Almost-Pareto decision sets in imprecise utility hierarchies

M. Farrow, *School of Mathematics and Statistics, Newcastle University, Newcastle upon Tyne, NE1 7RU, UK.*
Email: malcolm.farrow@newcastle.ac.uk
M. Goldstein, *Department of Mathematical Sciences, Durham University, Durham, DH1 3LE, UK.*
Email: Michael.Goldstein@durham.ac.uk

*Received: *** Revised: ****

Abstract

We develop methods for analysing decision problems based on multi-attribute utility hierarchies, structured by mutual utility independence, which are not precisely specified due to unwillingness or inability of an individual or group to agree on precise values for the trade-offs between the various attributes. Instead, our analysis is based on whatever limited collection of preferences we may assert between attribute collections. These preferences identify a class of Pareto optimal decisions. We show how to reduce the class further by combining rules which are almost equivalent and introduce general principles appropriate to selecting decisions in an imprecise hierarchy. The approach is illustrated by the design of a university course module.

Keywords: Robust decisions; Imprecise utilities; Pareto optimality; Utility hierarchies; Mutual utility independence.

1 Introduction

One of the most difficult steps in many decision analyses is the quantification of the relative importance of different types of risk. Therefore, it is of fundamental interest to develop methods for analysing multi-attribute utilities which do not require the specification of precise trade-offs between different risks. This paper is concerned with decision analysis using multi-attribute utilities which are not precisely specified, due to an unwillingness or inability to specify fixed risk trade-offs or from disagreement within a group with responsibility for the decision. We are particularly concerned with problems where the number of alternatives is large. Many real decision problems have very large spaces of possible decision rules. Relaxing the requirement for precise trade-off specification reduces our ability to eliminate rules by dominance and can leave us with a large class of rules, none of which is dominated by any other over the whole range of possible trade-offs allowed by the imprecise specification. We therefore need to find practical and tractable ways to reduce the decision space.

An earlier paper (Farrow and Goldstein, 2006) introduced our approach to constructing imprecise multi-attribute utility hierarchies and finding the Pareto optimal

rules in the context of experimental design. In section 3 of this paper, we describe this structure, based on a utility hierarchy with utility independence at each node. In section 4, we define imprecise utility trade-offs for such a hierarchy, based on limited collections of stated preferences between outcomes, and use Pareto optimality to reduce the set of alternatives. In section 5 we consider ways to reduce further the class of alternatives, by eliminating rules which are "ε-dominated" and combining rules which are "ε-equivalent." We explore the effects of different values of ε and of different parts of the hierarchy to identify when and why rules are eliminated. The approach is illustrated by an example which we describe in section 2.

2 Example: Designing a new course module at a university

To illustrate our approach, we consider the design of a new course module at a university. The example concerns the introduction into academic planning of utility-based consideration of attribute trade-offs between, for example, costs and the student learning experience. It is based on an actual course but some details have been changed and nothing in the example represents an official view or position at any university. However, in developing the example we benefited from discussions with Mr.W.Middleton of the University of Sunderland about his recent experiences of course design.

The module in question is for the first year of a degree course which contains mathematics, statistics and computing applicable to business as well as business studies topics. The aim of the module is to develop concepts and techniques in applicable mathematics which will give the student a solid foundation leading to successful study of mathematical modelling modules in subsequent years. The module is to contain six units, or topics, each of which may, for the purpose of this example, be considered to be of the same size in the sense that, given the same teaching method, they would require the same length of time. Each topic could be taught by any one of three teaching methods: (1) a traditional course of lectures and tutorials, (2) a laboratory-based course using a computer algebra package, (3) an "open learning" course in which students use prepared teaching material without lectures or formal laboratory sessions.

Thus we have $3^6 = 729$ possible choices of combinations of teaching methods. Observe that, in cases where we have to make a series of inter-related choices, the number of possible decisions can increase very quickly, and it is therefore important that our methods are able to cope with large numbers of alternatives. We can denote a choice $(\mu_1, \ldots, \mu_6)$ where $\mu_i = 1$,2 or 3 according to which method is used for unit i. The attributes which we consider in our analysis are

for students: short term learning as measured by grade at the end of the module; longer-term learning as measured by the grades achieved in later modules; satisfaction (measured by a feedback questionnaire);

for staff: satisfaction (measured through the module leader's annual report and the staff appraisal process); development (e.g. in terms of the acquisition of new skills and experience, measured through the annual staff appraisal process);

for the institution: benefit (e.g. in terms of being seen to be involved in "innovative" teaching methods and measured, for example, through the Teaching Quality Assessment); financial cost.

As for many problems, the attributes are in very different units and it may be difficult to establish precise attribute trade-offs on which to rank the teaching choices.

3 Utility Hierarchies

3.1 Mutually utility independent hierarchies

In this section we summarise the concepts and definitions, introduced in Farrow and Goldstein (2006), for a general class of multi-attribute utility functions. Attributes $\underline{Y} = (Y_1, ..., Y_k)$ are *utility independent* of $\underline{Z} = (Z_1, ..., Z_r)$ if conditional preferences over lotteries with differing values of $\underline{Y}$ but fixed values, $\underline{z}$, of $\underline{Z}$, do not depend on the particular choice of $\underline{z}$. Attributes $\underline{X} = (X_1, ..., X_s)$ are *mutually utility independent* if every subset of $\underline{X}$ is utility independent of its complement. If attributes $\underline{X}$ are mutually utility independent, then the utility function for $\underline{X}$ must be of the *multiplicative form*

$$(1 + kU(\underline{X})) = \prod_{i=1}^{s}(1 + ka_iU_i(X_i)), \tag{1}$$

or the *additive form*

$$U(\underline{X}) = \sum_{i=1}^{s} a_iU_i(X_i), \tag{2}$$

where $U_i(X_i)$ is a conditional utility function for attribute X_i, namely an evaluation of the utility of X_i for fixed values of the other attributes. (Mutual utility independence implies that $U_i(X_i)$ is unaffected by the choice of these fixed values.) The coefficients in (1) and (2) are the *trade-off parameters*; the a_i reflect the relative importance of the attributes and k reflects the degree to which rewards may be regarded as complementary, if $k > 0$, or as substitutes, if $k < 0$. See Keeney and Raiffa (1976).

We form an overall multi-attribute utility from marginal utilities for the various attributes by a hierarchical structure which may be represented by a graph. The marginal utilities are shown as *marginal nodes*, which have no predecessors. At each non-marginal node, i, several utilities, $U_{i,1}, \ldots, U_{i,s_i}$, are merged into a combined utility, U_i, as indicated by arcs from the "parent" nodes $i_1, \ldots i_{s_i}$ to the "child" node i. This combined utility is merged with others at a node in the next level until, finally, one overall utility function is formed. In each case we regard the children of a node as being at a "higher" level so that the overall utility node is at the "top" level. For each node i, we denote by $N(i)$, the *sub-hierarchy* under i, where $N(i)$ is the set of nodes containing i and all of its predecessors. If, at each node, we have mutual utility independence for the utilities combined at that node, then we term such a utility function a *Mutually Utility Independent Hierarchic (MUIH)* utility. Thus, in a MUIH utility, at each node we combine utilities using either (1) or (2). Note that we do not insist on mutual utility independence over the whole collection of attributes, only among the parents at each node. We divide the child nodes in the hierarchy into the following three types:
an *additive node*, where utilities are combined as in (2) with $\sum_{i=1}^{s} a_i \equiv 1$ and $a_i > 0$ for $i = 1, \ldots, s$;

a *binary node*, where exactly two utilities are combined, scaled as

$$U = a_1U_1 + a_2U_2 + hU_1U_2 \tag{3}$$

where $0 < a_i < 1$ and $-a_i \leq h \leq 1 - a_i$, for $i = 1, 2$, and $a_1 + a_2 + h \equiv 1$. Note that (3) is derived by setting $s = 2$ and $h = ka_1a_2$ in (1).
a *multiplicative node*, where more than two utilities are combined and the parameter k in (1) may be nonzero. We rescale the utility so that

$$U = \frac{\prod_{i=1}^{s}(1 + ka_iU_i) - 1}{\prod_{i=1}^{s}(1 + ka_i) - 1} \tag{4}$$

where $a_1 \equiv 1, k > -1$ and, for $i = 1, \ldots, s$, we have $a_i > 0$ and $ka_i > -1$.

For each child node i, we denote by $\underline{\phi}_i = (\phi_{i1}, \ldots, \phi_{im(i)})$ the collection of trade-off parameters which determine how the parent utilities at node i are combined to give the value at the child node. Thus, each ϕ_{ij} corresponds to an a_i in (2) an a_i or h term in (3), or an a_i or k in (4). If there are n child nodes, then we denote by $\underline{\theta} = (\underline{\phi}_1, \ldots, \underline{\phi}_n)$ the collection of all the trade-off parameters in the hierarchy.

As we shall vary the trade-off parameters, and thus the utilities at the child nodes, we require a *standard scale*, constructed as follows, for all utilities in the hierarchy, so that the interpretation does not depend on the choice of trade-off parameters. We norm all the marginal utilities to lie between 0, the worst outcome for the problem, and 1, the best outcome. The effect of our utility scalings is that, at each node i in the hierarchy, the utility is 1 for the outcome C_i when all marginal predecessor nodes have utility 1, and is zero for the outcome c_i when all marginal predecessor nodes have utility zero. Therefore, a utility value of u at node i may always be interpreted as the utility of a gamble giving C_i with probability u and c_i with probability $1 - u$, irrespective of the chain of trade-off parameters in the hierarchy. Throughout this paper, all utilities are assumed to be on the standard scale.

3.2 Example: Utility hierarchy

Continuing the example of section 2, the utility hierarchy is shown in figure 1.

The overall utility node U is a binary node, combining cost C and quality Q. The cost of teaching staff time is included in the financial cost although, in practice, there may be other considerations involved with the allocation of staff time. Staff time is used in both preparation and delivery. More preparation time is required in the first year of operation (very much more for open learning) but the cost considered here is the total over the five-year validation period. There are differences in the non-staff costs between teaching methods. While some resource costs, e.g. library, are common to all methods and can therefore be disregarded, there are extra costs for hardware and software, in the case of laboratory-based teaching, and for open learning material. The financial costs are summarised in table 1.

The utility for financial cost is $U_C = 1 - C/7500$ where C is the total cost, in £, of the module, calculated using table 1, and £7500 represents a "worst case" cost. This cost is traded against module quality Q so the overall utility is $U = a_QU_Q + a_CU_C + h_UU_QU_C$, where U_Q is the "module quality" utility. This is a binary node, given by

Figure 1: Utility hierarchy

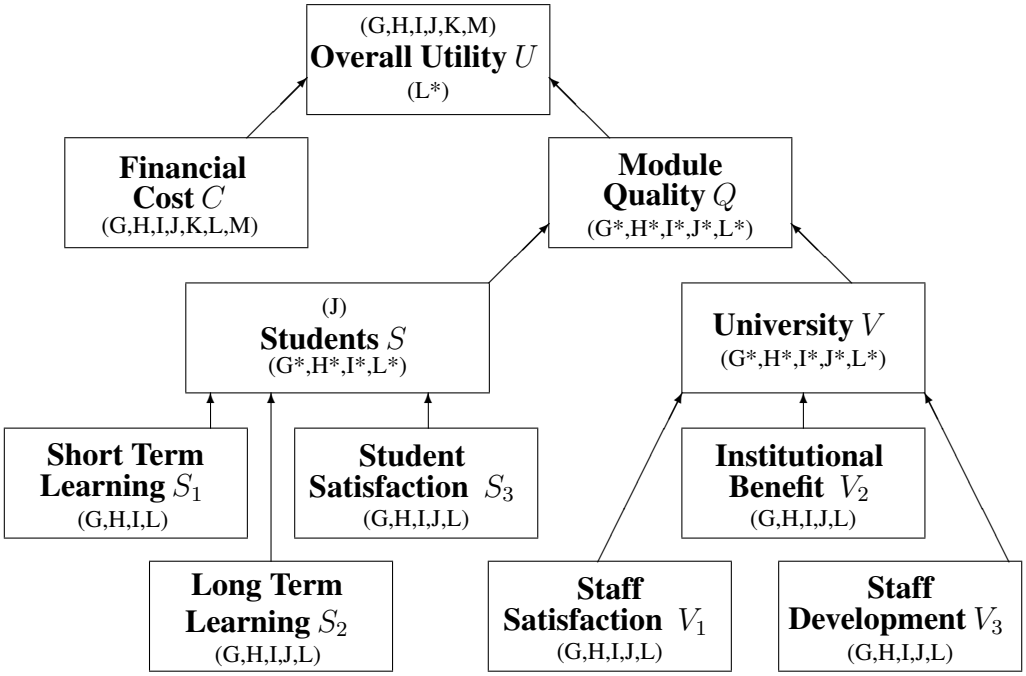

$U_Q = a_S U_S + a_V U_V + h_Q U_S U_V$, where U_S and U_V are the utilities for "Students" and "University". Each of these is an additive node which depends on three marginal attributes: $U_S = a_{S1}U_{S1} + a_{S2}U_{S2} + a_{S3}U_{S3}$, $U_V = a_{V1}U_{V1} + a_{V2}U_{V2} + a_{V3}U_{V3}$

The marginal attributes S_1, short term learning, and S_2, long term learning, are measured by the average final marks in this module and in a selection of second-year modules. The utilities for these in the example are simply the marks themselves on a $[0, 1]$ scale. Similarly, a score derived from the ratings given on student response forms, scaled to lie in $[0, 1]$, provides the marginal student satisfaction utility U_{S3}. For the university, the utility U_{V1}, representing staff satisfaction, is measured by the response to a rating scale in the annual staff self-appraisal. The institutional benefit V_2 is measured using the grades obtained in the Teaching Quality Assessment. Attribute V_3 represents the University's future ability to deploy staff experienced in the various teaching methods. An individual staff member's "experience level" for a particular method increases with the number of teaching hours in that method, but not linearly. The gain in "experience level" per hour is less when the accumulated experience is greater. Thus there is less to be gained in staff development terms from using a more traditional method. The utility function is fully specified when we assign values to all of the trade-off parameters in the above relations.

Table 1: Expected financial costs per unit per cohort. Suitable open learning material for Unit 2 had already been written.

	Staff	OL Material	Laboratory	Total
Lectures	£600	-	-	£600
Laboratory	£750	-	£300	£1050
Open Learning (Unit 2)	£300	£60	-	£360
Open Learning (other)	£600	£60	-	£660

4 Using Imprecise Trade-off Parameters

4.1 Imprecise utility trade-offs

Quantification of the various trade-off parameters between intrinsically different types of costs and benefits is a difficult task. Therefore, it is important to consider problems where we are unwilling to fix on particular trade-off values or where a group of individuals must make a joint decision, and there is broad agreement on the marginal utilities, but different members of the group have different priorities when trading risks.

The theory of imprecise probability can be built around the notion that, while we may be unwilling to specify a precise value for the probability of some event, there are various prices at which we would certainly buy a gamble on the outcome and other prices at which we would certainly sell such a gamble. These preferences may be used to construct upper and lower probabilities for the corresponding event. See e.g. Walley (1991). In the theory of imprecise utility we may apply a similar approach to the trade-offs between attribute values. There will be certain combinations of outcomes over which we are prepared to state preferences and these comparisons establish the region of the space of trade-off parameters which we must consider. We choose to elicit our imprecision in the values of the trade-off parameters $\underline{\theta}$ based on our stated preferences over utility combinations for outcomes, as this is usually more meaningful than considering directly the imprecision in the elements of $\underline{\theta}$. Farrow and Goldstein (2006) described the construction and properties of such imprecise utility hierarchies.

Some authors also consider imprecision in the marginal utility functions. Recent examples include Mateos et al. (2003) who describe a decision support system in which the imprecise multi-attribute utility function is additive and Jiménez et al. (2003), who allow a multiplicative function in which a range for the value of k in (1) is determined by considering the values implied by ranges given for $a_1, \ldots, a_s$. In both cases ranges for the trade-off parameters are combined to form a rectangular space. In this paper we only consider imprecision in trade-offs and assume that the necessary expectations of marginal utilities, and in some cases their products, can be agreed. However we do not impose an arbitrary probability distribution over ranges of imprecision, or over attributes, nor do we assume a rectangular shape for the space of trade-off parameters allowed by the imprecise specification resulting from a careful elicitation process.

If we allow imprecision in some of the elements of $\underline{\theta}$, then we refer to the resulting utility specification as an *imprecise independence hierarchy (IIH)*. If there are no mul-

tiplicative nodes in the hierarchy some features of the methodology become simpler and we refer to the specification as a *simple imprecise independence hierarchy (SIIH)*

For each child node, we make a collection of pairwise comparisons between vectors of values of parent utilities (or, equivalently, the corresponding vectors of attribute values). For example, in the hierarchy in figure 1, trade-off imprecision at node S is determined by eliciting pairwise preferences between values of the vector (S_1, S_2, S_3), while at node Q ranges are determined by preferences between vectors (S, V). In every case, preferences are expressed between lotteries over the corresponding marginal attributes. All utilities are expressed in the standard scale. For example, the utility S^* is the value of the lottery giving $(S_1 = 1, S_2 = 1, S_3 = 1)$ with probability S^*, and $(S_1 = 0, S_2 = 0, S_3 = 0)$ with probability $1 - S^*$.

At node i, we denote strict preference for utility vector $\underline{U} = (U_{1i}, U_{2i}, ..., U_{si})$ over utility vector $\underline{V} = (V_{1i}, V_{2i}, ..., V_{si})$ as $\underline{U} \succ^* \underline{V}$, and denote the weak preference, namely that we do not prefer $\underline{V}$ to $\underline{U}$ as $\underline{U} \succeq^* \underline{V}$. Each such preference places constraints on the allowable choices for the trade-off parameters $\underline{\phi}_i$. We term the collection, R, of all sets of trade-off parameters consistent with each of the stated preferences the *feasible* set of choices for the trade-off parameters. We say that the collection of pairwise comparisons is *consistent* if R is non-empty, i.e. there is at least one set of trade-off parameters satisfying all stated preferences. A comparison is *redundant* if its removal does not affect R. A change in any marginal utility will produce a change of the same sign in the overall utility. Therefore, when comparing two utility vectors where $A \succ^* B$, we consider whether we can decrease any of the utilities in A or increase any of the utilities in B while preserving the preference, to avoid redundancy and make stringent comparisons leading to the smallest feasible set that we can determine through our stated preferences. We describe a comparison as *sharp* if $(U_{11}, \ldots, U_{s1}) \succeq^* (U_{12}, \ldots, U_{s2})$ but there is no vector $(\delta_1, \ldots, \delta_s)$ such that $\delta_i \geq 0$, for $i = 1, \ldots, s$, and $\sum_{i=1}^{s} \delta_i > 0$ for which we are prepared to assert that $(U_{11} - \delta_1, \ldots, U_{s1} - \delta_s) \succeq^* (U_{12}, \ldots, U_{s2})$. We seek sharp comparisons to restrict R as far as possible. Note that the absolute limits for the values of tradeoff parameters provide bounds for the feasible set even before we specify any preferences. Specification of preferences results in the elimination of parts of this initial set.

For additive and binary child nodes, elicitation is unconstrained. For each such node, we state whichever preferences we wish between pairs of utility vectors for the parent nodes. Farrow and Goldstein (2006) showed that the shape of the resulting region of trade-off parameters for an SIIH is as follows. At each additive or binary node i, we obtain a convex polyhedron R_i for the allowable values of $\underline{\phi}_i$. The regions $R_1, \ldots, R_n$ together define a region R in the combined space of parameters $\underline{\theta}$, where $\underline{\theta} \in R$ if and only if $\underline{\phi}_i \in R_i$ for $i = 1, \ldots, n$. For each node i, we define P_i to be the set $\underline{\phi}_i = \{\phi_{i1}, \ldots, \phi_{ir(i)}\}$ of values at the $r(i) > m(i)$ vertices of R_i, and denote by P the set of overall vertex specifications for R, so that P is the subset of R with elements $\underline{\theta} = (\underline{\phi}_1, \ldots, \underline{\phi}_n) \in R$ such that $\underline{\phi}_i \in P_i$ for $i = 1, \ldots, n$.

For multiplicative nodes, the process of eliciting the feasible set for tradeoff parameters is modified. For additive and binary nodes we can express the combined utility linearly in the parameters and we constrain the sum of the parameters to be 1 without loss of generality. For a multiplicative node, we cannot express the combined utility linearly in the parameters but we still need to impose one constraint to make the param-

eters identifiable. Rather than attempt to constrain the sum of the weights on the parent utilities, which involve products of the parameters, we fix one of the a_i to be 1. Furthermore, if we allowed unconstrained comparisons at a multiplicative child node, then the boundary of the feasible set, even with respect to $a_1, \ldots, a_s$ for fixed k, would no longer be a convex polyhedron and comparisons over the region would become much more difficult. Therefore, we impose the following restrictions on elicited preferences at a multiplicative child node. First, at each such node, we choose one of the parent utilities to be a standard or reference utility. Suppose that this is U_1. We set $a_1 \equiv 1$. We then assess trade-off ranges for each of the other attributes at the node by comparing that attribute with the standard. Thus, in each comparison, we vary only one of the utilities at a time, so that we compare each attribute j with attribute 1. Denote by $\underline{U}_j^+(u)$ the vector $(U_1, \ldots, U_s)$ for which $U_j = u$ and $U_i = 0, j \neq i$. For each $j > 1$, we make at least one comparison of the form $\underline{U}_1^+(u_1) \succeq^* \underline{U}_j^+(u_j)$, to give each a_j an upper limit. We already have the constraint $a_j > 0$ but, optionally, we can make comparisons of the form $\underline{U}_j^+(u_j) \succeq^* \underline{U}_1^+(u_1)$ to impose tighter upper limits. Secondly, we elicit pairwise preferences between utility vectors of the form $\underline{U}_1^+(u)$ and vectors $(U_{12}, \ldots, U_{s2})$, where only $U_1 = U_{12}$ and $U_j = U_{j2}$ are positive. Such a comparison imposes a limit on a_j for fixed k and a limit on k for fixed a_j.

Farrow and Goldstein (2006) showed that the shape of the region of trade-off parameters resulting from the above elicitation scheme for an IIH is as follows. At each additive or binary node i, the shape is as for a SIIH. For each multiplicative node i, for each fixed value of k, we obtain a bounded rectangular region $R_i(k)$ for the remaining elements of $\underline{\phi}_i$. The region R_i of allowable specifications for $\underline{\phi}_i$ is the union of the collections $(k, R_i(k))$. For each fixed value of the remaining elements of $\underline{\phi}_i$, we obtain an interval for the value of k. The regions $R_1, \ldots, R_n$ together define a region R in the combined space of parameters $\underline{\theta}$, where $\underline{\theta} \in R$ if and only if $\underline{\phi}_i \in R_i$ for $i = 1, \ldots, n$.

The shape of the feasible region for a multiplicative node i is complex so we often choose to expand such a region to a more convenient shape $R^\star$. For additive and binary nodes we define $R_i^\star \equiv R_i$. For multiplicative nodes we define $R_i^\star$ as follows. Firstly, we identify the maximum and minimum values of k, denoted k_M, k_m, in the set R_i. For each trade-off parameter a_j, we define a_{jM}, a_{jm} to be the maximum and the minimum values of a_j respectively over the two sets $R_i(k_M), R_i(k_m)$. Denote by $R_i^\star$, the rectangular region bounded by the values $k_m < k < k_M, a_{jm} < a_j < a_{jM}, j = 2, \ldots$. Then $R_i^\star$ is the smallest rectangular region for which $R_i \subseteq R_i^\star$. When we replace each such R_i by $R_i^\star$, we denote the corresponding extended feasible region as $R^\star$. For each additive or binary node define $P_i^\star$ to be the set $\{\underline{\phi}_{i1}, \ldots, \underline{\phi}_{ir(i)}\}$ of the $r(i) > m(i)$ vertices of $R_i^\star$. For each multiplicative node with s parents define $P_i^\star$ to be the set of $\underline{\phi}_i$ such that $a_j = a_{jm}$ or $a_j = a_{jM}$, for $j = 2, \ldots, s$, and $k_m < k < k_M$. Then $P^\star$ is the set of overall specifications each element of which is the combination of one element from each of $P_1^\star, \ldots, P_n^\star$.

4.2 Example: specifying imprecise utility trade-offs

Consider node S in our example. We elicit preferences among vectors of attribute values. At this node, the utility values U_{Si} are equal to the attribute values S_i when known,

$i = 1, 2, 3$. For the vector (S_1, S_2, S_3), $(0.5, 0.65, 0.8) \succ^* (0.6, 0.5, 0.9)$. Further, reducing the value of U_{S2} in the first triple leads to uncertainty over the preference so that, for example, it is not agreed whether (0.5, 0.64, 0.8) should be preferred to (0.6, 0.5, 0.9). Therefore the comparison is sharp and we adopt $a_{S2} > 0.4$ as part of the boundary of our parameter region. Similarly, $(0.5, 0.5, 0.9) \succ^* (0.51, 0.51, 0.81)$ but we are unwilling to state preferences between (0.5, 0.5, 0.89) and (0.51, 0.51, 0.81). The definite preference here leads to $a_{S3} > 0.1$. Also $(0.58, 0.5, 0.7) \succ^* (0.5, 0.52, 0.72)$, but we do not state a preference between (0.57, 0.5, 0.7) and (0.5, 0.52, 0.72). Here, the definite preference leads to $a_{S2} + a_{S3} < 0.8$. These three inequalities define a triangular region. Note that we could impose more constraints but the region would always be the convex hull of a finite number of vertex points. In the absence of suitable information on definite preferences, the absolute limits would apply. For example, without the third inequality above, $a_{S2} + a_{S3} < 1$ would be the third side of the triangle.

As an example of a node higher in the hierarchy, consider node Q. Suppose that we offer a choice between the following alternatives: (1) with certainty, attribute values such that $U_S = 1$ and $U_V = 0$; (2) with probability α, attribute values such that $U_S = U_V = 1$ and, with probability $1 - \alpha$, attribute values such that $U_S = U_V = 0$. We find that (2) is preferred whenever $\alpha > 0.89$, so $h_Q \geq 0.11 - a_V$, and that (1) is preferred whenever $\alpha < 0.50$, so $h_Q \leq 0.50 - a_V$. Now suppose that we consider a choice between the following alternatives: (1) with probability $1/3$, attribute values such that $U_S = 1$ and $U_V = 0$ and with probability $2/3$, attribute values such that $U_S = 0$ and $U_V = 1$; (2) with probability α, attribute values such that $U_S = U_V = 1$ and, with probability $1 - \alpha$, attribute values such that $U_S = U_V = 0$. We find that (2) is preferred whenever $\alpha > 0.50$, so $h_Q \geq a_V - 0.5$, and that (1) is preferred whenever $\alpha < 0.37$, so $h_Q \leq a_V - 0.11$. These four constraints give the four sides of a quadrilateral region in the plane of a_V and h_Q. Table 2 gives the vertex set P_i for the elicited feasible set R_i, at each node i in our course module example.

4.3 Pareto optimal decisions

We have to choose from a set $\mathcal{D}$ of decisions. We denote the utility of a particular choice $A \in \mathcal{D}$, evaluated with trade-off parameters $\underline{\theta}$ as $U_{A\underline{\theta}}$. This is evaluated as the expected value of U_θ, with respect to the probability distribution, induced by the decision A, over the marginal attributes involved in U. For two alternatives, A, B, let $d_{AB}(\underline{\theta}) = U_{A\underline{\theta}} - U_{B\underline{\theta}}$. Different choices of trade-off parameters induce different preference orderings over the possible alternatives. A natural weak, partial preference ordering over allowable alternatives is that alternative A is at least as good as B, over feasible region R, written $A \succeq B$, if $U_{A\underline{\theta}} \geq U_{B\underline{\theta}}\ \forall \underline{\theta} \in R$, A is preferred to B, over R written $A \succ B$, if $A \succeq B$ and $U_{A\underline{\theta}} > U_{B\underline{\theta}}$ for some $\underline{\theta} \in R$, and A is equivalent to B, written $A \simeq B$, if $U_{A\underline{\theta}} = U_{B\underline{\theta}}\ \forall \underline{\theta} \in R$. We call alternative A *Pareto optimal* for R if there is no other allowable alternative B for which $B \succ A$ over R. We restrict attention to Pareto optimal alternatives. We form equivalence classes of equivalent decisions $A_1 \simeq A_2 \simeq ... \simeq A_r$, and restrict attention to only one representative member of each equivalence class. When we choose a decision, we may re-examine the corresponding equivalence class, to see whether there are any subsidiary criteria, not yet introduced into the problem formalism, which may distinguish between the members of the class.

Table 2: Trade-off parameter values.

Vertex	a_{S1}	a_{S2}	a_{S3}
$\underline{\phi}_{S1}$	0.2	0.4	0.4
$\underline{\phi}_{S2}$	0.2	0.7	0.1
$\underline{\phi}_{S2}$	0.5	0.4	0.1

(A) Parameter values at Node S (Students).

Vertex	a_{V1}	a_{V2}	a_{V3}
$\underline{\phi}_{V1}$	0.05	0.50	0.45
$\underline{\phi}_{V2}$	0.05	0.75	0.20
$\underline{\phi}_{V2}$	0.20	0.55	0.25

(B) Parameter values at Node V (University).

Vertex	a_S	a_V	h_Q
$\underline{\phi}_{Q1}$	0.890	0.110	0.000
$\underline{\phi}_{Q2}$	0.500	0.500	0.000
$\underline{\phi}_{Q3}$	0.890	0.305	-0.195
$\underline{\phi}_{Q4}$	0.500	0.305	0.195

(C) Parameter values at Node Q (Module Quality).

Vertex	a_C	a_Q	h_U
$\underline{\phi}_{U1}$	0.7	0.3	0.0
$\underline{\phi}_{U2}$	0.5	0.5	0.0
$\underline{\phi}_{U3}$	0.7	0.4	-0.1
$\underline{\phi}_{U4}$	0.5	0.4	0.1

(D) Parameter values at Node U (Overall Utility).

The operational research literature contains work concerned with mathematical, usually linear, programming methods for identifying Pareto optimal (or "non-dominated") and potentially optimal alternatives. In our terms, a potentially optimal alternative is an alternative A such that there exists $\underline{\theta}$ in the feasible set such that $d_{AB}(\underline{\theta}) \geq 0$ for all $B \neq A$. See, for example, Hazen (1986), Mateos et al. (2007). Our approach differs in that we express expectations over marginal utilities, we allow a more general form for the utility hierarchy and we allow a more general shape for the feasible set.

In Farrow and Goldstein (2006) we showed that, for any alternatives A, B,

$$\underset{R}{\operatorname{argmax}}\, d_{AB}(\underline{\theta}) \in P \qquad \text{and} \qquad \underset{R}{\operatorname{argmin}}\, d_{AB}(\underline{\theta}) \in P, \tag{5}$$

if we have an SIIH utility, and, for a general IIH utility, constructed as in section 4.1,

$$\underset{R^*}{\operatorname{argmax}}\, d_{AB}(\underline{\theta}) \in P^* \qquad \text{and} \qquad \underset{R^*}{\operatorname{argmin}}\, d_{AB}(\underline{\theta}) \in P^* \tag{6}$$

From these we showed that it is sufficient to check for Pareto decisions in the vertex set P or P^* since, if Q_S is the set of Pareto optimal alternatives in the set S, then

$$Q_R = Q_P, \tag{7}$$

for an SIIH utility, and, for a general IIH utility, constructed as in section 4.1,

$$Q_{R^*} = Q_{P^*}. \tag{8}$$

4.4 Example: Pareto decisions.

In the example there are two binary nodes, U and Q, and two additive nodes, S and V, each with $s = 3$. Thus, at every node in the example, the dimension of the parameter space is $m - 1 = 2$. Triangular regions are specified for nodes S and V with quadrilateral regions for nodes Q and U. Thus the overall parameter space has $3\times3\times4\times4 = 144$ vertices. We must choose between 729 alternatives, each being a particular selection of the teaching methods for the six units. For each decision, we elicit expectations of the utilities of the attributes. In this example, the attributes are considered to be stochastically independent, given a particular decision. Thus the expectations of products in utilities at binary nodes are just the products of the expectations at the parent nodes. The elicited expectations for the utilities of S_1 and S_2, if the module has n_1 lecture-based topics, n_2 laboratory-based topics and n_3 open-learning topics are

$$\begin{aligned}
\mathrm{E}(U_{S1}) &= \frac{55n_1 + 45n_2 + (60 - 5n_3)n_3}{600} - \frac{10\Delta_{1,2} + 2\delta_{62}\bar{\delta}_{52} + \delta_{52}\bar{\delta}_{42} + \delta_{42}\bar{\delta}_{32}}{100} \\
\mathrm{E}(U_{S2}) &= \frac{55n_1 + 60n_2 + (60 - 7n_3)n_3}{600} - \frac{5\Delta_{1,2} + 5\Delta_{3,2} - 5\delta_{62} - 4\delta_{12}}{100}
\end{aligned}$$

where $\delta_{ij} = 1$ if $m_i = j$, and $\delta_{ij} = 0$ otherwise, $\bar{\delta}_{ij} = 1 - \delta_{ij}$ and $\Delta_{j,2} = \prod_{i=j}^{6} \bar{\delta}_{i2}$.The quadratic terms in n_3 are present because open learning was judged to be more successful in small doses. The terms involving δ_{i2} and $\bar{\delta}_{i2}$ concern the laboratory-based elements of the course. The ability to use a computer algebra package is a learning outcome for the module so the complete absence of Method 2 is penalised. In some cases, using the computer algebra package in one unit helps with the following unit. Experience of laboratory-based work in certain units is important for Year 2.

Student satisfaction was expected to be good provided that not more than four of the units used the same method. The expectations, $\mathrm{E}(U_{S3})$, which were directly elicited, are given in Table 3. Staff satisfaction was judged to be worst with open learning and best with laboratory work so these were assigned expected utilities 0.1 and 0.9. Traditional lecturing was then given 0.7. Thus $\mathrm{E}(U_{V1}) = [0.7n_1 + 0.9n_2 + 0.1n_3]/6$. Similarly $\mathrm{E}(U_{V2}) = [0.1n_1 + 0.9n_2 + 0.7n_3]/6$ and $\mathrm{E}(U_{V3}) = [0.1n_1 + 0.9n_2 + 0.8n_3]/6$. To establish Pareto optimality we only need to make comparisons at the vertices. In our example we find that 50 of the original 729 alternatives are Pareto optimal. We can discard the 679 other alternatives, since each one is dominated by at least one of the Pareto optimal alternatives over the parameter range. Furthermore, we find that 37 of the Pareto optimal rules can be eliminated because they are equivalent to others which are retained, as, in this example, there are many pairs of decisions which have the same utility everywhere. The thirteen remaining rules are listed in Table 4 in which each is assigned a letter for identification later.

5 Almost Preferred Decisions

Given the imprecision in trade-off parameters, we seek a robust choice of decision. This necessitates a careful comparison between the Pareto decisions over the allowable

Table 3: Expected utilities for student satisfaction.

	n_2						
	0	1	2	3	4	5	6
$n_1 = 0$	0.3	0.5	0.9	0.9	0.9	0.7	0.6
1	0.4	0.8	0.9	0.9	0.9	0.8	-
2	0.8	0.8	0.9	0.9	0.9	-	-
3	0.8	0.8	0.9	0.9	-	-	-
4	0.8	0.8	0.9	-	-	-	-
5	0.6	0.7	-	-	-	-	-
6	0.5	-	-	-	-	-	-

range of trade-off parameters. In Farrow and Goldstein (2006) we introduced boundary linear utility as a means of choosing an alternative from among the Pareto optimal rules This is based on a linear combination of the utility functions at the vertices of the feasible set. Now, rather than move directly to a single criterion for selecting a single rule, we consider in more detail the comparisons between rules. If there are many alternative decisions, then there may be many Pareto optimal rules. Therefore, it is very helpful to reduce further the class of alternatives, both to simplify the subsequent comparisons of the remaining rules and also to allow us to make simple displays of the leading alternatives to support our formal arguments.

5.1 Almost dominance and almost equivalence

We may decide to eliminate decision A from consideration, even if A is Pareto optimal, provided that A is almost dominated, namely that there is an alternative B whose utility, over the entire feasible trade-off region, is never sufficiently less than the utility of A to be of any practical concern. We therefore define *ε-preference*, *ε-equivalence* and *ε-dominance* as follows, where $\varepsilon \geq 0$ is a value chosen to indicate a practical indifference between utility values. For two alternatives A and B, we say that A is ε-preferable to B, written $A \succeq_\varepsilon B$, over the set Q of parameter specifications if $\inf_Q(d_{AB}(\underline{\theta})) \geq -\varepsilon$. Two alternatives A, B are said to be *ε-equivalent*, written $A \simeq_\varepsilon B$, if both $A \succeq_\varepsilon B$ and $B \succeq_\varepsilon A$. Alternative A is said to *ε-dominate* alternative B, written $A \succ_\varepsilon B$, if $A \succeq_\varepsilon B$ but $B \not\succeq_\varepsilon A$, where the negation of the relationship is indicated in the usual way. Setting $\varepsilon = 0$, an alternative which is not 0-dominated by any other is Pareto optimal. We extend the preference notation to collections of alternatives as follows. The collection $\mathcal{A}$ is ε-preferable to the collection $\mathcal{B}$ of alternatives, written $\mathcal{A} \succeq_\varepsilon \mathcal{B}$ if, for each $B \in \mathcal{B}$, there is at least one $A \in \mathcal{A}$ for which $A \succeq_\varepsilon B$.

For any two alternatives A, B, if $A \succ_\varepsilon B$ it does not follow that $A \succ_\eta B$ where $\eta > \varepsilon$. However, if $A \succ_\varepsilon B$ but $A \not\succ_\eta B$, then it follows that $A \simeq_\eta B$. If $A \simeq_\varepsilon B$, then either there is an $\eta < \varepsilon$ for which $A \succ_\eta B$, in which case we might eliminate B, or a value for which $B \succ_\eta A$, in which case we might eliminate A, or there is no such value, when we might eliminate either. To determine which of these is the case, we may compare $\inf[d_{AB}(\underline{\theta})] = m_{AB}$ with $\inf[d_{BA}(\underline{\theta})] = M_{AB}$. We would

often prefer to eliminate B rather than A if $M_{AB} > -m_{AB}$ since this implies that, for $-m_{AB} < \varepsilon < M_{AB}$, $A \succ_\varepsilon B$. However this consideration does not give us a complete ordering of the alternatives as the following simple example illustrates. Suppose we have three alternatives, A, B, C, and three vertices and the utilities of A, B, C are 0.6, 0.4, 0.5 at the first vertex, 0.6, 0.7, 0.5 at the second and 0.6, 0.7, 0.8 at the third. Then, with $\varepsilon = 0.15$, $A \succ_\varepsilon B$, $B \succ_\varepsilon C$ and $C \succ_\varepsilon A$. Therefore, in any algorithm based on ε- preference, we will need to be careful as to which rules we choose to eliminate.

To define *ε-preference*, *ε-equivalence* and *ε-dominance* at a child node i we consider the utility at i under variation only of the parameters in $N(i)$, the sub-hierarchy under i. Consider first a SIIH. Just as we defined P and R for the whole hierarchy, we define $P(i)$ and $R(i)$ for $N(i)$. If $A \succeq_\varepsilon B$ over $P(i)$ we say that $A \succeq_\varepsilon B$ at node i. Clearly (5) and (7), applying to the whole hierarchy, also apply to any sub-hierarchy. Just as Pareto optimality over the whole feasible region corresponds to Pareto optimality at the vertex set, a similar equivalence holds for each ε comparison, at each level in the hierarchy. We have the following corollary, showing that to assess each comparison we only need to evaluate $d_{AB}(\underline{\theta})$ at the vertices, which is deduced directly from (5).

Corollary 1 *For alternatives A and B and a node $i, respectively$ $A \succeq_\varepsilon B$, $A \simeq_\varepsilon B$, $A \succ_\varepsilon B$, over $R(i)$ if and only if $A \succeq_\varepsilon B$ $A \simeq_\varepsilon B$, $A \succ_\varepsilon B$, over $P(i)$.*

In the case of a IIH we use (6) and (8) instead of (5) and (7) and replace $R(i)$ and $P(i)$ in Corollary 1 with $P^*(i)$ and $R^*(i)$.

5.2 Reducing the collection of alternative decisions

Having selected one representative from each equivalence class of the Pareto optimal rules, we now eliminate from further consideration alternatives which are almost dominated or almost equivalent to others by finding ε-Pareto decision sets for a range of values of ε. Let our set of Pareto optimal rules be $\mathcal{D}$. Then $\mathcal{A} \subseteq \mathcal{D}$ is an *ε-Pareto decision set* if $\mathcal{A} \succeq_\varepsilon \mathcal{B}$ where $\mathcal{A} \cup \mathcal{B} = \mathcal{D}$ and $\mathcal{A} \cap \mathcal{B} = \emptyset$. We choose values of ε on the standard utility scale for the hierarchy. We determine, for each ordered pair of rules A, B the value of $m_{AB} = \inf(d_{AB}(\underline{\theta}))$. From these values, we construct a list of decisions and the ε values at which they are just deleted by ε-preference. Increasing the value of ε eliminates progressively more alternatives, so that we must balance between preference for small ε and our aim of eliminating many decisions from consideration.

We form two lists of alternatives, $\mathcal{A}$ for those selected for retention and $\mathcal{B}$ for those selected for deletion, at the current value of ε. We start with $\varepsilon = 0$, and gradually increase ε, so initially all Pareto optimal alternatives are in $\mathcal{A}$ and $\mathcal{B}$ is empty. Suppose that there are n_P Pareto rules. We would like to find, for each $n = n_P - 1, n_P - 2, n_P - 3, \ldots, 1$ the smallest value of ε such that we may choose an ε-Pareto set $\mathcal{A}$ containing exactly n rules. Because ε-preference is not transitive, we impose a pragmatic constraint in the algorithm that we use to create these collections, namely that at each step, as we decrease n, rules already in $\mathcal{B}$ stay in $\mathcal{B}$. This constraint is imposed to simplify both the search procedure and the interpretation of the results. We use the following algorithm: repeat steps 1 to 3 until there is only one rule left in $\mathcal{A}$.
[1] For each rule $C \in \mathcal{A}$: (i) form $\mathcal{A}(C) = \mathcal{A} \backslash \{C\}$ and $\mathcal{B}(C) = \mathcal{B} \cup \{C\}$; (ii) for each $B \in \mathcal{B}(C)$ find $\max_{A \in \mathcal{A}(C)} m_{AB} = m_B(C)$; (iii) find $\min_{B \in \mathcal{B}(C)} m_B(C) = m^\star(C)$.

[2] Choose $C^\star$ which has the greatest value of $m^\star(C)$.
[3] Move $C^\star$ to $\mathcal{B}$. If there are now k rules in B, record the choices $C_k = C^\star$ and $\varepsilon_k = -m^\star(C_k)$.

For each $k = 1, \ldots, n_P - 1$, we have divided the Pareto rules into two sets, the eliminated rules $\mathcal{B}_k = \{C_1, C_2, ...C_k\}$, and the remaining rules retained in $\mathcal{A}_k$, where $\mathcal{A}_k \succeq_{\varepsilon_k} \mathcal{B}_k$. We can now see which alternatives are retained at each value of ε, and therefore consider how best to balance the conflicting requirements in our choice of ε.

While the above algorithm automatically suggests decisions to eliminate, we often want to compare the performance of the rules carefully before accepting such suggestions. We explore the elimination process in two ways, each referring to elimination at specific sub-nodes. This involves using the algorithm which we have just defined, suitably adapted to refer to the sub-hierarchy under a particular node.

First, we apply the selection algorithm at each node in the hierarchy, listing retained alternatives at each node. This allows us to examine the results of selection based only on particular subsets of the attributes and to see whether similar sets of rules are being selected in different parts of the hierarchy, and, if not, in which ways they differ. (It may be that lists differ simply because different members of ϵ-equivalence classes have been selected and this is checked by producing a list of equivalences for each alternative.)

Secondly we may look at stepwise elimination as we move through the hierarchy. Suppose that we have chosen a value for ϵ from the global elimination algorithm. Let the set of Pareto optimal rules be $\mathcal{A}$. Starting at the overall utility node and working down the hierarchy towards the marginal nodes we form two lists, $\mathcal{N}_1$ and $\mathcal{N}_2$, at each non-marginal node n. (At a marginal node n, there will be a single list $\mathcal{N}_1$. For notational simplicity, we set $\mathcal{N}_2 \equiv \mathcal{N}_1$ at n.) Lists $\mathcal{N}_1$ and $\mathcal{N}_2$ are formed as follows.
[1] From $\mathcal{A}$ we select those rules which would be eliminated at node n, using the chosen ϵ, for every possible combination of trade-off parameters. This is done by extending all parameter ranges at n to their maximum possible values. The eliminated rules are put in $\mathcal{N}_1$.
[2] We then introduce the restriction that the specified parameter ranges at n must fall within the feasible region that we have elicited for these parameters, and determine the further rules which are eliminated. Any additional rules eliminated, which are not in $\mathcal{N}_1$, form $\mathcal{N}_2$. (At a marginal node there are no trade-off parameters so we obtain a single list.)
[3] For ease of interpretation, we modify the selection algorithm, at both steps [1] and [2], so that any rules which are not eliminated at the child node of node n are automatically retained at node n.

The resulting lists can be displayed on a diagram of the utility hierarchy so that the elimination of a particular rule can be traced to its roots. Looking first at marginal nodes, and moving towards the overall utility, we see the effect of gradually reducing the range of possible parameter values by introducing ranges at the various nodes. This allows alternatives to be eliminated progressively and we can observe the effect on the retained list of each new node. We may eliminate A at a child node n if there is a retained choice B for which $B \succeq_\varepsilon A$. This dominance may be a consequence of restrictions imposed on the trade-off parameters at n or this may be because $B \succeq_{\varepsilon'} A$ at each parent node of n, where ε' is sufficiently small to ensure that $B \succeq_\varepsilon A$ at n, whatever the trade-off parameters at that node. In the latter case, we mark rule A with

an asterisk on the diagram at n, to indicate this. We now show that, when n is an additive node, the value $\varepsilon' = \varepsilon$ is sufficient for this purpose, if n is binary, $\varepsilon' = \varepsilon/2$ is sufficient and, if n is multiplicative with s components, $\varepsilon' = \varepsilon/s$ is sufficient.

5.3 Stepwise elimination

When we have many alternatives to compare, it is often computationally simplest to determine ε-preference at the lower levels of the hierarchy, as there are fewer trade-offs to consider. Therefore, it is natural to consider whether dominance at all parent nodes implies dominance for the corresponding child node, as this may suggest ways to reduce the number of competing alternatives that we must consider, leading to substantial simplifications in the algorithm of section 5.2.

For additive nodes, as we shall show, the relationship is very simple, namely that if $A \succeq_\varepsilon B$ at every parent of a node N, then $A \succeq_\varepsilon B$ at N. For binary and multiplicative child nodes the relation is more complicated, as the utility at the child node involves the product of the parent utilities and so, in general, is not deducible from the parent utilities. The special case where we may deduce the child utility from the parent utilities corresponds to the condition that the parents of any binary or multiplicative node are stochastically independent, given any decision in $\mathcal{D}$. We call an SIIH in which this condition holds a *simple imprecise stochastic independence hierarchy* (SISIH). Similarly an IIH with this additional property is called an *imprecise stochastic independence hierarchy* (ISIH). In a SISIH or ISIH the utility of a decision at any node may be calculated directly from the utilities of the decision at the parent nodes and the values of the trade-off parameters. The following theorem shows that ε-preference is preserved at additive nodes and that it is conservatively modified to $\varepsilon/2$-preference at binary nodes or ε/s-preference at multiplicative nodes.

Theorem 1 *If, for two alternatives A, B, we have $A \succeq_\varepsilon B$ at every parent of a node N with s parents, then (i) in a SISIH $A \succeq_{\varepsilon'} B$ at N where $\varepsilon' = \varepsilon$ in additive nodes and $\varepsilon' = 2\varepsilon - \varepsilon^2$ in binary nodes, (ii) in an ISIH, $A \succeq_{\varepsilon'} B$ at N where $\varepsilon' = \varepsilon$ in additive nodes and $\varepsilon' = 1 - (1-\varepsilon)^s$ in binary and multiplicative nodes.*

Proof: Let the utility of alternative D at parent i be U_{iD}. By the stochastic independence assumption expectations of all necessary products of parental utilities are the products of the expectations. Because the utility at N is increasing in every parental utility, and therefore, given independence where necessary, in their expectations, to attain the minimum for d_{AB} we must set each $U_{iA} - U_{iB}$ to its minimum value, i.e. $-\varepsilon$. Therefore we set $U_{iB} = U_{iA} + \varepsilon$.

1. Additive nodes. In this case $d_{AB} = \sum a_i U_{iA} - \sum a_i U_{iB}$. Setting $U_{iB} = U_{iA} + \varepsilon$ we obtain $d_{AB} = -\varepsilon$.

2. Binary nodes. In this case

$$d_{AB} = a_1(U_{1A} - U_{1B}) + a_2(U_{2A} - U_{2B}) + h(U_{1A}U_{2A} - U_{1B}U_{2B}).$$

Setting $U_{iB} = U_{iA} + \varepsilon$,

$$\begin{aligned} d_{AB} &= -a_1\varepsilon - a_2\varepsilon + h\{U_{1A}U_{2A} - (U_{1A}+\varepsilon)(U_{2A}+\varepsilon)\} \\ &= -a_1\varepsilon - a_2\varepsilon + h\{-U_{1A}\varepsilon - U_{2A}\varepsilon - \varepsilon^2\}. \end{aligned} \tag{9}$$

(a) $\underline{h < 0}$ With $h < 0$, (9) is minimised by making U_{1A}, U_{2A} as small as possible, i.e. $U_{1A} = U_{2A} = 0$. Hence (9) becomes $d_{AB} = -a_1\varepsilon - a_2\varepsilon - h\varepsilon^2$. The minimum of this must be at one of the vertices of the permissible region for a_1, a_2, h and it is easily confirmed that it occurs when $a_1 = a_2 = 1, h = -1$ in which case $d_{AB} = -2\varepsilon + \varepsilon^2$.

(b) $\underline{h > 0}$ With $h > 0$, (9) is minimised by making U_{1A}, U_{2A} as large as possible, i.e. $U_{1A} = U_{2A} = 1 - \varepsilon$. Hence (9) becomes $d_{AB} = -a_1\varepsilon - a_2\varepsilon - h(-2\varepsilon+\varepsilon^2)$. The minimum of this must be at one of the vertices of the permissible region for a_1, a_2, h and it is easily confirmed that it occurs when $a_1 = a_2 = 0, h = 1$ in which case $d_{AB} = -2\varepsilon + \varepsilon^2$.

3. Multiplicative nodes. With $k \neq 0$ we can write

$$d_{AB} = \frac{\prod(1 + ka_iU_{iA}) - \prod(1 + ka_iU_{iB})}{\prod(1 + ka_i) - 1}.$$

Setting $U_{iB} = U_{iA} + \varepsilon$ we obtain

$$d_{AB} = \frac{\prod[1 + ka_iU_{iA}] - \prod[1 + ka_i(U_{iA} + \varepsilon)]}{\prod[1 + ka_i] - 1}. \tag{10}$$

This is linear in any single U_{iA} with gradient

$$\frac{\partial d_{AB}}{\partial U_{iA}} = ka_i \left\{ \frac{\prod_{j\neq i}[1 + ka_iU_{iA}] - \prod_{j\neq i}[1 + ka_i(U_{iA} + \varepsilon)]}{\prod[1 + ka_i] - 1} \right\}.$$

Since $\prod_{j\neq i}[1+ka_iU_{iA}]-\prod_{j\neq i}[1+ka_i(U_{iA}+\varepsilon)] < 0$ the sign of this derivative is that of $-k$.

(a) $\underline{k < 0}$ With $k < 0$ (10) is minimised by making U_{iA} as small as possible, i.e. $U_{iA} = 0$. Hence (10) becomes

$$d_{AB} = \frac{1 - \prod[1 + ka_i\varepsilon]}{\prod[1 + ka_i] - 1}.$$

Choose any j in $1 \leq j \leq s$ and let $ka_j \to -1$, its upper limit. Then

$$d_{AB} \to (1 - \varepsilon)\prod_{i\neq j}(1 + ka_i\varepsilon).$$

This is clearly minimised by letting every other $ka_i \to -1$ so the minimum of d_{AB} is $(1 - \varepsilon)^s - 1$.

(b) $\underline{k > 0}$ With $k > 0$, (10) is minimised by making U_{iA} as large as possible, i.e. $U_{iA} = 1 - \varepsilon$. Hence (10) becomes

$$d_{AB} = \frac{\prod[1 + ka_i(1-\varepsilon)] - \prod[1 + ka_i]}{\prod[1 + ka_i] - 1}$$

Choose any j in $1 \leq j \leq s$ and hold ka_j constant while letting every other $ka_i \to \infty$. Then

$$d_{AB} \to \frac{[1 + ka_j(1-\varepsilon)](1-\varepsilon)^{s-1}}{(1 + ka_j)} - 1.$$

This is minimised by letting $ka_j \to \infty$. Hence the minimum of d_{AB} is $(1-\varepsilon)^s - 1$.

□

Observe that, for each node-type, the case $\varepsilon = 0$ implies that Pareto dominance for all parents always implies Pareto dominance for the corresponding child node.

We may therefore adjust the elimination procedure as follows. Our overall aim is as for the algorithm in section 5.2, namely to eliminate rules at stated levels of ε preference. If the overall node is an additive node, then a sufficient condition to ensure ε preference at this node is to require ε preference for each parent. Alternately, if the overall node is a binary node, then we must reduce the value of ε that we require at all parent nodes, so that the required value of ε will apply at the overall utility node. As a conservative approximation, we may divide ε by 2 to obtain, roughly, the value required. Simlarly, at a multiplicative node, we divide ε by s. We may continue in this fashion, working down from the top to the bottom of the hierarchy, and every time we pass down through a binary node, dividing ε by 2 and every time we pass down through a multiplicative node, dividing ε by s to obtain, conservatively, the required threshold for ε-preference at that node. Our algorithm now proceeds as follows (replacing P with P^* where there are multiplicative nodes).

[1] We evaluate the utility of each alternative at each of the marginal nodes.

[2] At any stage some of the nodes are considered to be *entered.* Initially the marginal nodes and no others are entered. At each subsequent stage we enter one more node, each one only after all of its predecessors have already been entered. Any node which is entered but whose successor is not is said to be *active.*

[3] At each stage, at each active node n we determine, for each ordered pair A, B of retained alternatives, the value of $m_{AB} = \inf_{P(n)}(d_{AB}(\underline{\theta}))$. We now consider ε-preference over all active nodes. This allows us to eliminate at this stage any alternative B to which another, retained, alternative A is ε-preferred at all active nodes. If $A \succeq_\varepsilon B$ at all of these nodes then it must be ε'-preferred at all succeeding nodes, where ε' may differ from ε because of binary or multiplicative nodes. The algorithm is the same as before except that, if nodes $n_1, \ldots, n_s$ are active, the values of M_{AB} and m_{AB} are taken over $\bigcup_{i=1}^{s} P(n_s)$. Much of the calculation required as each node is entered will already have been done when its predecessors were introduced. The list of alternatives eliminated as each node is entered is displayed.

[4] We continue until we reach the overall utility node.

Table 4: Selecting alternatives for retention: order of elimination.

	Order	Alternative	ε		Order	Alternative	ε
A	13	1, 3, 1, 1, 3, 2		H	6	1, 3, 3, 3, 3, 2	0.0105
B	12	2, 3, 1, 3, 3, 2	0.0403	I	5	2, 3, 3, 3, 3, 2	0.0105
C	11	1, 3, 1, 3, 3, 2	0.0205	J	4	1, 3, 1, 1, 1, 2	0.0103
D	10	1, 3, 1, 1, 1, 3	0.0180	K	3	2, 3, 1, 1, 1, 2	0.0079
E	9	1, 3, 1, 1, 3, 3	0.0145	L	2	1, 3, 1, 1, 1, 1	0.0040
F	8	2, 3, 1, 1, 3, 2	0.0135	M	1	2, 3, 3, 3, 2, 2	0.0032
G	7	1, 3, 1, 3, 3, 3	0.0105				

[5] It is possible that an alternative A may be eliminated at some stage because there is another retained alternative B which is ε-preferred to A, but at a later stage B may be also dropped. As ε-preference is not transitive, we check at the final stage whether we should reinstate any alternatives dropped at earlier stages by searching among the alternatives we have eliminated for any to which no retained alternative is now ε-preferred.

If the stochastic independence condition that we require for an SISIH or ISIH does not hold at all binary nodes in a SIIH or at all binary and multiplicative nodes in an IIH, then we need to modify the stepwise procedure. Starting from the overall utility node, we work downwards and determine the lowest node, down each branch, which can be reached without passing through any binary or multiplicative nodes which violate the independence condition. Where such a node is not a marginal node, we cannot assume that dominance and equivalence can be lifted from the marginal nodes up to this node and so we need to evaluate directly expected utilities at this node looking at the vertices of the region defined by the ranges at this node and its predecessors.

5.4 Example: eliminating alternatives

In the example, we have found 13 Pareto optimal rules. Applying the algorithm in section 5.2 we obtain the results in Table 4. In each case the value of ε given is the value at which it becomes possible to eliminate the rule. For example, it becomes possible to eliminate alternative $B = (2, 3, 1, 3, 3, 2)$ only when ε is increased to 0.0403.

A reasonable choice to balance our wish that ε should be small with our aim to eliminate many decisions is the value $\varepsilon = 0.012$. At the central value of the parameters at node U, this would correspond to a change of 0.02 in U_C, which corresponds to a change of £150 in the overall cost of the module. Such a change would be of little consequence. Using our chosen value of $\varepsilon = 0.012$ we retain only the first six choices. The fact that teaching method 3 is used for unit 2 in each of these is because the open learning material is already written for this unit, hence the lower cost.

The first five alternatives in the ε-preference ordering at node Q are M,F,K,B,J. Only the first three, M,F and K would be retained with $\varepsilon = 0.012$. One of these is also in the retained list at the overall utility node and the other two would be if we reduced the effective value of ε to 0.003. The extra retained rules in the overall list show the influence of Node C on the final choice. At Node S only rule K (2,3,1,1,1,2) is retained

and at node V only rule M (2,3,3,3,2,2) is retained when $\varepsilon = 0.012$. Each of these is retained at Node Q. The contrast between these two highlights an interesting difference between what might be done for the "student" and "University" interests.

The results of the top-down exploration of the elimination of the seven rules G, ..., M are displayed on figure 1, with $\varepsilon = 0.012$. At each node, list 1, of the rules eliminated regardless of the parameter values, is displayed below the node name and list 2, of the additional rules eliminated when the ranges are introduced, is shown above. We see that G,H,I are eliminated everywhere and we can discard them with confidence. Rule J is eliminated everywhere except at Node S_1. The financially more costly rules M,K are only eliminated at Node C. The most pronounced conflict appears to be between financial cost and the other attributes. Notice that it is necessary to introduce the parameter range at node U before any rules, other than L, are eliminated. In contrast, in the sub-hierarchy under node Q, rules tend to be eliminated without the need to specify the range. Indeed the asterisks indicate that, in many cases, there are rules which are ε-preferred across all of the parents of a node. Overall our exploration suggests that the choices in this example are relatively insensitive to the imposition of the parameter ranges, with the possible exception of node U.

Applying the stepwise procedure of section 5.3 over the SISIH, makes no difference in our example. None of the rules A-M is eliminated until the overall utility node U is entered. Then G-M are eliminated as before. This can be attributed to the fact that the main conflict of utilities among A-M is between financial cost and module quality.

The last rule to be eliminated in table 4 is Rule A. This is also the rule chosen if we apply the boundary linear utility with equal weights on all vertices. Table 5 gives the first eighteen lines of a list of all alternatives, whether Pareto optimal or not, ordered according to the smallest value of ε, given in the final column, at which they are just ε-preferred to A. Here A is denoted A1, Rules A1-A4 have the same utility everywhere in R. We can see that they only differ in the position of one of the two open-learning units, the other being Unit 2. A choice between these four rules might be made, for example, on timetabling grounds. No other rule has a critical ε value with respect to A of less than 0.012. The next group, N1-N2, the members of which are not Pareto optimal, have a critical ε value of 0.0138 and differ from A1-A2 only in that the methods for Units 3 and 6 are exchanged. Unexpected circumstances might make this advantageous and we see that little would be lost in terms of overall utility by this change. The first Pareto optimal rule to appear in the list after A1-A4 is Rule J in seventeenth place.

In more complicated applications we could repeat the calculation used to produce table 5 at each node, working our way up the hierarchy, and adjusting the value of ε according to theorem 1 to explore why particular rules are being selected. Another possibility might be a formal introduction of secondary criteria, in a second stage analysis, to help choose between rules which are almost equivalent in our initial analysis.

The example analysis leads us to conclude that one of A1-A4 would be a good choice, without having to collapse our trade-off ranges to give a fully specified utility function. Our example was chosen to be small enough to use as an illustration and many problems could involve much larger decision spaces and much greater potential for disagreement between stakeholders but this makes our approach to structuring the decision problem and reducing the number of potential choices even more important.

Table 5: Alternatives almost preferable to A.

	Order	Alternative	ε		Order	Alternative	ε
A1	1	1, 3, 1, 1, 3, 2	0.0000	O2	9	1, 3, 1, 2, 1, 3	0.0147
A2	1	1, 3, 1, 3, 1, 2	0.0000	O3	9	1, 3, 1, 2, 3, 1	0.0147
A3	1	1, 3, 3, 1, 1, 2	0.0000	O4	9	1, 3, 1, 3, 2, 1	0.0147
A4	1	3, 3, 1, 1, 1, 2	0.0000	O5	9	1, 3, 3, 1, 2, 1	0.0147
N1	5	1, 3, 2, 1, 1, 3	0.0138	O6	9	1, 3, 3, 2, 1, 1	0.0147
N2	5	1, 3, 2, 1, 3, 1	0.0138	O8	9	3, 3, 1, 1, 2, 1	0.0147
N3	5	1, 3, 2, 3, 1, 1	0.0138	O9	9	3, 3, 1, 2, 1, 1	0.0147
N4	5	3, 3, 2, 1, 1, 1	0.0138	J	17	1, 3, 1, 1, 1, 2	0.0169
O1	9	1, 3, 1, 1, 2, 3	0.0147	P1	18	2, 3, 1, 1, 1, 3	0.0169

Acknowledgement

We are grateful to Walter Middleton for his help in formulating the example.

References

Farrow M., Goldstein M., 2006. Trade-off sensitive experimental design: a multicriterion, decision theoretic, Bayes linear approach. *Journal of Statistical Planning and Inference*, 136, 498–526.

Hazen G.B., 1986. Partial information, dominance and potential optimality in multiattribute utility theory. *Operations Research*, 34, 296–310.

Jiménez A., Ríos-Insua S., Mateos A., 2003. A decision support system for multiattribute utility evaluation based on imprecise assignments. *Decision Support Systems*, 36, 65–79.

Keeney R.L., Raiffa H., 1976. *Decisions with Multiple Objectives: Preferences and Value Trade-offs*. John Wiley & Sons, New York.

Mateos A., Jiménez A., Ríos-Insua S., 2003. Modelling individual and global comparisons for multi-attribute preferences. *Journal of Multi-Criteria Decision Analysis*, 12, 177–190.

Mateos A., Jiménez A., Ríos-Insua S., 2007. Dominance, potential optimality and alternative ranking in imprecise multi-attribute decision making. *Journal of the Operational Research Society*, 58, 326–336.

Walley P., 1991. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London.