

An Influence Diagram Model for Detecting Credit Card Fraud

Barry R. Cobb
Virginia Military Institute
cobbbbr@vmi.edu

Abstract

A hybrid influence diagram is a compact graphical and numerical representation of a decision problem under uncertainty that includes both discrete and continuous chance variables. These models can be used by businesses to detect online credit card transactions that may be fraudulent. By creating decision rules based on merchandise value and additional address and product characteristics, the influence diagram model can be used to develop policies that help businesses decide when to investigate an order's legitimacy. The influence diagram establishes guidelines that minimize the sum of the costs of lost merchandise and order investigation.

1 Introduction

A major credit card issuer—Visa—encourages businesses to prevent “card-not-present fraud” by developing “...in-house fraud detection programs, such as guidelines for staff on how to spot and report suspected fraudulent transactions” (“Credit card fraud,” 2008) and lists the following common characteristics of falsified credit card orders: 1) first-time orders, 2) larger than normal orders, 3) orders consisting of several of the same item, 4) orders shipped rush or overnight, and 5) orders shipped to a foreign address.

Fraud detection methods have been implemented using a number of quantitative techniques in the fields of data mining, statistics, and artificial intelligence. Cobb (2010) provides a survey of several of these methods. A large portion of the previous research that adapts quantitative techniques to credit card fraud detection has focused on these problems from the perspective of banks and firms that issue credit cards. However, businesses that accept credit cards for purchases—particularly in online transactions—can also benefit from the application of such models.

Some credit card fraud can be prevented prior to the customer completing a purchase. For instance, when a credit card approval for the

purchase amount is requested, an order being submitted on a card that has been reported as stolen is denied. To further prevent fraud, a business can collect additional information, such as the three-digit card verification value on the back of the card. This prevents someone who has obtained a stolen credit card number, but not the actual card, from completing a fraudulent transaction. However, if a card has not been discovered as stolen, the fraud will not be detected at this point.

This paper proposes an influence diagram (ID) framework that can be employed to develop processes by which businesses can select transactions for further investigation as potentially falsified. The goal of the model is to appropriately balance the cost of shipping merchandise that will ultimately not be paid for because of a fraud-related chargeback, versus the cost of utilizing employee time and system resources to confirm and investigate potentially fraudulent orders. An ID is a probabilistic model that is a simultaneous graphical and numerical representation of a decision problem under uncertainty (Howard and Matheson, 1984). Recent innovations in IDs permit models with non-Gaussian continuous chance variables and discrete decision variables (Cobb and Shenoy, 2008). In this paper, a model that allows a business to develop an optimal decision rule for

whether or not to investigate a transaction for fraud based on the observation of both discrete and continuous variables is suggested.

The remainder of this paper is organized as follows. Section 2 provides notation and definitions. Section 3 describes the ID model. Section 4 illustrates the solution of the ID model for an example fraud detection problem. Section 5 summarizes the paper. This paper is derived from a longer working paper on this topic (Cobb, 2010).

2 Notation and Definitions

This section introduces notation and definitions used throughout the paper.

2.1 Notation

Variables are denoted by capital letters in plain text, e.g., A, B, C . Sets of variables are denoted by capital letters in boldface, e.g., $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$. If A and \mathbf{X} are one- and multi-dimensional variables, respectively, then a and \mathbf{x} represent specific values of those variables. The state space of \mathbf{X} is denoted by $\Omega_{\mathbf{X}}$.

A probability potential, ϕ , for \mathbf{X} is a function $\phi : \Omega_{\mathbf{X}} \rightarrow [0, 1]$. If A is discrete, the more intuitive notation $P(A)$ may be used to represent a discrete probability distribution. A utility potential, u , for a set of variables \mathbf{X} is a function $u : \Omega_{\mathbf{X}} \rightarrow \mathcal{R}$.

2.2 Mixtures of Truncated Exponentials

One difficulty associated with including continuous chance variables in IDs is that mathematical operations, such as integration, on probability density functions are difficult to perform in closed form. For the case where all chance variables are normally distributed and discrete variables do not have continuous parents, the technique of Madsen and Jensen (2005) can be applied to solve the ID.

For problems with continuous variables that are not normally distributed, the state spaces must be discretized to permit an ID solution or a mixture-of-Gaussians ID model (Poland and Shachter, 1993) can be used. Another approach is to approximate probability density functions

in the ID with mixtures of truncated exponentials (MTE) potentials, which are defined as follows.

Definition 1. (MTE Potential (Moral et al., 2001)). Let S be a continuous chance variable. Given a partition $\Omega_1, \dots, \Omega_n$ that divides Ω_S into hypercubes, an n -piece MTE potential $\phi : \Omega_S \mapsto \mathcal{R}^+$ has components

$$\phi_h(s) = a_0 + \sum_{i=1}^m a_i \exp\{b_i \cdot s\}$$

for $h = 1, \dots, n$, where $a_i, i = 0, \dots, m$ and $b_i, i = 1, \dots, m$ are real numbers.

MTE potentials can be used to approximate both probability distributions and utility functions. The optimization procedure outlined by Cobb et al. (2006) is used to determine the parameters (the values a_i and b_i) required to approximate probability density functions with MTE potentials.

2.3 Operations on MTE Potentials

In this paper, the operations of combination and marginalization are used to solve IDs where MTE potentials are used to represent probability density functions.

Definition 2. (Combination.) Combination of MTE potentials is pointwise multiplication. Let ϕ_1 and ϕ_2 be MTE potentials for \mathbf{X}_1 and \mathbf{X}_2 . The combination of ϕ_1 and ϕ_2 is a new MTE potential for $\mathbf{X} = \mathbf{X}_1 \cup \mathbf{X}_2$ defined as follows

$$\phi(\mathbf{x}) = (\phi_1 \otimes \phi_2)(\mathbf{x}) = \phi_1(\mathbf{x} \downarrow \Omega_{\mathbf{X}_1}) \cdot \phi_2(\mathbf{x} \downarrow \Omega_{\mathbf{X}_2})$$

for all $\mathbf{x} \in \Omega_{\mathbf{X}}$.

Combination of two MTE probability densities results in an MTE probability density. Combination of an MTE probability density and an MTE utility potential results in an MTE utility potential, as does the combination of two MTE utility potentials. Note that since a discrete probability distribution is a special case of an MTE potential where a_1, \dots, a_m in each component are equal to zero, this definition of combination applies to discrete probability distributions.

Definition 3. (Marginalization of Chance Variables.) Marginalization of a chance variable is summation over its state space. Let ϕ be an MTE potential for $\mathbf{X} = \mathbf{X}' \cup X$. The state space of X is $\Omega_X = \{x_1, \dots, x_n\}$. The marginal of ϕ for a set of variables \mathbf{X}' is an MTE potential computed as

$$\phi^{\downarrow \mathbf{X}'}(\mathbf{x}') = \phi^{-X}(\mathbf{x}') = \sum_{i=1}^n \phi(X = x_i, \mathbf{x}') \quad (1)$$

for all $\mathbf{x}' \in \Omega_{\mathbf{X}'}$. If the variable X is a continuous chance variable, the summation in Eq. (1) is replaced with integration as follows (assuming the state space of X is $\Omega_X = \{x : x_{min} \leq x \leq x_{max}\}$):

$$\phi^{\downarrow \mathbf{X}'}(\mathbf{x}') = \phi^{-X}(\mathbf{x}') = \int_{\Omega_X} \phi(\mathbf{x}) dx$$

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

Definition 4. (Marginalization of Decision Variables.) In this paper, all decision variables are discrete and binary. Assume I is a discrete decision variable with possible values $I = 0$ and $I = 1$ that has a continuous parent S with $\Omega_S = \{s : s_{min} \leq s \leq s_{max}\}$. Without loss of generality, arbitrarily assign $I = 0$ to the binary state of I that maximizes the value of u at s_{min} . To remove I from the ID, a threshold, Ψ , is determined as follows:

INPUT: $u, s_{min}, s_{max}, \epsilon$

OUTPUT: Ψ

INITIALIZATION: $\Psi = s_{min}$

DO WHILE $(u(I = 0, \Psi + \epsilon) \geq$
 $u(I = 1, \Psi + \epsilon)) \cap (\Psi \leq s_{max})$
 $\Psi = \Psi + \epsilon$

END DO

$\Psi = \Psi + \epsilon/2$

The parameter ϵ is an increment in S that can be assigned an appropriate value based on the application being addressed. The decision variable I is removed from the model by constructing the following MTE potential using the utility function u and the threshold value Ψ :

$$u^{\downarrow S}(s) = \begin{cases} u(I = 0, s) & \text{if } s_{min} \leq s < \Psi \\ u(I = 1, s) & \text{if } \Psi \leq s \leq s_{max} \end{cases} .$$

This definition applies when $u(I = 0, s) = u(I = 1, s)$ at one point and is a simpler version of the line search technique defined by Cobb and Shenoy (2008).

2.4 Fusion Algorithm

The ID is solved by applying the fusion algorithm (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 (according to Definition 2), then the variable is marginalized from the result. The appropriate marginalization operation depends on whether the variable being marginalized is a chance variable (see Definition 3) or a decision variable (see Definition 4).

3 ID Model

This section describes the ID model.

3.1 Graphical Representation

The ID model for the credit card detection problem is shown in Figure 1. The single-border ovals represent discrete chance variables. Fraud (F) indicates whether or not an order is fraudulent. The variables A_0 and P_0 reveal the number of suspicious characteristics in the address and product information on an order, respectively. The double-border oval for order Size (S) defines a continuous chance variable for the value (or cost) of the merchandise contained on an order. The arrows (or arcs) pointing from F to A_0 , F to P_0 , and F to S specify that the probability distributions for those variables are conditioned on F .

The rectangle in Figure 1 represents the firm's decision on whether or not to investigate an order. The arcs pointing from the variables A_0 to




Figure 1: ID Model.

I , P_0 to I , and S to I show that the firm will observe the values of these chance variables prior to making its decision on whether or not to investigate the order. The value $I = 0$ means that the firm does not investigate, while the value $I = 1$ means the firm investigates.

The diamond in the ID represents the joint utility function. The arcs pointing from F , S , and I to this node indicate that these variables are in the domain of the joint utility function. In this context, the firm's utility will be the cost to the firm of either investigating potential fraud or shipping unpaid merchandise.

3.2 Numerical Representation

This section describes the potentials in the ID.

3.2.1 Fraud (F)

The variable F has state space $\Omega_F = \{0, 1\}$, where $F = 0$ stands for legitimate and $F = 1$ signifies fraudulent. Thus, the probability of fraud is denoted by $P(F = 1) = \eta$.

3.2.2 Address Characteristics (A_0)

The variable A_0 has state space $\Omega_{A_0} = \{0, \dots, m\}$. The distribution $P(A_0|F)$ is constructed by using m variables representing the presence of specific address characteristics for credit card orders. For illustrative purposes, the remainder of the description of the potentials in the model will assume $m = 3$. Extension to the more general case is straightforward.

The variable A_0 "aggregates" the factors represented by the variables $\{A_1, A_2, A_3\}$; thus, the value of A_0 is determined by the number of variables in the set $\{A_1, A_2, A_3\}$ whose values equal one. The distribution $P(A_0|(A_1, A_2, A_3))$ is shown in Table 1. Combining the information in the variables $\{A_1, \dots, A_m\}$ into one vari-

Table 1: $P(A_0|(A_1, A_2, A_3))$.

			A_0			
A_1	A_2	A_3	0	1	2	3
0	0	0	1	0	0	0
1	0	0	0	1	0	0
0	1	0	0	1	0	0
0	0	1	0	1	0	0
1	1	0	0	0	1	0
0	1	1	0	0	1	0
1	0	1	0	0	1	0
1	1	1	0	0	0	1

able (A_0) reduces the number of decision rules determined when solving the ID from 2^{m+n} to $(m+1)(n+1)$, which significantly reduces the computational complexity of the solution. The resulting policies are also easier to implement.

The probability potential for the factor A_i given F is defined according to two parameters. The probabilities of the presence of the address inconsistency given the two states of F are $\alpha_{i,0} = P(A_i = 1|F = 0)$ and $\alpha_{i,1} = P(A_i = 1|F = 1)$, $\alpha_{i,0} < \alpha_{i,1}$, for $i = 1, 2, 3$.

Given $P(A_0|(A_1, A_2, A_3))$ and $P(A_i|F)$ for $i = 1, 2, 3$, the probabilities $P(A_0 = i|F = f)$ are determined as

$$P(A_0|F) = (P(A_0|(A_1, A_2, A_3)) \otimes P(A_1|F) \otimes P(A_2|F) \otimes P(A_3|F))^{-\{A_1, A_2, A_3\}}$$

according to Definitions 2 and 3. The result is shown in Table 2 and follows directly from the chain rule for Bayesian networks (Pearl, 1988).

3.2.3 Product Characteristics (P_0)

The variable P_0 has state space $\Omega_{P_0} = \{0, \dots, n\}$. The distribution $P(P_0|F)$ is constructed by using n variables representing the presence of specific product characteristics for credit card orders. The variables P_1, \dots, P_n are factors related to an order's product information that may be useful for distinguishing a legitimate order from a fraudulent order. For example, fraudulent orders are more likely than acceptable orders to have multiples of the same item.

Table 2: Probability Distribution for A_0 given F ($P(A_0 = i|F = f)$).

	$F = 0$	$F = 1$
$A_0 = 0$	$(1 - \alpha_{10})(\alpha_{20} - 1)(\alpha_{30} - 1)$	$(1 - \alpha_{11})(\alpha_{21} - 1)(\alpha_{31} - 1)$
$A_0 = 1$	$\alpha_{20} + \alpha_{30} - 2\alpha_{20}\alpha_{30}$ $+ \alpha_{10}(1 - 2\alpha_{30} + \alpha_{20}(3\alpha_{30} - 2))$	$\alpha_{21} + \alpha_{31} - 2\alpha_{21}\alpha_{31}$ $+ \alpha_{11}(1 - 2\alpha_{31} + \alpha_{21}(3\alpha_{31} - 2))$
$A_0 = 2$	$\alpha_{20}\alpha_{30} + \alpha_{10}(\alpha_{20} + \alpha_{30} - 3\alpha_{20}\alpha_{30})$	$\alpha_{21}\alpha_{31} + \alpha_{11}(\alpha_{21} + \alpha_{31} - 3\alpha_{21}\alpha_{31})$
$A_0 = 3$	$\alpha_{10}\alpha_{20}\alpha_{30}$	$\alpha_{11}\alpha_{21}\alpha_{31}$

The probabilities of the presence of the product characteristics given the two states of F are $\rho_{j,0} = P(P_j = 1|F = 0)$ and $\rho_{j,1} = P(P_j = 1|F = 1)$, $\rho_{j,0} < \rho_{j,1}$, $j = 1, \dots, n$. The variable P_0 summarizes the information in $\{P_1, \dots, P_n\}$ in much the same way as A_0 summarizes the information in the address characteristics for an order.

Calculation of $P(P_0|F)$ is accomplished in the same way as the determination of $P(A_0|F)$, so the details are omitted. More information is provided in (Cobb, 2010).

3.2.4 Order Size (S)

The chance variable S has $\Omega_S = \{s : s_{min} \leq s \leq s_{max}\}$. The probability potential ϕ for $\{F, S\}$ represents the conditional probability density functions for S given $F = 0$ and $F = 1$.

The parameters used in the example of the next section will be used to describe the potential ϕ . Suppose the natural log of S is normally distributed with mean $\mu = 2.5$ and variance $\sigma^2 = 0.5$ given that $F = 0$, i.e. $S|F = 0 \sim LN(2.5, 0.5)$. Also, assume the natural log of S given $F = 1$ is normally distributed with $\mu = 3.5$ and $\sigma^2 = 0.75$, i.e. $S|F = 1 \sim LN(3.5, 0.75)$. The MTE potential fragment representing the conditional distribution for S given $F = 0$ is defined as




Figure 2: MTE Lognormal Approximations.

$$\phi(F = 0, s) = \hat{f}_{F=0}(s) =$$

$$\left\{ \begin{array}{l} 39.78 + 5.71 \exp\{0.0326(s - 7.39)\} \\ -45.34 \exp\{0.0032(s - 7.39)\} \\ \quad \text{if } 1.46 \leq s < 2.72 \\ 21.72 - 1.49 \exp\{-0.0468(s - 7.39)\} \\ -20.17 \exp\{0.0035(s - 7.39)\} \\ \quad \text{if } 2.72 \leq s < 7.39 \\ \vdots \end{array} \right.$$

The full description of the function can be found in (Cobb, 2010). The MTE potential fragment $\phi(F = 1, s)$ representing the conditional distribution for S given $F = 1$ is defined similarly, and both MTE potentials are displayed in Figure 2, overlaid on the actual lognormal distributions. The prior distribution for order size is skewed farther to the right for fraudulent orders.

3.2.5 Utility Function (u_0)

The joint utility function u_0 has domain $\{F, S, I\}$. Assume that the cost of investigating an order for fraudulent activity is c (a constant). The values for u_0 are $u_0(F = 0, I = 0, s) = 0$,

$u_0(F = 0, I = 1, s) = -c$, $u_0(F = 1, I = 0, s) = -s$, and $u_0(F = 1, I = 1, s) = -c$.

If the firm investigates, it incurs the cost of the investigation, but avoids forfeiting the value of the merchandise when a fraudulent order is thwarted. If the firm fails to investigate a fraudulent order, it incurs a cost equal to the value of the merchandise shipped to fulfill the order.

The joint utility function u_0 is approximated by the MTE potential u_1 , which is identical to u_0 with the exception of one term, which is defined (Cobb and Shenoy, 2006) as $u_1(F = 1, I = 0, s) = (s_{max} - s_{min})(13.512870 \cdot \exp\{\frac{0.071387}{s_{max} - s_{min}} \cdot (s - s_{min})\} - 13.507018) - s_{min}$.

For the case where $F = 1$ and $I = 0$, the function u_1 is an MTE approximation to the linear function $g(s) = -s$. In the ID solution process, the MTE utility function will be combined via multiplication with the MTE density potential for S given F . Since the class of MTE potentials is closed under addition and multiplication, the result remains an MTE potential. This allows the resulting function to be integrated in closed form to determine the firm's maximum expected utility.

4 Example

This section describes an example where optimal decision rules are developed that allow the firm to decide when to investigate potentially fraudulent orders.

4.1 Problem Description

Assume $m = n = 3$, meaning that there are three address factors and three product factors used to determine the conditional distributions for A_0 given F and P_0 given F , respectively. These factors are:

- Shipping and billing addresses match (A_1)
- Untraceable e-mail (A_2) — the order originated from a free, web-based address.
- Foreign address (A_3)
- Leave at home (P_1) — the customer requests that the shipment be left at the door if no one is home.

Table 3: Parameters for the Example.

Variable	$F = 0$	$F = 1$
Fraud (F)	$1 - \eta = .99$	$\eta = .01$
Match (A_1)	$\alpha_{10} = .25$	$\alpha_{11} = .40$
Email (A_2)	$\alpha_{20} = .01$	$\alpha_{21} = .05$
Inter. (A_3)	$\alpha_{30} = .05$	$\alpha_{31} = .25$
Leave (P_1)	$\rho_{10} = .20$	$\rho_{11} = .30$
Rush (P_2)	$\rho_{20} = .10$	$\rho_{21} = .20$
Mult. (P_3)	$\rho_{30} = .05$	$\rho_{31} = .075$
Size (S)	$LN(2.5, .5)$	$LN(3.5, .75)$

Rush shipping (P_2)

Multiple units of the same item (P_3)

The presence of these factors is denoted by either $A_i = 1$ or $P_j = 1$ and corresponds to a higher incidence of fraud.

The potential representing the prior probability distribution for F has values $P(F = 0) = 1 - \eta = 0.99$ and $P(F = 1) = \eta = 0.01$. The conditional probability density functions for order Size (S) are those approximated by the MTE potential ϕ in Figure 2. The cost of investigating an order is $c = 10$, and the MTE potential fragment approximating $u_0(F = 1, I = 0, s)$ in the joint utility function u_0 is $u_1(F = 1, I = 0, s) = 5989.45 - 5993.51 \exp\{0.000161(s - 1.46)\}$.

The probability distribution $P(A_0|F)$ is determined using the result in Table 2, and $P(P_0|F)$ is calculated similarly. A summary of the parameters in the example problem is given in Table 3.

4.2 Solution

This section briefly describes the solution to the example using the fusion algorithm. In this problem, a possible deletion sequence is F, I, S, A_0, P_0 .

The potentials in the model at the outset are $P(F)$, $P(A_0|F)$, $P(P_0|F)$, ϕ for $\{S, F\}$, and u_1 for $\{F, S, I\}$. The first variable in the deletion sequence is F , and since all potentials contain F in their domain, all must be combined prior to the marginalization of F . The combination




Figure 3: u_2 where $A_0 = 2$ and $P_0 = 2$.

results in an MTE utility potential determined as

$$u'_2 = P(F) \otimes P(A_0|F) \otimes P(P_0|F) \otimes \phi \otimes u_1 .$$

The variable F is marginalized as

$$\begin{aligned} u_2(A_0 = i, P_0 = j, I = k, s) = \\ u'_2(F = 0, A_0 = i, P_0 = j, I = k, s) \\ + u'_2(F = 1, A_0 = i, P_0 = j, I = k, s) \end{aligned}$$

for all $(i, j, k, s) \in \Omega_{\{A_0, P_0, I, S\}}$. The function u_2 is shown graphically in Figure 3 for the case where the number of observed address inconsistencies and suspicious product characteristics are both two ($A_0 = 2$ and $P_0 = 2$). The expected utility that results from investigating the potential fraud ($I = 1$) is less than the expected utility associated with not investigating ($I = 0$) for smaller orders.

The objective of the firm is to decide optimally whether to investigate potential fraud after it observes the values of A_0 , P_0 , and S . Thus, for each configuration of states of the discrete variables $A_0 = i$ and $P_0 = j$, the firm must choose a threshold $\Psi_{i,j}$ for order size using the procedure in Definition 4. As an example, $\Psi_{2,2} = 31.96$ and is determined by finding the point where the two functions in Figure 3 are approximately equal. For values of S below this threshold, the firm will be better off in the long run not investigating the order for potential fraud. For values of S above this threshold, due to the potential loss of merchandise, the firm should investigate potential fraud on an order. Methods of investigating fraud include

Table 4: Decision Thresholds $\Psi_{i,j}$.

Ψ	$P_0 = 0$	$P_0 = 1$	$P_0 = 2$	$P_0 = 3$
$A_0 = 0$	95.56	89.86	76.96	61.96
$A_0 = 1$	83.46	67.16	52.56	44.06
$A_0 = 2$	47.16	38.96	31.96	26.86
$A_0 = 3$	28.16	23.16	19.06	16.16

validating the billing address, shipping address, e-mail address, and phone number, and contacting the customer to confirm the order. This investigation is carried out at an average cost of $c = 10$ per order.

The decision variable I is removed from the model by constructing the following MTE potential using the utility function u_2 and the threshold values $\Psi_{i,j}$:

$$u_3(A_0 = i, P_0 = j, s) = \begin{cases} u_2(A_0 = i, P_0 = j, I = 0, s) & \text{if } s_{min} \leq s < \Psi_{i,j} \\ u_2(A_0 = i, P_0 = j, I = 1, s) & \text{if } \Psi_{i,j} \leq s \leq s_{max} . \end{cases}$$

for $i = 0, 1, 2, 3$ and $j = 0, 1, 2, 3$.

4.3 Results

The decision thresholds $\Psi_{i,j}$ for the example problem are shown in Table 4. At higher numbers of address inconsistencies and suspicious product characteristics, the order size thresholds are lower, since the firm has more reason to believe that such orders are not legal.

The maximum expected utility from following the decision thresholds is calculated as

$$u_4 = \sum_{j=1}^n \left(\sum_{i=1}^m \left(\int_{\Omega_S} u_3(A_0 = i, P_0 = j, s) ds \right) \right) .$$

For the example problem, the maximum expected utility is $u_4 = -0.31173$, which represents the fraud investigation and loss expense stated on a per order basis.

Additional results and sensitivity analysis for this example are provided in (Cobb, 2010).

5 Conclusions

An ID model that can be used to detect potentially fraudulent online credit card transactions

was introduced. A business can employ such a model to establish decision policies that guide employees to investigate orders that are most likely to be suspect. By following such policies, a business can minimize its total fraud-related expenses, which include both the costs of lost merchandise and the expense of following up on suspicious orders.

In addition to considering order size as a criteria for identifying potentially fraudulent orders, the ID model allows a business to consider other characteristics of the address and product information on an order. When more of these factors indicate that fraud may be present, the order size threshold used to decide whether or not to investigate an order is lowered, because an illegal order becomes more likely. Using the ID to establish such rules allows a business to investigate the orders that are most likely to be fraudulent and save the cost of such inquiries on orders—even large ones—that are most likely legitimate.

Future research can incorporate additional complexities to make the model more realistic. For instance, an implicit assumption is that the investigation always concludes with certainty that an order is fraudulent. A node representing the result of the investigation can be added to the ID to relax this assumption. In cases where a good order is mistakenly canceled, the cost of the customer's dissatisfaction should be considered in the utility function. Also, the investigation cost, c , may not be a constant and can be modeled as a random variable, perhaps conditional on the number of suspicious address and product characteristics observed.

Acknowledgments

The author is grateful for the insightful comments of three anonymous reviewers. Support from the Spanish Ministry of Science and Innovation through project TIN2007-67418-C03-02 and by EFDR funds is gratefully acknowledged.

References

Cobb, B.R. 2010. Detecting Online Credit Card Fraud with Hybrid Influence Diagrams,

Working Paper, Virginia Military Institute, Lexington, VA. Available for download at: www.vmi.edu/fswebs.aspx?tid=24697&id=24791

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. *European J. Oper. Res.* **186**(1) 261–275.

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

Credit card fraud: A guide to help businesses recognize it, report it, stop it. 2008. www.visa.ca/en/merchant/pdfs/merchant_fraud.pdf. Accessed on 26 November 2008.

Howard, R.A., J.E. Matheson. 1984/2005. 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.

Pearl, J. 1988. *Probabilistic Reasoning in Expert Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Francisco, CA.

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.

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.