

RISK SENSITIVE GENERALIZATION OF MINIMUM VARIANCE ESTIMATION AND CONTROL *

John B Moore^{**}; Robert J Elliott^{***} and Subhrakanti Dey^{**}

**The authors wish to acknowledge the funding of the activities of the Cooperative Research Centre for Robust and Adaptive Systems by the Australian Commonwealth Government under the Cooperative Research Centres Program, and NSERC grant A7964*

***Department of Systems Engineering, Research School of Information Sciences and Engineering, The Australian National University, Canberra, ACT 0200 AUSTRALIA*

****Department of Mathematics, University of Alberta, Edmonton, Alberta, Canada, T6G 2G1*

Abstract. In this paper, the risk-sensitive nonlinear stochastic filtering problem is addressed in both continuous and discrete-time for quite general finite-dimensional signal models, including also discrete state hidden Markov models (HMMs). The risk sensitive estimates are expressed in terms of the so-called information state of the model given by the Zakai equation which is linear. In the linear Gaussian signal model case, the risk-sensitive (minimum exponential variance) estimates are identical to the minimum variance Kalman filter state estimates, and are thus given by a finite dimensional estimator. The estimates are also finite dimensional for discrete-state HMMs, but otherwise, in general, are infinite dimensional. In the small noise limit, these estimates (including the minimum variance estimates) have an interpretation in terms of a worst case deterministic noise estimation problem given from a differential game. The related control task, that is the risk-sensitive generalization of minimum-variance control is studied for the discrete-time models. This is motivated by the need for robustness in the widely used (risk neutral) minimum variance control, including adaptive control, of systems which are minimum phase, that is having stable inverses.

Key Words. Risk-sensitive filtering, Minimum variance estimation, Small noise limit, Minimum variance control

1. INTRODUCTION

In optimal filtering, the usual index optimized is a state estimation error variance. This can be referred to as H_2 filtering. This is appropriate when the (stochastic) signal model is known precisely but when there is uncertainty of the model dynamics and noise, there is a case for achieving robust filtering which is acceptable for a range of models. This is the motivation for so-called H_∞ filtering which has the interpretation in terms of a minimizing estimation error in a worst case noise scenario. Risk-sensitive filtering is a more general robust/optimal filtering approach than H_2 or H_∞ filtering. It minimizes the expected value of the exponential of an (typically quadratic) estimation error cost, weighted by a risk-sensitive parameter.

A linear risk sensitive problem has been studied in (Speyer *et al.*, 1992) and more general nonlinear problems studied in (Dey and Moore, 1995). The latter paper tackles the risk sensitive estimation problem using reference probability methods of (Elliott *et al.*, 1994) where the index works with the sum of quadratic estimation errors to the present.

In this paper, an alternative simpler risk-sensitive

index to that of (Dey and Moore, 1995) is studied. For this index the risk-sensitive filter for very general nonlinear stochastic models is seen to be a simple augmentation of the information state filter (the linear Zakai equations). For linear Gaussian models, the risk sensitive filter is identical to the H_2 Kalman filter and is finite dimensional. Otherwise, in the limit as the risk-sensitive parameter approaches zero, the risk sensitive filter becomes a risk-neutral filter, typically the minimum error variance H_2 filter. In the small noise limit, the risk-sensitive filtering problem has the interpretation of a deterministic differential game, as does the H_∞ problem.

Results are developed here for quite general nonlinear models in continuous time and for states in a continuous range. Corresponding results are included for the discrete-time case and the discrete-state case for completeness. Our model class requires the nonlinearities to be linearly bounded and the index class to be quadratically bounded, rather than simply bounded as in the work of (James *et al.*, 1994) for related control problems.

Closely related problems of smoothing and prediction are addressed briefly. We also present re-

sults for a robust version of minimum variance control (including adaptive control) for minimum phase plants, namely *risk-sensitive generalization of minimum variance control*. In Section 2, the theory is spelt out for risk-sensitive filtering, and in Section 3, results are presented for the risk-sensitive version of minimum variance control.

2. RISK-SENSITIVE ESTIMATION

Estimation Problem Formulation: Recall that minimum variance estimation of a state at time t , denoted $x_t \in \mathbb{R}^n$ based on measurements $y_t \in \mathbb{R}^m$ up to time t , denoted \mathcal{Y}_t , is given from the definition

$$\hat{x}_{t|t} \in \arg \min_{\xi \in \mathbb{R}^n} E\left[\frac{1}{2}(x_t - \xi)'Q(x_t - \xi) \mid \mathcal{Y}_t\right] \quad (1)$$

with $Q \geq 0$. (Of course $\hat{x}_{t|t} = E[x_t \mid \mathcal{Y}_t]$ with $Q > 0$).

Here we work with a risk-sensitive version of this estimation task, and define a *risk-sensitive estimate* with risk sensitive (scalar) parameter θ , suitably small to achieve existence of the expectation, as

$$\hat{x}_{t|t}^\theta \in \arg \min_{\xi \in \mathbb{R}^n} E\left[\exp\left\{\frac{\theta}{2}(x_t - \xi)'Q(x_t - \xi)\right\} \mid \mathcal{Y}_t\right] \quad (2)$$

The significance of θ is discussed after the optimal filter results are derived.

For simplicity of notation, and increased generality, let us work with the risk sensitive and risk neutral estimates, respectively,

$$\hat{x}_{t|t}^\theta \in \arg \min_{\xi \in \mathbb{R}^n} E\left[\exp\{\theta\Phi(x_t^\xi, \xi)\} \mid \mathcal{Y}_t\right] \quad (3)$$

(risk sensitive)

$$\hat{x}_{t|t}^\theta \in \arg \min_{\xi \in \mathbb{R}^n} E\left[\Phi(x_t^\xi, \xi) \mid \mathcal{Y}_t\right] \quad (4)$$

(risk neutral)

where $\Phi \in C(\mathbb{R}^{2n})$, the class of continuous functions, and $|\Phi(x, \xi)| \leq \kappa(1 + |x|^2 + |\xi|^2)$ for some $\kappa > 0$.

To proceed, let us work in the first instance with a continuous-time stochastic signal model, with $t \in [0, \infty)$

$$\begin{aligned} dx_t^\varepsilon &= a(x_t^\varepsilon)dt + dw_t^\varepsilon \\ dy_t^\varepsilon &= c(x_t^\varepsilon)dt + dv_t^\varepsilon \end{aligned} \quad (5)$$

where $w^\varepsilon, v^\varepsilon$ are standard independent Wiener processes scaled by $\sqrt{\varepsilon}$, and a, c are Lipschitz continuous with

$$|a(x)|, |c(x)| \leq \kappa(1 + |x|) \text{ for some } \kappa > 0 \quad (6)$$

Also, we require that the *a priori* density of x_0^ε , denoted $f(x)$, satisfy

$$\begin{aligned} \log f(x) \in \mathcal{D} &:= \{l(x) \in (\mathbb{R}^n) \\ &: l(x) \leq -\gamma_1 |x|^2 + \gamma_2, \\ &\gamma_1 > 0, \gamma_2 \geq 0\} \end{aligned} \quad (7)$$

Here the parameter ε is introduced so that subsequently we can conveniently consider our stochastic results taking a small noise limit with $\varepsilon \rightarrow 0$ to achieve deterministic worst case estimation.

We also generate results for a discrete-time model with $k \in \{0, 1, 2, \dots\}$.

$$\begin{aligned} x_{k+1}^\varepsilon &= a(x_k^\varepsilon) + w_{k+1}^\varepsilon \\ y_k^\varepsilon &= c(x_k^\varepsilon) + v_k^\varepsilon \end{aligned} \quad (8)$$

where w_k, v_k are white noise processes with densities $\psi(\cdot), \phi(\cdot)$, respectively. Here we assume Gaussian densities $N[0, \varepsilon I]$.

The special use of linear models are considered, namely

$$\begin{aligned} dx_t^\varepsilon &= Ax_t^\varepsilon dt + dw_t^\varepsilon \\ dy_t^\varepsilon &= Cx_t^\varepsilon dt + dv_t^\varepsilon \end{aligned} \quad (9)$$

for continuous time, and for the discrete-time case then we have

$$\begin{aligned} x_{k+1}^\varepsilon &= Ax_k^\varepsilon + w_k^\varepsilon \\ y_k^\varepsilon &= Cx_k^\varepsilon + v_k^\varepsilon \end{aligned} \quad (10)$$

The HMMs we consider are

$$\begin{aligned} X_{k+1} &= \mathcal{A}X_k + W_k \\ Y_k &= CX_k + v_k^\varepsilon \end{aligned} \quad (11)$$

where $X_k \in \{e_1, e_2, \dots, e_N\}$ with e_i being the unit vector with unity in the i th positions and zero elsewhere, and \mathcal{A} is the matrix of transitions probabilities such that $E[X_{k+1} | X_k] = \mathcal{A}X_k$. Again v_k is *i.i.d* with variance $\phi_\xi(\cdot) \sim N[0, \xi I]$, and estimates $\hat{x}_{k|k}^\theta$ are defined analogously. (Other HMM models with discrete range measurements and/or in continuous time as studied in (Elliott *et al.*, 1994) can be considered, but details are omitted here.)

Risk sensitive smoothing and prediction estimates are defined from a mild generalization of (3) as follows

$$\hat{x}_{t|\tau}^\theta \in \arg \min_{\xi \in \mathbb{R}^n} E[\exp(\theta\Phi(x_t, \xi)) \mid \mathcal{Y}_\tau]$$

For smoothing $\tau > t$, for prediction $\tau < t$, and for filtering $\tau = t$.

Measure Change: The continuous-time stochastic system (5) is defined on a probability space (Ω, \mathcal{F}, P) with $\mathcal{G}_t = \sigma(x_s^\varepsilon, y_s^\varepsilon; 0 \leq s \leq t)$ and $\mathcal{Y}_t = \sigma(y_s^\varepsilon; 0 \leq s \leq t)$.

Let \bar{P} denote the equivalent probability measure under which y_t^ε is a standard Wiener process independent of the state process (Elliott, 1982). Such a measure exists and (since v_t is Gaussian) is given from

$$\Lambda_{0,t}^\varepsilon := \frac{dP}{d\bar{P}} \Big|_{\mathcal{G}_t} = \exp \left(-\frac{1}{\varepsilon} \left[\frac{1}{2} \int_0^t |c(x_s^\varepsilon)|^2 ds - \int_0^t c(x_s^\varepsilon) dy_s^\varepsilon \right] \right) \quad (12)$$

For a proof of this Girsanov Theorem see (Elliott, 1982).

For the discrete-time models (8) and (11) defined on (Ω, \mathcal{F}, P) with $\mathcal{G}_k = \sigma(x_l^\varepsilon, y_l^\varepsilon; 0 \leq l \leq k)$ and $\mathcal{Y}_k = \sigma(y_l^\varepsilon; 0 \leq l \leq k)$, the corresponding measure change to yield y_k^ε *i.i.d.* is

$$\Lambda_{0,k}^\varepsilon := \frac{dP}{d\bar{P}} \Big|_{\mathcal{G}_k} = \prod_{i=0}^k \lambda_i(x_i) \quad (13)$$

$$\lambda_k(x_k) = \frac{\phi_\varepsilon(y_k - c(x_k))}{\phi_\varepsilon(y_k)}$$

For a proof see (Elliott *et al.*, 1994).

Information State: Let us denote the information state associated with the model (5) and measure change (12) as $q_{i|t}^\varepsilon(x)$. It satisfies the following defining equation for all $b: \mathbb{R}^n \rightarrow \mathbb{R}$ Borel test functions

$$\langle b, q_{i|t}^\varepsilon \rangle = \int_{\mathbb{R}^n} b(x) q_{i|t}^\varepsilon(x) dx = \bar{E} [\Lambda_{0,t} b(x_i^\varepsilon) | \mathcal{Y}_t] \quad (14)$$

where \bar{E} is the expectation under \bar{P} . Indeed, the information state satisfies the Zakai equation (Elliott, 1982).

$$dq_{i|t}^\varepsilon = \mathcal{B}^* q_{i|t}^\varepsilon dt + \frac{1}{\varepsilon} c'(x) q_{i|t}^\varepsilon dy_t^\varepsilon$$

$$q_{0|0}^\varepsilon(x) = \rho(x) \quad (15)$$

where $*$ denotes the adjoint and the operator \mathcal{B} is defined as

$$\mathcal{B}(b(x)) = \frac{\varepsilon}{2} \Delta b(x) + a(x) \frac{\partial b(x)}{\partial x}$$

Here Δ is the Laplace operator ($\Delta = \frac{\partial^2}{\partial x_1^2} + \dots + \frac{\partial^2}{\partial x_n^2}$). Also $\rho(x)$ is the *a priori* density of x . The Zakai equation is linear and infinite-dimensional in general.

For the discrete-time model (9), the information state defining equation is (14) with t replaced by k , and for the Zakai equation see (Elliott *et al.*, 1994). For the HMM (11), the information state is defined analogously but with $x \in \mathbb{R}^n$ replaced by $X \in \{e_1, \dots, e_N\}$. Also we denote $q_{k|k}^\varepsilon$ as a vector with i -th element $q_{k|k}^\varepsilon(e_i)$. Of course, now, $\langle b, q_{i|t}^\varepsilon \rangle = \sum_i b(e_i) q_{i|t}^\varepsilon(e_i)$ in (14).

The Zakai equation for the HMM model (11) is, (see (Elliott *et al.*, 1994))

$$q_{k+1|k+1}^\varepsilon = B_{k+1} A q_{k|k}^\varepsilon \quad (16)$$

where

$$B_{k+1} = \text{diag} \{ \phi_\varepsilon(y_{k+1} - Ce_1), \dots, \phi_\varepsilon(y_{k+1} - Ce_n) \}$$

With smoothing or prediction in mind, the information state is readily generalized to define $q_{i|\tau}$ from

$$\langle b, q_{i|\tau} \rangle = \bar{E} [\Lambda_{0,\tau} b(x_i^\varepsilon) | \mathcal{Y}_\tau] \quad (17)$$

Smoothed information state estimates with $\tau > t$ can be obtained from a forward filter for $q_{i|t}^\varepsilon$ and a backward filter from τ to t (see (Elliott *et al.*, 1994) for details).

Risk-sensitive Estimates in terms of Information State: The optimization task can be expressed in terms of the information state as follows. First apply a version of Bayes' Theorem (see (Elliott *et al.*, 1994)) as

$$E [\exp(\theta \Phi(x_i^\varepsilon, \xi)) | \mathcal{Y}_t] = \frac{\bar{E} [\Lambda_{0,t} \exp(\theta \Phi(x_i^\varepsilon, \xi)) | \mathcal{Y}_t]}{\bar{E} [\Lambda_{0,t} | \mathcal{Y}_t]} \quad (18)$$

Now application to (3) gives, for θ sufficiently small such that the expectation exists,

$$\hat{x}_{i|t}^\theta \in \arg \min_{\xi \in \mathbb{R}^n} \bar{E} [\Lambda_{0,t} \exp(\theta \Phi(x_i^\varepsilon, \xi)) | \mathcal{Y}_t] \in \arg \min_{\xi \in \mathbb{R}^n} \langle \exp(\theta \Phi(x, \xi)), q_{i|t}^\varepsilon(x) \rangle \quad (19)$$

Likewise, risk-neutral estimates are defined from

$$\hat{x}_{i|t} \in \arg \min_{\xi \in \mathbb{R}^n} \langle \Phi(x, \xi), q_{i|t}^\varepsilon(x) \rangle \quad (20)$$

Risk Neutral Filtering: Using a power series expansion for the exponential, it is straightforward to show that in the limit as $\theta \rightarrow 0$, then risk-neutral filter results are recovered from risk-sensitive filtering, that is, for the indices (19), (20)

$$\hat{x}_{i|t} = \lim_{\theta \rightarrow 0} \hat{x}_{i|t}^\theta \quad (21)$$

Linear Filters: Consider the special case of the

linear Gaussian signal models (9) and (10), and quadratic indices as in (1) and (2) with $Q > 0$. In this case, the information state is a scaled Gaussian with mean $\hat{x}_{t|t}$ and variance denoted $\Sigma_{t|t}$. Now the minimization (19) to achieve risk-sensitive filtering can be carried out analytically by completing-the-square arguments. Thus we have the key property that the risk-sensitive estimates $\hat{x}_{t|t}^\theta$ are identical to the risk-neutral (minimum variance) estimates $\hat{x}_{t|t}$

$$\hat{x}_{t|t}^\theta = \hat{x}_{t|t} \quad (22)$$

Trivially, also $\lim_{\theta \rightarrow 0} \hat{x}_{t|t}^\theta = \hat{x}_{t|t}$. This result applies in discrete-time also with t replaced by k .

Small Noise Limit: To consider small noise results as $\varepsilon \rightarrow 0$, express θ in terms of ε as

$$\theta = \frac{\mu}{\varepsilon} \quad (23)$$

Let us work first with the continuous time model (5) and define

$$S_t^\mu(q) \triangleq \inf_{\xi \in \mathbb{R}^n} \left\langle \exp\left(\frac{\mu}{\varepsilon} \Phi(x, \xi)\right), q_t^\varepsilon(x) \right\rangle \quad (24)$$

Also, as in (James *et al.*, 1993), recall that

$$\lim_{\varepsilon \rightarrow 0} \frac{\varepsilon}{\mu} \log q_t^\varepsilon(x) = p_t^\mu(x) \quad (25)$$

where, as long as $y_t^\varepsilon \rightarrow \int_0^t y_s ds$,

$$\frac{\partial p_t^\mu(x)}{\partial t} = \sup_{w \in \mathbb{R}^n} \left[-\frac{\partial p_t^\mu(x)}{\partial x} (a(x) + w) - \frac{1}{2\mu} |w|^2 \right] - \frac{1}{\mu} \left[\frac{1}{2} |c(x)|^2 - c'(x)y_t \right]$$

$$p_0^\mu(x) = \lim_{\varepsilon \rightarrow 0} \frac{\varepsilon}{\mu} \log \rho(x) \quad (26)$$

At this stage, recall a version of the Varadhan-Laplace limiting result from (James *et al.*, 1994)

$$\lim_{\varepsilon \rightarrow 0} \frac{\varepsilon}{\mu} \log \left\langle \exp\left(\frac{\mu}{\varepsilon} l(x)\right), \exp\left(\frac{\mu}{\varepsilon} m(x)\right) \right\rangle = \sup_x \{l(x) + m(x)\} \quad (27)$$

(27)

Mildly generalising the result of (James *et al.*, 1994), this result holds for $m(x) + l(x) \in \mathcal{D}$ and convergence is uniform on compact subsets of $\mathcal{D} \times \mathcal{D}$.

Now from (24)

$$S_t^\mu(\exp\left(\frac{\mu}{\varepsilon} p\right)) = \inf_{\xi} \left\langle \exp\left(\frac{\mu}{\varepsilon} \Phi(x, \xi)\right), \exp\left(\frac{\mu}{\varepsilon} p_t^\mu(x)\right) \right\rangle$$

So that, taking limits and applying the result (2), we have,

$$\lim_{\varepsilon \rightarrow 0} \frac{\varepsilon}{\mu} \log \{S_t^\mu(\exp\left(\frac{\mu}{\varepsilon} p\right))\} = \inf_{\xi} \sup_x \{p_t^\mu(x) + \Phi(x, \xi)\} \quad (28)$$

This holds if $p_t^\mu(x) + \Phi(x, \cdot) \in \mathcal{D}$, which in turn holds for $\log \rho(x) \in \mathcal{D}$ and $\mu > 0$ suitably small. At the limit of this condition holding as μ increases, there is a correspondence to H_∞ filtering, see (Glover, 1988).

Thus the small noise state estimate is, for $\mu > 0$ suitably small

$$\lim_{\varepsilon \rightarrow 0} \hat{x}_t^\mu \in \operatorname{argmin}_x \sup_{\xi} \{p_t^\mu(x) + \Phi(x, \xi)\} \quad (29)$$

which has the interpretation of a deterministic differential game in which state estimation is achieved in a worst case deterministic noise environment. An alternative nonrecursive interpretation of $p_t^\mu(x)$ is given from (see (James *et al.*, 1993))

$$p_t^\mu(x) = \sup_{\eta \in C(0,t)} \left\{ \alpha(\eta_0) - \frac{1}{\mu} \int_0^t \left(\frac{1}{2} |\dot{\eta}_s - a(\eta_s)|^2 + \frac{1}{2} |c(\eta_s)|^2 - c'(\eta_s)y_s \right) ds : \eta_t = x \right\} \quad (30)$$

We see that $p_t^\mu(x)$ serves the role of an information state in the deterministic setting, telling us as much as we can know about the states from the measurements in this setting. For the discrete-time model (8), the above results also hold with t replaced by k and (26) replaced by, (see (James *et al.*, 1994))

$$p_{k+1}^\mu(x) = \sup_{\nu \in \mathbb{R}^n} \left[-\frac{1}{2\mu} |x - a(\nu)|^2 - \frac{1}{\mu} \left[\frac{1}{2} |c(\nu)|^2 - c'(\nu)y_k \right] + p_k(\nu) \right] p_0^\mu(x) = \lim_{\varepsilon \rightarrow 0} \frac{\varepsilon}{\mu} \log \rho(x) \quad (31)$$

It is immediate from this equation that $p_0^\mu(x) \in \mathcal{D}$ implies $p_k^\mu(x) \in \mathcal{D}$ for all k , and $p_k^\mu(x) + \Phi(x, \xi) \in \mathcal{D}$ for μ suitably small. The nonrecursive version of (31) derived by successive application of (31) yields

$$p_k^\mu(z) = \sup_{x \in l_2((0,k), \mathbb{R}^n)} \left\{ p_0^\mu(z) - \frac{1}{\mu} \sum_{l=0}^{k-1} \frac{1}{2} |x_{l+1} - a(x_l)|^2 - \frac{1}{\mu} \sum_{l=0}^{k-1} \left[\frac{1}{2} |c(x_l)|^2 - c'(x_l)y_l \right] : x_k = z \right\} \quad (32)$$

We remark that in the linear Gaussian model case when the risk-sensitive estimate is identical to the minimum variance estimate, then of course this minimum variance estimate also has the interpretation of worst case estimate in the above sense. We believe this observation has not been made in earlier literature.

The results of this section can be summarised in the following theorem

Theorem 1 Continuous time: Suppose there is given the model (5) with the assumptions (6) and (7) holding and the risk-sensitive performance index given by (3). Then the optimum risk-sensitive estimate is given by (19). In the limit, as $\theta \rightarrow 0$, the risk-sensitive estimate approaches the risk-neutral estimate (20). Moreover, with $\theta = \frac{\mu}{\varepsilon}$, where ε is the noise variance, then in the limit as $\varepsilon \rightarrow 0$, the risk-sensitive estimate is given by (29) and (30).

Further, for linear signal model (9), (22) holds.

Discrete-time: For the discrete-time signal model (8), all the results for the continuous-time signal model hold with t replaced by k and (30) replaced by (32).

For the HMM case, only the measurement noise approaches zero as $\varepsilon \rightarrow 0$. In this case the differential game is not completely deterministic. Only the measurement noise is interpreted as deterministic, so that estimation of the stochastic discrete state X_k is achieved in the worst case deterministic measurement noise environment.

Of course, we would have studied the other signal models with only the process noise or measurement noise variance approaching zero; and in this case the limiting case is of a partially stochastic model in a deterministic noise environment.

Risk-sensitive Indices with Memory: In earlier works (Dey and Moore, 1994; Dey and Moore, 1995) the following risk-sensitive optimization is considered (in discrete-time)

$$\hat{x}_{k|k}^\theta \in \operatorname{argmin}_{\xi \in \mathbb{R}^n} E \left[\exp \left\{ \sum_{i=0}^{k-1} \theta \Phi(x_i^\varepsilon, \hat{x}_{i|i}^\theta) + \theta \Phi(x_k^\varepsilon, \xi) \right\} | \mathcal{Y}_k \right] \quad (33)$$

We observe here that this task can be tackled using the techniques of this paper by working with an augmented plant model and associated information state. Thus consider the augmented model (8)

$$\begin{aligned} x_{k+1}^\varepsilon &= a(x_k^\varepsilon) + w_{k+1}^\varepsilon \\ J_{k+1}^\varepsilon &= J_k^\varepsilon + \Phi(x_k^\varepsilon, \hat{x}_{k|k}^\theta) \end{aligned}$$

$$y_k^\varepsilon = c(x_k^\varepsilon) + v_k^\varepsilon \quad (34)$$

with states $(x_k^\varepsilon, J_k^\varepsilon)$. Now (33) can be rewritten as

$$\hat{x}_{k|k}^\theta \in \operatorname{argmin}_{\xi \in \mathbb{R}^n} \left\langle \exp \{ \theta J + \theta \Phi(x, \xi) \}, q_{k|k}^\varepsilon(x, J) \right\rangle$$

where $q_{k|k}^\varepsilon(x, J)$ is the information state associated with the augmented plant (34). Analysis can proceed using techniques of this section. The algorithms are formulated differently than in the earlier work. Here, the small noise limit results apply in a straightforward manner, giving new results.

We remark that in the linear Gaussian model case, the risk-sensitive estimate is no longer identical to the minimum variance estimate, and consequently this particular index is perhaps the more appealing one to work with. Both indices approach the minimum variance index in the risk-neutral case as $\theta \rightarrow 0$.

3. RISK-SENSITIVE GENERALIZATION OF MINIMUM VARIANCE CONTROL

Consider the discrete-time signal model

$$\begin{aligned} x_{k+1}^\varepsilon &= a(x_k^\varepsilon, u_k^\varepsilon) + w_{k+1}^\varepsilon \\ y_k^\varepsilon &= c(x_k^\varepsilon) + v_k^\varepsilon \end{aligned} \quad (35)$$

with control variable $u_k^\varepsilon \in \mathbb{R}^p$ is \mathcal{Y}_k measurable. Minimum variance control for such models is usually formulated as

$$u_{k|k} \in \operatorname{argmin}_{u \in \mathbb{R}^p} E \left[\frac{1}{2} y_{k+1}^\varepsilon{}' y_{k+1}^\varepsilon | \mathcal{Y}_k \right] \quad (36)$$

or more generally,

$$u_{k|k} \in \operatorname{argmin}_{u \in \mathbb{R}^p} E \left[\bar{\Psi}(y_{k+1}^\varepsilon, u) | \mathcal{Y}_k \right] \quad (37)$$

with $|\bar{\Psi}(y, u)| \leq \kappa(1 + |y|^2 + |u|^2)$ for some $\kappa > 0$.

The risk-sensitive generalization of minimum variance control is defined from

$$u_{k|k}^\theta \in \operatorname{argmin}_{u \in \mathbb{R}^p} E \left[\exp \{ \theta \bar{\Psi}(y_{k+1}^\varepsilon, u) \} | \mathcal{Y}_k \right] \quad (38)$$

Clearly, this problem is a special case of one-step-ahead prediction defined in its most general form in the previous section. Working under the measure \bar{P} ,

$$u_{k|k}^\theta \in \operatorname{argmin}_{u \in \mathbb{R}^p} \bar{E} \left[\Lambda_{0,k} \exp \{ \theta \bar{\Psi}(y_{k+1}^\varepsilon, u) \} | \mathcal{Y}_k \right]$$

Now substituting (35), taking expectation under \bar{E} and using inner product notation, we have

$$u_{k|k}^\theta \in \operatorname{argmin}_{u \in \mathbb{R}^p} \left\langle \exp \{ \theta \bar{\Psi}((a(x, u) + w) + v, u) \}, \phi_c(v) \psi_c(w) q_{k|k}^\varepsilon(x) \right\rangle \quad (39)$$

Linear Gaussian Model Case: The optimal control (39) can be solved analytically in the linear Gaussian model case since $v^\epsilon, w^\epsilon, q^\epsilon$ are normally distributed. Thus, applying completing-the-square arguments, as in derivation of (22), it can be shown that

$$u_{k|k}^\theta = \left[B'WB - B'WA(A'WA - \Sigma_{k|k}^{-1})^{-1}A'WB \right]^{-1} B'WA(A'WA - \Sigma_{k|k}^{-1})^{-1}\Sigma_{k|k}^{-1}\hat{x}_{k|k} \quad (40)$$

where

$$\begin{aligned} W &= V - V(V - \epsilon^{-1}I)V \\ V &= \theta C'C [1 - \theta(\theta - \epsilon^{-1})^{-1}] \end{aligned}$$

Indeed, observe that as $\theta \rightarrow 0$, $W \rightarrow \theta C'C$ and $u_{k|k}^\theta \rightarrow -[B'C'CB]^{-1}B'C'CA\hat{x}_{k|k}$ which is the risk-neutral minimum variance control

$$u_{k|k} = \operatorname{argmin}_{u \in \mathbf{R}^p} E\left[\frac{1}{2}y_{k+1}'y_{k+1} | \mathcal{Y}_k\right] \quad (41)$$

The results of this section can be summarised in the following theorem

Theorem 2 The risk-sensitive generalization of minimum variance control $u_{k|k}^\theta$ defined by (38) for the discrete-time model (35) is given in terms of the information state by (39).

For the linear Gaussian signal model, the risk-sensitive version of minimum variance control is given by (40). As $\theta \rightarrow 0$, the risk-sensitive version of minimum variance control approaches the risk-neutral minimum variance control (41).

Adaptive Risk-sensitive Version of Minimum Variance Controllers: An important application of the above risk-sensitive generalization of minimum variance control results is to indirect adaptive control, which is of course on-line. The control calculation is based on the most recent estimates of plant parameters, which in turn can be viewed as plant states and part of the state estimation process. As in minimum variance control of linear plants, for a closed-loop stability there is a severe restriction namely a minimum phase restriction on the plant. Equivalently, the inverse of the plant must be stable. Further details on this application, and extensions to more general situations will be omitted here. However, we stress that the motivation for using a risk-sensitive index is clear in the sense that in the small noise limit as $\epsilon \rightarrow 0$, given $\theta = \frac{\mu}{\epsilon}$ and μ suitably small, then the control is optimum for a worst case deterministic noise environment. Details follow closely the analysis of the previous section.

4. CONCLUSION

Risk-sensitive filtering, prediction and smoothing results have been developed for continuous-time and discrete-time stochastic models including hidden Markov models as augmentations to information state filtering, prediction and smoothing. In the small noise limit, for suitably small risk-sensitivity parameter, the risk-sensitive filtering and minimum variance filtering in the linear Gaussian case can be the interpretation of a deterministic estimation in a worst case deterministic noise environment and an information state has been derived for this case. Risk-sensitive estimation involving memory is achieved by application of the risk-sensitive estimation results to a signal model augmented by a state associated with the performance index.

Also, risk-sensitive versions of minimum variance controllers have been developed in terms of an optimization task involving the information state for discrete-time stochastic models. In the small noise limit the control is optimum for a solution with worst case deterministic noise. Application to adaptive risk-sensitive control is immediate.

5. REFERENCES

- [1] Dey, S. and J. B. Moore (1994). Risk-sensitive filtering and smoothing for Hidden Markov Models. *Systems and Control Letters*, accepted for publication.
- [2] Dey, S. and J. B. Moore (1995). Risk-sensitive filtering and smoothing via reference probability methods. *Proc. of American Control Conference*, Seattle, to appear.
- [3] Elliott, R.E. (1982) *Stochastic Calculus and Applications*. Springer-Verlag, Berlin-Heidelberg-New York.
- [4] Elliott, R.E., L. Aggoun and J. B. Moore (1994). *Hidden Markov Model Estimation and Control via Reference Probability Methods*. Springer-Verlag, New York.
- [5] James, M.R., J. S. Baras and R. J. Elliott (1993). Output feedback risk-sensitive control and differential games for continuous-time nonlinear systems. *Proc. 32nd IEEE CDC*.
- [6] James, M.R., J. S. Baras and R. J. Elliott (1994). Risk-sensitive control and dynamic games for partially observed discrete-time systems. *IEEE Transactions on Automatic Control*, **39**, 780-792.
- [7] Glover, K. and J. C. Doyle (1988). State-space formulae for all stabilizing controllers that satisfy an H_∞ norm bound and relations to risk-sensitivity. *Syst. Contr. Lett.*, **11**, 167-172.
- [8] Speyer, J.L., C. Fan and R. N. Banavar (1992). Optimal stochastic estimation with exponential cost criteria. *Proceedings of the 31st Conference on Decision and Control*, Vol. 2, pp.2293-2298.