TECHNICAL RESEARCH REPORT

Estimation of Hidden Markov Models: Risk-Sensitive Filter Banks and Qualitative Analysis of their Sample Paths

by Vahid Reza Ramezani, Steven I. Marcus

TR 2001-21

Estimation of Hidden Markov Models: Risk-Sensitive Filter Banks and Qualitative Analysis of their Sample Paths*

Vahid Reza Ramezani† and Steven I. Marcus‡
Department of Electrical and Computer Engineering
and Institute for Systems Research
University of Maryland
College Park, MD 20742

September 21, 2001

Abstract

A risk-sensitive generalization of the Maximum A Posterior Probability (MAP) estimation for partially observed Markov chains is presented. Using a change of measure technique, a cascade filtering scheme for the risk-sensitive state estimation is introduced. Structural results, the influence of the availability of information, mixing and non-mixing dynamics, and the connection with other risk-sensitive estimation methods are considered. A qualitative analysis of the sample paths clarifies the underlying mechanism.

Keywords: Hidden Markov Models, Risk-sensitive estimation.

*This work was supported in part by the National Science Foundation under Grant DMI-9988867, by the Air Force Office of Scientific Research under Grant F496200110161, and by ONR contract 01-5-28834 under the MURI Center for Auditory and Acoustics Research.

†Institute for Systems Research and the Department of Electrical and Computer Engineering University of Maryland College Park, MD 20742. 301-405-2942 rvahid@isr.umd.edu

‡Institute for Systems Research and the Department of Electrical and Computer Engineering University of Maryland College Park, MD 20742. 301-405-7589 marcus@isr.umd.edu (communicating author).

1 Introduction

A risk sensitive optimal controller or estimator differs from the standard or so called risk neutral one in its inclusion of the "higher order moments". This is typically accomplished by minimizing the expectation of the exponential of a positively scaled sum of costs whose minimization defines the optimality condition for the original risk-neutral problem. The Taylor expansion of the exponential function readily shows that this risk-neutral problem is recovered when the scaling parameter is allowed to approach zero. Whittle and others have extensively studied the linear exponential quadratic Gaussian (LEQG) controller [9]. The estimator counterpart was considered by Speyer [4]; more recently, the control problem for the discrete-time non-linear systems was treated by James, Baras and Elliot in [1] and [3]. Building on their work, structural results were given by E. Fernandez-Gaucherand and Marcus [10].

Dey and Moore [5]-[6] introduced a risk sensitive estimator for Hidden Markov Models (HMM's) with quadratic cost. The robustness of the risk sensitive estimation under certain types of model uncertainties was shown by Boel et al. [7].

In this paper, we introduce a filtering scheme for the risk sensitive Maximum A Posterior Probability (MAP) estimation for HMM's and analyze its structural properties. The previously introduced risk sensitive filters will be considered as special cases.

2 Risk sensitive filter banks

2.1 The estimation of Hidden Markov Models

We define an estimated Hidden Markov Model as a five tuple $< \mathbf{X}, \mathbf{Y}, \mathbf{X}, \mathcal{A}, Q >$; here $\mathcal{A}$ is the transition matrix, $\mathbf{Y} = \{1, 2, ..., N_{\mathbf{Y}}\}$ is the set of observations and $\mathbf{X} = \{1, 2, ..., N_{\mathbf{X}}\}$ is the finite set of (internal) states as well as the set of estimates or decisions. In addition, we have that $Q := [q_{x,y}]$ is the $N_{\mathbf{X}} \times N_{\mathbf{Y}}$ state/observation matrix, i.e., $q_{x,y}$ is the probability of observing y when the state is x. We consider the following information pattern. At decision epoch t, the system is in the (unobservable) state $X_t = i$ and the corresponding observation Y_t is gathered, such that:

$$P(Y_t = y | X_t = i) = q_{x,y} \tag{2.1}$$

The estimators V_t are functions of observations $(Y_0, ...Y_t)$ and are chosen according to some specified criterion. Through out this paper, we use upper case letters to denote estimators and script upper case letters to denote "estimation maps " from observations to the set $\mathbf{X}$. If Y_t is an observation and V_t an estimator: $V_t = \mathcal{V}_t \circ Y_t$. When it causes no confusion, we may use upper case letters for both.

2.2 Maximum A Posterior Probability Estimator (MAP) for HMM's

Consider a sequence of finite dimensional random variables X_t and the corresponding observations Y_t defined on the common probability space $(\mathbf{\Omega}, F, \mathbf{P})$. The Maximum A Posteriori Probability (MAP) estimator $\hat{X}_t$ is a Borel measurable function of the filtration generated by observations up to Y_t denoted by $\mathcal{Y}_t$ which satisfies for $\omega \in \Omega$:

$$\hat{X}_t(\omega) = argmin_{\zeta \in \mathbf{X}} \quad E[\rho(X_t, \zeta) \mid \mathcal{Y}_t = \mathcal{Y}_t(\omega)] \qquad t = 0, 1, ... \tag{2.2}$$

where

$$\rho(X_t, \zeta) = \begin{cases} 0 & \text{if } X_t = \zeta; \\ 1 & otherwise. \end{cases}$$

The usual definition of MAP as the argument with the greatest probability given the observation follows from the above [8]. We will need the following simple Lemma:

Lemma 2.2.1: MAP also results from the additive cost minimization:

$$(\hat{X}_0, ..., \hat{X}_N)(\omega) = argmin_{\zeta_0 ... \zeta_N \in \mathbf{X}^N} \quad E[\sum_{i=0}^{N} \rho(X_i, \zeta) \ )| \ \mathcal{Y}_i = \mathcal{Y}_i(\omega)] \tag{2.3}$$

where $\mathbf{X}^N$ is the product space and each $\hat{X}_i$ is $\mathcal{Y}_i$ measurable.

Proof:

The proof follows easily from the linearity of the conditional expectation and term by term minimization of the resulting sum.

2.3 Change of measure

To carry out the computations, we will use a change of measure technique introduced by Baras, Elliot and James in [1] and [2]. Let $(\mathbf{\Omega}, F, \mathbf{P})$ be the canonical probability space on which all of our time series are defined. 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. Then a new probability measure $\mathcal{P}^\dagger$ is defined by the restriction of the Radon-Nikodym derivative on $\mathcal{G}_t$ to

$$\frac{d\mathcal{P}}{d\mathcal{P}^\dagger}|_{\mathcal{G}_t} = \lambda_t := N_{\mathbf{Y}}^t \cdot \Pi_{k=1}^t q_{X_k, Y_k} \tag{2.4}$$

under which $\{Y_t\}$ is independently and identically distributed (i.i.d). Each distribution is uniform over the set $\mathbf{Y}$ and $\{Y_t\}$ is independent of $\{X_t\}$. That such a measure exists follows directly from Kolmogorov extension theorem.

Before we can introduce our filter, we need an optimization result. Let V_t be measurable functions of observations up to t taking values in $\{X_t\}$ and $\rho(\cdot,\cdot)$ as above. Fix $V_0,...,V_{k-1}$. We would like to find $\hat{V}_k,...,\hat{V}_{H-1}$ such that the following criterion is minimized :

$$S^\gamma(V_0,...,V_{H-1}) := E[exp(\gamma \cdot \mathcal{C}_H)], \tag{2.5}$$

where

$$\mathcal{C}_H := \sum_{t=0}^{H-1} \rho(X_t, V_t) \tag{2.6}$$

and γ is a strictly positive parameter. The minimum value will be denoted by $\bar{S}^\gamma(\hat{V}_k,...,\hat{V}_{H-1})$.

This optimization problem can be solved via dynamic programming. A detailed discussion is relegated to the Appendix. First we need to define recursively an *information state*

$$\sigma^\gamma_{t+1} = N_{\mathbf{Y}} \cdot \overline{Q}(Y_{t+1})\mathcal{D}^T(V_t) \cdot \sigma^\gamma_t, \tag{2.7}$$

where $\overline{Q}(y) := diag(q_{i,y})$, $\mathcal{A}^T$ denotes the transpose of the matrix $\mathcal{A}$ and the matrix $\mathcal{D}$ is defined by

$$[\mathcal{D}(v)]_{i,j} := a_{i,j} \cdot exp(\gamma\rho(i,v)). \tag{2.8}$$

σ^γ_0 is set equal to $N_{\mathbf{Y}} \cdot \overline{Q}(Y_0)p_0$, where p_0 is the initial distribution of the state and is assumed to be known. In the context of the risk-sensitive estimation of Markov chains, the meaning of σ^γ_t will become clear.

We define the matrix

$$L(v,y) := N_{\mathbf{Y}} \cdot \overline{Q}(y)\mathcal{D}^T(v). \tag{2.9}$$

It can be shown that (see the Appendix)

$$S^\gamma(V_0,...,V_{H-1}) = E[exp(\gamma \cdot \mathcal{C}_H)] = E^\dagger[\sum_{i=1}^{N_X} \sigma^\gamma_H(i)], \tag{2.10}$$

where $E^\dagger$ is the expectation with respect to the new measure. Define the value functions $J^\gamma(\cdot, H-j) : R_+^{N_{\mathbf{X}}} \to R$, $j = H,...,H-k$, as follows:

$$J^\gamma(\sigma, H-j) := \min_{V_{H-j}...V_{H-1}} \Big\{E^\dagger\{\sum_{i=1}^{N_X} \sigma^\gamma_H(i) \mid \sigma^\gamma_{H-j} = \sigma\}\Big\}.$$

Lemma 2.3.1: Let $V_0,...,V_{k-1}$ be given. Then the value functions defined above are obtained from the following dynamic programming equations:

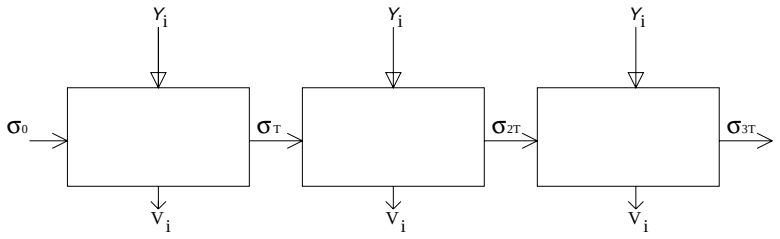

Figure 1: The T-step risk sensitive filter banks.

$$\begin{cases} \bar{J}^{\gamma}(\sigma, H) & = \sum_{i=1}^{N_{\mathbf{X}}} \sigma(i); \\ \bar{J}^{\gamma}(\sigma, H-j) & = \min_{v \in \mathbf{X}} \{E^{\dagger}[\bar{J}^{\gamma}(L(v, Y_{H-j+1}) \cdot \sigma, H-j+1)]\} \quad j = 1, 2, \ldots, H-k. \end{cases} \tag{2.11}$$

The estimation maps $\hat{\mathcal{V}}_k, ..., \hat{\mathcal{V}}_{H-1}$ obtained from (2.11) are risk optimal, i.e., $\hat{\mathcal{V}}_k(\sigma_k^{\gamma}), ..., \hat{\mathcal{V}}_{H-1}(\sigma_{H-1}^{\gamma})$ achieve the minimum in (2.5).

Proof:

We have:

$$S^{\gamma}(V_0, ..., V_{H-1}) = E^{\dagger}[\sum_{i=1}^{N_X} \sigma_H^{\gamma}(i)] = E^{\dagger}E^{\dagger}\{[\sum_{i=1}^{N_X} \sigma_H^{\gamma}(i)]|\sigma_k^{\gamma}\}.$$

Thus to obtain $\bar{S}^{\gamma}(\hat{V}_k, ..., \hat{V}_{H-1})$, it is sufficient to minimize the expectation $E^{\dagger}\{[\sum_{i=1}^{N_X} \sigma_M^{\gamma}(i)]|\sigma_k^{\gamma}\}$ which is given by the above dynamic programming equations (for a complete proof see the Appendix).

2.4 The T-step MAP risk-sensitive estimator (TMAP)

The T-step MAP risk-sensitive estimator (TMAP) is defined by the following criterion:

$$\hat{V}_{NT}, ..., \hat{V}_{(N+1)T-1} = argmin_{V_t \in X,\ t=NT,...,(N+1)T-1}\ E\ \{exp(\gamma \cdot [\sum_{t=0}^{NT-1} \rho(X_t, \hat{V}_t) + \sum_{t=NT}^{(N+1)T-1} \rho(X_t, V_t)])\}, \tag{2.12}$$

where V_t is $\mathcal{Y}_t$ measurable, T is the size of the filter and $N = 0, 1, ...,$ is the index of filtering segments. This exponential criterion is a generalization of the risk-sensitive filtering idea introduced in [5] for

the quadratic cost with the filtering performed in single steps, i.e., for T=1; we will look at this special case for TMAP in sections 3 and 4 and show that it is essentially a "greedy algorithm".

Theorem 2.4.1: The TMAP can be computed recursively by the following procedure:

1) Set $\sigma_0 = N_{\mathbf{Y}} \cdot \overline{Q}(Y_0)p_0$.

2) Given that $\sigma_{NT} = \sigma$, use the minimizing sequence of the value functions obtained from the following dynamic programming equations

$$\begin{cases} \bar{J}^{\gamma}(\sigma, T) & = \sum_{i=1}^{N_{\mathbf{X}}} \sigma(i); \\ \bar{J}^{\gamma}(\sigma, T-j) & = \min_{v \in \mathbf{X}} \{E^{\dagger}[\bar{J}^{\gamma}(L(v, Y_{NT+T-j+1}) \cdot \sigma, T-j+1)]\} \quad j = 1, 2, \ldots, T \end{cases} \tag{2.13}$$

to determine the value of the optimum estimates $\hat{V}_{NT}, ..., \hat{V}_{(N+1)T-1}$ as a function of the information state $\sigma_{NT}, ..., \sigma_{(N+1)T-1}$ obtained by (2.7).

3) Apply (2.7) once more to obtain $\sigma_{(N+1)T}$ and repeat steps (2) and (3) starting at (N+1)T.

Furthermore, for any given N as $\gamma \to 0$, TMAP (i.e. the above algorithm) reduces to MAP.

Proof:

The proof follows from repeated applications of Lemma 2.3.1. We will skip the details. The limiting result follows from the first order approximation of the exponential function and the observation that as $\gamma \to 0$, the matrix $\mathcal{D}(v) \to \mathcal{A}$ element wise. This implies that in the limit the input to each filtering step is the unnormalized conditional distribution and thus by Lemma 2.2.1 the filtering process reduces to the well known MAP estimation of HMM's.

Note that although the size of the sum $\sum_{i=0}^{NT-1} \rho(X_t, \hat{V}_t)$ increases with N, all we need to track is the information state, computed recursively. The optimal estimates $\hat{\mathcal{V}}_{NT}(\sigma_{NT}), ..., \hat{\mathcal{V}}_{(N+1)T-1}(\sigma_{(N+1)T-1})$ are measurable functions of the information state alone. Since our Markov chain is homogeneous and under the new measure the observations are i.i.d, (2.13) depends only on T and not on N. This justifies the cascade filter banks representation of Figure 1.

We point out that theorem 2.4.1 has no control counterpart. In this case, the estimators $\{V_t\}$ have no influence on the dynamics and thus estimation can be broken down into separate segments with the information state reinitialized. The same cannot be said for a controlled Markov chain due to the influence of the controller on the dynamics; the separate segments cannot be joined to represent the entire process in the above limiting sense as the "decoupling" Lemma 2.2.1 no longer holds. Note also that controlled Markov chains are not homogeneous.

3 Structural results: The filter banks and the information state

It is clear from the above that to describe the behavior of TMAP we must understand the operation of each filtering segment and understand the meaning of the information state. The key in understanding the filter's operation is the analysis of the value functions which are obtained via dynamic programming.

Lemma 3.1.1: The value functions are continuous and concave functions of the information state $\sigma \in R_+^{N_{\mathbf{X}}}$.

Proof:

Both statements are proved by induction. The continuity when j=0 is obvious. Since $\{Y_i\}$ is i.i.d., uniformly distributed and finite dimensional, then (2.13) is taking a minimum over the average of composition of functions, each of which is continuous by the continuity of the linear functions and the induction hypothesis, and is therefore continuous. For the second statement, once again the case j=0 is trivially verified. Assume concavity for j-1. Let $0 \leq \lambda \leq 1$ and $\sigma_1, \sigma_2 \in R_+^{N_{\mathbf{X}}}$; define $\tilde{\sigma} := \lambda\sigma_1 + (1-\lambda)\sigma_2$. Then by the induction hypothesis and by (2.13) we have (the first equality is shown in the Appendix):

$$\begin{aligned}
\bar{J}^{\gamma}(\tilde{\sigma}, T-j) &= \min_{v\in\mathbf{X}}\Big\{\frac{1}{N_{\mathbf{Y}}}\sum_{y=1}^{N_{\mathbf{Y}}}\bar{J}^{\gamma}(L(v,y)\cdot\tilde{\sigma}, T-j+1)\Big\} \\
&\geq \min_{v\in\mathbf{X}}\Big\{\frac{1}{N_{\mathbf{Y}}}\sum_{y=1}^{N_{\mathbf{Y}}}[\lambda\bar{J}^{\gamma}(L(v,y)\cdot\sigma_1, T-j+1) \\
&+ (1-\lambda)\bar{J}^{\gamma}(L(v,y)\cdot\sigma_2, T-j+1)]\Big\} \\
&\geq \lambda\bar{J}^{\gamma}(\sigma_1, T-j) + (1-\lambda)\bar{J}^{\gamma}(\sigma_2, T-j).
\end{aligned}$$

Next, for P a finite set of vectors in $R_+^{N_{\mathbf{X}}}$, denote by $\mathcal{O}(P)$ the set

$$\mathcal{O}(P) := \Big\{\frac{1}{N_{\mathbf{Y}}}\sum_{y=1}^{N_{\mathbf{Y}}}\alpha_y \cdot L(v,y) \mid \alpha_y \in P, v \in \mathbf{X}\Big\}. \tag{3.1}$$

Note that if P is finite so is $\mathcal{O}(P)$, since $|\mathcal{O}(P)| \leq |P|^{N_{\mathbf{Y}}} \cdot N_{\mathbf{X}}$.

Lemma 3.1.2: The value functions given by (2.13) are piecewise linear functions (hyper-planes through the origin) of $\sigma \in R_+^{N_{\mathbf{X}}}$, such that if P_{j-1} indicates the vectors in $R_+^{N_{\mathbf{X}}}$ which specify the set of hyper planes for $\bar{J}^{\gamma}(\sigma, T-j+1)$ then

$$\bar{J}^{\gamma}(\sigma, T-j+1) = \min_{\alpha\in P_{j-1}}\{\alpha\cdot\sigma\} \qquad \bar{J}^{\gamma}(\sigma, T-j) = \min_{\alpha\in\mathcal{O}(P_{j-1})}\{\alpha\cdot\sigma\}, \tag{3.2}$$

where $P_0 = \bar{\mathbf{1}} := (\sum_{k=1}^{N_\mathbf{X}} e_k)^\mathbf{T}$ and $\{e_k\}$ are the unit vectors in $R^{\mathbf{N_x}}$.

Proof:

The statement of the Lemma is readily verified for $j = 0$. Assume the Lemma holds for j-1 then piecewise linearity implies that for each vector α_0 in P_{j-1} there is a point $\sigma^* \in R^{\mathbf{N_x}}$ and a disk $d(\sigma^*, r)$ such that on this disk $\bar{J}^\gamma(\sigma, T-j+1) = \alpha_0 \cdot \sigma$. Consider a different point σ and $0 < t \leq 1$ small enough so that $t(\sigma - \sigma^*) + \sigma^* \in d(\sigma^*, r)$. Then by the concavity shown above

$$\bar{J}^\gamma(\sigma^* + t(\sigma - \sigma^*), T-j+1) = \bar{J}^\gamma((1-t)\sigma^* + t\sigma, T-j+1) \geq (1-t)\bar{J}^\gamma(\sigma^*, T-j+1) + t\bar{J}^\gamma(\sigma, T-j+1).$$

Since $\bar{J}^\gamma(\sigma^* + t(\sigma - \sigma^*), T-j+1) = \alpha_0 \cdot \sigma^* + t(\sigma - \sigma^*)$, after substitution and cancellations, we get

$$\bar{J}^\gamma(\sigma, T-j+1) \leq \alpha_0 \cdot \sigma.$$

But α_0 was arbitrary and so we have the first equality. To prove the second equality, note that

$$\begin{aligned} \bar{J}^\gamma(\sigma, T-j) &= \min_{v \in \mathbf{X}} \Big\{ \frac{1}{N_\mathbf{Y}} \sum_{y=1}^{N_\mathbf{Y}} \min_{\alpha \in P_{j-1}} \{\alpha \cdot L(v,y) \cdot \sigma\} \Big\} = \min_{v \in \mathbf{X}} \Big\{ \Big[\frac{1}{N_\mathbf{Y}} \sum_{y=1}^{N_\mathbf{Y}} \tilde{\alpha}(v,y,\sigma) \cdot L(v,y) \Big] \cdot \sigma \Big\} \\ &= \min_{\alpha \in \mathcal{O}(P_{j-1})} \{\alpha \cdot \sigma\}, \end{aligned} \tag{3.3}$$

where $\tilde{\alpha}(v, y, \sigma) \in P_{j-1}$ minimizes $\alpha \cdot L(v, y) \cdot \sigma$ in the first equality, and the last equality follows since $\alpha \cdot L(u, y) \cdot \sigma > \tilde{\alpha}(v, y, \sigma) \cdot L(u, y) \cdot \sigma$, for all $\alpha \in P_{j-1}$, $v \in \mathbf{X}$, $y \in \mathbf{Y}$, $\sigma \in R^+_{N_\mathbf{X}}$.

Lemma 3.1.3: The optimal estimates $\{\hat{V}_t\}$ are constant along rays through the origin, i.e., let $\sigma \in R_+^{N_\mathbf{X}}$ then $\hat{\mathcal{V}}_t(\sigma') = \hat{\mathcal{V}}_t(\sigma)$, for all $\sigma' = \lambda\sigma$, $\lambda > 0$.

Proof:

From Lemma 3.1.2, we see that $\bar{J}^\gamma(\sigma', T-j) = \lambda \bar{J}^\gamma(\sigma, T-j)$. The result follows from Theorem 2.4.1.

Definition: A cone in $R_+^{N_\mathbf{X}}$ is a set defined by $C_S := \{\sigma | \sigma = \lambda x, x \in S \subset R_+^{N_\mathbf{X}}, \lambda > 0\}$.

Definition: For j=1,2, ..., T and $v \in X$, let

$$\bar{J}_v^\gamma(\sigma, T-j) := E^\dagger[\bar{J}^\gamma(L(v, Y_{T-j+1}) \cdot \sigma, T-j+1)]$$

$$= \quad \frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} [\bar{J}^{\gamma}(L(v,y) \cdot \sigma, T-j+1)]. \tag{3.4}$$

The equality is shown in the Appendix.

Definition: The decision region $DR_v^j \subset R_+^{N_{\mathbf{X}}}$ for the estimate $v \in \mathbf{X}$, at the $T-j$ decision epoch, is defined as

$$DR_v^j := \{\sigma \mid \sigma \in R_+^{N_{\mathbf{X}}}, \bar{J}^{\gamma}(\sigma, T-j) = \bar{J}_v^{\gamma}(\sigma, T-j)\}. \tag{3.5}$$

It follows from the definition of $\hat{\mathcal{V}}_{NT+T-j}(\sigma)$ that $DR_v^j := \{\sigma \mid \sigma \in R_+^{N_{\mathbf{X}}}, \quad \hat{\mathcal{V}}_{NT+T-j}(\sigma) = v\}$. We say a decision is made "strictly", if it is the only possible decision.

Theorem 3.1.1: For each $v = i \in \{1, 2, ..., \mathbf{N_X}\}$ and for every $j = 1, 2, ..., T$, the decision region DR_i^j is always non-empty and includes a cone about the σ_i axis within which the decision made is (strictly) $\hat{\mathcal{V}}_{NT+T-j}(\sigma) = i$.

Proof:

We state the proof for $N_X = 2$, from which the proof for the general case will become evident for the reader. On the σ_1 axis, we have by definition

$$\bar{J}_1^{\gamma}(\sigma, T-j) = \frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} [\bar{J}^{\gamma}(N_{\mathbf{Y}} \cdot \overline{Q}(y) \mathcal{A}^T \begin{bmatrix} 1 & 0 \\ 0 & e^{\gamma} \end{bmatrix} \begin{bmatrix} \sigma_1 \\ 0 \end{bmatrix}, T-j+1)].$$

$$\bar{J}_2^{\gamma}(\sigma, T-j) = \frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} [\bar{J}^{\gamma}(N_{\mathbf{Y}} \cdot \overline{Q}(y) \mathcal{A}^T \begin{bmatrix} e^{\gamma} & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \sigma_1 \\ 0 \end{bmatrix}, T-j+1)].$$

$$\bar{J}_1^{\gamma}(\sigma, T-j) = \frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} [\bar{J}^{\gamma}(N_{\mathbf{Y}} \cdot \overline{Q}(y) \mathcal{A}^T \begin{bmatrix} \sigma_1 \\ 0 \end{bmatrix}, T-j+1)].$$

$$\bar{J}_2^{\gamma}(\sigma, T-j) = \frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} [\bar{J}^{\gamma}(N_{\mathbf{Y}} \cdot \overline{Q}(y) \mathcal{A}^T \begin{bmatrix} e^{\gamma}\sigma_1 \\ 0 \end{bmatrix}, T-j+1)].$$

Applying Lemma 3.1.2 to each term of the summation on the right-hand side of $\bar{J}_2^{\gamma}(\sigma, T-j)$, we get

$$\bar{J}_2^{\gamma}(\sigma, T-j) = \frac{e^{\gamma}}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} [\bar{J}^{\gamma}(N_{\mathbf{Y}} \cdot \overline{Q}(y) \mathcal{A}^T \begin{bmatrix} \sigma_1 \\ 0 \end{bmatrix}, T-j+1)].$$

Therefore, we can write

$$\bar{J}_2^{\gamma}(\sigma, T-j) - \bar{J}_1^{\gamma}(\sigma, T-j) = (e^{\gamma}-1)\{\frac{1}{N_{\mathbf{Y}}}\sum_{y=1}^{N_{\mathbf{Y}}}[\bar{J}^{\gamma}(N_{\mathbf{Y}} \cdot \overline{Q}(y)\mathcal{A}^T \begin{bmatrix} \sigma_1 \\ 0 \end{bmatrix}, T-j+1)]\}.$$

But $e^{\gamma} > 1$ since $\gamma > 0$ and for every j, the value functions are strictly positive being integrals of the exponential functions. Thus from the above on the σ_1 axis we have the strict inequality

$$\bar{J}_1^{\gamma}(\sigma, T-j) < \bar{J}_2^{\gamma}(\sigma, T-j)$$

which implies DR_1^j includes the σ_1 axis; fix σ on that axis, then by Lemma 3.1.1 and because sums and compositions of continuous functions are continuous, there exists a disk of positive radius in the R_+^2 metric, i.e., $d(r,\sigma) \bigcap R_+^2$, $r > 0$ such that $\bar{J}_1^{\gamma}(\sigma, T-j) < \bar{J}_2^{\gamma}(\sigma, T-j)$ for every $x \in d(r,\sigma) \bigcap R_+^2$. Therefore, Lemma 2.1.3 implies that on the cone $\mathbf{C}_{d(r,\sigma)\cap R_+^2} \subset DR_1^j$ the decision is strictly $v = 1$.

The same proof works in higher dimensions by fixing an axis σ_l and making pairwise comparisons between $\bar{J}_l^{\gamma}(\sigma, T-j)$ and $\bar{J}_k^{\gamma}(\sigma, T-j)$, $k \neq l$ along the σ_l axis. The "strict" cone around the σ_l axis will be the intersection of all cones obtained from pairwise comparisons.

In general, the boundaries among the decision regions are not of "threshold type" (unlike MAP). We give a two dimensional example. We consider the TMAP with $N_{\mathbf{X}} = 2$.

Remark: The transition cone $DR_1^j \bigcap DR_2^j$ is not, in general, of threshold type, i.e., the cone: $DR_1^j \bigcap DR_2^j$ does not degenerate to a line; we give a simple counter example.

Let $a_{11} = a_{22} = 1$ and $q_{xy} = 1/2$; then it can be shown (by straightforward induction) that the transition cones are not degenerate.

Let's look at the counter example more closely (Figure 2). The cone where the decision is strictly v=1, i.e., $R_+^2 \cap (DR_2^j)^c$ is given by $\sigma_1 > \sigma_2 \cdot exp(\gamma(j-1))$ and by symmetry $R_+^2 \cap (DR_1^j)^c$ is given by $\sigma_2 > \sigma_1 \cdot exp(\gamma(j-1))$. The transition cone (where either decision is acceptable) is given by the complement of the union of these two regions (the colored area). The value functions are given by $j+1$ hyper-planes: $\sigma_1 + exp(\gamma(j))\sigma_2$, $\sigma_1 exp(\gamma(1)) + exp(\gamma(j-1))\sigma_2$, $\sigma_1 exp(\gamma(2)) + exp(\gamma(j-2))\sigma_2$ ..., $\sigma_2 + exp(\gamma(j))\sigma_1$ on the $j+1$ cones beginning with $\sigma_1 > \sigma_2 \cdot exp(\gamma(j-1))$ and ending with $\sigma_2 > \sigma_1 \cdot exp(\gamma(j-1))$. The transition cone between them is the union of $j-1$ cones whose boundaries are lines: $exp(-(j-1)\gamma)\sigma_1 = \sigma_2$, $exp(-(j-2)\gamma)\sigma_1 = \sigma_2$,..., $\sigma_1 = \sigma_2$,..., $exp(-(j-2)\gamma)\sigma_2 = \sigma_1$, ..., $exp(-(j-1)\gamma)\sigma_2 = \sigma_1$. When j is odd, the the line $\sigma_1 = \sigma_2$ is a boundary (boundary in the sense of slope change in the cost and not decision); when j is even, it is not. (The solid cone which includes $\sigma_1 = \sigma_2$ for even values of j is meant to emphasize this). On the transition cone either decision is allowed. We can interpret this region as the zone of uncertainty. For MAP and TMAP with T=1 (we'll show this later) this region is the threshold $\sigma_1 = \sigma_2$, but as the above example shows, for TMAP with $T > 1$, it may be a non-degenerate cone. We could interpret this as reflecting the "conservative" nature of the risk-sensitive estimation. We are expanding the zone of "uncertainty" at the expense of the region of "certainty". We will show in the subsequent sections that this is

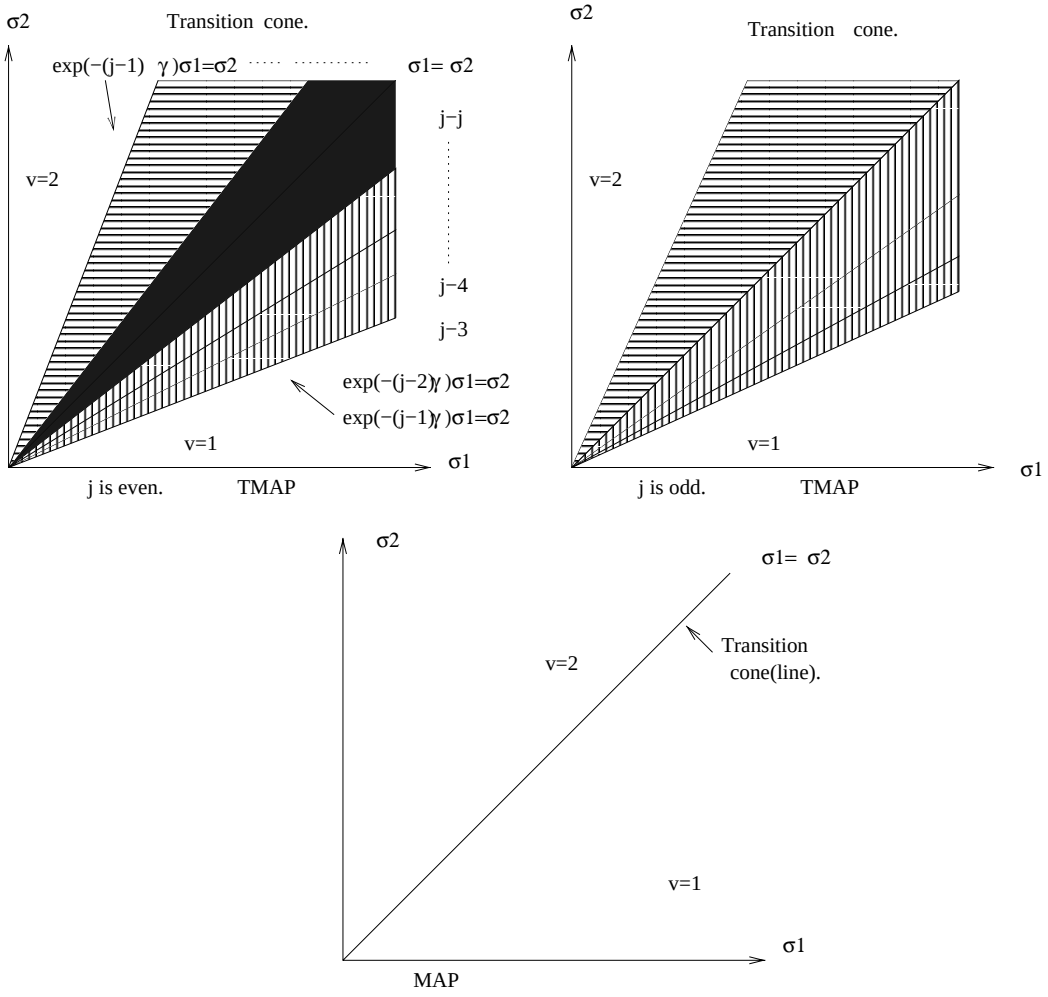

Figure 2: Decision cones for $a_{11} = a_{22} = 1, q_{xy} = 1/2$.

not always the manner in which risk-sensitivity manifests itself in the structure of decision regions. It is possible that the the transition cone remains degenerate and either of the two other cones expands at the expense of the other or the decision regions are not affected by risk sensitivity at all, i.e., they remain identical to that of MAP. In two dimensions for example, $DR_1^j = \{\sigma | \sigma_1 > \sigma_2\}$ and $DR_2^j = \{\sigma | \sigma_2 > \sigma_1\}$.

The above theorem only guarantees the existence of of non-degenerate cones around the σ_l axis but says nothing about their size. In fact, observe that in the above example the size of these cones becomes arbitrarily small as $\gamma \to \infty$ since the slopes of the lines $exp(-(j-1)\gamma)\sigma_1 = \sigma_2$ and $exp(-(j-1)\gamma)\sigma_2 = \sigma_1$, for every $j > 1$, converge to zero and infinity respectively.

Two special cases (N=0, T=M) and (T=1, N=0,..., M-1) are of interest. In both cases, the index t ranges from t=0 to t=M-1. In the first case, TMAP reduces to the exponential/sum criterion for HMM's which is the discrete and finite dimensional version of the risk-sensitive L^2 filter introduced by Speyer and others. The second case would be the MAP version of the quadratic cost risk-sensitive filtering introduced first to our best knowledge by Dey and Moore in [5]. Obviously, the structural results obtained so far apply to these special cases.

Theorem 3.1.2: Let $\mathcal{EX}(v)$ be the diagonal matrix $diag[exp(\gamma\rho(i,v))], \quad i = 1, ..., N_{\mathbf{X}}$. Then the one step TMAP decision regions are given by: $\hat{\mathcal{V}}_t(\sigma) = i$ if $\sigma_i \geq \sigma_j, \ \forall j \neq i$.

Proof:

From the definition, we have: $\mathcal{D}^T(v) = \mathcal{A}^T\mathcal{EX}(v)$ and that $\bar{J}^\gamma(\sigma, T) =< \sigma, \bar{\mathbf{1}} >$ and thus

$$\begin{aligned}
\bar{J}_v^\gamma(\sigma, T-1) &= \frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} [\bar{J}^\gamma(L(v,y)\cdot\sigma, \quad T)] \\
&= \frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} [\bar{J}^\gamma(N_{\mathbf{Y}} \cdot \overline{Q}(y)\mathcal{A}^T\mathcal{EX}(v)\cdot\sigma, T)] \\
&= \sum_{y=1}^{N_{\mathbf{Y}}} < \overline{Q}(y)(\mathcal{A}^T\mathcal{EX}(v)\cdot\sigma), \bar{\mathbf{1}} > \\
&= < (\sum_{y=1}^{N_{\mathbf{Y}}} \overline{Q}(y))(\mathcal{A}^T\mathcal{EX}(v)\cdot\sigma), \bar{\mathbf{1}} > \\
&= < I\mathcal{A}^T\mathcal{EX}(v)\cdot\sigma, \bar{\mathbf{1}} > \\
&= < \mathcal{EX}(v)\cdot\sigma, \mathcal{A}\bar{\mathbf{1}} > \quad = \quad < \mathcal{EX}(v)\cdot\sigma, \bar{\mathbf{1}} > .
\end{aligned} \tag{3.6}$$

A little calculation shows that given σ the above is minimized, if we set v equal to the index of the largest component of σ, i.e., if σ_l is the largest component of σ then $v = l$. This is precisely how the decision regions for MAP are defined.

Note that TMAP for T=1 is not reduced to MAP; although the decision regions are the same,

the information states are different. In the case of MAP, the information state is the conditional distribution, while in the case of TMAP the information state is given by (2.7). The conditional distribution has no memory of decisions made in the past while the TMAP's information state does depend on these decisions. On the other hand, when γ is very small, (2.7) becomes the unnormalized conditional distribution. We can think of TMAP's information state as the conditional distribution modified by the sequence of the decisions made in the past. This modified information state is then put through the next decision region which itself is calculated to minimize a certain cost structure based on the averaged behavior of the future sample paths: $(E^{\dagger}[\bar{J}^{\gamma}(L(v, Y_{NT+T-j+1})\cdot\sigma, T-j+1)])$. How far we look into the "averaged future" is determined by T. The information from the past is encapsulated in the information state. These ideas will become quantified in the next sections.

Theorem 3.1.3: The value function $\bar{J}^{\gamma}(\sigma, T-j)$ when $\forall(x,y),\ q_{yx} = 1/N_{\mathbf{Y}}$ is given by

$$\bar{J}^{\gamma}(\sigma, T-j) = \min_{v_{T-j},\ldots,v_{T-1}} < \sigma,\ \mathcal{EX}(v_{T-j})\cdot\mathcal{A}\cdot\mathcal{EX}(v_{T-j+1})\cdot\mathcal{A}\cdot\mathcal{EX}(v_{T-j+2})...\mathcal{A}\cdot\mathcal{EX}(v_{T-1})\cdot\bar{\mathbf{1}} > .$$

Proof:

The proof is based on the same technique used in the previous theorem and will be skipped.

In the above counter example when $\mathcal{A} = I_{2\times 2}$ and $q_{yx} = 1/2$, by the above theorem, we have

$$\bar{J}^{\gamma}(\sigma, T-j) = \min_{v_{T-j},\ldots,v_{T-1}} < \sigma, \mathcal{EX}(v_{T-j})\cdot ...\mathcal{EX}(v_{T-j})\cdot\bar{\mathbf{1}} > .$$

If we let the number of times we choose $x = 2$ be n_2, and likewise for $x = 1$, $n_1 = T - n_2$, a little algebra will show that the total cost $\bar{J}^{\gamma}(\sigma, 0)$ is given by

$$\sigma_1 exp\{\gamma n_2\} + \sigma_2 exp\{\gamma(T-n_2)\}.$$

By differentiation with respect to n_2, a few rearrangements and taking logarithms, the minimum cost is obtained when (modulo the integer parts)

$$T/2 - \frac{1}{2\gamma} log(\sigma_1/\sigma_2) = n_2;\quad 0 \le n_2 \le T. \tag{3.7}$$

This suggests that for large values of γ regardless of σ, we choose the two states an approximately equal number of times. To see why this limiting behavior occurs first let $\sigma_1 = \sigma_2 = 1/2$, then according to the decision regions, we could choose either 1 or 2. Suppose we choose 1, then the information state evolves to $(\sigma_1, \sigma_2 e^{\gamma}) = (1/2, e^{\gamma}/2)$. We could continue to choose v=1 and "move up" on the information state (successive) 2-dimensional planes. Note that the upper boundary of the transition cone is given by $\sigma_2 exp(\gamma) + exp(\gamma(j-1))\sigma_1$. Thus, going forward in time, the information state moves up while the upper boundary, $\sigma_2 exp(\gamma) + exp(\gamma(j-1))\sigma_1$, "moves down"

since going forward, $j-1$ in the above changes to $(j-1)-1=j-2$. Therefore, at about $j/2$ the information state moves out of the transition cone and enters the upper cone where we will have to choose 2 as our estimate. A similar argument shows that from that point on, the information state "moves forward" on the information state (successive) 2-dimensional planes and as the upper boundary $\sigma_2 exp(\gamma)+exp(\gamma(j-1))\sigma_1$ continues to decrease in slope, the information state σ remains on the decision region of the state 2. Hence, the decision sequence $\{1,1,1,....2,2,2\}$ with about the same number of 1's and 2's (exactly if j is even) is a possible solution. When γ is large, every point not on the axes falls into the transition cone and a similar argument shows that it follows the same decision pattern.

Now consider the case T=1. In this case the transition cone is reduced for all values of j, to the line $\sigma_1=\sigma_2$ and a similar argument shows that for large values of γ the decision is $\{1,2,1,2,1,2...\}$. This simple example provides insight as to how the two special cases differ. The decision regions for $T=M>1$ are more complex but it appears that this allows for a "smoother" solution. We will return to this subject in section 4.

Definition: A completely symmetric HMM is defined as an HMM whose transition probability matrix is symmetric in the sense that $a_{ii}=1-\epsilon$, $a_{ji}=\epsilon/(N_{\mathbf{X}}-1)$ for all $j\neq i$, $N_{\mathbf{X}}=N_{\mathbf{Y}}$ (so that Q is a square matrix) and $q_{yx}=q$ for $x=y$ and otherwise $q_{yx}=(1-q)/(N_{\mathbf{Y}}-1)$.

Note that the discrete metric which induces MAP (and in turn TMAP) is symmetric in the sense that $d(x_1,y_1)=d(x_2,y_2)\ \ \forall x_i\neq y_i$. Therefore, for a completely symmetric HMM, all criteria upon which the determination of the value functions at a particular σ depend are symmetrically defined and thus the value functions are unchanged under the permutation of components of σ.

In two dimensions, a completely symmetric HMM is given by

$$\mathcal{A}=\begin{bmatrix}1-\epsilon & \epsilon\\ \epsilon & 1-\epsilon\end{bmatrix}\quad Q=\begin{bmatrix}q & 1-q\\ 1-q & q\end{bmatrix}. \tag{3.8}$$

In the above example we have set $\epsilon=0$ and $q=1/2$.

Theorem 3.1.4: For a completely symmetric HMM, the value functions $\bar{J}^{\gamma}(\sigma,T-j)\}$ restricted to the simplex (the set $\{\sigma|\sigma_1+\sigma_2....+\sigma_{N_{\mathbf{X}}}=1\}$) have their global maximum at the center of the simplex, i.e., at $\sigma_1=\sigma_2=...=\sigma_{N_{\mathbf{X}}}$.

Proof:

The restriction of a concave function to a convex region is certainly concave. The simplex is a compact set, and so the value function restricted to the simplex is concave and has a maximum point. Let's begin with the 1-simplex. By the complete symmetry of the system, a value function will be symmetric with respect to the center of the simplex, namely around the point $\sigma_1=\sigma_2=1/2$. To see this let $\sigma=(\sigma_1,\sigma_2)$, $\sigma_1+\sigma_2=1$ be an arbitrary point on the simplex. Consider this point and its mirror image $\sigma^{'}=(\sigma_2,\sigma_1)$ obtained from σ by a permutation of its components $(\sigma_1\rightarrow\sigma_2,\sigma_2\rightarrow\sigma_1)$ which leaves the value function unchanged. Because the value functions are concave we can write

$$\bar{J}^{\gamma}(1/2, T-j) = J((\sigma+\sigma^{'})/2, T-j) \geq 1/2\{\bar{J}^{\gamma}(\sigma, T-j) + \bar{J}^{\gamma}(\sigma^{'}, T-j)\} = \bar{J}^{\gamma}(\sigma, T-j).$$

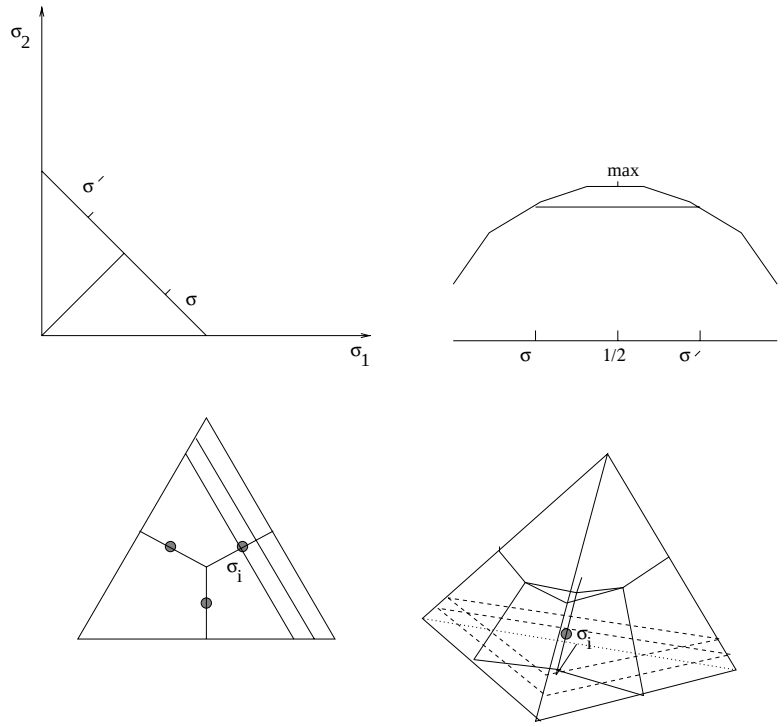

Figure 3: The value functions' global maximum under complete symmetry.

Therefore, the center of the simplex is the global maximum.

In general, for an $(N_{\mathbf{X}}-1)$-simplex, consider $N_{\mathbf{X}}$ permutations of σ

$$\sigma_1 \rightarrow \sigma_2, \sigma_2 \rightarrow \sigma_3 ..., \sigma_i \rightarrow \sigma_{i+1}, \quad \sigma_{N_{\mathbf{X}}} \rightarrow \sigma_1$$

repeated $N_{\mathbf{X}}$ times and generating $N_{\mathbf{X}}$ points $\sigma^k \quad k = 1, ..., N_{\mathbf{X}}$. Then by symmetry and by Jensen's inequality:

$$\begin{aligned} \bar{J}^{\gamma}([\sigma^1 + ... + \sigma^{N_{\mathbf{X}}}]/N_{\mathbf{X}}, T-j) &\geq 1/N_{\mathbf{X}} \cdot \bar{J}^{\gamma}(\sigma^1, T-j) + ... + 1/N_{\mathbf{X}} \cdot \bar{J}^{\gamma}(\sigma^{N_{\mathbf{X}}}, T-j) \\ &= \bar{J}^{\gamma}(\sigma, T-j). \end{aligned}$$

Let $\sigma^* = [\sigma^1 + ... + \sigma^{N_{\mathbf{X}}}]/N_{\mathbf{X}}$. Then $\sigma_i^* = 1/N_{\mathbf{X}} \sum_{l=1}^{l=N_{\mathbf{X}}} \sigma_l = 1/N_{\mathbf{X}}$ since σ is a point of the simplex. Therefore, σ^* is the center of the simplex and by the above the maximal point. This completes the proof of the theorem.

In fact a stronger version of the theorem holds. We will only outline the proof. Consider the 2-simplex which is bounded by three 1-simplices. Complete symmetry of system implies that along these simplices and all the line segments parallel to them, the value function has its maximum at the middle of such lines. Thus the maximum on the whole simplex can be found by considering only the values on the set of the mid-points of all these line segments joined at the center of the simplex.

Similarly, a $(N_{\mathbf{X}}-1)$-simplex is bounded by $N_{\mathbf{X}}$, $(N_{\mathbf{X}}-2)$-simplices. Along these sub-simplices and hyper-planes parallel to them and restricted to the simplex, the value functions have their maximum at the central points, reaching their global maximum at the center of the $(N_{\mathbf{X}}-1)$-simplex. Note that the maximum point need not be unique. In fact, in the above example for even values of j the maximum point is not unique.

If the assumption of complete symmetry is removed, Theorem 3.1.4 no longer holds. We will show this shortly.

4 Qualitative analysis of the sample paths and the information state

4.1 Sample path error smoothing

It is often said that risk sensitive estimators take into account "the higher order moments" of the cost. How do higher order moments manifest themselves in the behavior of the estimator? To understand what risk sensitivity take into account, we first explain what MAP does not.

Given a single observation and a single random variable, the MAP estimation is a reasonable thing to do: minimize the measure of the set where the estimate and the random variable disagree. But a problem arises when we consider a stochastic process (in our case a time series). The obvious thing to do then is to minimize the expected summation of the error functions for the whole time series. As shown before, this reduces back to finding the MAP estimate at each instant of time. Thus, at each instant of time our decision is not affected by our past or future decisions. This makes the MAP estimation insensitive to the accumulation of errors along sample paths. To see this evaluate X_k and the estimate $\hat{X}_k$ at some fixed value of $\omega \in \Omega$ to produce a sample path or a realization of the respective time series and its estimate. The sequence of decisions $\hat{X}_k \quad k = 0, 1, ...N$ partitions the sample space into 2^N subsets according to a binary tree with the branching criterion $X_k = \hat{X}_k$ or $X_k \neq \hat{X}_k$. Each ω belongs to exactly to one of these subsets. Some sample paths may end up on branches along which estimation errors accumulate for long stretches of time. Now consider TMAP with T=1 (which is the Dey-Moore filter equivalent for MAP). The exponential function turns the sum of the error functions into a product, the value of this product up to the last decision made for each fixed ω is $e^{\gamma \cdot m}$, where m counts the number of times an error has been made in estimation of that particular sample path. This product is then multiplied by either 1 or $e^{\gamma} > 1$ depending on the next decision made based on minimizing the measure of the error set at the present time *and* taking note of the errors made in the past. Thus the errors made in the past become more uniformly distributed over all sample paths. The information state condenses this error history and

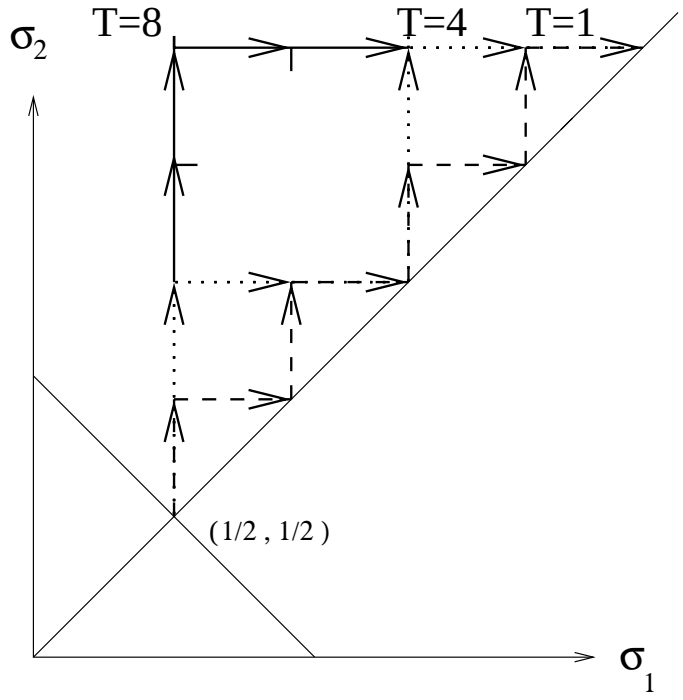

Figure 4: sample paths for $a_{11} = a_{22} = 1, q_{xy} = 1/2$ starting from (1/2,1/2).

makes the computation recursive. For this reason perhaps *sample path error smoothing estimators* could be an alternative name for the risk-sensitive estimators. Dey-Moore filter is (in a sense) a greedy algorithm which does not consider (on the average) the accumulation of errors in the future but has the important benefit of computational simplicity. The exponential/sum filter does so for the entire path and in general TMAP looks into the future for T steps. In Figure 4, observe that a particular sample path of our familiar example becomes more oscillatory as T is made smaller. Our simulations show that this is a general behavior no matter how complicated the chain: the smaller the filter size T, the bigger is the burden of error smoothing on the T next estimates. But making T large comes at the cost of increased computational complexity.

The underlying mechanism of all these estimators is the coupling of the estimation errors in the product form. In an upcoming paper, we will show that for HMM's this coupling is possible in a broader context.

4.2 Risk-sensitivity, information and mixing

Through the sample path perspective, we can explain the behavior of the risk-sensitive estimators. In HMM's all sample paths pass through a finite number of states. We can think of the transition probabilities as determining a flow in the system. So far the only example we have considered was a non-mixing dynamical system. The transition probabilities were set to be zero and there was no flow between the states. This is important from the error smoothing point of view, for as the flow passes through the states, so does the history of the errors. If sample paths that have

accumulated estimation error remain in a particular state, the estimator will be "attracted" to the state in order to relieve the error accumulated there. This explains the oscillatory behavior of our two state example in which no flow between the states was allowed. On the other hand, if the transition probabilities are non-zero this "attraction" is somewhat relieved; we have verified through simulations that mixing indeed inhibits the oscillatory behavior. This will also have an effect on the decision regions as we will see shortly. But if we go through a state "too quickly", we cannot use that state to smoothen the error accumulated in the path effectively. Both these cases lead to certain type of singularities in the decision regions.

The second issue is the role of information. If we expect to receive good information about the system in the future, which will in turn reduce the error accumulation, we are likely to be less conservative at the present about our decisions. This means that we expect TMAP's decision regions to become less conservative and look more like MAP's under increased availability of information. This too will be shown in the following example.

We will study the decision regions for T=2 TMAP for an HMM with $N_{\mathbf{X}} = N_{\mathbf{Y}} = 2$ and

$$\mathcal{A} = \begin{bmatrix} 1/2 & 1/2 \\ \delta & 1-\delta \end{bmatrix}; \;\; Q = \begin{bmatrix} 1/2+I & 1/2-I \\ 1/2-I & 1/2+I \end{bmatrix}. \tag{4.1}$$

The parameter I controls the availability of information. When $I = 0$, no information is available (the case of pure prediction) and as $I \to 1/2$, the HMM becomes perfectly observable. The parameter δ determines the transition probabilities of the second state and in particular $\delta = 0$ will make the Markov chain non-mixing.

As shown before for T=1, the decision regions are identical to those of MAP. First let $I = 0$, then it can be shown that for T=2, the decision regions of j=2 (the first stage of the two step filter) are of the threshold type determined by a transition line L with the slope $m(\delta)$ followed by the equi-partition decision regions (identical to the decisions regions of MAP). The decision regions of the first stage are given by $\sigma_2 < m(\delta)\sigma_1$ choose 1 and if $\sigma_2 > m(\delta)\sigma_1$ choose 2. The slope $m(\delta)$ is given by

$$m(\delta) = \begin{cases} \frac{e^\gamma+1}{2} \cdot \frac{1}{\delta e^\gamma + 1 - \delta} & \delta < 1/2; \\ \\ \frac{e^\gamma+1}{2} \cdot \frac{1}{(1-\delta)e^\gamma + \delta} & \delta > 1/2. \end{cases}$$

Simple calculations show that the slope is always greater than or equal to 1 (only when $\delta = 1/2$), so that the decision region of the first state is enlarged at the expense of the second. As expected when $\gamma \to 0$, the decision regions equalize. When $\gamma \to \infty$, the slope is given by

$$m(\delta) = \begin{cases} \frac{1}{2\delta} & \delta < 1/2; \\ \\ \frac{1}{2(1-\delta)} & \delta > 1/2. \end{cases}$$

When either $\delta = 0$ or $\delta = 1$, the slope becomes infinite. These are the singularities that we

mentioned earlier. The equalization of the two regions at $\delta = 1/2$ is a general property which holds true even when no constraint is put on the available information as the following theorem demonstrates:

Theorem 4.2.1: Consider the HMM described by (2.1). The risk-sensitive decision regions are equalized under uniform flow: $a_{ij} = 1/N_{\mathbf{X}} \ \ \forall(i,j)$. Furthermore, TMAP reduces to MAP for every choice of T and γ.

Proof:

Fix T and γ. Consider the calculations of the decision regions according to (2.13). For j=1, we already saw that the decision regions are equalized. For $j > 1$ we can write:

$$\begin{aligned}
\bar{J}^{\gamma}(\sigma, T-j) &= \min_{v\in\mathbf{X}}\{E[\bar{J}^{\gamma}(L(v, Y_{T-j+1})\cdot\sigma, T-j+1)]\} \\
&= \min_{v\in\mathbf{X}}\{\frac{1}{N_{\mathbf{Y}}}\sum_{y=1}^{N_{\mathbf{Y}}}[\bar{J}^{\gamma}(N_{\mathbf{Y}}\cdot\overline{Q}(y)\mathcal{A}^{T}\mathcal{E}\mathcal{X}(v)\cdot\sigma, T-j+1)]\} \\
&= \min_{v\in\mathbf{X}}\{\frac{1}{N_{\mathbf{Y}}}\sum_{y=1}^{N_{\mathbf{Y}}}[\bar{J}^{\gamma}(N_{\mathbf{Y}}\cdot\overline{Q}(y)\{<\mathcal{E}\mathcal{X}(v)\cdot\sigma, \bar{\mathbf{1}}>\}\mathcal{A}^{T}\cdot\bar{\mathbf{1}}, T-j+1)]\} \\
&= \min_{v\in\mathbf{X}}\{<\mathcal{E}\mathcal{X}(v)\cdot\sigma, \bar{\mathbf{1}}>\frac{1}{N_{\mathbf{Y}}}\sum_{y=1}^{N_{\mathbf{Y}}}[\bar{J}^{\gamma}(N_{\mathbf{Y}}\cdot\overline{Q}(y)\mathcal{A}^{T}\cdot\bar{\mathbf{1}}, T-j+1)]\} \\
&= \min_{v\in\mathbf{X}}\{<\mathcal{E}\mathcal{X}(v)\cdot\sigma, \bar{\mathbf{1}}>\}\{\frac{1}{N_{\mathbf{Y}}}\sum_{y=1}^{N_{\mathbf{Y}}}[\bar{J}^{\gamma}(N_{\mathbf{Y}}\cdot\overline{Q}(y)\mathcal{A}^{T}\cdot\bar{\mathbf{1}}, T-j+1)]\}.
\end{aligned}$$

The second term does not depend on v and so the minimization reduces to

$$\min_{v\in\mathbf{X}}\{<\mathcal{E}\mathcal{X}(v)\cdot\sigma, \bar{\mathbf{1}}>\}.$$

This shows that the decision regions are equalized. Similar techniques applied to the evolution of the information state show that TMAP's and MAP's information states, under uniform flow, are multiples of each other. This fact together with Lemma 3.1.3 and the above result regarding the decision regions completes the proof.

In the above result, the observation matrix $\overline{Q}(y)$ plays no role under the assumption of uniform flow in the determination of the decision regions. But this is the exception. In general, the availability of information appears to have an elegant relation to the structure of the decision regions as the following shows.

Proposition 4.2.2: In 4.1, let $\delta = 0$ and $I \geq \frac{1}{2(1+e^{-\gamma})}$; then the decision regions for TMAP, T=2 are equalized.

Proof:

The proof follows from the solution of a system of simultaneous inequalities defined by (2.13) with the constraints $\delta = 0$ and $I \geq \frac{1}{2(1+e^{-\gamma})}$. We skip the tedious algebra.

As we mentioned earlier, this does not imply that TMAP for T=2 reduces to MAP because the recursions governing the evolution of the information states for the two cases are different. But if for some T the decision regions are equalized then TMAP with filter size T does reduce to the TMAP with filter size T=1. This would be significant both conceptually and computationally if we could determine conditions under which the decision regions are equalized. Note in the above for computational reasons, we had constrained the observation matrix to be symmetric. This produces only a sufficient condition as stated in Proposition 4.2.2. The right measure of the minimum quantity of information needed must be free of such constraints (for example, Shannon's mutual information among the states and the observations). Observe that in the above example, the amount of needed information grows with increasing γ. Clearly $\frac{1}{2(1+e^{-\gamma})} \to 1/2$ as $\gamma \to \infty$ which implies under infinite risk, we need perfect observation to equalize the decision regions.

We proved earlier that the maximum expected error "cost to go" for a completely symmetric HMM occurs at the center of the simplex. This is not true in general. It can be shown that the maximum cost in the first case of the above example (I=0) occurs along the transition line L, with slope $m(\delta)$, which does not cross the center of the simplex $\sigma = (1/2, 1/2)$ unless $\delta = 1/2$, a special case of complete symmetry with $q = \epsilon = 1/2$. When $\delta = 0$ as $\gamma \to \infty$, this maximum point is moved arbitrarily closer to $\sigma = (0, 1)$. This is another example of what we termed a singularity under infinite risk.

We saw that under the condition of uniform flow TMAP reduces to MAP, i.e., the estimated process based on the discrete metric with the assumption of uniform flow is invariant under risk-sensitivity. We may be led to believe that perhaps for large values of γ, risk-sensitivity tends to move the estimator toward this stable invariance, making the estimator look more and more "uniform". One has to be careful about what this means. In fact, risk-sensitivity tends to increase oscillatory behavior and not relieve it. A more conservative estimator tends to move around the state space more rapidly from sample path to sample path and not for too long "trust" the correctness of the sample path it may be following. It is in this sense that the estimates are made more "uniform".

Finally, we point out that many of the results of this paper depend on the properties of the discrete metric (used to define MAP) which is not the natural metric for R^n. Therefore, our structural results do not directly illuminate the linear-Gaussian risk-sensitive estimation case. However, the intuition gained in the discrete finite dimensional setting about the behavior of the sample paths may lead to a better understanding of the linear-Gaussian risk-sensitive estimator as well.

Appendix (proof of Lemma 2.3.1)

We use the simplex representation of Markov chains developed in [1] by mapping the states of the chain onto the unit vectors of $R_{N_{\mathbf{X}}}$ denoted by $e_1 .... e_{N_{\mathbf{X}}}$. First, we define an information state and show that it evolves according to (2.7):

$$< \sigma_t, e_h > := E^{\dagger}[\lambda_t exp\{\gamma \textstyle\sum_{i=0}^{t-1} \rho(X_i, V_i)\} < X_t, e_h > |\mathcal{Y}_t].$$

Therefore,

$$\begin{aligned} < \sigma_{t+1}, e_h > &= E^{\dagger}[\ \lambda_{t+1}\ exp\{\gamma \sum_{i=0}^{t} \rho(X_i, V_i)\} < X_{t+1}, e_h > |\mathcal{Y}_{t+1}] \\ &= N_{\mathbf{Y}} E^{\dagger}[q_{X_{t+1},Y_{t+1}}\ \lambda_t\ exp\{\gamma \sum_{i=0}^{t} \rho(X_i, V_i)\} < \mathcal{A}^T X_t + Z_{t+1}, e_h > |\mathcal{Y}_{t+1}], \end{aligned}$$

where Z_{t+1} is a martingale increment defined by $Z_{t+1} := \mathcal{A}^T X_{t+1} - X_t$. Using the linearity of the inner-product and the expectation and the i.i.d properties of the observations under the new measure, it can be shown that Z_{t+1} drops out of the above expression to give

$$\begin{aligned} < \sigma_{t+1}, e_h > &= N_{\mathbf{Y}} E^{\dagger}[q_{e_h,Y_{t+1}}\ \lambda_t\ exp\{\gamma \sum_{i=0}^{t} \rho(X_i, V_i)\} < X_t, \mathcal{A}.e_h > |\mathcal{Y}_{t+1}] \\ &= N_{\mathbf{Y}} E^{\dagger}[\sum_{j=1}^{N_{\mathbf{X}}} q_{e_h,Y_{t+1}}\ \lambda_t\ exp\{\gamma \sum_{i=0}^{t} \rho(X_i, V_i)\} a_{jh} < X_t, e_j > |\mathcal{Y}_{t+1}] \\ &= N_{\mathbf{Y}} \sum_{j=1}^{N_{\mathbf{X}}} q_{e_h,Y_{t+1}} E^{\dagger}[\ \lambda_t\ exp\{\gamma \sum_{i=0}^{t} \rho(X_i, V_i)\} a_{jh} < X_t, e_j > |\mathcal{Y}_{t+1}] \\ &= N_{\mathbf{Y}} \sum_{j=1}^{N_{\mathbf{X}}} q_{e_h,Y_{t+1}} exp\{\rho(e_j, V_t)\} a_{jh} E^{\dagger}[\ \lambda_t\ exp\{\gamma \sum_{i=0}^{t-1} exp\{\rho(X_i, V_i)\} < X_t, e_j > |\mathcal{Y}_t, Y_{t+1}]. \end{aligned}$$

But since Y_{t+1} is independent of the sigma algebra generated jointly by $\mathcal{Y}_t$ and all the random variables within the expectation, it too drops out to give

$$N_{\mathbf{Y}} \sum_{j=1}^{N_{\mathbf{X}}} q_{e_h,Y_{t+1}} exp\{\gamma \rho(e_j, V_t)\} a_{jh} < \sigma_t, e_j >$$

which results in (2.7) when put in the matrix form. The initial condition is given by

$$\begin{aligned} < \sigma_0, e_h > &:= E^{\dagger}[\lambda_0 < X_0, e_h > |\mathcal{Y}_0] \\ &= N_{\mathbf{Y}} E^{\dagger}[q_{X_0,Y_0} < X_0, e_h > |Y_0] \\ &= N_{\mathbf{Y}} q_{e_h,Y_0} E^{\dagger}[< X_0, e_h > |Y_0] \\ &= N_{\mathbf{Y}} q_{e_h,Y_0} E^{\dagger}[< X_0, e_h >]. \end{aligned}$$

It is not hard to show that $E^{\dagger}[< X_0, e_h >] = E[< X_0, e_h >] =< p_0, e_h >$ where p_0 is the initial distribution of the state and so

$$\sigma_0 = N_{\mathbf{Y}} \cdot \overline{Q}(Y_0)p_0.$$

The information state is all we need to determine the optimal cost because of the following equality:

$$E[exp(\gamma \cdot \mathcal{C}_H)] = E^\dagger[\lambda_H exp(\gamma \cdot \mathcal{C}_H)] = E^\dagger[\sum_{i=1}^{N_X} \sigma_H^\gamma(i)].$$

The first equality is an application of discrete-time Girsanov Theorem ([see Theorem 3.2 in [1]]) and the second follows directly from the definition of the information state as a conditional expectation and the observation that $\sum_{i=1}^{N_X} < X_t, e_i >= 1$ since in the simplex representation $X_t \in \{e_1, ...e_{N_X}\}$.

Let $\mathcal{W}_{i,j}$ denote the set of estimates $V_i, ..., V_j$ where each V_i is $\mathcal{Y}_i$ measurable. With the notation and definitions of section 2.3 write

$$\hat{S}^\gamma(t) := inf_{\ \mathcal{W}_{t,H-1}} E^\dagger\{[\sum_{i=1}^{N_X} \sigma_H^\gamma(i)]|\mathcal{Y}_t\}. \qquad (*)$$

Define recursively the functions

$$\begin{cases} \bar{J}^\gamma(\sigma, H) & = \sum_{i=1}^{N_{\mathbf{X}}} \sigma(i); \\ \bar{J}^\gamma(\sigma, t) & = \min_{v \in \mathbf{X}}\{E^\dagger[\bar{J}^\gamma(L(V_t, Y_{t+1}) \cdot \sigma_t, t+1)|\sigma_t = \sigma, V_t = v]\} \quad t = k, k+1, \ldots, H-1. \end{cases} \qquad (**)$$

Now assume that for t+1,..., H-1, $(*)$ satisfies the above dynamic programming equation, i.e., assume that $inf_{\ \mathcal{W}_{t+1,H-1}} E^\dagger\{[\sum_{i=1}^{N_X} \sigma_H^\gamma(i)]|\mathcal{Y}_{t+1}\} = \bar{J}(\sigma_{t+1}^\gamma(\mathcal{Y}_{t+1}), t+1)$ with a choice of minimizing estimates $\hat{\mathcal{V}}_{t+1}(\sigma_{t+1}^\gamma)....\hat{\mathcal{V}}_{H-1}(\sigma_{H-1}^\gamma)$ obtained from $(**)$. We will show that the same statement holds true for t,...H-1:

$$\begin{aligned} \hat{S}^\gamma(t) &= inf_{\ \mathcal{W}_{t,H-1}} E^\dagger\{[\sum_{i=1}^{N_X} \sigma_H^\gamma(i)]|\mathcal{Y}_t\} \\ &= inf_{\ \mathcal{W}_{t,H-1}} E^\dagger\{E^\dagger\{[\sum_{i=1}^{N_X} \sigma_H^\gamma(i)]|\mathcal{Y}_{t+1}\}|\mathcal{Y}_t\} \\ &= inf_{\ \mathcal{W}_{t,t}} E^\dagger inf_{\ \mathcal{W}_{t+1,H-1}} E^\dagger\{\{[\sum_{i=1}^{N_X} \sigma_H^\gamma(i)]|\mathcal{Y}_{t+1}\}|\mathcal{Y}_t\} \\ &= inf_{\ \mathcal{W}_{t,t}} E^\dagger\{\bar{J}(\sigma_{t+1}^\gamma(\mathcal{Y}_{t+1}), t+1)|\mathcal{Y}_t\} \\ &= inf_{\ \mathcal{W}_{t,t}} E^\dagger\{\bar{J}(L(V_t, Y_{t+1})\sigma_t^\gamma(\mathcal{Y}_t),\ t+1)|\sigma_t^\gamma(\mathcal{Y}_t), V_t\} \end{aligned}$$

$$= \quad \bar{J}(\sigma_t^\gamma(\mathcal{Y}_t),\ t).$$

The interchange of the expectation and the infimum follows from Lemma 16.A.5 in [11]. (Since we are considering the finite dimensional setting, it could also be justified by writing out the conditional expectation as a finite sum in which case the exchange would be obvious.) The next two steps are true by the induction hypothesis and by the integral characterization of the conditional expectation from which one can show that in general $E[f(W, g(Z))|Z = z] = E[f(W, g(z))] = E[f(W, g(Z))|g(Z) = g(z)]$ if random vectors Z and W are independent (see for example chapter II of [12]). The last step is by $(**)$. The case t=H-1 is easily verified. Thus, by induction, the claim holds for k,...H-1. Showing that each minimizing V_t is σ_t^γ measurable can be done explicitly. The pair (σ_t^γ, V_t) is measurable with respect to $\mathcal{Y}_t$ which is independent of Y_{t+1} with a uniform distribution for Y_{t+1}. Thus $(**)$ reduces to the minimization of the following sum:

$\frac{1}{N_{\mathbf{Y}}} \sum_{y=1}^{N_{\mathbf{Y}}} [\bar{J}^\gamma(L(V_t, y) \cdot \sigma_t^\gamma, t+1)]$.

It is now obvious that the minimizing V_t must be a function of σ_t.

The above shows that the minimizing sequence of the estimates depend only on the information state. We can write

$$inf\ \mathcal{W}_{t,H-1} E^\dagger\{[\sum_{i=1}^{N_X} \sigma_H^\gamma(i)]|\sigma_t\} = E^\dagger\{inf\ \mathcal{W}_{t,H-1} E^\dagger\{[\sum_{i=1}^{N_X} \sigma_H^\gamma(i)]|\mathcal{Y}_t\}|\sigma_t\} = E^\dagger\{\bar{J}(\sigma_t^\gamma(\mathcal{Y}_t),\ t)|\sigma_t\} = \bar{J}(\sigma_t^\gamma(\mathcal{Y}_t))$$

which shows

$$inf\ \mathcal{W}_{t,H-1} E^\dagger\{[\sum_{i=1}^{N_X} \sigma_H^\gamma(i)]|\sigma_t\} = inf\ \mathcal{W}_{t,H-1} E^\dagger\{[\sum_{i=1}^{N_X} \sigma_H^\gamma(i)]|\mathcal{Y}_t\}.$$

It is straightforward to verify that the estimators, $\hat{\mathcal{V}}_k(\sigma_k^\gamma), ..., \hat{\mathcal{V}}_{H-1}(\sigma_{H-1}^\gamma)$ are risk optimal, i.e., they achieve the minimum in (2.5).

Changing the index of the dynamic programming according to t= H-j (for easing our notation in the later sections) finishes the proof of the Lemma.

References

[1] R. J. Elliot, L. Aggoun, and J. B. Moore, *Hidden Markov Models: Estimation and control,* Springer, Berlin, Heidelberg, New York, 1994.

[2] M. R. James, J. S. Baras, and R. J. Elliot, "Risk sensitive Control and dynamic games for partially Observed Discrete-Time Nonlinear Systems," *IEEE Trans. Aut. Control,* vol. 39 (4), pp. 780-792, 1994.

[3] J. S. Baras and M. R. James, "Robust and Risk-sensitive Output Feedback Control for Finite State Machines and Hidden Markov Models," *Institute for Systems Research, University of Maryland at college Park, Technical Research Reports* T. R. 94-63.

[4] J. L. Spyer, C. Fan and R. N. Banava, "Optimal Stochastic Estimation with Exponential Cost Criteria," *Proceedings of 31th conference on decision and control,* pp. 2293-2298, 1992.

[5] S. Dey and J. Moore, "Risk sensitive filtering and smoothing for hidden Markov models," *Systems and control letters,* vol. 25, pp. 361-366, 1995.

[6] S. Dey and J. Moore, "Risk sensitive filtering and smoothing via reference probability methods," *IEEE Trans. Aut. Control,* vol. 42, no.11 pp. 1587-91, 1997.

[7] R. K. Boel, M. R. James, and I. R. Petersen, "Robustness and risk-sensitive filtering," *Proceedings of the 36th IEEE Conference on Decision and Control,* Cat. No. 97CH36124, 1997.

[8] H. Vincent Poor, *An introduction to signal detection and estimation*, Springer Verlag, 1994.

[9] P. Whittle, *Risk-sensitive optimal control*, John Woley & sons, 1990.

[10] E. Fernandez-Gaucherand and S. I. Marcus, "Risk-sensitive optimal control of Hidden Markov Models: a case study," IEEE Transactions on Automatic Control, vol. 42, no. 10, pp. 1418-1422, 1997.

[11] R. J. Elliot, *Stochastic Calculus and Applications,* New York: Springer-Verlag, 1982.

[12] A. N. Shiryayev, *probability,* New York: Springer-Verlag, 1984.