

An Information-State Approach to Linear/Risk-Sensitive/Quadratic/Gaussian Control

Iain B. Collings, *student member, IEEE*, Matthew R. James, *member, IEEE*,
and John B. Moore, *Fellow, IEEE*

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

Tel: +61 6 249 2461, Fax: +61 6 279 8088, E-mail: iain@syseng.anu.edu.au

Abstract

In this paper we use an information-state approach to obtain the solution to the linear risk-sensitive quadratic Gaussian control problem. With these methods the solution is obtained without appealing to a certainty equivalence principle. Specifically we consider the case of tracking a desired trajectory. The result gives some insight to more general information-state methods for non-linear systems. Limit results are presented which demonstrate the link to standard linear quadratic Gaussian control. Also, a risk-sensitive filtering result is presented which shows the relationship between tracking and filtering problems. Finally, simulation studies are presented to indicate some advantages gained via a risk-sensitive control approach.

Keywords : Risk-Sensitive Control, LQG Control, Tracking

1 Introduction

Recently there has been much interest in risk-sensitive control techniques. Such control policies lead to an optimal solution for which the controller's sensitivity to risk can be varied. One application area for risk-sensitive control has been economics where risk-sensitivity is termed *hedging* or *risk-aversion*, for example Karp [10] and Caravani [6]. In these papers it is seen that advantages can be gained from the risk-sensitive approach, for problems such as dynamic trading and futures market prediction. In particular, the modified control policies are of a non-symmetric, exponential form.

The discrete-time risk-sensitive linear quadratic Gaussian (LQG) output feedback control problem was first solved by Whittle [13], where use was made of a risk-sensitive version of the certainty equivalence principle. This allowed the state estimation and control optimisation to be decoupled, solved separately and then re-coupled. The continuous-time case was solved by Bensoussan and van Schuppen [5] using a different technique, one which generalises to the non-linear case. Recent developments in risk-sensitive control have included a solution to the output feedback control problem for non-linear systems, using information-state techniques (James, Baras and Elliott [9]). While not requiring

the use of a certainty equivalence principle, the solution is of course infinite dimensional. In this paper we present the output feedback risk-sensitive LQG solution derived via the methods in [5, 9]. Specifically, we consider the case of tracking a desired trajectory. We show that, for regulation problems, the equations are consistent with those presented in [13] and that in the "risk-neutral" case, the standard LQG solution results.

The key to the technique used in this paper is that an information-state is chosen in such a way that it represents both a state estimate and the cost incurred to the time of the estimate. A change of reference probability measure is used to arrive at a linear recursive update equation for the information-state. Then dynamic programming methods are employed to obtain the solution to the control problem, having been re-formulated in terms of the information-state. This derivation is fundamentally different to Whittle's approach [13], being more closely linked to Bensoussan and van Schuppen [5].

An important feature of this paper is that it presents a finite dimensional solution to the risk-sensitive output feedback control problem. It therefore provides an example of the quite general methods derived in [9], and gives insight to the non-linear control solution.

This paper is, in part, an extension of the work presented in [1], where bi-linear systems were considered. In the work of this paper, however, the tracking solution is discussed in addition to regulation, and the control solution is solved explicitly. Simulation studies are also presented in an effort to demonstrate the effect of variations in the controller's sensitivity to risk. Various tracking problems are considered to show the advantages of the risk-sensitive approach. For a more thorough investigation into simulation results, the reader is directed to [7], where extensions to the system model are considered, such as the inclusion of integrators in the control design. Also in [7, 4] are results for risk-sensitive control of hidden Markov models, which have application to discrete-state systems.

Finally in this paper, a discussion on risk-sensitive filtering is presented, to demonstrate a link between filtering and tracking. The results presented are derived using the risk-sensitive control solution, with a modification to the cost criteria.

2 State Space Model

Consider the following discrete-time system on the probability space (Ω, \mathcal{F}, P) with complete filtration $\{\mathcal{F}_k\}$:

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k + v_k \\ y_{k+1} &= Cx_k + w_k \\ z_{k+1} &= Dx_k \end{aligned} \quad (1)$$

over the finite time interval $k = 0, 1, \dots, T$. The state of the system is represented by the process x . The observable part of the system is represented by the process y . In this paper we will consider the problem of output tracking, and denote the desired trajectory by \tilde{z} . The process which is to follow \tilde{z} is defined by z . The random variables v_k and w_k have normal densities $\psi \sim N(0, \Sigma)$ and $\phi \sim N(0, \Gamma)$ respectively, where Σ and Γ are $n \times n$ and $p \times p$ positive definite matrices. The control, u , takes values in \mathbb{R}^m . The complete filtration generated by (y_0, \dots, y_k) is denoted by \mathcal{Y}_k , and the admissible controls u are the set of \mathbb{R}^m -valued $\{\mathcal{Y}_k\}$ adapted processes. We write $U_{k,l}$ for the set of such control processes defined on the interval k, \dots, l .

In order to reformulate the system model (1), a new probability measure, \bar{P} , can be defined by setting

$$\Lambda_{0,k} = \frac{dP}{d\bar{P}} \Big|_{\mathcal{F}_k} = \prod_{\ell=1}^k \lambda_\ell, \quad (2)$$

where

$$\lambda_k = \frac{\psi(x_k - Ax_{k-1} - Bu_{k-1})\phi(y_k - Cx_{k-1})}{\psi(x_k)\phi(y_k)}. \quad (3)$$

Here, $\Lambda_{0,k}$ is an \mathcal{F}_k martingale, and $E[\Lambda_{0,k}] = 1$. Now, under \bar{P} , x_k and y_k are two sequences of independent, normally distributed random variables with densities ψ and ϕ respectively. This reformulated model results in a linear recursion for the unnormalised information-state, as in Section 4.

3 Cost

The cost function for the risk-sensitive control problem is given, for any admissible control $u \in U_{0,T-1}$, by

$$J(u) = E \left[\exp \theta \left\{ \Psi_{0,T-1} + \frac{1}{2} x_T' M_T x_T \right\} \right] \quad (4)$$

$$= \bar{E} \left[\Lambda_{0,T} \exp \theta \left\{ \Psi_{0,T-1} + \frac{1}{2} x_T' M_T x_T \right\} \right], \quad (5)$$

where

$$\begin{aligned} \Psi_{j,k} &\triangleq \sum_{\ell=j}^k \frac{1}{2} \left[x_\ell' M x_\ell + u_\ell' N u_\ell \right. \\ &\quad \left. + (\tilde{z}_{\ell+1} - Dx_\ell)' Q (\tilde{z}_{\ell+1} - Dx_\ell) \right]. \end{aligned} \quad (6)$$

Here, $\theta > 0$ is a real number and represents the amount of risk in the control policy. For small values of θ , approaching zero, the effect is to make control decisions assuming the stochastic disturbances are acting in an average manner. For larger values of θ , the control is effectively more conservative, or in other words, has a higher sensitivity to risk.

4 Information State

In this section we present finite dimensional recursions for the *information-state* which, as the name suggests, provides information about the state of the system [11](p. 81). In the case of risk-sensitive control, it is convenient to also include a component of the cost in the information-state. For the formulation presented here, the information-state is a probability distribution (it can be compared to the 'past stress' in [13]). For small values of θ , approaching zero, the mean and variance of the information-state become the state and covariance estimates for the linear Kalman filter.

For any admissible control u , consider the measure

$$\alpha_k(x) dx \triangleq \bar{E}[\Lambda_{0,k} \exp(\theta \Psi_{0,k-1}) I(x_k \in dx) | \mathcal{Y}_k] \quad (7)$$

where $I(\cdot)$ is the indicator function.

Lemma 4.1 *The information-state $\alpha_k(x)$, as defined in (7), obeys the following recursion:*

$$\begin{aligned} \alpha_{k+1}(x) &= \phi^{-1}(y_{k+1}) \int_{\mathbb{R}^n} \phi(y_{k+1} - C\xi) \exp(\theta \Psi_{k,k}) \\ &\quad \psi(x - A\xi - Bu) \alpha_k(\xi) d\xi \end{aligned} \quad (8)$$

Proof :

$$\begin{aligned} \alpha_{k+1}(x) dx &= \bar{E}[\Lambda_{0,k+1} \exp(\theta \Psi_{0,k}) I(x_{k+1} \in dx) | \mathcal{Y}_{k+1}] \\ &= \bar{E}[\lambda_{k+1} \Lambda_{0,k} \exp(\theta \Psi_{k,k}) \\ &\quad \exp(\theta \Psi_{0,k-1}) I(x_{k+1} \in dx) | \mathcal{Y}_{k+1}] \\ \alpha_{k+1}(x) &= \int_{\mathbb{R}^n} \dots \int_{\mathbb{R}^n} \frac{\phi(y_{k+1} - Cx_k)}{\phi(y_{k+1})} \\ &\quad \exp(\theta \Psi_{k,k}) \psi(x - Ax_k - Bu) \Lambda_{0,k} \\ &\quad \exp(\theta \Psi_{0,k-1}) d\bar{P}(x_0, \dots, x_k) \\ &= \phi^{-1}(y_{k+1}) \int_{\mathbb{R}^n} \phi(y_{k+1} - C\xi) \\ &\quad \exp(\theta \Psi_{k,k}) \psi(x - A\xi - Bu) \alpha_k(\xi) d\xi \end{aligned}$$

■

Theorem 4.1 *The information-state $\alpha_k(x)$ is an un-normalised Gaussian density given by*

$$\begin{aligned} \alpha_k(x) &= \alpha_k(x, \chi_k) \\ &= Z_k \exp(-1/2) [(x - \mu_k)' R_k^{-1} (x - \mu_k)] \end{aligned} \quad (9)$$

where $\chi_k = (\mu_k, R_k, Z_k)$, and μ_k , R_k^{-1} and Z_k are given by the following algebraic recursions:

$$\begin{aligned} \mu_{k+1} &= R_{k+1} \left[\Sigma^{-1} B u_k + \Sigma^{-1} A a_k^{-1} (R_k^{-1} \mu_k \right. \\ &\quad \left. - A' \Sigma^{-1} B u_k + C' \Gamma^{-1} y_{k+1} - \theta D' Q \tilde{z}_{k+1}) \right] \\ R_{k+1}^{-1} &= \Sigma^{-1} - \Sigma^{-1} A a_k^{-1} A' \Sigma^{-1} \\ Z_{k+1} &= Z_k |\Sigma|^{-\frac{1}{2}} |a_k|^{-\frac{1}{2}} \exp\left(-\frac{1}{2}\right) \left[\gamma_k - \mu_{k+1}' R_{k+1}^{-1} \mu_{k+1} \right] \end{aligned} \quad (10)$$

where

$$a_k = C' \Gamma^{-1} C - \theta (M + D' Q D) + A' \Sigma^{-1} A + R_k^{-1}, \quad (11)$$

$$\begin{aligned} \gamma_k &= u_k' (-\theta N + B' \Sigma^{-1} B) u_k + \mu_k' R_k^{-1} \mu_k - \theta \tilde{z}_{k+1}' Q \tilde{z}_{k+1} \\ &\quad - (\mu_k' R_k^{-1} - u_k' B' \Sigma^{-1} A + y_{k+1}' \Gamma^{-1} C - \theta \tilde{z}_{k+1}' Q D) a_k^{-1} \\ &\quad (R_k^{-1} \mu_k - A' \Sigma^{-1} B u_k + C' \Gamma^{-1} y_{k+1} - \theta D' Q \tilde{z}_{k+1}) \end{aligned} \quad (12)$$

under the condition that a_k and R_k be positive definite for all k .

Proof: Due to the linearity of the dynamics, and the fact that v_k and w_k are independent and normally distributed, we know that $\alpha_k(x)$ is an un-normalised Gaussian density. The recursions for μ_k , R_k^{-1} and Z_k are obtained by evaluating the integral in (8). The details are omitted. ■

Further matrix manipulations yield the following, more familiar, expressions:

$$\begin{aligned} \mu_{k+1} &= A\mu_k + Bu_k + A\tilde{K}_k [C'\Gamma^{-1}(y_{k+1} - C\mu_k \\ &\quad - \theta\Gamma Q\tilde{z}_{k+1}) + \theta(M + D'QD)\mu_k] \\ \tilde{K}_k &\triangleq (R_k^{-1} + C'\Gamma^{-1}C - \theta(M + D'QD))^{-1} \\ R_{k+1} &= \Sigma + A\tilde{K}_kA' \end{aligned} \quad (13)$$

which can be compared to the result presented in [13] for the case where $Q = 0$.

Limit Result :

Equations (13) can be re-expressed in the following form:

$$\begin{aligned} \mu_{k+1} &= A\mu_{k|k} + Bu_k \\ \mu_{k|k} &\triangleq \mu_k + K_k[y_{k+1} - C\mu_k \\ &\quad - \theta(\Gamma Q\tilde{z}_{k+1} - \Gamma C^{-1}(M + D'QD)\mu_k)] \\ K_k &\triangleq (R_k^{-1} - \theta(M + D'QD))^{-1}C' \\ &\quad [C(R_k^{-1} - \theta(M + D'QD))^{-1}C' + \Gamma]^{-1} \\ R_{k+1} &= \Sigma + AR_{k|k}A' \\ R_{k|k} &\triangleq R_k - K_kCR_k \end{aligned} \quad (14)$$

In the case when θ approaches zero, it can easily be seen that the equations in (14) reduce to the standard Kalman filter equations [2] (p. 40).

5 Alternate Cost Representation

In this section we show that the cost function can be expressed in terms of the information-state. This allows the optimisation problem to be solved by dynamic programming, without any appeal to a certainty equivalence principle.

Theorem 5.1 For any admissible control u , the risk sensitive cost can be expressed in the form

$$J(u) = \bar{E} [\langle \alpha_T(\cdot, \chi_T), \beta_T \rangle] \quad (15)$$

where $\langle f(\cdot), q(\cdot) \rangle = \int_{\mathbb{R}^n} f(z)q(z)dz$, and

$$\beta_T(x) \triangleq \exp(\frac{\theta}{2}x'M_Tx). \quad (16)$$

Proof: We have from (5) that

$$\begin{aligned} J(u) &= \bar{E} [\Lambda_{0,T} \exp(\theta\Psi_{0,T-1}) \exp(\frac{\theta}{2}x'_T M_T x_T)] \\ &= \bar{E} [\bar{E} [\Lambda_{0,T} \exp(\theta\Psi_{0,T-1}) \beta_T(x_T) | \mathcal{Y}_T]] \\ &= \bar{E} [\int_{\mathbb{R}^n} \beta_T(x) \alpha_T(x) dx] \\ &= \bar{E} [\langle \alpha_T(\cdot, \chi_T), \beta_T \rangle]. \end{aligned}$$

■

6 Dynamic Programming

Following [9] we know that the alternative control problem can be solved using dynamic programming. Suppose that at some time k , $0 < k < T$, the information-state χ_k is $\chi = (\mu, R, Z)$.

The value function for this control problem is [1, 9]:

$$V(\chi, k) = \inf_{u \in U_{k,T-1}} \bar{E}[\langle \alpha_k, \beta_k \rangle | \alpha_k = \alpha(\chi)], \quad (17)$$

where β_k is an adjoint process defined by

$$\beta_k(x) = \bar{E}[\Lambda_{k+1,T} \exp(\theta\Psi_{k,T-1}) \exp(\frac{\theta}{2}x'_T M_T x_T) | x_k = x, \mathcal{Y}_T]. \quad (18)$$

The adjoint process is different to the 'future stress' in [13], as it relates to output feedback, not state feedback.

Theorem 6.1 [1, 9] The value function satisfies the recursion

$$V(\chi, k) = \inf_{u \in U_{k,k}} \bar{E}[V(\chi_{k+1}(\chi_k, u, y_{k+1}), k+1) | \chi_k = \chi] \quad (19)$$

and $V(\chi, T) = \langle \alpha_T(\cdot, \chi), \beta_T \rangle$. ■

7 Dynamic Programming Solution

Theorem 7.1 The value function is the exponential of a quadratic in μ ,

$$V(\chi, k) = Z_k \exp(\theta/2)[\mu'_k S_k^a \mu_k + 2S_k^{b'} \mu_k + S_k^c], \quad (20)$$

and the optimal control is linear in μ .

$$u_k^{min} = -(N + B'\tilde{S}_{k+1}B)^{-1}B'[\tilde{S}_{k+1}\tilde{A}\mu_k + S_{k+1}^b + \theta K_k^b] \quad (21)$$

where

$$\begin{aligned} \tilde{S}_{k+1} &= ((S_{k+1}^a)^{-1} - \theta\tilde{\Gamma}_k\tilde{\Gamma}_k')^{-1}, \\ K_k^a &= (N + B'\tilde{S}_{k+1}B)^{-1}B'\tilde{S}_{k+1}\tilde{A}, \\ K_k^b &= S_{k+1}^a\tilde{\Gamma}_k\delta^{-1}(\tilde{\Gamma}_k'S_{k+1}^b - \theta D'Q\tilde{z}_{k+1}), \\ \tilde{\Gamma}_k &= A\tilde{K}_kC'\Gamma^{-1}C\Theta, \\ \Theta &= [(C'\Gamma C)^{-1} + \rho^{-1}R_k]^{1/2}, \\ \tilde{A} &= A\rho^{-1}, \\ \tilde{M} &= (M + D'QD)\rho^{-1}, \\ \sigma &= N + B'\tilde{S}_{k+1}B, \\ \delta &= I - \theta\tilde{\Gamma}_k'S_{k+1}^a\tilde{\Gamma}_k, \\ \rho &= I - \theta R_k(M + D'QD). \end{aligned}$$

Also, S_k^a and S_k^b are given by the following backwards recursions:

$$\begin{aligned} S_k^a &= \tilde{M} + \tilde{A}'S_{k+1}^a(I + BN^{-1}B'S_{k+1}^a - \theta\tilde{\Gamma}_k\tilde{\Gamma}_k'S_{k+1}^a)^{-1}\tilde{A} \\ S_k^b &= \tilde{A}'S_{k+1}^b - (I + \theta\tilde{M}R_k)D'Q\tilde{z}_{k+1} \\ &\quad + \tilde{A}'\tilde{S}_{k+1}B\sigma^{-1}B'(S_{k+1}^b + \theta K_k^b) + \theta\tilde{A}'K_k^b \end{aligned} \quad (22)$$

under the condition that $(I - \theta\tilde{\Gamma}_k'S_{k+1}^a\tilde{\Gamma}_k)$ is positive definite for all k , and C is positive definite except in the cases where $C = D$ or $D = 0$.

Proof : By evaluating the dynamic programming equation (19) for $V(\chi, T-1)$ it can be seen that the value function is the exponential of a quadratic in μ . The remainder of the proof is too long for presentation in this paper, but is essentially an evaluation of the dynamic programming equation (19), with appropriate variable transformations. ■

Remark : The condition that C be positive definite, is a manifestation of the variable transformation used in order to present the results in a form which more readily demonstrates the link to standard LQG results. As can be seen from the exclusion when $D = 0$, the condition only applies to the tracking part of the solution, (ie. S_k^b and K_k^b recursions). It is possible to solve the dynamic programming problem without such a variable transformation and thus remove the condition on C . □

In order to demonstrate consistency with the results presented in [13], where an appeal was made to a certainty equivalence principle, and $Q \equiv 0$, we now set

$$\Pi_k = S_k^a [I + \theta R_k S_k^a]^{-1} \quad (23)$$

and $Q = 0$, which results in the following recursion for Π_k :

$$\Pi_k = M + A' [\Pi_{k+1}^{-1} + B N^{-1} B' - \theta \Sigma]^{-1} A \quad (24)$$

under the condition that $(I - \theta R_k \Pi_k)$ is positive definite for all k .

Substitution of equation (24) into (21), yields

$$u_k^{min} = -N^{-1} B' (\Pi_{k+1}^{-1} + B N^{-1} B' - \theta \Sigma)^{-1} A [I - \theta R_k \Pi_k]^{-1} \mu_k \quad (25)$$

where the term $[I - \theta R_k \Pi_k]^{-1} \mu_k$ is sometimes referred to as the minimum stress estimate.

Limit Result :

In the case where θ approaches zero, it can easily be seen that \hat{S}_{k+1} , \hat{A} and \hat{M} approach S_{k+1}^a , A and $M + D' Q D$ respectively, and the following equations result from manipulations to (21) and (22):

$$\begin{aligned} u_k^{min} &= -(N + B' S_{k+1}^a B)^{-1} B' [S_{k+1}^a A \mu_k + S_{k+1}^b] \\ S_k^a &= M + D' Q D + A' [S_{k+1}^a \\ &\quad - S_{k+1}^a B (N + B' S_{k+1}^a B)^{-1} B' S_{k+1}^a] A \\ S_k^b &= (A - B (N + B' S_{k+1}^a B)^{-1} B' S_{k+1}^a A)' S_{k+1}^b \\ &\quad - D' Q \hat{z}_{k+1} \end{aligned} \quad (26)$$

These are the standard LQG equations, as presented for example in [3] (p. 32 and p. 81).

8 Risk-Sensitive Filtering Interpretations

In this section the risk-sensitive filtering problem is presented in order to demonstrate its connection with the control problem considered in the preceding sections. The risk-sensitive filtering problem is shown to be solved by the same equations derived previously for the tracking problem, but with a slight reinterpretation of the cost function. These results for the linear filtering

case have already been solved, without the control interpretation, in [12], and for the non-linear case in [8].

To see the connection to the control problem, the risk-sensitive filtering cost function is now presented.

$$J_k(\hat{x}_k) = \bar{E} [\Lambda_{0,k} \exp \theta \Psi_{0,k}(\hat{x}_k) | \mathcal{Y}_k, \hat{x}_{k-1}] , \quad (27)$$

where

$$\Psi_{j,k}(\hat{x}_k) \triangleq \sum_{\ell=j}^k \frac{1}{2} [(x_\ell - \hat{x}_\ell)' Q (x_\ell - \hat{x}_\ell)] , \quad (28)$$

and \hat{x}_ℓ is the risk-sensitive state estimate of x_ℓ . Note here that the cost is an expectation conditioned on the set of observations. This is due to the fact that the filtering problem is one of optimization in the forward direction (as opposed to the control problem which is an optimization in the backwards direction), and as such, the previous observations will be available when the optimization procedure is carried out at each iteration.

Comparing (27) and (28) to (5) and (6) it can easily be seen that (28) can be obtained from (6) by replacing $\hat{z}_{\ell+1}$ by \hat{x}_ℓ , and setting $D \equiv I$ and $B \equiv M \equiv N \equiv 0$. Therefore, with the same definition of information-state given in (7), the result presented in Theorem 4.1 holds here as well, with the appropriate replacement of symbols. Equation (10) can be compared to the equations in [12] which differ only slightly since the prediction problem was considered, rather than the filtering solution presented here. From the information-state, α_k , the risk-sensitive state estimate, \hat{x}_k , is obtained, as in [8], by solving the forward optimization task of the following theorem.

Theorem 8.1 The risk-sensitive state estimate, \hat{x}_k , defined by

$$\hat{x}_k = \underset{\zeta}{\operatorname{argmin}} J_k(\zeta) , \quad (29)$$

is given by the mean, μ_k , of the information-state, α_k , defined in (7), where $\Psi_{j,k}$ is given in (28), $\hat{z}_{\ell+1}$ is replaced by \hat{x}_ℓ , $D \equiv I$, and $B \equiv M \equiv N \equiv 0$.

Proof :

$$\begin{aligned} \hat{x}_k &= \underset{\zeta}{\operatorname{argmin}} \bar{E} [\Lambda_{0,k} \exp \theta \Psi_{0,k}(\zeta) | \mathcal{Y}_k, \hat{x}_{k-1}] \\ &= \underset{\zeta}{\operatorname{argmin}} \int_{\mathbf{R}^n} \dots \int_{\mathbf{R}^n} \exp \theta \Psi_{k,k}(\zeta) \Lambda_{0,k} \\ &\quad \exp \theta \Psi_{0,k-1} d\bar{P}(x_0, \dots, x_k) \\ &= \underset{\zeta}{\operatorname{argmin}} \int_{\mathbf{R}^n} \exp \theta \Psi_{k,k}(\zeta) \alpha_k(z) dz \\ &= \underset{\zeta}{\operatorname{argmin}} \int_{\mathbf{R}^n} \exp \frac{\theta}{2} (z - \zeta)' Q (z - \zeta) Z_k \\ &\quad \exp(-1/2) (z - \mu_k)' R_k^{-1} (z - \mu_k) dz \\ &= \underset{\zeta}{\operatorname{argmin}} C_1 \exp(1/2) \{ \theta \zeta' Q \zeta - \mu_k' R_k^{-1} \mu_k \\ &\quad - (\mu_k' R_k^{-1} - \theta \zeta' Q) a^{-1} (R_k^{-1} \mu_k - \theta Q \zeta) \} \\ &= -(\theta Q - \theta^2 Q a^{-1} Q)^{-1} \theta Q a^{-1} R_k^{-1} \mu_k \\ &= -(a - \theta Q)^{-1} R_k^{-1} \mu_k \\ &= \mu_k , \end{aligned}$$

where $a = (\theta Q - R_k^{-1})$, C_1 is constant with respect to ζ , and a is positive definite. It can now be seen that the risk-sensitive state estimate is given by the mean, μ_k , of the information-state, α_k . ■

The recursion for this estimate is easily obtained by rearranging the expression for μ_k from (13), with the appropriate substitutions for the filtering problem. The following equation results:

$$\mu_{k+1} = A\mu_k + A(R_k^{-1} + C'\Gamma^{-1}C - \theta Q)^{-1} [C'\Gamma^{-1}(y_{k+1} - C\mu_k) - \theta(\mu_k - \hat{x}_k)] \quad (30)$$

The final recursive equation is obtained by setting \hat{x}_k equal to μ_k , from Theorem 8.1, as follows:

$$\hat{x}_{k+1} = A\hat{x}_k + A(R_k^{-1} + C'\Gamma^{-1}C - \theta Q)^{-1} [C'\Gamma^{-1}(y_{k+1} - C\hat{x}_k)] \quad (31)$$

This filtering result corresponds, as noted before, to the prediction result presented in [12].

9 Simulation Studies

We now present simulation studies to demonstrate the effect of variations to the risk-sensitive parameter θ .

Example 1 : In this example we demonstrate a case where modelling errors are present. The true system is given by the following parameters:

$$\begin{aligned} A &= \begin{bmatrix} -0.2 & 1 \\ -0.2 & 0 \end{bmatrix}, & B &= \begin{bmatrix} 0.9 \\ -0.6 \end{bmatrix}, & C &= [1 \ 0], \\ \Sigma &= \begin{bmatrix} 0.01 & 0 \\ 0 & 0.01 \end{bmatrix}, & T &= 100, & D &= [1 \ 0], \\ M &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, & N &= 0.1, & \Gamma &= 0.01, \\ & & & & Q &= 100, \end{aligned}$$

and the trajectory to be followed, \hat{z}_k , is a unit step at $k = 20$. The modelling error is introduced by assuming in the design that A is given by

$$A = \begin{bmatrix} -0.8 & 1 \\ -0.8 & 0 \end{bmatrix}.$$

Table 1 gives values of the LQG, minimum variance, cost function (ie. $\Psi_{0,T-1} + 0.5x_T'M_Tx_T$) averaged over 100 simulation runs. It can be seen that in the case where no modelling error is present, of course $\theta = 0$ gives lowest cost. However, when the error is introduced, a higher value of θ gives a lower minimum variance cost. This example displays an advantage of the risk-sensitive approach in the presence of modelling errors.

Unfortunately, the sample path properties may not improve with a lower minimum variance cost, as one would wish, especially if θ is too large. Here, *too large* will depend on the type of modelling error, and will of course be unknown to the designer. Figure 1 shows a typical sample run for the case of no modelling errors. It shows that the cost function chosen for the tracking task considered, results in little difference in tracking errors between the LQG and risk-sensitive policies. Figure 2 shows a typical sample run for the case where modelling errors are present. Even though the minimum variance cost is lower for the risk-sensitive policy, the tracking performance might not be as desirable, having much greater oscillations in the transient response. Therefore the desirability of a risk-sensitive approach cannot be measured purely by the minimum variance cost.

10 Conclusion

In this paper we have presented the solution to the linear risk-sensitive quadratic Gaussian control problem. Results have been derived for the case of tracking a desired trajectory. The solution to the dynamic programming problem has been achieved without the need to appeal to a certainty equivalence principle, and hence gives insight to the solution for non-linear systems. Limit results have been presented which demonstrate the link to standard linear quadratic Gaussian control. Also, a risk-sensitive filtering solution is presented to demonstrate the link between filtering and tracking problems. Simulation studies were presented in order to show some advantages of the risk-sensitive approach.

Acknowledgement

The authors wish to acknowledge the funding of the Australian Telecommunications and Electronics Research Board (ATERB), and the activities of the Cooperative Research Centre for Robust and Adaptive Systems ((CR)²ASys) by the Australian Commonwealth Government under the Cooperative Research Centre Program.

$\times 10^2$	$\theta = 0$ (LQG)	$\theta = 0.1$	$\theta = 0.15$
No model error	4.714	4.715	4.716
With model error	9.363	6.076	6.593

Table 1: Error analysis for risk-sensitive control




Figure 1: No modelling errors

References

- [1] L. Aggoun, A. Bensoussan, R. J. Elliott, and J. B. Moore. Finite-dimensional quasi-linear risk-sensitive control. *Systems and Control Letters*, to appear, 1994.
- [2] B. D. O. Anderson and J. B. Moore. *Optimal filtering*. Prentice-Hall, New Jersey, 1979.
- [3] B. D. O. Anderson and J. B. Moore. *Optimal control : linear quadratic methods*. Prentice-Hall, New Jersey, 1989.
- [4] J. S. Baras and M. R. James. Robust and risk-sensitive output feedback control for finite-state machines and hidden Markov models. Technical report, Institute for Systems Research, Maryland, 1994.
- [5] A. Bensoussan and J. H. van Schuppen. Optimal control of partially observable stochastic systems with an exponential-of-integral performance index. *SIAM Jour. on Control and Optimization*, 23:599–613, 1985.
- [6] P. Caravani. On extending linear quadratic control theory to non-symmetric risky objectives. *Jour. Economic Dynamics & Control*, 10:83–88, 1986.
- [7] I. B. Collings, M. R. James, and J. B. Moore. An information state approach to risk-sensitive tracking problems. submitted to *Jour. Mathematical Systems, Estimation, and Control*, 1994.
- [8] S. Dey and J. B. Moore. Risk-sensitive filtering and smoothing via reference probability method. submitted, 1994.
- [9] M. R. James, J. S. Baras, and R. J. Elliott. Risk-sensitive control and dynamic games for partially observed discrete-time systems. *IEEE Trans. on Automatic Control*, AC-39(4):780–792, 1994.
- [10] L. S. Karp. Dynamic hedging with uncertain production. *Int. Economic Review*, 29:621–637, 1988.
- [11] P. R. Kumar and P. Varaiya. *Stochastic systems*. Prentice-Hall, New Jersey, 1986.
- [12] J. L. Speyer, C. Fan, and R. N. Banavar. Optimal stochastic estimation with exponential cost criteria. In *Proc. of the 31st Conf. on Decision and Control*, pages 2293–2298, December 1992.
- [13] P. Whittle. Risk-sensitive linear/quadratic/gaussian control. *Advances in Applied Probability*, 13:764–777, 1981.




Figure 2: With modelling errors