Aalborg Universitet

AALBORG UNIVERSITY
DENMARK

Multi-currency Influence Diagrams

Nielsen, Søren Holbech; Nielsen, Thomas Dyhre; Jensen, Finn V.

Published in:
Advances in Probabilistic Graphical Models

Publication date:
2007

Document Version
Early version, also known as pre-print

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Nielsen, S. H., Nielsen, T. D., & Jensen, F. V. (2007). Multi-currency Influence Diagrams. In P. Lucas, J. Gamez, & A. Salmeron (Eds.), *Advances in Probabilistic Graphical Models* (Vol. 213, pp. 275-294). Springer. Studies in Fuzziness and Soft Computing Vol. 213

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Multi-currency Influence Diagrams

Søren Holbech Nielsen, Thomas D. Nielsen, and Finn V. Jensen

Aalborg University, Aalborg, Denmark

Abstract. When using the influence diagrams framework for solving a decision problem with several different quantitative utilities, the traditional approach has been to convert the utilities into one common currency. This conversion is carried out using a tacit transformation, under the assumption that the converted problem is equivalent to the original one. In this paper we present an extension of the influence diagram framework. The extension allows for these decision problems to be modelled in their original form. We present an algorithm that, given a linear conversion function between the currencies of the original utilities, discovers a characterisation of all other such functions, which induce the same optimal strategy. As this characterisation can potentially be very complex, we give methods to present it in an approximate way.

1 Introduction

Influence diagrams (IDs) were introduced by [4] as a compact modelling language for decision problems with a single decision maker (DM). When a decision problem is represented using the ID framework, the specification rests on two principal components: A graphical structure for capturing the qualitative part of the domain, and quantitative information in the form of probabilities for representing uncertainty and utilities for representing preferences.

The separation of the qualitative and quantitative part of the ID is one of the appealing properties of IDs when considered as a modelling tool. First, it helps the modeller to focus on structure rather than calculations, and second, the structure emphasises the local relations, which govern the specification of the probabilities. Unfortunately, this locality principle does not completely extend to the specification of the utility function: The utility function is usually specified through a collection of local utility functions, which appear as the additive components of the global utility function. This implies that all local utility functions should appear on the same scale. For instance, in a medical domain money and discomfort would need to be transformed onto a common scale. For decision problems where several parties are affected by the decision making process, the utility functions for the individual stakeholders would also have to be transformed into a single utility function before reasoning can take place. Unfortunately, it is usually difficult to elicit the parameters, which governs such transformations (e.g. what is the monetary cost

of one unit of discomfort?).¹ Moreover, the nature of such a transformation has a direct impact on the solution of the ID, but this effect cannot be made transparent when the transformation is tacit. Furthermore, this type of uncertainty, or ignorance, is not easily represented in the model.

In this paper we propose a framework, termed multi-currency IDs (MCIDs), for representing decision problems with local utility functions of different currencies.² An MCID can be seen as an ID augmented with currency information for the local utility functions. We propose an algorithm that, based on an MCID representation of a decision problem and a solution corresponding to a given set of currency transformation parameters, provides a characterisation of all combinations of parameters, which would give rise to the same solution – a solution being an optimal strategy composed of a policy for each decision. The result of such an analysis thus provides an indication of how robust the optimal strategy is in terms of the tacit conversion parameters. Regular sensitivity analysis (see e.g. [3] and [8]), on the other hand, provides robustness measures in terms of isolated deviances in one or more of the actual utility values in the ID, with no regard to how the uncertainties in these parameters are related. By encoding each value for which regular sensitivity analysis is to be performed by its own currency, the analysis could be performed by the method we propose here, and it can therefore be seen as a generalization of regular sensitivity analysis.

As the result of the analysis may be quite complex, we provide, in addition to this algorithm, methods for presenting the result to a DM in a comprehensible manner.

Example 1 (A Motivating Example). The ID in Fig. 1 models a decision problem where a doctor is faced with a patient. The health $Health_1$ of the patient at the time of the initial consultation is revealed only indirectly by the symptoms $Symptoms$ exhibited by the patient. Based on the symptoms, the doctor must decide whether to perform a test $Test$. The test will produce a result $Result$, which the doctor observes before deciding on a treatment (if any) $Treat$. The health of the patient, after a possible treatment has been administered, is represented by the variable $Health_2$. As medical supplies are expensive both the $Test$ and $Treat$ decisions are associated with a monetary cost represented by the utility nodes U_2 and U_3 . Furthermore, if a test is performed, it might be associated with a degree of pain, risk of mortality, or other side effects on behalf of the patient. This is represented by the utility node U_1 . The $Health_2$ variable also has a preferred state (corresponding to the patient being well), which is encoded by the utility node U_4 .

¹ [12] considers utilities that are defined as a linear combination of cost, insult and risk, for instance.

² Even though the word currencies is used here, we are not restricting ourselves to monetary currencies, but consider human lives, spare time etc. as currencies also. Also, in decision problems with several stakeholders, each currency can be seen as the utility for one specific stakeholder.

Now, before the doctor calculates an optimal strategy for his decision prob-

Fig. 1. Example of an ID

lem, he needs to transform U_1 , U_2 , U_3 , and U_4 onto a common scale – say, dollars. This involves estimating the monetary equivalents of the patient being ill and of him being subjected to a painful test. However, different transformations might produce differing optimal strategies. Therefore it would be advantageous to know:

1. What choice of conversion parameters has been the basis of the calculated optimal strategy?
2. What other conversion parameters would produce the same optimal strategy? If another stakeholder, such as the patient in this example, disagrees with the parameters, we could guarantee that even though there is disagreement on the exact choice of parameters, the identified set of parameters all render the same strategy optimal.

None of these questions can be answered from the ID alone.

As an example of how decision problems involving several stakeholders can be interpreted as a multi-currency problem, the ID introduced above can be seen as describing a conflict between the interests of the patient and the hospital (and/or the patient’s medical insurance company) mediated by the doctor: By letting the patient specify utility values corresponding to utility nodes U_1 and U_4 , and the hospital the ones for U_2 and U_3 , we can investigate how the decisions of the doctor relate to these two stakeholders; he may be indulgent to please the patient more than the hospital or vice versa.

2 Influence Diagrams

An ID is a directed acyclic graph consisting of *chance* ($\mathbf{V}_C$), *decision* ($\mathbf{V}_D$), and *utility nodes* ($\mathbf{V}_U$), with the two constraints that a node has no children if and only if (iff) it is a utility node, and that there is a directed path encompassing all decision nodes in the diagram. A chance node (drawn as an ellipse) represents a *chance variable*, which is a discrete variable outside

the DM's direct control. A decision node (drawn as a rectangle) represents a *decision variable* (or simply a *decision*), which is a discrete variable under the DM's direct control. A utility node (drawn as a diamond) represents a local utility function, and the set of local utility functions constitute the components in an additive factorisation of the global utility function [13]. When the meaning is obvious from the context, we use the terms node and variable interchangeably and do not distinguish between a variable or local utility function and the node representing it.

When speaking of IDs we denote by $\mathbf{pa}(V)$ the set of nodes, which are parents of the node V , and by $\mathbf{sp}(V)$ we denote the set of *states* of variable V – states being outcomes for chance variables and decision options for decisions. For a set of variables, $\mathbf{S}$, we denote by $\mathbf{sp}(\mathbf{S})$ the configurations $\times_{V \in \mathbf{S}} \mathbf{sp}(V)$.

An arc in an ID represents either functional dependence, probabilistic dependence, or temporal precedence, depending on the type of node it goes into. In particular, an arc emanating from a node X going into a decision node D is a temporal precedence arc and states that X is observed or decided upon immediately before D is decided upon. *No-forgetting* is assumed so that variables observed or decided upon immediately before previous decisions are remembered at subsequent decisions. A chance variable, which is not a parent of any decision in the ID, is either never observed or observed after the last decision. The temporal precedence arcs impose a partial temporal ordering $\prec$ on $\mathbf{V}_C \cup \mathbf{V}_D$. Together with the requirement on a directed path through all decisions, this ordering induces a total temporal ordering on $\mathbf{V}_D$. For notational convenience, we assume that the decisions are labelled $D_1, \dots, D_n$, such that $i < j$ implies $D_i \prec D_j$. Furthermore, we use the notation $\mathbf{C}_{i-1}$ to mean the set of chance variables observed immediately before deciding on D_i . By $\mathbf{C}_n$ we refer to the set of chance variables never observed (or observed after the last decision D_n has been decided upon). In summary we have

$$\mathbf{C}_0 \prec D_1 \prec \mathbf{C}_1 \prec \dots \prec D_n \prec \mathbf{C}_n ,$$

and no-forgetting amounts to the assumption that for any D_k all observations of variables in $\cup_{i < k} \mathbf{C}_i$ and decisions taken for $D_1, \dots, D_{k-1}$ are remembered when the DM decides on D_k . We define the *past* of decision D to be $\mathbf{past}(D) = \{V \in \mathbf{V}_C \cup \mathbf{V}_D \mid V \prec D\}$. We encode the quantitative aspects of the modelled decision problem as a set Φ of conditional probability distributions and a set Ψ of local utility functions:

$$\begin{aligned} \Phi &= \{P(C \mid \mathbf{pa}(C)) \mid C \in \mathbf{V}_C\}, \text{ and} \\ \Psi &= \{U(\mathbf{pa}(U)) \mid U \in \mathbf{V}_U\} . \end{aligned}$$

A pair (Φ, Ψ) is called a *realisation* for the ID. Here, and henceforth, we have used $f(V_1, \dots, V_k)$ to denote a function of the type $f : \mathbf{sp}(V_1) \times \dots \times \mathbf{sp}(V_k) \rightarrow \mathbb{R}$. Such a function is called a *potential*; we distinguish between *probability potentials*, denoted by ϕ 's, and *utility potentials*, denoted by ψ 's. Furthermore, the set of variables $\{V_1, \dots, V_k\}$ is referred to as the *domain* of f , denoted

$\text{dom}(f)$.

Given an ID and a decision D a function, $\delta_D : \mathbf{sp}(\mathbf{past}(D)) \rightarrow \mathbf{sp}(D)$, is called a *policy* for D . A collection of policies for each decision in an ID,

$$\Delta = \{\delta_D : \mathbf{sp}(\mathbf{past}(D)) \rightarrow \mathbf{sp}(D) \mid D \in V_D\},$$

is called a *strategy* for the ID. Given a policy δ_D for a decision D we define the *chance variable policy*[2] for D $P_{\delta_D}(D|\mathbf{past}(D))$ as $P_{\delta_D}(d|\mathbf{c}) = 1$ if $\delta_D(\mathbf{c}) = d$ and 0 otherwise. An *optimal strategy* Δ^* for an ID is a strategy that fulfills

$$\Delta^* = \arg \max_{\Delta} \sum_{\mathbf{c} \in \mathbf{sp}(V_C \cup V_D)} \left(\prod_{D \in V_D} P_{\delta_D}(\mathbf{c}) \prod_{\phi \in \Phi} \phi(\mathbf{c}) \sum_{\psi \in \Psi} \psi(\mathbf{c}) \right). \quad (1)$$

The individual policies in an optimal strategy are referred to as *optimal policies*. To denote that a policy for a decision D is a part of an optimal strategy, we write it as δ_D^* . The quantity that is maximised in (1) is the *expected utility* of the decision problem given the strategy Δ , and it is denoted $\text{eu}(\Delta)$. Not all variables in the past of a decision D are necessarily relevant for D . Therefore, we call a variable V *required* for D , if there exists a realisation and a configuration $\mathbf{c}$ over the variables in $\mathbf{past}(D) \setminus \{V\}$, such that $\delta_D^*(\mathbf{c}, v_i) \neq \delta_D^*(\mathbf{c}, v_j)$ for two states v_i and v_j in $\mathbf{sp}(V)$. The set of required variables for a decision D we denote by $\mathbf{req}(D)$.³ We may then redefine a policy to be a function $\delta_D : \mathbf{sp}(\mathbf{req}(D)) \rightarrow \mathbf{sp}(D)$.

When optimal policies are to be identified, it is usually easier to work with a recursive expression for the maximum expected utility instead of (1). An example is the *variable elimination* algorithm [5]: Define

$$\Phi_{(n)} = \{\phi \in \Phi \mid \text{dom}(\phi) \cap (C_i \cup \{D_i\}) \neq \emptyset\},$$

and similarly for $\Psi_{(n)}$. Then

$$\delta_{D_n}^*(\mathbf{c}) = \arg \max_{d \in \mathbf{sp}(D_n)} \sum_{\mathbf{e} \in \mathbf{sp}(C_n)} \prod_{\phi \in \Phi_{(n)}} \phi(\mathbf{c}, d, \mathbf{e}) \sum_{\psi \in \Psi_{(n)}} \psi(\mathbf{c}, d, \mathbf{e}), \quad (2)$$

and we set

$$\phi_n(\mathbf{c}) = \sum_{\mathbf{e} \in \mathbf{sp}(C_n)} \prod_{\phi \in \Phi_{(n)}} \phi(\mathbf{c}, \delta_{D_n}^*(\mathbf{c}), \mathbf{e}), \quad (3)$$

and

$$\psi_n(\mathbf{c}) = \left(\sum_{\mathbf{e} \in \mathbf{sp}(C_n)} \prod_{\phi \in \Phi_{(n)}} \phi(\mathbf{c}, \delta_{D_n}^*(\mathbf{c}), \mathbf{e}) \sum_{\psi \in \Psi_{(n)}} \psi(\mathbf{c}, \delta_{D_n}^*(\mathbf{c}), \mathbf{e}) \right) / \phi_n(\mathbf{c}), \quad (4)$$

³ [9] provides an operational method for determining $\mathbf{req}(D)$ for any decision D in an ID.

for all configurations $\mathbf{c}$ over $\mathbf{req}(D_n)$. For all $i < n$ we recursively define

$$\Phi_{(i)} = \{\phi \in \Phi \cup \{\phi_{i+1}, \dots, \phi_n\} \mid \text{dom}(\phi) \cap (C_i \cup \{D_i\}) \neq \emptyset\} \setminus \cup_{j>i} \Phi_{(j)} \quad (5)$$

and similarly for $\Psi_{(i)}$ with $\psi_{i+1}, \dots, \psi_n$. We then get

$$\delta_{D_i}^*(\mathbf{c}) = \arg \max_{d \in \mathbf{sp}(D_i)} \sum_{\mathbf{e} \in \mathbf{sp}(C_i)} \prod_{\phi \in \Phi_{(i)}} \phi(\mathbf{c}, d, \mathbf{e}) \sum_{\psi \in \Psi_{(i)}} \psi(\mathbf{c}, d, \mathbf{e}), \quad (6)$$

and set

$$\phi_i(\mathbf{c}) = \sum_{\mathbf{e} \in \mathbf{sp}(C_i)} \prod_{\phi \in \Phi_{(i)}} \phi(\mathbf{c}, \delta_{D_i}^*(\mathbf{c}), \mathbf{e}), \quad (7)$$

and

$$\psi_i(\mathbf{c}) = \left(\sum_{\mathbf{e} \in \mathbf{sp}(C_i)} \prod_{\phi \in \Phi_{(i)}} \phi(\mathbf{c}, d, \mathbf{e}) \left(\sum_{\psi \in \Psi_{(i)}} \psi(\mathbf{c}, d, \mathbf{e}) \right) \right) / \phi_i(\mathbf{c}), \quad (8)$$

for all configurations $\mathbf{c}$ over $\mathbf{req}(D_i)$. We may then write the maximum expected utility of the ID as

$$\text{eu}(\Delta^*) = \sum_{\mathbf{e} \in \mathbf{sp}(C_0)} \prod_{\phi \in \Phi_{(0)}} \phi(\mathbf{e}) \sum_{\psi \in \Psi_{(0)}} \psi(\mathbf{e}).$$

Given an ID and a realisation, an optimal strategy may be found through the use of any one of a number of algorithms including [5], [7], [10], and [11], which all utilise the distributive and associative law on the expressions in (2) to (8).

3 Multi-currency Influence Diagrams

From the summations of utility potentials in (2) to (8), it is clear that these must be of the same type, i.e. defined over the same currency; in the ID framework this is tackled by transforming the different currencies into one currency during construction of the ID. We now introduce a framework capable of handling decision problems involving utilities of several currencies. We call models in this framework Multi-currency Influence Diagrams (MCIDs). Basically, an MCID is just an ID where each of the utility nodes is annotated with the currency of the corresponding local utility function. Formally, the syntax and semantics of IDs described in Sec. 2 carry over to MCIDs, except for the requirement that the local utility functions must be an additive decomposition of the global utility function. By assuming some arbitrary, but fixed, order of the currencies in the MCID $s_1, \dots, s_m$ we may refer to the currency of a local utility function by a natural number $i \in \{1, \dots, m\}$.

In what follows we refer to the number of different currencies of an MCID

as the *dimension* of the MCID, and throughout we assume this to be m . A realisation of an MCID is therefore a tuple $(\Phi, \Psi_1, \dots, \Psi_m)$, where Ψ_i is the set of local utility functions of currency i . We require that, for each Ψ_i , its elements form an additive decomposition of a utility function, which encodes the same preference ordering as if the DM had disregarded all consequences measured in currencies different from i , and that there is some linear combination of these utility functions describing the decision makers preferences in full. Note that it follows that an MCID of dimension 1 is an ID, and hence from this point on we assume m to be larger than 1.

A strategy for an MCID is the same as for an ID: A prescription of choices given the required variables in the past. However, if we want to compute an optimal strategy for an MCID, we need a method for comparing amounts of one currency with amounts of another. This can be seen from the following example.

Example 2 (A Simple Example of an MCID). Consider a simple two-dimensional MCID, over the currencies A and B , with only a single binary decision D and two local utilities U_1 and U_2 defined as $U_1(D = d_1) = 1A$, $U_1(D = d_2) = 5A$, $U_2(D = d_1) = 2B$, and $U_2(D = d_2) = 1B$.

Both choices of D can be optimal choices depending on how much the DM values amounts of currency A relative to amounts of currency B . If $D = d_1$ should be an optimal strategy then $\text{eu}(D = d_1) \geq \text{eu}(D = d_2)$, which is equivalent to

$$1A + 2B \geq 5A + 1B \Leftrightarrow -4A + B \geq 0,$$

If we regard the currency name A as a real variable, representing the DM's degree of appreciation of amounts of A , and similarly for currency name B , then $-4 \cdot A + B$ corresponds to an amount of appreciation equivalent to -4 A 's and one B . The set of all values for A and B , where $-4 \cdot A + B \geq 0$, then corresponds to the possible attitudes of the DM rendering d_1 the optimal choice of the decision problem modelled by the MCID. The state space of $A \times B$, viz. $\mathbb{R}^2$, is thus partitioned into two regions, each corresponding to an optimal strategy.

To sum up, we see that the payoff of following a specific strategy is an element of $\mathbb{R}^m$ (e.g. $(5, 1)$ in Ex. 2) rather than a scalar value as is the case with strategies for IDs. The means we use for comparing amounts of different currencies are called *currency mappings*:

Definition 1. Let $\mathcal{I}$ be an MCID of dimension m and $\alpha = (\alpha_1, \dots, \alpha_m)$ a point in $\mathbb{R}^m$, then α is a currency mapping (CM) for $\mathcal{I}$.

The semantics of a CM α , *reflecting a DM's preferences*, is that, for any two amounts x_i and x_j of currencies i and j , respectively, we have that $\alpha_i x_i$ equals $\alpha_j x_j$ iff the DM values x_i of currency i as much as x_j of currency j . We also say that the DM *adheres* to α . This way α_i becomes a measure of appreciation for the DM of one unit of currency i . In a multi-stakeholder

setting, α_i becomes a weight of importance attributed to satisfying the i 'th stakeholder compared to the other stakeholders. In any case, the objective of the DM is to maximize the *expected global utility* of a strategy Δ given by

$$\alpha \cdot \text{eu}(\Delta) = \sum_{i=1}^m \alpha_i \sum_{\mathbf{c} \in \mathbf{sp}(V_C \cup V_D)} \left(\prod_{D \in V_D} P_{\delta_D}(\mathbf{c}) \prod_{\phi \in \Phi} \phi(\mathbf{c}) \sum_{\psi_i \in \Psi_i} \psi_i(\mathbf{c}) \right). \quad (9)$$

It follows that we assume the preferential relationship between currencies is a linear one. Thus, in Ex. 2, a DM adhering to the CM $(3, 2)$, would have an expected global utility of $3 \cdot 1 + 2 \cdot 2 = 7$ for choosing d_1 and $3 \cdot 5 + 2 \cdot 1 = 17$ for choosing d_2 . In order to maximize the expected global utility he should therefore choose d_2 .

We use bold face Latin letters ($\mathbf{f}$, $\mathbf{q}$, etc.) to denote arbitrary points in $\mathbb{R}^m$, and Greek letters (α and β) to denote points when we want to emphasise that they are CMs. In general, we use q_i to refer to the i 'th coordinate of a point $\mathbf{q}$. We shall not distinguish between a point $\mathbf{q}$ and the corresponding vector going from the origin to $\mathbf{q}$. In what follows we furthermore use $\mathbf{f} \cdot \mathbf{q}$ to denote the scalar product $\sum_i f_i q_i$.

In this paper we assume, without loss of generality, that each element of a currency mapping is positive, i.e. a CM is an element of $\mathbb{R}_+^m$ rather than $\mathbb{R}^m$, where $\mathbb{R}_+$ denotes the set of strictly positive reals. This assumption implies that everybody should be able to agree on whether amounts of each currency, i , is beneficial to be had or not, and that no-one would contest the relevance of amounts of any currency. If a currency i is disadvantageous to be had (meaning that α_i should be negative) we expect the modelled decision problem to have negative utilities specified for positive amounts of i , as is usually done when costs are specified in IDs.

If for a strategy, Δ , we have that $\alpha \cdot \text{eu}(\Delta)$ is greater than or equal to $\alpha \cdot \text{eu}(\Delta')$ for all other strategies Δ' , it means that a DM adhering to α appreciates the expected utility of following Δ at least as much as that of following any other strategy, and we consequently say that Δ is optimal given the CM α . As can be seen from (9), and the distributive law of the scalar product, the optimal strategy for α is equivalent to the one obtained by solving an ID resulting from multiplying the individual utilities in the MCID with α beforehand. We denote an optimal strategy for an MCID given a CM α as Δ_α^* . As several strategies might give rise to the same expected utility, the set of all optimal strategies corresponding to α is denoted as $\overline{\Delta_\alpha^*}$.

4 Support Analysis of MCIDs

Clearly, if we are given a CM α in addition to an MCID we can solve it by means of simply converting the MCID into an ID, through multiplying each local utility function of currency i by α_i , and then solving the resulting ID. This simple solution method allows for optimal strategies to be computed for any α , and the distinction between MCIDs and CMs emphasises

the assumptions leading to the results. However, we would not be closer to answering the second question in the motivating example, viz. what other currency mappings would lead to the same optimal strategy? In order to do this, we must render the effect of α on the optimal strategy transparent. We would then be able to reason about the universality of the applicability of that optimal strategy. We obtain this transparency by postponing the conversion of utilities until it is needed, so that we can analyse the requirements these conversions bestow upon those CMs giving rise to the same optimal strategy.

4.1 Preliminaries

Given an MCID and a strategy Δ , we call the set $\text{su}(\Delta) = \{\beta \in \mathbb{R}_+^m \mid \Delta \in \overline{\Delta_\beta^*}\}$ the *support* of Δ . Intuitively, $\text{su}(\Delta)$ is the set of CMs for which Δ is an optimal strategy. We refer to the process of calculating the support of an optimal strategy as performing support analysis of the MCID. In Ex. 2 we actually found the support of both strategies $D = d_1$ and $D = d_2$ (the two partitions of $\mathbb{R}_+^2$ defined by $-4A + B = 0$). Later it will become apparent that any such support can be described as an intersection of partitions of $\mathbb{R}_+^m$ – each partition described by a linear inequality.

As mentioned in the beginning of this section, we postpone the conversion of utility potentials until needed. Hence, we introduce a new type of potential, which can represent utility functions of several currencies:⁴ A *multi-currency utility potential* (MCUP) of dimension m over the variables in a set S is a function $\theta : \text{sp}(S) \rightarrow \mathbb{R}^m$ attributing to each configuration c over the variables in S a measure of utility $(\theta(c))_i$ of each currency i . In Ex. 2 we could have exchanged the two utility functions U_1 and U_2 with the two MCUPs θ_1 and θ_2 , where $\theta_1(d_1) = (1, 0)$, $\theta_1(d_2) = (5, 0)$, $\theta_2(d_1) = (0, 2)$, $\theta_2(d_2) = (0, 1)$. For a DM whose preferences are reflected by the CM β , $\beta \cdot \theta(S)$ is a utility potential that for any configuration, c , over variables in S yields the value $\beta_1(\theta(c))_1 + \dots + \beta_m(\theta(c))_m$, which to the DM is equivalent to the amounts $(\theta(c))_1, \dots, (\theta(c))_m$ of currencies $1, \dots, m$, respectively. When adding MCUPs or multiplying probability potentials onto them, we simply treat each dimension of the MCUP as a regular utility potential, and perform the operation on each dimension separately. For instance, in Ex. 2 we have that $\theta_1 + \theta_2$ is the MCUP θ_+ where $\theta_+(d_1) = (1, 2)$ and $\theta_+(d_2) = (5, 1)$, which to a DM adhering to $\beta = (3, 2)$ would be equivalent to a utility potential ψ , where $\psi(d_1) = 7$ and $\psi(d_2) = 17$.

4.2 Support Analysis

We are now ready to give a procedure for performing support analysis of an MCID. We describe the method first, and give an example of its application

⁴ Such potentials are not part of the MCID framework as such, but rather data structures used by the proposed method for doing support analysis.

afterwards. We assume the existence of an optimal strategy Δ_α^* determined by some initial currency mapping α as well as a realisation $(\Phi, \Psi_1, \dots, \Psi_m)$, and we look for the support of Δ_α^* for this realisation. The method is inspired by that of Lazy evaluation presented in [7] and basically traces the steps of this method while recording the requirements on α for Δ_α^* to be an optimal strategy. The method consists of an initialisation phase and an identification phase. The initialisation phase consists of two steps: First, two empty sets Ξ and Θ are generated, where Ξ will hold inequalities defining $\text{su}(\Delta_\alpha^*)$, and Θ is a container for MCUPs used in the identification phase. Second, for each currency i each potential ψ in Ψ_i is converted to a MCUP θ and put into Θ , such that $\theta_k = \psi$ if $k = i$ and 0 otherwise, where 0 denotes the function yielding the zero value for all input.

The identification phase follows the Lazy evaluation method, except for the steps normally carried out when a variable is eliminated (see Algorithm 1). The major difference between these steps and the corresponding steps in Lazy evaluation is that we do not perform a maximisation to uncover an optimal strategy. Instead we look for a set of linear inequalities (Step 4) that need to be fulfilled if Δ_α^* is to be optimal. We refer to the inequalities in Ξ as *constraints*, since they constrain the support set.

Algorithm 1: *The elimination steps of the identification phase. The strategy Δ_α^* is assumed to be given apriori.*

1. Let V be the variable to be eliminated, and let Φ_V denote the set of probability potentials in Φ with V in their domain, and Θ_V denote the set of MCUPs in Θ with V in their domain.

2. Let

$$\phi_V = \prod_{\phi \in \Phi_V} \phi, \quad \text{and} \quad \theta_V = \sum_{\theta \in \Theta_V} \theta.$$

3. If V is a chance variable, then set

$$\Phi \leftarrow (\Phi \setminus \Phi_V) \cup \left\{ \sum_V \phi_V \right\},$$

and

$$\Theta \leftarrow (\Theta \setminus \Theta_V) \cup \left\{ \frac{\sum_V (\phi_V \theta_V)}{\sum_V \phi_V} \right\}.$$

4. If V is a decision variable, then set

$$\Phi \leftarrow (\Phi \setminus \Phi_V) \cup \{ \phi(V = v) \mid \phi \in \Phi_V \},$$

where v is some arbitrary state of V ⁵, and

$$\Theta \leftarrow (\Theta \setminus \Theta_V) \cup \{ \theta_V(\delta_V^*) \},$$

⁵ Note that any potential in Φ_V must be constant over V

where δ_V^* is the appropriate element of Δ_α^* . For each configuration $\mathbf{c}$ over the variables in $\mathbf{req}(V)$ and state $v \neq \delta_V^*(\mathbf{c})$ of V , set

$$\Xi \leftarrow \Xi \cup \{f_{\mathbf{c},v} \cdot \gamma \geq 0\}, \quad (10)$$

where $f_{\mathbf{c},v}$ denotes $\theta(\mathbf{c}, \delta_V^*(\mathbf{c})) - \theta(\mathbf{c}, v)$.

Before proceeding, we illustrate the workings of the algorithm by an example:

Example 3. We specify the ID from Ex. 1 as an MCID. We assume that the state space of the variables are as follows:

$$\begin{aligned} \mathbf{sp}(\text{Health}_1(H_1)) &= \{bad(b_1), healthy(h_1)\}, \\ \mathbf{sp}(\text{Symptoms}(S)) &= \{symptoms(sy), none(no)\}, \\ \mathbf{sp}(\text{Test}(Te)) &= \{test(te), no_test(nte)\}, \\ \mathbf{sp}(\text{Results}(R)) &= \{positive(po), negative(ne), no_result(nr)\}, \\ \mathbf{sp}(\text{Treat}(Tr)) &= \{treat(tr), no_treatment(ntr)\}, \text{ and} \\ \mathbf{sp}(\text{Health}_2(H_2)) &= \{bad(b_2), healthy(h_2)\}. \end{aligned}$$

The three currencies we work with are, and are ordered as *comfort*(c), *dollars*($\$$), and *health*(h). The realization that we work with is defined by the following parameters:

$$\begin{aligned} P(b_1) &= 0.5, \quad P(sy|b_1) = 0.95, \quad P(sy|h_1) = 0.1, \quad P(nr|nte, \cdot) = 1, \\ P(nr|te, \cdot) &= 0, \quad P(po|te, b_1) = 0.99, \quad P(po|te, h_1) = 0.01, \\ P(b_2|b_1, tr) &= 0.001, \quad P(b_2|b_1, ntr) = 0.999, \quad P(b_2|h_1, tr) = 0, \\ P(b_2|h_1, ntr) &= 0.0001, \quad U_1(te, nte) = (-1c, 0), \quad U_2(te, nte) = (-\$1000, 0), \\ U_3(tr, ntr) &= (-\$10000, 0), \text{ and } U_4(b_2, h_2) = (-1h, 0). \end{aligned}$$

We assume that the doctor is adhering to a CM $\alpha = (10, 1, 100000)$ meaning that he regards discomfort on behalf of the patient as bad as the loss of \$10 and the death of the patient as bad as a loss of \$100000. The CM corresponds to a strategy prescribing *treatment* either if a *test* was conducted, and a *positive* test result was gotten, or *no test* was conducted and there was *symptoms*, and *no treatment* otherwise.

In the initialization part of the algorithm, we construct an empty set Ξ and convert the four utility potentials into a set of four MCUPs

$$\begin{aligned} \Theta &= \{\theta_1(te, nte) = ([-1, 0, 0], [0, 0, 0]), \theta_2(te, nte) = ([0, -1000, 0], [0, 0, 0]), \\ &\quad \theta_3(tr, ntr) = ([0, -10000, 0], [0, 0, 0]), \theta_4(b_2, h_2) = ([0, 0, -1], [0, 0, 0])\}. \end{aligned}$$

Next we perform variable elimination in an order that respects the $\prec$ -ordering of the MCID: First *Health*₂, then *Health*₁, *Treat*, *Results*, *Test*, and lastly *Symptoms*. Marginalizing out *Health*₂, according to Steps 2 and 3 results in

dropping $P(H_2|H_1, Tr)$ from the set of probability potentials and replacing θ_4 with a new MCUP:

$$\theta_{H_1, Tr} = \frac{\sum_{H_2} P(H_2|H_1, Tr) \cdot \theta_4}{\sum_{H_2} P(H_2|H_1, Tr)},$$

in which $\theta_{H_1, Tr}(b_1, tr) = [0, 0, -10^{-3}]$, $\theta_{H_1, Tr}(b_1, ntr) = [0, 0, -0.999]$, $\theta_{H_1, Tr}(h_1, tr) = [0, 0, 0]$, and $\theta_{H_1, Tr}(h_1, ntr) = [0, 0, -10^{-4}]$. Further marginalizing out $Health_1$ (again according to Steps 2 and 3) replaces all probability potentials with

$$\phi_{R, S, Te} = \sum_{H_1} P(H_1) \cdot P(S|H_1) \cdot P(R|H_1, Te)$$

and $\theta_{H_1, Tr}$ with

$$\theta_{R, S, Te, Tr} = \frac{\sum_{H_1} P(H_1) \cdot P(S|H_1) \cdot P(R|H_1, Te) \cdot \theta_{H_1, Tr}}{\phi_{S, R, Te}}.$$

This last potential is shown in Table 1.

When marginalizing *Treat*, we construct the MCUP $\theta_{R, S, Te, Tr}^+ = \theta_3 + \theta_{R, S, Te, Tr}$ (shown in Table 2), according to Step 2, and by marginalizing *Treat* out of this, according to Step 4 and Δ_{α}^* , we end up with a new MCUP $\theta_{R, S, Te}$ (not shown), but also add 12 constraints to Ξ – one for each configuration over *Results*, *Symptoms*, and *Test*. For instance, for the configuration (po, sy, te) , we construct the constraint $(\gamma = (\gamma_c, \gamma_s, \gamma_h))$

$$\begin{aligned} ((0, -10000, -9.99 \cdot 10^{-4}) - (0, 0, -9.98 \cdot 10^{-1}))\gamma &\geq 0 \\ -10000\gamma_s + 9.97 \cdot 10^{-1}\gamma_h &\geq 0, \end{aligned}$$

which is satisfied by α , as can easily be verified. After the other 11 constraints have been calculated the algorithm continues with elimination of the remaining variables.

Table 1. $\theta_{R, S, Te, Tr}$

		te		nte	
		tr	ntr	tr	ntr
po	sy	$[0, 0, -9.99 \cdot 10^{-4}]$	$[0, 0, -9.98 \cdot 10^{-1}]$	$[0, 0, 0]$	$[0, 0, 0]$
	no	$[0, 0, -8.46 \cdot 10^{-4}]$	$[0, 0, -8.45 \cdot 10^{-1}]$	$[0, 0, 0]$	$[0, 0, 0]$
ne	sy	$[0, 0, -8.76 \cdot 10^{-4}]$	$[0, 0, -8.76 \cdot 10^{-2}]$	$[0, 0, 0]$	$[0, 0, 0]$
	no	$[0, 0, -5.61 \cdot 10^{-4}]$	$[0, 0, -6.60 \cdot 10^{-4}]$	$[0, 0, 0]$	$[0, 0, 0]$
nr	sy	$[0, 0, 0]$	$[0, 0, 0]$	$[0, 0, -9.05 \cdot 10^{-4}]$	$[0, 0, -9.04 \cdot 10^{-1}]$
	no	$[0, 0, 0]$	$[0, 0, 0]$	$[0, 0, -5.26 \cdot 10^{-5}]$	$[0, 0, -5.27 \cdot 10^{-2}]$

Table 2. $\theta_{R,S,Te,Tr}^+$

		te		nte	
		tr	ntr	tr	ntr
po	sy	$[0, -10^4, -9.99 \cdot 10^{-4}]$	$[0, 0, -9.98 \cdot 10^{-1}]$	$[0, -10^4, 0]$	$[0, 0, 0]$
	no	$[0, -10^4, -8.46 \cdot 10^{-4}]$	$[0, 0, -8.45 \cdot 10^{-1}]$	$[0, -10^4, 0]$	$[0, 0, 0]$
ne	sy	$[0, -10^4, -8.76 \cdot 10^{-4}]$	$[0, 0, -8.76 \cdot 10^{-2}]$	$[0, -10^4, 0]$	$[0, 0, 0]$
	no	$[0, -10^4, -5.61 \cdot 10^{-4}]$	$[0, 0, -6.60 \cdot 10^{-4}]$	$[0, -10^4, 0]$	$[0, 0, 0]$
nr	sy	$[0, -10^4, 0]$	$[0, 0, 0]$	$[0, -10^4, -9.05 \cdot 10^{-4}]$	$[0, 0, -9.04 \cdot 10^{-1}]$
	no	$[0, -10^4, 0]$	$[0, 0, 0]$	$[0, -10^4, -5.26 \cdot 10^{-5}]$	$[0, 0, -5.27 \cdot 10^{-2}]$

As presented here, the algorithm presupposes that Δ_α^* has been computed beforehand. Alternatively, the Lazy evaluation algorithm itself can easily be interleaved by inserting the following step prior to Step 4 in Algorithm 1:

* For each configuration $\mathbf{c}$ over the variables in $\mathbf{req}(V)$ set

$$\delta_V^*(\mathbf{c}) = \arg \max_{v \in \mathbf{sp}(V)} \alpha \cdot \theta_V(\mathbf{c}, v) .$$

Although the method, as presented here, follows the structure of the Lazy evaluation method, it can easily be adapted to follow the structure of any other solution method that is based on the expressions in (2) to (8). We conjecture that any such adaptation would identify the support set as long as the constraints are identified and stored. With or without this modification, though, we have the following important result:

Theorem 1. *Let β be in $\mathbb{R}_+^m$, Δ a strategy, and Ξ the result of running the method described above on Δ . Then β is an element of $\mathbf{su}(\Delta)$ iff β satisfies all inequalities in Ξ .*

Proof. The “if” part of the theorem is obvious. We therefore only show the “only if” part, viz. that if β fails to satisfy at least one constraint in Ξ , then Δ cannot be an optimal strategy for a DM adhering to β , and hence that β is not an element of $\mathbf{su}(\Delta)$.

Without loss of generality, assume that $f \in \Xi$ is the first constraint not satisfied by β that is identified by the algorithm. This happens during elimination of some decision D_i in Step 4 of the algorithm, with some MCUP θ_{D_i} having been calculated during Step 2. For some configuration $\mathbf{c}$ over $\text{dom}(\theta_{D_i}) \setminus \{D_i\}$ and state $v \neq \delta_{D_i}$ we must have that f is

$$(\theta_{D_i}(\mathbf{c}, \delta_{D_i}(\mathbf{c})) - \theta_{D_i}(\mathbf{c}, v)) \cdot \gamma \geq 0 .$$

Since β fails to satisfy this constraint, it follows that

$$(\theta_{D_i}(\mathbf{c}, \delta_{D_i}(\mathbf{c})) - \theta_{D_i}(\mathbf{c}, v)) \cdot \beta \not\geq 0 ,$$

which is equivalent to

$$\theta_{D_i}(\mathbf{c}, \delta_{D_i}(\mathbf{c})) \cdot \boldsymbol{\beta} < \theta_{D_i}(\mathbf{c}, v) \cdot \boldsymbol{\beta}.$$

We construct the strategy $\boldsymbol{\Delta}'$ obtained from $\boldsymbol{\Delta}$ by letting $\delta'_{D_i}(\mathbf{c}) = v$ and leaving all other policies intact. We then have that $\text{eu}(\boldsymbol{\Delta}') \cdot \boldsymbol{\beta}$ is strictly greater than $\text{eu}(\boldsymbol{\Delta}) \cdot \boldsymbol{\beta}$, and hence that $\boldsymbol{\Delta}$ cannot be an optimal strategy for a DM adhering to $\boldsymbol{\beta}$.

Glancing over the constraints identified in Ex. 3 it is clear that each constraint $f : a\gamma_1 + b\gamma_2 + c\gamma_3 \geq 0$ defines a hyperplane in $\mathbb{R}^3$ given by the equation $a\gamma_1 + b\gamma_2 + c\gamma_3 = 0$. As the zero vector $\mathbf{0} = (0, 0, 0)$ is a satisfying solution to each equation, it follows that all these hyperplanes must pass through the origin of $\mathbb{R}^3$, and hence that the points satisfying all constraints must lie in a bottomless pyramid extending from the origin, as illustrated in Fig. 2(a). A corresponding visualization for a two-dimensional MCID is shown in Fig. 2(b). That the support of a strategy extends indefinitely from the origin as a pyramid, also makes sense from a purely semantical point of view: If a CM is scaled by multiplying each entry by some positive constant, the relative difference in appreciation between amounts of the individual currencies, for a DM adhering to this CM, stays the same. Hence, if a CM renders some strategy optimal, that strategy should also be optimal for any positively scaled version of this CM. From a more formal point of view, we also have that, for any CM $\boldsymbol{\alpha}$ and two strategies $\boldsymbol{\Delta}_i$ and $\boldsymbol{\Delta}_j$, if $\boldsymbol{\alpha} \cdot \text{eu}(\boldsymbol{\Delta}_i) \geq \boldsymbol{\alpha} \cdot \text{eu}(\boldsymbol{\Delta}_j)$ then it must necessarily be the case that $c\boldsymbol{\alpha} \cdot \text{eu}(\boldsymbol{\Delta}_i) \geq c\boldsymbol{\alpha} \cdot \text{eu}(\boldsymbol{\Delta}_j)$ for any $c > 0$, and hence that the pyramidal forms of support areas are what we should expect.

Fig. 2. Supports for a three-dimensional (a) and a two-dimensional MCID (b)

5 Finding a Minimal Support Set

The procedure described above finds the support of a strategy for a given CM and stores it as a set of constraints Ξ . The cardinality of Ξ is given by

$$|\Xi| = \sum_{D \in \mathcal{V}_D} |\mathbf{sp}(\mathbf{req}(D))|(|\mathbf{sp}(D)| - 1) ,$$

and storing it requires an amount of memory proportional to $m \cdot |\Xi|$. This size can be problematic for larger decision problems – both in terms of memory requirements and in terms of representing the resulting Ξ to a human DM in a comprehensible manner.

We have devised a method for keeping the size of Ξ minimal during the analysis. By minimal we mean that Ξ does not include a constraint f , such that the set $\Xi \setminus \{f\}$ defines the same volume as Ξ . Conversely, if such a constraint f is in Ξ , we call it *superfluous*. Traditionally, such a task would be carried out by repeated applications of linear programming, but this can be done more efficiently when the number of dimensions m is significantly smaller than the number of decisions in the MCID. This seems to be the case in most examples of multi-currency decision problems in the literature, see e.g. [1] and [12].

The approach is purely geometric and rests on the already mentioned fact that all constraints in Ξ define a pyramid extending from the origin of $\mathbb{R}^m$ (see Figs. 2(a) and 3(a)), and a black box view of the support analysis as a simple constraint generating process. For sake of clarity, we first describe the method informally in the intuitively understood three dimensional setting, and then give a formal presentation of the more general m -dimensional setting.

As the support of a strategy is a pyramid extending from the origin, we can represent the volume defined by Ξ as a set of lines extending from the origin, each corresponding to the intersection of two hyperplanes defined by constraints in Ξ (see Fig. 3(a), which illustrates the support defined by the initial set of constraints $\gamma_1 \geq 0$, $\gamma_2 \geq 0$, and $\gamma_3 \geq 0$, as well as an additional constraint f). We refer to such lines as *edges* of the pyramid, and we denote the set of them as $\mathbf{E}_\Xi$. When a new constraint g is to be added to Ξ (see Fig. 3(b)) it may be superfluous. With the alternative representation this question can be restated as whether the points on any of the edges in $\mathbf{E}_\Xi$ fail to satisfy g . If this is the case, g is not superfluous. In Fig. 3(b) we have that there was originally a pyramid defined by the edges e_1 , e_2 , e_3 , and e_4 , but as all points (but the origin) on e_1 and e_2 fail to satisfy the new constraint g , it is not superfluous. Therefore it is added to Ξ , which now defines a new pyramid whose edges are e_3 , e_4 , e_5 , and e_6 . Thus $\mathbf{E}_\Xi$ is updated by dropping e_1 and e_2 and adding e_5 and e_6 . This means that the original constraint $\gamma_1 \geq 0$ no longer participates in defining any of the edges in $\mathbf{E}_\Xi$, and it is therefore superfluous now. Consequently, $\gamma_1 \geq 0$ must be dropped from Ξ to keep it minimal (see Fig. 3(c)). That is, in this alternative representation, we can

also detect when old constraint are rendered superfluous by new constraints. There is one main hurdle to be overcome by an implementation: Whenever

Fig. 3. Keeping Ξ minimal when a new constraint g is added. Edges are shown as dashed lines

a new constraint h is identified, the edges in $\mathbf{E}_\Xi$ need to be checked for points not satisfying h . By storing the pyramid as a list of edges sorted according to their angular distance to α (see Fig. 4(a)), and calculating the minimal angular distance from α to the hyperplane defined by h , we can quickly determine those edges that can contain points not satisfying h , i.e. those with an angular distance greater than the distance to the hyperplane defined by h . If no such edges exist (as in Fig. 4(b)), h is clearly superfluous. If such edges do exist, they have to be checked against h one at the time, and only if all of them satisfy h , it must be superfluous. However, as the set of constraints grow, the girth of the pyramid becomes smaller, and we would therefore expect that more new constraints fail the initial check of being closer to α than any of the edges in Ξ , and that this further check by enumeration can be avoided.

To describe the approach precisely, we formalize and generalize the discussion above: First, if $f(\gamma) = \mathbf{f} \cdot \gamma \geq 0$ is a constraint in Ξ and β is some point in $\mathbb{R}^m$, then we use $f(\beta)$ to state that $\mathbf{f} \cdot \beta \geq 0$ and $\neg f(\beta)$ to state that $\mathbf{f} \cdot \beta < 0$. Furthermore, we denote the hyperplane $\{\beta \in \mathbb{R}^m \mid \mathbf{f} \cdot \beta = 0\}$ as $\mathbf{H}(f)$, and talk of Ξ as consisting of hyperplanes when it introduces no ambiguity. A constraint f in Ξ is then defined to be *superfluous* in Ξ if there exists no point β in $\mathbb{R}^m$, such that $\neg f(\beta)$ and $g(\beta)$ for all $g \neq f$ in Ξ . If no constraint in a set of constraints Ξ is superfluous, we say that Ξ is *minimal*. The *angular distance* from β to α is given by

$$\angle \beta \alpha = \cos^{-1} \frac{\beta \cdot \alpha}{\|\beta\| \cdot \|\alpha\|}.$$

Here and henceforth $\|q\|$ denotes the Euclidean length of the vector q . Moreover, we say that the set of constraints Ξ constitutes a *pyramid*, if for

Fig. 4. Rejecting a constraint, h , because of its angular distance to α

all points β in $\mathbb{R}^m$, where the angular distance to α is greater than 90° , there exists at least one constraint f in Ξ , such that $\neg f(\beta)$. The previously introduced assumption of only considering CMs in $\mathbb{R}_+^m$, we restate as an inclusion of the m constraints $\gamma_i \geq 0$ in Ξ from the outset of the support analysis. This ensures that Ξ is a pyramid at the beginning of the support analysis. The triangle inequality ensures that the points in a pyramid having the largest angular distance to α must lie at the intersection of a number of hyperplanes in Ξ . We need to consider intersections of exactly $m - 1$ hyperplanes. The reason is that an intersection of k non-parallel hyperplanes in $\mathbb{R}^m$ describe an $m - k$ dimensional subspace of $\mathbb{R}^m$, and the points in this subspace will have varying angular distances to α unless the subspace has dimension 1, and this is only the case when k is 1. If part of such an intersection lies within Ξ , we refer to that part as an *edge*. An edge determined by constraints $f_1, \dots, f_{m-1}$ we represent by a pair $e = (\mathbf{I}, \mathbf{p})$, where $\mathbf{I} = \{f_1, \dots, f_{m-1}\}$ and $\mathbf{p} \in \mathbb{R}_+^m \cap \mathbf{H}(f_1) \cap \dots \cap \mathbf{H}(f_{m-1})$. The set of edges of a pyramid Ξ is denoted $\mathbf{E}_\Xi$, and we assume it to be kept sorted such that $(\mathbf{I}_i, \mathbf{p}_i)$ is stored before $(\mathbf{I}_j, \mathbf{p}_j)$ only if $\angle \mathbf{p}_i \alpha \leq \angle \mathbf{p}_j \alpha$. Finally, for any constraint f , we denote by $\alpha^{\perp f}$ the projection of α onto the subspace $\mathbf{H}(f)$. With these terms specified, we can now present the proposed method in Algorithm 2.

Algorithm 2: Takes as input a minimal pyramid Ξ with a sorted set of edges $\mathbf{E}_\Xi$ and a constraint f . Outputs a new minimal pyramid Ξ' , describing the same volume as $\Xi \cup \{f\}$, along with a sorted set of edges $\mathbf{E}_{\Xi'}$.

1. Partition $\mathbf{E}_\Xi$ into two sets $\mathbf{E}_+$ and $\mathbf{E}_-$, such that each edge in $\mathbf{E}_+$ has angular distance to α less than or equal to $\angle \alpha \alpha^{\perp f}$.
2. Move each edge $(\mathbf{I}, \mathbf{p})$ in $\mathbf{E}_-$, where $f(\mathbf{p})$, from $\mathbf{E}_-$ to $\mathbf{E}_+$.
3. If $\mathbf{E}_- = \emptyset$ then stop and return $\Xi' = \Xi$ and $\mathbf{E}_{\Xi'} = \mathbf{E}_\Xi$.

4. Let the new set of constraints be

$$\Xi' = \{f\} \cup \bigcup_{(I, \mathbf{p}) \in E_+} I,$$

and the set of constraints possibly defining new edges be

$$\Xi_N = \left(\Xi' \cap \bigcup_{(I, \mathbf{p}) \in E_-} I \right) \cup \{f\}.$$

Then let

$$E_{\Xi'} = E_+ \cup \{(I, \mathbf{p}_I) \mid I \subseteq \Xi_N, f \in I, |I| = m-1, \text{ and } g(\mathbf{p}_I) \forall g \in \Xi'\},$$

where $\mathbf{p}_I$ is the unique point in $\mathbb{R}_+^m \cap (\cap_{h \in I} H(h))$ for which $\|\mathbf{p}_I\| = 1$.

5. Return Ξ' and $E_{\Xi'}$.

Proposition 1. *Let Ξ constitute a minimal pyramid and f be a constraint not in Ξ . Then the constraints in $\Xi \cup \{f\}$ define the same region as those in Ξ' obtained from using Algorithm 2 to add f to Ξ . Furthermore, Ξ' is a minimal pyramid.*

The complexity of inserting an element in the sorted set E_{Ξ} is $O(\log |E_{\Xi}|)$, and we have a maximum of $\binom{|\Xi|}{m-2}$ new edges to add to E_{Ξ} in Step 4, in effect yielding a worst-case complexity of $O(\binom{|\Xi|}{m-2} \log |E_{\Xi}|)$ for insertion of a constraint. This is an improvement with respect to the complexity of $(|\Xi| + 1)m^{|\Xi|+1}$, offered by using simplex repeatedly [6], especially when the dimension m is low.

6 Presenting the Support to a Human DM

Given that Ξ has been identified as in Fig. 5(a), we provide two compact abstractions, which are useful for presenting the support to a DM.⁶

An immediate approach to representing Ξ in a compact manner is to define an m -dimensional ball centered at α with radius equal to the minimum Euclidean distance from α to any one of the hyperplanes defined by constraints in Ξ (see Fig. 5(b)).

The approach allows for a highly compact representation of Ξ during computation too: Only a single scalar value (the radius of the ball) needs to be stored. Whenever a constraint closer to α is identified, we replace the old radius with the distance from α to this new constraint. Unfortunately this representation can be a rather crude abstraction, as seen in Fig. 5(b).

Another, more accurate, representation technique is to present Ξ by the

⁶ These techniques do not presuppose that Ξ is minimal.

Fig. 5. Three ways of representing the result of support analysis

largest ball that will fit into the support defined by it. That is, abandon α as a fixed center of the representation. As Ξ describes an infinite pyramid, no such largest ball exists, though, so we propose to make a cut through the computed pyramid Ξ , by forcing one of the γ_i s to be 1, and then representing the resulting intersection by the largest ball, which can fit into it (see Fig. 5(c)). The first step of this approach corresponds to identifying a base currency i that all other currencies are compared to.

For the second step to be successfully completed it is necessary that the intersection of the pyramid and the cut is a bounded volume. This is the case if the hyperplane defined by $\gamma_i = 1$ intersects all hyperplanes in Ξ . This is equivalent to that there exists no f in Ξ , such that the hyperplane defined by f is parallel to the hyperplane defined by $\gamma_i = 1$. As all hyperplanes pass through the origin, it is sufficient to choose an i where the constraint $\gamma_i \geq 0$ is not in Ξ . If no such i exists, additional linear constraints on parameters will have to be put into Ξ , e.g. $\gamma_i < k$ for some currency i and positive constant k .

Once a base currency i has been chosen, each constraint $f : f_1\gamma_1 + \dots + f_m\gamma_m \geq 0$ in Ξ is replaced with the constraint $f_1\gamma_1 + \dots + f_{i-1}\gamma_{i-1} + f_i + f_{i+1}\gamma_{i+1} + \dots + f_m\gamma_m \geq 0$, corresponding to f 's effect on points on the hyperplane defined by $\gamma_i = 1$. The resulting set of constraints we denote $\Xi^{\gamma_i=1}$. To find the largest ball enclosed in the volume of $\mathbb{R}^{m-1}$ defined by the constraints in $\Xi^{\gamma_i=1} = \{f^1, \dots, f^k\}$ we solve the linear program

$$f^1(\mathbf{y}), \dots, f^k(\mathbf{y}), z \leq \text{dist}(\mathbf{y}, \mathbf{H}(f^1)), \dots, z \leq \text{dist}(\mathbf{y}, \mathbf{H}(f^k)),$$

where $\mathbf{y}$ is in $\mathbb{R}^{m-1}$ and $\text{dist}(\mathbf{y}, \mathbf{H}(f^i))$ denotes the Euclidean distance from $\mathbf{y}$ to the hyperplane $\mathbf{H}(f^i)$. The function that is to be optimised is z , and upon resolution the ball centered at $\mathbf{y}$ having radius z is the largest ball inscribed in the part of the support corresponding to γ_i being 1.

Acknowledgments

We would like to thank Martin Raussen of the Department of Mathematical Sciences at Aalborg University for useful discussions and two anonymous reviewers for helpful comments and suggestions. Furthermore, we are grateful for access to the Hugin software, which has been used in the implementation of the algorithms described in the paper. The implementation can be downloaded from <http://www.cs.aau.dk/~holbech/mcids.tgz>

References

1. Magnus Boman, Paul Davidsson, and Håkan L. Younes. Artificial decision making under uncertainty in intelligent buildings. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, pages 65–70, 1999.
2. Gregory F. Cooper. A method for using belief networks as influence diagrams. In *Proceedings of the Fourth Conference on Uncertainty in Artificial Intelligence*, pages 55–63, 1988.
3. James C. Felli and Gordon B. Hazen. Do sensitivity analysis really capture problem sensitivity? an empirical analysis based on information value. *Risk, Decision and Policy*, 4(2):79–98, 1999.
4. Ronald A. Howard and James E. Matheson. Influence diagrams. *Readings on the Principles and Applications of Decision Analysis*, pages 720–763, 1984.
5. Frank Jensen, Finn V. Jensen, and Søren L. Dittmer. From influence diagrams to junction trees. In R. Lopez de Mantaras and D. Poole, editors, *Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence*, pages 367–373. Morgan Kaufmann, 1994.
6. V. Klee and G. J. Minty. How good is the simplex algorithm? In *Inequalities III*, pages 159–175. Academic Press Inc., 1972.
7. Anders L. Madsen and Finn V. Jensen. Lazy evaluation of symmetric Bayesian decision problems. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, pages 382–390, 1999.
8. Thomas Dyhre Nielsen and Finn Verner Nielsen. Sensitivity analysis in influence diagrams. *IEEE Transactions on Systems, Man, and Cybernetics*, 33(2):223–234, 2003.
9. Ross D. Shachter. Bayes-ball: The rational pastime (for determining irrelevance and requisite information in belief networks and influence diagrams. In G. F. Cooper and S. Moral, editors, *Uncertainty in Artificial Intelligence: Proceedings of the Fourteenth Conference*, pages 480–487. Morgan Kaufmann.
10. Ross D. Shachter. Evaluating influence diagrams. *Operations Research*, 34:871–882, 1986.
11. Prakash P. Shenoy. Valuation-based systems for Bayesian decision analysis. *Operations Research*, 40:463–484, 1991.
12. Claus Skaaning. A knowledge acquisition tool for Bayesian-network troubleshooters. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, pages 549–557, 2000.
13. Joseph A. Tatman and Ross D. Shachter. Dynamic programming and influence diagrams. *IEEE Transactions on Systems Management and Cybernetics*, 20:265–279, 1990.

Index

- currency, 2, 6
- currency mapping, 7
- expected global utility, 8
- expected utility, 5
- influence diagram, 1, 3
- multi-currency influence diagram, 1, 2, 6
 - dimension of, 7
 - realisation, 7
 - support analysis, 8, 9
- multi-currency utility potential, 9
- multiple stakeholders, 3, 8
- optimal strategy, 2, 5, 8
- support of strategy, 9
 - abstraction, 18
 - constraint, 15
 - constraints, 10
 - minimal, 15
 - representation, 18
 - superfluous constraint, 15
- utility potentials, 4, 6