{\mathbf{P}\,send_to_A(), \mathbf{Do}\,send_to_A(), \\ & \mathbf{Do}\,send_to_B(), \mathbf{O}send_to_B()\} \end{array}$$

is deontically consistent (there are no **W** and **F** atoms at all, and the action preconditions are true), but not action consistent.

The deontic and action closure of a probabilistic status set $\mathcal{PS}$ is defined in exactly the same way (see appendix, Definition A.3 on page 39) as in the non-probabilistic case.

5.4 Probabilistic State Consistency

The final requirement of a feasible probabilistic status set ensures that the new state that results after concurrently executing a set of actions is consistent with the integrity constraints.

$\mathcal{O}$ *satisfies* the integrity constraint $\psi \Rightarrow \chi$ iff either $\mathcal{O} \not\models^{[1,1]} \psi$ or $\mathcal{O} \models^{[1,1]} \chi$.

Definition 5.7 (Probabilistic State Consistency) *A probabilistic status set $\mathcal{PS}$ is* probabilistically state consistent *w.r.t. $\mathcal{O}_\mathcal{S}$ iff the new state, $\mathcal{O}'_\mathcal{S} = \boldsymbol{conc}(\mathbf{Do}\,(\mathcal{PS}), \mathcal{O}_\mathcal{S})$ obtained after concurrently executing all actions of the form $\mathbf{Do}\,\alpha \in \mathcal{PS}$ satisfies all integrity constraints.*

The following example illustrates the concept of probabilistic state consistency.

Example 5.8 Suppose we have a vehicle coordination agent that tracks vehicle on a road (line), and makes sure that two vehicles do not collide. Such an agent may have the integrity constraint

$$\mathbf{in}(\mathtt{X}, \mathtt{geo}: getposition(\mathtt{a})) \ \&\ \mathbf{in}(\mathtt{Y}, \mathtt{geo}: getposition(\mathtt{b})) \Rightarrow X \neq Y$$

It may be able to perform an action $move_forward(a)$:

$$\begin{array}{ll} Pre(move_forward(a)) & = \mathbf{in}(\mathtt{X}, \mathtt{geo}: getposition(\mathtt{a})) \\ Del(move_forward(a)) & = \mathbf{in}(\mathtt{X}, \mathtt{geo}: getposition(\mathtt{a})) \\ Add(move_forward(a)) & = \mathbf{in}(\mathtt{X}+1, \mathtt{geo}: getposition(\mathtt{a})) \end{array}$$

In a state $\mathcal{O}$ where $\mathtt{geo}: getposition(\mathtt{a})$ returns 200, and $\mathtt{geo}: getposition(\mathtt{b})$ returns 201, the status set

$$\mathcal{PS} = \{\mathbf{P}move_forward(a), \mathbf{Do}\, move_forward(a)\}$$

is not state consistent, as executing $\mathbf{Do}\,(\mathcal{PS})$ leads to where both agent $\mathtt{a}$ and agent $\mathtt{b}$ are in position 201, violating the above integrity constraint.

5.5 Feasible Probabilistic Status Sets

The meaning of a pap (w.r.t. a given state) may be characterized via those probabilistic status sets that satisfy the conditions of closure under program rules, deontic/action consistency and probabilistic state consistency. Such probabilistic status sets are said to be *feasible*.

Definition 5.9 (Feasible Probabilistic Status Set) *Suppose $\mathcal{PP}$ is an agent program and $\mathcal{O}_\mathcal{S}$ is an agent state. A probabilistic status set $\mathcal{PS}$ is* feasible *for $\mathcal{PP}$ on $\mathcal{O}_\mathcal{S}$ if the following conditions hold:*

($\mathcal{PS}$ 1): $\boldsymbol{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}) \subseteq \mathcal{PS}$ (closure under the program rules)*;*

($\mathcal{PS}$ 2): *$\mathcal{PS}$ is deontically and action consistent* (deontic/action consistency)*;*

($\mathcal{PS}$ 3): *$\mathcal{PS}$ is action closed and deontically closed* (deontic/action closure)*;*

($\mathcal{PS}$ 4): *$\mathcal{PS}$ is state consistent* (state consistency)*.*

paps may have zero, one or many feasible status sets, as seen via the following examples.

Example 5.10 Consider the following agent program.

$$\begin{array}{lcl} \mathbf{P}send_warn(t80) & \leftarrow & . \\ \mathbf{F}send_warn(t80) & \leftarrow & . \end{array}$$

In any agent state $\mathcal{O}_\mathcal{S}$ such that $\mathcal{O}_\mathcal{S} \models^{[1,1]} Pre(send_warn(t80))$, the above program cannot have any feasible probabilistic status set $\mathcal{PS}$. This is because closure under program rules requires that $\mathbf{P}send_warn(t80), \mathbf{F}send_warn(t80)$ are both in $\mathcal{PS}$, but this causes $\mathcal{PS}$ to violate deontic consistency.

In contrast, consider the following one-rule program for checking the power level of surveillance equipment.

$$\mathbf{O}power_warn() \leftarrow \mathbf{in}(\mathtt{X}, \mathtt{surv}: powerlevel()) \;\&\; X < 2000.$$

Suppose $\mathsf{surv}: \mathit{powerlevel}()$ returns 1000 in some state $\mathcal{O}_{\mathcal{S}}$, and suppose $\mathit{power_warn}()$ has no preconditions. If no integrity and action constraints are present, then this pap has exactly one feasible status set, viz. for $\mathit{power_warn}()$, and without action

$$\{\mathbf{O}\mathit{power_warn}(), \mathbf{Do}\,\mathit{power_warn}(), \mathbf{P}\mathit{power_warn}()\}$$

in $\mathcal{O}_{\mathcal{S}}$.

Now let us consider a pap which says that one of the two agents a, b must be warned (if it is active). Furthermore, if b is to be warned, its regular (non emergency) channel must not be on.

$$\begin{array}{rcl}
\mathbf{F}\mathit{open_ch}(b) & \leftarrow & \neg\mathbf{F}\mathit{open_ch}(b)\ \&\ \mathbf{Do}\,\mathit{warn_ag}(b). \\
\mathbf{Do}\,\mathit{warn_ag}(a) & \leftarrow & \mathbf{in}(\mathtt{a}, \mathsf{surv}: \mathit{activeagents}())\ \&\ \neg\mathbf{Do}\,\mathit{warn_ag}(b). \\
\mathbf{O}\mathit{warn_ag}(b) & \leftarrow & \mathbf{in}(\mathtt{b}, \mathsf{surv}: \mathit{activeagents}())\ \&\ \neg\mathbf{Do}\,\mathit{warn_ag}(a).
\end{array}$$

We assume the absence of integrity constraints, and preconditions for all actions. However, the following action constraint is present:

$$\{\mathit{warn_ag}(a), \mathit{warn_ag}(b)\} \hookleftarrow .$$

If $\mathsf{surv}: \mathit{activeagents}()$ returns $\{a, b\}$ in state $\mathcal{O}_{\mathcal{S}}$, then the above program has several feasible status sets:

$$\begin{array}{l}
\{\mathbf{F}\mathit{open_ch}(b), \mathbf{O}\mathit{warn_ag}(b), \mathbf{Do}\,\mathit{warn_ag}(b), \mathbf{P}\mathit{warn_ag}(b)\} \\
\{\mathbf{F}\mathit{open_ch}(b), \mathbf{Do}\,\mathit{warn_ag}(a), \mathbf{P}\mathit{warn_ag}(a)\} \\
\{\mathbf{Do}\,\mathit{warn_ag}(a), \mathbf{P}\mathit{warn_ag}(a)\}
\end{array}$$

Notice that no feasible status set contains both $\mathbf{Do}\,\mathit{warn_ag}(a)$ and $\mathbf{Do}\,\mathit{warn_ag}(b)$.

5.6 Rational Probabilistic Status Sets

As seen from the above examples, feasible status sets may contain action status atoms that are not required for feasibility. Rational probabilistic status sets refine this definition.

Definition 5.11 (Groundedness; Rational Probabilistic Status Set) *A probabilistic status set $\mathcal{PS}$ is* grounded *if there is no probabilistic status set $\mathcal{PS}' \neq \mathcal{PS}$ such that $\mathcal{PS}' \subseteq \mathcal{PS}$ and $\mathcal{PS}'$ satisfies conditions ($\mathcal{PS}1$)–($\mathcal{PS}3$) of a feasible probabilistic status set.*

$\mathcal{PS}$ is rational *iff it is feasible and grounded.*

Example 5.12 Consider the last case in Example 5.10. Only two of the listed feasible status sets are rational, viz.

$$\begin{array}{l}
\{\mathbf{F}\mathit{open_ch}(b), \mathbf{O}\mathit{warn_ag}(b), \mathbf{Do}\,\mathit{warn_ag}(b), \mathbf{P}\mathit{warn_ag}(b)\} \text{ and} \\
\{\mathbf{Do}\,\mathit{warn_ag}(a), \mathbf{P}\mathit{warn_ag}(a)\}
\end{array}$$

5.7 Reasonable Probabilistic Status Sets

As we can see from the preceding example, certain action status atoms may be true in a rational status set even though there is no rule whose head contains (or implies) that action status atom. The concept of a reasonable status set (which is derived from the well known stable model semantics of logic programs (Gelfond and Lifschitz 1988)) prevents this.

Definition 5.13 (Reasonable Probabilistic Status Set) *Suppose $\mathcal{PP}$ is a* pap*, $\mathcal{O}_{\mathcal{S}}$ is an agent state, and $\mathcal{PS}$ is a probabilistic status set.*

1. *If $\mathcal{PP}$ is a positive (i.e. $B^-_{as}(r) = \emptyset$ for all $r \in \mathcal{PP}$), then $\mathcal{PS}$ is a* reasonable probabilistic status set *for $\mathcal{PP}$ on $\mathcal{O}_{\mathcal{S}}$, if, by definition, $\mathcal{PS}$ is a rational probabilistic status set for $\mathcal{PP}$ on $\mathcal{O}_{\mathcal{S}}$.*

2. *Otherwise, the reduct of $\mathcal{PP}$ w.r.t. $\mathcal{PS}$ and $\mathcal{O}_{\mathcal{S}}$, denoted by $red^{\mathcal{PS}}(\mathcal{PP}, \mathcal{O}_{\mathcal{S}})$, is the program which is obtained from the ground instances of the rules in $\mathcal{PP}$ over $\mathcal{O}_{\mathcal{S}}$ as follows.*

   (a) *First, remove every rule r such that $B^-_{as}(r) \cap \mathcal{PS} \neq \emptyset$;*

   (b) *Remove all atoms in $B^-_{as}(r)$ from the remaining rules.*

   Then $\mathcal{PS}$ is a reasonable probabilistic status set *for $\mathcal{PP}$ w.r.t. $\mathcal{O}_{\mathcal{S}}$, if it is a reasonable probabilistic status set of the program $red^{\mathcal{PS}}(\mathcal{PP}, \mathcal{O}_{\mathcal{S}})$ with respect to $\mathcal{O}_{\mathcal{S}}$.*

The following example illustrates the concept of a reasonable status set.

Example 5.14 Consider again the last case in Example 5.10. Only one of the listed feasible status sets is reasonable, viz.

$$\mathcal{PS} = \{\mathbf{Do}\, warn_ag(a), \mathbf{P}\, warn_ag(a)\}$$

To see why this probabilistic status set is feasible, note that the reduct of $\mathcal{PP}$ w.r.t. $\mathcal{PS}$ is:

$$\mathbf{Do}\, warn_ag(a) \quad \leftarrow \quad \mathbf{in}(\mathtt{a}, \mathtt{surv} : activeagents()).$$

whose (unique) rational status set is obviously $\mathcal{PS}$.

5.8 Semantical Properties

In this section, we prove some properties about the different semantics described above.

Proposition 5.15 (Properties of Feasible Status Sets) *Let $\mathcal{PS}$ be a feasible probabilistic status set. Then,*

1. *If $\mathbf{Do}\,(\alpha) \in \mathcal{PS}$, then $\mathcal{O}_{\mathcal{S}} \models^{[1,1]} Pre(\alpha)$;*

2. *If* $\mathbf{P}\alpha \notin \mathcal{PS}$, *then* $\mathbf{Do}\,(\alpha) \notin \mathcal{PS}$;

3. *If* $\mathbf{O}\alpha \in \mathcal{PS}$, *then* $\mathcal{O}_{\mathcal{S}} \models^{[1,1]} Pre(\alpha)$;

4. *If* $\mathbf{O}\alpha \in \mathcal{PS}$, *then* $\mathbf{F}\alpha \notin \mathcal{PS}$.

The following theorem says that reasonable status sets are rational.

Theorem 5.16 (Reasonable Status Sets are Rational) *Let $\mathcal{PP}$ be a probabilistic agent program and $\mathcal{O}_{\mathcal{S}}$ an agent state. Every reasonable probabilistic status set of $\mathcal{PP}$ on $\mathcal{O}_{\mathcal{S}}$ is a rational probabilistic status set of $\mathcal{PP}$ on $\mathcal{O}_{\mathcal{S}}$.*

Given any pap $\mathcal{PP}$ and agent state $\mathcal{O}_{\mathcal{S}}$, we may define an operator that maps probabilistic status sets to probabilistic status sets as follows. We use then notation $\mathbf{D\text{-}Cl}(\mathcal{PS})$ to denote the closure of $\mathcal{PS}$ under the rule $\mathbf{O}\alpha \in \mathcal{PS} \Rightarrow \mathbf{P}\alpha \in \mathcal{PS}$ and $\mathbf{A\text{-}Cl}(\mathcal{PS})$ to denote the closure of $\mathcal{PS}$ under the rules $\mathbf{O}\alpha \in \mathcal{PS} \Rightarrow \mathbf{Do}\,\alpha \in \mathcal{PS}$ and $\mathbf{Do}\,\alpha \in \mathcal{PS} \Rightarrow \mathbf{O}\alpha \in \mathcal{PS}$.

Definition 5.17 ($\mathbf{T}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}$ Operator) *Suppose $\mathcal{PP}$ is a probabilistic agent program and $\mathcal{O}_{\mathcal{S}}$ is an agent state. Then, for any probabilistic status set $\mathcal{PS}$,*

$$\boldsymbol{T}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}(\mathcal{PS}) = \boldsymbol{App}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}(\mathcal{PS}) \cup \mathbf{D\text{-}Cl}(\mathcal{PS}) \cup \mathbf{A\text{-}Cl}(\mathcal{PS}).$$

Note that as $\mathbf{D\text{-}Cl}(\mathcal{PS}) \subseteq \mathbf{A\text{-}Cl}(\mathcal{PS})$, we may equivalently write this as

$$\mathbf{T}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}(\mathcal{PS}) = \mathbf{App}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}(\mathcal{PS}) \cup \mathbf{A\text{-}Cl}(\mathcal{PS}).$$

The following property of feasible probabilistic status sets is easily seen.

Lemma 5.18 ($\mathcal{PS}$ as Prefixpoint of $\mathbf{T}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}$) *Let $\mathcal{PP}$ be a probabilistic agent program, $\mathcal{O}_{\mathcal{S}}$ be any agent state, and $\mathcal{PS}$ be any probabilistic status set. If $\mathcal{PS}$ satisfies conditions ($\mathcal{PS}1$) and ($\mathcal{PS}3$) of feasibility, then $\mathcal{PS}$ is pre-fixpoint of $\boldsymbol{T}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}$, i.e., $\boldsymbol{T}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}(\mathcal{PS}) \subseteq \mathcal{PS}$.*

Proof: Suppose $\mathsf{Op}(\alpha) \in \mathbf{T}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}(\mathcal{PS}) = \mathbf{App}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}(\mathcal{PS}) \cup \mathbf{A\text{-}Cl}(\mathcal{PS})$. Then we have either $\mathsf{Op}(\alpha) \in \mathbf{App}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}(\mathcal{PS})$ or $\mathsf{Op}(\alpha) \in \mathbf{A\text{-}Cl}(\mathcal{PS})$. By condition ($\mathcal{PS}$ 1) defining a feasible probabilistic status set, we know that $\mathbf{App}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}(\mathcal{PS}) \subseteq \mathcal{PS}$. By condition ($\mathcal{PS}$ 3), $\mathcal{PS} = \mathbf{A\text{-}Cl}(\mathcal{PS})$ and hence, $\mathbf{A\text{-}Cl}(\mathcal{PS}) \subseteq \mathcal{PS}$. Therefore, $\mathbf{T}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}(\mathcal{PS}) \subseteq \mathcal{PS}$. ∎

The following theorem says that in the absence of integrity constraints, a pap has a rational probabilistic status set *if and only if* it has a feasible one.

Theorem 5.19 (Existence of Rational Probabilistic Status Sets) *Let $\mathcal{PP}$ be a probabilistic agent program. If $\mathcal{IC} = \emptyset$, then $\mathcal{PP}$ has a rational probabilistic status set if and only if $\mathcal{PP}$ has a feasible probabilistic status set.*

6 Computing Probabilistic Status Sets of Positive paps

In this section, we present a sound and complete algorithm to compute the unique reasonable status set of a positive pap. For this purpose, we use a variant of the $\mathbf{T}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$ operator introduced earlier. This operator, denoted $\mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$, is defined as

$$\mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}) = \mathbf{A\text{-}Cl}(\mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS})).$$

Computationally, we may compute the operator $\mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$ using algorithm 6.1 below.

Algorithm 6.1 ($\mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$ Computation for Positive paps)
Compute-$\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$ ($\mathcal{O}_\mathcal{S}$: agent state, $\mathcal{PS}$: probabilistic status set)

(⋆ the probabilistic agent program $\mathcal{PP}$ is positive; ⋆)
*(⋆ **Input:** an agent state $\mathcal{O}_\mathcal{S}$, and a prob. status set $\mathcal{PS}$ ⋆)*
*(⋆ **Output:** a deontically and action consistent set $\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS})$ (if existent) ⋆)*
(⋆ or "no consistent set exists" *⋆)*

1. $X := \mathcal{PS}$;
2. ***for each*** *rule* $r \in \mathcal{PP}$
3. ***for each*** *ground instance* $r\theta$ *of* r
4. ***if*** $r\theta = \mathsf{Op}\,\alpha \leftarrow \Gamma, L_1, \ldots, L_n$ ***and***
   $\mathcal{O}_\mathcal{S} \models \Gamma$ *and* $\{L_1, \ldots, L_n\} \subseteq \mathcal{PS}$ ***and***
   for every atom $\mathsf{Op}'(\beta) \in \{L_1, \ldots, L_n\} \cup \{\mathsf{Op}(\alpha)\}$
   such that $\mathsf{Op}' \in \{\mathbf{P}, \mathbf{O}, \mathbf{Do}\}$: $\mathcal{O}_\mathcal{S} \models^{[1,1]} Pre(\beta)$
5. ***then*** $X := X \cup \mathbf{A\text{-}Cl}(\{\mathsf{Op}\,\alpha\})$,
6. ***if*** X *contains* ($\mathbf{O}\alpha$ *and* $\mathbf{W}\alpha$) *or* ($\mathbf{P}\alpha$ *and* $\mathbf{F}\alpha$)
7. ***then Return*** "no consistent set exists".
8. ***Return*** X.

end.

The behavior of Algorithm 6.1 is illustrated by the following example.

Example 6.1 Consider the following program, saying that whenever a (probably) enemy vehicle Y is detected, a warning message about Y is sent to a friendly source and the agent perfroming the detection is not allowed to move.

$$\begin{aligned}
\mathbf{O}send_warn(Y) &\leftarrow \mathbf{in}(\mathtt{F}, \mathrm{surv}: file(\mathtt{imagedb}))\,\&\\
&\quad \mathbf{in}(\mathtt{Y}, \mathrm{surv}:_{\mathbf{RV}} identify(\mathtt{F}))\langle[0.5, 1.0], \otimes_{ig}\rangle\,\&\\
&\quad \mathbf{O}send_warn(X).\\
\mathbf{F}move() &\leftarrow \mathbf{Do}\,send_warn(X).\\
\mathbf{O}send_warn(X) &\leftarrow \mathbf{in}(\mathtt{F}, \mathrm{surv}: file(\mathtt{imagedb}))\,\&\\
&\quad \mathbf{in}(\mathtt{X}, \mathrm{surv}:_{\mathbf{RV}} identify(\mathtt{F}))\,\&\\
&\quad \mathbf{in}(\mathtt{X}, \mathrm{surv}: enemyvehicles()) : \langle[0.5, 1.0], \otimes_{ig}\rangle.
\end{aligned}$$

Moreover, assume that in the current state $\mathcal{O}_\mathcal{S}$, $\mathtt{surv}: \mathit{file}(\mathtt{imagedb})$ returns $\mathit{image}1$, $\mathtt{surv}: \mathit{identify}(\mathtt{image1})$ returns the random variables $\langle\{t80\}, \{\langle t80, 0.6\rangle\}\rangle$ and $\langle\{t72\}, \{\langle t72, 0.5\rangle\}\rangle$, and $\mathtt{surv}: \mathit{enemyvehicles}()$ returns $t80$.

Now we apply Algorithm 6.1 to compute $\mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{O}_\mathcal{S}, \emptyset)$: Step 1 sets $X = \emptyset$, step 2 selects the first rule, while step 3 considers all ground instances (of the first rule) whose body's truth is checked in step 4: no instance satisfies it (because $\mathcal{PS}$ is empty), so nothing happens. The same result is obtained when step 2 considers the second rule. Eventually, step 2 considers the third rule, which satisfies the condition of step 4 with its head instantiated to $\mathbf{O}\mathit{send_warn}(t80)$. Step 5 inserts $\mathbf{O}\mathit{send_warn}(t80)$, $\mathbf{Do}\,\mathit{send_warn}(t80)$, $\mathbf{P}\mathit{send_warn}(t80)$ into X. There is no deontic inconsistency, so the check n step 6 fails and then we jump to step 8, which returns the result:

$$X = \{\ \mathbf{O}\mathit{send_warn}(t80), \mathbf{Do}\,\mathit{send_warn}(t80), \mathbf{P}\mathit{send_warn}(t80)\ \}.$$

The operator $\mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$ may be iteratively applied as follows.

$$\begin{aligned}
\mathbf{S}^0_{\mathcal{PP},\mathcal{O}_\mathcal{S}} &= \emptyset. \\
\mathbf{S}^{i+1}_{\mathcal{PP},\mathcal{O}_\mathcal{S}} &= \mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathbf{S}^i_{\mathcal{PP},\mathcal{O}_\mathcal{S}}). \\
\mathbf{S}^\omega_{\mathcal{PP},\mathcal{O}_\mathcal{S}} &= \bigcup_{i=0}^{\infty} \mathbf{S}^i_{\mathcal{PP},\mathcal{O}_\mathcal{S}}.
\end{aligned}$$

The following theorem says that for positive paps, operator $\mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$ is monotonic, continuous, and has a (unique) least fixpoint.

Lemma 6.2 (Monotonicity and Continuity of $\mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$) *Suppose $\mathcal{PP}$ is a positive pap. Then the operator $\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$ is monotone and continuous, i.e.*

1. $\mathcal{PS}_1 \subseteq \mathcal{PS}_2 \Rightarrow \boldsymbol{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}_1) \subseteq \boldsymbol{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}_2)$,
2. $\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\bigcup_{i=0}^{\infty} \mathcal{PS}_0) = \bigcup_{i=0}^{\infty} \boldsymbol{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}_i)$ *for any chain* $\mathcal{PS}_0 \subseteq \mathcal{PS}_1 \subseteq \mathcal{PS}_2 \subseteq \cdots$ *of probabilistic status sets,*
3. $\boldsymbol{S}^\omega_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$ *is a fixpoint of* $\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$. *Moreover, it is the least fixpoint of* $\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$.

Proof: By the well known Knaster/Tarski theorem, *3.* follows from *1.* and *2.*.

To show *1.*, let $\mathcal{PS}_1 \subseteq \mathcal{PS}_2$. But then

$$\mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}_1) \subseteq \mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}_2),$$

because of the monotonicity of $\mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}()$ (see (Lloyd 1987; Apt 1990)). This implies $\mathbf{A\text{-}Cl}(\mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}_1)) \subseteq \mathbf{A\text{-}Cl}(\mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}_2))$.

2. follows similarly from the continuity of $\mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}()$ and the fact that

$$\mathbf{A\text{-}Cl}(\bigcup_{i=0}^{\infty} \mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}_i)) = \bigcup_{i=0}^{\infty} \mathbf{A\text{-}Cl}(\mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}_i)).$$

∎

The following example shows the computation of $\mathbf{S}^{\omega}_{\mathcal{PP},\mathcal{O}_S}$.

Example 6.3 Consider the program in Example 6.1. Applying the operator $\mathbf{S}_{\mathcal{PP},\mathcal{O}_S}$ iteratively, we obtain:

$$\begin{aligned}
\mathbf{S}^0_{\mathcal{PP},\mathcal{O}_S} &= \emptyset. \\
\mathbf{S}^1_{\mathcal{PP},\mathcal{O}_S} &= \{\mathbf{O}send_warn(t80), \mathbf{Do}\, send_warn(t80), \mathbf{P}send_warn(t80)\} \\
\mathbf{S}^2_{\mathcal{PP},\mathcal{O}_S} &= \mathbf{S}^1_{\mathcal{PP},\mathcal{O}_S} \cup \{\mathbf{O}send_warn(t72), \mathbf{Do}\, send_warn(t72), \\
&\quad \mathbf{P}send_warn(t72), \mathbf{F}move()\} \\
\mathbf{S}^3_{\mathcal{PP},\mathcal{O}_S} &= \mathbf{S}^2_{\mathcal{PP},\mathcal{O}_S} = \mathbf{S}^{\omega}_{\mathcal{PP},\mathcal{O}_S}
\end{aligned}$$

The following results tell us that Lemma 5.18, which holds for arbitrary programs, can be strengthened to the case of positive probabilistic programs.

Theorem 6.4 (Rational Probabilistic Status Sets as Least Fixpoints) *Suppose $\mathcal{PP}$ is a positive* pap*, and $\mathcal{O}_S$ is an agent state. Then: $\mathcal{PS}$ is a rational probabilistic status set if and only if $\mathcal{PS} = lfp(\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_S})$ and $\mathcal{PS}$ is a feasible probabilistic status set. Recall that lfp stands for least fixpoint.*

Corollary 6.5 *Let $\mathcal{PP}$ be a positive probabilistic agent program. Then, on every agent state $\mathcal{O}_S$, the rational probabilistic status set of $\mathcal{PP}$ (if one exists) is unique, i.e., if $\mathcal{PS}, \mathcal{PS}'$ are rational probabilistic status sets for $\mathcal{PP}$ on $\mathcal{O}_S$, then $\mathcal{PS} = \mathcal{PS}'$.*

An important corollary of this theorem is that to compute a reasonable feasible status set of a pap, all we need to do it to compute $lfp(\mathbf{S}_{\mathcal{PP},\mathcal{O}_S})$. This may be done via Algorithm **Compute-*lfp*** below.

Algorithm 6.2 (Reas. Prob. Status Set Computation for Positive paps)

Compute-lfp ($\boldsymbol{T}_{\mathcal{PP},\mathcal{O}_S}$): agent state, $\mathcal{PP}$: probabilistic agent program)

(⋆ the probabilistic agent program $\mathcal{PP}$ is positive; ⋆)
*(⋆ **Input:** an agent state $\mathcal{O}_S$, and a pap $\mathcal{PP}$ ⋆)*
*(⋆ **Output:** a reasonable probabilistic status set ⋆)*

1. *change* := **true**; $X := \emptyset$;
2. ***while*** *change* ***do***
3. *newX* = ***Compute-$\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_S}$ (X)***;
4. ***if*** *newX* := "no consistent set exists"
5. ***then return*** *no reasonable prob. status set exists.*
6. ***if*** $X \neq newX$ ***then*** $X := newX$
7. ***else*** *change* := **false**.
8. ***end while***;

9. ***if*** *X satisfies all the following conditions*
10. • $\mathbf{Do}\,\alpha \in X \Rightarrow \mathcal{O} \models^{[1,1]} Pre(\alpha)$;
11. • *The new state obtained by executing* ***conc***$(\{\mathbf{Do}\,\alpha \mid \mathbf{Do}\,\alpha \in X\}$
12. *satisfies the integrity constraints;*
13. • $\{\mathbf{Do}\,\alpha \mid \mathbf{Do}\,\alpha \in X\}$ *satisfies the action constraints.*
14. ***then return*** *X*
15. ***else return*** *no reasonable prob. status set exists.*
end.

Theorem 6.6 (Polynomial Data Complexity) *Algorithm* **Compute-***lfp has polynomial data-complexity.*

Proof: It is easy to see that the **while** loop of the algorithm can be executed in polynomial time (data-complexity). Checking if X satisfies the three conditions at the end of the algorithm are each polynomial time checks (assuming the existence of a polynomial oracle to compute code call conditions). ∎

The following example walks the reader through the detailed working of this algorithm on the motivating example pap introduced earlier on in this paper.

Example 6.7 We apply the above algorithm to the program (and agent state) of Example 6.1:

- Step 1 initialize X to $\emptyset$, i.e. to $\mathbf{S}^0_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$, while steps 2–8 iteratively apply the procedure which implements the operator $\mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$.
- At the first iteration, in step 3 $newX$ becomes $\mathbf{S}^1_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$ (shown in Example 6.3); since there are no deontic inconsistencies, the test in step 4 fails, and then we skip to step 6 which will assign $X := newX$, and then the cycle starts again.
- At the second iteration $newX$ becomes $\mathbf{S}^2_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$, then it is still inconsistency-free and different from X, so that we go on for another iteration.
- The third iteration is also the last one, since $\mathbf{S}^3_{\mathcal{PP},\mathcal{O}_\mathcal{S}} = \mathbf{S}^2_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$, and so we skip to the tests in steps 10–13.
- In our example, the preconditions of all actions are empty and then satisfied, there are not integrity constraint and then X is trivially integrity consistant, and eventually, there are no action constraints and then X is also action consistent. X is then returned as the (unique) reasonable status set of the program.

6.1 Agent Programs are Probabilistic Agent Programs

In this section, we show that the concept of an agent program defined by Eiter, Subrahmanian, and Pick (1999) is a special case of the framework defined here. Hence, paps generalize the Eiter et. al. semantics. Furthermore,

algorithm **Compute-*lfp*** may be used to compute reasonable status sets of positive agent programs. First, we show how agent programs may be captured as paps.

Definition 6.8 *Let $\mathcal{PP}$ be a probabilistic agent program, $\mathcal{PS}$ a probabilistic status set and $\mathcal{O}$ a probabilistic agent state. Assume further that each random variable contains exactly one object with probability 1. Then we can define the following mappings:*

Red$_1(\cdot)$**,** *which maps every probabilistic code call of the form $\langle\{o\},1\rangle$ to o:*

$$Red_1(\langle\{o_{\mathbf{RV}}\},1\rangle) = o.$$

Red$_2(\cdot)$**,** *which maps annotated code call conditions to code call conditions by simply removing the annotations and the conjunction strategy:*

$$Red_2(\chi : \langle[\mathsf{ai}_1, \mathsf{ai}_2], \otimes\rangle) = \chi.$$

We can easily extend $Red_2(\cdot)$ to a mapping from arbitrary conjunctions of annotated code calls to conjunctions of code calls.

Red$_3(\cdot)$**,** *which maps every probabilistic agent program to a non-probabilistic agent program: it clearly suffices to define $Red_3(\cdot)$ on probabilistic agent rules. This is done as follows*

$$Red_3(A \leftarrow \Gamma, L_1, \ldots, L_n) = A \leftarrow Red_2(\Gamma), L_1, \ldots, L_n.$$

Under the above assumptions, the following theorem holds.

Theorem 6.9 (Semantics of Agent Programs as an Instance of paps) *Suppose all random variables have the form*

$$\langle\{object_{\mathbf{RV}}\},1\rangle.$$

Then: ($\chi : \langle[\mathsf{ai}_1, \mathsf{ai}_2], \otimes\rangle$ is a ground annotated code call condition, $\mathcal{O}_{\mathcal{S}}$ an agent state)

Satisfaction: *the satisfaction relations coincide, i.e.*

$$\mathcal{O} \models^{[\mathsf{ai}_1,\mathsf{ai}_2]} \chi : \langle[\mathsf{ai}_1, \mathsf{ai}_2], \otimes\rangle \text{ if and only if } \mathcal{O} \models Red_2(\chi : \langle[\mathsf{ai}_1, \mathsf{ai}_2], \otimes\rangle).$$

App-Operators: *the App-Operators coincide, i.e.*

$$\boldsymbol{App}_{Red_3(\mathcal{PP}),\mathcal{O}_{\mathcal{S}}}(\mathcal{PS}) = \boldsymbol{App}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}(\mathcal{PS}),$$

where the operator on the left hand side is the one introduced in Definition A.4 on page 39.

Feasibility: *Feasible probabilistic status sets coincide with feasible status sets under our reductions, i.e. $\mathcal{PS}$ is a feasible probabilistic status set w.r.t. $\mathcal{PP}$ if and only if $\mathcal{PS}$ is a feasible status set w.r.t. $Red_3(\mathcal{PP})$.*

Rational: *Rational probabilistic status sets coincide with rational status sets under our reductions, i.e. $\mathcal{PS}$ is a rational probabilistic status set w.r.t. $\mathcal{PP}$ if and only if $\mathcal{PS}$ is a rational status set w.r.t. $Red_3(\mathcal{PP})$.*

Reasonable: *Reasonable probabilistic status sets coincide with reasonable status sets under our reductions, i.e. $\mathcal{PS}$ is a reasonable probabilistic status set w.r.t. $\mathcal{PP}$ if and only if $\mathcal{PS}$ is a reasonable status set w.r.t. $Red_3(\mathcal{PP})$.*

Computation of Status Sets: *The computations of probabilistic status sets given in Algorithms 6.1 on page 20 and 6.2 on page 22 for a* pap *$\mathcal{PP}$ reduce to the computation of status sets for $Red_3(\mathcal{PP})$.*

Proof: The first two statements are immediate. Feasibility requires checking conditions ($\mathcal{PS}1$)–($\mathcal{PS}4$), and therefore reduces to the first two statements. Rational and reasonable status sets are handled in a completely analogous manner.

That our algorithms reduce to the non-probabilistic case under our general assumption is trivial: the difference is only the satisfaction relation $\models^{[1,1]}$ which, by the first statement, coincides with $\models$. ∎

7 Probabilistic Agent Programs: Kripke Semantics

The definition of a feasible status set given in Section 5 makes several simplifying assumptions. First (see Definition 5.3), it assumes that an action can be executed only if its precondition is believed by the agent to be true in the agent state with probability 1. Second (see Definition 5.5), every action that is permitted must also have a precondition that is believed to be true with probability 1. In this section, we propose a Kripke-style semantics for agent programs that removes these conditions.

To do this, we will start by noting that in a probabilistic state $\mathcal{O}^p$, the agent returns a set of random variables for each code call. Every probabilistic state implicitly determines a set of (ordinary) states that are "compatible" with it. We use the notation $\mathbf{eval}(\mathtt{a}:f(\mathtt{d}_1,\ldots,\mathtt{d}_n),\mathcal{O})$ to denote the result of evaluating the code call $\mathtt{a}:f(\mathtt{d}_1,\ldots,\mathtt{d}_n)$ w.r.t. the state $\mathcal{O}$.

Definition 7.1 (Compatibility of State w.r.t. a Probabilistic State) *Let $\mathcal{O}^p$ be a probabilistic agent state. An (ordinary) agent state $\mathcal{O}$ is said to be* compatible *with $\mathcal{O}^p$ iff for every ground code call $\mathtt{a}:f(\mathtt{d}_1,\ldots,\mathtt{d}_n)$, it is the case that for every object $o \in \mathbf{eval}(\mathtt{a}:f(\mathtt{d}_1,\ldots,\mathtt{d}_n),\mathcal{O})$, there exists a random variable $(X,\wp) \in \mathbf{eval}(\mathtt{a}:f(\mathtt{d}_1,\ldots,\mathtt{d}_n),\mathcal{O}^p)$ such that $o \in X$ and $\wp(o) > 0$, and there is no other object $o' \in X$ such that $o' \in \mathbf{eval}(\mathtt{a}:f(\mathtt{d}_1,\ldots,\mathtt{d}_n),\mathcal{O})$.*

The following example illustrates this concept.

Example 7.2 *Consider a probabilistic agent state $\mathcal{O}^p$ with only two code calls* $\mathtt{surv}:\mathit{identify}(\mathtt{image1})$ *and* $\mathtt{surv}:\mathit{location}(\mathtt{image1})$*, which respectively return the random variables*

$$\langle\{t80, t72, t70\}, \{\langle t80, 0.3\rangle, \langle t72, 0.7\rangle, \langle t70, 0.0\rangle\}\rangle$$

and $\langle\{Loc2\}, \{\langle Loc2, 0.8\rangle\}\rangle$*. The agent states compatible w.r.t. $\mathcal{O}^p$ are described in the following table:*

State	***Vehicle***	***Location***
1	*none*	*none*
2	*t80*	*none*
3	*t72*	*none*

State	***Vehicle***	***Location***
4	*none*	*Loc2*
5	*t80*	*Loc2*
6	*t72*	*Loc2*

The object "t70" in the first random variable has a null probability, and hence it does not appear in any compatible agent state. In states 1–3, the location is unknown. In states 1 and 4, the vehicle in the image is unknown.

We use the notation $\mathsf{COS}(\mathcal{O}^p)$ to denote the set of all ordinary agent states that are compatible with a $\mathcal{O}^p$. We now define the notion of a probabilistic Kripke structure.

Definition 7.3 (Probabilistic Kripke Structure)
A probabilistic Kripke structure *is a pair $(\mathcal{S}, \wp)$ where $\mathcal{S}$ is a set of ordinary states, and $\wp : \mathcal{S} \to [0, 1]$ is a mapping such that $\sum_{\mathcal{O}\in\mathcal{S}} \wp(\mathcal{O}) = 1$.*

Definition 7.4 (Compatible Probabilistic Kripke Structure) *Let $\mathcal{O}^p$ be a probabilistic agent state. A coherent probabilistic Kripke structure* $(\mathsf{COS}(\mathcal{O}^p), \wp)$ *is said to be* compatible *with $\mathcal{O}^p$ iff for every ground code call $\mathtt{a}:f(\mathtt{d}_1, \ldots, \mathtt{d}_n)$, for every random variable $(X, \wp') \in \mathbf{eval}(\mathtt{a}:f(\mathtt{d}_1, \ldots, \mathtt{d}_n), \mathcal{O}^p)$, and for each object o, it is the case that:*

$$\sum_{o\in\mathbf{eval}(\mathtt{a}:f(\mathtt{d}_1,\ldots,\mathtt{d}_n),\mathcal{O})} \wp(\mathcal{O}) = \begin{cases} \wp'(o) & \text{if } o \in X \\ 0 & \text{otherwise} \end{cases}$$

By definition, two distinct objects from the same random variable cannot appear in the same compatible state. If such a random variable has a complete probability distribution, then in any compatible Kripke structure, the sum of the probabilities of the states containing one of its objects is equal to 1. This means that any (compatible) agent state containing no such objects will have a null probability, avoiding the intuitive inconsistency pointed in example 7.2.

Example 7.5 *Considering the situation in example 7.2 on the previous page, a probabilistic Kripke structure compatible with $\mathcal{O}^p$ is $\langle\mathsf{COS}(\mathcal{O}^p), \wp\rangle$, with the following probability distribution:*

State	*Probability*
1	*0*
2	*0.1*
3	*0.1*

State	*Probability*
4	*0*
5	*0.2*
6	*0.6*

The following result says that compatible Kripke structures always exist.

Theorem 7.6 (Existence of Compatible Probabilistic Kripke Structure) *Suppose $\mathcal{O}^p$ is a probabilistic agent state. Then there is at least one probabilistic Kripke structure which is compatible with it.*

Hence, given a probabilistic agent state $\mathcal{O}^p$, we use the notation $\mathsf{PKS}(\mathcal{O}^p)$ to denote the set of all probabilistic Kripke structures compatible with $\mathcal{O}^p$—this set is guaranteed to be nonempty by the preceding result. However, in almost all cases, $\mathsf{PKS}(\mathcal{O}^p)$ contains an infinite number of elements.

Theorem 7.7 (Existence of Infinitely Many Kripke Structures) *If a probabilistic state $\mathcal{O}^p$ contains at least two random variables (returned by the same code call or by two distinct ones) containing at least one object with associated probability $0 < p < 1$, then there exist an infinite number of probabilistic Kripke structures compatible with $\mathcal{O}^p$.*

We are now in a position to specify what it means to execute an action in a probabilistic Kripke structure.

Definition 7.8 (Action Execution) *Suppose $\mathcal{O}^p$ is a probabilistic agent state. A ground action α is said to be* possibly executable in $\mathcal{O}^p$ *iff there is at least one probabilistic Kripke structure $(\mathsf{COS}(\mathcal{O}^p), \wp) \in \mathsf{PKS}(\mathcal{O}^p)$, and an ordinary agent state $\mathcal{O}$ in $\mathsf{COS}(\mathcal{O}^p)$ in which the precondition of α is true such that $\wp(\mathcal{O}) > 0$. In this case, $\mathcal{O}$* witnesses *the executability of α.*

We say that α is executable with probability p *in $\mathcal{O}^p$ if, by definition,*

$$\mathsf{p} \quad = \quad min\{\wp(\mathcal{O}) \mid \mathcal{O} \in \mathsf{COS}(\mathcal{O}^p) \textit{ witnesses the executability of } \alpha\}.$$

The following example illustrates this definition.

Example 7.9 *Let us consider the probabilistic agent state of examples 7.2 and 7.5 and the following actions:*

$$\begin{array}{lll} \alpha_1: & Pre(\alpha_1) = & \mathbf{in}(\mathtt{t70}, \mathrm{surv}: \mathit{identify}(\mathtt{image1})) \\ \alpha_2: & Pre(\alpha_2) = & \mathbf{in}(\mathtt{X}, \mathrm{surv}: \mathit{location}(\mathtt{image1})) \ \& \ \mathtt{X} \neq \mathit{Loc2} \\ \alpha_3: & Pre(\alpha_3) = & \mathbf{in}(\mathtt{Loc2}, \mathrm{surv}: \mathit{location}(\mathtt{image1})) \ \& \\ & & \mathbf{in}(\mathtt{t80}, \mathrm{surv}: \mathit{identify}(\mathtt{image1})) \end{array}$$

As stated before, the object "t70" cannot appear in any compatible agent state, and hence the action α_1 is not possibly executable. On the other hand, the precondition of α_2 requires the existence of an object other than "Loc2", which is known to be the only one possibly returned by the corresponding code call,

and hence α_2 is not possibly executable either. Eventually, the precondition of α_3 requires the presence of both objects "Loc2" and "t80", which is true in the agent state number 5 described in the above examples. As this state has a non null probability, α_3 is possibly executable and the agent state witnesses its executability.

We are now ready to define the new probabilistic Kripke structure that results when an action is executed in it.

Definition 7.10 (Result of Action (θ, γ)-Execution) *Suppose $\mathcal{O}^p$ is a probabilistic agent state, $\alpha(\vec{X})$ is an action and γ is a ground substitution for all variables occurring in the precondition, add, and delete list of $\alpha(\vec{X})\theta$. Let $(\mathsf{COS}(\mathcal{O}^p), \wp)$ be a probabilistic Kripke structure that contains a witness $\mathcal{O}$ to the possible executability of $\alpha(\vec{X})\theta$. The* result of executing *action $\alpha(\vec{X})$ under substitutions θ, γ in probabilistic Kripke structure $\mathsf{PKS}(\mathcal{O}^p) = (\mathcal{S}, \wp)$ is a new Kripke structure $(\mathcal{S}', \wp')$ defined in the following way:*

1. *$\mathcal{S}' = \{map_{\alpha(\vec{X}),\theta,\gamma}(\mathcal{O}) \mid \mathcal{O} \in \mathcal{S}\}$ where map is defined as follows:*

$$map_{\alpha(\vec{X}),\theta,\gamma}(\mathcal{O}) = \begin{cases} apply(\alpha(\vec{X}), \theta, \gamma, \mathcal{O}) & \text{if } \mathcal{O} \in W \\ \mathcal{O} & \text{otherwise} \end{cases}$$

2. *$\wp'$ is defined as follows:*

$$\wp'(\mathcal{O}') = \sum\{\wp(\mathcal{O}) \mid \mathcal{O}' = map_{\alpha(\vec{X}),\theta,\gamma}(\mathcal{O})\}$$

In the above definitions, W is the set of all witnesses in $\mathcal{S}$ to the executability of $\alpha(\vec{X})\theta\gamma$.

The result of executing *action $\alpha(\vec{X})$ under substitutions θ, γ in a set $\mathcal{K}$ of probabilistic Kripke structures is $\{\mathcal{S}' \mid \mathcal{S}'$ is obtained by executing $\alpha(\vec{X})$ under substitutions θ, γ on $\mathcal{S} \in \mathcal{K}\}$.*

The definition causes the agent states in which α's precondition is true to change, while those in which α's precondition is false stay unchanged. The probability of each final state is the sum of the probabilities of the corresponding (old) states. This is illustrated by the following example.

Example 7.11 *Let us consider the compatible probabilistic Kripke structure in Example 7.5, and the action erase(X):*

Pre: in(`X`, surv : *identify*(`image1`))

Del: in(`X`, surv : *identify*(`image1`))

Add: $\emptyset$

The result of executing action erase(X) under substitutions $\{X/t80\}, \epsilon$, is the probabilistic Kripke structure $\langle \mathcal{S}', \wp' \rangle$, briefly described by the following table:

State	***Vehicle***	***Location***	***Probability***
a	*none*	*none*	*0.1*
b	*t72*	*none*	*0.1*
c	*none*	*Loc2*	*0.2*
d	*t72*	*Loc2*	*0.6*

i.e., the states 1 and 2 merge together yielding the new state "a" and their probabilities are summed. Similarly, states 4 and 5 yield the new state "c".

The following result states that our definitions are coherent.

Proposition 7.12 (Closure of Probabilistic Kripke Structures) *The result of (θ, γ)-execution of an action in a probabilistic Kripke structure is also a probabilistic Kripke structure.*

Proof: Let $\langle \mathcal{S}, \wp \rangle$ be the original Kripke structure, and $\langle \mathcal{S}', \wp' \rangle$ the result of executing the action. We just need to show that $\sum\{\wp'(\mathcal{O}') \,|\, \mathcal{O}' \in \mathcal{S}'\} = 1$. Using Definition 7.10 on the preceding page:

$$\sum_{\mathcal{O}' \in \mathcal{S}'} \wp'(\mathcal{O}') = \sum_{\mathcal{O}' \in \mathcal{S}'} \sum_{\mathcal{O}' = \mathrm{map}(\mathcal{O})} \wp(\mathcal{O}) = \sum_{\mathcal{O} \in \mathcal{S}} \wp(\mathcal{O}) = 1$$

∎

8 p-Feasible Status Sets

Probabilistic Feasible Status Sets prevent an action from being executed unless its precondition is known for sure to be true (which is exactly the intuitive reading of $\mathcal{O}^p \models^{[1,1]} Pre(\alpha)$, for a probabilistic agent state $\mathcal{O}^p$ and an action α). Analogously, an action constraint has to be checked only if its precondition is certainly true. Finally, state consistency requires that the execution of the actions does not corrupt the consistency of the original agent state, i.e. it has to lead to an agent state where the integrity constraints are with 100% probability,

In this section, we define the concept of p-feasibility, where p is a probability. p-feasibility weakens the above requirements to only requiring that preconditions are true with probability p (or higher).

Definition 8.1 (Operator p-App$_{\mathcal{PP},\mathcal{O}^p}(\mathcal{PS})$) *Suppose $\mathcal{PP}$ is a probabilistic agent program, $\mathcal{O}^p$ is a probabilistic agent state, and $\mathcal{PS}$ is a probabilistic status set. Then*

$$\textbf{\textit{p-App}}_{\mathcal{PP},\mathcal{O}^p}(\mathcal{PS}) = \{Op\,\alpha \mid \quad Op\,\alpha \text{ is the head of a ground instance of a rule } r \text{ in } \mathcal{PP} \text{ and:}$$

1. $B^+_{as}(r) \subseteq \mathcal{PS}$ *and* $\neg.B^-_{as}(r) \cap \mathcal{PS} = \emptyset$;
2. *For every annotated code call condition* $\chi : \langle [\mathsf{ai}_1, \mathsf{ai}_2], \otimes \rangle$ *in the body of* r, *it is the case that* $\mathcal{O}^p \models^{[\mathsf{ai}_1, \mathsf{ai}_2]} \chi : \langle [\mathsf{ai}_1, \mathsf{ai}_2], \otimes \rangle$;

3. *if* $\mathsf{Op} \in \{\mathbf{P}, \mathbf{O}, \mathbf{Do}\}$*, then the preconditions of* α *are true with probability* p *or greater, i.e.* $\mathcal{O}^p \models^{[p,1]} Pre(\alpha)$*;*

4. *for every action status atom of the form* $\mathsf{Op}\,\beta$ *in* $B^+_{as}(r)$ *such that* $\mathsf{Op} \in \{\mathbf{P}, \mathbf{O}, \mathbf{Do}\}$*, the preconditions of* β *are true with probability* p *or greater, i.e.* $\mathcal{O}^p \models^{[p,1]} Pre(\beta)$ *}*

The only difference between this definition and that of $\mathbf{App}_{\mathcal{PP},\mathcal{O}_S}(\mathcal{PS})$ is that the entailment $\models^{[1,1]}$ is replaced by the more general $\models^{[p,1]}$.

Definition 8.2 (Deontic and Action p-Consistency) *A probabilistic status set* $\mathcal{PS}$ *is* deontically p-consistent *with respect to a probabilistic agent state* $\mathcal{O}^p$ *if, by definition, it satisfies the following rules for any ground action* α*:*

- *If* $\mathbf{O}\alpha \in \mathcal{PS}$*, then* $\mathbf{W}\alpha \notin \mathcal{PS}$*.*
- *If* $\mathbf{P}\alpha \in \mathcal{PS}$*, then* $\mathbf{F}\alpha \notin \mathcal{PS}$*.*
- *If* $\mathbf{P}\alpha \in \mathcal{PS}$*, then the preconditions of* α *are true with probability* p *or greater, i.e.* $\mathcal{O}^p \models^{[p,1]} Pre(\alpha)$*.*

A probabilistic status set $\mathcal{PS}$ *is* action p-consistent *with respect to an agent state* $\mathcal{O}^p$ *iff for every action constraint of the form*

$$\{\alpha_1(\vec{X}_1), \ldots, \alpha_k(\vec{X}_k)\} \leftarrow \chi \tag{3}$$

either $\mathcal{O}^p \not\models^{[p,1]} \chi$ *or* $\{\alpha_1(\vec{X}_1), \ldots, \alpha_k(\vec{X}_k)\} \not\subseteq \mathcal{PS}$*.*

Generalizing probabilistic state consistency to p-probabilistic state consistency may be done in two ways.

Definition 8.3 (Probabilistic State p-Consistency) *A probabilistic status set* $\mathcal{PS}$ *is* weakly probabilistically state p-consistent *w.r.t. state* $\mathcal{O}^p$ *iff the new state,* $\mathcal{O}^{p'} = \boldsymbol{conc}(\mathbf{Do}\,(\mathcal{PS}), \mathcal{O}^p)$ *obtained after concurrently executing all actions of the form* $\mathbf{Do}\,\alpha \in \mathcal{PS}$ *satisfies all integrity constraints with probability greater than or equal to* p*, i.e. for every integrity constraint* $\psi \Rightarrow \chi$ *either* $\mathcal{O}^{p'} \not\models^{[p,1]} \psi$ *or* $\mathcal{O}^{p'} \models^{[p,1]} \chi$*.*
We say that a probabilistic status set $\mathcal{PS}$ *is* strongly probabilistically state p-consistent *w.r.t. state* $\mathcal{O}^p$ *iff the new state* $\mathcal{O}^{p'}$ *satisfies the following condition: if* $\mathcal{O}^p$ *satisfies the integrity constraints with probability* $\geq q$ *(*$q \in [0,1]$*) then also* $\mathcal{O}^{p'}$ *does so.*

These definitions induce two types of feasibility for arbitrary probabilities p.

Definition 8.4 (Weak (resp. Strong) p-Feasibility) *Let* $\mathcal{PP}$ *be an agent program and let* $\mathcal{O}^p$ *be an agent state. Then, a probabilistic status set* $\mathcal{PS}$ *is a* p-feasible probabilistic status set *for* $\mathcal{PP}$ *on* $\mathcal{O}^p$*, if the following conditions hold:*

(p-$\mathcal{PS}$ 1): ***p-App***$_{\mathcal{PP},\mathcal{O}^p}(\mathcal{PS}) \subseteq \mathcal{PS}$ *(closure under the program rules);*

(p-$\mathcal{PS}$ 2): $\mathcal{PS}$ *is deontically and action p-consistent (deontic and action p-consistency);*

(p-$\mathcal{PS}$ 3): $\mathcal{PS}$ *is action closed and deontically closed (deontic and action closure);*

(p-$\mathcal{PS}$ 4): $\mathcal{PS}$ *is weakly (resp. strong) state p-consistent (state p-consistency).*

Remark 8.1 If S is a p-feasible probabilistic status set for $\mathcal{PP}$ on $\mathcal{O}^p$ and $0 \leq q \leq p$, then S *is not always* q-feasible. Indeed, $[q,1] \supseteq [p,1]$, and then for any formula ϕ $\mathcal{O}^p \models^{[p,1]} \phi$ implies that $\mathcal{O}^p \models^{[q,1]} \phi$ (and analogously for $\not\models^{[p,1]}$ and $\not\models^{[q,1]}$). This means that all preconditions of actions, preconditions of action constraints and integrity constraints which are verified for p are also verified for q.
The problem is that **p-App**$_{\mathcal{PP},\mathcal{O}^p}(\mathcal{PS})$ is anti-monotonic w.r.t. p, as a smaller value for p may allow a larger set of rules to be firable. Then the closure under the program rules is not guaranteed any more.

The following example illustrates this point.

Example 8.5 Consider the following trivial program:

$$\mathbf{Do}\,\alpha \leftarrow .$$

where $Pre(\alpha) = \mathbf{in}(\mathrm{a}, \mathrm{d}:f())$, and $\mathbf{eval}(\mathrm{d}:f(), \mathcal{O}^p) = \{\langle\{a\}, \langle 0.7\rangle\rangle\}$. Suppose $\mathcal{PS} = \emptyset$. Conditions p-$\mathcal{PS}$ 2–4 are true for any value of p. Note that 0.8-**App**$_{\mathcal{PP},\mathcal{O}_S}(\mathcal{PS}) = \emptyset$ nd hence, $\mathcal{PS}$ is 0.8-feasible. In contrast, we see that 0.6-**App**$_{\mathcal{PP},\mathcal{O}_S}(\mathcal{PS}) = \{\mathbf{Do}\,\alpha\} \not\subseteq \mathcal{PS}$, and hence $\mathcal{PS}$ is not 0.6-feasible.

We can easily see that probabilistic feasibility is a particular case p-feasibility:

Proposition 8.6 *Let $\mathcal{PP}$ be a probabilistic agent program and let $\mathcal{O}^p$ be a consistent (or equivalently 1-consistent) agent state. Then, a probabilistic status set $\mathcal{PS}$ is a feasible probabilistic status set if and only if it is weakly 1-feasible if and only if it is strongly 1-feasible.*

Proof: All definitions for weak p-feasibility trivially coincide with those for feasibility if p=1. The distinction between weak and strong p-feasibility is just in the definition of p-consistency, and we can easily see that for p=1 they coincide, since a probability *greater or equal to* 1 cannot be but equal to 1. ∎

8.1 p-Rational and p-Reasonable Status Sets

The notions of Rational and Reasonable Status Sets can be straightforwardly extended to those of p-Rational and p-Reasonable status sets.

Definition 8.7 (p-Rational Status Set) *A probabilistic status set $\mathcal{PS}$ is a* p-rational probabilistic status set*, if $\mathcal{PS}$ is a p-feasible probabilistic status set and there exists no probabilistic status set $\mathcal{PS}' \subset \mathcal{PS}$ satisfying conditions $(p-\mathcal{PS}1)$–$(p-\mathcal{PS}3)$ of a p-feasible probabilistic status set.*

Obviously, in the case that $\mathcal{IC} = \emptyset$ (i.e., there are no integrity constraints) p-rational status sets are simply inclusion-minimal feasible status sets.

Definition 8.8 (p-Reasonable Probabilistic Status Set) *Let $\mathcal{PP}$ be a probabilistic agent program, let $\mathcal{O}_\mathcal{S}$ be an agent state, and let $\mathcal{PS}$ be a probabilistic status set.*

1. *If $\mathcal{PP}$ is a positive probabilistic agent program, then $\mathcal{PS}$ is a* p-reasonable probabilistic status set *for $\mathcal{PP}$ on $\mathcal{O}_\mathcal{S}$, if, by definition, $\mathcal{PS}$ is a p-rational probabilistic status set for $\mathcal{PP}$ on $\mathcal{O}_\mathcal{S}$.*
2. *Exploiting the definition of $red^{\mathcal{PS}}(\mathcal{PP}, \mathcal{O}_\mathcal{S})$ (see Definition 5.13 on page 18), $\mathcal{PS}$ is a* p-reasonable probabilistic status set *for $\mathcal{PP}$ w.r.t. $\mathcal{O}_\mathcal{S}$, if it is a p-reasonable probabilistic status set of the program $red^{\mathcal{PS}}(\mathcal{PP}, \mathcal{O}_\mathcal{S})$ with respect to $\mathcal{O}_\mathcal{S}$.*

It is easy to verify that all p-reasonable probabilistic status sets are p-rational probabilistic status sets:

Proposition 8.9 (p-Reasonable Status Sets are p-Rational) *Let $\mathcal{PP}$ be a probabilistic agent program and $\mathcal{O}_\mathcal{S}$ an agent state. Then, every p-reasonable probabilistic status set of $\mathcal{PP}$ on $\mathcal{O}_\mathcal{S}$ is a p-rational probabilistic status set of $\mathcal{PP}$ on $\mathcal{O}_\mathcal{S}$.*

Proof: Identical to the proof of Proposition 5.16 on page 19. ∎

As in Section 6 on page 20, we can define a fixpoint operator and build an algorithm on its top to compute p-reasonable status sets for positive programs.

Definition 8.10 (Operator $\mathrm{p}-\mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$)

$$\mathrm{p}-\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS}) = \mathbf{A\text{-}Cl}(\mathrm{p}-\mathbf{App}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}(\mathcal{PS})),$$

Operator $\mathrm{p}-\mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$ can be computed by an algorithm identical to Algorithm 6.1 on page 20, but for step 4, where the entailment $\models^{[1,1]}$ has to be replaced by $\models^{[p,1]}$. $\mathrm{p}-\mathbf{S}_{\mathcal{PP},\mathcal{O}_\mathcal{S}}$ is monotonic and continuous and has a unique least fixpoint.

Lemma 8.11 (Monotonicity and Continuity of $\mathrm{p}{-}\mathbf{S}_{\mathcal{PP},\mathcal{O}_S}$) *Suppose $\mathcal{PP}$ is a positive* pap*. Then the operator $\mathrm{p}{-}\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_S}$ is monotone and continuous, i.e.*

1. $\mathcal{PS}_1 \subseteq \mathcal{PS}_2 \Rightarrow p - \boldsymbol{S}_{\mathcal{PP},\mathcal{O}_S}(\mathcal{PS}_1) \subseteq p - \boldsymbol{S}_{\mathcal{PP},\mathcal{O}_S}(\mathcal{PS}_2)$,
2. $\mathrm{p}{-}\boldsymbol{S}^{\omega}_{\mathcal{PP},\mathcal{O}_S}$ *is a fixpoint of* $\mathrm{p}{-}\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_S}$*. Moreover, it is the least fixpoint of* $\mathrm{p}{-}\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_S}$*. (We assume the iterations of* $\mathrm{p}{-}\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_S}$ *are defined in the same way as the iterations of* $\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_S}$*).*

The following result now follows immediately and has a proof similar to that of Theorem 6.4 on page 22.

Theorem 8.12 (p-Rational Probab. Status Sets as Least Fixpoints) *Let $\mathcal{PP}$ be a positive probabilistic agent program, and let $\mathcal{O}_S$ be an agent state. Then, $\mathcal{PS}$ is a p-rational probabilistic status set of $\mathcal{PP}$ on $\mathcal{O}_S$, if and only if $\mathcal{PS} = \mathit{lfp}(\mathrm{p}{-}\boldsymbol{S}_{\mathcal{PP},\mathcal{O}_S})$ and $\mathcal{PS}$ is a p-feasible probabilistic status set.*

Uniqueness of the p-reasonable status set (if it exists) holds too, and then we can compute it by Algorithm 6.2 on page 22, replacing–as usual—the entailment $\models^{[1,1]}$ with $\models^{[p,1]}$.

Unfortunately, the resulting algorithm is not polynomial because of the integrity constraint check in steps (11)–(12). This will be discussed in detail in Section 8.2 below. However, when no integrity constraints are present, the algorithm is still polynomial.

Theorem 8.13 (Polynomial Data Complexity) *The problem of computing p-reasonable probabilistic status sets of positive* paps *without Integrity Constraints (i.e., $\mathcal{IC} = \emptyset$) has polynomial data-complexity.*

8.2 Checking p-Consistency of Integrity Constraints

In this section, we provide an algorithm to check p-consistency of an integrity constraint IC after an action have been executed in state $\mathcal{O}^p$, leading to a new state $\mathcal{O}^{p\prime}$ assuming that all integrity constraints are true in the original state $\mathcal{O}^p$. Suppose $O_1, \ldots, O_N$ are all states compatible with $\mathcal{O}^p$ while $O'_1, \ldots, O'_{N'}$ are those compatible with the new state $\mathcal{O}^{p'}$. Let p_i, p'_i denote the probabilities of O_i, O'_i respectively. Then consider the following system of constraints:

minimize $\Sigma_{O'_i \models IC} p'_i$ **such that:**

(K)	$\sum_{i=1}^{N} p_i = 1$
(CK)	$\forall i \in \{1, \ldots, E\}. \sum_{j:Obj_i \in O_j} p_j = p(Obj_i)$
(IC)	$\forall IC \in \mathcal{IC}. 1 \geq \sum_{O_i:O_i \models IC} p_i \geq p$
(K → K')	$\forall i \in \{1, \ldots, N'\}. p'_i = \sum_{O_j:O_j \xrightarrow{\alpha} O'_i} p_j$
(IG)	$\forall i \in \{1, \ldots, N\}. \max\{0, \sum_{Obj_k \in O_i} p(Obj_k) + 1 - \lvert O_i \rvert\} \leq$ $\leq p_i \leq \min_{Obj_k \in O_i}\{p(Obj_k)\}$

The objective function captures the probability of IC being true in the new Kripke structure. **(K)** and **(CK)** define any arbitrary compatible Kripke structure over N states w.r.t. $\mathcal{O}^p$ (which cointains E objects), **(IC)** expresses the fact that our actual state has to be p-consistent, while **(K→ K')** defines the Kripke structure obtained after the execution of action α. Eventually, **(IG)** gives an upper and a lower bound to the probability of worlds (it extends the Bool expression for conjunction of events of unknown inter-relation). It is easy to see that a straightforward implementation of this algorithm requires exponential time and space.

9 Related Work

There has been an incredible amount of work on uncertainty in knowledge based and database systems (Shafer and Peal 1990). However, almost all this work assumes that we are reasoning with logic or with Bayesian nets (Koller 1998) and most work proceeds under strong assumptions about the relationships between events (e.g. most Bayesian approaches assume conditional independence between events, while other approaches such as (Fagin, Halpern, and Megiddo 1990; Ng and Subrahmanian 1993b) assume that we have no knowledge of the dependencies between events).

This paper introduces techniques to allow an agent developer to encode different assumptions about the relationships between events, when writing probabilistic agent programs. The idea of conjunction strategies to facilitate this was first introduced in the ProbView system (Lakshmanan, Leone, Ross, and Subrahmanian 1997) in an attempt to allow users querying probabilistic relational databases to express in their query, their knowledge of the dependencies between events. Later, (Dekhtyar and Subrahmanian 1997) extended the use of conjunction and disjunction strategies to the case of logic programs. In this paper, the idea of conjunction strategies are applied in the context of deontic-logic based agent programs. We are not aware of any extant work on allowing flexible dependency assumptions in the context of logics and actions.

Research on epistemic logic (e.g., (Morgenstern 1988; Moore 1985; Kraus and Lehmann 1988)) enables reasoning about what is known and is not known at a given time. However, epistemic logics have not been used as a representation in decision making and in automated planning systems, perhaps, because the richness of these languages makes efficient reasoning very difficult. In contrast, our framework has polynomial data complexity.

Halpern and Tuttle (1992) study the semantics of reasoning about distributed systems when uncertainty is present. They develop a logic where a process has knowledge about the probability of events which facilitates decision-making by the process. We, on the other hand, consider probabilistic states, and as argued in (Dix, Subrahmanian, and Pick 2000) this also allows us to reason about probabilistic beliefs, i.e. probabilities are assigned to the agents' beliefs about events, rather than to the events themselves. That is, in Halpern's work (Halpern and Tuttle 1992), the beliefs of the agent are CERTAIN, but in

our framework, the beliefs of the agent may themselves be uncertain (with the phenomenon when they are certain being a special case of our framework).

Poole (1997) presented a framework that allows a natural specification of multi-agent decision problems. It extends logic with a new way to handle and think about non-determinism and uncertainty in terms of independent choices made by various agents, including nature and a logic program that gives the consequence of choices. It has general independence assumption. This work is more expressive than ours, but its generality leads to complexity problems and to difficulties in using the framework.

Haddawy (1991) developed a logic that allows to write sentences that describe uncertainty in the state of the world, uncertainty of action effects, combine possibility and chance, distinguish between truth and chance and express information about probability distributions. He uses model theoretic semantics and demonstrates how his logic can be used to specification various reasoning and planning problems. The main purpose of the specification is to prove correctness, and not for programming of agents. Kushmerick, Hanks, and Weld (1995) model uncertainty about the true state of the world with a probability distribution over the state space. Actions have uncertain effects, and each of these effects is also modeled with a probability distribution. They seek plans whose probability of success exceeds the threshold. They describe BURIDAN, an implemented algorithm for probabilistic planning. In contrast, we focus on programming agents, rather than on how agents will construct plans. Other researchers extended Kushmerick et al.'s model to increase the efficiency of the planning (Haddawy, Doan, and Goodwin 1996) or to more realistic domains (Doan 1996). Thiébaux, Hertzberg, Shoaff, and Schneider (1995) developed a framework for anytime generation of plans under incomplete and ambiguous knowledge and actions with alternative and context dependent effects.

Kaelbling, Littman, and Cassandra (1998) propose using *partially observable Markov decisionprocesses* (POMDPs) for planning under uncertainty. Similar to BURIDAN they use a probability distributions over states to express uncertainty about the situation of the agent. They also consider the problem of non-deterministic actions and getting feedback from the environment which we mentioned only briefly.

10 Conclusions

Agents are programs that autonomously react to changes in their environment by taking appropriate actions. In (Eiter, Subrahmanian, and Pick 1999), the authors have proposed a framework within which agents may be built on top of an existing body of legacy code, and/or on top of specialized data structures appropriate for the intended functionality of the agent being built.

However, there are an increasing number of applications where agents are uncertain about what is true in their state (or environment). Such situations occur all the time in image identification programs, in programs that predict future events (such as battlefield events, stock market events, etc.), and in

scenarios where an agent a attempts to predict what an agent b will do.

In this paper, we first introduce the concept of a probabilistic code call, which is a mechanism to describe uncertain application program interfaces for arbitrary functions over arbitrary data types. Based on this concept, we define probabilistic agent programs — sets of rules that encode the operating principles of an agent. Such rules encode the *probabilistic* conditions under which an agent is obliged to take some actions, permitted to take some actions and/or forbidden to take some actions.

We then provide two broad classes of semantics for such "probabilistic agents." In the first class of semantics, actions that are permitted, obligatory or done, must have preconditions that are true with 100% probability in the current agent state. In the second class of semantics (which use probabilistic variants of Kripke structures), the actions that are permitted, obligatory or done, must have preconditions that are true with at least a given probability. This latter class of semantics allows reasoning by cases. We provide complexity arguments showing that though the second family of semantics is perhaps epistemologically more appealing than the first, the second family of semantics is also computationally more complex.

Finally, the paper includes algorithms to compute the semantics of probabilistic agent programs, as long as such programs are negation free.

Future work on probabilistic agent programs will focus on computing the semantics of paps that contain negation. A detailed study of computational complexity is also envisaged. We are also interested in identifying polynomially computable fragments of paps and implementing them on top of the current *IMPACT* implementation. Last, but not least, *IMPACT* has been used in a battlefield monitoring application where there is considerable uncertainty in predicting tactical enemy movements. We hope to build an application of paps addressing this problem, once the implementation of paps is complete.

References

Apt, K. (1990). Logic Programming. In J. van Leeuwen (Ed.), *Handbook of Theoretical Computer Science*, Volume B, Chapter 10, pp. 493–574. Elsevier Science Publishers B.V. (North-Holland).

Arisha, K., F. Ozcan, R. Ross, V. S. Subrahmanian, T. Eiter, and S. Kraus (1999, March/April). IMPACT: A Platform for Collaborating Agents. *IEEE Intelligent Systems 14*, 64–72.

Boole, G. (1854). *The Laws of Thought.* Macmillan, London.

Cattell, R. G. G., et al. (Ed.) (1997). *The Object Database Standard: ODMG-93.* Morgan Kaufmann.

Dekhtyar, A. and V. S. Subrahmanian (1997). Hybrid Probabilistic Logic Programs. In L. Naish (Ed.), *Proceedings of the 14th International Conference on Logic Programing*, Leuven, Belgium, pp. 391–405. MIT Press.

Extended version accepted for publication in Journal of Logic Programming, |http://www.cs.umd.edu/TRs/authors/Alex_Dekhtyar.html—.

Dix, J., V. S. Subrahmanian, and G. Pick (2000). Meta Agent Programs. accepted for publication in Journal of Logic Programming, to appear 2000.

Doan, A. (1996). Modeling Probabilistic Actions for Practical Decision-Theoretic Planning. In *Proceedings of the Third International Concerence on Artificial-Intelligence Planning Systems*, Edinburgh, Scotland, UK.

Eiter, T., V. Subrahmanian, and G. Pick (1999). Heterogeneous Active Agents, I: Semantics. *Artificial Intelligence 108*(1-2), 179–255.

Eiter, T., V. Subrahmanian, and T. Rogers (1999, May). Heterogeneous Active Agents, III: Polynomially Implementable Agents. Technical Report INFSYS RR-1843-99-07, Institut für Informationssysteme, Technische Universität Wien, A-1040 Vienna, Austria.

Eiter, T. and V. S. Subrahmanian (1999). Heterogeneous Active Agents, II: Algorithms and Complexity. *Artificial Intelligence 108*(1-2), 257–307.

Fagin, R., J. Y. Halpern, and N. Megiddo (1990, July/August). A logic for reasoning about probabilities. *Information and Computation 87*(1/2), 78–128.

Gelfond, M. and V. Lifschitz (1988). The Stable Model Semantics for Logic Programming. In *Logic Programming: Proceedings Fifth International Conference and Symposium*, Cambridge, Massachusetts, pp. 1070–1080. MIT Press.

Haddawy, P. (1991). *Representing Plans under Uncertainty: A Logic of Time, Chance and Action*. Ph. D. thesis, University of Illinois. Technical Report UIUCDCS-R-91-1719.

Haddawy, P., A. Doan, and R. Goodwin (1996). Efficient Decision-Theoretic Planning: Techniques and Empirical Analysis. In *Proceedings of the Third International Concerence on Artificial-Intelligence Planning Systems*, Edinburgh, Scotland, UK.

Halpern, J. Y. and M. Tuttle (1992). Knowledge, Probability and Adversaries. Technical report, IBM. IBM Research Report.

Kaelbling, L. P., M. L. Littman, and A. R. Cassandra (1998). Planning and Acting in Partially Observable Stochastic Domains. *Artificial Intelligence 101*, 99–134.

Kifer, M. and V. S. Subrahmanian (1992). Theory of Generalized Annotated Logic Programming and its Applications. *Journal of Logic Programming 12*(4), 335–368.

Koller, D. (1998). Structured Probabilistic Models: Bayesian Networks and Beyond. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence, AAAI'98*, Madison, Wisconsin. AAAI Press/MIT Press. |http://robotics.Stanford.EDU/∼koller/Invited98—.

Kraus, S. and D. Lehmann (1988). Knowledge, Belief and Time. *Theoretical Computer Science 58*, 155–174.

Kushmerick, N., S. Hanks, and D. Weld (1995). An Algorithm for probabilistic planning. *Artificial Intelligence 76*(1-2), 239–286.

Lakshmanan, V. S., N. Leone, R. Ross, and V. S. Subrahmanian (1997, September). ProbView: A Flexible Probabilistic Database System. *ACM Transactions on Database Systems 22*(3), 419–469.

Lloyd, J. (1984, 1987). *Foundations of Logic Programming.* Berlin, Germany: Springer-Verlag.

Moore, R. (1985). A Formal theory of Knowledge and Action. In J. Hobbs and R. Moore (Eds.), *Formal Theories of the Commonsesnse World.* Norwood, N.J.: ABLEX publishing.

Morgenstern, L. (1988). *Foundations of a Logic of Knowledge, Action, and Communication.* Ph. D. thesis, New York University.

Ng, R. and V. S. Subrahmanian (1993a). A Semantical Framework for Supporting Subjective and Conditional Probabilities in Deductive Databases. *Journal of Automated Reasoning 10*(2), 191–235.

Ng, R. and V. S. Subrahmanian (1993b). Probabilistic Logic Programming. *Information and Computation 101*(2), 150–201.

Poole, D. (1997). The independent choice logic for modelling multiple agents under uncertainty. *Artificial Intelligence 94*(1-2), 7–56.

Ross, S. (1997). *A First Course in Probability.* Prentice-Hall.

Shafer, G. and J. Peal (Eds.) (1990). *Readings in uncertain reasoning.* Morgan Kaufmann.

Siegal, J. (1996). *CORBA Fundementals and Programming.* New York: John Wiley & Sons.

Subrahmanian, V. S. (1987, September). On the Semantics of Quantitative Logic Programs. In *Proceedings of the 4th IEEE Symposium on Logic Programming*, pp. 173–182. Computer Society Press.

Thiébaux, S., J. Hertzberg, W. Shoaff, and M. Schneider (1995). A stochastic model of actions and plans for anytime planning under uncertainty. *International Journal for Intelligent Systems 10*(2).

A Agent Programs without Uncertainty

The following definitions are taken from (Eiter, Subrahmanian, and Pick 1999).

A.1 Feasible, Rational and Reasonable Semantics

Definition A.1 (Status Set) *A* status set *is any set S of ground action status atoms over $\mathcal{S}$. For any operator $Op \in \{\mathbf{P}, \mathbf{Do}, \mathbf{F}, \mathbf{O}, \mathbf{W}\}$, we denote by $Op(S)$ the set $Op(S) = \{\alpha \mid Op(\alpha) \in S\}$.*

Definition A.2 (Deontic and Action Consistency) *A status set S is called* deontically consistent*, if, by definition, it satisfies the following rules for any ground action α:*

- *If $\mathbf{O}\alpha \in S$, then $\mathbf{W}\alpha \notin S$*
- *If $\mathbf{P}\alpha \in S$, then $\mathbf{F}\alpha \notin S$*
- *If $\mathbf{P}\alpha \in S$, then $\mathcal{O}_S \models \exists^* Pre(\alpha)$, where $\exists^* Pre(\alpha)$ denotes the existential closure of $Pre(\alpha)$, i.e., all free variables in $Pre(\alpha)$ are governed by an existential quantifier. This condition means that the action α is in fact executable in the state $\mathcal{O}_S$.*

A status set S is called action consistent*, if $S, \mathcal{O}_S \models \mathcal{AC}$ holds.*

Besides consistency, we also wish that the presence of certain atoms in S entails the presence of other atoms in S. For example, if $\mathbf{O}\alpha$ is in S, then we expect that $\mathbf{P}\alpha$ is also in S, and if $\mathbf{O}\alpha$ is in S, then we would like to have $\mathbf{Do}\,\alpha$ in S. This is captured by the concept of deontic and action closure.

Definition A.3 (Deontic and Action Closure) *The* deontic closure *of a status S, denoted $\mathbf{D}\text{-}\mathbf{Cl}(S)$, is the closure of S under the rule*

If $\mathbf{O}\alpha \in S$, then $\mathbf{P}\alpha \in S$

where α is any ground action. We say that S is deontically closed*, if $S = \mathbf{D}\text{-}\mathbf{Cl}(S)$ holds.*

The action closure *of a status set S, denoted $\mathbf{A}\text{-}\mathbf{Cl}(S)$, is the closure of S under the rules*

If $\mathbf{O}\alpha \in S$, then $\mathbf{Do}\,\alpha \in S$

If $\mathbf{Do}\,\alpha \in S$, then $\mathbf{P}\alpha \in S$

where α is any ground action. We say that a status S is action-closed, if $S = \mathbf{A}\text{-}\mathbf{Cl}(S)$ holds.

The following straightforward results shows that status sets that are action-closed are also deontically closed, i.e.

Definition A.4 (Operator $\mathbf{App}_{\mathcal{P},\mathcal{O}_S}(S)$) *Suppose $\mathcal{P}$ is an agent program, and $\mathcal{O}_S$ is an agent state. Then, $\boldsymbol{App}_{\mathcal{P},\mathcal{O}_S}(S)$ is defined to be the set of all ground action status atoms A such that there exists a rule in P having a ground instance of the form $r : A \leftarrow L_1, \ldots, L_n$ such that*

1. *$B^+_{as}(r) \subseteq S$ and $\neg.B^-_{as}(r) \cap S = \emptyset$, and*
2. *every code call $\chi \in B^+_{cc}(r)$ succeeds in $\mathcal{O}_S$, and*
3. *every code call $\chi \in \neg.B^-_{cc}(r)$ does not succeed in $\mathcal{O}_S$, and*

4. *for every atom* $\mathsf{Op}(\alpha) \in B^+(r) \cup \{A\}$ *such that* $\mathsf{Op} \in \{\mathbf{P}, \mathbf{O}, \mathbf{Do}\}$*, the action* α *is executable in state* $\mathcal{O}_\mathcal{S}$*.*

Note that part (4) of the above definition only applies to the "positive" modes $\mathbf{P}, \mathbf{O}, \mathbf{Do}$. It does not apply to atoms of the form $\mathbf{F}\alpha$ as such actions are not executed, nor does it apply to atoms of the form $\mathbf{W}\alpha$, because execution of an action might be (vacuously) waived, if its prerequisites are not fulfilled.

Our approach is to base the semantics of agent programs on consistent and closed status sets. However, we have to take into account the rules of the program as well as integrity constraints. This leads us to the notion of a feasible status set.

Definition A.5 (Feasible Status Set) *Let* $\mathcal{P}$ *be an agent program and let* $\mathcal{O}_\mathcal{S}$ *be an agent state. Then, a status set* S *is a* feasible status set *for* $\mathcal{P}$ *on* $\mathcal{O}_\mathcal{S}$*, if the following conditions hold:*

(S1): *(closure under the program rules)* $\boldsymbol{App}_{\mathcal{P},\mathcal{O}_\mathcal{S}}(S) \subseteq S$*;*

(S2) *(deontic and action consistency)* S *is deontically and action consistent;*

(S3) *(deontic and action closure)* S *is action closed and deontically closed;*

(S4) *(state consistency)* $\mathcal{O}'_\mathcal{S} \models \mathcal{IC}$*, where* $\mathcal{O}'_\mathcal{S} = apply(\mathbf{Do}\,(S), \mathcal{O}_\mathcal{S})$ *is the state which results after taking all actions in* $\mathbf{Do}\,(S)$ *on the state* $\mathcal{O}_\mathcal{S}$*.*

Definition A.6 (Groundedness; Rational Status Set) *A status set* S *is* grounded*, if there exists no status set* $S' \neq S$ *such that* $S' \subseteq S$ *and* S' *satisfies conditions* $(S1)$–$(S3)$ *of a feasible status set.*

A status set S *is a* rational status set*, if* S *is a feasible status set and* S *is grounded.*

Definition A.7 (Reasonable Status Set) *Let* $\mathcal{P}$ *be an agent program, let* $\mathcal{O}_\mathcal{S}$ *be an agent state, and let* S *be a status set.*

1. *If* $\mathcal{P}$ *is a positive agent program, then* S *is a* reasonable status set *for* $\mathcal{P}$ *on* $\mathcal{O}_\mathcal{S}$*, if and only if* S *is a rational status set for* $\mathcal{P}$ *on* $\mathcal{O}_\mathcal{S}$*.*

2. *The reduct of* $\mathcal{P}$ *w.r.t.* S *and* $\mathcal{O}_\mathcal{S}$*, denoted by* $red^S(\mathcal{P}, \mathcal{O}_\mathcal{S})$*, is the program which is obtained from the ground instances of the rules in* $\mathcal{P}$ *over* $\mathcal{O}_\mathcal{S}$ *as follows.*

   (a) *First, remove every rule* r *such that* $B^-_{as}(r) \cap S \neq \emptyset$*;*

   (b) *Remove all atoms in* $B^-_{as}(r)$ *from the remaining rules.*

   Then S *is a* reasonable status set *for* $\mathcal{P}$ *w.r.t.* $\mathcal{O}_\mathcal{S}$*, if it is a reasonable status set of the program* $red^S(\mathcal{P}, \mathcal{O}_\mathcal{S})$ *with respect to* $\mathcal{O}_\mathcal{S}$*.*

B Proofs of Theorems

Proof: (of Proposition 5.15)

1. Suppose $\mathbf{Do}\,\alpha \in \mathcal{PS}$. Then, as $\mathcal{PS}$ is feasible, we know that $\mathcal{PS} = \mathbf{A\text{-}Cl}(\mathcal{PS})$, and hence $\mathbf{P}\alpha \in \mathcal{PS}$. As $\mathcal{PS}$ is feasible, and hence deontically consistent, the third condition of deontic consistency specifies that α's precondition is true in state $\mathcal{O}_{\mathcal{S}}$.

2. This follows immediately because as $\mathcal{PS}$ is feasible, we have $\mathcal{PS} = \mathbf{A\text{-}Cl}(\mathcal{PS})$. The second condition defining $\mathbf{A\text{-}Cl}(\mathcal{PS})$, when written in contrapositive form, states that $\mathbf{P}\alpha \notin \mathcal{PS}$ implies that $\mathbf{Do}\,\alpha \notin \mathcal{PS}$.

3. As $\mathcal{PS}$ is feasible, $\mathcal{PS} = \mathbf{A\text{-}Cl}(\mathcal{PS})$. The first condition specifying $\mathbf{A\text{-}Cl}(\mathcal{PS})$ allows us to infer that $\mathbf{O}\alpha \in \mathcal{PS}$ implies that $\mathbf{Do}\,\alpha \in \mathcal{PS}$. The result follows immediately from part (1) of this proposition.

4. From the above argument, as $\mathcal{PS} = \mathbf{A\text{-}Cl}(\mathcal{PS})$, we can conclude that $\mathbf{O}\alpha \in \mathcal{PS}$ implies that $\mathbf{P}\alpha \in \mathcal{PS}$. By the deontic consistency requirement, $\mathbf{F}\alpha \notin \mathcal{PS}$.

Proof: (of Theorem 5.16)
In order to show that a reasonable probabilistic status set $\mathcal{PS}$ of $\mathcal{PP}$ is a rational status of $\mathcal{PP}$, we have to verify (1) that $\mathcal{PS}$ is a feasible probabilistic status set and (2) that $\mathcal{PS}$ is grounded.

Since $\mathcal{PS}$ is a reasonable probabilistic status set of $\mathcal{PP}$, it is a rational probabilistic status set of $\mathcal{PP}' = red^{\mathcal{PS}}(\mathcal{PP}, \mathcal{O}_{\mathcal{S}})$, i.e., a feasible and grounded probabilistic status set of $\mathcal{PP}'$. Since the conditions ($\mathcal{PS}2$)–($\mathcal{PS}4$) of the definition of feasible probabilistic status set depend only on $\mathcal{PS}$ and $\mathcal{O}_{\mathcal{S}}$ but not on the program, this means that for showing (1) it remains to check that ($\mathcal{PS}1$) (closure under the program rules) is satisfied.

Let thus r be a ground instance of a rule from $\mathcal{PP}$. Suppose the body $B(r)$ of r satisfies the conditions 1.–4. of ($\mathcal{PS}1$). Then, by the definition of $red^{\mathcal{PS}}(\mathcal{PP}, \mathcal{O}_{\mathcal{S}})$, we have that the reduct of the rule r, obtained by removing all literals of $B^-_{as}(r)$ from the body, is in $\mathcal{PP}'$. Since $\mathcal{PS}$ is closed under the rules of $\mathcal{PP}'$, we have $H(r) \in \mathcal{PS}$. Thus, $\mathcal{PS}$ is closed under the rules of $\mathcal{PP}$, and hence ($\mathcal{PS}1$) is satisfied. As a consequence, (1) holds.

For (2), we suppose $\mathcal{PS}$ is not grounded, i.e., that some smaller $\mathcal{PS}' \subset \mathcal{PS}$ satisfies ($\mathcal{PS}1$)–($\mathcal{PS}3$) for $\mathcal{PP}$, and derive a contradiction. If $\mathcal{PS}'$ satisfies ($\mathcal{PS}1$) for $\mathcal{PP}$, then $\mathcal{PS}'$ satisfies ($\mathcal{PS}1$) for $\mathcal{PP}'$. For, if r is a rule from $\mathcal{PP}'$ such that 1.–4. of ($\mathcal{PS}1$) hold for $\mathcal{PS}'$, then there is a ground rule r' of $\mathcal{PP}$ such that r is obtained from r' in the construction of $red^{\mathcal{PS}}(\mathcal{PP}, \mathcal{O}_{\mathcal{S}})$ and, as easily seen, 1.–4. of ($\mathcal{PS}1$) hold for $\mathcal{PS}'$. Since $\mathcal{PS}'$ satisfies ($\mathcal{PS}1$) for $\mathcal{PP}$, we have $H(r) \in \mathcal{PS}'$. It follows that $\mathcal{PS}'$ satisfies ($\mathcal{PS}1$) for $\mathcal{PP}'$. Furthermore, since ($\mathcal{PS}2$) and ($\mathcal{PS}3$) do no depend on the program, also ($\mathcal{PS}2$) and ($\mathcal{PS}3$) are satisfied for $\mathcal{PS}'$ w.r.t. $\mathcal{PP}'$. This means that $\mathcal{PS}$ is not a rational probabilistic status set of $\mathcal{PP}'$, which is the desired contradiction.

Thus, (1) and (2) hold, which proves the result. ∎

Proof: (of Theorem 5.19)
By definition of rationality, we know that if $\mathcal{PS}$ is a rational status set of $\mathcal{PP}$ then it must be a feasible probabilistic status set as well.

Suppose $\mathcal{PP}$ has a feasible probabilistic status set. Then the set of all feasible probabilistic status sets of $\mathcal{PP}$ on $\mathcal{O}_{\mathcal{S}}$ has a non-empty set of inclusion-minimal elements. Indeed, from the grounding of the probabilistic agent program, we can remove all rules which violate the conditions 2.-4. of the operator $\mathbf{App}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}(\mathcal{PS})$, and can remove literals involving code calls from the remaining rules. Moreover, the deontic and action closure conditions can be incorporated into the program via rules. Thus, we end up with a set T of propositional clauses, whose models are feasible probabilistic status sets of $\mathcal{PP}$. Since $\mathcal{PP}$ has a feasible probabilistic status set, T has a model, i.e., an assignment to the propositional atoms which satisfies all clauses in T. Now, each satisfiable set of clauses in a countable language posseses at least one minimal model (w.r.t. inclusion, i.w., a $\subseteq$-minimal set of atoms is assigned the value **true**); this can be shown applying the same technique which proves that every such set of clauses can be extended to a maximal satisfiable set of clauses. Thus, T has at least one minimal model. As easily seen, any such model is a minimal feasible probabilistic status set of $\mathcal{PP}$.

Suppose now $\mathcal{PS}'$ is one of the minimal feasible probabilistic status sets of $\mathcal{PP}$ on $\mathcal{O}_{\mathcal{S}}$. Then (as we show below) $\mathcal{PS}'$ is grounded, and hence a rational probabilistic status set.

To show that $\mathcal{PS}'$ is grounded, we need to show that $\mathcal{PS}'$ satisfies conditions ($\mathcal{PS}$ 1)–($\mathcal{PS}$ 3) of the definition of feasible probabilistic status set—this is true because $\mathcal{PS}'$ is feasible. In addition, we need to show that no strict subset $\mathcal{PS}^\star$ of $\mathcal{PS}$ satisfies conditions ($\mathcal{PS}$ 1)–($\mathcal{PS}$ 3).

Suppose there is a strict subset $\mathcal{PS}^\star$ of $\mathcal{PS}$ satisfying conditions ($\mathcal{PS}$ 1)–($\mathcal{PS}$ 3). Then, as $\mathcal{IC} = \emptyset$, $\mathcal{PS}^\star$ also satisfies condition ($\mathcal{PS}$ 4) of the definition of feasibility, and hence $\mathcal{PS}^\star$ is a feasible probabilistic status set. But this contradicts the inclusion minimality of $\mathcal{PS}'$, and hence, we may infer that $\mathcal{PS}'$ has no strict subset $\mathcal{PS}^\star$ of $\mathcal{PS}$ satisfying conditions ($\mathcal{PS}$ 1)–($\mathcal{PS}$ 3). Thus, $\mathcal{PS}'$ is grounded, and we are done. ∎

Proof: (of Theorem 7.6)
For each random variable $V_i = (X_i, \wp_i)$ returned by some ground code call condition in the probabilistic state $\mathcal{O}^p$, let us define its *normalized* version $V_i' = (X_i', \wp_i')$ where:

$$X_i' = \{x \mid x \in X_i \text{ and } \wp_i(x) > 0\} \cup \{\epsilon \mid \textstyle\sum_{x \in X_i} \wp_i(x) < 1\};$$
$$\wp_i'(x) = \begin{cases} \wp_i(x) & \text{if } x \in X_i' \setminus \{\epsilon\} \\ 1 - \sum_{x \in X_i} \wp_i(x) & \text{if } x = \epsilon \text{ and } \epsilon \in X_i' \end{cases}$$

i.e., we delete the zero-probability elements and add the extra one ϵ (which stands for "none of the above") whenever the distribution $\wp_i$ is incomplete. Now we can see that each tuple $\overline{x} = \langle x_1, \ldots, x_n \rangle$ in the Cartesian product

$\overline{X'} = X'_1 \times \cdots \times X'_n$ corresponds to a distinct compatible state $\mathcal{O}$ w.r.t. $\mathcal{O}^p$. In $\mathcal{O}$, a ground code call returns an object o iff in the probabilistic state $\mathcal{O}^p$ it returns a variable V_i such that $x_i = o$. Let associate to each state $\mathcal{O}$ of this kind the value

$$\wp^*(\mathcal{O}) = \wp'_1(x_1) \cdots \wp'_n(x_n)$$

and set $\wp^*(\mathcal{O}) = 0$ for all the other states. We can easily verify that $\langle \mathsf{COS}(\mathcal{O}^p), \wp^* \rangle$ is a compatible probabilistic Kripke structure for $\mathcal{O}^p$: for each random variable V_i returned in state $\mathcal{O}^p$ and each object $o \in X_i$ if $\wp_i(o) = 0$ then $o \notin X'$, so it could appear only in the zero-probability states; otherwise:

$$\sum_{o \in \mathcal{O}} \wp^*(\mathcal{O}) = \sum_{\overline{x} \in \overline{X'}, x_i = o} \prod_{j=0}^{n} \wp'_j(x_j) = \wp'_i(o) \sum_{\overline{x} \in \overline{X'}, x_i = o} \prod_{j \neq i} \wp'_j(x_j) = \wp'_i(o) = \wp_i(o)$$

Both cases satisfy the condition for compatibility. Finally, it is easy to verify that $\langle \mathsf{COS}(\mathcal{O}^p), \wp^* \rangle$ is really a probabilistic Kripke structure:

$$\sum_{\mathcal{O}} \wp^*(\mathcal{O}) = \sum_{\overline{x} \in \overline{X'}} \prod_{j=0}^{n} \wp'_j(x_j) = \sum_{x_1 \in X'_1} \wp'_1(x_1) = 1$$

∎

Proof: (of Theorem 7.7)
Let us consider the compatible Kripke structure described in the proof of Proposition 7.6 on page 27, and let us assume that V_1 and V_2 are the variables required in the thesis. The corresponding *completed* versions V'_1 and V'_2 will then contain at least two non zero-probability objects (one of them could be the extra object ϵ), respectively a_1, b_1 and a_2, b_2. Now let choose an arbitrary real number δ such that:

$$0 \leq \delta \leq \min_{x \in \{a_1, b_1\}, y \in \{a_2, b_2\}} \{\wp'_1(x)\wp'_2(y)\}$$

We can build a Kripke structure $\langle \mathsf{COS}(\mathcal{O}^p), \wp^\delta \rangle$, where $\wp^\delta$ is defined in the same way as $\wp^*$ but replacing $\wp'_1(x_1)\wp'_2(x_2)$ by $\phi(x_1, x_2)$, which in turn is defined in the following way:

$$\phi(x_1, x_2) = \begin{cases} \wp'_1(x_1)\wp'_2(x_2) - \delta & \text{if } \langle x_1, x_2 \rangle \in \{\langle a_1, a_2 \rangle, \langle b_1, b_2 \rangle\} \\ \wp'_1(x_1)\wp'_2(x_2) + \delta & \text{if } \langle x_1, x_2 \rangle \in \{\langle a_1, b_2 \rangle, \langle b_1, a_2 \rangle\} \\ \wp'_1(x_1)\wp'_2(x_2) & \text{otherwise} \end{cases}$$

It is easy to verify that it is a compatible Kripke structure. Since δ can be arbitrarily chosen within a non-point interval, we can obtain an infinite number of distinct compatible Kripke structures. ∎

Proof: (of Theorem 8.12)
($\Rightarrow$) Suppose $\mathcal{PS} = \mathit{lfp}(\mathbf{S}_{\mathcal{PP}, \mathcal{O}_S})$ a rational probabilistic status set of $\mathcal{PP}$ on

$\mathcal{O}_{\mathcal{S}}$. Then, $\mathcal{PS}$ is feasible by definition of rational probabilistic status set. By Lemma 5.18, $\mathcal{PS}$ is a pre-fixpoint of $\mathbf{S}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}$. Since $\mathbf{S}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}$ is monotone, it has by the Knaster-Tarski Theorem a least pre-fixpoint, which coincides with $\mathit{lfp}(\mathbf{S}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}})$ (see (Apt 1990; Lloyd 1987)). Thus, $\mathit{lfp}(\mathbf{S}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}) \subseteq \mathcal{PS}$. Clearly, $\mathit{lfp}(\mathbf{S}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}})$ satisfies $(\mathcal{PS}1)$ and $(\mathcal{PS}3)$; moreover, $\mathit{lfp}(\mathbf{S}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}})$ satisfies $(\mathcal{PS}2)$, as $\mathcal{PS}$ satisfies $(\mathcal{PS}2)$ and this property is hereditary. By the definition of rational probabilistic status set, it follows $\mathit{lfp}(\mathbf{S}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}) = \mathcal{PS}$.

($\Leftarrow$) Suppose $\mathcal{PS} = \mathit{lfp}(\mathbf{S}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}})$ is a feasible probabilistic status set. Since every probabilistic status set $\mathcal{PS}'$ which satisfies $(\mathcal{PS}1)$–$(\mathcal{PS}3)$ is a pre-fixpoint of $\mathbf{S}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}}$ and $\mathit{lfp}(\mathbf{S}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}})$ is the least prefix point, $\mathcal{PS}' \subseteq \mathcal{PS}$ implies $\mathcal{PS} = \mathcal{PS}'$. It follows that $\mathcal{PS}$ is rational.

Notice that in case of positive programs, $\mathit{lfp}(\mathbf{S}_{\mathcal{PP},\mathcal{O}_{\mathcal{S}}})$ always satisfies the conditions $(\mathcal{PS}1)$ and $(\mathcal{PS}3)$ of a feasible probabilistic status set (i.e., all closure conditions), and thus is a rational probabilistic status set if it satisfies $(\mathcal{PS}2)$ and $(\mathcal{PS}4)$, i.e., the consistency criteria. The uniqueness of the rational probabilistic status set is immediate from the previous theorem. ∎