Measuring Efficiency in Influence Diagram Models

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

Abstract

A measure of efficiency for influence diagram models with continuous decision variables that considers both the accuracy and complexity of the representation and solution technique is presented. Accuracy is determined by calculating the mean squared error between influence diagram decision rules and an analytical solution. Complexity is assessed by calculating the size of the functions in the numerical representation at each stage of the solution to reflect both storage requirements and potential computational complexity of downstream mathematical operations. The resulting efficiency score considers the preferences of an individual decision maker for accuracy and complexity. Three influence diagram models proposed for use with continuous decision variables are compared using the efficiency measurement. The best model for a given problem may vary based on a decision maker's willingness to substitute accuracy and complexity.

1 Introduction

The *influence diagram* (ID) was introduced by Howard and Matheson (1984) as a graphical and numerical representation for a decision problem under uncertainty. The ID model is composed of a directed acyclic graph that shows the relationships among chance and decision variables in the problem, as well as a set of conditional probability distributions for chance variables and a joint utility function. In addition to providing a tractable, intuitive view that facilitates communication about the decision problem, the ID solution provides an optimal strategy and maximum expected utility.

Since their invention, most subsequent improvements to exact solution procedures for solving IDs (see, e.g., citations in Cobb (2008)) assume that all chance and decision variables in the model are discrete. Cobb and Shenoy (2008) introduce mixtures of truncated exponentials (MTE) influence diagrams (MTEIDs), which are influence diagrams where probability density functions (pdfs) and utility functions are represented by MTE potentials (Moral et al. 2001). Each piece of an MTE function is composed of a sum of exponential terms where the exponent contains a linear function of the independent variables.

Discrete IDs and MTEIDs can only accomodate continuous decision variables if their state spaces are limited to a countable number of discrete values. In contrast, Shachter and Kenley (1989) introduce Gaussian IDs, where all continuous chance variables are normally distributed, all decision variables are continuous, and utility functions are quadratic. The mixture-of-Gaussians ID (Poland and Shachter 1993, Madsen and Jensen 2005) requires continuous chance variables to be modeled as mixtures of normal distributions and allows continuous decision variables.

Cobb (2007) introduces an ID model, the continuous decision MTE influence diagram (CDMTEID), which allows continuous decision variables with one continuous parent and continuous chance variables having any pdf.

The aim of this paper is to define a measurement that can be used to compare the efficiency of ID models with continuous decision variables.

This measurement captures both the accuracy and complexity of the ID solution, then weights these in accordance with the preferences of an individual decision maker. Using the efficiency metric, the competency in performance of ID solutions in three models—discrete IDs, MTEIDs, and CDMTEIDs—are compared. Using the measurements proposed in this paper, a decision maker can attempt to answer the question, "Is a more accurate model worth the additional computational complexity?"

The remainder of the paper is organized as follows. Section 2 describes notation and definitions. Section 3 provides details of the accuracy, complexity, and efficiency measurements. Section 4 provides an example of calculating accuracy and complexity. Section 5 compares efficiency results for the three ID models under consideration. Section 6 concludes the paper. This contribution is extracted from a longer working paper (Cobb 2008).

2 Notation and Definitions

2.1 Graphical Representation

Chance and decision variables in IDs are depicted as ovals and rectangles, respectively. Utility nodes appear as diamonds. An arrow pointing to a chance node indicates that the distribution for this chance node is conditional on the variable at the head of the arrow. An arrow pointing to a decision node means that the value of the variable at the head of the arrow will be known at the time the decision is made.

Example 1. Fig. 1 shows an ID model for a capacity planning and pricing decision problem under uncertainty (Göx 2002). Capacity (K) and price (P) are decision variables, the random demand "shock" (Z) is a chance variable, and u_0 is the joint utility function.

2.2 Numerical Representation

In this paper, we assume all decision and chance variables take values in bounded, continuous (non-countable) state spaces. All variables are denoted by capital letters in plain text, e.g., A, B, C, with sets of variables denoted by capital letters in boldface, e.g., $\mathbf{X}$. If A and $\mathbf{X}$ are

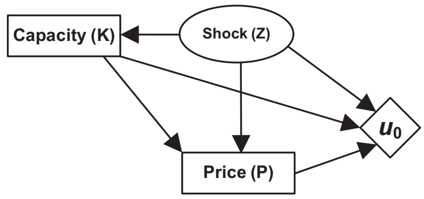

Figure 1: An influence diagram model.

one- and multi-dimensional variables, respectively, then a and $\mathbf{x}$ represent specific values of those variables.

The finite, continuous state space of $\mathbf{X}$ is denoted by $\Omega_{\mathbf{X}}$. The state space for a single variable B is defined as $\Omega_B = \{b : b_{min} \leq b \leq b_{max}\}$. At certain points in the ID representation and solution, a variable B's continuous state space, Ω_B, may be replaced by a discrete approximation, $\Omega_B^{(d)}$.

A probability potential, ϕ, for a set of variables $\mathbf{X}$ is a function $\phi : \Omega_{\mathbf{X}} \to [0,1]$. A utility potential, u, for a set of variables $\mathbf{X}$ is a function $\Omega_{\mathbf{X}} \to \mathcal{R}$.

All piecewise functions are implicitly understood to equal zero in undefined regions.

Example 2. In the ID shown in Fig. 1, product demand is determined as $Q(p,z) = 12 - p + z$. Assume $Z \sim N(0,1)$ and that the firm's utility (profit) function is

$$u_0(k,p,z) = \begin{cases} (p-1)\cdot(12-p+z)-k & \text{if } (12-p+z) \leq k \\ (p-1)\cdot k - k & \text{if } (12-p+z) > k\ . \end{cases} \quad (1)$$

The state spaces of the variables are: $\Omega_K = \{k : 0 \leq k \leq 14\}$; $\Omega_P = \{p : 1 \leq p \leq 9\}$; and $\Omega_Z = \{z : -3 \leq z \leq 3\}$.

2.3 Combination

Combination of potentials is pointwise multiplication. Let ψ_1 and ψ_2 be probability and/or utility potentials for $\mathbf{X}_1$ and $\mathbf{X}_2$. The combination of ψ_1 and ψ_2 is a new potential for $\mathbf{X} = \mathbf{X}_1 \cup \mathbf{X}_2$ defined as

$$\psi(\mathbf{x}) = (\psi_1 \otimes \psi_2)(\mathbf{x}) = \psi_1(\mathbf{x}^{\downarrow\Omega_{\mathbf{X}_1}}) \cdot \psi_2(\mathbf{x}^{\downarrow\Omega_{\mathbf{X}_2}})$$

for all $\mathbf{x} \in \Omega_{\mathbf{X}}$. If either ψ_1 or ψ_2 is a utility potential, the result of the combination will be a utility potential; otherwise, the result is a probability potential.

2.4 Marginalization of Chance Variables

Marginalization of chance variables corresponds to integrating over the chance variable to be removed. Let ψ be a potential for $\mathbf{X} = \mathbf{X}' \cup Z$, where Z is a chance variable. The marginal of ψ for $\mathbf{X}'$ is a potential computed as

$$\psi^{\downarrow \mathbf{X}'}(\mathbf{x}') = \int_{\Omega_Z} \psi(\mathbf{x}) \, dz$$

for all $\mathbf{x}' \in \Omega_{\mathbf{X}'}$, where $\mathbf{x} = (\mathbf{x}', z)$.

2.5 Marginalization of Decision Variables

Marginalization with respect to a decision variable is only defined for utility potentials. Let u be a utility potential for $\mathbf{X} \cup D$, where D is a decision variable. The marginal of u for $\mathbf{X}$ is a utility potential computed as

$$u^{\downarrow \mathbf{X}}(\mathbf{x}) = \max_{d \in \Omega_D} u(\mathbf{x}, d) \qquad (2)$$

for all $\mathbf{x} \in \Omega_{\mathbf{X}}$. The mechanics of performing the maximization operation in Eq. (2) vary with each ID model compared in this paper.

2.6 Fusion Algorithm

IDs are solved in this paper by applying the fusion algorithm of Shenoy (1993). This algorithm involves deleting the variables in an elimination sequence that respects the information constraints in the problem. The sequence is chosen so that decision variables are eliminated before chance or decision variables that are immediate predecessors. When a variable is to be deleted from the model, all probability and/or utility potentials containing this variable in their domains are combined, then the variable is marginalized from the result.

3 Measuring Accuracy, Complexity, and Efficiency

3.1 Analytical Solution

In the problem from Examples 1 and 2, the firm knows the true value, $Z = z$, of the demand shock Z when it chooses capacity, so it would logically set $K = 12 - P + z$. Göx (2002) finds an analytical solution to the problem with optimal values for P and K of

$$p^* = \Theta_1^*(z) = 2 + \frac{10 + z}{2} \qquad (3)$$

and

$$k^* = \Theta_2^*(z) = \frac{10 + z}{2} \,. \qquad (4)$$

3.2 Accuracy

Since analytical decision rules are available for K and P, the mean squared error (MSE) (Winkler and Hays 1970) can be used as a measure of the difference between the analytical and ID decision rules. For instance, define Θ_2 as a decision rule for K as a function of Z determined using an ID method. The MSE of this function is calculated as

$$\begin{aligned} MSE &= E\left[\left(\Theta_2(z) - \Theta_2^*(z) \right)^2 \right] \\ &= \int_{\Omega_Z} \phi(z) \cdot \left(\Theta_2(z) - \Theta_2^*(z) \right)^2 \, dz \,, \end{aligned} \qquad (5)$$

where ϕ is the pdf shown in Figure 2. The MSE between the decision rule Θ_1 developed in the ID models for P as a function of K and Z and the analytical decision rule Θ_1^* is similarly calculated as

$$\begin{aligned} MSE &= E\left[\left(\Theta_1(\Theta_2(z), z) - \Theta_1^*(z) \right)^2 \right] \\ &= \int_{\Omega_Z} \phi(z) \cdot \left(\Theta_1(\Theta_2(z), z) - \Theta_1^*(z) \right)^2 \, dz \,. \end{aligned} \qquad (6)$$

The accuracy of a given ID model for this example will be denoted by $\mathcal{A}$ and will be defined as the sum of the MSEs calculated using Eqs. (5) and (6).

3.3 Complexity

The ID models in this paper are solved using Mathematica software (`www.wolfram.com`). This package provides a function called `LeafCount` that gives the total "size" of an expression defined using the `Piecewise` representation, based on applying the `FullForm` function (Wolfram 2003). `LeafCount` (denoted by $\mathcal{L}$) will be used to measure the complexity of potentials in the ID solution procedures.

Example 3. Consider the expression

$$f(z) = \begin{cases} -84.0 + 81.1 \exp\{0.0119(z+3)\} \\ \quad \text{if} \ -3 \leq z \leq 3 \ . \end{cases} \tag{7}$$

Using the `Piecewise` environment, this function is defined in Mathematica as

$f[z\_] :=$
`Piecewise`$\{\{-84.0$
$+81.1 \exp\{0.0119(z+3)\}, \ -3 \leq z \leq 3\}\}$.

Applying the `FullForm` function in Mathematica to this expression yields

`Piecewise`[`List`[`List`[`Plus`[-84, `Times`[81.1, `Power`[E, `Times`[0.0119, `Plus`[3, z]]]]], `LessEqual`[-3, z, 3]]], 0] .

Each word, number, or variable in the `FullForm` expression increases the `LeafCount` of the expression by one. In this case, $\mathcal{L}\{f\} = 19$.

`LeafCount` captures the total size of all pieces of an MTE approximation or decision rule, including both the parameters of the function and the inequalities required to define the domains of each piece.

The complexity of the various ID methods will be determined by measuring the `LeafCount` of the potentials stored in memory after each combination or marginalization operation (or sub-operation thereof) performed in the solution technique. This measure of complexity is used because the size of the potentials at each step of the solution technique affects both the storage required and the subsequent number of calculations needed to solve the ID model.

Suppose the ID solution procedure for a particular problem requires n operations (or sub-operations) and denote the probability potentials, utility potentials, and decision rules remaining after operation i as ϕ_{ij}, $j = 1, \ldots, m_i$. The total complexity, $\mathcal{C}$, of the ID solution is the sum of the complexity measurements, $\mathcal{C}_i$, $i = 0, \ldots, n$, taken after each operation in the procedure. The total complexity of the solution for a particular ID model is determined as

$$\mathcal{C} = \sum_{i=0}^{n} \mathcal{C}_i = \sum_{i=0}^{n} \sum_{j=1}^{m_i} \mathcal{L}\{\phi_{ij}\} \ .$$

The value $\mathcal{C}_0$ represents the complexity of the potentials in the initial ID model.

3.4 Normalized Measurements

Since the MSE accuracy measurement, $\mathcal{A}$, and the complexity measurement determined by compiling `LeafCount` formulas, $\mathcal{C}$, are stated on different numerical scales, it is advantageous to normalize these two measurements onto a common scale to determine the trade-off between accuracy and complexity.

Select any two positive real numbers, $\mathcal{N}_{min}$ and $\mathcal{N}_{max}$. Throughout the remainder of the paper, we assume $\mathcal{N}_{min} = 1$ and $\mathcal{N}_{max} = 2$. When comparing the accuracy and complexity of ID solutions for multiple models, we denote by $\underline{\mathcal{A}}$ and $\underline{\mathcal{C}}$ the measurements for the least accurate and most complex models, respectively, measured for all models under consideration. Likewise, we denote by $\overline{\mathcal{A}}$ and $\overline{\mathcal{C}}$ the measurements for the most accurate and least complex models, respectively. In the case of both accuracy and complexity, note that smaller measurements are desirable. The normalized accuracy measurement for a given model is determined as

$$\hat{\mathcal{A}} = \mathcal{N}_{min} + \frac{(\mathcal{N}_{max} - \mathcal{N}_{min}) \cdot (\mathcal{A} - \underline{\mathcal{A}})}{\overline{\mathcal{A}} - \underline{\mathcal{A}}} \ . \tag{8}$$

Similarly, the normalized complexity measurement for a given model is calculated as

$$\hat{\mathcal{C}} = \mathcal{N}_{min} + \frac{(\mathcal{N}_{max} - \mathcal{N}_{min}) \cdot (\mathcal{C} - \underline{\mathcal{C}})}{\overline{\mathcal{C}} - \underline{\mathcal{C}}} \ . \tag{9}$$

3.5 Efficiency

Once the normalized accuracy and complexity measurements are determined, the *efficiency* of the model is calculated as

$$\mathcal{E} = \hat{\mathcal{A}}^{\alpha} \cdot \hat{\mathcal{C}}^{1-\alpha} \,, \tag{10}$$

The exponent α is a parameter assigned by the decision maker that conveys an individual preference for solutions that are either more accurate or less complex. If $\alpha > 0.5$, the decision maker values accuracy over complexity, and vice versa. Two properties of the functional form in Eq. (10) that make the expression a useful model for production and consumer utility in economics (see, e.g., Baye (2006)) also make it valuable for measuring the efficiency of ID solutions:

1. If two ID models have equivalent accuracy, the model with a better complexity score will have greater efficiency, and vice versa.
2. There is a diminishing marginal rate of substitution between accuracy and complexity.

4 Example

This section describes the calculation of the accuracy and efficiency measurements in the context of the CDMTEID solution to the problem from Examples 1 and 2. Details of calculations for the discrete ID and MTEID can be found in (Cobb 2008).

4.1 Representation

The MTE potential ϕ with $\mu = 0$ and $\sigma^2 = 1$ that approximates the normal distribution (as defined by Cobb and Shenoy (2006)) for the random demand shock (Z) is shown in Fig. 2, overlaid on the actual $N(0, 1)$ distribution.

The decision variable K is limited to discrete outcomes. A v-point discrete approximation to a continuous decision variable K with $\Omega_K = \{k : k_{min} \leq k \leq k_{max}\}$ has values $k_t = k_{min} + (t - 0.5) \cdot (k_{max} - k_{min})/v$ for $t = 1, \ldots, v$. Thus, the discrete state space is defined as $\Omega_K^{(k)} = \{k_1, k_2, \ldots, k_v\}$. To illustrate this example, we assume $v = 6$.

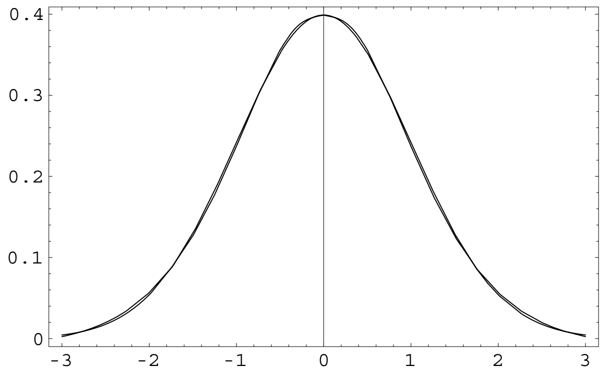

Figure 2: MTE probability density.

The function $f_1(p) = p$ on the interval $[1, 9]$ is modeled by the MTE potential $u_P(p) = -107.056144 + 108.102960 \exp\{0.0089234(p-1)\}$ for all $p \in \Omega_P$. Note that $u_P(1) = 1.047$, $u_P(5) = 4.975$, and $u_P(9) = 9.046$, so the MTE approximation fits $f_1(p)$ reasonably well. More accurate approximations can be obtained by dividing the state space of P and defining separate approximations over each region, at the expense of increasing the representation's complexity measurement (as defined in Section 3.3). The function $f_2(z) = z$ on $[-3, 3]$ is modeled with a similar approximation u_Z. With K assigned v discrete values, the MTE utility function is defined as

$$u_1(k_t, p, z) = \begin{cases} (u_P(p) - 1) \cdot \\ (12 - u_P(p) + u_Z(z)) - k_t \\ \qquad \text{if } (12 - p + z) \leq k_t \\ (u_P(p) - 1) \cdot k_t - k_t \\ \qquad \text{if } (12 - p + z) > k_t \,, \end{cases} \tag{11}$$

for $t = 1, \ldots, v$. For instance, with $v = 6$ and $K = k_3 = 5.83$, the MTE utility function is defined as

$$u_1(5.83, p, z) = \begin{cases} -3789.32 + 15328.9 \exp\{0.00892338p\} \\ -11479.5 \exp\{0.0178468p\} \\ +9002.5 \exp\{0.00892338p + 0.0118978z\} \\ -9079.3 \exp\{0.0118978z\} \\ \qquad \text{if } p - z \geq 6.17 \\ -636.2 + 625.0 \exp\{0.00892338p\} \\ \qquad \text{if } p - z < 6.17 \,. \end{cases}$$

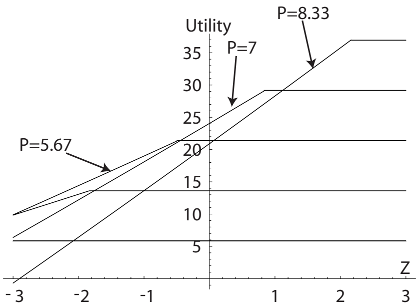

Figure 3: The utility functions $u_1(k_3, p_u, z)$ for $u = 1, \ldots, 6$.

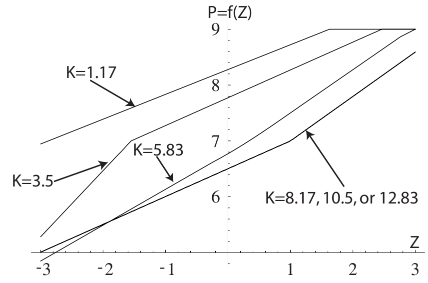

Figure 4: The piecewise linear decision rules $\Theta_{1,t}$ corresponding to values k_t, $t = 1, \ldots, 6$.

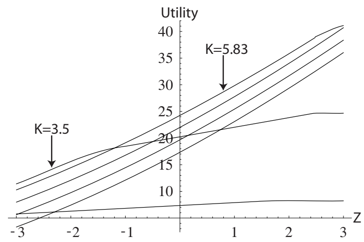

Figure 5: The utility functions $u_2(k_t, z)$.

The CDMTEID solution has initial complexity $\mathcal{C}_0 = \mathcal{L}\{\phi\} + \mathcal{L}\{u_1\} = 517$.

4.2 Solution

The elimination sequence employed in the fusion algorithm is P, K, Z.

Price (P) is a continuous decision variable; however, the first step in marginalizing this variable is accomplished by using the discrete approximation $\Omega_P^{(d)}$. The values p_u, $u = 1, \ldots, 6$ are inserted in the utility potential u_1 to form the utility functions $u_1(k_t, p_1, z)$, ..., $u_1(k_t, p_6, z)$ for $t = 1, \ldots, 6$. After this step, both these new potentials and the existing MTE utility function u_1 remain, so the complexity is

$$\mathcal{C}_1 = \mathcal{C}_0 + \sum_{t=1}^{6}\sum_{u=1}^{6} \mathcal{L}\{u_1(k_t, p_u, z)\} = 1313 \ .$$

The second step in removing P is to create a piecewise linear decision rule for P as a function of Z for each value k_t, $t = 1, \ldots, 6$. For $K = k_3 = 5.83$, the utility functions $u_1(k_3, p_u, z)$ for $u = 1, \ldots, 6$ are shown in Fig. 3 and we can conclude that $P = 5.67$ is optimal over $[-3, -0.45)$, $P = 7$ is optimal over $[-0.45, 1.15)$, and $P = 8.33$ is optimal over $[1.15, 3]$. These values are used to create the piecewise linear decision rule

$$P(z) = \Theta_{1,3}(z) =$$

$$\begin{cases} 6.775100 + 0.642570z & \text{if } -3 \le z < -0.35 \\ 6.729469 + 0.772947z & \text{if } 0.35 \le z < 2.9375 \\ 9 & \text{if } 2.9375 \le z \le 3 \ . \end{cases}$$

Similar decision rules $\Theta_{1,1}$, ..., $\Theta_{1,6}$ are determined corresponding to values k_t, $t = 1, \ldots, 6$ (see Fig. 4). When combined, these functions form the decision rule Θ_1 with complexity $\mathcal{L}\{\Theta_1\} = 259$. Since ϕ and u_1 also remain after this step, the complexity of the model is now $\mathcal{C}_2 = 517 + 259 = 776$.

The last step in removing P from the model is to substitute the values of the decision rule Θ_1 into the utility function u_1 to form the utility functions, $u_2(k_t, z) = u_1(k_t, \Theta_{1,t}(z), z)$, for $t = 1, \ldots, 6$. With ϕ and u_2 as the remaining potentials in the network, the complexity stands at

$$\mathcal{C}_3 = \mathcal{L}\{\phi\} + \sum_{t=1}^{6} \mathcal{L}\{u_2(k_t, z)\} = 61 + 530 = 591 \ .$$

A plot of the functions $u_2(k_1, z)$, ..., $u_2(k_6, z)$ (Fig. 5) shows that $u_2(3.5, z) \approx u_2(5.83, z)$ at $Z = -1.25$. The resulting decision rule Θ_2 specifies that $K = 3.5$ if $-3 \le z < -1.25$ and $K = 5.83$ if $-1.25 \le z \le 3$. After creating this decision rule, the complexity of the model is

$$\mathcal{C}_4 = \mathcal{L}\{\phi\} + \mathcal{L}\{u_2\} + \mathcal{L}\{\Theta_2\} = 610 \ .$$

To complete the marginalization of K, we create a new utility function $u_3(z) = u_2\left(\Theta_2(z), z\right)$.

The complexity after this operation is $\mathcal{C}_5 = 278$, which captures the `LeafCount` of the potentials ϕ and u_3.

To remove Z, the potentials ϕ and u_3 are combined, with the resulting complexity $\mathcal{C}_6 = \mathcal{L}\{(\phi \otimes u_3)\} = 569$. Integrating the result over the state space of Z completes the solution. The total complexity of the ID model is

$$\mathcal{C} = \sum_{i=0}^{6} \mathcal{C}_i = 4654 \ .$$

The MSE of the CDMTEID solution is calculated according to Eqs. (5) and (6) as $\mathcal{A} = 0.7760$.

5 Results

This section discusses the effects on model efficiency of changing the number of states in the discrete approximations to continuous decision variables used in each of the methods

In each of the three ID methods illustrated, the state space of continuous variables is either permanently or temporarily discretized. To investigate the efficiency of models with a varying number of pieces in the discrete approximation, we consider the three ID models with approximations of six through twelve pieces. Thus, the best and worst solutions in terms of accuracy and complexity are chosen from among 21 models when calculating the values of $\underline{\mathcal{A}}$, $\overline{\mathcal{A}}$, $\underline{\mathcal{C}}$, and $\overline{\mathcal{C}}$.

Figs. 6 and 7 show efficiency scores for accuracy parameters of $\alpha = 0.1$ and $\alpha = 0.9$. When accuracy is a low priority, the efficiency of the models decreases with additional discrete pieces in the approximations as computational complexity overburdens the solution. In this case, both the CDMTEID and MTEID models provide comparable efficiency. When accuracy is a high priority, the efficiency of the models generally increases with additional pieces in the discrete approximations and the CDMTEID provides the best efficiency. In some cases, the placement of the mid-points of the discrete bins within the state space of the continuous decision variable adversely affects accuracy, which explains the low efficiency of the solutions with eight-piece approximations.

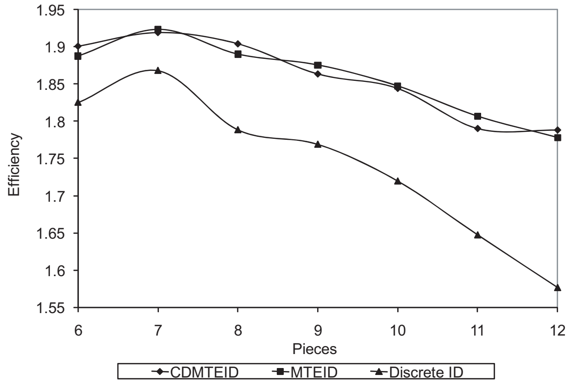

Figure 6: Efficiency scores with $\alpha = 0.1$.

Fig. 8 displays the efficiency scores for the ID solutions over the entire range of possible accuracy values. To make the graph simpler to comprehend, only the efficiency scores for models that gave the optimal efficiency over some range of the accuracy parameter α are shown. The un-normalized accuracy (MSE) and complexity values for these models are also displayed on the chart. For very low values of α, the MTEID solution with six discrete states is optimal. However, as the decision maker's preference for accuracy increases somewhat, a model with seven discrete states provides the best compromise between accuracy and complexity. Once desired accuracy increases beyond $\alpha \approx 0.25$, the CDMTEID model with an increasing number of discrete states in the temporary approximation

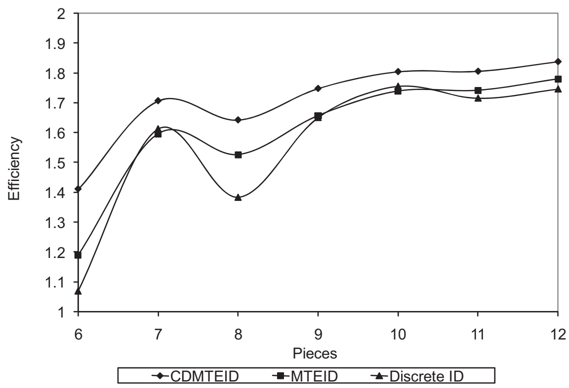

Figure 7: Efficiency scores with $\alpha = 0.9$.

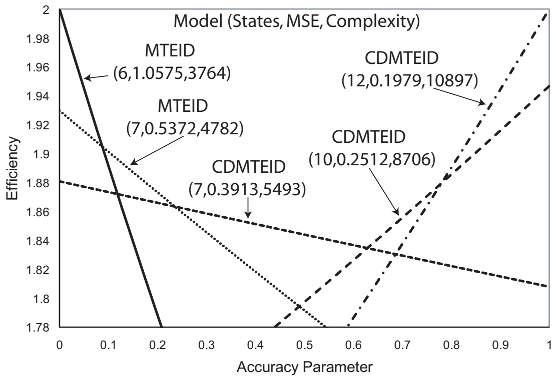

Figure 8: Efficiency values for ID solutions.

becomes optimal.

6 Conclusions

This paper has defined a measure of efficiency for ID models with continuous decision variables that considers both the accuracy and complexity of the representation and solution. Accuracy is measured by calculating the mean squared error between the final decision rule determined using the ID and corresponding analytical decision rules. Complexity is determined by the size of the potentials in the initial ID representation, and after each subsequent operation involved in solving the ID. The efficiency measurement combining accuracy and complexity is able to consider the preferences of an individual decision maker for both accuracy and complexity.

The details of additional comparisons and discussion of further results can be found in (Cobb 2008).

Acknowledgments

Support from the Spanish Ministry of Science and Innovation through project TIN2007–67418–C03–02 and from a Virginia Military Institute grant-in-aid is gratefully acknowledged.

References

Baye, M.R. 2006. *Managerial Economics and Business Strategy.* McGraw–Hill/Irwin, New York.

Cobb, B.R., 2007. Influence diagrams with continuous decision variables and non-Gaussian uncertainties. *Decision Anal.* **4**(3) 136–155.

Cobb, B.R., 2008. Efficiency of Influence Diagram Models with Continuous Decision Variables. Working Paper, Virginia Military Institute, Lexington, VA. Available for download at: `http://new.vmi.edu/media/ecbu/cobb/WP08b.pdf`

Cobb, B.R., P.P. Shenoy. 2006. Inference in hybrid Bayesian networks using mixtures of truncated exponentials. *Internat. J. Approx. Reason.* **41**(3) 257–286.

Cobb, B.R., P.P. Shenoy. 2008. Decision making with hybrid influence diagrams using mixtures of truncated exponentials. *Eur. J. Oper. Res.* **186**(1) 261–275.

Göx, R.F. 2002. Capacity planning and pricing under uncertainty. *J. Management Accounting Res.* **14**(1) 59–78.

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

Madsen, A.L., F. Jensen. 2005. Solving linear-quadratic conditional Gaussian influence diagrams. *Internat. J. Approx. Reason.* **38**(3) 263–282.

Moral, S., R. Rumí, A. Salmerón. 2001. Mixtures of truncated exponentials in hybrid Bayesian networks. P. Besnard, S. Benferhart, eds. *Symbolic and Quantitative Approaches to Reasoning under Uncertainty: Lecture Notes in Artificial Intelligence*, Vol. 2143. Springer-Verlag, Heidelberg, 156–167.

Poland, W.B., R.D. Shachter. 1993. Mixtures of Gaussians and minimum relative entropy techniques for modeling continuous uncertainties. D. Heckerman, E.H. Mamdani, eds. *Uncertainty in Artificial Intelligence: Proc. Ninth Conf.*, Morgan Kaufmann, San Francisco, CA, 183–190.

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

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

Winkler, R.L., W.L. Hays. 1970. *Statistics: Probability, Inference, and Decisions.* Holt, Rinehart, and Winston, New York.

Wolfram, S. 2003. *The Mathematica Book, 5th ed.* Wolfram Media, Champaign, IL.