

TECHNICAL RESEARCH REPORT

Risk-Sensitive Optimal Control of Hidden Markov Models: A Case Study

by E. Fernandez-Gaucherand, S.I. Marcus

T.R. 94-71



*Sponsored by
the National Science Foundation
Engineering Research Center Program,
the University of Maryland,
Harvard University,
and Industry*

**RISK-SENSITIVE OPTIMAL CONTROL OF
HIDDEN MARKOV MODELS:
STRUCTURAL RESULTS ***

Emmanuel Fernández-Gaucherand [†]
Systems & Industrial Engineering Department
The University of Arizona
Tucson, AZ 85721.

Steven I. Marcus [‡]
Institute for Systems Research and
Electrical Engineering Department
The University of Maryland
College Park, MD 20742.

* To appear in *IEEE Transactions on Automatic Control*.

[†] Research supported in part by a grant from the University of Arizona Foundation and the Office of the Vice President for Research; in part by the National Science Foundation under grant NSF-INT 9201430; and in part by The Engineering Foundation under grant RI-A-93-10. Email: emmanuel@sie.arizona.edu.

[‡] Research partially supported by the National Science Foundation under Grant EEC 9402384. Email: marcus@src.umd.edu.

1. INTRODUCTION

We consider a risk-sensitive optimal control problem for hidden Markov models (HMM), i.e., controlled Markov chains where state information is only available to the decision-maker (DM) or controller via an output (message) process. The optimal control of HMM under standard, risk-neutral performance criteria, e.g., discounted and average costs, has received much attention in the past. Many basic results and numerous applications have been reported in the literature in this subject; see [1], [2], [14], and references therein. Controlled Markov chains with full state information and a risk-sensitive performance criterion have also received some attention [4], [6], [12].

On the other hand, quite the opposite is the situation for HMM under risk-sensitive criteria, e.g., expected value of the exponential of additive costs. Whittle and others (see [19] and references therein) have extensively studied the risk-sensitive optimal control of partially-observable linear exponential quadratic Gaussian (LEQG) systems; see also [5]. More recently, James, Baras and Elliott [13], [3], have treated the risk-sensitive partially-observable optimal control problem of discrete-time non-linear systems.

The paucity of results in this subject area can be mostly attributed to the lack in the past of appropriate *sufficient statistics*, or *information states*. As is well known, if the cost criterion being considered is of the type “expected value of additive costs,” then the posterior probability density, given all available information up to the present, constitutes a sufficient statistic for control (or information state); see [1], [2], [14]. The latter result was originally proved by Shiryaev in the early sixties, who also proved that this was not the case for non-additive cost criteria; see [16] and references therein. In particular, the posterior probability density is not a *sufficient* statistic for HMM under an “exponential of sum of costs” type of criterion, which is non-additive. This fact was overlooked in [11].

Recently, James, Baras, and Elliott [3], [13] have derived information states for HMM under an “exponential of additive costs” criterion, and have also given dynamic programming equations from which optimal values and controls can be computed, for problems with a finite horizon. Building upon their work, we report in this paper results of an investigation on the nature and structure of risk-sensitive controllers for HMM. We pose the following question: **How does risk-sensitivity manifest itself in the structure of a controller?**

Whittle [19] has addressed a similar question for the LEQG problem, and he has shown that much insight can be gained from a comparison of the risk-neutral (i.e., the classical LQG) and risk-sensitive equations describing the optimal controller. In our context, one difficulty encountered is that optimal controllers are defined in terms of different information states for the risk-neutral and risk sensitive cases; see also [3], [13].

The paper is organized as follows. In section 2 we present our model, and recall the main results on information states from [3], [13] that will be needed

for our developments. Section 3 contains several general structural results, and in section 4 we present a particular case study of a popular benchmark problem. We obtain structural results for the optimal risk-sensitive controller, and compare it to that of the risk-neutral case. Furthermore, we show that indeed the risk-sensitive controller and its corresponding information state converge to the known solutions for the risk-neutral situation, as the risk factor goes to zero. We also study the infinite and general risk aversion cases.

2. THE CONTROLLED HIDDEN MARKOV MODEL

A *controlled hidden Markov model*, or *partially observable Markov decision processes* is given by a five-tuple $\langle \mathbf{X}, \mathbf{Y}, \mathbf{U}, \{P(u) : u \in \mathbf{U}\}, \{Q(u) : u \in \mathbf{U}\} \rangle$; here $\mathbf{X} = \{1, 2, \dots, N_{\mathbf{X}}\}$ is the finite set of (internal) states, $\mathbf{Y} = \{1, 2, \dots, N_{\mathbf{Y}}\}$ is the set of observations (or messages), $\mathbf{U} = \{1, 2, \dots, N_{\mathbf{U}}\}$ is the set of decisions (or controls). In addition, we have that $P(u) := [p_{i,j}(u)]$ is the $N_{\mathbf{X}} \times N_{\mathbf{X}}$ state transition matrix, and $Q(u) := [q_{x,y}(u)]$ is the $N_{\mathbf{X}} \times N_{\mathbf{Y}}$ state/message matrix, i.e., $q_{x,y}(u)$ is the probability of receiving message y when the state is x and action u has been selected. Different types of information patterns are possible (see [9], [10]); we consider the following information pattern.

Information Pattern (IP):

At decision epoch t , the system is in the (unobservable) state $X_t = i$, a decision $U_t = u$ is taken, and the state evolves to $X_{t+1} = j$ with probability $p_{i,j}(u)$. Once the state has evolved to X_{t+1} , an observation Y_{t+1} is gathered, such that:

$$\text{Prob}\{Y_{t+1} = y \mid X_{t+1} = i, U_t = u\} = q_{x,y}(u). \quad (2.1)$$

Hence, based on $\mathcal{I}_t := (Y_0, U_0, Y_1, \dots, U_t, Y_{t+1})$, a new decision U_{t+1} is selected.

Given an expected cost per stage $(i, u) \mapsto c(i, u)$, the sum of costs for the finite horizon M is given by

$$\mathcal{C}_M := \sum_{t=0}^{M-1} c(X_t, U_t). \quad (2.2)$$

The *risk-sensitive optimal control* problem is that of finding a control policy $\pi = \{\pi_0, \pi_1, \dots, \pi_{M-1}\}$, with $\mathcal{I}_t \mapsto \pi_t(\mathcal{I}_t) \in \mathbf{U}$, such that the following criterion is minimized:

$$J^\gamma(\pi) := \text{sgn}(\gamma) \mathbb{E}^\pi [\exp(\gamma \cdot \mathcal{C}_M)], \quad (2.3)$$

where $\gamma \neq 0$ is the *risk-factor*, and $\text{sgn}(\gamma)$ is the sign of γ ; here \mathbb{E}^π denotes the expectation induced by policy π and, implicitly, the initial distribution of the state. By computing the Taylor series expansion of $J^\gamma(\pi)$, when γ is sufficiently small, the risk sensitivity of the above criterion becomes evident in that, in addition to the

standard expected sum of costs, a second order term in the expansion measures the variance of \mathcal{C}_M [19]. If $\gamma > 0$, then the DM or controller is *risk-averse* or *pessimistic*, whereas if $\gamma < 0$ then the DM or controller is *risk-prefering* or *optimistic* [10], [19]

2.1 INFORMATION STATES

As for the risk-neutral case [1], [2], [14], an equivalent stochastic optimal control problem can be formulated in terms of *information states* and *separated policies*. Here we follow the work of Baras, Elliott, and James [3], [13]. Let \mathcal{Y}_t be the filtration generated by the available observations up to decision epoch t , and let \mathcal{G}_t be the filtration generated by the sequence of states and observations up to that time as given by (IP). Then the probability measure induced by a policy π is equivalent to a canonical distribution \mathcal{P}^\dagger , under which $\{Y_t\}$ is independently and identically distributed (i.i.d), uniformly distributed, independent of $\{X_t\}$, and $\{X_t\}$ is a controlled Markov chain with transition matrix as above. We have that

$$\frac{d\mathcal{P}^\pi}{d\mathcal{P}^\dagger} \Big|_{\mathcal{G}_t} = \lambda_t^\pi := N_{\mathbf{Y}}^t \cdot \prod_{k=1}^t q_{X_k, Y_k}(U_{k-1}). \quad (2.4)$$

Then, the cost incurred by using the policy π is given by

$$J^\gamma(\pi) := \text{sgn}(\gamma) \mathbb{E}^\pi[\exp(\gamma \cdot \mathcal{C}_M)] = \text{sgn}(\gamma) \mathbb{E}^\dagger[\lambda_M^\pi \cdot \exp(\gamma \cdot \mathcal{C}_M)]. \quad (2.5)$$

Following [3], [13], the information state for our problem is given by

$$\sigma_t^\gamma(i) := \mathbb{E}^\dagger[\mathbf{1}[X_t = i] \exp(\gamma \cdot \mathcal{C}_t) \cdot \lambda_t^\pi \mid \mathcal{Y}_t], \quad (2.6)$$

where $\mathbf{1}[A]$ is the indicator function of the event A , and $\sigma_0^\gamma(i) = p_0$, where p_0 is the initial distribution of the state and is assumed to be known. Notice that $\sigma_t^\gamma \in \mathbb{R}_+^{N_{\mathbf{x}}} := \{\sigma \in \mathbb{R}^{N_{\mathbf{x}}} \mid \sigma(i) \geq 0, \forall i\}$. With this definition of information state, similar results as in the risk-neutral case can be obtained. In particular, one obtains a recursive updating formula for $\{\sigma_t^\gamma\}$, which is driven by the output (observation) path and evolves forward in time. Moreover, the value functions can be expressed in terms of the information state only, and dynamic programming equations give necessary and sufficient optimality conditions for *separated policies*, i.e., maps $\sigma_t^\gamma \mapsto \tilde{\pi}_t(\sigma_t^\gamma) \in \mathbf{U}$; see [3], [13]. In particular we have that:

$$J^\gamma(\pi) = \text{sgn}(\gamma) \mathbb{E}^\dagger \left[\sum_{i=1}^{N_{\mathbf{x}}} \sigma_M^\gamma(i) \right], \quad (2.7)$$

where $\{\sigma_M^\gamma\}$ is obtained from (2.4)-(2.6) under the action of policy π . Hence, the original partially observed problem is equivalently expressed as one with complete state information, i.e., $\{\sigma_t^\gamma\}$. For ease of presentation, we consider hereafter the risk-averse case only ($\gamma > 0$); the risk-seeking case is treated similarly.

3. GENERAL RESULTS

As in the completely observed case [12], define the *disutility contribution matrix* as:

$$[\mathcal{D}(u)]_{i,j} := p_{i,j}(u) \cdot \exp(\gamma c(i, u)). \quad (3.1)$$

The following lemma gives the recursions that govern the evolution of the information state; its proof follows easily from [3] and [13].

Lemma 3.1: The information state process $\{\sigma_t^\gamma\}$ is recursively computable as:

$$\sigma_{t+1}^\gamma = N_{\mathbf{Y}} \cdot \overline{Q}(Y_{t+1}, U_t) \mathcal{D}^T(U_t) \cdot \sigma_t^\gamma, \quad (3.2)$$

where $\overline{Q}(y, u) := \text{diag}(q_{i,y})(u)$, and A^T denotes the transpose of the matrix A .

Remark 3.1: Observe that as $\gamma \rightarrow 0$, $\mathcal{D}(u) \rightarrow P(u)$ (elementwise). Therefore, we see that (3.2) is the “natural” extrapolation of the (unnormalized) conditional probability distribution of the (unobservable) state, given the available observations, which is the standard risk-neutral information state [1], [2], [14].

As in [3], [13] define value functions $J^\gamma(\cdot, M - k) : \mathbb{R}_+^{N_{\mathbf{x}}} \rightarrow \mathbb{R}$, $k = 1, \dots, M$, as follows:

$$J^\gamma(\sigma, M - k) := \min_{\pi_{M-k} \dots \pi_{M-1}} \left\{ \mathbb{E}^\dagger \left\{ \sum_{i=1}^{N_{\mathbf{x}}} \sigma_M^\gamma(i) \mid \sigma_{M-k}^\gamma = \sigma \right\} \right\}. \quad (3.3)$$

Denote by $T(u, y)$ the matrix

$$T(u, y) := N_{\mathbf{Y}} \cdot \overline{Q}(y, u) \mathcal{D}^T(u). \quad (3.4)$$

The next result follows directly from [3], [13].

Lemma 3.2: The dynamic programming equations for the value functions in this problem are given as:

$$\begin{cases} J^\gamma(\sigma, M) &= \sum_{i=1}^{N_{\mathbf{x}}} \sigma(i); \\ J^\gamma(\sigma, M - k) &= \min_{u \in \mathbf{U}} \left\{ \mathbb{E}^\dagger [J^\gamma(T(u, Y_{M-k+1}) \cdot \sigma, M - k + 1)] \right\} \quad k = 1, 2, \dots, M. \end{cases} \quad (3.5)$$

Furthermore, a separated policy $\pi^* = \{\pi_0^*, \dots, \pi_{M-1}^*\}$ that attains the minimum in (3.5) is risk-sensitive optimal.

Next, we present several general results for the risk-sensitive case that have similar counterparts in the standard risk-neutral case [1], [2], [7], [14], [17].

Lemma 3.3: The value functions given by (3.5) are concave functions of $\sigma \in \mathbb{R}_+^{N_{\mathbf{X}}}$.

Proof:

We proceed by induction in k , with the case $k = 0$ being trivially verified from (3.5). Assume that the claim holds true for $0 \leq \bar{k} = k - 1 < M$. Let $0 \leq \lambda \leq 1$ and $\sigma_1, \sigma_2 \in \mathbb{R}_+^{N_{\mathbf{X}}}$, and define $\tilde{\sigma} := \lambda\sigma_1 + (1 - \lambda)\sigma_2$. Then we have that:

$$\begin{aligned}
J^\gamma(\tilde{\sigma}, M - k) &= \min_{u \in \mathbf{U}} \left\{ \frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} J^\gamma(T(u, y) \cdot \tilde{\sigma}, M - k + 1) \right\} \\
&\geq \min_{u \in \mathbf{U}} \left\{ \frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} [\lambda J^\gamma(T(u, y) \cdot \sigma_1, M - k + 1) \right. \\
&\quad \left. + (1 - \lambda) J^\gamma(T(u, y) \cdot \sigma_2, M - k + 1)] \right\} \\
&\geq \lambda J^\gamma(\sigma_1, M - k) + (1 - \lambda) J^\gamma(\sigma_2, M - k),
\end{aligned} \tag{3.6}$$

where the first inequality follows due to the induction hypothesis, and the second inequality due to (3.5). \square

Next, define recursively sets of vectors in $\mathbb{R}_+^{N_{\mathbf{X}}}$ as follows:

$$\begin{aligned}
A_0 &:= \{\mathbf{1} = (1, 1, \dots, 1)\}, \\
A_k &:= \left\{ \frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} \alpha_y \cdot T(u, y) \mid \alpha_y \in A_{k-1}, u \in \mathbf{U} \right\}.
\end{aligned} \tag{3.7}$$

Note that the cardinality of the sets defined in (3.7) obeys the recursion $|A_k| \leq |A_{k-1}|^{N_{\mathbf{Y}}} \cdot N_{\mathbf{U}}$. In the risk-neutral case, the counterpart of the following result has been shown to have important computational implications [1], [7], [17]. It will play a key role in our subsequent developments.

Lemma 3.4: The value functions given by (3.5) are piecewise linear functions in $\sigma \in \mathbb{R}_+^{N_{\mathbf{X}}}$, such that:

$$J^\gamma(\sigma, M - k) = \min_{\alpha \in A_k} \{\alpha \cdot \sigma\}. \tag{3.8}$$

Proof:

We proceed by induction in k , with the case $k = 0$ being trivially verified from (3.5). Assume that the claim holds true for $0 \leq \bar{k} = k - 1 < M$, then from (3.5) above we have:

$$\begin{aligned}
J^\gamma(\sigma, M - k) &= \min_{u \in \mathbf{U}} \left\{ \frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} \min_{\alpha \in A_{k-1}} \{ \alpha \cdot T(u, y) \cdot \sigma \} \right\} \\
&= \min_{u \in \mathbf{U}} \left\{ \left[\frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} \tilde{\alpha}(u, y, \sigma) \cdot T(u, y) \right] \cdot \sigma \right\}, \quad (3.9) \\
&= \min_{\alpha \in A_k} \{ \alpha \cdot \sigma \}
\end{aligned}$$

where $\tilde{\alpha}(u, y, \sigma) \in A_{k-1}$ denotes a minimizer in the expression on the right of the first equality above. The last equality follows since $\alpha \cdot T(u, y) \cdot \sigma > \tilde{\alpha}(u, y, \sigma) \cdot T(u, y) \cdot \sigma$, for all $\alpha \in A_{k-1}$, $u \in \mathbf{U}$, $y \in \mathbf{Y}$, $\sigma \in \mathbb{R}_{N_{\mathbf{x}}}^+$. \square

Lemma 3.5: Optimal separated policies $\{\pi_t^*\}$ are constant along rays through the origin, i.e., let $\sigma \in \mathbb{R}_+^{N_{\mathbf{x}}}$ then $\pi_t^*(\sigma') = \pi_t^*(\sigma)$, for all $\sigma' = \lambda\sigma$, $\lambda \geq 0$.

Proof:

From Lemma 3.4 we see that $J^\gamma(\sigma', M - k) = \lambda J^\gamma(\sigma, M - k)$. Hence, the result follows from Lemma 3.2. \square

Definition 3.1: From (3.5), for $u \in \mathbf{U}$ and $k = 1, 2, \dots, M$, let

$$\begin{aligned}
J_u^\gamma(\sigma, M - k) &:= \mathbb{E}^\dagger [J^\gamma(T(u, Y_{M-k+1}) \cdot \sigma, M - k + 1)] \\
&= \frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} [J^\gamma(T(u, Y_{M-k+1}) \cdot \sigma, M - k + 1)]. \quad (3.10)
\end{aligned}$$

The *control region* $CR_u^k \subseteq \mathbb{R}_+^{N_{\mathbf{x}}}$ for action $u \in \mathbf{U}$, at the $M - k$ decision epoch, is defined as:

$$CR_u^k := \{ \sigma \mid \sigma \in \mathbb{R}_+^{N_{\mathbf{x}}}, J^\gamma(\sigma, M - k) = J_u^\gamma(\sigma, M - k) \}. \quad (3.11.a)$$

Furthermore by Lemma 3.2 if π_{M-k}^* is an optimal separated policy for stage $M - k$ then, for $u \in \mathbf{U}$,

$$CR_u^k := \{ \sigma \in \mathbb{R}_+^{N_{\mathbf{x}}} \mid \pi_{M-k}^*(\sigma) = u \}. \quad (3.11.b)$$

Definition 3.2: An action $\bar{u} \in \mathbf{U}$ is said to be a *resetting* action if there exists $j^* \in \mathbf{X}$ such that $p_{i,j^*}(\bar{u}) = 1$, for all $i \in \mathbf{X}$. Therefore from (3.1)-(3.2) and (3.4) we note that, for any $\sigma \in \mathbb{R}_+^{N_{\mathbf{X}}}$ and $y \in \mathbf{Y}$,

$$T(\bar{u}, y) \cdot \sigma = q_{j^*, y} \sum_{\ell=1}^{N_{\mathbf{X}}} [\exp(c(\ell, \bar{u})) \cdot \sigma(\ell)] \cdot \nu_{j^*}, \quad (3.12)$$

where $\nu_1 = (1, 0, 0, \dots, 0)^T$, ..., $\nu_{N_{\mathbf{X}}} = (0, 0, 0, \dots, 1)^T$. We further denote

$$\Lambda(\sigma, \bar{u}) := q_{j^*, y} \sum_{\ell=1}^{N_{\mathbf{X}}} [\exp(c(\ell, \bar{u})) \cdot \sigma(\ell)]. \quad (3.13)$$

Theorem 3.1: Let $\bar{u} \in \mathbf{U}$ be a resetting action. Then $CR_{\bar{u}}^k$ is a convex subset of $\mathbb{R}_+^{N_{\mathbf{X}}}$.

Proof:

Recall from Lemma 3.3 that the optimal cost-to-go functions $J^\gamma(\cdot, M - k)$ are concave. Since the maps $T(u, y)$ in (3.4) are linear then $J_u^\gamma(\cdot, M - k)$ are also concave, for all $u \in \mathbf{U}$. Furthermore, $\mathbb{R}_+^{N_{\mathbf{X}}}$ is a convex domain. Then, by Lemma 1 in [15] we have that: if $J_{\bar{u}}^\gamma(\cdot, M - k)$ is a linear function in $\mathbb{R}_+^{N_{\mathbf{X}}}$, then $CR_{\bar{u}}^k$ is convex. Thus all that remains to be proven is the linearity of $J_{\bar{u}}^\gamma(\cdot, M - k)$.

Let $\sigma \in \mathbb{R}_+^{N_{\mathbf{X}}}$, we have by (3.2), (3.4), Definitions 3.1-3.2, and Lemma 3.4 that:

$$\begin{aligned} J_{\bar{u}}^\gamma(\sigma, M - k) &= \frac{1}{N_{\mathbf{X}}} \sum_{y=0}^1 \min_{\alpha \in A_{k-1}} \{ \alpha \cdot T(\bar{u}, y) \sigma \} \\ &= \Lambda(\sigma, \bar{u}) \cdot \min_{\alpha \in A_{k-1}} \{ \alpha \cdot \nu_{j^*} \} \\ &= \Lambda(\sigma, \bar{u}) \cdot \alpha^*(j^*), \end{aligned} \quad (3.14)$$

where $\alpha^*(j^*) := \min \{ \alpha(j^*) \mid (\alpha(1), \alpha(2), \dots, \alpha(j^*), \dots, \alpha(N_{\mathbf{X}})) \in A_{k-1} \}$. Hence, since by (3.13) $\Lambda(\cdot, \bar{u})$ is linear in $\mathbb{R}_+^{N_{\mathbf{X}}}$, so is $J_{\bar{u}}^\gamma(\cdot, M - k)$. \square

4. A CASE STUDY

We consider a popular benchmark problem for which much is known in the risk-neutral case. This is a two-state replacement problem which models failure-prone units in production/manufacturing systems, communication systems, etc. The underlying state of the unit can either be *working* ($X_t = 0$) or *failed* ($X_t = 1$), and the available actions are to *keep* ($U_t = 0$) the current unit or *replace* ($U_t = 1$) the unit by a new one. The cost function $(x, u) \mapsto c(x, u)$ is as follows: let $R > C > 0$, then $c(0, 0) = 0$, $c(1, 0) = C$, $c(x, 1) = R$. The messages received have probability $1/2 < q < 1$ of coinciding with the true state of the unit. The state transition matrices are given as:

$$P(0) = \begin{bmatrix} 1 - \theta & \theta \\ 0 & 1 \end{bmatrix}; \quad P(1) = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}, \quad (4.1)$$

with $0 < \theta < 1$; see [7], [8], [18] for more details. With the above definitions, the matrices used to update the information state vector are given by:

$$T(0, y) = 2 \begin{bmatrix} q_y(1 - \theta) & 0 \\ (1 - q_y)\theta & (1 - q_y)e^{\gamma C} \end{bmatrix}; \quad T(1, y) = 2 \begin{bmatrix} q_y e^{\gamma R} & q_y e^{\gamma R} \\ 0 & 0 \end{bmatrix}, \quad (4.2)$$

where $q_y := q(1 - y) + (1 - q)y$, $y = 0, 1$. For this case $\sigma = (\sigma(1), \sigma(2))^T \in \mathbb{R}_+^2$, and the dynamic programming recursions (3.5) take the form:

$$\begin{cases} J^\gamma(\sigma, M) & = \sigma(1) + \sigma(2); \\ J^\gamma(\sigma, M - k) & = \min\{J_0^\gamma(\sigma, M - k); J_1^\gamma(\sigma, M - k)\}. \end{cases} \quad (4.3)$$

Define the *replace* control region $CR_{replace}^k$ and the *keep* control region CR_{keep}^k in the obvious manner, c.f. Definition 3.1. The next result follows from (4.2), Lemma 3.5, and Theorem 3.1.

Lemma 4.1. For all decision epochs the *replace* control region is a (possibly empty) conic segment in \mathbb{R}_+^2 .

The next result establishes an important *threshold* structural property of the optimal control policy. This is similar to well known results for the risk neutral case [7], [8], [15], [18].

Theorem 4.1. If $CR_{replace}^k$ is nonempty, then it includes the $\sigma(2)$ -axis, i.e., \mathbb{R}_+^2 is partitioned by a line through the origin such that for values of $\sigma \in \mathbb{R}_+^2$ above the line it is optimal to *replace* the unit, and it is optimal to *keep* the unit otherwise.

Proof:

We proceed to show that if it is optimal to *keep* the unit in the $\sigma(2)$ -axis (see Lemma 3.5), then the optimal policy is to *keep* the unit for all values of $\sigma \in \mathbb{R}_+^2$. Hence, by contradiction, we can then conclude from Lemma 4.1 that if $CR_{replace}^k$ is nonempty, then it must include the $\sigma(2)$ -axis, and the statement of the theorem then follows. Let $\sigma' = (0, \sigma(2))^T$, $\sigma(2) > 0$. Then, for $0 < k \leq M$, we have from (4.2), (4.3), and Lemma 3.4 that:

$$J_0^\gamma(\sigma', M - k) = e^{\gamma C} \sigma(2) \alpha^*(2), \quad J_1^\gamma(\sigma', M - k) = e^{\gamma R} \sigma(2) \alpha^*(1),$$

where $\alpha^*(i)$ denotes the componentwise minimum over A_{k-1} . Suppose that

$$J_0^\gamma(\sigma', M - k) < J_1^\gamma(\sigma', M - k) \iff e^{\gamma C} \alpha^*(2) < e^{\gamma R} \alpha^*(1). \quad (4.4)$$

Now, for any other $\sigma \in \mathbb{R}_+^2$,

$$J_1^\gamma(\sigma, M - k) = e^{\gamma R} (\sigma(1) + \sigma(2)) \alpha^*(1),$$

and, since $\gamma > 0$ and $R > C > 0$, then

$$\begin{aligned} J_0^\gamma(\sigma, M - k) &= \sum_{y=0}^1 \min_{\alpha \in A_{k-1}} \{q_y(1 - \theta)\sigma(1)\alpha(1) + (1 - q_y)\theta\sigma(1)\alpha(2) \\ &\quad + (1 - q_y)\sigma(2)e^{\gamma C}\alpha(2)\} \\ &< \sum_{y=0}^1 \min_{\alpha \in A_{k-1}} \{q_y(1 - \theta)\sigma(1)e^{\gamma R}\alpha(1) + (1 - q_y)\theta\sigma(1)e^{\gamma C}\alpha(2) \\ &\quad + (1 - q_y)\sigma(2)e^{\gamma R}\alpha(2)\}. \end{aligned} \quad (4.5)$$

Now, define $\tilde{A}_{k-1} \subseteq A_{k-1}$ as:

$$\tilde{A}_{k-1} := \left\{ \alpha \in A_{k-1} \mid e^{\gamma C} \alpha(2) < e^{\gamma R} \alpha^*(1) \right\}, \quad (4.6)$$

which is nonempty by (4.4). Then by minimizing over \tilde{A}_{k-1} the terms on the right hand side in (4.5) we obtain an upper-bound for this expression, and we finally get that:

$$J_0^\gamma(\sigma, M - k) < e^{\gamma R} \sigma(1) \alpha^*(1) + e^{\gamma R} \sigma(2) \alpha^*(1) = J_1^\gamma(\sigma, M - k),$$

and therefore it is optimal to *keep* the unit at all $\sigma \in \mathbb{R}_+^2$. \square

Using the dynamic programming recursions (4.3), the structure of optimal policies can be further elucidated. First we need a simple technical result; see [9], [10] for a proof. Let $\alpha_0 = 1$, and for $k = 0, 1, \dots, M$ define:

$$\alpha_{k+1} := (1 - \theta)\alpha_k + \theta e^{k\gamma C}. \quad (4.7)$$

Lemma 4.2: $\alpha_{k+1} > \alpha_k$, and $e^{\gamma R}\alpha_k > \alpha_{k+1}$; $k = 1, 2, \dots, M$.

The following theorem gives more precise results on the structure of optimal policies. Its proof is by backwards induction using (4.3); see [9], [10].

Theorem 4.2: Let $0 < \bar{K} \leq M$ be given.

- (i) The necessary and sufficient condition for the policy with $\pi_{M-1}^*(\cdot) = \dots = \pi_{M-\bar{K}}^*(\cdot) = 0$ (i.e., always *keep* the unit in the last \bar{K} stages) to be optimal is that:

$$\frac{e^{\bar{K}\gamma C}}{\alpha_{\bar{K}-1}} \leq e^{\gamma R} \iff R \geq \bar{K}C - \frac{\ln(\alpha_{\bar{K}-1})}{\gamma}. \quad (4.8)$$

- (ii) If (4.8) holds, then:

$$\begin{aligned} J^\gamma(\sigma, M - \bar{K}) &= J_0^\gamma(\sigma, M - \bar{K}) = \alpha_{\bar{K}}\sigma(1) + e^{\bar{K}\gamma C}\sigma(2); \\ J_1^\gamma(\sigma, M - \bar{K}) &= \alpha_{\bar{K}-1}e^{\gamma R}(\sigma(1) + \sigma(2)). \end{aligned} \quad (4.9)$$

- (iii) If $1 \leq \bar{K} \leq M$ is the smallest integer for which (4.8) fails to hold, then $\pi_{M-\bar{K}}^*(\cdot)$ is of threshold type, with \mathbb{R}_+^2 being partitioned by the line:

$$\frac{e^{\gamma R}\alpha_{\bar{K}-1} - \alpha_{\bar{K}}}{e^{\bar{K}\gamma C} - e^{\gamma R}\alpha_{\bar{K}-1}}\sigma(1) = \sigma(2), \quad (4.10)$$

such that the region to the left (above) the line is the *replace* control region.

Remark 4.1. Note that the simplest nontrivial decision process corresponds to the case $M = 2$, since (4.8) is always satisfied for $\bar{K} = 1$.

4.1. SMALL AND LARGE RISK LIMITS

Infinite Risk Aversion Case ($\gamma \rightarrow \infty$).

Consider the situation $\gamma \rightarrow \infty$ in (4.7). Note that for γ large enough:

$$\alpha_{\overline{K}-1} \approx \theta \cdot e^{(\overline{K}-2)\gamma C}.$$

Therefore, as $\gamma \rightarrow \infty$, we have from (4.8) that the necessary and sufficient condition for it to be always optimal to *keep* the unit in the last \overline{K} stages approaches $R \geq 2C$, which is the same condition for it to be always optimal to *keep* the unit in the last two stages. Furthermore, as is readily verified from (4.10), if $R < 2C$ and $\gamma \rightarrow \infty$ then it is always optimal to *replace* the unit at stage $M - k$, for all $2 \leq k \leq M$, i.e., the threshold line tends to the $\sigma(1)$ -axis. Hence the DM becomes *myopic* in the sense that, perhaps except for the last one, all decision epochs appear to be the same. The DM appears to always face a *two*-stage decision process, the simplest one possible. In the jargon of Whittle [19], it could be then said that an infinitely risk averse DM exhibits “neurotic” behavior, his optimal strategy being of the “*bang bang*” type with respect to the parameter R : if $R \geq 2C$, then $\pi_{M-k}^*(\cdot) = 0$, and otherwise $\pi_{M-k}^*(\cdot) = 1$, for all $2 \leq k \leq M$. This behavior can be partly explained by noting that at most one change will then occur in the stream of costs, thus achieving least variability in the cumulative cost.

Small Risk Aversion Case ($\gamma \rightarrow 0$).

Next, we examine the question: **How do the results in Theorems 3.1, 4.1-4.2 compare to known results for the risk-neutral case?** The answer is that the risk-sensitive controller obtained here has as its small risk limit the known risk-neutral controller, and both controllers have in general a similar structure. Similarly as in [18], the dynamic programming equations for the risk-neutral case can be written, with the conditional probability distribution of the state as the information state. Then, it can be shown that the optimal risk-neutral controller has a structure similar to the risk-sensitive controller given in Theorem 4.1. Furthermore, it can be shown that the necessary and sufficient condition in the risk-neutral case for the separated policy $\pi_{M-1}^*(\cdot) = \dots = \pi_{M-\overline{K}}^*(\cdot) = 0$ to be optimal is:

$$R > \overline{K}C - \alpha'_{\overline{K}-1}, \tag{4.11}$$

where $\alpha'_{\overline{K}-1}$ is obtained as the derivative with respect to γ , evaluated as $\gamma \rightarrow 0$, of (4.7). As can be easily verified, the above is nothing but the small risk limit (i.e., as $\gamma \rightarrow 0$) of (4.8).

General Risk Aversion Case ($\gamma > 0$).

The following result helps bring to light a manifestation of aversion to risk in the DM or controller; its proof is given in the Appendix.

Lemma 4.3. Let $\gamma > 0$, then for all $k > 1$:

$$\ln(\alpha_k) > \gamma \cdot \alpha'_k. \quad (4.12)$$

□

Notice that the decision to *replace* a unit involves an uncertain, and therefore a risky, investment in that the unit being replaced may actually be in *working* condition, or it may subsequently fail. This is reflected in (4.8), (4.11) and (4.12) in that a risk neutral DM or controller may decide to *replace* a unit for values of R higher than a risk averse DM or controller would.

REFERENCES

- [1] A. Arapostathis, V. Borkar, E. Fernández-Gaucherand, M.K. Ghosh and S.I. Marcus, Discrete-Time Controlled Markov Processes with Average Cost Criterion: A Survey, *SIAM Journal on Control & Optimization* **31** (1993) 282-344.
- [2] D.P. Bertsekas, *Dynamic Programming: Deterministic and Stochastic Models*, Prentice-Hall, Englewood Cliffs, 1987.
- [3] J.S. Baras and M.R. James, Robust and Risk-sensitive Output Feedback Control for Finite State Machines and Hidden Markov Models, preprint, August (1994).
- [4] M. Bouakiz and M.J. Sobel, Inventory Control with an Exponential Utility Criterion, *Operations Research* **40** (1992) 603-608.
- [5] A. Bensoussan and J.H. Van Schuppen, Optimal Control of Partially Observable Stochastic Systems with an Exponential-of-Integral Performance Index, *SIAM Journal on Control & Optimization* **23** (1985) 599-613.
- [6] K-J. Chung and M.J. Sobel, Discounted MDP's: Distribution Functions and Exponential Utility Maximization, *SIAM Journal on Control & Optimization* **25** (1987) 49-62.
- [7] E. Fernández-Gaucherand, A. Arapostathis, and S.I. Marcus, On the Average Cost Optimality Equation and the Structure of Optimal Policies for Partially Observable Markov Decision Processes, *Annals of Operations Research* **29** (1991) 439-470.
- [8] E. Fernández-Gaucherand, A. Arapostathis and S.I. Marcus, Analysis of an Adaptive Control Scheme for a Partially Observed Controlled Markov Chain, *IEEE Transactions on Automatic Control* **38** (1993) 987-993.
- [9] E. Fernández-Gaucherand and S.I. Marcus, Risk-sensitive Optimal Control of Hidden Markov Models: a Case Study, in *Proceedings 33rd IEEE Conf. Decision and Control*, Orlando, FL, (1994) 1657-1662.

- [10] E. Fernández-Gaucherand and S.I. Marcus, Risk-sensitive Optimal Control of Hidden Markov Models: Structural Results. University of Maryland at College Park, Institute for Systems Research report, March 1996.
- [11] A. Gheorghe, Partially Observable Markov Processes with a Risk-sensitive Decision Maker, *Rev. Roum. Math. Pures et Appl.* **22** (1977) 461-482.
- [12] R.A. Howard and J.E. Matheson, Risk-sensitive Markov Decision Processes, *Management Science* **18** (1972) 356-370.
- [13] M.R. James, J.S. Baras and R.J. Elliott, Risk-sensitive control and dynamic games for Partially Observed Discrete-time Nonlinear Systems, *IEEE Transactions on Automatic Control* **39** (1994) 780-792.
- [14] P.R. Kumar and P. Varaiya, *Stochastic Systems: Estimation, Identification and Adaptive Control*, Prentice-Hall, Englewood Cliffs, 1986.
- [15] W.S. Lovejoy, On the Convexity of Policy Regions in Partially Observed Systems, *Operations Research*, **35** (1987) 619-621.
- [16] A. N. Shiryaev, On Markov Sufficient Statistics in Non-additive Bayes Problems of Sequential Analysis, *Theory Probab. Appl.* **9** (1964) 604-618.
- [17] R.D Smallwood and E.J. Sondik, The Optimal Control of Partially Observable Markov Processes Over a Finite Horizon, *Operations Research* **21** (1973) 1071-1088.
- [18] C.C. White, A Markov Quality Control Process Subject to Partial Observation, *Management Science* **23** (1977) 843-852.
- [19] P. Whittle, *Risk-sensitive Optimal Control*, Wiley, England, 1990.

Appendix

Proof of Lemma 4.3. From the recursion (4.7) it follows that:

$$\alpha_k = \sum_{l=0}^{k-2} (1-\theta)^l \theta e^{(k-1-l)\gamma C} + (1-\theta)^{k-1}, \quad (A.1.a)$$

$$\alpha'_k = \sum_{l=0}^{k-2} (1-\theta)^l \theta (k-1-l)C. \quad (A.1.b)$$

Now, it also follows from (4.7) that $\alpha_k|_{\gamma=0} = 1$, and therefore (A.1.a) is a convex combination of exponentials (the last term corresponding to $e^{0 \cdot \gamma C}$). Therefore, since $\ln(\cdot)$ is a strictly convex function, we have that:

$$\begin{aligned} \ln(\alpha_k) &= \ln\left(\sum_{l=0}^{k-2} (1-\theta)^l \theta e^{(k-1-l)\gamma C} + (1-\theta)^{k-1}\right) \\ &> \sum_{l=0}^{k-2} (1-\theta)^l \theta \ln(e^{(k-1-l)\gamma C}) = \gamma \alpha'_k. \end{aligned} \quad (A.2)$$

□