Fundamenta Informaticae 57:2–4 (2003) 323–344
IOS Press

Modal Probability, Belief, and Actions

Andreas Herzig*
Institut de Recherche en Informatique de Toulouse (CNRS-INP-UPS)
Toulouse, France

Abstract. We investigate a modal logic of probability with a unary modal operator expressing that a proposition is more probable than its negation. Such an operator is not closed under conjunction, and its modal logic is therefore non-normal. Within this framework we study the relation of probability with other modal concepts: belief and action. We focus on the evolution of belief, and propose an integration of revision. For that framework we give a regression algorithm.

Keywords: Modal probability, belief, action, epistemic logic, dynamic logic, successor state axioms, regression

1. Introduction

Several researchers have investigated modal logics of probability. Some have added probability measures to possible worlds semantics, most prominently Fagin, Halpern and colleagues [10]. They use modal operators of knowledge, $\mathcal{K}\phi$ expressing that the agent knows that ϕ, and they introduce modal operators of the kind $w(\phi) \geq b$ expressing that "according to the agent, formula ϕ holds with probability at least b".

Others have studied the properties of comparative probability, following Kraft, Pratt, and Seidenberg, and Segerberg [38]. They use a relation $\phi > \psi$ (that can also be viewed as a binary modal construction) expressing "ϕ is more probable than ψ".

Only few have studied a still more qualitative notion, viz. the modal logic of constructions of the kind $\mathcal{P}\phi$ expressing that ϕ is more probable than $\neg\phi$ (or dually, at least as probable as $\neg\phi$). Among

Address for correspondence: Andreas Herzig, IRIT, 118 route de Narbonne, F-31062 Toulouse Cedex 04, France; `herzig@irit.fr`, `www.irit.fr/~Andreas.Herzig`
*The present paper continues work started in [26]. Thanks to Dominique Longin, Didier Dubois, Henri Prade, Jérôme Lang, Luis Godo, Ivan Varzinczak, and Andrzej Szalas for useful comments and discussions. Thanks to Joseph Halpern for indicating an error in the axiomatics of a previous version, and explaining the link with his work in [12].

those are Hamblin [21], Burgess [5], and T. Fine [40]. Lenzen [29, 30] has identified weak belief with being probable, and he has argued that a comprehensive account of knowledge and belief requires such a notion of weak belief. Halpern and colleagues have studied the similar notion of likelihood [20, 19].

Probably one of the reasons for the lack of interest in such approaches is that the corresponding logical systems are very poor, and are difficult to characterize w.r.t. the underlying probability measures. (Things are different if we do not take probability theory but possibility theory: as shown in [6, 7], Lewis' operator of comparative possibility [32] provides a simple and complete axiomatization of qualitative possibility relations.)

A related but different research topic is the investigation of modal logics allowing to count accessible worlds [11, 28]. Yet another option is to interpret $\mathcal{P}\phi$ not as a two-valued modal proposition but as a many-valued modal proposition, as it is done in [18, Chapter 8] and [17]. There, the truth degree of $\mathcal{P}\phi$ is taken as $Prob(\phi)$, so the bigger is the probability of ϕ, the 'more true' is the proposition $\mathcal{P}\phi$. One can then express that for the agent ϕ is more probable than $\neg\phi$ by a non-classical implication $\mathcal{P}\neg\phi \Rightarrow \mathcal{P}\phi$.

In this paper we investigate the logic of the modal operator $\mathcal{P}$ and its interactions with belief and action. After some preliminary definitions and notions (section 2) we introduce a standard modal logics of action (section 3) and belief (section 4). In the same section we also study the modal logic of 'probable' and its relation to belief. Contrarily to comparative possibility, $\mathcal{P}$ is not closed under conjunction. Therefore its modal logic is non-normal, and our models are minimal models in the sense of [8], that are based on neighborhood functions (as opposed to probability distributions).

We then turn to the relation of probability with action (section 5), and propose principles for the interplay between action, belief, and probability. As we are speaking about belief and not knowledge, we have to integrate belief revision into our framework (section 6). Our simple account allows us to identify so-called successor state axioms for both belief and probability (section 6.3), which enable us to use regression (section 7).

2. Preliminaries

In this article we only consider the case of a single agent for the sake of simplicity, the main reason being that for the time being there are no relevant interactions between several agent in our approach.

2.1. Atomic Formulas, Atomic Actions

We have a set of atomic formulas $Atm = \{p, q, \ldots\}$. One of our running example is the Yale Shooting Problem (YSP), having atoms $Loaded$, $Alive$, $HasGun$. Another one will be in terms of playing dice; we thus consider atomic formulas $d_1, d_2, \ldots$, respectively expressing "the die shows 1", etc. The set of classical formulas $Class$ is constructed from Atm with the connectives $\neg, \wedge, \vee$. Literals are atomic formulas or negations of atomic formulas. Implications $\phi \rightarrow \psi$ and equivalences $\phi \leftrightarrow \psi$ are considered to be abbreviations as usual.

We have a set of atomic actions $Act = \{\alpha, \beta, \ldots\}$. In the YSP we have $load$, $shoot$, and $wait$. In the other example we have the $throw$ action of throwing the die, the action $wait$ of waiting, and the actions $obs\, d_1$, $obs\, d_2, \ldots$ of the agent observing that the die shows 1, etc.

Actions are not necessarily executed by the agent under concern, but may be executed by other agents or by nature. (So we might as well speak about events instead of actions.)

We do not consider complex actions here. We just note that tests and nondeterministic compositions cannot behave in the standard dynamic logic way if we want to reason about knowledge or belief, and refer to [23, 24] for a more detailed analysis.

From these ingredients complex formulas will be built together with modal operators in the standard way.

2.2. Modal Operators

We have a standard doxastic modal operator $\mathcal{B}$, and the formula $\mathcal{B}\phi$ is read "the agent believes that ϕ", or "ϕ is true for the agent". For example, $\mathcal{B}d_6$ expresses that the agent believes that the die shows "6". The formula $\mathcal{B}(d_1 \vee d_2 \vee d_3 \vee d_4 \vee d_5 \vee d_6)$ expresses that the agent believes the die shows one of 1, 2, 3, 4, 5, or 6.

Moreover we have a modal operator $\mathcal{P}$ where $\mathcal{P}\phi$ is read "ϕ is probable for the agent".[1] The dual $\neg\mathcal{P}\neg\phi$ expresses that ϕ is not improbable. (This operator has been considered primitive in some papers in the literature.) For example, $\mathcal{P}(d_1 \vee d_2 \vee d_3 \vee d_4)$ expresses that it is probable for the agent that the die shows one of 1, 2, 3, or 4. $\neg\mathcal{P}d_6$ expresses that it is improbable for the agent that the die shows "6".

Finally, for every action $\alpha \in Act$ we have a dynamic logic operator $[\alpha]$. The formula $[\alpha]\phi$ is read "ϕ holds after every execution of α". For example, $[throw]d_6$ expresses that the die shows 6 after the throwing action. $\neg[throw]\bot$ expresses that $throw$ cannot be executed, and $\neg[throw]\neg d_6$ expresses that the die may show 6 after the throwing action. $\neg\mathcal{P}[throw]d_6$ expresses that this is improbable for the agent. $[throw]\mathcal{P}\neg d_6$ expresses that after throwing the die it is probable for the agent that it did not fall 6. $[throw][obs\ d_6]\mathcal{B}d_6$ expresses that after throwing the die and observing that it fell 6 the agent believes that it fell 6.

2.3. Models for Modal Operators

A *frame* is composed of a nonempty set of possible worlds W together with some structure on W. For normal modal logics such structures are accessibility relations, while for non-normal modal logics we will have neighborhood functions.

A *model* is a frame together with a valuation Val, where $Val : W \longrightarrow 2^{Atm}$ associates to every world w the set of atoms that are true in w. Models are denoted by $\mathcal{M}$.

The truth conditions are formulated in terms of the *forcing relation* $\models$ linking models, worlds, and formulas. Given a model $\mathcal{M}$ and a world w of $\mathcal{M}$, the truth condition for an atomic formula $\phi \in Atm$ is:

- $\models_{\mathcal{M},w} \phi$ if $\phi \in Val(w)$

The truth conditions for the standard operators of classical logic are as usual. When defining those for the modal operators we shall make use of the truth set for ϕ, defined as:

$$\|\phi\|_{\mathcal{M}} = \{w \in W : \models_{\mathcal{M},w} \phi\}$$

We define truth in a model $\models_{\mathcal{M}} \phi$ as $\models_{\mathcal{M},w} \phi$ for every $w \in W$. And ϕ is valid in the class of models $\mathcal{C}$, noted $\models_{\mathcal{C}} \phi$, if $\models_{\mathcal{M}} \phi$ for every $\mathcal{M} \in \mathcal{C}$.

[1] As several authors noted, we get the same account if we read $\mathcal{P}\phi$ as "probability of ϕ is high", and interpret 'high' as "greater than b", for $1 > b \geq 0.5$.

2.4. Relations Agreeing with a Probability Measure

$\mathcal{P}$ can also be viewed as a property of formulas. Let $Prob$ be any subjective probability measure defined on formulas. When it holds that

$$\mathcal{P}\phi \text{ iff } Prob(\phi) > Prob(\neg\phi)$$

we say that $\mathcal{P}$ *agrees with* $Prob$.

2.5. Hypotheses about Action

We make some hypotheses about actions and their perception by the agent. They permit to simplify the theory.

2.5.1. Public Action Occurrences

We suppose that the agent perceives action occurrences completely and correctly. For example, whenever a die is thrown the agent is aware of that, and whenever the agent believes a die is thrown then indeed such an action has occurred. One might imagine that action occurrences are publicly announced.

2.5.2. Public Action Laws

We suppose that the agent knows the laws governing the actions. Hence the agent knows that after throwing a die the effect always is that 1, 2, 3, 4, 5, or 6 show up, and that 1 and 2 cannot show up simultaneously, etc.

2.5.3. Non-informativity

We suppose that all actions are *non-informative*. Non-informative actions are actions whose instances are not observed by the agent beyond their mere occurrence. In particular the agent does not observe the outcome of nondeterministic actions such as that of throwing a die. Upon learning that such an action has occurred the agent updates his beliefs: he computes the new beliefs from his previous beliefs and his knowledge about the action laws. Hence the new beliefs neither depend on the state of the world before the action occurrence, nor on the state of the world after the action occurrence.

In our example we suppose that the $throw$ action is non informative: the agent throws the die without observing the outcome. If the agent learns that the action of throwing a die has been executed then he does not learn which side shows up.

Clearly, the action of checking the outcome of the $throw$ action is informative: the new beliefs depends on the position of the die in the real world. Other examples of informative actions are that of *looking up* a phone number, *testing if* a proposition is true, *informing whether* a proposition is true, etc.

Nevertheless, to suppose all actions are uninformative does not mean that the agent is disconnected from the world: he may learn that some proposition is true, or more precisely that some action of observing that some proposition has some value has occurred. For example, when he learns that it has been *observed that* the die fell 6 (i.e., he learns that the action of observing 6 has been executed) then he is able to update his beliefs accordingly. Indeed, the $obs\, d_i$ actions are non-informative according to our definition: when the agent learns that $obs\, d_i$ has occurred then this information enables him to adjust his beliefs in a complete way, and there is no need for further observation of the world. Other examples of uninformative actions are that of *learning that* the phone number of another agent is N, *testing that* a proposition is true (in the sense of dynamic logic tests), *informing that* a proposition is true, etc.

3. Models of Action

Our logic of action is a simple version of dynamic logic [22]. Within that logic we recast Reiter's solution to the frame problem of [34].

3.1. A-Models

Actions are viewed as transition systems. In such a system, truth of a formula $[\alpha]\phi$ in a state (alias possible world) means truth of ϕ in all states possibly resulting from the execution of α.

A frame is a tuple $\langle W, \{R_\alpha : \alpha \in Act\}\rangle$ such that for every $\alpha \in Act$:

- $R_\alpha : W \longrightarrow 2^W$ maps worlds to sets of worlds

$R_\alpha(w)$ is the set of worlds accessible from w. We call A-models the set of models of action defined in this way.

The truth condition for the modal operator is:

- $\models_{\mathcal{M},w} [\alpha]\phi$ if $R_\alpha(w) \subseteq \|\phi\|_{\mathcal{M}}$

Validity of ϕ in A-models is noted $\models_A \phi$.

Theorem 3.1. ([22, 8])
$\models_A \phi$ iff ϕ is provable from the inference rule

$$\text{if } \phi \to \psi \text{ then } [\alpha]\phi \to [\alpha]\psi \qquad (\mathrm{RM}_\alpha)$$

and the axioms:

$$[\alpha]\top \qquad (\mathrm{N}_\alpha)$$

$$([\alpha]\phi \wedge [\alpha]\psi) \to [\alpha](\phi \wedge \psi) \qquad (\mathrm{C}_\alpha)$$

3.2. Reiter Action Theories

The reasoning about actions community focuses on deductions from a theory describing a given set of actions in terms of preconditions and effects. We here present the case of action theories in Reiter's style, as transposed to dynamic logic in [9].

In such theories, all actions are deterministic, and it is supposed that the following pieces of information are given, about which some assumptions of complete information are made.

3.2.1. Action preconditions

For each atomic action α there is a classical formula $Poss(\alpha)$ describing the action preconditions of α, i.e., the condition under which α can be executed. For example $Poss(shoot) = HasGun$, $Poss(toss) = HasCoin$, and $Poss(wait) = \top$. In terms of dynamic logic, this means that for every $\alpha \in Act$ we have a global axiom $Poss(\alpha) \to \neg[\alpha]\bot$.

It is supposed that *the action preconditions are complete*: α is executable if and only if $Poss(\alpha)$ is true. Hence the global axiom can be strengthened from an implication to an equivalence $Poss(\alpha) \leftrightarrow \neg[\alpha]\bot$.

3.2.2. Set of possible causes

For each literal L there is a *finite* set of actions $Cause(L)$ describing the causes of L. (In Reiter's presentation this function is not given explicitly, but it can be retrieved from his functions γ^+ and γ^-.) $Cause(L)$ contains the actions in Act which under some circumstances cause L to become true. For example, $Cause(Alive) = \emptyset$ (no action makes an agent alive), $Cause(\neg Alive) = \{shoot, strangle\}$, and $Cause(\neg Loaded) = \{shoot\}$. As $wait$ is an action without any effect, there is no L such that $wait \in Cause(L)$. By convention, if $L = \neg p$ is a negative literal then we identify $Cause(\neg L)$ with $Cause(p)$.

We (and Reiter) suppose that $Cause(L)$ is not only finite (while Act might be infinite) but also small, in the sense that the cardinality of $Cause(L)$ is much smaller than that of Act.

Moreover, we suppose that every $Cause(L)$ is *complete*: whenever $\alpha \notin Cause(L)$ then the execution of α can never make L true. In terms of dynamic logic, causal completeness means that we have a global axiom $\neg L \to [\alpha]\neg L$ in that case. Axioms of that form are called *frame axioms*. In our example, as $strangle \notin Cause(\neg Loaded)$, we have $Loaded \to [strangle]Loaded$.

The next piece of information specifies the causal relation in more detail.

3.2.3. Effect preconditions

For every literal L and every atomic action $\alpha \in Cause(L)$ there is a classical formula $Cond(\alpha, L)$ describing the *effect preconditions* of action α. (This function corresponds to Reiter's γ^+ and γ^-.) Again by convention, if $L = \neg p$ then $Cond(\alpha, \neg L) = Cond(\alpha, p)$. For example, $Cond(strangle, \neg Alive) = \top$, and $Cond(shoot, \neg Alive) = Loaded$. In terms of dynamic logic, to every effect precondition $Cond(\alpha, L)$ there is associated a global axiom $Cond(\alpha, L) \to [\alpha]L$. As an example, consider the formula $Loaded \to [shoot]\neg Alive$.

It is supposed that *the effect preconditions are complete*: in situations where the formula $Cond(\alpha, L)$ does not hold the execution of α can never make L true. This means that we moreover have a global axiom $(\neg Cond(\alpha, L) \land \neg L) \to [\alpha]\neg L$ for every $\alpha \in Cause(L)$. For example, we have $(\neg Loaded \land Alive) \to [shoot]Alive$.

3.2.4. Models for a Reiter action theory $\mathcal{R}$

A Reiter action theory $\mathcal{R}$ is made up of the three functions $Poss(.)$, $Cause(.)$, $Cond(.,.)$. We call *models for* $\mathcal{R}$ the set of A-models satisfying the following *constraints*:

($\delta_{[\alpha]}$) every R_α is a partial function: for every $w \in W$, $card(R_\alpha(w)) \leq 1$;

(Indep) if $p \in Atm$ and $w' \in R_\alpha(w)$ then

- $\alpha \notin \mathit{Cause}(p)$ and $p \notin \mathit{Val}(w)$ implies $p \notin \mathit{Val}(w')$, and
- $\alpha \notin \mathit{Cause}(\neg p)$ and $p \in \mathit{Val}(w)$ implies $p \in \mathit{Val}(w')$;

(Exec) for every $w \in W$, $\models_w \mathit{Poss}(\alpha) \leftrightarrow \neg[\alpha]\bot$;

(Effect) for every $w \in W$,

- $\models_w \mathit{Cond}(\alpha, L) \to [\alpha]L$, and
- $\models_w (\neg \mathit{Cond}(\alpha, L) \wedge \neg L) \to [\alpha]\neg L$.

We define $\mathcal{R} \models_A \phi$ as $\models_{\mathcal{M}} \phi$ for every model $\mathcal{M}$ for $\mathcal{R}$. We have e.g. determinism of every action: $\mathcal{R} \models_A [\alpha]\phi \vee [\alpha]\neg\phi$

Remark 3.1. Note that although we suppose that every α is deterministic, we can nevertheless represent incomplete knowledge about the effects of α. For example, for the action $toss$ of tossing a coin we can write $\mathit{Poss}(toss) = \top$, $\mathit{Cause}(\mathit{Heads}) = \mathit{Cause}(\neg \mathit{Heads}) = \{toss\}$, $\mathit{Cond}(toss, \mathit{Heads}) = \top$, $\mathit{Cond}(toss, \neg \mathit{Heads}) = \top$.

3.2.5. Reiter's Regression Theorem

Lemma 3.1. Let $\mathcal{R}$ a Reiter action theory. Then the following are valid in $\mathcal{R}$–models.

1. $[\alpha]L \leftrightarrow \neg \mathit{Poss}(\alpha) \vee L$ if $\alpha \notin \mathit{Cause}(L)$ and $\alpha \notin \mathit{Cause}(\neg L)$;
2. $[\alpha]L \leftrightarrow \neg \mathit{Poss}(\alpha) \vee (L \wedge \neg \mathit{Cond}(\alpha, \neg L))$ if $\alpha \notin \mathit{Cause}(L)$ and $\alpha \in \mathit{Cause}(\neg L)$;
3. $[\alpha]L \leftrightarrow \neg \mathit{Poss}(\alpha) \vee \mathit{Cond}(\alpha, L) \vee L$ if $\alpha \in \mathit{Cause}(L)$ and $\alpha \notin \mathit{Cause}(\neg L)$;
4. $[\alpha]L \leftrightarrow \neg \mathit{Poss}(\alpha) \vee \mathit{Cond}(\alpha, L) \vee (L \wedge \neg \mathit{Cond}(\alpha, \neg L))$ if $\alpha \in \mathit{Cause}(L)$ and $\alpha \in \mathit{Cause}(\neg L)$.

Proof:
The proof can be found in [9]. We just sketch the proof of the second item from the left to the right. It suffices to show that $\mathcal{R} \models ([\alpha]L \wedge \neg L) \to \neg \mathit{Poss}(\alpha)$ and $\mathcal{R} \models ([\alpha]L \wedge \mathit{Cond}(\alpha, \neg L)) \to \neg \mathit{Poss}(\alpha)$.

First, $\mathcal{R} \models \neg L \to [\alpha]\neg L$ because $\alpha \notin \mathit{Cause}(L)$. It follows by standard modal principles that $\mathcal{R} \models (\neg L \wedge [\alpha]L) \to [\alpha]\bot$, from which $\mathcal{R} \models (\neg L \wedge [\alpha]L) \to \neg \mathit{Poss}(\alpha)$ by the laws in $\mathcal{R}$.

Second, $\mathcal{R} \models \mathit{Cond}(\alpha, \neg L) \to [\alpha]\neg L$ by the laws in $\mathcal{R}$, from which $\mathcal{R} \models ([\alpha]L \wedge \mathit{Cond}(\alpha, \neg L)) \to \neg \mathit{Poss}(\alpha)$ follows by standard modal principles.

□

For example, $\mathcal{R} \models_A [shoot]\neg \mathit{Alive} \leftrightarrow \neg \mathit{HasGun} \vee \mathit{Loaded} \vee \neg \mathit{Alive}$.

Theorem 3.2. Let $\mathcal{R}$ be a Reiter action theory. For every formula ϕ there exists a classical formula ϕ' such that $\mathcal{R} \models_A \phi \leftrightarrow \phi'$.

Proof:
First we put ϕ in negation normal form. Second, as all α are deterministic $[\alpha]$ distributes over $\wedge$ and $\vee$. Hence we end up with the innermost $[\alpha]$ having only literals in their scope. Finally, by the above lemma 3.1 these can be eliminated. Iterating this we obtain an equivalent formula without modal operators. $\square$

For example, consider the formula $(\neg Loaded \wedge Alive) \rightarrow [wait][shoot]Alive$ of our running example. Its negation normal form is $Loaded \vee \neg Alive \vee [wait][shoot]Alive$. In the first regression step $[shoot]Alive$ is replaced by $\neg HasGun \vee (Alive \wedge \neg Loaded)$. The negation normal form of the result is $Loaded \vee \neg Alive \vee [wait]\neg HasGun \vee ([wait]Alive \wedge [wait]\neg Loaded)$. The three modal subformulas are then regressed to the formula $Loaded \vee \neg Alive \vee \neg HasGun \vee (Alive \wedge \neg Loaded)$. And the latter is valid in classical propositional logic.

To sum it up, in order to decide whether $\mathcal{R} \models_A \phi$ it suffices to regress ϕ to some classical formula ϕ', and then to apply any decision procedure for classical logic.

4. Models of Belief and Probability

4.1. Principles for Belief

Full belief, or strong belief, can be identified with having probability 1. Following [27] we suppose standard principles for $\mathcal{B}$: we have the inference rule

$$\text{if } \phi \rightarrow \psi \text{ then } \mathcal{B}\phi \rightarrow \mathcal{B}\psi \tag{RM$_\mathcal{B}$}$$

and the following axioms:

$$\mathcal{B}\top \tag{N$_\mathcal{B}$}$$

$$(\mathcal{B}\phi \wedge \mathcal{B}\psi) \rightarrow \mathcal{B}(\phi \wedge \psi) \tag{C$_\mathcal{B}$}$$

$$\mathcal{B}\phi \rightarrow \mathcal{B}\mathcal{B}\phi \tag{4$_\mathcal{B}$}$$

$$\neg\mathcal{B}\phi \rightarrow \mathcal{B}\neg\mathcal{B}\phi \tag{5$_\mathcal{B}$}$$

Hence the set of beliefs is closed under logical consequences (RM$_\mathcal{B}$) and contains tautologies (N$_\mathcal{B}$). Moreover beliefs are closed under conjunction (C$_\mathcal{B}$), and agents are aware of their beliefs (4$_\mathcal{B}$) and disbeliefs (5$_\mathcal{B}$). The latter two principles are called positive and negative introspection. It will follow from the axioms linking belief and probability that beliefs are consistent, i.e. we will deduce $\neg(\mathcal{B}\phi \wedge \mathcal{B}\neg\phi)$.

4.2. Principles for Probability

There are much less principles for $\mathcal{P}$. Following [5] we accept the inference rule

$$\text{if } \phi \rightarrow \psi \text{ then } \mathcal{P}\phi \rightarrow \mathcal{P}\psi \tag{RM$_\mathcal{P}$}$$

and the axioms:

$$\mathcal{P}\top \tag{N$_\mathcal{P}$}$$

$$\neg(\mathcal{P}\phi \wedge \mathcal{P}\neg\phi) \tag{D$_\mathcal{P}$}$$

Hence just as for belief, the set of probabilities is closed under logical consequences, contains tautologies, and is consistent. In opposition to belief, probabilities are not closed under conjunctions: $(\mathcal{P}\phi \wedge \mathcal{P}\psi) \to \mathcal{P}(\phi \wedge \psi)$ is not always the case. We will give introspection principles for $\mathcal{P}$ in section 4.4.

Our axioms match those that have been put forward in the literature, e.g. those in [40]. As stated there, it seems that there are no other natural properties of 'being probable' that could be formulated using $\mathcal{P}$. These properties are only sound w.r.t. probability measures (cf. section 4.5).

4.3. Principles for Belief and Probability

What is the relation between $\mathcal{P}$ and $\mathcal{B}$? According to our reading things that are believed should also be probable for an agent, i.e., we expect $\mathcal{B}\phi \to \mathcal{P}\phi$ to hold. The following main axiom will allow us to derive that:

$$(\mathcal{B}\phi \wedge \mathcal{P}\psi) \to \mathcal{P}(\phi \wedge \psi) \tag{C-MIX}$$

Just as for the case of beliefs and disbeliefs, agents are aware of probabilities and improbabilities. This is expressed by the following two axioms:

$$\mathcal{P}\phi \to \mathcal{B}\mathcal{P}\phi \tag{4-MIX}$$

$$\neg\mathcal{P}\phi \to \mathcal{B}\neg\mathcal{P}\phi \tag{5-MIX}$$

4.4. BP-Models

Truth of the formula $\mathcal{B}\phi$ means truth of ϕ in all worlds that are possible for the agent.

In what concerns the formula $\mathcal{P}\phi$, one might imagine that to every possible world there is associated a probability measure over the set of epistemically accessible worlds, and that $Prob(\phi) > Prob(\neg\phi)$. Sometimes the intuition is put forward that among the set of accessible worlds there are more worlds where ϕ is true than worlds where ϕ is false. We shall show in section 5.4 that such an explanation is not entirely correct.

A frame is a tuple $\langle W, B, P\rangle$ such that

- $B : W \longrightarrow 2^W$
- $P : W \longrightarrow 2^{2^W}$

Thus for every possible world $w \in W$, $B(w)$ is a set of accessible worlds as usual. $P(w)$ is a *set of sets* of possible worlds, alias neighborhoods. Intuitively P collects 'big' subsets of B (in the sense that for $V \in P(w)$, V contains more elements than its complement w.r.t. W, $W \setminus V$). Nevertheless, there is no formal requirement reflecting this. (Condition ($\mathrm{d}_{\mathcal{P}}$) below is in this spirit but is weaker.)

Every frame must satisfy some *constraints*: for every $w \in W$,

($45_{\mathcal{B}}$) if $w' \in B(w)$ then $B(w') = B(w)$

($\mathrm{n}_{\mathcal{P}}$) $P(w) \neq \emptyset$

($\mathrm{d}_{\mathcal{P}}$) if $V_1, V_2 \in P(w)$, then $V_1 \cap V_2 \neq \emptyset$

(c-mix) if $V \in P(w)$ then $V \subseteq B(w)$

(45-mix) if $w' \in B(w)$ then $P(w') = P(w)$

We call BP*-models* the set of models of belief and probability defined in this way. Note that it follows from ($n_{\mathcal{P}}$), ($d_{\mathcal{P}}$), and (c-mix) that $B(w) \neq \emptyset$.

The truth conditions are:

- $\models_{\mathcal{M},w} \mathcal{B}\phi$ if $B(w) \subseteq \|\phi\|_{\mathcal{M}}$
- $\models_{\mathcal{M},w} \mathcal{P}\phi$ if there is $V \in P(w)$ such that $V \subseteq \|\phi\|_{\mathcal{M}}$

Validity of ϕ in BP-models is noted $\models_{BP} \phi$.

Theorem 4.1. $\models_{BP} \phi$ iff ϕ is provable from the inference rules ($RM_{\mathcal{B}}$) and ($RM_{\mathcal{P}}$), and the following axioms: ($N_{\mathcal{B}}$), ($C_{\mathcal{B}}$), ($4_{\mathcal{B}}$), ($5_{\mathcal{B}}$), ($N_{\mathcal{P}}$), ($D_{\mathcal{P}}$), (C-MIX), (4-MIX), (5-MIX).

The proof is essentially Burgess' in [5]. The logic he investigates in this paper differs from ours only in that he considers knowledge instead of belief, which is irrelevant for his central results.

Burgess has also established a S5-like normal form theorem, whose proof is just the same in our case. Call a formula *primitive* if it is either of the form $\mathcal{B}$ or $\neg\mathcal{B}$ followed by a disjunction of literals, or of the form $\mathcal{P}$ or $\neg\mathcal{P}$ followed by a disjunction of conjunction of literals.

Theorem 4.2. For every formula ϕ there exists a formula ϕ' such that $\models_{BP} \phi \leftrightarrow \phi'$, and ϕ' is a disjunction of conjunctions of primitive formulas.

Burgess has given a decision procedure for his logic that is based on this normal form, which carries over almost unchanged:

Theorem 4.3. BP-validity is decidable.

Provability of ϕ in our axiomatics for BP is noted $\vdash_{BP} \phi$. Here are some provable formulas:

1. if $\vdash_{BP} \phi \leftrightarrow \psi$ then $\vdash_{BP} \mathcal{P}\phi \leftrightarrow \mathcal{P}\psi$
2. $\vdash_{BP} \neg\mathcal{P}\bot$
3. $\vdash_{BP} \mathcal{B}\phi \to \mathcal{P}\phi$
4. $\vdash_{BP} \neg(\mathcal{B}\phi \wedge \mathcal{B}\neg\phi)$
5. $\vdash_{BP} (\mathcal{B}\phi \wedge \neg\mathcal{P}\neg\psi) \to \neg\mathcal{P}\neg(\phi \wedge \psi)$
6. $\vdash_{BP} \mathcal{P}\phi \to \neg\mathcal{B}\neg\phi$
7. $\vdash_{BP} (\mathcal{P}\phi \wedge \mathcal{P}\psi) \to \neg\mathcal{B}\neg(\phi \wedge \psi)$
8. $\vdash_{BP} (\mathcal{B}(\phi \to \psi) \wedge \mathcal{P}\phi) \to \mathcal{P}\psi$
9. $\vdash_{BP} \mathcal{P}\phi \leftrightarrow \mathcal{B}\mathcal{P}\phi$
10. $\vdash_{BP} \neg\mathcal{P}\phi \leftrightarrow \mathcal{B}\neg\mathcal{P}\phi$

11. $\vdash_{BP} \mathcal{P}\phi \leftrightarrow \mathcal{P}\mathcal{P}\phi$

12. $\vdash_{BP} \neg\mathcal{P}\phi \leftrightarrow \mathcal{P}\neg\mathcal{P}\phi$

13. $\vdash_{BP} \mathcal{P}\mathcal{B}\phi \rightarrow \mathcal{P}\phi$

And here are some formulas that cannot be proved:

1. $\not\vdash_{BP} \mathcal{P}\phi \rightarrow \mathcal{B}\phi$

2. $\not\vdash_{BP} \mathcal{P}\phi \rightarrow \mathcal{P}\mathcal{B}\phi$

3. $\not\vdash_{BP} (\mathcal{P}\phi \wedge \mathcal{P}\psi) \rightarrow \mathcal{P}(\phi \wedge \psi)$

4. $\not\vdash_{BP} (\mathcal{P}\phi \wedge \mathcal{P}(\phi \rightarrow \psi)) \rightarrow \mathcal{P}\psi$

Note that neither of them is valid w.r.t. probability measures.

4.5. Incompleteness w.r.t. Probability Measures

BP-models differs from the standard semantics in terms of probability measures. What is the relation between them?

It is straightforward to establish that our axiomatics is sound w.r.t. probability measures: whenever we define $\mathcal{P}\phi$ by $Prob(\phi) > 0.5$ then $\mathcal{P}$ satisfies (RM$_\mathcal{P}$), (N$_\mathcal{P}$), (D$_\mathcal{P}$).

Nevertheless and as announced in Section 4.2, our axioms are incomplete w.r.t. probability measures.

Theorem 4.4. ([40])
Let $Atm = \{a, b, c, d, e, f, g\}$ and $Act = \emptyset$. Take a model $\mathcal{M}$ where

- $W = 2^{Atm}$

- for every $w \in W$, $P(w) = \{efg, abg, adf, bde, ace, cdg, bcf\}$, where efg is a shorthand for $\{e, f, g\}$, etc.

- for $p \in Atm$, $V(p) = \{w \in W : p \in w\}$

Then $\mathcal{M}$ is a BP-model, but there is no agreeing probability measure.

Therefore our neighborhood semantics is strictly weaker than probability measures. We note that based on Segerberg's axiomatization of comparative probability, Lenzen [29] has proposed complex but complete axiomatizations of the concept 'having probability greater than 0.5'.

5. ABP-Models of Action, Belief, and Probability

We recall the hypotheses we have stated in Section 2.5:

- the agent perceives action occurrences completely and correctly;
- the agent knows the laws governing the actions;
- actions are non-informative, i.e., the agent does not learn about particular effects of actions beyond what is stipulated in the action laws.

As action effects are not observed, when the agent learns that the action of throwing a die has been executed then he does not learn whether it fell 6 or not.

5.1. Observations

We suppose that Act contains observation actions of the form $obs\ \phi$, where ϕ is any formula.

We postulate that ϕ can be observed if and only if ϕ is true:

(pre$_{obs\ \phi}$) $R_{obs\ \phi}(w) \neq \emptyset$ iff $\models_{\mathcal{M},w} \phi$.

This validates

$$\neg[obs\ \phi]\bot \leftrightarrow \phi \qquad (\text{Pre}_{obs\ \phi})$$

Idealizing a bit we consider that observations have no physical effect. This means that the valuations before and after $obs\ \phi$ should be identical:

(eff$_{obs\ \phi}$) if $w' \in R_{obs\ \phi}(w)$ then $Val(w') = Val(w)$.

This validates:

$$\phi \to [obs\ \psi]\phi \quad \text{for every classical formula } \phi \qquad (\text{Eff}_{obs\ \phi})$$

It follows from (Pre$_{obs\ \phi}$) and (Eff$_{obs\ \phi}$) that $[obs\ \phi]\phi$ whenever ϕ is a classical formula. It also follows that $\mathcal{B}[obs\ \phi]\phi$. We shall prove in section 6.3 that $[obs\ \phi]\mathcal{B}\phi$.

We do not say anything about the effects of observations on the agent's beliefs. This will follow from the more general principles below.

5.2. Principles for Action and Belief

A lot of researchers have proposed principles for the interaction between belief and action that might be called "no forgetting" (NF) and "no learning" (NL). Starting with [33] and [35], there is work by [14, 16, 15, 4, 3, 36, 37]. We here follow our exposition in [23, 25]. The basic idea is that the agent's new beliefs only depends on the previous beliefs and the event whose occurrence he has learned. Being confined to his set of possibilities, when the agent learns that an event α has occurred then he 'mentally executes' α in each of his possible worlds. He then collects the resulting set of worlds to form his new set of possibilities.

For sets of possible worlds $V \subseteq W$ we define $R_\alpha(V) = \bigcup_{v \in V} R_\alpha(v)$. This allows us to define more concisely the following *constraint*: for every $w \in W$, (nf-nl$_\mathcal{B}$) if $w' \in R_\alpha(w)$ and $R_\alpha(B(w)) \neq \emptyset$, then $B(w') = R_\alpha(B(w))$.

Such models validate:

$$(\neg[\alpha]\bot \wedge [\alpha]\mathcal{B}\phi) \to \mathcal{B}[\alpha]\phi \tag{NL$_\mathcal{B}$}$$

$$(\neg\mathcal{B}[\alpha]\bot \wedge \mathcal{B}[\alpha]\phi) \to [\alpha]\mathcal{B}\phi \tag{NF$_\mathcal{B}$}$$

$$[\alpha]\mathcal{B}\phi \vee [\alpha]\neg\mathcal{B}\phi \tag{e$\Delta_\mathcal{B}$}$$

For the "no learning" axiom, we must suppose that the action α is executable ($\neg[\alpha]\bot$), because else $[\alpha]\mathcal{B}\phi)$ would follow from $[\alpha]\bot$, and we could not deduce anything relevant. Similarly, for the "no forgetting" axiom we must suppose that the agent does not believe α to be inexecutable ($\neg\mathcal{B}[\alpha]\bot$). The last principle e$\Delta_\mathcal{B}$ of epistemic determinism says that the agent has no means to distinguish the different possible outcomes of a nondeterministic action whose occurrence he is informed of.

If ϕ is a classical formula then $\neg\mathcal{B}\neg\phi \to [obs\ \phi]\mathcal{B}\phi$ follows from (Pre$_{obs\ \phi}$), (Eff$_{obs\ \phi}$), and (NF$_\mathcal{B}$): the agent believes observations except if he believes the contrary.

When the agent believes α to be inexecutable and nevertheless learns that it has occurred then he must revise his beliefs. This is investigated in section 6.

Remark 5.1. We can illustrate now what we have said in Section 2.1 about complex actions. First, we cannot have the standard dynamic logic equivalence for tests $[\phi?]\psi \leftrightarrow (\phi \to \psi)$ if we suppose that tests are performed by agents. In this case we expect $[\phi?]\mathcal{B}\phi$, to hold, while $\phi \to \mathcal{B}\phi$ is in general false. Second, we cannot have $[\alpha \cup \beta]\psi \leftrightarrow ([\alpha]\phi \wedge [\beta]\phi)$ either: for example, $[throw_1 \cup \ldots \cup throw_6]\neg\mathcal{B}d_6$ should not be equivalent to $[throw_1]\neg\mathcal{B}d_6 \wedge \ldots \wedge [throw_6]\neg\mathcal{B}d_6$: while the former is plausible given that we have supposed the agent does not learn about the outcome of $throw_1 \cup \ldots \cup throw_6$, the latter should not hold under our hypothesis of public action.

5.3. Principles for Action and Probability

Similar considerations lead to the following constraint: for every $w \in W$,

(nf-nl$_\mathcal{P}$) if $w' \in R_\alpha(w)$ and $R_\alpha(V) \neq \emptyset$ for all $V \in P(w)$, then $P(w') = \{R_\alpha(V) : V \in P(w)\}$

This validates a no forgetting axiom:

$$(\neg\mathcal{P}[\alpha]\bot \wedge \mathcal{P}[\alpha]\phi) \to [\alpha]\mathcal{P}\phi \tag{NF$_\mathcal{P}$}$$

To illustrate this suppose before learning that a die has been thrown it is probable for the agent that the die will not fall 6: $\mathcal{P}[throw]\neg d_6$. When the agent learns that the die-throwing action has been executed (without learning the outcome, cf. our hypotheses) then it is probable for him that the die does not show 6. Therefore we should have $\mathcal{P}[throw]\neg d_6 \to [throw]\mathcal{P}\neg d_6$. We have to condition this: if it is probable that $throw$ is inexecutable (because e.g. it is probable that there is no die around), i.e., $\mathcal{P}[throw]\bot$, then $\mathcal{P}[throw]\neg d_6$ follows by principles of monotonic modal logics. This should be excluded.

The other way round, when it is probable for the agent that the die fell 6 after $throw$ then (as we have supposed that he does not observe the outcome of throwing) it was already probable for the agent that 6 would show up *before* learning that the action has been executed. This is expressed by the following no learning axiom:

$$(\neg[\alpha]\bot \wedge [\alpha]\mathcal{P}\phi) \to \mathcal{P}[\alpha]\phi \qquad (\mathrm{NL}_{\mathcal{P}})$$

The last principle that is validated by (nf-nl$_{\mathcal{P}}$) is similar to epistemic determinism for belief:

$$[\alpha]\mathcal{P}\phi \vee [\alpha]\neg\mathcal{P}\phi \qquad (\mathrm{e}\Delta_{\mathcal{P}})$$

5.4. Example

Let us illustrate that the intuition of $P(w)$ 'collecting more than 50% of the accessible worlds' is misleading.

Let the agent learn in w_0 that a die has been thrown. Then we might suppose that after $throw$ the situation is described by a possible world w where $B(w) = \{v_1, \ldots, v_6\}$ such that $v_i \in V(d_j)$ iff $i = j$, and where $P(w)$ is the set of all subsets of $B(w)$ containing more than half of the worlds in $B(w)$, i.e., $P(w) = \{V \subseteq B(w) : card(V) > 3\}$.

Now suppose we are in a game where a player is entitled to throw his die a second time if (and only if) his first throw was a 6. Let $throw2$ describe that complex action. (In terms of dynamic logic it could be written $throw$; `if` d_6 `then` $throw$ `else` $wait$.) We have thus

$$R_{throw2}(w) = R_{throw}(w) \cup \{v'_{6_1}, \ldots, v'_{6_6}\} = \{v_1, \ldots, v_5, v'_{6_1}, \ldots, v'_{6_6}\},$$

with $v'_{6_i} \in V(d_j)$ iff $i = j$. According to our semantics, the situation after $throw2$ might be described by a possible world $w' \in R_{throw2}(w)$ such that

- $B(w') = R_{throw2}(w)$, i.e., $card(B(w')) = 11$;
- The neighborhood $P(w')$ of w' contains in particular $\{v_1, \ldots, v_5\}$, i.e., $P(w')$ has sets containing much less than half of the worlds in $B(w')$.

5.5. ABP-Models

An *ABP-model* is a tuple $\langle W, B, P, \{R_\alpha : \alpha \in Act\}, Val\rangle$ such that $\langle W, \{R_\alpha : \alpha \in Act\}, Val\rangle$ is an A-model, $\langle W, B, P, Val\rangle$ is a BP-model, and the constraints (nf-nl$_{\mathcal{B}}$), (nf-nl$_{\mathcal{P}}$), (pre$_{obs}$ ϕ), (eff$_{obs}$ ϕ) are satisfied. Note that observations are insensitive to the syntactical form of the formula under concern: if $\|\phi\|_{\mathcal{M}} = \|\psi\|_{\mathcal{M}}$ then $R_{obs\,\phi} = R_{obs\,\psi}$.

Validity of ϕ in ABP-models is noted $\models_{ABP} \phi$. Our axioms are sound w.r.t. ABP-validity:

Theorem 5.1. If ϕ is provable from (1) the inference rules and axioms of the logic of actions A, (2) the inference rules and axioms of the logic of belief and probability BP, and (3) the axioms (NL$_{\mathcal{B}}$), (NF$_{\mathcal{B}}$), (NL$_{\mathcal{P}}$), (NF$_{\mathcal{P}}$), (e$\Delta_{\mathcal{B}}$), (e$\Delta_{\mathcal{P}}$), (Pre$_{obs\,\phi}$), (Eff$_{obs\,\phi}$), then $\models_{ABP}\ \phi$.

We conjecture that we have completeness, too, but do not pursue this here. Instead, in the rest of the paper we will focus on the case of revision and its integration into Reiter action theories.

First, note that from the axioms (NF$_{\mathcal{B}}$) and (NL$_{\mathcal{B}}$) we get:

$$\vdash_{ABP} \neg\mathcal{B}[\alpha]\bot \ \to\ ([\alpha]\mathcal{B}\phi \leftrightarrow ([\alpha]\bot \vee \mathcal{B}[\alpha]\phi)).$$

From the latter we get with (e$\Delta_{\mathcal{B}}$):

$$\vdash_{ABP} \neg\mathcal{B}[\alpha]\bot \;\rightarrow\; ([\alpha]\neg\mathcal{B}\phi \leftrightarrow ([\alpha]\bot \vee \neg\mathcal{B}[\alpha]\phi)),$$

This is almost a so-called successor state axiom (SSA) as studied in cognitive robotics. Such axioms allow us to permute $[\alpha]$ and $\mathcal{B}$, and they are important because they enable the regression proof technique [34, 31] in the case of deterministic actions without indirect effects.

So what can we say about the contexts $\mathcal{B}[\alpha]\bot$? This will be the subject of the next section.

6. $ABP\star$-Models and Revision

Most of the approaches in the literature have worked with modal operators of knowledge $\mathcal{K}$. For this concept there are no surprising events because $\mathcal{K}[\alpha]\bot$ implies $[\alpha]\bot$. It follows that the SSA for knowledge is just $[\alpha]\mathcal{K}\phi \leftrightarrow ([\alpha]\bot \vee \mathcal{K}[\alpha]\phi)$. But for the weaker belief operator, $\mathcal{B}[\alpha]\bot$ does not imply $[\alpha]\bot$, and we therefore have to investigate the case $\mathcal{B}[\alpha]\bot$ if we want to speak about belief (and a fortiori about weak belief, alias probability). Can we define a permutation principle for such contexts?

In the case of $\mathcal{B}[\alpha]\bot$ the event α comes as a surprise for the agent, while in the case of $\mathcal{P}[\alpha]\bot$ the occurrence of α was perhaps just improbable for him before he learned about it. In the first case, the agent has to revise his beliefs, and in the second case he has to revise his probabilities.

In [39] and [25] it has been investigated how belief revision in the style of AGM [13] can be integrated. We here follow the ideas in [25].

6.1. Enabling Actions

Let $v \in B(w)$. If $R_\alpha(v) = \emptyset$ then the preconditions of α do not hold in v. When the agent subsequently learns that α has nevertheless occurred, he cannot just mentally execute α to form his new beliefs: he has to revise his beliefs about α's preconditions $Poss(\alpha)$ first. We suppose that the agent does this by adjusting v so as to make the preconditions true, and that he achieves this by mentally executing a particular action whose effect is $Poss(\alpha)$.

Formally, we associate to every atomic action α a new construction $\star\alpha$, and we say that $\star\alpha$ makes α executable.

First, we postulate that for every action, $\star\alpha$ can always be executed: for every $w \in W$,

(d$_{\star\alpha}$) $\quad R_{\star\alpha}(w) \neq \emptyset$

Second, what are the effects of $\star\alpha$? As $\star\alpha$ makes α executable, we must have:

(d2$_{\star\alpha}$) $\quad R_\alpha(R_{\star\alpha}(w)) \neq \emptyset$

Finally, if α is already executable then $\star\alpha$ does nothing:

(triv$_{\star\alpha}$) $\quad$ if $R_\alpha(w) \neq \emptyset$ then for every $w' \in R_{\star\alpha}(w)$, $Val(w') = Val(w)$.

Remark 6.1. Actions that are never executable cannot be enabled. This is our main reason to exclude dynamic logic action composition. To witness, consider a slightly modified version of our tossing example where $Poss(toss) = HasCoin$, $Cause(\neg HasCoin) = \{toss\}$, and $Cond(toss, \neg HasCoin) = \top$. Then the complex action $toss; toss$ of executing $toss$ twice cannot be enabled.

6.2. $ABP\star$-Models

Using enable-actions we are now ready to define supplementary constraints of no forgetting:

(nf-nl$\star_{\mathcal{B}}$) if $w' \in R_\alpha(w)$, then $B(w') = R_\alpha(R_{\star\alpha}(B(w)))$.

(nf-nl$\star_{\mathcal{P}}$) if $w' \in R_\alpha(w)$, then $P(w') = \{R_\alpha(R_{\star\alpha}(V)) : V \in P(w)\}$

They say that when the agent learns that α has occurred then he first mentally enables executability of α, and then mentally executes α and collects the resulting states.

Figure 1 illustrates the kind of models we obtain by a variation of the tossing example where the coin falls tails and the agent wrongly believes that the coin fell heads, and then observes that it fell tails.

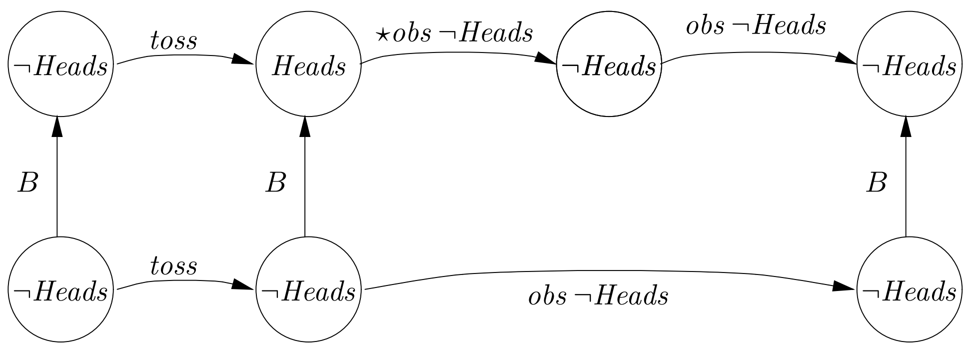

Figure 1. the toss example

We call *$ABP\star$-models* the set of ABP-models satisfying (nf-nl$\star_{\mathcal{B}}$) and (nf-nl$\star_{\mathcal{P}}$), and the constraints (d$_{\star\alpha}$), (d2$_{\star\alpha}$), (triv$_{\star\alpha}$) for enabling actions. Beyond the axioms for ABP-models, $ABP\star$-models validate principles for enabling:

$$\neg[\star\alpha]\bot \qquad (\mathrm{D}_{\star\alpha})$$

$$[\star\alpha]\neg[\alpha]\bot \qquad (\mathrm{D2}_{\star\alpha})$$

$$(\neg[\alpha]\bot \wedge \phi) \to [\star\alpha]\phi \quad \text{if } \phi \text{ is a classical formula} \qquad (\mathrm{Triv}_{\star\alpha})$$

and two new no forgetting principles:

$$\mathcal{B}[\star\alpha][\alpha]\phi \to [\alpha]\mathcal{B}\phi \qquad (\star\mathrm{NF}_{\mathcal{B}})$$

$$\mathcal{P}[\star\alpha][\alpha]\phi \to [\alpha]\mathcal{P}\phi \qquad (\star\mathrm{NF}_{\mathcal{P}})$$

As announced in section 5.1, we can establish now that if ϕ is a classical formula then $\models_{ABP\star}$ $[obs\ \phi]\mathcal{B}\phi$. For example, for a classical formula ϕ we have seen in section 5.1 that $\vdash_{ABP}\ [obs\ \phi]\phi$. Hence $\vdash_{ABP\star}\ [obs\ \phi]\phi$, too. By ABP-principles we obtain $\vdash_{ABP\star} \mathcal{B}[\star obs\ \phi][obs\ \phi]\phi$. It follows from our successor state axioms that $\vdash_{ABP\star}\ [obs\ \phi]\mathcal{B}\phi$.

Theorem 6.1. If ϕ is provable from the inference rules and axioms of the logic ABP together with the axioms (D$_{\star\alpha}$), (D2$_{\star\alpha}$), (Triv$_{\star\alpha}$), ($\star$NF$_{\mathcal{B}}$), ($\star$NF$_{\mathcal{P}}$), then $\models_{ABP\star} \phi$.

As for ABP, we conjecture that we have completeness, but do not pursue this issue.

What is the relation with AGM belief revision? Let us define revision by ϕ as enabling the observation of ϕ: $\star(obs\ \phi)$, and identify the deduction problem of whether the revision $\psi \star \phi$ of ψ by ϕ implies ξ with $\psi \rightarrow [\star obs\ \phi]\xi$, then it can be shown that most of the AGM postulates are satisfied [25].

6.3. Successor State Axioms for Belief and Probability

Putting ($\star$NF$_{\mathcal{B}}$) and ($\star$NF$_{\mathcal{P}}$) together with the no learning principles we obtain:

$$\vdash_{ABP\star} [\alpha]\mathcal{B}\phi \leftrightarrow ([\alpha]\bot \vee \mathcal{B}[\star\alpha][\alpha]\phi)$$

$$\vdash_{ABP\star} [\alpha]\mathcal{P}\phi \leftrightarrow ([\alpha]\bot \vee \mathcal{P}[\star\alpha][\alpha]\phi)$$

From the latter we get with (e$\Delta_{\mathcal{B}}$):

$$\vdash_{ABP\star} [\alpha]\neg\mathcal{B}\phi \leftrightarrow ([\alpha]\bot \vee \neg\mathcal{B}[\star\alpha][\alpha]\phi)$$

$$\vdash_{ABP\star} [\alpha]\neg\mathcal{P}\phi \leftrightarrow ([\alpha]\bot \vee \neg\mathcal{P}[\star\alpha][\alpha]\phi)$$

Thus we have succeeded in finding successor state axioms for belief and probability.

For example, $[shoot]\mathcal{P}\neg Alive$ is equivalent to $[shoot]\bot \vee \mathcal{P}[\star shoot][shoot]\neg Alive$.

In the rest of the paper we focus on the integration of $\mathcal{B}$ and $\mathcal{P}$ into Reiter action theories.

7. Regression

In this section we investigate how Reiter's regression can be extended to $ABP\star$-logic. Therefore we restrict our attention to *deterministic actions*. In order to guarantee determinism of enable-actions we need two hypotheses:

- for every α, $Poss(\alpha)$ is a conjunction of literals;
- the argument ϕ of $obs\ \phi$ is a conjunction of literals.

7.1. Augmented Reiter Action Theories

The first thing we do is to augment our Reiter action theories $\mathcal{R}$ of section 3.2 by laws for observations and enable-actions. We hence do 'double work' because the preconditions and effects of these two particular classes of actions have already been established in sections 5.1 and 6.1, but we prefer to do so for the sake of uniformity.

Given a theory $\mathcal{R}$ without obs and $\star$, we add:

- $Poss(\star\alpha) = \top$;
- $Poss(obs\ \phi) = \phi$;
- $\star\alpha \in Cause(L)$ iff L appears in $Poss(\alpha)$;
- $obs\ \phi \notin Cause(L)$ for any literal L;

- $Cond(\star\alpha, L) = \top$ for every L appearing in $Poss(\alpha)$;

Note that the condition $Cond(obs\,\phi, L)$ need not be defined because $Cause(L)$ never contains $obs\,\phi$.

For example, $Cause(HasGun) = \star shoot$, and $Cause(HasCoin) = \star toss$. The items we have added guarantee that the expected preconditions and effects already follow in ABP-logic, such as $\neg[\star obs\,\phi]\bot \leftrightarrow \phi$, for observations, and $\neg[\star\alpha]\bot$ for enabling.

Let us call $\mathcal{R}\star$ the resulting augmentation of $\mathcal{R}$. Then $\mathcal{R}\star$-models are defined to be particular $ABP\star$-models just in the same way as $\mathcal{R}$-models are defined from A-models.

7.2. Separation Theorem

The successor state axioms for belief and probability of section 6.3 allow us to 'move in' action operators.

Theorem 7.1. For any formula ϕ there exists a formula ϕ' such that $\models_{ABP\star} \phi \leftrightarrow \phi'$, and ϕ' has no action operators in the scope of $\mathcal{B}$ and $\mathcal{P}$.

Proof:
As every α is deterministic, $[\alpha]$ distributes over $\wedge$ and $\vee$. And it can be moved inwards for formulas of the form $[\alpha]\mathcal{B}\phi$, $[\alpha]\mathcal{P}\phi$, $[\alpha]\neg\mathcal{B}\phi$, $[\alpha]\neg\mathcal{P}\phi$ by the successor state axioms for belief and probability of section 6.3. □

7.3. Regression Theorem

Now we are ready for the grand finale.

Theorem 7.2. For any formula ϕ there exists a formula without action operators ϕ' such that

$$\mathcal{R}\star \models_{ABP\star} \phi \leftrightarrow \phi'.$$

Proof:
This follows from the previous separation theorem and Reiter's regression theorem (theorem 3.2 of section 3.2.5). □

As the regressed formula ϕ' contains no action operators, we have $\mathcal{R}\star \models_{ABP\star} \phi'$ iff $\models_{BP} \phi'$. Now whether $\models_{BP} \phi'$ is the case can be established using the decision procedure for BP-validity (theorem 4.3 section 4.4).

Theorem 7.3. Logical consequence from a Reiter action theory $\mathcal{R}\star$ in $ABP\star$ is decidable.

For example, consider the formula

$$\mathcal{P}Loaded \rightarrow [shoot]\mathcal{P}\neg Alive.$$

Its negation normal form is $\neg\mathcal{P}Loaded \vee [shoot]\mathcal{P}\neg Alive$. Using the separation theorem this is transformed into

$$\neg\mathcal{P}Loaded \vee \neg HasGun \vee \mathcal{P}[\star shoot][shoot]\neg Alive.$$

Then the innermost $[shoot]\neg Alive$ is replaced by $\neg HasGun \vee Loaded \vee \neg Alive$. The negation normal form of the result is

$$\neg\mathcal{P}Loaded \vee \neg HasGun \vee \mathcal{P}([\star shoot]\neg HasGun \vee [\star shoot]Loaded \vee [\star shoot]\neg Alive).$$

Now $[\star shoot]\neg HasGun$ is regressed to $\bot$, $[\star shoot]Loaded$ to $Loaded$, and $[\star shoot]\neg Alive$ to $\neg Alive$. We thus obtain

$$\neg \mathcal{P}Loaded \lor \neg HasGun \lor \mathcal{P}(\bot \lor Loaded \lor \neg Alive).$$

The latter is valid in BP-logic. Therefore we have shown

$$\mathcal{R}\star \models_{ABP\star} \mathcal{P}Loaded \to [shoot]\mathcal{P}\neg Alive$$

as expected.

8. Conclusion

We have investigated a 'very qualitative' notion of probability, that of a formula being more probable than its negation. We have presented principles governing its relation with belief, and we have proposed principles for the interplay between action, belief, and probability. The resulting logic is a fairly standard non-normal, monotonic modal logic. Hence our semantics is rather far away from probabilities. This might be felt to be at odds with intuitions, but as a matter of fact it matches the formal properties of $\mathcal{P}$ that we have put forward. We have stated regression within our modal logic, taking over Reiter's solution.

In our general framework we have allowed for nondeterministic actions (while we have restricted Reiter action theories and thus the regression algorithm to deterministic actions). In fact the general account is not entirely satisfying because we did not consider probabilistic action: in our framework uncertainty can only diminish as actions occur. This means that probability will converge towards belief on the long run, in the sense that we will end up with $P(w) = \{B(w)\}$. The problem is the same for belief in the existing logics of belief and action of the literature, and just as in their case, this is unsatisfactory.

Concerning related work, as noted in [2] in AI there is a lot of work on uncertainty within probabilistic formalisms on the one hand, and on reasoning about actions, update and revision on the other hand. But only very little approaches offer an integrated account of uncertainty and action. (For example, while BDI-style approaches provide a fine-grained analysis of agents' proactive attitudes in terms of desires and intentions, in most of the approaches there is only one mental attitude, viz. that of full belief.)

Friedman and Halpern [12] have generalized both probability and possibility measures to plausibility measures. In terms of our approach, for each possible world $w \in W$ there is a measure $Plaus_w$ mapping subsets of $B(w)$ to some partially ordered set such that $Plaus_w(V) \leq Plaus_w(V')$ whenever $V \subseteq V'$. Such plausibility orderings provide truth conditions for their conditional operator $\Rightarrow$:

$$\models_{\mathcal{M},w} \phi \Rightarrow \psi \quad \text{if} \quad Plaus_w(\|\phi \land \psi\|_{\mathcal{M}} \cap B(w)) > Plaus_w(\|\phi \land \neg\psi\|_{\mathcal{M}} \cap B(w))$$

If we moreover suppose that there are V and V' such that $Plaus_w(V) > Plaus_w(V')$, then $\mathcal{P}\phi$ can be expressed in their framework as $\top \Rightarrow \phi$. Semantically, the neighborhood $P(w)$ can be defined from $Plaus_w$ by $P(w) = \{V : Plaus(V) > Plaus(B(w) \setminus V)\}$. Less obviously,[2] a given $P(w)$ defines a plausibility ordering $Plaus_w$ by:

$$Plaus_w(V) > Plaus_w(V') \quad \text{iff there is } V'' \in P(w) \text{ such that } V'' \subseteq V \text{ and } V' \subseteq B(w) \setminus V''$$

Albeit this semantical correspondence it seems that there is no syntactical way to express Friedman and Halpern's conditional $\Rightarrow$ by means of $\mathcal{P}$. Friedman and Halpern go on to define updating in terms of conditioning. When doing this they constrain their plausibility measures $Plaus$ in a way such that they are closed under conjunction.

[2] Joseph Halpern has pointed this out to me.

In [1, 2] Bacchus et al. have integrated probability measures into a logic of action. While being thus more quantitative in spirit, their approach is similar to ours in what concerns the integration of action. The differences are as follows. First, they have probabilistic actions, and allow for (at least some form of) misperception by means of a special predicate $OI(\alpha, \alpha')$ similar to Baltag's account. $OI(\alpha, \alpha')$ expresses that action α is observation indistinguishable from action α' by the agent. We note that α must be among the actions that are observation indistinguishable from α because the approach is in terms of knowledge. This leads us to the second difference: as we have said, an approach in terms of knowledge simplifies the account a lot. If we considered knowledge then our approach can be translated to theirs by stating (1) that the only action that is observation indistinguishable from α is α itself, and (2) that for every state there is exactly one α whose likelihood is 1, while that of all the other actions is 0. There seems to be no other similar work relating modal probability to belief and action.

In future work we would like to improve our account by introducing misperception and studying probabilistic actions. The former has been done for beliefs in [25], following the ideas of [4, 3]. Such a solution should apply here without difficulties. The latter is a much harder task, which probably requires abandoning our 'extremely qualitative' account.

References

[1] Bacchus, F., Halpern, J., Levesque, H.: Reasoning about noisy sensors in the situation calculus, *Proc. 14th Int. Joint Conf. on Artificial Intelligence (IJCAI'95)*, 1995.

[2] Bacchus, F., Halpern, J., Levesque, H.: Reasoning about noisy sensors in the situation calculus, *Artificial Intelligence*, **111**, 1999, 131–169.

[3] Baltag, A.: *A Logic of Epistemic Actions*, Technical report, CWI, 2000, `http://www.cwi.nl/~abaltag/papers.html`.

[4] Baltag, A., Moss, L. S., Solecki, S.: The Logic of Public Announcements, Common Knowledge, and Private Suspicions, *Proc. TARK'98*, Morgan Kaufmann, 1998.

[5] Burgess, J. P.: Probability logic, *J. of Symbolic Logic*, **34**, 1969, 264–274.

[6] Fariñas del Cerro, L., Herzig, A.: A modal analysis of possibility theory, *Proc. European Conf. on Symbolic and Quantitative Approaches to Uncertainty (ECSQAU'91)*, number 548 in LNCS, Springer Verlag, 1991, (short version; long version published in FAIR'91).

[7] Fariñas del Cerro, L., Herzig, A.: A modal analysis of possibility theory (Invited Paper), *Proc. of the Int. Workshop on Foundations of AI Research (FAIR 91)* (P. Jorrand, J. Kelemen, Eds.), number 535 in LNAI, Springer Verlag, September 1991, (short version published in ECSQAU'91).

[8] Chellas, B.: *Modal logic: An introduction*, Cambridge University Press, 1980.

[9] Demolombe, R., Herzig, A., Varzinczak, I.: Regression in modal logic, *Journal of Applied Non-Classical Logics*, **13**(2), 2003, 165–185.

[10] Fagin, R., Halpern, J. Y.: Reasoning about knowledge and probability, *Journal of the ACM*, **41**(2), 1994, 340–367.

[11] Fattorosi-Barnaba, M., de Caro, F.: Graded modalities I, *Studia Logica*, **44**, 1985, 197–221.

[12] Friedman, N., Halpern, J. Y.: Modeling belief in dynamic systems. Part I: foundations, *Artificial Intelligence*, **95**(2), 1997, 257–316.

[13] Gärdenfors, P.: *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, MIT Press, 1988.

[14] Gerbrandy, J.: *Dynamic epistemic logic*, Technical report, ILLC, Amsterdam, 1997.

[15] Gerbrandy, J.: *Bisimulations on Planet Kripke*, Ph.D. Thesis, University of Amsterdam, 1999.

[16] Gerbrandy, J., Groeneveld, W.: Reasoning about information change, *J. of Logic, Language and Information*, **6**(2), 1997.

[17] Godo, L., Hajek, P., Esteva, F.: A fuzzy modal logic for belief functions, *Proc. 17th Int. Joint Conf. on Artificial Intelligence (IJCAI'01)*, 2001.

[18] Hajek, P.: *Metamathematics of fuzzy logic*, Kluwer, 1998.

[19] Halpern, J., McAllester, D.: Likelihood, probability, and knowledge, *Computational Intelligence*, **5**, 1989, 151–160.

[20] Halpern, J., Rabin, M.: A logic to reason about likelihood, *Artificial Intelligence J.*, **32**(3), 1987, 379–405.

[21] Hamblin, C.: The modal 'probably', *Mind*, **68**, 1959, 234–240.

[22] Harel, D.: Dynamic Logic, in: *Handbook of Philosophical Logic* (D. M. Gabbay, F. Günthner, Eds.), vol. II, D. Reidel, Dordrecht, 1984, 497–604.

[23] Herzig, A., Lang, J., Longin, D., Polacsek, T.: A logic for planning under partial observability, *Proc. Nat. (US) Conf. on Artificial Intelligence (AAAI'2000)*, Austin, Texas, August 2000.

[24] Herzig, A., Lang, J., Polacsek, T.: A modal logic for epistemic tests, *Proc. Eur. Conf. on Artificial Intelligence (ECAI'2000)*, Berlin, August 2000.

[25] Herzig, A., Longin, D.: Sensing and revision in a modal logic of belief and action, *Proc. ECAI2002* (F. van Harmelen, Ed.), IOS Press, 2002.

[26] Herzig, A., Longin, D.: On modal probability and belief, *Proc. ECSQARU2003* (N. L. Zhang, T. D. Nielsen, Eds.), 2711, Springer Verlag, 2003.

[27] Hintikka, J. K. K.: *Knowledge and belief*, Cornell University Press, Ithaca, N.Y., 1962.

[28] van der Hoek, W.: On the semantics of graded modalities, *J. of Applied Non-classical Logics (JANCL)*, **2**(1), 1992.

[29] Lenzen, W.: *Recent work in epistemic logic*, North Holland Publishing Company, Amsterdam, 1978.

[30] Lenzen, W.: On the semantics and pragmatics of epistemic attitudes, in: *Knowledge and belief in philosophy and AI* (A. Laux, H. Wansing, Eds.), Akademie Verlag, Berlin, 1995, 181–197.

[31] Levesque, H. J., Reiter, R., Lespérance, Y., Lin, F., Scherl, R.: GOLOG: A logic programming language for dynamic domains, *J. of Logic Programming*, 1997, Special issue on reasoning about action and change.

[32] Lewis, D.: *Counterfactuals*, Basil Blackwell, Oxford, 1973.

[33] Moore, R. C.: A formal theory of knowledge and action, in: *Formal Theories of the Commonsense World* (J. Hobbs, R. Moore, Eds.), Ablex, Norwood, NJ, 1985, 319–358.

[34] Reiter, R.: The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression, in: *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy* (V. Lifschitz, Ed.), Academic Press, San Diego, CA, 1991, 359–380.

[35] Scherl, R., Levesque, H. J.: The frame problem and knowledge producing actions, *Proc. Nat. Conf. on AI (AAAI'93)*, AAAI Press, 1993.

[36] Schmidt, R. A., Tishkovsky, D.: Multi-Agent Logic of Dynamic Belief and Knowledge, *Proc. 8th Eur. Conf. on Logics in AI (JELIA)* (S. Flesca, S. Greco, N. Leone, G. Ianni, Eds.), 2424, Springer, 2002, `http://www.cs.man.ac.uk/~schmidt/publications/SchmidtTishkovsky02c.html%`.

[37] Schmidt, R. A., Tishkovsky, D.: Combining Dynamic Logic with Doxastic Modal Logics, *Advances in Modal Logic, Volume 4* (P. Balbiani, N.-Y. Suzuki, F. Wolter, M. Zakharyaschev, Eds.), King's College London Publications, 2003, To appear, `http://www.cs.man.ac.uk/~schmidt/publications/SchmidtTishkovsky03a.html%`.

[38] Segerberg, K.: Qualitative probability in a modal setting, *Proc. of the 2nd Scandinavian Logic Symp.* (J. Fenstad, Ed.), North Holland Publ. Company, Amsterdam, 1971.

[39] Shapiro, S., Pagnucco, M., Lespérance, Y., Levesque, H. J.: Iterated Belief Change in the Situation Calculus, *Proc. KR2000*, 2000.

[40] Walley, P., Fine, T. L.: Varieties of modal (classificatory) and comparative probability, *Synthese*, **41**, 1979, 321–374.