MATHEMATICAL ENGINEERING TECHNICAL REPORTS

Risk-sensitive mixed discrete-Gaussian networks

Kei KOBAYASHI and Fumiyasu KOMAKI

METR 2004–36 June 2004

DEPARTMENT OF MATHEMATICAL INFORMATICS
GRADUATE SCHOOL OF INFORMATION SCIENCE AND TECHNOLOGY
THE UNIVERSITY OF TOKYO
BUNKYO-KU, TOKYO 113-8656, JAPAN

WWW page: http://www.i.u-tokyo.ac.jp/mi/mi-e.htm

Risk-sensitive mixed discrete-Gaussian networks

Kei Kobayashi*and Fumiyasu Komaki†
Department of Mathematical Informatics
Graduate School of Information Science and Technology
University of Tokyo

June, 2004

Abstract

The paper presents a class of decision network models with a risk-sensitive parameter. The risk-sensitive parameter represents the decision maker's degree of optimism or pessimism. We propose a decision algorithm for exponential quadratic mixed discrete-Gaussian networks and an LEQG optimal control for mixed discrete-Gaussian networks. The proposed class of mixed discrete-Gaussian networks includes various models for decision problems. Risk-sensitive versions of various decision making problems can be analyzed using the introduced algorithm.

Next, the H_∞^Δ optimal decision is proposed as a generalization of H_∞ optimal control. The H_∞^Δ optimal decision has robustness similar to that of the H_∞ optimal control. It is proved that the H_∞^Δ optimal decision is a limit of the optimal decision for the corresponding exponential quadratic mixed discrete-Gaussian networks. A direct calculation method for the H_∞^Δ optimal decision is proposed.

Key words: graphical model, linear exponential quadratic Gaussian model, optimal control, risk-sensitivity, mixed discrete-Gaussian networks, H_∞ optimal control.

1 Introduction

1.1 The classical model

The state-space linear exponential quadratic Gaussian (LEQG) model has a controlled stochastic dynamics represented by state and observation equations as

$$x_t = Ax_{t-1} + Bu_{t-1} + \epsilon_t, \tag{1}$$
$$y_t = Cx_{t-1} + \eta_t \quad (t = 1, \dots, h), \tag{2}$$

*kkoba@stat.t.u-tokyo.ac.jp
†komaki@mist.i.u-tokyo.ac.jp

where x_t, u_t, and y_t are the process, control, and observation variables, respectively, at discrete time t, and A, B and C are appropriately chosen matrices. Here, x_t and y_t are finite dimensional real-valued vectors and

$$\begin{bmatrix} \epsilon \\ \eta \end{bmatrix}_t \sim N\left(0, \begin{bmatrix} N & L \\ L' & M \end{bmatrix}_t\right), \tag{3}$$

where N, L, and M are appropriately chosen matrices. Let W_t be $(W_0, u_{0:t-1}, y_{1:t})$, where W_0 indicates the initial information, and $u_{0:t-1} = (u_0, u_1, \dots, u_{t-1})$ and $y_{1:t} = (y_1, y_2, \dots, y_t)$ indicate the control and observation histories. A policy, π, is a specification of the control variable, u_t, as a function of W_t $(t = 1, \dots, h)$.

Let Q be a quadratic utility function of the process path $\{x_t, u_t\}$ as

$$Q = -\frac{1}{2}\sum_{t=0}^{h-1}\left[\begin{bmatrix} x \\ u \end{bmatrix}' \begin{bmatrix} R & S' \\ S & W \end{bmatrix} \begin{bmatrix} x \\ u \end{bmatrix}\right]_t - \frac{1}{2}[x'\Pi_h x]_h, \tag{4}$$

with appropriately chosen matrices R, S, W, and Π_h. In the classical linear quadratic Gaussian (LQG) optimal control, the optimal control policy π is the one maximizing the criterion $E_\pi[Q]$.

The classical LEQG optimal control has the criterion

$$\gamma_\pi(\theta) = \theta^{-1} \log E_\pi[\exp(\theta Q)], \tag{5}$$

where θ is a scalar variable. We term such a criterion as *an exponential quadratic criterion* or *a risk sensitive criterion*. When $\theta = 0$, $\gamma_\pi(0)$ is defined by $\lim_{\theta\to 0} \gamma_\pi(\theta)$. The variable θ is the risk-sensitivity parameter. The decision becomes more optimistic as θ increases and more pessimistic as θ decreases. The terms "optimistic" and "pessimistic" are explained as follows. When θ is close to zero, $\gamma_\pi(\theta)$ is expanded as

$$\gamma_\pi(\theta) = E_\pi[Q] + \frac{\theta}{2} Var_\pi[Q] + o(\theta).$$

If θ is positive, the optimizer wants to increase the variance $Var_\pi[Q]$ as well as $E_\pi[Q]$. This means he is optimistic and risk-seeking because the variance seems to be a good option for him. If θ is negative, the optimizer wants to decrease the variance; he is pessimistic.

In particular, when $\theta = 0$, $\gamma_\pi(0)$ is equal to $E_\pi[Q]$. Thus, we call such a case "risk-neutral." We can easily check $\lim_{\theta\to\infty} \gamma_\pi(\theta) = \sup_x Q$ and $\lim_{\theta\to-\infty} \gamma_\pi(\theta) = \inf_x Q$. These cases can be recognized as the most optimistic and pessimistic cases, respectively.

1.2 The general linear exponential quadratic Gaussian model

We say that the optimization problem over $[0, h]$, where h is the horizon, has the *LQG structure* if the following five conditions [12] hold:

- (LQG1) For each t, the control variable, u_t, takes values in some finite-dimensional real-valued vector space, and may take any value in this space.

- (LQG2) The horizon h is fixed, and the utility function $\mathcal{U}$, whose expectation $E_\pi[\mathcal{U}]$ is to be maximized, can be expressed in the reduced form as $\mathcal{U} = Q(u_{0:h-1}, \xi)$, where ξ is a vector-valued noise vector.

- (LQG3) The function Q is quadratic in all arguments, and negative definite in $u_{0:h-1}$.

- (LQG4) Conditional on the state of information, W_0, at time $t = 0$ and whatever the control policy, the noise vector, ξ, is normally distributed with zero mean and covariance matrix, V, independent of the policy.

- (LQG5) The observations, y_t, are reducible to a policy-independent linear function of ξ, without loss of information.

The original notation by [12] uses the cost function $\mathcal{C} = -\mathcal{U}$ instead of the utility function $\mathcal{U}$.

The *LEQG optimal control* model satisfies all the LQG assumptions (LQG1-5), except for the replacement of the maximization of the criterion $E_\pi(Q)$ by the maximization of the criterion $\gamma_\pi(\theta)$.

1.3 The risk-sensitive certainty-equivalence principle (RSCEP)

We define the *total negative stress* by

$$S(u_{0:h-1}, \xi) := Q - \theta^{-1}\xi' V^{-1}\xi/2.$$

We use the term *θ-extremization* to denote an operation that one maximizes when $\theta \geq 0$ and minimizes when $\theta < 0$. This operation is simply denoted by "ext" (analogously to "max" or "sup").

The optimization in LEQG model can be effectively evaluated by using the following theorem:

Theorem 1.1 (The risk-sensitive certainty-equivalence principle (RSCEP) [12]) *Assume conditions LQG1-5, and assume that the policy, π, is to be chosen to maximize the criterion $\gamma_\pi(\theta)$ defined by (5). Also, suppose that θ has a value such that $-\infty < \max_{u_{0:h-1}} \mathrm{ext}_\xi\, S(u_{0:h-1}, \xi) < \infty$ when $\theta < 0$. Then,*

1. *The optimal value of u_t is determined by simultaneously maximizing S with respect to $u_{t:h-1}$ and θ-extremizing it with respect to $y_{t+1:h}$. In other words, one obtains an optimal current decision by maximizing negative stress with respect to all currently unmade decisions and θ-extremizing it with respect to all currently unobservable quantities.*

2. *Define the value function as $G(W_t) = \mathrm{ext}_\pi[f(y_{1:t}|u_{0:t-1})E_\pi \exp(\theta Q)]$. The value function $G(W_t)$ of the problem is calculated as*

$$G(W_t) = \exp[\theta(k_t + S_t(W_t))].$$

Here $S_t(W_t)$ obeys the recursion

$$S_t(W_t) = \max_{u_t} \operatorname*{ext}_{y_{t+1}} S_{t+1}(W_{t+1}), \quad (0 \le t < h)$$

with the terminal condition

$$S_h(W_h) = S(W_h),$$

and k_t is a constant, independent of W_t.

If we use the RSCEP (Theorem 1.1), the optimal control for the LEQG model is directly calculable by the Kalman filter. On the other hand, the RSCEP is not directly applicable to graphical models including both of discrete and continuous variables. In the present paper, we introduce a decision theory based on the risk-sensitive criterion (5) for such network models.

In section 2, exponential quadratic mixed discrete-Gaussian network (EQM) models are proposed. This model is a risk-sensitive decision network model with both discrete and continuous nodes. The calculation can be carried out effectively under certain conditions by using the strong decomposition technique introduced in section 2.

In section 3, the H_∞^Δ optimal decision theory for EQM models is introduced. The H_∞^Δ optimal decision has some robustness as the H_∞ optimal control does. It is proved that the H_∞^Δ optimal decision is a limit of the optimal decision for the corresponding EQM models. An efficient algorithm to obtain the H_∞^Δ optimal decision is also proposed.

In section 4, some experimental results are presented.

2 Risk-sensitive decision networks

2.1 Mixed discrete-Gaussian networks

In this section, we deal with a graphical model that includes both quantitative (continuous) and qualitative (discrete) variables. For simplicity, the distinction between a node and the associated variable is ignored in this paper.

The following notations are used: $V = \{v_1, \ldots, v_n\}$ is the set of all nodes, and Γ and Δ are the sets of all continuous and discrete nodes, respectively. Throughout this paper, we assume that discrete variables take a finite number of values. The set of all random variables is denoted by $x = (x_\alpha)_{\alpha \in V} = (i, \zeta) = ((i_\delta)_{\delta \in \Delta}, (\zeta_\gamma)_{\gamma \in \Gamma})$, where i_δ is a discrete variable and ζ_γ is a continuous variable.

We say that x follows a *conditional Gaussian distribution (CG-distribution)* if the simultaneous distribution has a density

$$f(x) = f(i, \zeta) = \exp\{g(i) + h(i)'\zeta - \zeta' K(i)\zeta/2\} \tag{6}$$

where $g(i), h(i)$, and $K(i)$ are scalar, vector, and positive-definite matrices, respectively [6]. We term network models with a CG-distribution as *mixed discrete-Gaussian networks.*

If an operation on function $f(x)$ does not change the form of (6) but only changes the values of $g(i), h(i)$, and $K(i)$, we say that the operation is *CG-calculable*. If an algorithm consists of a sequence of CG-calculable operations, we say the algorithm is CG-calculable. The CG-calculablitiy is an important property because every CG-calculable operation can be represented as a transformation of $g(i), h(i)$, and $K(i)$. For example, integration and maximization of $f(x)$ with respect to continuous variables are CG-calculable, see lemmas A.1 and A.2 in the appendix.

Example 2.1 (The CG-decomposable distribution model)

Here, we consider decomposable graphical models.[1] Let $\mathcal{C}$ and $\mathcal{S}$ denote the set of all cliques and the set of all separators in the triangulated graph, respectively. We call the model has a decomposable distribution if the density of the simultaneous distribution is decomposed as

$$f(x) \propto \frac{\prod_{C\in\mathcal{C}} \phi_C}{\prod_{S\in\mathcal{S}} \phi_S}. \tag{7}$$

Here, each ϕ_C is represented as

$$\phi_C = \phi_C(i,\zeta) = \exp\{g_C(i) + h_C(i)'\zeta - \zeta' K_C(i)\zeta/2\} \quad \text{for} \quad C = (i,\zeta), \tag{8}$$

where $g_C(i)$, $h_C(i)$, and $K_C(i)$ are scalar, vector, and symmetric matrices, respectively. We term such a ϕ_C as *a CG-potential*. We note that $K_C(i)$ is symmetric but not necessarily positive-definite. The CG-potential ϕ_S for a separator S is defined in the same manner. The right-hand side of (7) is interpreted as 0 whenever any term in the denominator is 0.

It is evident that the joint density function (7) can be represented in the form of (6).

Example 2.2 (The DAG model)

The directed acyclic graphs (DAG) is an important and useful class of CG-distributed graphical models. Here, we assume that continuous nodes do not have discrete child nodes.

Let $pa(x_n)$ be the set of all the parent nodes of x_n. Let $\Gamma_n^{(p)} := pa(x_n) \cap \Gamma$ and $\Delta_n^{(p)} := pa(x_n) \cap \Delta$. In DAG models with a CG-distribution, the distribution of each continuous node $x_n \in \Gamma$ is given by

$$x_n | pa(x_n) = \sum_{k: x_k \in \Gamma_n^{(p)}} B_{nk} x_k + m_n(\Delta_n^{(p)}) + v_n^{1/2}(\Delta_n^{(p)})\epsilon_n, \tag{9}$$

$$\epsilon_n \sim N(0, I_p), \tag{10}$$

where, x_n and x_k are the p and q dimensional real-valued vectors respectively, B_{nk} is a $q \times p$ real-valued matrix, ans $m_n(\Delta_n^{(p)})$ and $v_n(\Delta_n^{(p)})$ are a p-dimensional real-vector valued function and a $p \times p$ positive-definite matrix-valued function, respectively.

[1] See [2] about CG-decomposable graphs and other terminologies of ordinary graphical models (for example, parent, child, ancestor, descendent, neighbor, moralization, and triangulation).

All the parents of each discrete node are discrete by the assumption. The conditional distribution of a discrete variable x_n is given by

$$p(x_n|\Delta_n^{(p)}) = M(x_n; q(\Delta_n^{(p)})), \tag{11}$$

where $M(x_n; q)$ is a multinomial density with parameter q. When the model distribution is given by (9) and (11), it is easy to prove that the joint distribution of all variables is a CG-distribution and can be decomposed as (7) and (8).

2.2 Exponential quadratic mixed discrete-Gaussian networks

In this subsection, we consider mixed discrete-Gaussian networks with the exponential quadratic criterion (5). We assume that graphs for decision networks have observation variables $Y = \{y_1, \ldots, y_h\}$, decision variables $D = \{d_0, d_1, \ldots, d_{h-1}\}$ and the other variables N. We call N "nuisance variables."

The conditional distribution $p(Y, N|D)$ is assumed to be decomposed as

$$\begin{aligned} p(Y, N|D) = p(y_1|d_0)p(y_2|d_{0:1}, y_1) \ldots p(y_{t+1}|d_{0:t}, y_{1:t}) \ldots \\ p(y_h|d_{0:h-1}, y_{1:h-1})p(N|d_{0:h-1}, y_{1:h}). \end{aligned} \tag{12}$$

The decomposition (12) corresponds to the decisions and observations whose chronological order is

$$d_0, y_1, d_1, y_2, \ldots, d_{h-1}, y_h. \tag{13}$$

If we consider some of elements in Y and D as empty, all the chronological orders of decisions and observations are represented as (13).

Let $\{N_i\}_{i=1}^{h+1}$ be a partition of N and $N_{0:i} := \bigcup_{j=0}^{i} N_j$. If $p(Y, N|D)$ is decomposed as

$$\begin{aligned} p(Y, N|D) = p(N_0)p(y_1, N_1|d_0, N_0)p(y_2, N_2|d_{0:1}, y_1, N_{0:1}) \ldots p(y_{t+1}, N_{t+1}|d_{0:t}, y_{1:t}, N_{0:t}) \ldots \\ p(y_h, N_h|d_{0:h-1}, y_{1:h-1}, N_{0:h-1})p(N_{h+1}|d_{0:h-1}, y_{1:h}, N_{0:h}), \end{aligned} \tag{14}$$

then

$$\begin{aligned} p(Y, N|D) &= p(y_1|d_0)p(y_2|d_{0:1}, y_1) \ldots p(y_h|d_{0:h-1}, y_{1:h-1}) \\ &\quad p(N_0)p(N_1|d_0, y_1, N_0) \ldots p(N_h|d_{0:h-1}, y_{1:h}, N_{0:h-1})p(N_{h+1}|d_{0:h-1}, y_{1:h}, N_{0:h}) \\ &= p(y_1|d_0)p(y_2|d_{0:1}, y_1) \ldots p(y_h|d_{0:h-1}, y_{1:h-1})p(N|d_{0:h-1}, y_{1:h}). \end{aligned} \tag{15}$$

Therefore, the conditional distribution (14) satisfies the assumption (12).

The decision policy $\mathcal{D}$, corresponding to the control π of an LEQG model, is represented by a constant d_0 and a set of functions $\{d_t(\cdot) \mid t = 1, \ldots, h-1\}$, where $d_t(\cdot)$ is a function of $y_{1:t}$ and $d_{0:t-1}$.

Similar to LEQG models, we assume that the criterion function is exponential quadratic function

$$\gamma_{\mathcal{D}}(\theta) = \theta^{-1} \log E_{\mathcal{D}}[\exp(\theta Q(x))], \tag{16}$$

where

$$Q(x) = Q(\zeta, i) = g^{(v)}(i) + h^{(v)}(i)'\zeta - \zeta' K^{(v)}(i)\zeta/2. \tag{17}$$

Here, ζ is a vector that consists of the all elements of Γ, and $K^{(v)}(i)$ is a positive semi-definite matrix.

Definition 2.3 (The exponential quadratic mixed discrete-Gaussian network model)

Assume that the conditional distribution (12) is a CG-distribution (6). Here, $K(i)$ in (6) is positive-semidefinite and the submatrix of $K(i)$ corresponding to all the nondecision variables $(\zeta_\gamma)_{\gamma \in \Gamma \cap (Y \cup N)}$ is strictly positive-definite. The risk-sensitive function (16) is adopted as a criterion function.

We term the decision model as an *EQM* model. If we emphasize that an EQM model has a risk-sensitivity parameter θ, we denote it as an EQM(θ) model.

An EQM model including only continuous variables is an LEQG model. In this case, the criterion function can be optimized by the RSCEP. However, in general, EQM models include both continuous and discrete variables. We introduce the following generalized RSCEP.

Remark 2.4 (The RSCEP for EQM models) *Consider an EQM model. We define value functions V_t for $t = 0, \ldots, h-1$ recursively as*

$$V_{h-1}(d_{0:h-1}, y_{1:h-1}; \theta) = \int \exp(\theta Q(x)) p(N, y_h | d_{0:h-1}, y_{1:h-1}) d\mu(y_h) d\mu(N),$$

and

$$V_t(d_{0:t}, y_{1:t}; \theta) = \int p(y_{t+1}|d_{0:t}, y_{1:t}) \max_{d_{t+1}(d_{0:t}, y_{1:t+1})} V_{t+1}(d_{0:t+1}, y_{1:t+1}; \theta) d\mu(y_{t+1}) \tag{18}$$

for $t = 0, \ldots, h-2$. Suppose that the decision results $d_{0:t-1}$ and observations $y_{1:t}$ are given. The optimal decision d_t^* for $t = 0, \ldots, h-1$ is given by

$$d_t^*(d_{0:t-1}, y_{1:t}) = \arg\max_{d_t} V_t(d_{0:t}, y_{1:t}; \theta).$$

Proof. Each optimal decision policy function $d_t(d_{0:t-1}, y_{1:t})$ for $t = 0, \ldots, h$ satisfies the following maximization:

$$\max_{d_0} \int_{y_1} p(y_1|d_0) \max_{d_1} \int_{y_2} p(y_2|d_{0:1}, y_1) \ldots \max_{d_t} \int_{y_{t+1}} p(y_{t+1}|d_{0:t}, y_{1:t}) \ldots$$
$$\max_{d_{h-1}} \int_{y_h} p(y_h|d_{0:h-1}, y_{1:h-1}) \int_N p(N|d_{0:h-1}, y_{1:h}) \exp(\theta Q(x)) d\mu(N) d\mu(y_{1:h}).$$

The remark is proved by induction. □

We cannot simplify the recursion formula (18) as in the regular RSCEP. This is because i) maximization with respect to discrete variables and integration are not permutable and ii) results of expectation with respect to discrete variables take the mixed Gaussian form, whose analytical maximization is difficult. We propose an efficient algorithm for calculating V_t for EQM models in the next subsection.

Although the optimal control for LEQG models is unique, there may be more than one optimal decision for CG-distributed models. If d_t is a discrete decision, $\max_{d_t} V_t(d_{0:t}, y_{1:t}; \theta)$ may be attained for more than one value of d_t. We recognize $d_t(\cdot)$ as a function of $d_{0:t-1}, y_{1:t}$, and let $\boldsymbol{d}_t(\cdot)$ be a set of all such functions. We choose an arbitrary function $d_t^*(\cdot)$ from $\boldsymbol{d}_t(\cdot)$, and continue the calculation for the rest of the decisions $d_{0:t-1}$. Then, the direct product of $\boldsymbol{d}_t(\cdot)$ for $t = 0, \dots, h-1$ corresponds to the set of all optimal decisions, i.e., for every $\hat{d}_0 \in \boldsymbol{d}_0, \hat{d}_1(\cdot) \in \boldsymbol{d}_1(\cdot), \dots, \hat{d}_{h-1}(\cdot) \in \boldsymbol{d}_{h-1}(\cdot)$, $(\hat{d}_0, \hat{d}_1(\hat{d}_0, y_1), \dots, \hat{d}_{h-1}(\hat{d}_0, \hat{d}_1(\hat{d}_0, y_1), \dots))$ is an optimal decision. The algorithmic description is as follows:

For $t = h-1, \dots, 0$

For each value of $d_{0:t-1}, y_{1:t}$,

1. $\boldsymbol{d}_t(d_{0:t-1}, y_{1:t}) \leftarrow \arg\max_{d_t} V_t(d_{0:t}, y_{1:t}; \theta)$.
2. If $t \geq 1$, choose an arbitrary value $d_t \in \boldsymbol{d}_t(d_{0:t-1}, y_{1:t})$ and $d_t^* \leftarrow d_t$.
3. $V_{t-1}(d_{0:t-1}, y_{1:t-1}; \theta) \leftarrow \int p(y_t | d_{0:t-1}, y_{1:t-1}) V_t(d_{0:t-1}, d_t^*, y_{1:t}; \theta) d\mu(y_t)$.

Example 2.5 (The CG-decomposable decision model)

Let $\mathcal{C}$ and $\mathcal{S}$ be the sets of cliques and separators of the triangulated graph, respectively. Assume that the conditional density function $p(Y, N|D)$ is CG-decomposed as $f(x)$ in (7) and that the exponential utility function is decomposed as

$$\exp(\theta Q(x)) \propto \frac{\prod_{C \in \mathcal{C}} \phi_C^{(v)}}{\prod_{S \in \mathcal{S}} \phi_S^{(v)}}. \tag{19}$$

Here each $\phi_C^{(v)}$ is described as

$$\phi_C^{(v)} = \phi_C^{(v)}(i, \zeta) = \exp\{\theta(g_C^{(v)}(i) + h_C^{(v)}(i)'\zeta - \zeta' K_C^{(v)}(i)\zeta/2)\} \quad \text{for} \quad C = (i, \zeta),$$

where $g_C^{(v)}(i)$, $h_C^{(v)}(i)$, and $K_C^{(v)}(i)$ are scalar, vector and symmetric matrices, respectively. $\phi_S^{(v)}$ of separator $S \in \mathcal{S}$ are defined in the same manner. The right-hand side term of (19) is interpreted as 0 whenever any term in the denominator is 0.

We call such a model *a CG-decomposable decision model*. In this model,

$$\exp(\theta Q(x)) p(Y, N|D) \propto \frac{\prod_{C \in \mathcal{C}} \phi_C \phi_C^{(v)}}{\prod_{S \in \mathcal{S}} \phi_S \phi_S^{(v)}}, \tag{20}$$

and the calculation of the optimal decision can be decomposed into local calculations as we will see in the next subsection.

If the utility function is not decomposed as in (19), we can make the decision model CG-decomposable by adding some dummy edges to the original graph.

2.3 CG-calculable algorithm for EQM models

When we use the RSCEP for EQM models, the integrations (18) with respect to discrete variables are not necessarily CG-calculable. In this subsection, we assume that the EQM model is CG-decomposable. The existence of a CG-calculable algorithm for EQM models is proven under the assumptions "strong root for decision" and "adequate sequence" defined below. These assumptions are satisfied when all the nodes are continuous. Therefore, the CG-calculable algorithm can be applied to some mixed discrete-Gaussian network models including the Gaussian networks.

First, we explain junction trees. By moralization, any directed graphs can be undirected. Therefore, we consider only undirected graphs. Let $\mathcal{G}$ be an undirected graph and $\mathcal{C}_i$ $(i = 1, \ldots, n)$ be sets of nodes in $\mathcal{G}$ such that $\bigcup_{i=1}^{n} \mathcal{C}_i = \mathcal{G}$. Consider a tree whose nodes are $\mathcal{C}_i$ $(i = 1, \ldots, n)$. We term the tree as *a junction tree* of $\mathcal{G}$ if every node $\mathcal{C}_k$ on the path between every pair of nodes $\mathcal{C}_i$ and $\mathcal{C}_j$ satisfies

$$(\mathcal{C}_i \cap \mathcal{C}_j) \subset \mathcal{C}_k.$$

It has been proven that for every triangulated graph $\mathcal{G}$, there is a junction tree $\mathcal{T}$ of $\mathcal{G}$ such that all the nodes of $\mathcal{T}$ are the cliques of $\mathcal{G}$. See [2] for details.

In Lauritzen (1992) [5], the notion of junction tree with strong roots is introduced.

Definition 2.6 (Strong root [5])

A node R of a junction tree is a *strong root if any pair A, B of neighbors on the junction tree with A closer to R than B satisfies*

$$(B \backslash A) \subset \Gamma \text{ or } (B \cap A) \subset \Delta.$$

For each clique C, ϕ_C^* represents the *CG-potential with an observation* $Y^* = y^*$ defined by

$$\phi_C^* := \phi_C I_C(Y^* = y^*),$$

where

$$I_C(Y^* = y^*) := \begin{cases} 1 & \text{if the value of } Y^* \cap C \text{ is consistent with } y^*, \\ 0 & \text{otherwise.} \end{cases}$$

For each separator S, ϕ_S^* is defined in the same manner.

Lemma 2.7 *(i) Assume that a mixed discrete-Gaussian network has a CG-decomposable distribution. Then for each clique $\tilde{C} \in \mathcal{C}$,*

$$\begin{aligned} p(\tilde{C}|y^*) &= \int p(x|y^*) d\mu(x \backslash (\tilde{C} \cup Y^*)) \\ &\propto \int \frac{\prod_{C \in \mathcal{C}} \phi_C^*}{\prod_{S \in \mathcal{S}} \phi_S^*} d\mu(x \backslash (\tilde{C} \cup Y^*)). \end{aligned} \tag{21}$$

(ii) Assume that after a triangulation, the junction tree has at least one strong root. Then, there is a CG-calculable algorithm for evaluating the right-hand side of (21) for each clique $\tilde{C}$.

Proof. The first part of the lemma is obvious. From theorem 7.17 of [5], the right-hand side of (21) is calculated by the Collect Evidence and Distribute Evidence algorithm in [5]. The algorithm is CG-calculable. □

Definition 2.8 (Strong root for decision)

Let T_1 be the smallest t such that d_t or y_t is continuous. A node R of a junction tree is a strong root for decision if any pair A, B of neighbors on the tree with A closer to R than B satisfies

$$(B \backslash A) \subset (\Gamma \cup d_{0:T_1-1} \cup y_{1:T_1-1}) \text{ or } (B \cap A) \subset \Delta.$$

Definition 2.9 (Adequate sequence of decisions and observations)

We say that a sequence of the decisions and observations (13) is adequate if it satisfies the following two conditions:

(i) there is no continuous observation between two discrete decisions, i.e.

$$t_1 < t \leq t_2 \text{ and } d_{t_1}, d_{t_2} \in \Delta \Rightarrow y_t \notin \Gamma,$$

(ii) there is no continuous decision before any discrete decisions, i.e.

$$t_1 < t_2 \text{ and } d_{t_2} \in \Delta \Rightarrow d_{t_1} \notin \Gamma.$$

Lemma 2.10 *Assume that the sequence of decisions and observations (13) is adequate. Assume that there is at least one discrete decision node. Let T_0 be the smallest t such that d_t is discrete. Let T_1 be the smallest t such that $t \geq T_0 + 1$, and d_t or y_t is continuous, if such a t exists. If such a t does not exist, set $T_1 = h$.*

Then, the sequence is divided into the following three parts:

(A) $d_0, y_1, \ldots, y_{T_0}$ *whose decisions are all empty.*

(B) $d_{T_0}, y_{T_0+1}, \ldots, d_{T_1-1}$ *that includes no continuous variable.*

(C) $y_{T_1}, d_{T_1}, \ldots, y_h$ *that includes no discrete decision.*

Proof. From the definition of T_0 and T_1, $T_0 \leq T_1$:

T_0 is the smallest t such that d_t is a discrete decision. Thus, from definition (ii) of the adequate sequence, part (A) of the sequence includes no nonempty decision.

From the definition of T_1, part (B) includes no continuous observation or decision.

Assume that there is a discrete decision d_t in part (C). From the definition of T_1, y_{T_1} or d_{T_1} is continuous. If y_{T_1} is continuous, y_{T_1} is a continuous observation between discrete decisions d_{T_0} and d_t. This contradicts definition (i) in Definition 2.9 of an adequate sequence. If d_{T_1} is continuous, d_{T_1} is a continuous decision before the discrete decision d_t. This contradicts definition (ii) in Definition 2.9 of an adequate sequence. Thus there is no discrete decision in part (C). □

We propose a CG-calculable algorithm for EQM models under the following five conditions.

- (Condition 1): The risk-sensitive parameter θ satisfies $\sup_{\mathcal{D}} \gamma_{\mathcal{D}}(\theta) < \infty$.
- (Condition 2): The sequence of decisions and observations (13) is adequate.

After a triangulation of the graph,

- (Condition 3): The junction tree has at least one strong root for decision.
- (Condition 4): The decision model is CG-decomposable.
- (Condition 5): For each continuous decision d_t, there is a clique C that contains d_t and contains neither discrete nuisance nodes nor discrete observation nodes after t. In other words,

$$\forall d_t \in (D \cap \Gamma), \exists C,\ d_t \in C,\ \text{and } (C \cap \Delta) \subset (D \cup y_{1:t}).$$

Algorithm 2.11

Step 1: If there is at least one discrete decision,
$T_0 \leftarrow \min\{t \mid d_t \in \Delta\}$
$T_1 \leftarrow \min\{t, h \mid d_t \in \Gamma \text{ or } y_t \in \Gamma\}$
else
$T_0 \leftarrow 0$, $T_1 \leftarrow 0$
go to step 6.

Step 2: $y_{1:T_0} \leftarrow$ observation value $y^*_{1:T_0}$

Step 3: For $d_{T_0:T_1-1}, y_{T_0+1:T_1-1} \leftarrow$ each discrete value $\tilde{d}_{T_0:T_1-1}, \tilde{y}_{T_0+1:T_1-1}$,
3.1: $\phi^*_C, \phi^{(v)*}_C, \phi^*_S$ and $\phi^{(v)*}_S$
$\leftarrow \phi_C, \phi^{(v)}_C, \phi_S$ and $\phi^{(v)}_S$ with $\tilde{d}_{T_0:T_1-1}, y^*_{1:T_0}$ and $\tilde{y}_{T_0+1:T_1-1}$.
3.2:

$$\tilde{V}_{T_1-1}(\tilde{d}_{T_0:T_1-1}, y^*_{1:T_0}, \tilde{y}_{T_0+1:T_1-1}) \leftarrow \int \frac{\prod_{C\in\mathcal{C}} \phi^*_C \phi^{(v)*}_C}{\prod_{S\in\mathcal{S}} \phi^*_S \phi^{(v)*}_S} d\mu(N \cup d_{T_1:h-1} \cup y_{T_1:h}) \qquad (22)$$

by Collect Evidence and Distribute Evidence [5].

Step 4: For $t \leftarrow T_1 - 1, T_1 - 2, \ldots, T_0 + 1$,

$$\tilde{V}_{t-1}(\tilde{d}_{T_0:t-1}, y^*_{1:T_0}, \tilde{y}_{T_0+1:t-1}) \leftarrow \sum_{\tilde{y}_t} \max_{\tilde{d}_t} \tilde{V}_t(\tilde{d}_{T_0:t}, y^*_{1:T_0}, \tilde{y}_{T_0+1:t}).$$

Step5: Optimization of d_t ($T_0 \le t < T_1$):[2]
for $t \leftarrow T_0, T_0 + 1, \ldots, T_1 - 1$
(an observation value $y^*_{1:t}$ is given)

$$\hat{d}_t \leftarrow \arg\max_{d_t} \tilde{V}_t(\hat{d}_{T_0:t-1}, d_t, y^*_{1:t}).$$

[2]As stated in section 2.2, the optimal decision $\hat{d}_t$ for each $t = 0, \ldots, h-1$ may not be unique. If we select an arbitrary optimal decision for each t, $(\hat{d}_0, \ldots, \hat{d}_{h-1})$ becomes an optimal decision sequence.

Step 6: Optimization of d_t ($t \geq T_1$):

for $t \leftarrow T_1, \ldots, h-1$,

(an observation value $y^*_{1:t}$ and the optimal decision $\hat{d}_{T_0:t-1}$ are given)

6.1: $\phi^*_C, \phi^{(v)*}_C, \phi^*_S$, and $\phi^{(v)*}_S$
$\leftarrow \phi_C, \phi^{(v)}_C, \phi_S$, and $\phi^{(v)}_S$ with $\hat{d}_{T_0:t-1}$ and $y^*_{1:t}$.

6.2: Select a clique $\tilde{C} \ni d_t$ such that $(\tilde{C} \cup \Delta) \subset (D \cup y_{1:t})$,

$$\tilde{V}^*_t(\tilde{C}) \leftarrow \int \frac{\prod_{C \in \mathcal{C}} \phi^*_C \phi^{(v)*}_C}{\prod_{S \in \mathcal{S}} \phi^*_S \phi^{(v)*}_S} d\mu((N \cup d_{t+1:h-1} \cup y_{t+1:h}) \backslash \tilde{C})$$

by Collect Evidence and Distribute Evidence [5].

6.3:

$$\tilde{V}_t(d^*_{0:t-1}, d_t, y^*_{1:t}) \leftarrow \int \tilde{V}^*_t(\tilde{C}) d\mu((N \cup d_{t+1:h-1} \cup y_{t+1:h}) \cap \tilde{C}). \tag{23}$$

6.4:

$$\hat{d}_t \leftarrow \arg\max_{d_t} \tilde{V}_t(d^*_{0:t-1}, d_t, y^*_{1:t}).$$

Theorem 2.12 *Assume (condition 1-5), then $\hat{d}_t$ $(t = 0, \ldots, h-1)$ by algorithm 2.11 is the optimal decision. Algorithm 2.11 is CG-calculable.*

Proof. The proof consists of the following four parts:

(a) optimality of $\hat{d}_t$.

(b) CG-calculability of Step 3.2 and 6.2.

(c) CG-calculability of Step 4.

(d) CG-calculability of Step 6.3 and 6.4.

(a) Proof of optimality of $\hat{d}_t$

We first prove that $\tilde{V}_t(d^*_{0:t}, y^*_{1:t})$ calculated by (22) is proportional to the value function $V_t(d^*_{0:t}, y^*_{1:t})$ defined by (18) for $t \geq T_1 - 1$. We abbreviate the conditional density function $p(\cdot | d^*_{0:t}, y^*_{1:t})$ to $p^*(\cdot)$.

For $t \geq T_1 - 1$, all $d_{t+1:h-1}$ are not discrete (continuous or empty). If we represent the summation with respect to discrete observation $y \in Y \cup \Delta$ as $\int_y$ instead of $\sum_y$,

$$
\begin{aligned}
&V_t(d^*_{0:t}, y^*_{1:t}) \\
&= \int_{y_{t+1}} p^*(y_{t+1}) \max_{d_{t+1}} \int_{y_{t+2}} p^*(y_{t+2}|d_{t+1}, y_{t+1}) \dots \\
&\quad \max_{d_{h-1}} \int_{y_h} p^*(N|d_{t+1:h-1}, y_{t+1:h}) p^*(y_h|d_{t+1:h-1}, y_{t+1:h-1}) d\mu(N \cup y_{t+1:h}) \\
&\propto \int_{y_{t+1}} p^*(y_{t+1}) \int_{d_{t+1}} \int_{y_{t+2}} p^*(y_{t+2}|d_{t+1}, y_{t+1}) \dots \\
&\quad \int_{d_{h-1}} \int_{y_h} \int_N p^*(N|d_{t+1:h-1}, y_{t+1:h}) p^*(y_h|d_{t+1:h-1}, y_{t+1:h-1}) d\mu(N \cup d_{t+1:h-1} \cup y_{t+1:h}) \\
&= \int_{y_{t+1:h}} \int_{d_{t+1:h-1}} \int_N p^*(y_{t+1:h}, N|d_{t+1:h}) d\mu(N \cup d_{t+1:h-1} \cup y_{t+1:h}) \\
&= \tilde{V}_t(d^*_{0:t}, y^*_{1:t}). \qquad (24)
\end{aligned}
$$

Therefore, for $t \geq T_1$, $\hat{d}_t$ in Step 6.4 is the optimal decision.
For $T_0 \leq t < T_1$, the optimal decision is

$$
\hat{d}_t = \arg\max_{d^*_t} \sum_{y^*_{t+1}} \max_{d^*_{t+1}} \cdots \sum_{y^*_{T_1-1}} \max_{d^*_{T_1-1}} \tilde{V}_{T_1-1}(d^*_{0:T_1-1}, y^*_{1:T_1-1}).
$$

Step 4 and Step 5 calculate this $\hat{d}_t$.

(b) Proof of CG-calculability of Step 3.2 and 6.2.
Here $d^*_{0:t-1}$ and $y^*_{1:t}$ are fixed. Thus, the strong root for decision (Condition 3) corresponds to the normal strong root (Definition 2.6). Set $\tilde{\phi}^*_C := \phi^*_C \phi^{(v)*}_C$ and $\tilde{\phi}^*_S := \phi^*_S \phi^{(v)*}_S$, then $\tilde{\phi}_C$ and $\tilde{\phi}_S$ take the CG-potential form (8).
Then, the same algorithm in lemma 2.7 is applied to the calculation of

$$
\tilde{V}_t(\tilde{C}) = \int \frac{\prod_{C \in \mathcal{C}} \tilde{\phi}^*_C}{\prod_{S \in \mathcal{S}} \tilde{\phi}^*_S} d\mu(x \backslash (\tilde{C} \cup Y^*))
$$

for $t \geq T_1 - 1$. Therefore, Step 6.2 is CG-calculable.
In the same way, calculate $\tilde{V}_{T_1-1}(R)$ for the strong root $R \in \mathcal{C}$. Then,

$$
\tilde{V}_{T_1-1}(d^*_{0:T_1-1}, y^*_{1:T_1-1}) \propto \int \tilde{V}_{T_1-1}(R) d\mu(R),
$$

and the integration is CG-calculable if the integration with respect to $R \cap \Gamma$ is first, and $R \cap \Delta$ is second. Therefore, Step 3.2 is CG-calculable.

(c) Proof of CG-calculability of Step 4
All of nonempty $d_{T_0:T_1-1}$ and $y_{T_0:T_1-1}$ are discrete. Thus, $\tilde{V}_{T_1}(d^*_{0:T_1-1}, y^*_{1:T_1-1})$ is a scalar function of discrete variables. Maximization and summation of such a function are just comparisons and summations of a finite number of scalar values, respectively. Therefore, Step 4 is CG-calculable.

(d) Proof of CG-calculability of Step 6.3 and 6.4.
From (Condition 5), there is no discrete nuisance variable or discrete observation after t in clique $\tilde{C}$. From the definition of T_1, for $t \geq T_1$, $d_{t+1:h-1}$ are continuous or empty. Thus,

$$((N \cup d_{t+1:h-1} \cup y_{t+1:h}) \cap \tilde{C}) \subset \Gamma.$$

Therefore, integration in (23) is CG-calculable by lemma A.1. Then, Step 6.3 is CG-calculable.

Maximization with respect to the continuous variable d_t is CG-calculable by lemma A.2. Then, Step 6.4 is CG-calculable. □

If the model is described as a DAG with CG-distribution, there is a simpler sufficient condition for CG-calculability.

Corollary 2.13 *Consider the EQM(θ) model represented as a DAG. Assume the following four conditions, then algorithm 2.11 is CG-calculable.*

- *(Condition 1 for the DAG): The risk sensitive parameter θ satisfies $\sup_{\mathcal{D}} \gamma_{\mathcal{D}}(\theta) < \infty$.*
- *(Condition 2 for the DAG): The sequence of decisions and observations is adequate.*
- *(Condition 3 for the DAG): The total utility is decomposed as $Q(x) = \sum_{v \in V} q_v(v, pa(v))$, where q_v is a function defined for each node v.*
- *(Condition 4 for the DAG): Every discrete node with some continuous children is a decision or an observation node until the first continuous observation time T_1, i.e.,*

$$(\Delta \cap pa(\Gamma)) \subset (D \cup y_{1:T_1}).$$

3 H_∞^Δ optimal decision networks

In this section, we propose H_∞^Δ optimal decision networks. The H_∞^Δ optimal decision networks are a generalization of H_∞ optimal controls. H_∞^Δ optimal decision networks have robustness similar to that of H_∞ optimal controls. Details of the H_∞ optimal control theory are given in Francis (1987) [3]. The main result of this section is that the optimal decision for each H_∞^Δ optimal decision network is calculated as a limit of the optimal decision for the corresponding EQM models.

We first define a *linear decision*. We describe all the continuous and discrete variables in an EQM model as $\zeta = (\zeta_\gamma)_{\gamma \in \Gamma}$ and $i = (i_\delta)_{\delta \in \Delta}$, respectively. The decision values $(\zeta_\gamma)_{\gamma \in D \cap \Gamma}$ and $(i_\delta)_{\delta \in D \cap \Delta}$ are the optimal decision determined by a decision policy, $\mathcal{D}$, and an observation, Y.

Definition 3.1 (Linear decision)

If a decision policy, $\mathcal{D}$, is fixed, the conditional distribution of all continuous variables, ζ, given all discrete variables, i, is uniquely determined. Denote the conditional distribution by $p(\zeta|i, \mathcal{D})$. For some $\mathcal{D}$, there exists $p \times q$ $(p \geq q)$-dimensional matrix, $H(i)$, and a q-dimensional vector, $\mu(i)$, such that $p(\zeta|i, \mathcal{D})$ corresponds to

$$\zeta = H(i)(\xi + \mu(i)), \tag{25}$$

where ξ is a random vector distributed according to $N(0, I_q)$. We say that such a decision policy, $\mathcal{D}$, is a linear decision.

We note that every optimal decision policy for $EQM(\theta)$ is a linear decision. This is because when all discrete variables are given, the optimal decision for each continuous decision variable is a linear function of the continuous observation and nuisance variables.

In (condition 1) of theorem 2.12, we assumed that the risk-sensitive parameter θ satisfies $\sup_{\mathcal{D}} \gamma_{\mathcal{D}}(\theta) < \infty$. The following lemmas give a more explicit representation of (condition 1).

Lemma 3.2 *Assume that a linear decision policy, $\mathcal{D}$, is represented as (25). The risk sensitive criterion $\gamma_{\mathcal{D}}(\theta)$ is finite if and only if $I_q + \theta H(i)' K^{(v)}(i) H(i)$ is positive definite for any discrete value i.*

This condition is described as $\theta > \theta_0(\mathcal{D})$ for a negative constant $\theta_0(\mathcal{D})$.

Lemma 3.3 *In an EQM(θ) model, (Condition 1) in theorem 2.12 is satisfied if and only if $I_q + \theta H(i)' K^{(v)}(i) H(i)$ is positive-definite for any discrete value i and any linear decision policy $\mathcal{D}$.*

This condition is described as $\theta > \theta_0$ for a negative constant θ_0.

The proof is omitted. In this sense, if θ approaches θ_0, EQM(θ) nearly becomes the most pessimistic model.

We define the H_∞^Δ norm and the H_∞^Δ optimal decision.

Definition 3.4 (H_∞^Δ norm)

Let $G(i)$ be an arbitrary matrix-valued function of finite discrete variable i. H_∞^Δ norm $|| \cdot ||_\infty^\Delta$ is defined as

$$||G||_\infty^\Delta := \max_i \sqrt{\text{The maximum eigen value of } G(i)'G(i)}.$$

In the following, we consider $G(i)$ defined as

$$G(i) := K^{(v)}(i)^{1/2} H(i).$$

The H_∞^Δ optimal decision is defined as follows:

Definition 3.5 (H_∞^Δ optimal decision)

If a linear decision minimizes $||G||_\infty^\Delta$, we say it is an H_∞^Δ optimal decision.

If a decision network is a state-space model described as (1) and (2), the H_∞^Δ optimal decision corresponds to the H_∞ optimal control. Thus, the H_∞^Δ optimal decision is an extension of the H_∞ optimal control.

The next theorem shows an important relation between the H_∞^Δ optimal decision and EQM(θ) models.

Theorem 3.6 *Consider a decision network model with CG-distribution and the corresponding EQM(θ) model. The H_∞^Δ optimal decisions for a CG-distributed model is equivalent to the optimal decisions $\{\mathcal{D}\}$ for the EQM(θ_0) model.*

Proof. Let $\mathcal{D}$ be a linear decision,

$$\begin{aligned}\gamma_{\mathcal{D}}(\theta) < \infty \Leftrightarrow \max_i \int N(\xi; \mu(i), I_q) \exp(\theta(g^{(v)}(i)\\ + h^{(v)}(i)' H(i)\xi - \xi' H(i)' K^{(v)}(i) H(i)\xi/2)) d\mu(\xi) < \infty\\ \Leftrightarrow \theta > \max_i \sup\{\theta || I_q + \theta H(i)' K^{(v)}(i) H(i)| = 0\}.\end{aligned} \quad (26)$$

From the definition of $\theta_0(\mathcal{D})$,

$$\begin{aligned}\theta_0(\mathcal{D}) &= \max_i \sup\{\theta || I_q + \theta H(i)' K^{(v)}(i) H(i)| = 0\}\\ &= \max_i \sup\{\theta || I_q + \theta G(i)' G(i)| = 0\}\\ &= -(||G||_\infty^\Delta)^{-1}.\end{aligned} \quad (27)$$

From (26) and (27),

$$\begin{aligned}&\mathcal{D} \in \text{optimal decision for } EQM(\theta_0)\\ &\quad \Leftrightarrow \gamma_{\mathcal{D}}(\theta_0) = \infty\\ &\quad \Leftrightarrow \theta_0(\mathcal{D}) \geq \theta_0 = \inf_{\mathcal{D}} \theta_0(\mathcal{D})\\ &\quad \Leftrightarrow \mathcal{D} = \arg\min \theta_0(\mathcal{D})\\ &\quad \Leftrightarrow \mathcal{D} = \arg\min ||G||_\infty^\Delta\\ &\quad \Leftrightarrow \mathcal{D} \in H_\infty^\Delta \text{ optimal decision.}\end{aligned}$$

□

From theorem 3.7 given below, we can construct a calculation algorithm for exact H_∞^Δ optimal decisions. Using theorem 3.6, we can calculate an approximation of an H_∞^Δ optimal decision by calculating the optimal decision for EQM($\theta_0 + \epsilon$), where ϵ is a small positive value. However, this method can not be applied directly to the exact evaluation because algorithm 2.11 is valid only when $\gamma_{\mathcal{D}}(\theta) < \infty$ (condition 1). Moreover, in certain cases of integration or maximization as in lemmas A.1 or A.2, $K_{11}(i) + \theta K_{11}^{(v)}(i)$ becomes almost singular when θ approaches θ_0. Thus, numerical computation becomes unreliable.

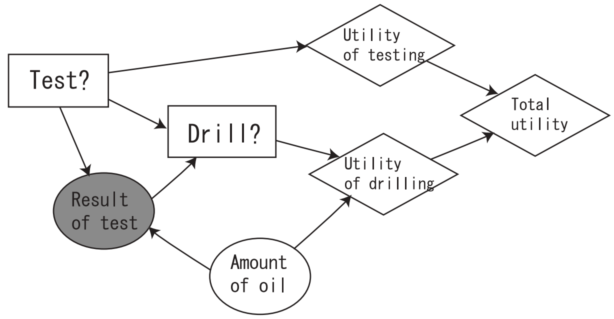

Figure 1: The influence diagram of Oil Wildcatter problem.

Theorem 3.7 *If we use (46), (48), and (49) in lemmas A.4 and A.5 instead of (40), (42), and (43) in lemmas A.1 and A.2, the optimal decision for EQM(θ_0) model is H_∞^Δ optimal.*

Proof. Let the optimal decision for the EQM(θ_0) model be $\hat{\mathcal{D}}$. Because $\gamma_{\hat{\mathcal{D}}}(\theta_0) = \infty$, it is proved that $\hat{\mathcal{D}}$ is H_∞^Δ optimal in the same manner of the proof of theorem 3.6. □

4 Experimental result

4.1 Oil Wildcatter (discrete case)

The oil wildcatter problem is a widely known example in decision analysis, posed by Raiffa (1968) [9]. We deal with the variation given by Shenoy (1992) [11] as follows:

> An oil wildcatter must decide either to drill (yes) or not drill (no). He is uncertain whether the hole is dry (Dry), wet (Wet) or soaking (Soak).... At a cost of 10,000, the wildcatter could take seismic soundings which help determine the geological structure at the site. The soundings will disclose whether the terrain below has no structure (NoS) – that's bad, or open structure (OpS) – that's so-so, or closed structure (ClS) – that's really hopeful.

The influence diagram of the oil wildcatter problem is shown in Fig. 1. In the influence diagram, white oval nodes, shaded oval nodes, and rectangular nodes represent the nuisance, observation, and decision variables, respectively. In order to represent the relation between each node and the utility value, we additionally set the diamond utility nodes. The influence diagram for the oil wildcatter problem has seven nodes in total: two utility nodes U_T (Utility of testing) and U_D (Utility of drilling), and the total utility node represented as the sum of two components of utility; two decision nodes, T (Test?) and D (Drill?); a nuisance node, O (Amount of oil); and an observation node, R (Result of test).

Table 1: Four decision policies of different degree of optimism (Oil Wildcatter example).

	Region of θ	Test?	Drill? NoS	OpS	ClS
A	$-1000 \sim -800$	no		no	
B	-750	yes	no	no	yes
C	$-700 \sim 150$	yes	no	yes	yes
D	$200 \sim 1000$	no		yes	

We use the same conditional probabilities and utility functions as [2]:

$$P(O = \text{Dry}, \text{Wet}, \text{Soak}) = 0.5, 0.3, 0.2. \tag{28}$$

$$P(R|O) = \left[\begin{array}{c|ccc} O & \text{NoS} & \text{OpS} & \text{ClS} \\ \hline \text{Dry} & 0.6 & 0.3 & 0.1 \\ \text{Wet} & 0.3 & 0.4 & 0.3 \\ \text{Soak} & 0.1 & 0.4 & 0.5 \end{array}\right], \tag{29}$$

$$U(T = \text{yes}, \text{no}) = -10, 0, \text{ and} \tag{30}$$

$$U(O, D) = \left[\begin{array}{c|ccc} D & \text{Dry} & \text{Wet} & \text{Soak} \\ \hline \text{yes} & -70 & 50 & 200 \\ \text{no} & 0 & 0 & 0 \end{array}\right]. \tag{31}$$

We maximize the criterion

$$\gamma_{\mathcal{D}}(\theta) = \theta^{-1} \log E_{\mathcal{D}} \exp(\theta(U(T) + U(O, D))) \tag{32}$$

for each $\theta \in \{-1000 + 50i \mid i = 0, \ldots, 40\}$. The result is listed in Table 1: there are four different optimal decision policies (A, B, C, and D) depending on the value of θ. You can observe that the policies become more optimistic in the order of A, B, C, and D. For example, policy A, with the minimum θ values, does not test or drill because the testing and drilling seem like a loss of cost for the most pessimistic decision-maker. On the contrary, policy D, with the maximum θ values, does not test but drills. This is because the most optimistic decision-maker believes that the result of drilling must be the best and the test is only a loss of cost.

In some competitive situations, as follows, risk-sensitivity plays an important role. We assume that there are ten oil wildcatters including you. The nine wildcatters other than you independently take one of the decision policies A, B, C, or D with probability 1/4. We know the probability but do not know the decision policy or the observation result of the other wildcatters. The objective is to find the best decision policy.

Distributions of the rank of the utility value in the case of each decision policy are shown in Fig. 2. If the number of survivors is one in ten wildcatters, the most optimistic policy, A, is much better than other policies. In contrast, if only the last one loses the competition, the most pessimistic policy, D, presents the largest probability of winning.

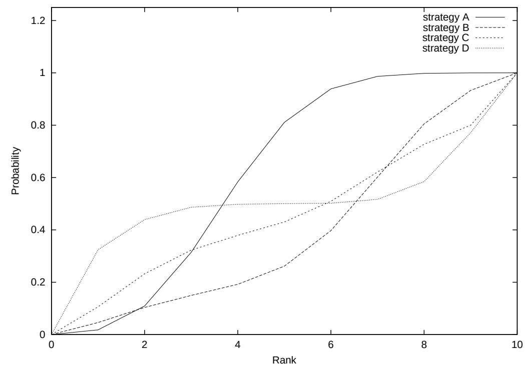

Figure 2: Distributions of the ranking by the four decision policies (A, B, C, D).

4.2 Waste Incinerator

The waste incinerator example is a fictious example of a conditional Gaussian model. The original problem is illustrated as follows [2]:

> The emission from a waste incinerator differs because of compositional differences in incoming waste. Another important factor is the waste burning regime which can be monitored by measuring the concentration of CO2 in the emission. The filter efficiency depends on the technical state of the electrofilter and the amount and composition of waste. The emission of heavy metal depends both on the concentration of metal in the incoming waste and the emission of dust particulates in general. The emission of dust is monitored through measuring the penetrability of light.

The variables in the model are Type of waste (W), Filter state (F), Burning regime (B), Metal in waste (M_{in}), Filter efficiency (E), CO2 emission (C), Emission of dust (D), Emission of metal (M_{out}), and Light penetrability (L).

The conditional probability structure in the model is represented as a DAG in Fig. 3. As in Fig. 1, the oval, rectangular and diamond nodes represent nuisance, decision, and utility nodes, respectively. The symbol "$*$" near a node indicates that the node is discrete.

The original problem is not a decision problem but an estimation problem. We assume that three discrete nodes (Type of waste, Filter state, and Burning regime) are decision nodes. We set utility values for these three discrete nodes and other two continuous nodes (Filter efficiency and Emission of metal). The total utility value is the sum of the five utility values. The optimal decision does not depend on the order of decisions because there is no observation variable in the model.

The probability structure of the model is the same as [2]. Burning regime (B) takes two values $\{stabe, unstable\}$. If $B = stable$, we denote $B = 1$, and if $B = unstable$, we denote $B = 2$, as a short form. Similarly, $F = intact$ or $defective$ is represented as $F = 1$ or 2, respectively, and $W = industrial$ or $household$ is represented as $W = 1$ or 2,

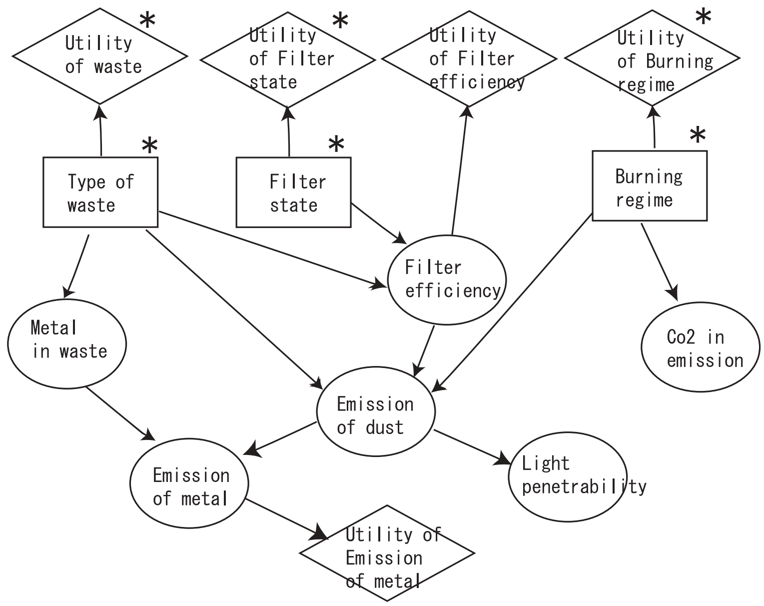

Figure 3: Influence diagram of the waste incinerator example.

respectively. The conditional density of the Filter efficiency (E) is defined as

$$\begin{aligned} p(E|F=1,W=1) &= N(-3.9, 0.00002), \\ p(E|F=1,W=2) &= N(-3.2, 0.00002), \\ p(E|F=2,W=1) &= N(-0.4, 0.0001), \text{ and} \\ p(E|F=2,W=2) &= N(-0.5, 0.0001). \end{aligned}$$

The conditional density of the Emission of dust (D) is defined as

$$\begin{aligned} p(D|B=1,W=1,E=e) &= N(6.5+e, 0.03), \\ p(D|B=1,W=2,E=e) &= N(6.0+e, 0.04), \\ p(D|B=2,W=1,E=e) &= N(7.5+e, 0.1), \text{ and} \\ p(D|B=2,W=2,E=e) &= N(7.0+e, 0.1). \end{aligned}$$

The conditional density of the concentration of CO2 is defined as

$$p(C|B=1) = N(-2, 0.1), \quad p(C|B=2) = N(-1, 0.3).$$

The conditional density of the penetrability of dust is defined as

$$p(L|D=d) = N(3-d/2, 0.25).$$

The conditional density of the concentration of heavy metal in the waste (M_{in}) is defined as

$$p(M_{in}|W=1) = N(0.5, 0.01), \quad p(M_{in}|W=2) = N(-0.5, 0.005).$$

Table 2: Four decision policies with different degrees of optimism (the waste incinerator example).

	Region of θ	Burning (B)	Filter (F)	Waste (W)
-	$-9.9 \sim -9.0$	Too pessimistic		
A	$-8.9 \sim -7.3$	$stable$ (1)	$intact$ (1)	$household$ (2)
B	$-7.2 \sim 0.4$	$stable$ (1)	$defective$ (2)	$household$ (2)
C	$0.5 \sim 9.5$	$instable$ (2)	$defective$ (2)	$household$ (2)
D	$9.6 \sim 15.0$	$instable$ (2)	$defective$ (2)	$industrial$ (1)

The conditional density of the concentration of emission of metal is defined as

$$p(M_{out}|D = d, M_{in} = m_{in}) = N(d + m_{in}, 0.002).$$

Here, the utility value is set as

$$U(B = 1, 2) = 0.0, 5.0, \quad (33)$$

$$U(F = 1, 2) = 0.0, 10.0, \quad (34)$$

$$U(W = 1, 2) = 10.0, 5.0, \text{ and} \quad (35)$$

$$U(E = e) = -e^2 \text{ and } U(M_{out} = m_{out}) = -m_{out}^2.$$

We maximize the criterion

$$\gamma_{\mathcal{D}}(\theta) = \theta^{-1} \log E_{\mathcal{D}} \exp[\theta(U(B) + U(F) + U(W) + U(E) + U(M_{out}))] \quad (36)$$

for each $\theta \in \{-10+0.1i \mid i = 1, \ldots, 250\}$. This example satisfies the condition in corollary 2.13. Therefore, Algorithm 2.11 is CG-calculable for the moralized and triangulated model.

The result is given in Table 2: there are four different optimal decision policies depending on the value of θ. From theorem 3.7, the most pessimistic value is $\theta_0 = -8.92059949615108$ and the H_∞^Δ optimal decision is $B = 1$ (*stable*), $F = 1$ (*intact*), and $W = 2$ (*household*).

In this decision network, all decision variables are discrete. Next, we deal with a model with a continuous decision node. We set the Filter efficiency node (E) as a decision node and remove all the arrows into this node. Thus, the influence diagram is as shown in Fig. 4.

The result of the optimization is listed in Table 3 and the graph is as shown in Fig. 5. Here, the most pessimistic value is $\theta_0 = -8.92857142857142$, and the H_∞^Δ optimal decision is $B = 1$ (*stable*), $F = 2$ ($defective$), $W = 2$ (*household*) and $E = -3.48022598870056$.

Table 3: Decision policies with different degrees of optimism (the waste incinerator example with a continuous decision node).

θ	Burning (B)	Filter (F)	Waste (W)	Efficiency (E)
−9.0000	Too pessimistic			
−8.0000	*stable* (1)	*defective* (2)	*household* (2)	−3.3867
−7.0000	*stable* (1)	*defective* (2)	*household* (2)	−3.2914
−6.0000	*stable* (1)	*defective* (2)	*household* (2)	−3.2014
−5.0000	*stable* (1)	*defective* (2)	*household* (2)	−3.1161
−4.0000	*stable* (1)	*defective* (2)	*household* (2)	−3.0353
−3.0000	*stable* (1)	*defective* (2)	*household* (2)	−2.9586
−2.0000	*stable* (1)	*defective* (2)	*household* (2)	−2.8856
−1.0000	*stable* (1)	*defective* (2)	*household* (2)	−2.8162
0.0000	*stable* (1)	*defective* (2)	*household* (2)	−2.7498
1.0000	*instable* (2)	*defective* (2)	*household* (2)	−3.0850
2.0000	*instable* (2)	*defective* (2)	*household* (2)	−2.9359
3.0000	*instable* (2)	*defective* (2)	*household* (2)	−2.8005
4.0000	*instable* (2)	*defective* (2)	*household* (2)	−2.6771
5.0000	*instable* (2)	*defective* (2)	*household* (2)	−2.5641
6.0000	*instable* (2)	*defective* (2)	*household* (2)	−2.4603
7.0000	*instable* (2)	*defective* (2)	*household* (2)	−2.3645
8.0000	*instable* (2)	*defective* (2)	*household* (2)	−2.2759
9.0000	*instable* (2)	*defective* (2)	*household* (2)	−2.1937
10.0000	*instable* (2)	*defective* (2)	*industrial* (1)	−2.5641
11.0000	*instable* (2)	*defective* (2)	*industrial* (1)	−2.4752
12.0000	*instable* (2)	*defective* (2)	*industrial* (1)	−2.3923
13.0000	*instable* (2)	*defective* (2)	*industrial* (1)	−2.3148
14.0000	*instable* (2)	*defective* (2)	*industrial* (1)	−2.2422
15.0000	*instable* (2)	*defective* (2)	*industrial* (1)	−2.1739

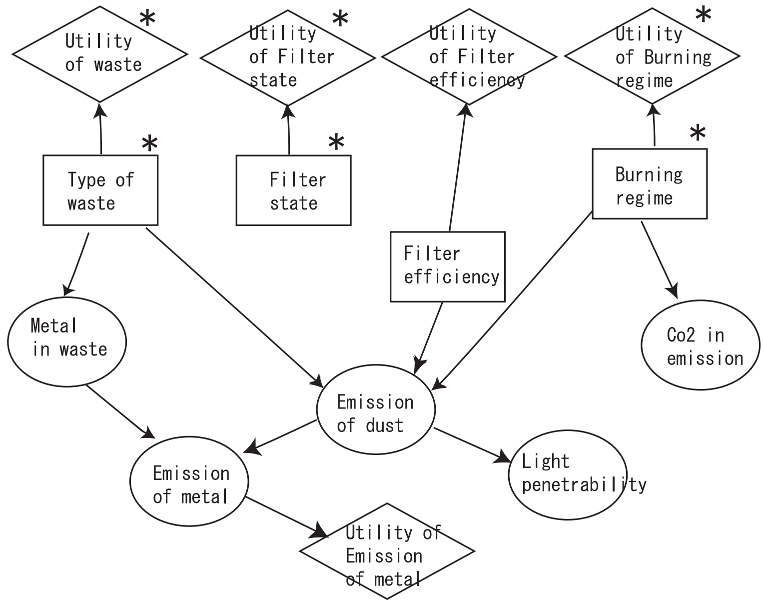

Figure 4: Influence diagram of the waste incinerator example with a continuous decision node.

5 Some discussions

This paper presents the calculation method for the optimal decision with an optimistic parameter that represents the decision maker's degree of optimism or pessimism. The exponential quadratic criterion function is used for modeling the degree of optimism.

The EQM model is introduced. This model includes both discrete and continuous nodes and has the risk-sensitive criterion function. Under several assumptions (CG-decomposability, strong root for decision, adequate sequence, etc.), we propose a CG-calculable algorithm for the optimal decision for EQM models.

There are several studies on exponential quadratic decision networks. Although Shachter and Kenley (1989) [10] present a method for decision making for Gaussian models with exponential quadratic criterion, their algorithm cannot be applied to the mixed discrete-Gaussian model.

Poland (1994) [8] deals with a type of mixed discrete-Gaussian case with the exponential quadratic criterion. In his paper, it is assumed that the mixed discrete-Gaussian node has no children other than the utility nodes. His algorithm includes some calculations on mixed normal distribution (for example, maximization by the EM algorithm), while our algorithm avoids them in order to save on computational cost. Therefore, the motivation of his research is different from ours. However, when the number of necessary calculations on a mixed normal distribution is small, the incorporation of the two algorithms should be efficient.

Our algorithm uses the graph decomposition technique while the ones in [10] and [8]

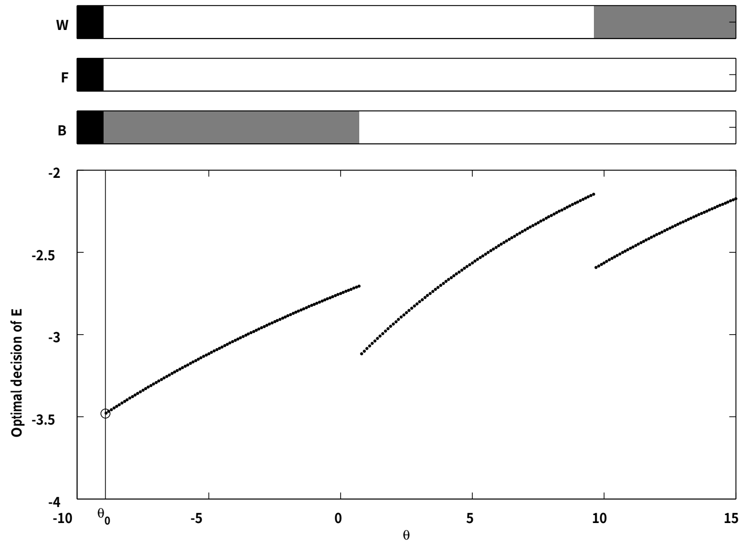

Figure 5: The optimal decision policies with different degrees of optimism for the waste incinerator example with a continuous decision node. W (Type of waste), F (Filter state), and B (Burning regime) are the discrete decisions and E (Filter efficiency) is the continuous decision. For discrete decisions, the gray part of bar graph represents decision 1 and the white part represents decision 2. The black part is the singular part because θ_0 becomes too small. The circle in the graph of E represents the H_∞^Δ decision.

use repeat "arc reversals" and "node reductions." This is another difference between their algorithms and ours.

Next, an H_∞^Δ optimal decision is proposed in this paper. The H_∞^Δ optimal decision is a generalization of the H_∞ optimal control. Glover and Doyle (1988) [4] proved that the H_∞ optimal control is equivalent to a limit of the optimal control for corresponding the LEQG models. Whittle (1990) [12]; (1996) [13] explains further details of this fact. The H_∞ optimal control is a deterministic control. On the other hand, the LEQG model is a stochastic control with the expectation of exponential quadratic utility. The close relationship between these two apparently different controls is a breakthrough.

Their argument is restricted to the control theory; this paper generalizes it to the network decision theory with discrete variables. It is proved that the H_∞^Δ optimal decision is a limit of the optimal decision for the corresponding exponential quadratic mixed discrete-Gaussian decision models. In the limit model, the optimism parameter, θ, corresponds to the most negative value, θ_0. We introduced an algorithm for deriving the exact H_∞^Δ optimal decision with a matrix decomposition technique.

In this paper, we assume that DAGs represent the definition of conditional probabilities. The theory of causality in Pearl (2000) [7] gives a different meaning to DAGs. In his theory, DAGs represent the causal relationship between variables. If we combine our work with the theory of causality, the range of applications of our research will become much wider.

A Appendix

Lemma A.1 (Lemma 7.5 of Cowell etc. (1999) [2])

Let

$$\phi(i,\zeta_1,\zeta_2)=\exp(g(i)+\begin{bmatrix}h_1(i)\\h_2(i)\end{bmatrix}'\begin{bmatrix}\zeta_1\\\zeta_2\end{bmatrix}-\frac{1}{2}\begin{bmatrix}\zeta_1\\\zeta_2\end{bmatrix}'\begin{bmatrix}K_{11}(i)&K_{12}(i)\\K_{12}(i)'&K_{22}(i)\end{bmatrix}\begin{bmatrix}\zeta_1\\\zeta_2\end{bmatrix}) \tag{37}$$

and

$$\phi_\theta^{(v)}(i,\zeta_1,\zeta_2)=\exp(\theta(g^{(v)}(i)+\begin{bmatrix}h_1^{(v)}(i)\\h_2^{(v)}(i)\end{bmatrix}'\begin{bmatrix}\zeta_1\\\zeta_2\end{bmatrix}-\frac{1}{2}\begin{bmatrix}\zeta_1\\\zeta_2\end{bmatrix}'\begin{bmatrix}K_{11}^{(v)}(i)&K_{12}^{(v)}(i)\\K_{12}^{(v)}(i)'&K_{22}^{(v)}(i)\end{bmatrix}\begin{bmatrix}\zeta_1\\\zeta_2\end{bmatrix})). \tag{38}$$

If $K_{11}(i)+\theta K_{11}^{(v)}(i)$ *is positive-definite,*

$$\int\phi(i,\zeta_1,\zeta_2)\phi_\theta^{(v)}(i,\zeta_1,\zeta_2)d\zeta_1=\exp(\tilde{g}_\theta(i)+\tilde{h}_\theta(i)'\zeta_2-\zeta_2{}'\tilde{K}_\theta(i)\zeta_2/2) \tag{39}$$

holds, where

$$\begin{aligned}\tilde{g}_\theta(i)&=g(i)+\theta g^{(v)}(i)+\{(\dim\zeta_1)\log 2\pi-\log(K_{11}(i)+\theta K_{11}^{(v)}(i))\\&\quad+(h_1(i)+\theta h_1^{(v)}(i))'(K_{11}(i)+\theta K_{11}^{(v)}(i))^{-1}(h_1(i)+\theta h_1^{(v)}(i))\}/2,\\\tilde{h}_\theta(i)&=h_2(i)+\theta h_2^{(v)}(i)-(K_{12}(i)+\theta K_{12}^{(v)}(i))'(K_{11}(i)+\theta K_{11}^{(v)}(i))^{-1}(h_1(i)+\theta h_1^{(v)}(i)),\ \text{and}\\\tilde{K}_\theta(i)&=K_{22}(i)+\theta K_{22}^{(v)}(i)-(K_{12}(i)+\theta K_{12}^{(v)}(i))'(K_{11}(i)+\theta K_{11}^{(v)}(i))^{-1}(K_{12}(i)+\theta K_{12}^{(v)}(i)).\end{aligned} \tag{40}$$

If $K_{11}(i)+\theta K_{11}^{(v)}(i)$ *is positive-definite,* $\int\phi(i,\zeta_1,\zeta_2)\phi_\theta^{(v)}(i,\zeta_1,\zeta_2)d\zeta_1$ *is infinite.*

Lemma A.2

We make the same assumption as lemma A.1. If $K_{11}(i)+\theta K_{11}^{(v)}(i)$ *is positive-definite,*

$$\max_{\zeta_1}\phi(i,\zeta_1,\zeta_2)\phi_\theta^{(v)}(i,\zeta_1,\zeta_2)=\exp(\tilde{g}_\theta(i)+\tilde{h}_\theta(i)'\zeta_2-\zeta_2{}'\tilde{K}_\theta(i)\zeta_2/2) \tag{41}$$

holds, where

$$\begin{aligned}\tilde{g}_\theta(i)&=g(i)+\theta g^{(v)}(i)+(h_1(i)+\theta h_1^{(v)}(i))'(K_{11}(i)+\theta K_{11}^{(v)}(i))^{-1}(h_1(i)+\theta h_1^{(v)}(i))/2,\\\tilde{h}_\theta(i)&=h_2(i)+\theta h_2^{(v)}(i)-(K_{12}(i)+\theta K_{12}^{(v)}(i))'(K_{11}(i)+\theta K_{11}^{(v)}(i))^{-1}(h_1(i)+\theta h_1^{(v)}(i)),\ \text{and}\\\tilde{K}_\theta(i)&=K_{22}(i)+\theta K_{22}^{(v)}(i)-(K_{12}(i)+\theta K_{12}^{(v)}(i))'(K_{11}(i)+\theta K_{11}^{(v)}(i))^{-1}(K_{12}(i)+\theta K_{12}^{(v)}(i)).\end{aligned} \tag{42}$$

The maximum is attained as

$$\hat{\zeta}_\theta=(K_{11}(i)+\theta K_{11}^{(v)}(i))^{-1}[h_1(i)+\theta h_1^{(v)}(i)-(K_{12}(i)+\theta K_{12}^{(v)}(i))\zeta_2]. \tag{43}$$

If $K_{11}(i)+\theta K_{11}^{(v)}(i)$ *is not positive-definite,* $\max_{\zeta_1}\phi(i,\zeta_1,\zeta_2)\phi_\theta^{(v)}(i,\zeta_1,\zeta_2)$ *is infinite.*

Lemma A.3 (Theorem 10.1 of Ansley and Kohn (1985) [1]) *We assume that* $\Omega^{(1)}$ *and* $\Omega^{(0)}$ *are* $d \times d$ *positive semidefinite-matrices. Thus, there are* $d \times d$ *matrices* $\Lambda(\theta)$ *and* $L(\theta)$ *such that*

$$\Omega^{(1)} + \theta\Omega^{(0)} = L(\theta)\Lambda(\theta)L(\theta)', \tag{44}$$

where $\Lambda(\theta)$ *is a diagonal matrix and* $L(\theta)$ *is a lower triangle matrix with 1s on the diagonal.*

We omit the derivation of $\Lambda(\theta)$ and $L(\theta)$. You can see an efficient algorithm of computing Λ and L in Corollary 10.2 of [1].

Lemma A.4 *Let* K_{11} *and* $K_{11}^{(v)}$ *be positive-semidefinite matrices such that* $|K_{11}+\theta_0 K_{11}^{(v)}| = 0$*. Let* $\theta > \theta_0$*. There exist* Λ *and* L *such that*

(i) $K_{11} + \theta K_{11}^{(v)} = L\Lambda L' + O(\theta - \theta_0)$.

(ii) Λ *is a diagonal matrix and* L *is a lower triangle matrix with 1s on the diagonal.*

(iii) Λ *and* L *are independent of* θ.

Let the i*-th diagonal element of* Λ *be* λ_i *and* $\check{\Lambda} := diag(\check{\lambda}_i)$*, where* $\check{\lambda}_i = \lambda_i^{-1}$ *if* $\lambda_i \neq 0$ *and* $\check{\lambda}_i = 0$ *if* $\lambda_i = 0$*. Let* d' *the number of* i *such that* $\lambda_i \neq 0$*.*

Using such $\check{\Lambda}$*,*L*, and* d'*,* $\tilde{g}_\theta, \tilde{h}_\theta$*, and* $\tilde{K}_\theta$ *in (38) are described as follows:*

$$\begin{aligned} \tilde{g}_\theta &= g^* + O(\theta - \theta_0), \\ \tilde{h}_\theta &= h^* + O(\theta - \theta_0), \text{ and} \\ \tilde{K}_\theta &= K^* + O(\theta - \theta_0), \end{aligned} \tag{45}$$

where

$$\begin{aligned} g^* &= g + \theta_0 g^{(v)} + \{d' \log 2\pi - \sum_i \log \check{\lambda}_i + (h_1 + \theta_0 h_1^{(v)})' L^{-1\prime} \check{\Lambda} L^{-1} (h_1 + \theta_0 h_1^{(v)})\}/2, \\ h^* &= h_2 + \theta_0 h_2^{(v)} - (K_{12} + \theta_0 K_{12}^{(v)})' L^{-1\prime} \check{\Lambda} L^{-1} (h_1 + \theta_0 h_1^{(v)}), \\ K^* &= K_{22} + \theta_0 K_{22}^{(v)} - (K_{12} + \theta_0 K_{12}^{(v)})' L^{-1\prime} \check{\Lambda} L^{-1} (K_{12} + \theta_0 K_{12}^{(v)}). \end{aligned} \tag{46}$$

Proof. The first part of the lemma is proved by $K_{11} + \theta K_{11}^{(v)} = (K_{11} + \theta_0 K_{11}^{(v)}) + (\theta - \theta_0) K_{11}^{(v)}$ and lemma A.3. The formulas of g^*, h^*, and K^* are derived by the same manner as the results of Section 6 of [1]. □

Lemma A.5 *Let* K_{11} *and* $K_{11}^{(v)}$ *be positive-semidefinite matrices such that* $|K_{11}+\theta_0 K_{11}^{(v)}| = 0$*. Let* $\theta > \theta_0$*. Using the same definition of* $\check{\Lambda}$*,* L*, and* d' *as in lemma A.4,* $\tilde{g}_\theta, \tilde{h}_\theta$*, and* $\tilde{K}_\theta$ *in (42) and* $\hat{\zeta}_\theta$ *in (43) are described as follows:*

$$\begin{aligned} \tilde{g}_\theta &= g^* + O(\theta - \theta_0), \\ \tilde{h}_\theta &= h^* + O(\theta - \theta_0), \\ \tilde{K}_\theta &= K^* + O(\theta - \theta_0), \text{ and} \\ \hat{\zeta}_\theta &= \zeta^* + O(\theta - \theta_0), \end{aligned} \tag{47}$$

where

$$g^* = g + \theta_0 g^{(v)} + (h_1 + \theta_0 h_1^{(v)})' L^{-1\prime} \check{\Lambda} L^{-1} (h_1 + \theta_0 h_1^{(v)})/2,$$
$$h^* = h_2 + \theta_0 h_2^{(v)} - (K_{12} + \theta_0 K_{12}^{(v)})' L^{-1\prime} \check{\Lambda} L^{-1} (h_1 + \theta_0 h_1^{(v)}),$$
$$K^* = K_{22} + \theta_0 K_{22}^{(v)} - (K_{12} + \theta_0 K_{12}^{(v)})' L^{-1\prime} \check{\Lambda} L^{-1} (K_{12} + \theta_0 K_{12}^{(v)}), \text{ and} \tag{48}$$
$$\zeta^* = \check{\Lambda} L^{-1} (h_1 + \theta_0 h_1^{(v)} - (K_{12} + \theta_0 K_{12}^{(v)}) \zeta_2). \tag{49}$$

References

[1] Ansley, C. F. and Kohn, R., Estimation, filtering, and smoothing in state space models with incompletely specified initial conditions, *The Annals of Statistics*, Vol. 13, No. 4, pp. 1286-1316, 1985.

[2] Cowell, R., Dawid, P., Lauritzen, S. and Spiegelhalter, D., *Probabilistic Networks and Expert Systems*, Springer, New York, 1999.

[3] Francis, A. B., *A Course in H_∞ Control Theory*, Lecture Notes in Control and Information Sciences, Springer-Verlag, 1987.

[4] Glover, K. and Doyle, J. C., State-space formulae for all stabilizing controllers that satisfy an H_∞-norm bound and relations to risk-sensitivity, *Systems & Control Letters*, Vol. 11, pp. 167-172, 1988.

[5] Lauritzen, S. L., Propagation of probabilities, means, and variances in mixed graphical association models, *Journal of the American Statistical Association*, Vol. 87, pp. 1098-1108, 1992.

[6] Lauritzen, S. L. and Wermuth, N., Graphical models for associations between variables, some of which are qualitative and some quantitative, *Annals of Statistics*, Vol. 17, pp. 31-57, 1989.

[7] Pearl, J., *Causality: Models, Reasoning, and Inference*, Cambridge, 2000.

[8] Poland, W. B., *Decision Analysis with Continuous and Discrete Variables: A Mixture Distribution Approach*, Ph.D. thesis, Department of Engineering-Economic Systems, Stanford University, Stanford, California, 1994.

[9] Raiffa, H., *Decision Analysis*, Addison-Wesley, Reading, Massachusetts, 1968.

[10] Shachter, R. D. and Kenley, C. R., Gaussian influence diagrams, *Management Science*, Vol. 35, No. 5, pp. 527-550, 1989.

[11] Shenoy, P. P., Valuation-based systems for Bayesian decision analysis, *Operations Research*, Vol. 40, pp. 463-484, 1992.

[12] Whittle, P., *Risk-sensitive optimal control*, John Wiley & Sons, Inc., New York, 1990.

[13] Whittle, P., *Optimal control*, John Wiley & Sons, Inc., New York, 1996.