A Probabilistic Logic for Sequences of Decisions*

Magdalena Ivanovska and **Martin Giese**
Department of Informatics
University of Oslo, Norway

Abstract

We define a probabilistic propositional logic for making a finite, ordered sequence of decisions under uncertainty by extending an existing probabilistic propositional logic with expectation and utility-independence formulae. The language has a relatively simple model semantics, and it allows a similarly compact representation of decision problems as influence diagrams. We present a calculus and show that it is complete at least for the type of reasoning possible with influence diagrams.

1 Introduction

Decision making under uncertainty is a central topic of artificial intelligence, and a number of approaches have been suggested to deal with it, some based on logic [Boutilier, 1994], some on graphical representations like influence diagrams [Howard and Matheson, 1981], some on Markov chains etc. Our research in this area was initially motivated by our work in the CODIO project on COllaborative Decision Support for Integrated Operations.[1] As part of that project, we developed a support system for operational decisions in petroleum drilling using Bayesian networks (BN) modeling [Giese and Bratvold, 2010]. We discovered a number of features that would have been useful in this application, but which are missing in BN-based approaches:

- the possibility to work with incomplete information about probabilities and utilities, e.g. intervals instead of precise numbers
- an explicit treatment of time
- a well-integrated treatment of continuous values instead of just discrete values
- the ability to reason about several decision makers with different knowledge of the situation

These observations prompted us to consider logic and logical deduction as a basis for decision support. First, the semantics of a logic ensures that any unknown information can simply be omitted. Nothing is ever deduced from something that is not explicitly stated. Second, logics are known to be relatively easy to combine. Although we have not done this yet, it is natural to consider combinations of our approach with first-order logic (for reasoning about continuous values), temporal logic, knowledge logic, etc. Additionally, we consider the problem of a logical axiomatization of decision making to be an interesting (theoretical) problem in its own right.

Our first contribution in this spirit was a probabilistic logic with conditional independence formulae [Ivanovska and Giese, 2011] extending the probabilistic logic of [Fagin, Halpern, and Megiddo, 1990]. Expressing (conditional) independence is a prerequisite for a compact representation of probabilistic models, and one of the main reasons for the success of Bayesian networks. We showed that similar compactness and equivalent reasoning can be achieved with a purely logical notation. That work was not concerned with decisions, but only with the modelling of uncertainty.

The present paper extends our previous work by presenting a logic to describe and reason about a fixed, finite sequence of decisions under uncertainty with the aim of maximizing the expected utility of the outcome.

The most closely related existing approach is that of influence diagrams (IDs) [Howard and Matheson, 1981], probably the most successful formalism for modelling decision situations. We show that our logic, together with a suitable calculus, allows to derive all conclusions that an ID model permits. It goes slightly beyond the possibilities of IDs in that it allows more fine-grained statements about conditional independence, namely that some variables are independent only if certain options are chosen.

Compared to other logic-based approaches (see Sect. 8) to treating decision making with a logic, our approach is relatively simple: it incorporates reasoning about multiple decisions, observations, and independence, with a fairly straightforward model semantics, no need for frame axioms, and a rather small inference system. This makes it a good candidate for future combination with other types of reasoning.

*A version of this logic has been published at NIK(Norsk informatikkonferanse) 2011. This new version distinguishes between two different notions of expected utility, and adds reasoning about independence of the expected utility on decision options and uncertainties. Moreover, in this paper we provide a more detailed technical discussion of the logic.

[1] Partially funded by RCN Petromaks, 2007–2009.

2 Sequences of Decisions

To avoid confronting the reader with the technicalities of petroleum engineering, our original application domain, we will use a running example from a more familiar field:

Example 1 *(Organic garden) Consider a situation in which a gardener has to decide whether to apply an antifungal or an antibacterial treatment to her organic garden that shows symptoms of both fungal and bacterial diseases. Based on some examinations, she only knows that there is a 40% chance that the garden is affected by a fungal disease. (We assume that the garden has either a fungal or a bacterial disease, but not both.) Even if she treats the plants with a wrong treatment, there is still a 20% chance that their condition will get better. If the choice of the treatment is right, there is a 90% chance of improvement. A week after applying the chosen treatment the gardener observes the condition of the plants, after which she decides whether to continue with the same treatment, or to switch to the other one. Each decision she makes affects the utility (which is the profit – the earned income decreased by the cost of the treatments) in different ways depending on the actual condition of the garden, which is uncertain. In addition, let's assume that the income of a healthy garden (with two correct treatments applied) is 20; the income of a partially healed garden is 12 when the first treatment is correct and the second is wrong, and is 10 in the opposite case. A garden treated with two wrong treatments gives no income. Also assume that the cost of the antifungal treatment is 2, and the cost of the antibacterial one is 1.*

In general, we consider the scenario that a fixed, finite sequence of n decisions has to be taken. Each of the decisions requires the decision maker to commit to exactly one of a finite set of *options.* We can therefore represent the decisions by a sequence $\mathbf{A} = (A_1, \ldots, A_n)$ of n finite sets that we call *option sets*. For instance, in the above example we have a sequence of two decisions, hence two option sets, $\mathbf{A} = (\{t_1, t_2\}, \{c_1, c_2\})$, one containing two elements representing the two treatment options at the beginning: t_1 – the antifungal, and t_2 – the antibacterial one; and the other containing an element for each of the two possible continuations of the treatment at the second decision point, i.e. c_1 – continue with the same treatment, and c_2 – switch to the other one. There are also elements of uncertainty in the scenario described above, namely the unknown disease of the plants, and the condition of the plants after the first treatment that depends in a non-deterministic way on the disease and the applied treatment. Both of these uncertainties are of a binary type, i.e. they have only two states (fungi or bacteria for the disease, better or worse for the condition), so for each element of the uncertainty, we use a propositional letter to represent one of its states and its negation to represent the other one. Ex: F stands for "fungi", and $\neg F$ stands for "bacteria." In general, we assume we have a set of propositional letters $\mathbf{P} = \{X_1, X_2, \ldots\}$.

Before each of the decisions is taken, some observations might be available to guide the decision maker. E.g., before taking the second decision in the above example, it can be observed whether the condition of the plants is better or worse. Observations are in general represented by subsets of propositional letters, or, more precisely by fixing a sequence $\mathbf{O} = (O_1, \ldots, O_n)$ where each $O_k \subseteq \mathbf{P}$ is a set of *observable propositional letters*, i.e. a set of letters whose value is known before taking the k-th decision. We require this sequence to be monotonic, $O_1 \subseteq \cdots \subseteq O_n$, to reflect that everything that can be observed before each decision, can be observed later. Later, the semantics of expectation formulae (and the EXP rules based on it), will be defined in a way that ensures that observations made before some decision do not change at later decisions, i.e. we model a "non-forgetting decision maker".

We call $\Omega = (\mathbf{P}, \mathbf{A}, \mathbf{O})$ a *decision signature*. In the following, we show how we build our formulae over a given decision signature.

3 Syntax

To express the element of chance in our logic, we follow the approach of [Fagin, Halpern, and Megiddo, 1990]. They define a probabilistic propositional logic by augmenting the propositional logic with *linear likelihood formulae*

$$b_1\ell(\varphi_1) + \cdots + b_k\ell(\varphi_k) \geq b,$$

where $b_1, \ldots, b_k, b$ are real numbers, and $\varphi_1, \ldots, \varphi_k$ are *pure propositional formulae*, i.e. formulae which do not themselves contain likelihood formulae. The term $\ell(\varphi)$ represents the probability of φ being true, and the language allows expressing arbitrary linear relationships between such probabilities. [2]

In Example 1 we can use the formula $\ell(F) = 0.4$ to represent that the probability of the plants having fungal disease is 0.4 (this formula is an abbreviation for $(\ell(F) \geq 0.4) \wedge (-\ell(F) \geq -0.4)$. The probability of the plants' condition after the first week of treatment however is something that we can not express with such formulae. To be able to express probabilistic statements that depend on the decisions that are taken, our logic uses likelihood terms indexed by sequences of decision options. The intention is that these likelihood terms represent the likelihoods of propositional statements being true after some decision making (choosing of options) has taken place. We define general likelihood terms and formulae with the following definitions.

Definition 1 *Given a sequence of option sets* $\mathbf{A} = (A_1, \ldots, A_n)$*, and a subsequence* $\mathbf{S} = (A_{i_1}, \ldots, A_{i_k})$ *for some* $1 \leq i_1 < \cdots < i_k \leq n$*, an* $\mathbf{S}$-option sequence *is a sequence* $\sigma = a_{i_1} \ldots a_{i_k}$ *with* $a_{i_j} \in A_{i_j}$ *for* $j = 1 \ldots k$*. An* $\mathbf{A}$*-option sequence is also called a* full option sequence*.*

In the following text, we will use σ to denote option sequences, and δ for full option sequences.

We introduce the likelihood term $\ell_\delta(\varphi)$ to represent the likelihood of φ after the options in δ (all decisions) have

[2] They actually use w (as "weight") instead of ℓ in [Fagin, Halpern, and Megiddo, 1990] since they use it to represent both probability and inner probability there. We use ℓ as in [Halpern, 2003] since, as we will see later in the semantics, we only consider the measurable probability structures and in that case ℓ stands for likelihood (probability).

taken place. Sometimes the likelihood of a statement does not depend on all the choices one makes, but just on a subset of them, so we give a more general definition of a likelihood term and likelihood formulae:

Definition 2 *A* general likelihood term *is defined as:*

$$\ell_\sigma(\varphi),$$

where σ is an option sequence, and φ is a pure propositional formula. A linear likelihood formula has the following form:

$$b_1\ell_{\sigma_1}(\varphi_1) + \cdots + b_k\ell_{\sigma_k}(\varphi_k) \geq b, \quad (1)$$

where $\sigma_1, \ldots, \sigma_k$ are $\mathbf{S}$-option sequences for the same subsequence $\mathbf{S}$ of $\mathbf{A}$, $\varphi_1, \ldots, \varphi_k$ are pure propositional formulae, and $b, b_1, \ldots, b_k$ are real numbers. [3]

A general likelihood term represents the likelihood (probability) of φ being true, if the options in σ are chosen; the linear likelihood formula represents a linear relationship between such likelihoods, and implies that that relationship holds independently of the options taken for any decision not mentioned in the σ_is. The definition is restricted to option sequences for the same sequence of option sets $\mathbf{S}$, since it is difficult to define a sensible semantics without this restriction. For instance, in the context of the organic garden example, the formula $2\ell_{t_1}(B \wedge F) + 0.5\ell_{t_2}(B) \geq 2$ is a well-formed likelihood formula; whilst $2\ell_{t_1}(B \wedge F) + 0.5\ell_{t_1c_1}(B) \geq 2$ is not.

We can also define conditional likelihood formulae as abbreviations, like [Fagin, Halpern, and Megiddo, 1990] do: $\ell_\sigma(\varphi|\psi) \geq (\leq) c$ iff $\ell_\sigma(\varphi \wedge \psi) - c\ell_\sigma(\psi) \geq (\leq) 0$ where σ is an option sequence, and φ and ψ are pure propositional formulae. $\ell_\sigma(\varphi|\psi) = c$ is defined as a conjunction of the corresponding two inequality formulae. Now we can represent the statement about the probability of the plants getting better after applying treatment t_1 conditional on the fact that they had fungal disease with the formula $\ell_{t_1}(B|F) = 0.9$.

It is well-known, e.g. from the literature on Bayesian networks, that the ability to express conditional independence between events can lead to very compact representations of the joint probability distributions of sets of events. Therefore, to the language of propositional and linear likelihood formulae defined so far, we add conditional independence formulae (CI-formulae) like the ones proposed by [Ivanovska and Giese, 2011], but indexed with option sequences. Their general form is the following:

$$I_\sigma(\mathbf{X_1}, \mathbf{X_2}|\mathbf{X_3}), \quad (2)$$

where $\mathbf{X_i}$, for $i = 1, 2, 3$ are sets of propositional letters, and σ is an option sequence. It expresses that knowledge about the propositions in $\mathbf{X_2}$ does not add knowledge about the propositions in $\mathbf{X_1}$ whenever the value of the propositions in $\mathbf{X_3}$ is known and the options in σ are chosen.

Since our logic is intended to describe decision problems that contain an element of uncertainty, we follow the standard approach of decision theory, which is to model a rational decision maker as an expected utility maximizer. To reason about the expected utility, we need to introduce a new kind of formulae. Halpern in [Halpern, 2003] shows how reasoning about the expected values of random variables can be included in a logic similarly to linear likelihood terms. We cannot use this approach directly however, since we need to include (1) the possibility to condition on observations made before taking decisions, and (2) the principle of making utility maximizing decisions. On the other hand, we only need to consider the expected value of one random variable, namely the utility.

For a full option sequence δ, we introduce the term $e_\delta(\varphi)$ to represent the expected utility conditional on a fact φ. But sometimes the expected utility can be independent of some of the decisions. We represent that by introducing more general expectation terms that do not necessarily include a full option sequence in the subscript. We give the following formal definition of expectation formulae:

Definition 3 *An* expectation formula *is a formula of type:*

$$e_\sigma(\varphi) = c, \quad (3)$$

where σ is an option sequence, φ is a pure propositional formula, and c is a real number.

E.g., for representing the expected utility from a garden that had a fungal disease, and received two weeks of antifungal treatment, we can write $e_{t_1c_1}(F) = 16$ (which is the profit of the garden in this case, i.e. the income of 20 decreased by the cost of the two antifungal treatments). Missing decisions (i.e. non-full option sequences) in the expectation formulae of Def. 3 are intended to express that the expected utility is *independent* of the options chosen for the missing decisions.

To reason about optimal decisions, we also need a different notion, namely the expected utility under the assumption that *optimal* (i.e. utility maximizing) decisions will be made for all future decisions. The following definition introduces "optimal expectation" terms $\bar{e}_{a_1\ldots a_k}(\varphi)$ to capture this idea. They denote the expected utility, conditional on φ, after the initial options $a_1, \ldots, a_k$ have been chosen, assuming that all future choices are made in such a way that the expected utility is maximized. Unfortunately, it turns out to be difficult to define the semantics of such formulae for arbitrary φ. To obtain a useful semantics, the formula φ that is conditioned upon has to be required to be an "observable" formula.

Definition 4 *Given a propositional letter X, an X-*literal *is either X or $\neg X$. An S-*atom *for some set $S \subseteq \mathbf{P}$ is a conjunction of literals containing one X-literal for each $X \in S$.*

An optimal expectation formula *is a formula of type:*

$$\bar{e}_{a_1\ldots a_k}(\varphi) = c, \quad (4)$$

where $a_i \in A_i$, $i = 1, \ldots, k$, φ is an O_k-atom, and c is a real number.

For instance, $e_{t_1}(\top) = 4$ means that the expected utility when treatment t_1 is chosen is 4, no matter how the treatment is continued.[4] This is clearly not the case in our example. On the other hand, $\bar{e}_{t_1}(\top) = 4$ means that the expected

[3] the resulting logic up to these formulae is structurally similar to the multi-agent probabilistic logics of [Halpern, 2003], but with option sequences instead of agents as modalities.

[4] $\top$ is an O_k-atom, when $O_k = \emptyset$ (empty conjunction).

utility for t_1 is 4, if it is followed by the best choice for $\{c_1, c_2\}$, given the observation of B.

As already noted, expectation formulae (3) have the concept of independence of the expected utility on some of the decisions already embedded in their index. In some contexts, like we will see later in Sect. 5, it's useful to be able to express the independence of the expected utility on some facts (propositional letters) as well. As in the case of conditional-independence of likelihood, we choose to talk about utility independence directly by introducing utility independence formulae in the logical language:

$$UI(\mathbf{X}|\mathbf{Y}) \tag{5}$$

where $\mathbf{X}$, $\mathbf{Y}$ are disjoint sets of propositional letters. It says that the (expected) utility is independent of (the "truth condition" of) the set of "facts" $\mathbf{X}$, knowing a set of "facts" $\mathbf{Y}$.

In the gardening example, we can say that the expected utility is independent of the condition of the garden after the first treatment, if we know what the disease was, since it is the disease, not the observation, that determines the outcome of the possible continuations of the treatment. We represent this with the formula $UI(B|F)$.

We conclude this section with the following definition.

Definition 5 *Let the* decision signature $\Omega = (\mathbf{P}, \mathbf{A}, \mathbf{O})$ *be given. The language consisting of all of the propositional formulae, linear likelihood formulae type* (1)*, conditional-independence formulae type* (2)*, expectation formulae type* (3)*, optimal expectation formulae type* (4)*, utility independence formulae* (5) *over the decision signature* Ω*, as well as any Boolean combination of the above, will be denoted by* $\mathbf{L}(\Omega)$.

4 Semantics

In the following, we give a *model semantics* for our logic. It is built around a notion of *frames* which capture the mathematical aspects of a decision situation independently of the logical language used to talk about it. These frames are then extended to *structures* by adding an interpretation function for the propositional letters.

Our semantics is based on the probabilistic structures of [Halpern, 2003], with two modifications: a) the probability measure on the set of worlds (i.e. the possible outcomes) depends on the sequence of options chosen, and b) each world is assigned a utility.

Definition 6 *Let the sequence of* n *option sets* $\mathbf{A}$ *be given, and let* Δ *be the set of all full option sequences. A* probabilistic decision frame (for reasoning about n decisions) *is a triple*

$$(W, (\mu_\delta)_{\delta\in\Delta}, u)$$

where W *is a set of worlds,* μ_δ*, for every* $\delta \in \Delta$*, is a probability measure on* 2^W*, and* $u : W \to R$ *is a utility function.*

To interpret linear likelihood formulae (1), conditional independence formulae (2), expectation formulae (3), and utility-independence formulae (5), we add an interpretation function to these frames. A further restriction will be needed for the interpretation of optimal expectation formulae (4), see Def.(13).

Definition 7 *A* probabilistic decision structure *is a tuple*

$$M = (W, (\mu_\delta)_{\delta\in\Delta}, u, \pi)$$

where $(W, (\mu_\delta)_{\delta\in\Delta}, u)$ *is a probabilistic decision frame, and* π *is an* interpretation function *which assigns to each element* $w \in W$ *a truth-value function* $\pi_w : \mathbf{P} \to \{0, 1\}$.

The interpretation of the linear likelihood formulae (1) is defined in the following way:

$\pi_w(b_1\ell_{\sigma_1}(\varphi_1) + \cdots + b_k\ell_{\sigma_k}(\varphi_k) = b) = 1$ iff

$b_1\mu_{\delta_1}(\varphi_1^M) + \cdots + b_k\mu_{\delta_k}(\varphi_k^M) = b$ for *every* choice of full option sequence δ_j, $j = 1, \ldots, k$, satisfying the conditions:

- σ_j is a subsequence of δ_j;
- if σ_j are $\mathbf{S}$-option sequences, for a subsequence $\mathbf{S}$ of $\mathbf{A}$, then all δ_j agree on the options belonging to sets not in $\mathbf{S}$.

In other words, the linear relationship between the likelihoods has to hold independently of the choices made for any decisions not mentioned in the formula, and which therefore are not contained in $\mathbf{S}$.

Also note that the interpretation of likelihood formulae does not depend on the world w, since statements about likelihood always refer to the entire set of worlds rather than any particular one. Nevertheless, to keep the interpretation general, we will always define the validity of a formula in a certain structure as dependent on a certain world in that structure.

For the semantics of the conditional-independence formulae (2) we extend the definition of the standard probabilistic notion of (conditional) independence of events to define (conditional) independence of sets of events:

Definition 8 *Given a probability space* $(W, 2^W, \mu)$*, we say that events* A *and* B *are* independent conditional on *an event* C*,* $I_\mu(A, B|C)$*, iff*

$\mu(B \cap C) \neq 0$ *implies* $\mu(A|C) = \mu(A|B \cap C)$

or, equivalently,

$\mu(A \cap C) \neq 0$ *implies* $\mu(B|C) = \mu(B|A \cap C)$.

The sets of events $\mathbf{A}_i$, $i = 1, 2$ are conditionally independent given the set of events $\mathbf{A}_3$*,* $I_\mu(\mathbf{A}_1, \mathbf{A}_2|\mathbf{A}_3)$*, iff*

$I_\mu(B_1, B_2|B_3)$ *for all intersections* $B_i = \bigcap_{A\in\mathbf{A}_i} A^{(C)}$ *of possibly complemented events from* $\mathbf{A}_i$*,* $i = 1, 2, 3$.

The interpretation of CI-formulae is then defined by:

$\pi_w(I_\sigma(\mathbf{X}_1, \mathbf{X}_2|\mathbf{X}_3)) = 1$ iff

$I_{\mu_\delta}(\mathbf{X}_1^M, \mathbf{X}_2^M|\mathbf{X}_3^M)$, where $\mathbf{X}_i^M := \{X^M \mid X \in \mathbf{X}_i\}$, for *every* full option sequence δ extending σ.

Before we can give the interpretation of the expectation and optimal expectation formulae, we have to define some semantic concepts within the probabilistic decision frames. We start by recalling the definition of *(conditional) expectation* from probability theory:

Definition 9 *Let (W, F, μ) be a probability space, and $X : W \to R$ be a random variable. The expected value of X (the expectation of X) with respect to the probability measure μ, $E_\mu(X)$, is defined as:*

$$E_\mu(X) = \sum_{w \in W} \mu(w) X(w). \tag{6}$$

For $B \in F$, such that $\mu(B) \neq 0$, the conditional expectation of X with respect to μ conditional on B is given by $E_\mu(X|B) = E_{\mu|B}(X)$.

This notion is sufficient to interpret the expectation formulae (3):

$\pi_w(e_\sigma(\varphi) = c) = 1$ iff

$\mu_\delta(\varphi^M) = 0$ or $E_{\mu_\delta}(u|\varphi^M) = c$, for every δ that extends σ;

And the utility independence formulae (5):

$\pi_w(UI(\mathbf{X}|\mathbf{Y})) = 1$ iff $\mu_\delta(\varphi^M \cap \psi^M) \neq 0$ or $E_{\mu_\delta}(u|\psi^M) = E_{\mu_\delta}(u|\varphi^M \cap \psi^M)$, for every $\mathbf{X}$-atom φ, $\mathbf{Y}$-atom ψ, and for every $\delta \in \Delta$.

To be able to interpret the optimal expectation formulae $\bar{e}_{a_1 \ldots a_k}(\varphi) = c$, where only some initial number of options is fixed, we need to incorporate the idea that the decision maker will pick the best (i.e. expected utility maximizing) option for the remaining decisions. This is captured by the notion of *optimal expected value* which is defined below. The definition relies on a notion of successively refined observations, such that 1. the conditional expectations may only be conditional on observed events, and 2. the probability of an observation is not influenced by decisions taken after the observation. We give the formal definitions in what follows:

Definition 10 *Given a set of worlds W, an* event matrix *of length n for W is a sequence $\mathbf{B} = (B_1, \ldots, B_n)$ where each $B_i \subseteq 2^W$ is a partition of W, and B_{i+1} is a refinement of B_i for $i = 1, \ldots, n-1$.*

The successive refinement captures the idea of an increasing amount of observed information on a semantic level. To capture the fact that observations are not influenced by future decisions, we require $\mathbf{B}$ to be *regular* with respect to the frame F:

Definition 11 *Given a frame $F = (W, (\mu_\delta)_{\delta \in \Delta}, u)$, we call an event matrix $\mathbf{B} = (B_1, \ldots, B_n)$ for W* regular *w.r.t. F if*

$$\mu_{a_1 \ldots a_{k-1} a_k \ldots a_n}(B) = \mu_{a_1 \ldots a_{k-1} a'_k \ldots a'_n}(B), \tag{7}$$

for every $k = 1, \ldots, n$, every $B \in B_k$, and for every $a_i \in A_i$, $i = 1, ..., n$, and $a'_i \in A_i$, $i = k, \ldots, n$.

If (7) holds, we can define new probability measures on B_k, for $k = 1, \ldots, n$, as restrictions:

$$\mu_{a_1 \ldots a_{k-1}}(B) := \mu_{a_1 \ldots a_{k-1} a_k \ldots a_n}(B), \tag{8}$$

for every $B \in B_k$.

Definition 12 *Let $F = (W, (\mu_\delta)_{\delta \in \Delta}, u)$, be a probabilistic decision frame and $\mathbf{B} = (B_1, \ldots, B_n)$ an event matrix for W that is regular w.r.t. F.*

Now, the optimal expected value *of the option sequence $a_1 \ldots a_k$ under an event $B \in B_k$, with respect to F and $\mathbf{B}$, is defined in the following recursive way:*

For $k = n$:

$$\bar{E}^{F,\mathbf{B}}_{a_1 \ldots a_n}(B) := E_{\mu_{a_1 \ldots a_n}}(u|B) \tag{9}$$

For $k = n-1, \ldots, 0$:

$$\bar{E}^{F,\mathbf{B}}_{a_1 \ldots a_k}(B) := \sum_{\substack{B' \in B_{k+1} \\ B' \subseteq B}} \mu_{a_1 \ldots a_k}(B'|B) \cdot \max_{a \in A_{k+1}} \{\bar{E}^{F,\mathbf{B}}_{a_1 \ldots a_k a}(B')\} \tag{10}$$

where $\mu_{a_1 \ldots a_k}$, for every $k = 0, \ldots, n-1$ are the probability measures defined in (8) above.

To complete this definition, we define the optimal expected value in the following special cases:

- *If $\mu_{a_1 \ldots a_k}(B) = 0$ then $\bar{E}^{F,\mathbf{B}}_{a_1 \ldots a_k}(B)$ is not defined and it doesn't count in (10);*
- *If $B_{k+1} = B_k$ then we have: $\bar{E}^{F,\mathbf{B}}_{a_1 \ldots a_k}(B) = \max_{a \in A_{k+1}} \{\bar{E}^{F,\mathbf{B}}_{a_1 \ldots a_k a}(B)\}$.*

We now have the tools we need to give the interpretation of the optimal expectation formulae.

Lemma 1 *For any decision signature $(\mathbf{P}, \mathbf{A}, \mathbf{O})$ and any probabilistic decision structure M, the sequence of sets $\mathbf{O}^M(\mathbf{A}) = (O_1^M, \ldots, O_n^M)$, where $O_k^M := \{\psi^M \mid \psi$ is an O_k-atom$\}$, is an event matrix.*

The proof of Lemma 1 follows immediately from the nesting $O_1 \subseteq \cdots \subseteq O_n$ of observable propositional letters.

Definition 13 *Let a decision signature $(\mathbf{P}, \mathbf{A}, \mathbf{O})$ be given. Then a probabilistic decision structure $M = (W, (\mu_\delta)_{\delta \in \Delta}, u, \pi)$ is called* regular *if $\mathbf{O}^M(\mathbf{A})$ is a regular event matrix for W.*

We interpret the optimal expectation formulae (4) in a regular structure $M = (F, \pi)$, $F = (W, (\mu_\delta)_{\delta \in \Delta}, u)$, in the following way:[5]

[5]The case with zero probability is similar to the corresponding case of conditional likelihood formulae. Namely, in the latter case we also have $\ell(\psi|\varphi) \geq c$ vacuously true for any c, when $\mu(\varphi^M) = 0$, i.e. when $\ell(\varphi) = 0$ is true. Whilst in that case it can be interpreted as "when conditioning on impossible, anything is possible," in the case of expectation formulae, it can be read as "if we know something that is not possible, then we can expect anything." Note also that BNs and IDs exhibit the same behavior when probabilities conditional to impossible events are given.

$\pi_w(\bar{e}_{a_1 \ldots a_k}(\varphi) = c) = 1$ iff

$\mu_{a_1 \ldots a_k}(\varphi^M) = 0$ or $\bar{E}^{F,\mathbf{O^M(A)}}_{a_1 \ldots a_k}(\varphi^M) = c.$

This completes the model semantics for our logic. To summarize, we have defined a notion of (regular probabilistic decision) structures, and shown how the truth value of any type of formula of our logic can be determined at any world of any such structure. The interpretation of the pure propositional formulae as well as of Boolean combinations of formulae is defined in the usual way.

Definition 14 *A formula f is a* logical consequence of *or* entailed by *a set of formulae Φ, $\Phi \models f$, if for every structure $M = (W, (\mu_\delta)_{\delta \in \Delta}, u, \pi)$, and every $w \in W$, we have:*

$(M, w) \models \Phi$ *implies* $(M, w) \models f$.

To see why this notion is sufficient for decision making, let's assume that Φ contains all the formulae that represent the given facts about the organic garden problem. And that treatment t_1 has already been chosen, and the observation B ("better") has been made before deciding between c_1 and c_2. If we determine that $\Phi \models \bar{e}_{t_1 c_1}(B) = 11$ and $\Phi \models \bar{e}_{t_1 c_2}(B) = 8.5$, (we show how to infer these facts syntactically in the following sections) then we know that the expected utility of taking c_1 is larger than that of c_2, and therefore c_1 is the optimal decision option in this case.

5 Influence Diagrams

Influence diagrams [Howard and Matheson, 1981] are the most prominent formalism for representing and reasoning about fixed sequences of decisions. IDs consist of a qualitative graph part, which is complemented by a set of tables giving quantitative information about utilities and conditional probabilities. We will show that our formalism subsumes influence diagrams in the sense that it allows to represent problems given as IDs as sets of formulae, using a similar amount of space as required by the ID. In Sect. 6, we will give a calculus for our logic that allows to derive the same statements about expected utilities as would be derived by reasoning on the ID. Conversely, our logic allows expressing and reasoning about some kinds of facts that are not supported by influence diagrams: inequalities on probabilities and utilities, and (utility) independence restricted to certain option sequences.

The graph part of an ID is a directed acyclic graph in which three different kinds of nodes can occur. The *chance nodes* (drawn as ovals) represent random variables and are associated with the given conditional probability distributions. *Decision nodes* (drawn as rectangles) represent the decisions to be taken. *Value nodes* (drawn as diamonds) are associated with real-valued utility functions. Arcs between decision nodes determine the order in which decisions are taken, and arcs from chance nodes to decision nodes represent that the value of the chance node is known (observed) when the decision is taken. Arcs into chance and value nodes represent (probabilistic) dependency.

$T = t_1$	B
F	0.9
$\neg F$	0.2

$T = t_2$	B
F	0.2
$\neg F$	0.9

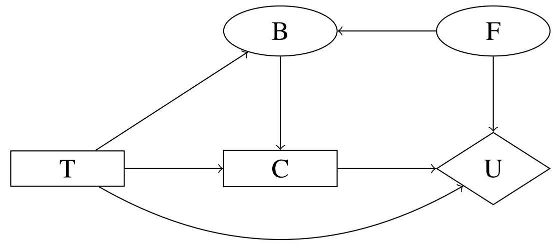

F
0.4

U	t_1		t_2	
	c_1	c_2	c_1	c_2
F	16	9	-2	7
$\neg F$	-4	7	18	9

Figure 1: The influence diagram of the gardening example.

Example 2 *Let us consider an influence diagram on Figure 1, which is a graphical representation of the situation described in example 1. We have two decision nodes to represent the two decisions, T – the choice of a treatment, and then C – the choice for the continuation of the treatment. And two chance nodes, F – for the disease and B – for the condition of the plants after the first week of treatment. The corresponding probability and utility values are given in the tables next to the nodes.*

An influence diagram is said to be *regular* [Shachter, 1986] if there is a path from each decision node to the next one. It is *no-forgetting* if each decision has an arc from any chance node that has an arc to a previous decision. If all the chance nodes of an influence diagram represent binary variables, then we call it a *binary influence diagram*. We can identify a binary chance node X with a propositional letter and denote its two states by X and $\neg X$. We consider here only binary, regular and no-forgetting influence diagrams with only one value node.[6]

We denote the set of parent nodes of a node X by $Pa(X)$ and the set of non-descendants with $ND(X)$. We use $NP(X)$ to denote the set of all nodes that are not parents of X. If we want to single out parents or non-descendants of a certain type, we use a corresponding subscript, for example with $Pa_\circ(X)$ we denote the set of all parent nodes of a node X that are chance nodes, and the set of all parent nodes of X that are decision nodes we denote by $Pa_\Box(X)$.

We can use the formulae of the logical language defined in Sect. 3 to encode influence diagrams.

[6]The restriction to binary nodes is not essential: it is possible to encode nodes with more states by using several propositional letters per node. It would also be straightforward to extend our logic to represent several utility nodes. It may also be possible to extend our framework to allow dropping the regularity and no-forgetting conditions, but we have not investigated this yet.

Example 3 *The decision problem given by the ID in Example 2 can be represented by the following formulae over the sequence of option sets* $\mathbf{A} = (\{t_1, t_2\}, \{c_1, c_2\})$, *with corresponding sets of observables* $\mathbf{O} = (\emptyset, \{B\})$*:*

$\ell_\lambda(F) = 0.4,\ \ell_{t_1}(B|F) = 0.9,\ \ell_{t_1}(B|\neg F) = 0.2,$

$\ell_{t_2}(B|F) = 0.2,\ \ell_{t_2}(B|\neg F) = 0.9,$

where λ *is the empty option sequence, and*

$e_{t_1c_1}(F) = 16,\ e_{t_1c_1}(\neg F) = -4,$

$e_{t_1c_2}(F) = 9,\ e_{t_1c_2}(\neg F) = 7,$

$e_{t_2c_1}(F) = -2,\ e_{t_2c_1}(\neg F) = 18,$

$e_{t_2c_2}(F) = 7,\ e_{t_2c_2}(\neg F) = 9$

$UI(B|F)$ ◁

In general, we encode an influence diagram with a set of formulae that we call its *specific axioms*:

Definition 15 *Let an influence diagram* I *with* n *decision nodes be given. Let* $\mathbf{A}_I$ *be the sequence of option sets determined by the decision nodes of* I, *i.e.* $\mathbf{A}_I = (A_1, \ldots, A_n)$, *and* $\mathbf{O}_I = (O_1, \ldots, O_n)$ *with* $O_i = Pa_\circ(A_i)$, *for every* $i = 1, \ldots, n$. *We define the set of* specific axioms *of* I, $Ax(I)$, *to be the set consisting of the following formulae:*

- $\ell_\sigma(X|\varphi) = c$, *for every chance node* X, *every* $Pa_\circ(X)$*-atom* φ, *and every* $Pa_\Box(X)$*-option sequence* σ, *where* $c = p(X|\varphi, \sigma)$*;*
- $I_\lambda(X, ND_\circ(X)|Pa_\circ(X))$, *for every chance node* X*;*
- $e_\sigma(\varphi) = b$, *for every* $(A_{i_1}, \ldots, A_{i_k})$*-option sequence* σ, $k \leq n$, *where* $\{A_{i_1}, \ldots, A_{i_k}\} = Pa_\Box(U)$, *and every* $Pa_\circ(U)$*-atom* φ, *where* $b = U(\varphi, \sigma)$*;*
- $UI(NP_\circ(U)|Pa_\circ(U))$.

Note that in the encoding in Example 3, we omit the conditional independence formulae $I_\lambda(B, \emptyset|F)$ and $I_\lambda(F, \emptyset|\emptyset)$. They are easily seen to be tautologies, so omitting them makes no difference semantically.

6 Axioms and Inference Rules

While Sect. 4 defines entailment in terms of the model semantics, it does not say how entailment may be checked algorithmically.

In this section, we present an *inference system* that contains separate axioms and inference rules for the different types of reasoning that we want to capture in our logic. This inference system is not complete with respect to the model semantics for the whole logic, but it is sufficient for the entailment of statements of the kind needed for decision making, at least to the same extent as that supported by influence diagrams.

For the propositional reasoning, and reasoning about likelihood, inequalities, and independence, we have the following axiomatic schemes and rules adapted from the ones given by [Fagin, Halpern, and Megiddo, 1990] and [Ivanovska and Giese, 2011]:

Prop All the substitution instances of tautologies in propositional logic,

QU1 $\ell_\sigma(\varphi) \geq 0$

QU2 $\ell_\sigma(\top) = 1$

QU3 $\ell_\sigma(\varphi) = \ell_\sigma(\varphi \wedge \psi) + \ell_\sigma(\varphi \wedge \neg\psi)$, for every pure prop. formulae φ and ψ.

Ineq All substitution instances of valid linear inequality formulae,

MP From f and $f \Rightarrow g$ infer g for any formulae f, g.

QUGen From $\varphi \Leftrightarrow \psi$ infer $\ell_\sigma(\varphi) = \ell_\sigma(\psi)$, for every pure prop. formulae φ and ψ.

SYM From $I_\sigma(\mathbf{X}_1, \mathbf{X}_2|\mathbf{X}_3)$ infer $I_\sigma(\mathbf{X}_2, \mathbf{X}_1|\mathbf{X}_3)$.

DEC From $I_\sigma(\mathbf{X}_1, \mathbf{X}_2 \cup \mathbf{X}_3|\mathbf{X}_4)$ infer $I_\sigma(\mathbf{X}_1, \mathbf{X}_2|\mathbf{X}_4)$.

IND From $I_\sigma(\mathbf{X}_1, \mathbf{X}_2|\mathbf{X}_3)$ and $\ell_\sigma(\varphi_1|\varphi_3) \leq (\geq) a$ infer $\ell_\sigma(\varphi_1|\varphi_2 \wedge \varphi_3) \leq (\geq) a$, where φ_i is an arbitrary $\mathbf{X}_i$-atom, for $i \in \{1, 2, 3\}$.

The rule **Prop** captures the propositional reasoning and can be replaced by any complete set of axioms for propositional logic stated for the formulae of our logical language. **Q1**–**Q3** are based on the defining properties of probability measures. One complete axiomatization of the reasoning about inequalities is given in [Fagin, Halpern, and Megiddo, 1990]; a substitution instances of those axioms with formulae of this logic can replace **Ineq**. **QUGen** reflects the property that equivalent statements have equal likelihoods. Conditional independence reasoning rules for symmetry and decomposition are based on basic properties of conditional independence of sets of random variables, whilst the **IND** rule is based on the definition of conditional independence. As we show in [Ivanovska and Giese, 2011], the above rules are necessary and sufficient for conveying complete reasoning about statements as ones inferred by Bayesian networks.

Here we add the following new rules for reasoning about preservation of likelihood and independence, and about (optimal) expected utilities.

PP From $\ell_\sigma(\varphi|\psi) = b$ infer $\ell_{\sigma'}(\varphi|\psi) = b$, for every option sequence σ' containing σ.[7]

PI From $I_\sigma(\mathbf{X}_1, \mathbf{X}_2|\mathbf{X}_3)$ infer $I_{\sigma'}(\mathbf{X}_1, \mathbf{X}_2|\mathbf{X}_3)$, for every option sequence σ' containing σ.

E1 From $e_\sigma(\varphi) = b$ infer $e_{\sigma'}(\varphi) = b$, for every σ' extending σ.

The last three rules capture the idea that once the probability or independence statement is determined to hold for some choice of decision options, it holds independently of any decisions that are not mentioned in it.

GRE From $e_{a_1\ldots a_k}(\varphi) = c$ infer $\bar{e}_{a_1\ldots a_k}(\varphi) = c$, for $a_i \in A_i$, and φ an $O(A_k)$-atom.

EE $e_\delta(\varphi) = b$ iff $\bar{e}_\delta(\varphi) = b$, for $\delta \in \Delta$ and φ an $O(A_n)$-atom.

GRE and **EE** give the connection between the two types of expectation formulae.

[7] This rule can be extended to arbitrary linear likelihood formulae, if care is taken to extend all occurring option sequences by the same additional options.

EXP1 From $e_\sigma(\varphi_i) = b_i$, $\psi = \varphi_1 \vee \ldots \vee \varphi_m$, $\varphi_i \wedge \varphi_j \leftrightarrow \bot$, for $i, j = 1 \ldots m$, and $b_1\ell_\sigma(\varphi_1) + \cdots + b_m\ell_\sigma(\varphi_m) - b\ell_\sigma(\psi) = 0$, infer $e_\sigma(\psi) = b$.

EXP2 Let ψ be an O_k-atom and $\{\varphi_1, \ldots, \varphi_m\}$ be the set of all O_{k+1}-atoms, such that ψ is a sub-atom of φ_i, $i = 1, \ldots, m$. From $\bar{e}_{a_1 \ldots a_k a}(\varphi_i) = b_{i,a}$, for every $a \in A_{k+1}$, and $b_i = \max_a\{b_{i,a}\}$, for every $i = 1, \ldots, m$, and $b_1\ell_{a_1 \ldots a_k}(\varphi_1) + \cdots + b_m\ell_{a_1 \ldots a_k}(\varphi_m) - b\ell_{a_1 \ldots a_k}(\psi) = 0$, infer $\bar{e}_{a_1 \ldots a_k}(\psi) = b$.

The first expectation rule is based on the definition of mathematical expectation, and the second is the one we use to propagate knowledge about optimal expected utility from further to a closer decision point.

We also introduce the following utility-independence rules:

UI1 From $UI(\mathbf{X}|\mathbf{Y})$ and $e_\sigma(\psi) = b$, infer $e_\sigma(\varphi \wedge \psi) = b$, where $\mathbf{X}$ and $\mathbf{Y}$ are (mutually disjoint) sets of propositional letters, φ and ψ are arbitrary $\mathbf{X}$ and $\mathbf{Y}$-atoms correspondingly, and σ is an arbitrary option sequence.

UI2 From $UI(\mathbf{X_1}|\mathbf{Y})$ infer $UI(\mathbf{X_2}|\mathbf{Y})$, where $\mathbf{X_2} \subset \mathbf{X_1}$.

The soundness of the given axioms and rules mostly follows easily from the semantics of the formulae.

Example 4 *We can use this calculus to derive some conclusions about the situation described in Example 1, and axiomatized in Example 3. If we want to determine the optimal expected utility of taking the option t_1, $\bar{e}_{t_1}(\top)$, we can use the following derivation:*

1. $e_{t_1c_1}(F) = 16$, $e_{t_1c_1}(\neg F) = -4$ .... *(premises)*
2. $UI(B|F)$ .... *(premise)*
3. $e_{t_1c_1}(B \wedge F) = 16$, $e_{t_1c_1}(B \wedge \neg F) = -4$ .. *(1,2, and UI1)*
4. $\ell_{t_1}(F) = 0.4$, $\ell_{t_1}(B|F) = 0.9$, $\ell_{t_1}(B|\neg F) = 0.2$ *(premises and PP)*
5. $\ell_{t_1}(\neg F) = 0.6$ .... *(4, and QU3)*
6. $\ell_{t_1}(B \wedge F) = 0.36$, $\ell_{t_1}(B \wedge \neg F) = 0.12$ ...... *(4,5, def of cond. likelihood, Ineq)*
7. $\ell_{t_1}(B) = 0.48$ .... *(6, QU3)*
8. $\ell_{t_1}(\neg B) = 0.52$ .... *(7, QU3)*
9. $\ell_{t_1c_1}(B \wedge F) = 0.36$, $\ell_{t_1c_1}(B \wedge \neg F) = 0.12$, $\ell_{t_1c_1}(B) = 0.48$ .... *(6, 7, and PP)*
10. $16\ell_{t_1c_1}(B \wedge F) + 7\ell_{t_1c_1}(B \wedge \neg F) - ((16 \cdot 0.36 + (-4) \cdot 0.12)/0.48)\ell_{t_1c_1}(B) = 0$ .... *(9, and Ineq)*
11. $e_{t_1c_1}(B) = 11$ .... *(10, Prop, and EXP1)*
12. $e_{t_1c_2}(B) = 8.5$, $e_{t_1c_1}(\neg B) = -2.46154$, $e_{t_1c_2}(\neg B) = -5.76923$ .... *(obtained similarly to step 11)*
13. $\bar{e}_{t_1c_1}(B) = 11$, $\bar{e}_{t_1c_2}(B) = 8.5$, $\bar{e}_{t_1c_1}(\neg B) = -2.46154$, $\bar{e}_{t_1c_2}(\neg B) = -5.76923$ .... *(12, and GRE)*
14. $\ell_{t_1}(\top) = 1$ .... *(QU2, PP)*
15. $11\ell_{t_1}(B) + (-2.46154)\ell_{t_1}(\neg B) - (11 \cdot 0.48 + (-2.46154) \cdot 0.52)\ell_{t_1}(\top) = 0$ .... *(8,9,13,14 and Ineq)*
16. $\bar{e}_{t_1}(\top) = 3.99999$ .... *(15,Prop, and EXP2)* ◁

The given calculus is not complete in general, which is in part due to the incompleteness of its likelihood and conditional independence part. That issue is discussed in our paper about the combination of likelihoods and conditional independence [Ivanovska and Giese, 2011] and uses the fact shown by Studeny, [Studený, 1992], that a "standard" axiomatization of conditional independence cannot be complete. As observed by [Fagin, Halpern, and Megiddo, 1990], the reasoning about independence requires reasoning about polynomial and not just linear inequalities. The terms $b_i\ell_{a_1 \ldots a_k}(\varphi_i)$ in the EXP2 rule indicate that polynomial inequality reasoning is also required in general to reason about conditional expectation, when no concrete values for the b_i can be derived. Hence we could possibly achieve complete reasoning in our logic by incorporating polynomial instead of only linear inference, but that would make the complexity increase: [Fagin, Halpern, and Megiddo, 1990] show that reasoning about (quantifier-free) polynomial likelihood formulae is possible in PSPACE, by application of the theory of real closed fields, and we know of no better bound. On the other hand, reasoning in our calculus lies within NP, as explained below.

We can however prove the following restricted completeness theorem for entailments corresponding to those possible with an ID.

Theorem 1 *Let I be a given influence diagram with n decision nodes and $Ax(I)$ its set of specific axioms. Then for every $k \in \{1, \ldots, n\}$, every $(A_1, \ldots, A_k)$-option sequence $a_1 \ldots a_k$, and every O_k-atom ψ, there is a real number b such that*

$$Ax(I) \vdash \bar{e}_{a_1 \ldots a_k}(\psi) = b \quad .$$

Proof: We use a backward induction on the length of the option sequence k.

For $k = n$, let $a_1 \ldots a_n$ be a fixed $(A_1, \ldots, A_n)$-option sequence and ψ be an O_n-atom. Let $\varphi_1, \ldots, \varphi_m$ be all of the $Pa_\circ(U)$-atoms. Then $Ax(I)$ contains the formulae $e_{a_1 \ldots a_n}(\varphi_i) = b_i$, for some real numbers b_i, $i = 1, \ldots, m$. Let ψ_i, $i = 1, \ldots, s$ be all $O_n \cup Pa_\circ(U)$-atoms that contain ψ as a subatom. And we have the following derivation steps:

1. $e_{a_1 \ldots a_n}(\varphi_i) = b_i$ .... (premise)(for $i = 1, \ldots, m$)
2. $UI(NP_\circ(U)|Pa_\circ(U))$ .... (premise)
3. $UI(O_n|Pa_\circ(U))$ .... (2, and UI2)
4. $e_{a_1 \ldots a_n}(\psi_i) = b_i$ ..... (1, 3, and UI1)(for $i = 1, \ldots, s$)
5. $\ell_{a_1 \ldots a_n}(\psi_i) = c_i$ ...... (Q1–Q3, PP)(for $i = 1, \ldots, s$)
6. $\ell_{a_1 \ldots a_n}(\psi) = c$ .... (Q1–Q3, PP)

Depending on c we then have the following two groups of possible final steps of this derivation:

For $c \neq 0$:

7. $b_1\ell_{a_1 \ldots a_n}(\psi_1) + \cdots + b_m\ell_{a_1 \ldots a_n}(\psi_s) - (b_1c_1 + \cdots + b_sc_s)/c\ell_{a_1 \ldots a_n}(\psi) = 0$ .... (4, 5, 6, and Ineq)
8. $e_{a_1 \ldots a_n}(\psi) = (b_1c_1 + \cdots + b_sc_s)/c$ .... (7 and EXP1)
9. $\bar{e}_{a_1 \ldots a_n}(\psi) = (b_1c_1 + \cdots + b_sc_s)/c$ ..... (8 and GRE)

For $c = 0$:

7′. $c_i = 0$ .... (3, 4, QU3)(for $i = 1, \ldots, m$)

8′. $b_1\ell_{a_1 \ldots a_n}(\varphi_1) + \cdots + b_s\ell_{a_1 \ldots a_n}(\varphi_s) - b\ell_{a_1 \ldots a_n}(\psi) = 0$ (4, 7′ and Ineq)

9′. $e_{a_1 \ldots a_n}(\psi) = b$ .... (8′ and EXP1)

10′. $\bar{e}_{a_1 \ldots a_n}(\psi) = b$ .... (9′ and GRE)

where b is any real number.

For $k < n$, let us suppose that the assumption holds for every $k+1, \ldots, n$. Let $a_1, \ldots, a_k$ be an arbitrary option sequence such that $a_i \in A_i$, $i = 1, \ldots, k$ and ψ be an arbitrary O_k-atom. Let $\{\varphi_1, \ldots, \varphi_m\}$ be the set of all O_{k+1}-atoms such that ψ is a subatom of φ_i, $i = 1, \ldots, m$. Then we have the following derivation steps:

1. $\bar{e}_{a_1 \ldots a_k a}(\varphi_i) = b_{i,a}$ (IS) (for $i = 1, \ldots, m$, $a \in A_{k+1}$)
2. $b_i = \max_a\{b_{i,a}\}$ .... (Ineq)(for $i = 1, \ldots, m$)

And then we proceed with steps similar to those in the case $k = n$, using (EXP2) instead of (EXP1). **Q.E.D.**

From soundness, and an inspection of the axiomatization $Ax(I)$, we can conclude that the value b must clearly be the same as what would be derived from the ID.

[Fagin, Halpern, and Megiddo, 1990] show that reasoning in probabilistic propositional logic lies within NP. The effort of propagating optimal expected values is comparable to that in influence diagrams. Since precise calculations for influence diagrams are known to be NP-hard, we conclude that our calculus is worst-case optimal for these problems.

7 Decisions under incomplete information

Our logic can easily be extended to allow inequalities for (optimal) expected utility formulas. Unfortunately, the obvious modification of the model semantics does *not* correspond to the notion of decision making under incomplete information about the decision situation: inequalities in the problem specification lead to the selection of a *class* of structures, and the logical consequences are those that hold in all of them. But within each structure, the utilities and likelihoods are precise values, and Def. 12 will reflect optimal decisions for each structure individually. To describe decisions under incomplete information, the information available is important for the decision strategy.

E.g., given $1 \le \bar{e}_{t_1} \le 2$ and $0 \le \bar{e}_{t_2} \le 3$, one decision maker might choose t_1 to minimise the worst-case risk, giving $1 \le \bar{e} \le 2$, another t_2 to maximize the best-case gain, giving $0 \le \bar{e} \le 3$. But choosing the optimal strategy individually for each structure, based on precise values for the optimal expected values, leads to the unsound result $1 \le \bar{e} \le 3$.

It turns out that

- the consequence relation defined in this paper is in general not suited for reasoning about decisions under incomplete information.
- reasoning about decisions under incomplete information is non-monotonic: additional knowledge about expected utilities, etc., can change which decision is preferred by a decision maker, and make previous conclusions unsound.
- this issue can be treated by using as models, instead of the structures described here, sets of such structures. These sets correspond to the decision maker's beliefs about the decision situation, and a suitable semantics can be built in the style of a "logic of only knowing," see e.g. [Chen, 1993].

We will discuss this approach properly in a future publication.

8 Related Work

In our logic, all likelihoods are conceptually indexed by full option sequences, although the formalism allows writing only a subset of the options in formulae. It is tempting to try to reduce the conceptual complexity of the formalism by using propositions to represent the decisions. This has been suggested already by [Jeffrey, 1965], and is taken up e.g. by [Bacchus and Grove, 1996]. However, it requires keeping track of "controlled" versus "non-controlled" variables, and some mechanism is needed to express preference of one option over another. It also gives no immediate solution for the description of observations, and there is an issue with frame axioms. Ultimately, keeping decisions separate from propositions seems to lead to a simpler framework.

Another related line of work in this direction is based on Markov decision processes (MDPs). A MDP is a complete specification of a stochastic process influenced by actions and with a "reward function" that accumulates over time. In contrast to our formalism, MDPs can accommodate unbounded sequences of decisions. [Kwiatkowska, 2003] has investigated model checking of formulae in a probabilistic branching time logic over MDPs. Our approach is not as general, but significantly simpler. We also describe the decision problem itself by a set of formulae and reason about entailment instead of model checking.

Another approach that embeds actions into the logical formalism is the situation calculus, a probabilistic version of which has been described by [Mateus et al., 2001]. This is a very general approach, but the situation calculus is based on second-order logic. Our approach is based on propositional logic, and is therefore conceptually simpler, although it is less general.

We should also point out that our formalism allows more compact representation than most other logic-based approaches, since, similar to IDs, it gives the possibility of expressing independence on both uncertainties and decisions.

9 Conclusion and Future Work

We have argued that a logic-based approach can have advantages over the more common graphical approaches, in particular in combination with reasoning about time, knowledge, continuous values, etc.

As a possible basis for such a logic-based approach, we have described a propositional logic designed to specify a fixed, finite sequence of decisions, to be taken with the aim of maximizing expected utility. Our approach is to let each complete sequence of actions impose a separate probability measure on a common set of worlds equipped with a utility function. The formulae of the logic may refer to only a subset of the decisions, which allows for a more compact representation in the presence of independencies. We have shown how influence diagrams can be mapped into our logic, and we have given a calculus which is complete for the type of inferences possible with IDs. A step beyond IDs would be to allow (optimal) expectation formulae with inequalities, which provides a representation equivalent to a "credal influence diagram" (an analog of a credal network). As discussed in Sect. 7, we are in the process of developing an appropriate extension of our work in the spirit of a logic of only knowing.

The main contribution here is the definition of a suitable syntax, model semantics and calculus. By "suitable" we mean the defined formulae, rules, and semantics are necessary and sufficient to accommodate decision problems we are concerned with (and encode influence diagrams), while still representing general statements about probability, utility, and independence, and reflecting the main aspects of the corresponding types of reasoning.

We consider the main appeal of our logic over other logic-based approaches to be its relative simplicity: it incorporates reasoning about multiple decisions, observations, and independence, with a fairly straightforward model semantics, no need for frame axioms, and a rather small inference system.

The presented logic is intended as a basis for treating more general problems, rather than treating the known ones more efficiently. For that to be achieved, it's necessary to find a way to utilize the linear likelihood inequalities in the decision making inference, and also to include expectation inequality formulae in the logic. In future work, we will also consider the effect of including polynomial and not only linear inequality reasoning. (As we have previously stated, including polynomial reasoning into our logic would improve completeness but at a high price.) This should make it possible to design a calculus that is complete for arbitrary entailment between formulae as given in this paper, and also extensions allowing e.g. comparisons between different expected utility terms. This will put reasoning in the style of "qualitative influence diagrams"[Renooij and van der Gaag, 1998] within the range of our framework.

References

Bacchus, F., and Grove, A. J. 1996. Utility independence in a qualitative decision theory. In *KR*, 542–552.

Boutilier, C. 1994. Toward a logic for qualitative decision theory. In *KR*, 75–86.

Chen, J. 1993. The logic of only knowing as a unified framework for non-monotonic reasoning. In Komorowski, J., and Ras, Z., eds., *Methodologies for Intelligent Systems*, volume 689 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg. 152–161.

Fagin, R.; Halpern, J. Y.; and Megiddo, N. 1990. A logic for reasoning about probabilities. *Information and Computation* 87:78–128.

Giese, M., and Bratvold, R. B. 2010. Probabilistic modelling for decision support in drilling operations, SPE 127761. In *Proc. SPE Intelligent Energy Conference*.

Halpern, J. Y. 2003. *Reasoning about Uncertainty*. MIT Press.

Howard, R., and Matheson, J. 1981. Influence diagrams. In *Readings on the Principles and Applications of Decision Analysis*, volume II. Strategic Decisions Group. 721–762.

Ivanovska, M., and Giese, M. 2011. Probabilistic logic with conditional independence formulae. In *STAIRS 2010*, 127–139. IOS Press.

Jeffrey, R. C. 1965. *The Logic of Decision*. McGraw-Hill.

Kwiatkowska, M. Z. 2003. Model checking for probability and time: from theory to practice. In *LICS*, 351–360.

Mateus, P.; Pacheco, A.; Pinto, J.; Sernadas, A.; and Sernadas, C. 2001. Probabilistic situation calculus. *Ann. Math. Artif. Intell.* 32(1-4):393–431.

Renooij, S., and van der Gaag, L. C. 1998. Decision making in qualitative influence diagrams. In Cook, D. J., ed., *Proc. FLAIRS*, 410–414.

Shachter, R. 1986. Evaluating influence diagrams. *Operations Research* 34(2):871–882.

Studený, M. 1992. Conditional independence relations have no finite complete characterization. In *11th Prague Conf. on Information Theory, Statistical Decision Functions and Random Processes*, volume B, 377–396. Kluwer.