Continuous Decision MTE Influence Diagrams

Barry R. Cobb
cobbbr@vmi.edu
Department of Economics and Business
Virginia Military Institute
Lexington, VA, 24450 U.S.A.

Abstract

Mixtures of truncated exponentials (MTE) influence diagrams can represent decision problems with discrete decision variables without limitations on the distributions of continuous chance variables or the nature of the utility functions. This paper presents an extension of MTE influence diagrams that allows both discrete and continuous decision variables. This new model—the continuous decision MTE influence diagram—develops decision rules for continuous decision variables as a function of their continuous parents. In this model, a continuous decision node is marginalized by replacing it with a deterministic chance node and applying an operation derived from the method of convolutions in probability theory. In experiments, continuous decision MTE influence diagrams provide an increase in maximum expected utility when compared to existing methods.

1 Introduction

An *influence diagram* is a compact graphical representation for a decision problem under uncertainty. Initially, influence diagrams were proposed by Howard and Matheson (1984) as a front-end for decision trees. Subsequently, Olmsted (1983) and Shachter (1986) developed methods for evaluating an influence diagram directly without converting it to a decision tree. These methods assume that all uncertain variables in the model are represented by discrete probability mass functions (PMF's) and that decision variables have discrete state spaces.

Influence diagram models that permit continuous chance and decision variables include Gaussian influence diagrams (Shachter and Kenley, 1989), mixtures of Gaussians influence diagrams (Poland and Shachter, 1993), and linear-quadratic conditional Gaussian influence diagrams (Madsen and Jensen, 2005). Each of these models exploits multivariate normal probability theory to provide the decision strategy and expected utility that solves the influence diagram.

Cobb and Shenoy (2004) introduce mixtures of truncated exponentials (MTE) influence diagrams, which are influence diagrams where probability distributions and utility functions are represented by MTE potentials. Decision variables in MTE influence diagrams must have discrete state spaces; however, many decision problems in practice have continuous decision variables.

Consider the following decision problem from economics (see Figure 1):

> A monopolist faces uncertain demand dependent on favorable ($M = 1$) or unfavorable ($M = 0$) market conditions. Demand is reflected in the price (P) it receives for its output (Q). The monopolist cannot directly observe demand, but rather relies on the results (R) of a market survey to gauge demand, and thus predict price (P). Based on the survey results (R), the monopolist must determine its optimal output (Q).

The monopolist can produce from 0 to 70 units and has the following utility function:

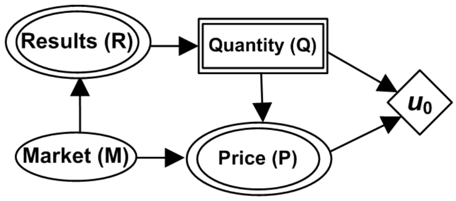

Figure 1: An influence diagram for the monopolist's decision problem.

$u_0(p, q) = (p - 5) \cdot q - 200000$. The probability of an unfavorable market ($M = 0$) is 0.40. The mean of the conditional Gaussian probability density function (PDF) for price (P) is a linear function of quantity (Q) produced, which is based on the market conditions as follows: $P \mid \{Q, M = 0\} \sim N(5000 - 70q, 100^2)$ and $P \mid \{Q, M = 1\} \sim N(10000 - 50q, 1000^2)$. Test results ($R$) are given as a percentage of consumers surveyed who intend to buy the product and are modeled by the following beta PDF's: $R \mid M = 0 \sim Beta(1.3, 2.7)$ and $R \mid M = 1 \sim Beta(2.7, 1.3)$.

Influence diagrams with continuous decision variables and non-Gaussian continuous chance variables are difficult to solve using current methodology. This paper introduces the continuous decision MTE influence diagram (CDMTEID), a model which allows any combination of discrete and continuous chance variables with no restrictions on the type of probability distribution, as well as discrete and/or continuous decision variables. CDMTEID's can also accommodate conditionally deterministic chance variables. An operation derived from the method of convolutions in probability theory is used to eliminate a continuous decision node during the solution phase in a CDMTEID.

The remainder of this paper is organized as follows. Section 2 gives notation and definitions. Section 3 describes operations used to solve CDMTEID's. Section 4 presents the CDMTEID solution to the example problem. Section 5 provides conclusions. This paper is extracted from a longer working paper (Cobb, 2006).

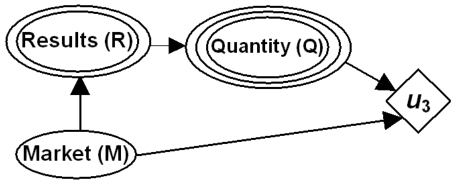

Figure 2: Influence diagram for Example 1.

2 Notation and Definitions

2.1 Notation

Variables will be denoted by capital letters, e.g., A, B, C. Sets of variables will be denoted by boldface capital letters, $\mathbf{Y}$ if all are discrete chance variables, $\mathbf{Z}$ if all are continuous chance variables, $\mathbf{D}$ if all are decision variables, or $\mathbf{X}$ if the components are a mixture of discrete chance, continuous chance, and decision variables. If $\mathbf{X}$ is a set of variables, $\mathbf{x}$ is a configuration of specific states of those variables. The discrete, continuous, or mixed state space of $\mathbf{X}$ is denoted by $\Omega_{\mathbf{X}}$.

MTE probability potentials, discrete probability potentials, and deterministic potentials are denoted by lower-case greek letters, e.g., α, β, γ. MTE utility potentials are denoted by u_i.

In graphical representations, decision variables are represented by rectangular nodes (with a single border for discrete and a double border for continuous), chance variables are represented by ovals (with a single border for discrete, a double border for continuous, and a triple border if the chance variable is deterministic), and utility functions are represented by diamonds.

Example 1. Consider the influence diagram in Figure 2. In this model, M is a discrete variable which can take on values $M = 0$ or $M = 1$. The variables R and Q are continuous chance variables, with Q a deterministic chance variable (Q is conditionally deterministic given a value of R). The utility function (u_3) is a continuous function of Q.

2.2 Mixtures of Truncated Exponentials

A mixture of truncated exponentials (MTE) potential in an influence diagram has the following definition, which is a modification of the definition proposed by Rumí and Salmerón (2005).

MTE potential. Let $\mathbf{X}$ be a mixed n-dimensional variable. Let $\mathbf{Y} = (Y_1, \ldots, Y_d)$, $\mathbf{Z} = (Z_1, \ldots, Z_c)$, and $\mathbf{D} = (D_1, \ldots, D_f)$ be the discrete chance, continuous chance, and decision variable parts of $\mathbf{X}$, respectively, with $c + d + f = n$. A function $\phi : \Omega_{\mathbf{X}} \mapsto \mathcal{R}^+$ is an MTE potential if one of the next three conditions holds:

1. $\mathbf{Y} \cup \mathbf{D} = \emptyset$ and ϕ can be written as

$$\phi(\mathbf{x}) = \phi(\mathbf{z}) = a_0 + \sum_{i=1}^{m} a_i \exp\left\{ \sum_{j=1}^{c} b_i^{(j)} z_j \right\} \quad (1)$$

   for all $\mathbf{x} \in \Omega_{\mathbf{X}}$, where $a_i, i = 0, \ldots, m$ and $b_i^{(j)}$, $i = 1, \ldots, m$, $j = 1, \ldots, c$ are real numbers.

2. $\mathbf{Y} \cup \mathbf{D} = \emptyset$ and there is a partition $\Omega_1, \ldots, \Omega_k$ of $\Omega_{\mathbf{Z}}$ into hypercubes such that ϕ is defined as

$$\phi(\mathbf{x}) = \phi_h(\mathbf{x}) \qquad \text{if } \mathbf{x} \in \Omega_h, \quad (2)$$

   where each $\phi_h, h = 1, ..., k$ can be written in the form of equation (1) (i.e. each ϕ_h is an MTE potential on Ω_h).

3. $\mathbf{Y} \cup \mathbf{D} \neq \emptyset$ and for each fixed value $(\mathbf{y}, \mathbf{d}) \in \Omega_{\mathbf{Y} \cup \mathbf{D}}$, $\phi_{\mathbf{y},\mathbf{d}}(\mathbf{z})$ can be defined as in (2).

In the definition above, k is the number of *pieces*, and m is the number of exponential *terms* in each piece of the MTE potential. In the third case, the potential *fragments* $\phi_{\mathbf{y},\mathbf{d}}(\mathbf{z})$ for all $(\mathbf{y}, \mathbf{d}) \in \Omega_{\mathbf{Y},\mathbf{D}}$ constitute the MTE potential for $\{\mathbf{Y}, \mathbf{Z}, \mathbf{D}\}$. In this paper, all MTE probability and utility potentials are equal to zero in unspecified regions. In CDMTEID's, all probability distributions and utility functions are approximated by MTE potentials.

The definition presented here assumes that decision variables are discrete. This paper presents a method for developing a decision rule for a continuous decision variable as a function of its continuous parents; however, the MTE representation of this method first uses a discrete approximation to the continuous decision variable.

2.3 MTE Probability Densities

Suppose ϕ' is an input MTE potential for $\mathbf{X} = \mathbf{Y} \cup \mathbf{Z} \cup \mathbf{D}$ representing a PDF for $Z \in \mathbf{Z}$ given its parents $\mathbf{X} \setminus \{Z\}$. If we can verify that

$$\int_{\Omega_Z} \phi'(\mathbf{x}, z)\, dz = 1\,, \quad (3)$$

for all $\mathbf{x} \in \Omega_{\mathbf{X} \setminus Z}$, we state that ϕ' is an MTE density for Z. We assume that all input MTE probability potentials in a CDMTEID are normalized prior to the solution phase.

2.4 Deterministic Potential

A *deterministic potential* describes the linear deterministic relationship between a set of variables $\mathbf{Z} = \{Z_1, \ldots, Z_c\}$. A deterministic potential (Cobb and Shenoy, 2005) for $\mathbf{Z}$ is defined as an equation

$$[g(\mathbf{z}) = 0] = \{w_p \cdot [g_p(\mathbf{z}) = 0]\}_{p=1}^{P}\ , \quad (4)$$

where w_p, $p = 1, \ldots, P$ are constants. The equation $g_p(\mathbf{z}) = 0$ defines a linear deterministic relationship, where $g_p(\mathbf{z}) = a_{p1} z_1 + \cdots + a_{pc} z_c + b_p$, and where $a_{p1}, \ldots, a_{pc}$ and b_p are real numbers. The coefficient a_{pi} on $Z_i \in \mathbf{Z}$ is equal to 1 when the deterministic potential is specified as a conditional potential for Z_i given $\mathbf{Z} \setminus Z_i$. The *factors* $w_p \cdot [g_p(\mathbf{z}) = 0]$ are maintained as a decomposed set of weighted equations, with w_p representing the *weights* for all $p = 1, \ldots P$. The weights (w_p) typically result from the marginalization of a discrete variable, as shown in Example 3 in Section 3.2.1.

We term $[g(\mathbf{z}) = 0]$ a potential in the sense that any variable Z_i takes on the value $z_i = (-a_{p1} z_1 - \cdots - a_{p,i-1} z_{i-1} - a_{p,i+1} z_{i+1} - \cdots - a_{pc} z_c - b_p) / a_{pi}$ with probability 1 and any other value with probability 0. As with MTE potentials, multiple fragments of the form in (4) can define a deterministic potential for a set of variables $\mathbf{X}$.

The fragments can be parameterized by either sets of discrete chance and decision variables, or by partitions of hypercubes of continuous chance variables, as shown below.

Example 2. In the influence diagram of Figure 2, Q is conditionally deterministic given R. This relationship is represented by the deterministic potential fragment $[g_0(q,r) = 0]$ or $[q - 46.667r - 30.112 = 0]$ if $0 \leq r \leq 0.8547$, and by the deterministic potential fragment $[g_1(q,r) = 0]$ or $[q - 0r - 70 = 0]$ if $0.8547 < r \leq 1$.

3 Solving CDMTEID's

CDMTEID's are solved by using the fusion algorithm developed by Shenoy (1993) and the operations presented in this section. The fusion algorithm involves deleting variables from the network in a sequence which respects the information constraints (represented by arcs pointing to decision variables in influence diagrams) in the problem. This condition ensures that unobserved chance variables are deleted before decision variables. The fusion method applies to problems where there is only one utility function (or a joint utility function which factors multiplicatively into several utility potentials).

Combination of two MTE potentials is pointwise multiplication. Marginalization of chance variables from MTE potentials is integration over continuous chance variables being removed and summation over discrete chance variables being removed. Details of these two operations can be found in (Cobb and Shenoy, 2004). Other combination and marginalization operations required for CDMTEID's are described below.

3.1 Combination of MTE and Deterministic Potentials

Let $\phi_{\mathbf{y},\mathbf{d}}(\mathbf{z}_1)$ be an MTE probability potential on $\mathbf{X} = \mathbf{Y} \cup \mathbf{Z} \cup \mathbf{D}$ and let $[g_{\mathbf{y},\mathbf{d}}(\mathbf{z}_2) = 0]$ be a deterministic potential on $\mathbf{X} = \mathbf{Y} \cup \mathbf{Z} \cup \mathbf{D}$, where $\mathbf{Z} = \mathbf{Z}_1 \cup \mathbf{Z}_2$. The combination of $\phi_{\mathbf{y},\mathbf{d}}$ and $g_{\mathbf{y},\mathbf{d}}$ is a potential ζ for $\mathbf{X}$ defined as

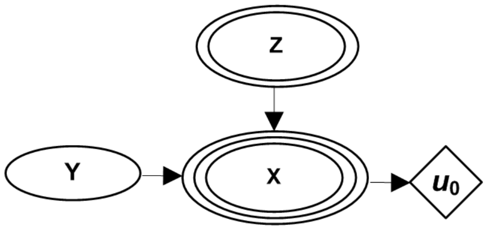

Figure 3: Influence diagram for Example 3.

$$\begin{aligned} \zeta_{\mathbf{y},\mathbf{d}}(\mathbf{z}) &= (\phi_{\mathbf{y},\mathbf{d}} \otimes g_{\mathbf{y},\mathbf{d}})_{\mathbf{y},\mathbf{d}}(\mathbf{z}) \\ &= (\{\phi_{\mathbf{y},\mathbf{d}}(\mathbf{z}_1), [g_{\mathbf{y},\mathbf{d}}(\mathbf{z}_2) = 0]\}) \ , \end{aligned} \tag{5}$$

for all $\mathbf{z} \in \Omega_{\mathbf{Z}}$. According to the LAZY propagation scheme (Madsen and Jensen, 1999) the potentials are not combined, but rather maintained as a decomposed set of potentials during combination.

3.2 Marginalization

3.2.1 Discrete Chance Variables from MTE and Deterministic Potentials

This operation is illustrated by example (see (Cobb, 2006) for a formal definition).

Example 3. Consider the influence diagram in Figure 3, where $P(Y = 0) = 0.5$ and $P(Y = 1) = 0.5$. The variable X is conditionally deterministic given $\{Y, Z\}$, a relationship represented by the deterministic potential fragments $g_0(x, z, Y = 0) = [x - 2z + 1 = 0]$ and $g_1(x, z, Y = 1) = [x - 0.25z - 1 = 0]$. The removal of Y from the combination of g for $\{X, Y, Z\}$ and the PMF for Y results in the following deterministic potential: $\{0.5 \cdot [x - 2z + 1 = 0], 0.5 \cdot [x - 0.25z - 1 = 0]\}$.

The values for the deterministic potential are weighted with the probability values upon removal of the discrete variable.

3.2.2 Continuous Chance Variables from MTE and Deterministic Potentials

Marginalization of a continuous variable Z_i from the combination of an MTE potential and a deterministic potential is substitution of the inverse of the equation(s) in the deterministic potential into the MTE potential. This operation is derived from the method of convolutions in probability theory, as explained in (Cobb and Shenoy, 2005).

Let $\phi_{\mathbf{y},\mathbf{d}}(\mathbf{z}_1)$ be an MTE potential on $\mathbf{X}_1 = \mathbf{Y} \cup \mathbf{Z}_1 \cup \mathbf{D}$ and let $[g_{\mathbf{y},\mathbf{d}}(\mathbf{z}_2) = 0]$ be a deterministic potential on $\mathbf{X}_2 = \mathbf{Y} \cup \mathbf{Z}_2 \cup \mathbf{D}$. Assume $\mathbf{Z} = \mathbf{Z}_1 \cup \mathbf{Z}_2$, $\mathbf{X} = \mathbf{X}_1 \cup \mathbf{X}_2$, and that each $[g_{\mathbf{y},\mathbf{d},p}(\mathbf{z}_2) = 0]$, $p = 1, \ldots, P$, is invertible in $Z_i \in (\mathbf{Z}_1 \cap \mathbf{Z}_2)$. The marginal of $\{\phi_{\mathbf{y},\mathbf{d}}, g_{\mathbf{y},\mathbf{d}}\}$ for a set of variables $\mathbf{X}' = \mathbf{Y} \cup (\mathbf{Z} \setminus Z_i) \cup \mathbf{D} \subseteq \mathbf{X}$ is computed as

$$\begin{aligned}&(\phi_{\mathbf{y},\mathbf{d}} \otimes g_{\mathbf{y},\mathbf{d}})^{\downarrow \mathbf{X}'}_{\mathbf{y},\mathbf{d}}(\mathbf{z}') \\ &= \sum_{p=1}^{P} K_p \cdot w_p \cdot \phi_{\mathbf{y},\mathbf{d}}\left(h_{\mathbf{y},\mathbf{d},p}\left(\mathbf{z}'_2\right), \mathbf{z}'_1\right) \quad ,\end{aligned} \tag{6}$$

for all $\mathbf{x}' \in \Omega_{\mathbf{X}'}$ where $\mathbf{z} = (\mathbf{z}', z_i)$, $\mathbf{z}_1 = (\mathbf{z}'_1, z_i)$, $\mathbf{z}_2 = (\mathbf{z}'_2, z_i)$, and

$$\begin{aligned}&h_{\mathbf{y},\mathbf{d},p}(\mathbf{z}'_2) = (-a_{p1}z_1 - \cdots \\ &-a_{p,i-1}z_{i-1} - a_{p,i+1}z_{i+1} - \cdots - a_{pc}z_c - b_p) / a_{pi} \ .\end{aligned}$$

The constant K_p is $\dfrac{1}{|\ a_{pi}\ |}$ if $\phi_{\mathbf{y},\mathbf{d}}$ is an MTE probability potential and 1 if $\phi_{\mathbf{y},\mathbf{d}}$ is an MTE utility potential, where a_{pi} represents the coefficient on variable Z_i in $g_{\mathbf{y},\mathbf{d},p}(\mathbf{z}_2)$.

3.2.3 Discrete Decision Variables

Marginalization with respect to a discrete decision variable is only defined for MTE utility potentials. Let u be an MTE utility potential for $\mathbf{X} = \mathbf{Y} \cup \mathbf{Z} \cup \mathbf{D}$, where $D \in \mathbf{D}$. The marginal of u for a set of variables $\mathbf{X} - \{D\}$ is an MTE utility potential computed as

$$u^{\downarrow(\mathbf{X}-\{D\})}(\mathbf{y}, \mathbf{z}, \mathbf{d}') = \max_{d \in \Omega_D} u(\mathbf{y}, \mathbf{z}, \mathbf{d}) \tag{7}$$

for all $(\mathbf{y}, \mathbf{z}, \mathbf{d}') \in \Omega_{\mathbf{X}-\{D\}}$ where $\mathbf{d} = (\mathbf{d}', d)$. Marginalization of decision variables results in an MTE potential. For details, see Cobb and Shenoy (2004).

3.2.4 Continuous Decision Variables

Eliminating a continuous decision variable from a CDMTEID is a three-step process:

Step 1: Create a discrete approximation to the continuous decision variable and apply the procedure in Section 3.2.3.

Step 2: Using least squares regression, create a decision rule for the continuous decision variable as a function of its continuous parent(s).

Step 3: Construct a deterministic potential from the decision rule developed in Step 2, convert the continuous decision variable to a deterministic chance variable, and marginalize the variable using the procedure in Section 3.2.2.

This operation will be described for the case where a decision variable has one continuous parent, but multivariate regression can be used to extend the operation for the case where a decision variable has multiple continuous parents. Suppose we want to eliminate a continuous decision variable D with continuous parent Z from the influence diagram using the fusion algorithm. Let u be an MTE utility potential for $\{Z, D\}$. First, a v-point discrete approximation is created for D. Each qualitative state of the discrete approximation corresponds with a real numbered value in the continuous state space of the variable, $\Omega_D = \{d : d_{min} \leq d \leq d_{max}\}$. The real numbered values associated with the qualitative states are determined as $d_t = d_{min} + (t - 0.5) \cdot (d_{max} - d_{min}) / v$ for $t = 1, \ldots, v$. For simplicity, the qualitative states will also be referred to as $d_1, \ldots, d_v$.

Marginalization of D from the utility potential u involves finding the value of d_t that maximizes u for each point in the continuous domain of Z. The following k-piece component MTE potential for Z is derived from the discretized continuous decision variable D using the procedure in Section 3.2.3:

$$u_{max}(z) = \begin{cases} u_1(z) & \text{if } e_0 \leq z < e_1 \\ u_2(z) & \text{if } e_1 \leq z < e_2 \\ \vdots & \\ u_k(z) & \text{if } e_{k-1} \leq z \leq e_k \end{cases}$$

where e_{j-1} and e_j are the lower and upper bounds of the domain of the variable Z in j-

th piece of the MTE utility potential u_{max}, as determined by the procedure in Section 3.2.3. These bounds are defined such that $(e_{i-1}, e_i) \cap (e_{j-1}, e_j) = \emptyset$ for all $i, j = 1, \ldots, k$, $i \neq j$ and $\bigcup_{j=1}^{k} (e_{j-1}, e_j) = (z_{min}, z_{max})$, where the state space of Z is given as $\Omega_Z = \{z : z_{min} \leq z \leq z_{max}\}$. Let $\Psi(z) = \{d^{(1)}, ..., d^{(k)}\}$ be a vector of real-numbered values d_t of the decision variable D that correspond with $u_1, \ldots, u_k$ in u_{max}, i.e. Ψ is the decision rule obtained upon the removal of D. The decision rule is a function where $\Psi(z) = d^{(j)}$ if $e_{j-1} \leq z < e_j$ for all $j = 1, \ldots, k$.

The decision rule Ψ is transformed to a decision rule $\hat{\Psi}$ that is stated as a function of Z. Using least squares regression, we obtain a linear equation of the form $f(z) = b_0 + b_1 z$. The values $\mathbf{b} = [b_0, b_1]$ are calculated as $\mathbf{b} = \left(\Delta^\top \Delta\right)^{-1} \Delta^\top \theta$ where $\Theta = [d^{(1)}, d^{(2)}, \ldots, d^{(k)}]^\top$ and

$$\Delta = \begin{bmatrix} 1 & \frac{e_0+e_1}{2} \\ 1 & \frac{e_1+e_2}{2} \\ \vdots & \vdots \\ 1 & \frac{e_{k-1}+e_k}{2} \end{bmatrix}.$$

The decision rule is then stated as

$$\hat{\Psi}(z) = \begin{cases} d_{min} & \text{if } b_0 + b_1 z < d_{min} \\ b_0 + b_1 z & \\ & \text{if } d_{min} \leq b_0 + b_1 z \leq d_{max} \\ d_{max} & \text{if } b_0 + b_1 z > d_{max} \end{cases},$$

where the first and third regions of $\hat{\Psi}(z)$ are added to ensure D takes on a value within its state space. The decision rule above is used to construct a deterministic potential $[g(z, d) = 0]$. Variable D has been converted to a deterministic chance variable and is eliminated from the model by using the operation described in Section 3.2.2.

The class of MTE potentials is closed under each of the operations required to solve MTE influence diagrams (Cobb and Shenoy, 2004) and the convolution operation used to eliminate continuous decision variables (Cobb and Shenoy, 2005).

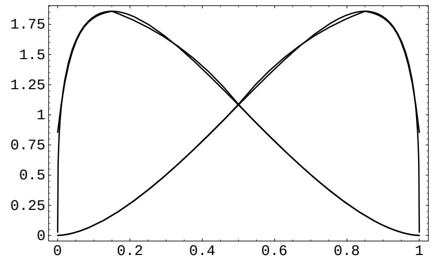

Figure 4: The MTE potential fragments which constitute the MTE potential β for $\{R, M\}$.

4 Example

This section presents the CDMTEID solution to the monopolist's decision problem. Additional numerical details of MTE potentials can be found in (Cobb, 2006).

4.1 Representation

The MTE potential for P given $\{M, Q\}$ is a 2-piece MTE approximation to the normal PDF (Cobb and Shenoy 2006a) with $\mu = 5000 - 70q$ and $\sigma^2 = 100^2$ given $M = 0$, and $\mu = 10000 - 50q$ and $\sigma^2 = 1000^2$ given $M = 1$. The potential γ for $\{P, M, Q\}$ has potential fragments $\gamma(p, q, M = 0)$ and $\gamma(p, q, M = 1)$.

The MTE approximation to the utility function u_0 for $\{P, Q\}$ is denoted by u_1 and is determined using the procedure described in (Cobb and Shenoy, 2004). MTE approximations to the beta PDF's for R given M are constructed using the procedure in (Cobb *et al.*, 2006). These MTE potential fragments—which constitute the MTE potential β for $\{R, M\}$—are displayed graphically in Figure 4 overlayed on the corresponding beta PDF's. The PMF for M is represented by the potential α where $\alpha(0) = P(M = 0) = 0.4$ and $\alpha(1) = P(M = 1) = 0.6$.

Although Q is a continuous decision variable, we will initially use a three-point ($v = 3$) discrete approximation with $q_1 = 11.667$, $q_2 = 35$, and $q_3 = 58.333$ to apply the solution procedure.

4.2 Solution

The CDMTEID solution proceeds by eliminating the variables in the sequence P, M, Q, R.

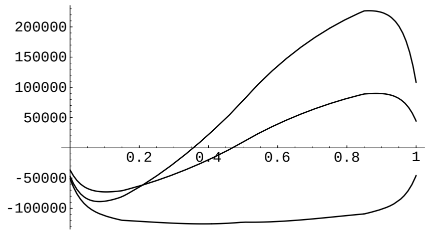

Figure 5: The utility potential fragments which constitute the utility potential u_3.

To remove P, we calculate $u_2 = (u_1 \otimes \gamma)^{\downarrow\{M,Q\}}$. The values for u_2 are calculated as

$$u_2(q, M=0) = \int_{\Omega_P} u_1(p,q) \cdot \gamma(p,q,M=0)\ dp\ ,$$
$$u_2(q, M=1) = \int_{\Omega_P} u_1(p,q) \cdot \gamma(p,q,M=1)\ dp\ .$$

To remove M, we calculate $u_3 = (u_2 \otimes \alpha \otimes \beta)^{\downarrow\{Q,R\}}$. The utility potential fragments $u_3(r, Q=q_1)$, $u_3(r,Q=q_2)$, and $u_3(r,Q=q_3)$ are shown graphically in Figure 5. Removing Q involves first finding $\text{Max}\{u_3(r,Q=q_1), u_3(r,Q=q_2), u_3(r,Q=q_3)\}$ at each point in the domain of R. In this case, we find $u_3(r,Q=q_2) \approx u_3(r,Q=q_3)$ at $R=0.2095$. Next, the decision rule is developed by using this value.

4.2.1 Creating a Decision Rule for a Continuous Decision Variable

Using the procedure in Section 3.2.3, we determine the optimal strategy of producing $Q = q_2 = 35$ if $0 < r < 0.2095$, and $Q = q_3 = 58.333$ if $0.2095 \le r < 1$. Using least squares regression (see Section 3.2.4), the following decision rule is obtained:

$$\hat{\Psi}(r) = \begin{cases} 46.667 \cdot r + 30.112 & \text{if } 0 \le r \le 0.8547 \\ 70 & \text{if } 0.8547 < r \le 1 \end{cases} .$$

To convert Q to a deterministic chance variable, we use $\hat{\Psi}(r)$ to define the deterministic potential g for $\{Q,R\}$ as in Example 2 of Section 2.4 (the domain of R precludes the possibility of producing zero units).

Table 1: Decision strategy for an CDMTEID solution with eight states for Q.

Test Results (R)	Quantity Produced (Q)
$0 \le r \le 0.1655$	Q=39.375
$0.1655 < r \le 0.2895$	Q=48.125
$0.2895 < r \le 0.3975$	Q=56.875
$0.3975 < r \le 1$	Q=65.625

A CDMTEID solution with eight discrete states for the decision variable Q yields the decision rule in Table 1 (some states of Q are not optimal for any values of R). Define $\Theta = [39.375, 48.125, 56.875, 65.625]^\top$ and

$$\Delta = \begin{bmatrix} 1 & \frac{0+0.1655}{2} \\ 1 & \frac{0.1655+0.2895}{2} \\ 1 & \frac{0.2895+0.3975}{2} \\ 1 & \frac{0.3975+1}{2} \end{bmatrix} = \begin{bmatrix} 1 & 0.08275 \\ 1 & 0.2275 \\ 1 & 0.3435 \\ 1 & 0.69875 \end{bmatrix} .$$

The resulting decision rule is based on the linear equation $f(r) = 41.403r + 38.501$ and the deterministic potential $[g(r,q) = 0]$ for $\{Q,R\}$ has fragments $[g_0(r,q) = 0]$ or $[q - 41.403r - 38.501 = 0]$ if $0 \le r \le 0.7608$ and $[g_1(r,q) = 0]$ or $[q - 0r - 70 = 0]$ if $0.7608 < r \le 1$.

4.2.2 Converting the Decision Variable to a Deterministic Chance Variable

To eliminate the decision variable Q, we consider the revised influence diagram in Figure 2 that replaces the decision node with a deterministic probability node. The potentials remaining after Q is converted to a deterministic chance variable are u_3 for $\{Q,R\}$ and g for $\{Q,R\}$. To marginalize a continuous chance variable from the combination of a utility potential and a deterministic potential, we substitute an expression obtained from the deterministic potential for the variable to be removed into the utility potential, as detailed in Section 3.2.2. In this case, a new MTE utility potential is:

$$u_4(r) = \begin{cases} u_3(46.667 \cdot r + 30.112, r) & \text{if } 0 \le r \le 0.8547 \\ u_3(70, r) & \text{if } 0.8547 < r \le 1 \end{cases} .$$

Integrating u_4 over the domain of R gives the final expected utility of 83729. Creating the decision rule specified by the function $\hat{\Psi}(r)$ yields an increase in expected value of 9210 over the decision rule created using an MTE influence diagram with a discrete decision variable. Using the same procedure, the maximum expected utility for the monopolist using a CDMTEID with eight discrete states for Q is 89686, which is 3053 higher than an MTE influence diagram with eight states for Q. Cobb (2006) gives a detailed comparison of the CDMTEID solution to similar MTE and discrete influence diagrams.

5 Conclusions

Continuous decision MTE influence diagrams—which determine a decision rule for a continuous decision variable as a function of its continuous parents—show potential to improve decision-making at a lower computational cost than discrete influence diagrams or MTE influence diagrams. Previous methods for handling continuous decision variables in influence diagrams require either normal distributions for continuous chance variables and/or discrete approximations to continuous decision variables. The continuous decision MTE influence diagram allows a more precise decision rule and the ability to exploit the continuous nature of test results.

Acknowledgments

The author is grateful for the insightful comments of three anonymous reviewers which significantly improved the paper.

References

Cobb, B.R. (2006), "Refining decision rules in influence diagrams," Working Paper, Virginia Military Institute. Available for download at: www.vmi.edu/media/ecbu/cobb/CONTID.pdf.

Cobb, B.R. and P.P. Shenoy (2004), "Hybrid influence diagrams using mixtures of truncated exponentials," in M. Chickering and J. Halpern (eds.), *Uncertainty in Artificial Intelligence*, 20, 85–93.

Cobb, B.R. and P.P. Shenoy (2005), "Hybrid Bayesian networks with linear deterministic variables," in F. Bacchus and T. Jaakkola (eds.), *Uncertainty in Artificial Intelligence*, 21, 136–144.

Cobb, B.R. and P.P. Shenoy (2006a), "Inference in hybrid Bayesian networks using mixtures of truncated exponentials," *International Journal of Approximate Reasoning*, 41(3), 257–286.

Cobb, B.R., Shenoy, P.P. and R. Rumí (2006), "Approximating probability density functions in hybrid Bayesian networks with mixtures of truncated exponentials," *Statistics and Computing*, 16(3), 293–308.

Howard, R.A. and J.E. Matheson (1984), "Influence diagrams," in R.A. Howard and J.E. Matheson (eds.), *Readings on the Principles and Applications of Decision Analysis*, 2, 719–762, Strategic Decisions Group, Menlo Park, CA.

Madsen, A.L. and F. Jensen (2005), "Solving linear-quadratic conditional Gaussian influence diagrams," *International Journal of Approximate Reasoning*, 38(3), 263–282.

Madsen, A.L. and F.V. Jensen (1999), "Lazy evaluation of symmetric Bayesian decision problems," in K.B. Laskey and H. Prade (eds.), *Uncertainty in Artificial Intelligence*, 15, 382–390.

Olmsted, S.M. (1983), "On representing and solving decision problems," Ph.D. Thesis, Department of Engineering–Economic Systems, Stanford University, Stanford, CA.

Poland, W.B. and R.D. Shachter (1993), "Mixtures of Gaussians and minimum relative entropy techniques for modeling continuous uncertainties," in D. Heckerman and E.H. Mamdani (eds.), *Uncertainty in Artificial Intelligence*, 9, 183–190.

Rumí, R. and A. Salmerón (2005), "Penniless propagation with mixtures of truncated exponentials," in L. Godo (ed.), *Symbolic and Quantitative Approaches to Reasoning under Uncertainty*, Lecture Notes in Artificial Intelligence, 3571, 39–50.

Shachter, R.D. (1986), "Evaluating influence diagrams," *Operations Research*, 34(6), 871–882.

Shachter, R.D. and C.R. Kenley (1989), "Gaussian influence diagrams," *Management Science*, 35(5), 527–550.

Shenoy, P.P. (1993), "A new method for representing and solving Bayesian decision problems," in D.J. Hand (ed.), *Artificial Intelligence Frontiers in Statistics: AI and Statistics III*, Chapman and Hall, London, 119–138.