

A Logic-based Approach to Decision Making

Magdalena Ivanovska* and Martin Giese*

Abstract

We propose a novel approach to the well-studied problem of making a finite, ordered sequence of decisions under uncertainty. Most existing work in this area concentrates on graphical representations of decision problems, and requires a complete specification of the problem. Our approach is based on a formal logic, such that the choice of an optimal decision can be treated as a problem of logical deduction. The logical formalism allows to leave unknown aspects of the problem unstated, and it lends itself to various extensions that are difficult to incorporate into graphical approaches. On the other hand, we can show that our formalism can deduce at least the same consequences as are possible with the most popular graphical approach.

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 (e.g. Boutilier, 1994), some on graphical representations like influence diagrams (IDs, 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). One of the difficulties of the modelling endeavour was to elicit precise quantitative assessments of probabilities and potential financial risks and benefits. A BN model requires concrete numbers for every detail however, and there is no way of saying, e.g., “we don’t know this probability” or “the expected cost is between 1 MNOK and 3 MNOK.” Also other shortcomings of the BN approach became evident from this application:

- time is an essential factor when dealing with operational decisions, and BN technology offers only very limited support for temporal modelling.
- In many cases, modeling with continuous values would have been more natural than the discrete states imposed by BN technology.
- An essential aspect of the project was the differing knowledge of the involved decision makers. It would have been useful to reason about which knowledge needs to be available to which actor to arrive at the right decision. Also this is far outside the scope of BN-based methods.

*Dept. of Informatics, University of Oslo, Norway

This paper was presented at the NIK-2011 conference; see <http://www.nik.no/>.

¹Partially funded by the RCN Petromaks programme, 2007–2009.

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, 2011a) extending the probabilistic logic of Fagin et al. (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. Since the rigorous definition of the logic requires an amount of technical detail that would not have been desirable for the current paper, we have decided to produce a separate “extended version” of this paper that includes more rigorous definitions, theorem statements, and proofs (Ivanovska and Giese, 2011b).

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.

The paper is structured as follows: We introduce the syntax of the logic in Sect. 2 and give the model semantics in Sect. 3. In Sect. 4, we briefly review influence diagrams and show how a decision problem described as an ID can be translated into a set of formulae of our logic. Sect. 5 gives an inference system and states a partial completeness result. Sect. 6 reviews related work, and Sect. 7 concludes the paper with some observations about future work.

2 Syntax

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 $\mathbf{A} = (A_1, \dots, A_n)$ of n finite sets that we call *option sets*. For instance, $\mathbf{A} = (\{r_1, r_2\}, \{d_1, d_2\})$ represents a sequence of two decisions, each of which requires to choose one of two options.

The consequences of the decisions taken depend partly on chance, and partly on the state of the world. E.g. the result of drilling an oil well depends on the success of the drilling operation (chance), and whether there actually is oil to be found (state). We can encode the state of the world using a set of Boolean variables $\mathbf{P} = \{X_1, X_2, \dots\}$, which are traditionally called *propositional letters* in formal logic.

Before each of the decisions is taken, some observations might be available to guide the decision maker. These observations also depend on the state of the world and an element of chance. E.g. a seismic test might be available to judge the presence of oil

before drilling, but it is not 100% reliable. Observations can also be represented by propositional letters by fixing a sequence $\mathbf{O} = (O_1, \dots, O_n)$ 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, $O_1 \subseteq \dots \subseteq O_n$, to reflect that everything that can be observed before each decision, can be observed later. Later, the semantics of expectation formulae (and the (EXP) rule 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."

To express the element of chance in our logic, we follow the approach of Fagin et al. (1990). They define a probabilistic propositional logic by augmenting propositional logic with *linear likelihood formulae*

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

where b_1, \dots, b_k, b are real numbers, and $\varphi_1, \dots, \varphi_k$ 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.

To be able to express probabilistic statements that depend on the decisions that are taken, our logic uses likelihood terms indexed by sequences of 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.

Definition 1 *Given the option sets sequence \mathbf{A} , an \mathbf{A} -sequence $a_1 \dots a_n$, ($a_i \in A_i$, $i = 1, \dots, n$) is called a full option sequence and is denoted in general by δ . An option sequence, denoted in general by σ , is an \mathbf{S} -sequence for a sub-sequence \mathbf{S} of \mathbf{A} .*

A general likelihood term is defined to have the form $\ell_\sigma(\varphi)$, where σ is an option sequence, and φ is a pure propositional formula. A linear likelihood formula is:

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

where $\sigma_1, \dots, \sigma_k$ are \mathbf{S} -option sequences for the same sub-sequence \mathbf{S} of \mathbf{A} , $\varphi_1, \dots, \varphi_k$ are pure propositional formulae, and b, b_1, \dots, b_k are real numbers.

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_i s. The definition is restricted to option sequences for the same sequence of option sets \mathbf{S} , since it is unclear what meaning such formulae would have otherwise. For instance, for $\mathbf{A} = (\{r_1, r_2\}, \{d_1, d_2\})$, the formula $2\ell_{r_1}(\varphi_1) + 0.5\ell_{r_2}(\varphi_2) \geq 2$ is a well-formed likelihood formula; whilst $2\ell_{r_1}(\varphi_1) + 0.5\ell_{r_1 d_1}(\varphi_2) \geq 2$ is not.

We can also define conditional likelihood formulae as abbreviations, like Halpern (2003) does:

$$\ell_\sigma(\varphi|\psi) \geq (\leq) c \quad \text{iff} \quad \ell_\sigma(\varphi \wedge \psi) - c\ell_\sigma(\psi) \geq (\leq) 0$$

$\ell_\sigma(\varphi|\psi) = c$ is defined as a conjunction of the corresponding two inequality formulae.

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 (2011a), 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 (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 could introduce formulae of type $e_{\delta} = c$, for a real number c , to represent the expected utility of taking the decisions described by δ . But we are actually interested in representing expected utilities *conditional* on some observations made before taking a decision. We therefore use the more general form $e_{\delta}(\varphi) = c$ to represent the expected utility conditional on the observation φ .

In order to reason about the optimal choice after some, but not all of the decisions have been made, this needs to be generalized further to the form $e_{a_1 \dots a_k}(\varphi) = c$, which will express the expected utility, conditional on φ , after the initial options $a_1 \dots a_k$ 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 2 *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 expectation formula is a formula of type:

$$e_{a_1 \dots a_k}(\varphi) = c, \quad (3)$$

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

As stated before, the term on the left-hand side of (3) represents the expected utility after committing to the options a_1, \dots, a_k , conditional on the observation that φ is true, and then deciding upon the rest of the decisions represented by the option sets A_{k+1}, \dots, A_n such that the expected utility is maximized.

The logical language considered in the remainder of this paper consists of all of propositional formulae, linear likelihood formulae (1), conditional-independence formulae (2), and expectation formulae (3), as well as arbitrary Boolean combinations of the above.

3 Semantics

In the previous section, we have defined the formulae of our logic, but we have yet to give it a semantics that says whether some set of formulae is a consequence of another.

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.

Given any structure M and any formula f , we then go on to define whether f is true in that particular structure. At the end of the current section, we will see how this can be used to define a notion of logical consequence, and how this ultimately can be used to make optimal decisions.

Definition 3 *Let the sequence of n option sets \mathbf{A} be given, and Δ be the set of full option sequences. A probabilistic decision frame (for reasoning about n decisions) is a triple $F = (W, (\mu_\delta)_{\delta \in \Delta}, u)$, where W is a set of worlds, μ_δ , for every $\delta \in \Delta$, is a probability measure on 2^W , and $u : W \rightarrow \mathbb{R}$ is a utility function. A probabilistic decision structure is a quadruple*

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

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 $\pi_w : \mathbf{P} \rightarrow \{0, 1\}$.

We will see later on that the notion of probabilistic decisions structures has to be restricted for the interpretation of expectation formulae (3).

It is straightforward to extend the interpretation π_w from propositional letters to arbitrary propositional formulae. We define the extension $\varphi^M := \{w \mid \pi_w(\varphi) = 1\}$ for any propositional formula φ as the set of worlds where φ is true. The interpretation of the linear likelihood formulae (1) is defined in the following way:

$$\begin{aligned} \pi_w(b_1 \ell_{\sigma_1}(\varphi_1) + \dots + b_k \ell_{\sigma_k}(\varphi_k) \geq b) = 1 \quad \text{iff} \\ b_1 \mu_{\delta_1}(\varphi_1^M) + \dots + b_k \mu_{\delta_k}(\varphi_k^M) \geq b \text{ for every choice of } \delta_j \text{ extending } \sigma_j, \text{ such} \\ \text{that } \delta_j \text{'s agree on all the options from sets not in } S, \text{ where } \sigma_j \text{'s are } S\text{-atoms,} \\ j = 1, \dots, k. \end{aligned}$$

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.²

The interpretation of conditional independence formulae (2) is defined by:

$$\begin{aligned} \pi_w(I_\sigma(\mathbf{X}_1, \mathbf{X}_2 \mid \mathbf{X}_3)) = 1 \quad \text{iff} \\ I_{\mu_\delta}(\mathbf{X}_1^M, \mathbf{X}_2^M \mid \mathbf{X}_3^M), \text{ where } \mathbf{X}_i^M := \{X^M \mid X \in \mathbf{X}_i\}, \text{ for every } \delta \text{ extending } \sigma. \end{aligned}$$

See e.g. Ivanovska and Giese (2011a) for the definition of I_μ , which denotes stochastic independence of two sets of events, conditional on a third set of events.

For the interpretation of the expectation formulae, we define a series of semantic concepts within the probabilistic decision frames. We start by recalling the definition of (*conditional*) *expectation* from probability theory:

Definition 4 *Let (W, F, μ) be a probability space, and $X : W \rightarrow R$ 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:*

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

²Also note that the interpretation of likelihood formulae does not depend on the state w , since statements about likelihood always refer to the entire set of worlds rather than any particular one.

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

$$E_\mu(X|B) = E_{\mu/B}(X).$$

where μ/B is the probability measure that results from μ by conditioning on B , i.e. $\mu/B(X) = \mu(X \cap B)/\mu(B)$.

This notion is sufficient to interpret the expectation formulae indexed by full option sequences. Recall that they should express the expected utility conditional on some formula, given one particular option for each of the decisions to be taken. We define:

$$\pi_w(e_\delta(\varphi) = c) = 1 \text{ iff } \mu_\delta(\varphi^M) = 0 \text{ or } E_{\mu_\delta}(u|\varphi^M) = c.$$

To be able to interpret general expectation formulae $e_{a_1 \dots a_k}(\varphi) = c$, 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 a brief account of the definition here, referring the reader to the long version (Ivanovska and Giese, 2011b) for a more detailed discussion. An *event matrix* is a sequence $\mathbf{B} = (B_1, \dots, B_n)$ where each $B_i \subseteq 2^W$ is a partition of W , and B_{i+1} is a refinement of B_i for $i = 1, \dots, n-1$. 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 , which we define to be the case if

$$\mu_{a_1 \dots a_{k-1} a_k \dots a_n}(B) = \mu_{a_1 \dots a_{k-1} a'_k \dots a'_n}(B), \quad (4)$$

for every $k = 1, \dots, n$, every $B \in B_k$, and for every $a_i \in A_i$, $i = 1, \dots, n$, and $a'_i \in A_i$, $i = k, \dots, n$. If \mathbf{B} is regular w.r.t. F , it is easy to see that a new set of probability measures $\mu_{a_1 \dots a_{k-1}}$ on B_k , for $k = 1, \dots, n$ can be defined as

$$\mu_{a_1 \dots a_{k-1}}(B) := \mu_{a_1 \dots a_{k-1} a_k \dots a_n}(B), \quad (5)$$

for every $B \in B_k$.

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

For $k = n$:

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

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

$$\bar{E}_{a_1 \dots a_k}^{F, \mathbf{B}}(B) := \sum_{B' \in B_{k+1}, B' \subseteq B} \mu_{a_1 \dots a_k}(B'|B) \cdot \max_{a \in A_{k+1}} \{\bar{E}_{a_1 \dots a_k a}^{F, \mathbf{B}}(B')\}, \quad (6)$$

where $\mu_{a_1 \dots a_k}$, for every $k = 0, \dots, n-1$ are the probability measures defined in (5) above.

Now for any probabilistic decision structure M , it is easy to see that the sequence of sets $\mathbf{O}^M = (O_1^M, \dots, O_n^M)$, where $O_k^M := \{\psi^M \mid \psi \text{ is an } O_k\text{-atom}\}$ is an event matrix.

To be able to interpret expectation formulae, we restrict ourselves to probabilistic decision structures $M = (W, (\mu_\delta)_{\delta \in \Delta}, u, \pi)$ where \mathbf{O}^M is a regular event matrix for W , and we call such structures *regular*.

We interpret the expectation formulae in a regular structure $M = (F, \pi)$, $F = (W, (\mu_\delta)_{\delta \in \Delta}, u)$, in the following way:³

$$\pi_w(e_{a_1 \dots a_k}(\varphi) = c) = 1 \quad \text{iff} \quad \mu_{a_1 \dots a_k}(\varphi^M) = 0 \text{ or } \bar{E}_{a_1 \dots a_k}^{F, \mathbf{O}^M}(\varphi^M) = c.$$

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 an arbitrary formula of our logic can be determined in any such structure.

As usual in formal logic, we now define when a formula f is a *logical consequence* of or *entailed* by some set of formulae Φ . Namely, this is the case if every structure M that makes all formulae in Φ true, also makes f true. This is written $\Phi \models f$.

To see why this notion is sufficient for decision making, assume that we have a decision situation for the decisions $(\{r_1, r_2\}, \{d_1, d_2\})$, specified by a set of formulae Φ . Assume that r_1 has already been chosen, and that some observation B has been made before deciding between d_1 or d_2 . If we can now determine that

$$\Phi \models e_{r_1 d_1}(B) = 100 \quad \text{and} \quad \Phi \models e_{r_1 d_2}(B) = 150 \quad ,$$


then we know that the expected utility of taking d_2 is larger than that of d_1 , and therefore d_2 is the optimal decision in this case.

4 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 allows to represent problems given as IDs as sets of formulae, using a similar amount of space as required by the ID. In the next section, 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.

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.

Example 1 Let us consider an influence diagram, which is an abstraction of the well known ‘‘oil wildcatter’’ example (Reiffa, 1968), given by the following figure:



³See the long version (Ivanovska and Giese, 2011b) for a discussion of the case where conditioning on an event with zero probability occurs.

The situation modelled is that an oil wildcatter has to decide whether or not to drill an exploration well. Before making that decision, he may or may not perform a seismic test, which may give a better clue as to whether there is oil to be found, but which also incurs an additional cost. In the diagram, R is the decision whether to perform the test, which is followed by D , the decision whether to drill. The seismic condition T is available as observation for D . The presence of oil is represented by C , which is not directly observed, but which has a probabilistic influence on the seismic condition T . The total utility is determined by the presence of oil, whether a well is drilled, and whether the test is performed.

For the remaining examples, we will work with the following concrete probabilities $p(C) = 0.8$, $p(T|C, r_1) = 0.3$, $p(T|\neg C, r_1) = 0.9$, $p(T|C, r_2) = 0.4$, $p(T|\neg C, r_2) = 0.2$ ⁴ and utilities: $U(r_1, d_1, C) = 10$, $U(r_1, d_1, \neg C) = 7$, $U(r_1, d_2, (\neg)C) = 9$, $U(r_2, d_i, C) = 11$, $U(r_2, d_i, (\neg)C) = 8$.

An influence diagram is said to be *regular* (Schachter, 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 $\neg X$ and X . We consider here only binary, regular and no-forgetting influence diagrams with only one value node.

We denote the set of parent nodes of a node X by $Pa(X)$ and the set of non-descendants with $ND(X)$. 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 $Pa_{\square}(X)$.

We can use the formulae of the logical language defined in Sec. 2 to encode influence diagrams. In what follows we explain how we do the encoding.

In influence diagrams, the utility function can depend directly on some of the decisions and some of the chance variables. In order to achieve the same compactness of representation as in an ID, we would have to represent that the utility is (conditionally) independent on all other nodes. It turns out that this is not quite possible in our formalism. We can however represent independence on those chance nodes which are never observed.

In order to express the utility dependent on some state φ of the chance nodes using an expectation formula $e_{\dots}(\varphi) = c$ of our logic, φ has to be an observable. But typically, the utility will depend on some non-observable nodes. We address this by using an option set sequence \mathbf{A} which contains, in addition to one element for each decision node in the ID, a final “dummy” option set with only one element, written $U = \{\diamond\}$. O_U will be the last element of the monotonic sequence of other sets of observables, with $O_U \setminus O_n$ being a set of propositional letters that represent the chance variables (facts) that have an arc into the utility node.

Example 2 For the decision problem given by the influence diagram in Example 1, we can use the sequence of option sets $\mathbf{A} = (R, D, U)$, with corresponding sets of “observables” $\mathbf{O} = (\emptyset, \{T\}, \{T, C\})$, so we can represent the problem with the following formulae:

⁴The probabilities differ dependent on the options chosen. We write this using “conditioning” on those options.

$$\ell_\lambda(C) = 0.8, \ell_{r_1}(T|C) = 0.3, \ell_{r_1}(T|\neg C) = 0.9, \ell_{r_2}(T|C) = 0.4, \ell_{r_2}(T|\neg C) = 0.2,$$

where λ is the empty option sequence, and

$$\begin{aligned} e_{r_1 d_1 \diamond}(C \wedge \varphi) &= 10, & e_{r_1 d_1 \diamond}(\neg C \wedge \varphi) &= 7, & e_{r_1 d_2 \diamond}((\neg)C \wedge \varphi) &= 9, \\ e_{r_2 d_i \diamond}(C \wedge \varphi) &= 11, & e_{r_2 d_i \diamond}((\neg)C \wedge \varphi) &= 8, & \text{for } \varphi \in \{T, \neg T\}. & \triangleleft \end{aligned}$$

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

Definition 5 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 and the utility node of I , i.e. $\mathbf{A}_I = (A_1, \dots, A_n, U)$, $U = \{\diamond\}$; and $\mathbf{O}_I = (O_1, \dots, O_n, O_U)$ be such that $O_i = Pa_\circ(A_1) \cup \dots \cup Pa_\circ(A_i)$, for every $i = 1, \dots, n$, and $O_U = O_n \cup Pa_\circ(U)$. 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 $Pa_\square(X)$ -option sequence σ , where $c = p(X|\varphi, \sigma)$
- $I_\lambda(X, ND_\circ(X)|Pa_\circ(X))$, for every chance node X ;
- $e_{\sigma \diamond}(\varphi \wedge \psi) = b$, for every (A_1, \dots, A_n) -option sequence σ and every $Pa_\circ(U)$ -atom φ , where $b = U(\varphi, \sigma)$ and ψ is any O_n -atom.

5 Axioms and Inference Rules

While Sect. 3 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* i.e. a set of universally valid formulae (*axioms*) and *inference rules* that allow to infer more valid formulae that are entailed by a given set of formulae. The set of reasoning rules allows to infer the entailment of statements of the kind needed for decision making, at least to the same extent as that supported by influence diagrams.

For propositional reasoning, reasoning about likelihood and reasoning about inequalities, we have the following axiomatic schemes and rules adapted from the ones given by Fagin et al. (1990) and Ivanovska and Giese (2011a):

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 $i \in \{1, 2, 3\}$.

We add the following new rules for reasoning about preservation of likelihood and independence, and about expected utilities.

PP From $\ell_\sigma(\phi|\psi) = b$ infer $\ell_{\sigma'}(\phi|\psi) = b$, for every option sequence σ' containing σ .⁵

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 .

EXP Let ψ be an O_k -atom and $\{\phi_1, \dots, \phi_m\}$ be the set of all O_{k+1} -atoms, such that ψ is a sub-atom of ϕ_i , $i = 1, \dots, m$. From $e_{a_1 \dots a_k a}(\phi_i) = b_{i,a}$, for every $a \in A_{k+1}$, and $b_i = \max_a \{b_{i,a}\}$, for every $i = 1, \dots, m$, and $b_1 \ell_{a_1 \dots a_k}(\phi_1) + \dots + b_m \ell_{a_1 \dots a_k}(\phi_m) - b \ell_{a_1 \dots a_k}(\psi) = 0$, infer $e_{a_1 \dots a_k}(\psi) = b$.

The soundness of the given axioms and rules mostly follows easily from the semantics. The long version (Ivanovska and Giese, 2011b) gives a soundness proof for the EXP rule.

Example 3 We can use this calculus to derive some conclusions about the influence diagram given in Example 1, and axiomatized in Example 2. If we want to determine the expected utility of taking the option r_1 , $e_{r_1}(\top)$, we can use the following derivation:

1. $e_{r_1 d_1} \diamond (C \wedge T) = 10$, $e_{r_1 d_1} \diamond (\neg C \wedge T) = 7$ (premises)
2. $\ell_{r_1}(C) = 0.8$, $\ell_{r_1}(T|C) = 0.3$, $\ell_{r_1}(T|\neg C) = 0.9$ (premises and PP)
3. $\ell_{r_1}(\neg C) = 0.2$ 2, (QU3)
4. $\ell_{r_1}(C \wedge T) = 0.24$, $\ell_{r_1}(\neg C \wedge T) = 0.18$ (2, 3, def of cond. likelihood, Ineq)
5. $\ell_{r_1}(T) = 0.42$ 4 and (QU3)
6. $\ell_{r_1}(\neg T) = 0.58$ 5 and (QU3)
7. $\ell_{r_1}(C \wedge \neg T) = 0.56$ 2, 4, (QU3)
8. $\ell_{r_1 d_1}(C \wedge T) = 0.24$, $\ell_{r_1 d_1}(\neg C \wedge T) = 0.18$, $\ell_{r_1 d_1}(T) = 0.42$ (4, 6, and PP)
9. $10\ell_{r_1 d_1}(C \wedge T) + 7\ell_{r_1 d_1}(\neg C \wedge T) - ((10 \cdot 0.24 + 7 \cdot 0.18)/0.42)\ell_{r_1 d_1}(T) = 0$ (8, Ineq)
10. $e_{r_1 d_1}(T) = 8.71$ (9, Prop, and EXP 2)
11. $e_{r_1 d_2}(T) = 9$, $e_{r_1 d_1}(\neg T) = 9.90$, $e_{r_1 d_2}(\neg T) = 9$ (obtained similarly to step 10)
12. $\ell_{r_1}(\top) = 1$ (QU2, PP)
13. $9\ell_{r_1}(T) + 9.90\ell_{r_1}(\neg T) - (9 \cdot 0.42 + 9.90 \cdot 0.58)\ell_{r_1}(\top) = 0$ (5, 6, 11, 12 and Ineq)
14. $e_{r_1}(\top) = 6.19$ (13, Prop, and EXP 2) ◁

The given calculus is not complete in general, i.e. there are sets of formulas Φ which entail some formula f , but the calculus does not allow to derive f from Φ . This is in part due to the fact that complete reasoning about independence requires reasoning about polynomial and not just linear inequalities. But the terms $b_i \ell_{a_1 \dots a_k}(\phi_i)$ in the **EXP** 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.

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 $k \in \{1, \dots, n\}$, every (A_1, \dots, A_k) -option sequence $a_1 \dots a_k$, and every O_k -atom ψ , there is a real number b such that $e_{a_1 \dots a_k}(\psi) = b$ can be inferred from $Ax(I)$ with the given inference system.*

⁵This rule can be extended to arbitrary linear likelihood formulae, if care is taken to extend all occurring option sequences by the same additional options.

The proof is again given in the long version (Ivanovska and Giese, 2011b). 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.

6 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. This can be an advantage in particular if complete information about the decision problem is not available.

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.

7 Conclusion and Future Work

We have argued how a logic-based approach can have advantages over the more common graphical approaches, in particular when combined with elicitation problems, and 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 towards going beyond IDs would be to extend the current framework by allowing expectation formulae with inequalities with the obvious semantics, which would provide a representation equivalent to a “credal influence diagram”(an analog of a credal network).

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. In future work, we will consider the effect of including polynomial and not only linear inequality reasoning. 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. We will also investigate improved ways of expressing utilities to make better use of independence of the utility of arbitrary decisions and propositional letters.

References

- F. Bacchus and A. J. Grove. Utility independence in a qualitative decision theory. In *KR*, pages 542–552, 1996.
- C. Boutilier. Toward a logic for qualitative decision theory. In *KR*, pages 75–86, 1994.
- R. Fagin, J. Y. Halpern, and N. Megiddo. A logic for reasoning about probabilities. *Information and Computation*, 87:78–128, 1990.
- M. Giese and R. B. Bratvold. Probabilistic modelling for decision support in drilling operations, SPE 127761. In *Proc. SPE Intelligent Energy Conference*, 2010.
- J. Y. Halpern. *Reasoning about Uncertainty*. MIT Press, 2003.
- R. Howard and J. Matheson. Influence diagrams. In *Readings on the Principles and Applications of Decision Analysis*, volume II, pages 721–762. 1981.
- M. Ivanovska and M. Giese. Probabilistic logic with conditional independence formulae. In *STAIRS 2010*, pages 127–139. IOS Press, 2011a.
- M. Ivanovska and M. Giese. A logic-based approach to decision making (long version). Technical Report 411, Dept. of Informatics, Univ. of Oslo, Norway, 2011b.
- R. C. Jeffrey. *The Logic of Decision*. McGraw-Hill, 1965.
- M. Z. Kwiatkowska. Model checking for probability and time: from theory to practice. In *LICS*, pages 351–360, 2003.
- P. Mateus, A. Pacheco, J. Pinto, A. Sernadas, and C. Sernadas. Probabilistic situation calculus. *Ann. Math. Artif. Intell.*, 32(1-4):393–431, 2001.
- H. Reiffa. *Decision analysis*. Addison-Wesley, 1968.
- S. Renooij and L. C. van der Gaag. Decision making in qualitative influence diagrams. In D. J. Cook, editor, *Proc. FLAIRS*, pages 410–414, 1998.
- R. Schachter. Evaluating influence diagrams. *Operations Research*, 34(2):871–882, 1986.