INTERNATIONAL JOURNAL OF ROBUST AND NONLINEAR CONTROL, VOL. 7, 1047–1055 (1997)

RISK-SENSITIVE DUAL CONTROL

SUBHRAKANTI DEY AND JOHN B. MOORE

Cooperative Research Centre for Robust and Adaptive Systems, Department of Systems Engineering, Research School of Information Sciences and Engineering, Australian National University, Canberra ACT 0200, Australia

SUMMARY

In this paper, we develop new results concerning the risk-sensitive dual control problem for output feedback nonlinear systems, with unknown time-varying parameters. These results are not merely immediate specializations of known risk-sensitive control theory for nonlinear systems, but rather, are new formulations which are of interest in their own right. A dynamic programming equation solution is given to an optimal risk-sensitive dual control problem penalizing outputs, rather than the states, for a reasonably general class of nonlinear signal models. This equation, in contrast to earlier formulations in the literature, clearly shows the dual aspects of the risk-sensitive controller regarding control and estimation. The computational task to solve this equation, as has been seen for the risk-neutral dual control problem, suffers from the so-called 'curse of dimensionality'. This motivates our study of the risk-sensitive version for a suboptimal risk-sensitive dual controller. Explicit controllers are derived for a minimum phase single-input, single-output auto-regressive model with exogenous input and unknown time-varying parameters. Also, simulation studies are carried out for an integrator with a time-varying gain. They show that the risk-sensitive suboptimal dual controller is more robust to uncertain noise environments compared with its risk-neutral counterpart. © 1997 by John Wiley & Sons, Ltd.

Int. J. Robust Nonlinear Control, **7**, 1047–1055 (1997)
No. of Figures: 2 No. of Tables: 0 No. of References: 21

Key words: dual control; dynamic programming; nonlinear; estimation

1. INTRODUCTION

The concept of dual control is generally attributed to Fel'dbaum.[1] In the case of a partially observable system, it has been shown[2,3] that the dynamic programming equation solution to the optimal control problem is computationally more difficult than for the complete information case. The additional computational effort is attributed to the dual aspects of the control; the controller must first obtain reasonable information about the states of the system before having a chance to achieve control objectives. In the case of a system with unknown (possibly time-varying) parameters, the task of the control actions is therefore twofold, probing for achieving information concerning the states, and feedback of this information to achieve control objectives. Probing for state estimation needs more aggressive control than for the case when the states are

This paper was recommended for publication by editor A. Sideris

* Correspondence to: S. Dey, Cooperative Research Centre for Robust and Adaptive Systems, Department of Systems Engineering, Research School of Information Sciences and Engineering, Australian National University, Canberra ACT 0200, Australia

CCC 1049-8923/97/121047–09$17.50 *Received 8 February 1995*
Revised 20 March 1996

known, and hence good control and good estimation are conflicting objectives. The optimal control in the case of partially observable systems achieves a trade-off between these two conflicting demands.

The computational effort in so-called dual control is quite formidable and it has been found[2] that the optimal solution to this control problem can be obtained only for a handful of very simple systems, and sometimes not even analytically but numerically.[4–7] To avoid the computational burden of the dynamic programming equations, researchers have considered a single-step horizon cost function instead of a multi-step cost function and have termed the optimizing control as cautious control[2] since it decreases the feedback gain when the parameter estimates are uncertain. Unfortunately, the solution to this one-step horizon control problem does not introduce any probing feature and thus, does not have the desired dual aspects. Various suboptimal strategies have been studied therefore to obtain an algorithm where the control would achieve a good balance between control and estimation.[8–14]

The control strategies so far studied in detail, aim at optimizing costs which are quadratic, involving the control and/or estimation energy. These problems have been termed *risk-neutral* control problems[15] as opposed to *risk-sensitive* control problems which optimize an exponential of a quadratic criteria weighted by a risk-sensitive parameter (usually > 0). The risk-sensitive control problem for discrete-time partially observed systems has been solved in Reference 15. A related control and tracking problem for linear discrete-time systems has been solved in Reference 16. Also, risk-sensitive filtering and smoothing problems have been solved for a class of general nonlinear systems in Reference 17 and for hidden Markov models with finite-discrete states.[18] It has been seen that risk-sensitive controllers and filters are more robust in the presence of plant and noise uncertainties than their risk-neutral counterparts. Also, they make connection to worst case control and estimation problems in a deterministic noise scenario (H_∞ control and filtering problems for linear systems).[15,19] Risk-sensitive problems also specialize to risk-neutral problems as the risk-sensitive parameter tends to zero. These facts establish the general nature of the risk-sensitive problems.

In this paper, we study the risk-sensitive version of the dual control problem. Although Reference 15 actually addresses the risk-sensitive optimal control problem for partially observable systems and achieves a dynamic programming equation by applying change of probability measure technique, it is difficult to interpret the dual aspects of risk-sensitive control from these results. Therefore, by considering a cost function which penalizes the system output, we achieve a dynamic programming equation which achieves the same objectives, without resorting to the measure change technique of Reference 15. This result is of interest in its own right (see Remark 2.1). In addition, we present a suboptimal risk-sensitive dual controller, the risk-neutral version of which has been considered in Reference 14. We also present some simulation studies illustrating the robustness of the risk-sensitive suboptimal dual controller to uncertain noise environments. The risk-sensitive version of the cautious control problem with a single-step cost criterion has been addressed in Reference 20.

We present the optimal risk-sensitive dual control problem and the dynamic programming equation solution to it for a certain class of nonlinear systems in Section 2. General nonlinear systems can be addressed without much difficulty using the same techniques, but are not discussed in this paper. In Section 3, we consider a particular extension of the one-step horizon control cost to obtain a suboptimal risk-sensitive dual controller. This controller is obtained for a single input single output (SISO) auto-regressive with exogenous input (ARX) model with time-varying unknown parameters and simulation studies are carried out to show the superiority of this controller to its risk-neutral counterpart. Section 4 presents some concluding remarks.

2. RISK-SENSITIVE DUAL CONTROL

In this section we introduce the risk-sensitive dual control problem for a certain class of nonlinear systems. We describe the signal model, introduce the cost criterion and give a dynamic programming equation solution to the optimal control problem assuming separability between estimation and control.

2.1. Signal model

We consider the following discrete-time stochastic nonlinear state space model defined on a probability space $(\Omega, \mathcal{F}, P)$:

$$x_{k+1} = A_k(x_k) + B_k(u_k) + w_{k+1}$$
$$y_k = C_k(x_k) + v_k \tag{1}$$

where $x_k, w_k \in \mathbb{R}^n$, $y_k, v_k \in \mathbb{R}^p$, $u_k \in \mathbb{R}^m$. Here, x_k denotes the augmented state of the system including the unknown system parameters, u_k denotes the control input, y_k denotes the measurement, w_k and v_k are the process noise and the measurement noise respectively. The vectors A_k, B_k and C_k are nonlinear functions in general. We assume that $w_k, k \in \mathbb{N}$ has a density function ψ_k and v_k, $k \in \mathbb{N}$ has a strictly positive density function ϕ_k. The initial state x_0 or its density is assumed to be known and w_k is independent of v_k.

2.2. Cost criterion

Define $Y_k \triangleq (y_0, y_1, \ldots, y_k)$, the σ-field generated by Y_k as $\mathcal{Y}_k^0$ and the corresponding complete filtration by $\mathcal{Y}_k$. Also define $U_{m,n}$ to be the set of the admissible controls u_k in the interval $m \leqslant k \leqslant n$, where u_k is $\mathcal{Y}_k$ measurable. The risk-sensitive cost criterion for the dual control problem is given as, for $u \in U_{k-1,T-1}$,

$$J(u) = E\left[\exp\left\{\theta\left(\sum_{i=k}^{T} L(y_i, u_{i-1}, r_i)\right)\right\}\right] \tag{2}$$

The problem objective is to find $u^* \in U_{k-1,T-1}$ such that

$$u^* = \underset{u \in U_{k-1,T-1}}{\operatorname{argmin}} E\left[\exp\left\{\theta\left(\sum_{i=k}^{T} L(y_i, u_{i-1}, r_i)\right)\right\}\right] \tag{3}$$

Here, $r_i \in \mathbb{R}^p$, $i \in \mathbb{N}$ is the reference output that is supposed to be tracked by y_i. We also assume that $L \in C(\mathbb{R}^p \times \mathbb{R}^m \times \mathbb{R}^p)$ is non-negative, bounded and uniformly continuous. $\theta(>0)$ is the risk-sensitive parameter.

Using a fundamental result of stochastic control,[2] the problem objective is to find u^* such that

$$u^* = \underset{u \in U_{k-1,T-1}}{\operatorname{argmin}} E\left[\exp\left\{\theta\left(\sum_{i=k}^{T} L(y_i, u_{i-1}, r_i)\right)\right\} \middle| \mathcal{Y}_{k-1}\right] \tag{4}$$

Remark 2.1

The cost criterion could have been expressed in terms of the state x_i, rather than the output y_i, as

$$J(u) = E\left[\exp\left\{\theta\left(\sum_{i=k}^{T-1} L(x_i, u_{i-1}) + \Phi(x_T)\right)\right\}\right] \tag{5}$$

where $L \in C(\mathbb{R}^n \times \mathbb{R}^m)$ is non-negative, bounded and uniformly continuous in (x, u) and $\Phi \in C(\mathbb{R}^n)$ is non-negative, bounded, and uniformly continuous. This risk-sensitive control problem has been solved in References 15 and 16 using change of probability measure techniques. But the dual aspects of the control are not so evident from the dynamic programming equation obtained in References 15 and 16 and so this case is not studied further here.

2.3. Dynamic programming

We have separability between estimation and control as in Reference 15. The estimation problem is solved by evaluating the information state, which in this case is a conditional probability density function of the state given the observations.

Definition 2.1

Define the information state $\alpha_{k|k-1}(x)$ such that

$$\alpha_{k|k-1}(x)\mathrm{d}x = E[I(x_k \in \mathrm{d}x)|\mathcal{Y}_{k-1}] \tag{6}$$

Definition 2.2

Let us define the value function $V(\alpha_{k|k-1}, k)$ such that

$$V(\alpha_{k|k-1}, k) = \inf_{u \in U_{k-1,T-1}} E\left[\exp\left\{\theta\left(\sum_{i=k}^{T} L(y_i, u_{i-1}, r_i)\right)\right\}\middle|\mathcal{Y}_{k-1}\right] \tag{7}$$

Remark 2.2

We assume here that $\exp\{\theta(\sum_{i=k}^{T} L(y_i, u_{i-1}, r_i))\}$ is integrable.

We state the following theorem without proof.

Theorem 2.1

The value function $V(\alpha_{k|k-1,} k)$ satisfies the following recursive dynamic programming equation

$$V(\alpha_{k|k-1}, k) = \inf_{u_{k-1}} E[\exp\{\theta(L(y_k, u_{k-1}, r_k))\} V(\alpha_{k+1|k}, K+1)|\mathcal{Y}_{k-1}] \tag{8}$$

$$V(\alpha_{T|T-1}, T) = \inf_{u_{T-1}} \int_{\mathbb{R}^p} \int_{\mathbb{R}^n} \exp\{\theta(L(c_T(x) + v, u_{T-1}, r_T))\} \alpha_{T|T-1}(x)\phi_T(v)\mathrm{d}x\mathrm{d}v \tag{9}$$

Remark 2.3

Note that considering the cost criterion (2) instead of (5) results in the dynamic programming equation (8) (without applying change of probability measure techniques) which involves computing the expectation of the product of two terms. The first term denotes the immediate risk-sensitive control cost. The second term is a function of $\alpha_{k+1|k}(x)$ which itself is a function of Y_k and $u_{k-1}, \ldots, u_0$. This implies therefore, that u_{k-1} not only affects the immediate risk-sensitive control cost but also influences the future information state. This clearly shows the dual nature of the risk-sensitive control. Unfortunately, just like the optimal solution to the risk-neutral dual control problem,[2] the optimal risk-sensitive dual control cannot be computed analytically.

Numerical solutions are probable in a few cases, but are computationally expensive because the computational complexity increases exponentially with the dimension of the information state.

3. ROBUST (RISK-SENSITIVE) SUBOPTIMAL DUAL CONTROLLER

In Section 2, we found that the optimal risk-sensitive dual control cannot be achieved analytically. Owing to similar difficulties encountered in the risk-neutral optimal dual control problem, researchers have considered other suboptimal strategies which could substantially simplify the computational procedure. Since the cautious controller (which optimizes a single-step cost criterion), is not a dual controller, adding perturbation signals to the cautious controller has been considered in References 8 and 9. In References 10 and 11, constrained one step minimization techniques have been considered, the constraint being on the minimum value of the control signal or on the variance of the parameter estimates. Several works[12,13] have considered different extensions of the single-step cost criterion (i.e., the cost criterion for the cautious control problem) in the risk-neutral case.

In this section, we consider a similar extension of the single-step risk-sensitive cost criterion.[20] The corresponding extension in the risk-neutral case has been studied in Reference 14. We first present a generalized extended cost-criterion for a risk-sensitive suboptimal dual controller, followed by a specific cost-criterion for a SISO minimum phase ARX model. We then present an analytical solution for the control that optimizes this specific cost-criterion. This controller is suboptimal in the sense that it does not achieve the optimal risk-sensitive dual control but, by optimizing a cost that includes an extra term penalizing the estimation error, it tries to achieve a reasonable balance between control and estimation. We also present some simulation studies which illustrate that in the presence of uncertainties in the model dynamics, the risk-sensitive suboptimal dual controller incurs less cost than its corresponding risk-neutral counterpart and is thus robust to uncertainties in the noise dynamics.

3.1. Cost criterion

Consider a generalized cost-criterion for a risk-sensitive suboptimal dual controller for the system (1) given by

$$J_{\text{sub}}(u_{k-1}) = E[\exp\{\theta_1(L(y_k, u_{k-1}, r_k) + \theta_2 f(x_k, \hat{x}_{k|k-1}))\} | \mathscr{Y}_{k-1}] \tag{10}$$

where θ_1, θ_2 are risk-sensitive parameters and $f: \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ is a convex function reflecting a measure of the estimation error energy ($\hat{x}_{k|k-1}$ being an estimate of x_k given $\mathscr{Y}_{k-1}$), so that both the control and estimation costs are penalized.

3.2. Risk-sensitive suboptimal dual controller for a SISO ARX model

Consider the discrete-time minimum phase SISO ARX model

$$y_k + a_k^1 y_{k-1} + \cdots + a_k^n y_{k-n} = b_k^1 u_{k-1} + \cdots + b_k^n u_{k-n} + v_k \tag{11}$$

where y_k, u_k, v_k are output, input and measurement noise respectively at the kth time instant. The noise sequence $\{v_k\}, k \in \mathbb{N}$ is assumed to be Gaussian distributed with a density $\phi_v \sim N(0, \sigma_v^2)$. v_k is also assumed to be independent of y_i, $i \in \{1, 2, \ldots, k-1\}$ and a_i^j, b_i^j, $i \in \{1, 2, \ldots, k\}, j \in \{1, 2, \ldots, n\}$. It is further assumed that $b_k^1 \neq 0\ \forall k$.

The state of the system is denoted by $x_k = [b_k^1 b_k^2 \cdots b_k^n a_k^1 \cdots a_k^n]'$ and the state dynamics is given by

$$x_{k+1} = A_k x_k + w_k \tag{12}$$

where A_k is a known matrix and $\{w_k\}$ is a sequence of i.i.d random vectors distributed with a density function $\phi_w \sim N(0, \Sigma_w)$, $\forall k \in \mathbb{N}$.

With this state description, the output dynamics are given by

$$y_k = \psi'_{k-1} x_k + v_k \tag{13}$$

where

$$\psi'_{k-1} = [u_{k-1} \cdots u_{k-n} - y_{k-1} \cdots y_{k-n}]$$

The initial state x_0 or its distribution is assumed to be known.

Cost criterion. Let us consider the following cost criterion for the SISO ARX model described above, given by

$$J_{\text{sub}}^{\text{SISO}}(u_{k-1}) = E[\exp\{\theta_1((y_k - r_k)^2 + \lambda e_k^2)\} | \mathscr{Y}_{k-1}] \tag{14}$$

where $e_k = y_k - \psi'_{k-1}\hat{x}_{k|k-1}$ and $\lambda = \theta_2/\theta_1$.

Therefore, the problem objective is to find u_{k-1}^* such that

$$u_{k-1}^* = \underset{u_{k-1} \in U_{k-1,k-1}}{\operatorname{argmin}} \; J_{\text{sub}}^{\text{SISO}}(u_{k-1}) \tag{15}$$

Remark 3.1

It should be noted that the minimum phase assumption on (11) is not restrictive. Non-minimum phase systems can be treated by including a term penalizing the control cost in the cost index described above. This, of course, would result in a more complicated stability criterion.

Separability of estimation and control applies as before and the estimation is carried out by a Kalman filter. Details can be found in Reference [20]. The following theorem gives the result for the risk-sensitive suboptimal dual controller for the SISO ARX model (11).

Theorem 3.1

The risk-sensitive suboptimal dual control that optimizes the cost criterion (14) is given by

$$u_{k-1}^* = \underset{u_{k-1} \in U_{k-1,k-1}}{\operatorname{argmin}} \; \beta_k \exp\left[\frac{\theta_1(1 - \lambda\theta_1\sigma_{yk}^2)}{2(1 - \theta_1(1+\lambda)\sigma_{yk}^2)}(\psi'_{k-1}\hat{x}_{k|k-1} - r_k)^2\right] \tag{16}$$

where $\beta_k = 1/\sqrt{(1 - \theta_1(1+\lambda)\sigma_{y_k}^2)}$. Also, σ_{yk}^2 is the variance of the process y_k given $\mathscr{Y}_{k-1}$ and $\hat{x}_{k|k-1}$ is the mean of the conditional Gaussian density $\alpha_{k|k-1}(x)$.

Remark 3.2

(16) has to be solved numerically.

Remark 3.3

We assume $\theta_1 < 1/(1+\lambda)\sigma_{y_k}^2$, $\forall k \in \mathbb{N}$. The choice of θ_1 and λ is dependent on the trade-off between good control and good estimation.

3.3. *Simulation studies*

Here, we present a brief simulation study to show how risk-sensitive suboptimal dual control can perform better than the risk-neutral suboptimal dual control in uncertain noise environments. Consider an integrator in discrete-time with a time-varying gain given by

$$\begin{aligned} b_{k+1} &= A_k b_k + w_k \\ y_k &= y_{k-1} + b_k u_{k-1} + v_k \end{aligned} \tag{17}$$

We assume $w_k \sim N(0, \sigma_w^2)$, $v_k \sim N(0, \sigma_v^2)$. Choose $A_k = A = 0{\cdot}95$, $\forall k \in \mathbb{N}$ $\sigma_w^2 = 1{\cdot}0$, $\sigma_v^2 = 0{\cdot}49$. Also, let $r_k = r = 1{\cdot}0$, $\forall k \in \mathbb{N}$. We implement the risk-neutral suboptimal dual controller given in Reference 14 and our risk-sensitive suboptimal dual controller given by Theorem 3.1. The performance measure used to compare the two schemes is $\frac{1}{N}\sum_{k=1}^{N} y_k^2$.

We consider two types of uncertainties.

Coloured process noise. In realistic environments, the process noise is often coloured. Let us take a particular case where the dynamics of the gain parameter b_k is given by

$$b_{k+1} = 0{\cdot}95 b_k + w_k - 0{\cdot}4 w_{k-1} + 0{\cdot}7 w_{k-2}$$

The risk-neutral suboptimal dual controller studied in Reference 14 is given by

$$u_{r.n} = \frac{\hat{b}_k (r_k - y_{k-1})}{\hat{b}_k^2 + \sigma_{b_k}^2 (1 + \lambda)} \tag{18}$$

We run the risk-neutral controller with $\lambda = -0{\cdot}5$ and the risk-sensitive controller (16) with $\theta_1 = 4 \times 10^{-4}$ and $\lambda = -0{\cdot}5$ assuming the process noise is white. Figure 1 shows the cost

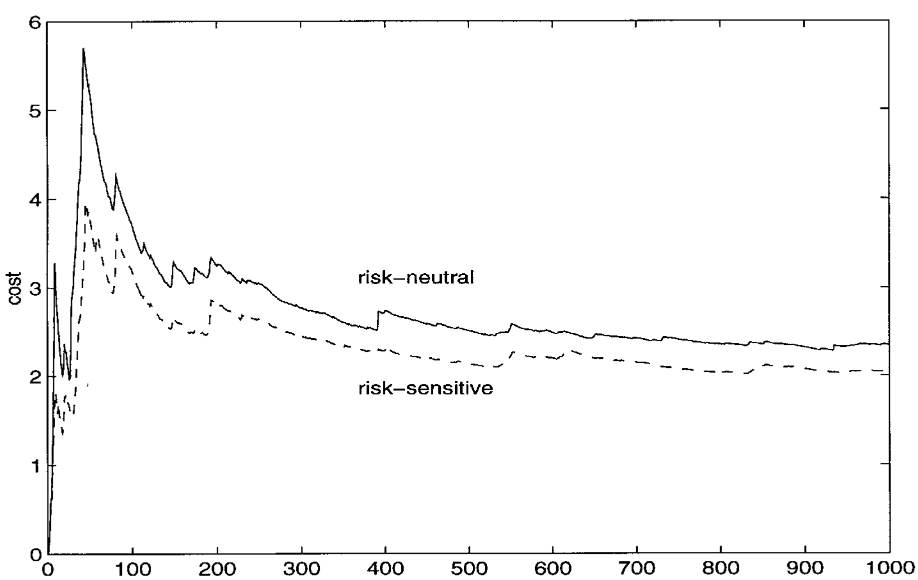

$\theta_1 = 4 \times 10^{-4}$, $\lambda = -0.5$

Figure 1. Robustness of risk-sensitive dual controller against coloured noise

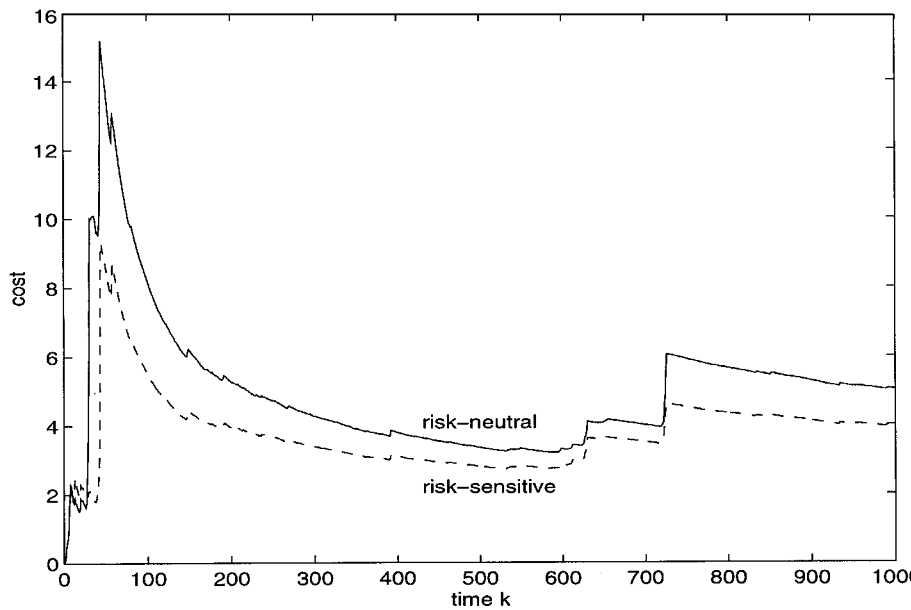

$\theta_1 = 4 \times 10^{-3}$, $\lambda = -0.5$

Figure 2. Robustness of risk-sensitive dual controller against high noise

accumulated over 1 000 time points using simulated data. It is clear that the risk-sensitive controller yields a lower cost.

Unexpectedly high process noise. In this case, we take the actual $\sigma_w^2 = 4$ whereas both the controllers run assuming $\sigma_w^2 = 1$. For this example, we take $\theta_1 = 4 \times 10^{-3}$ and $\lambda = -0{\cdot}5$. Figure 2 shows the cost incurred by the risk-neutral and the risk-sensitive dual controllers. It is seen that even in such hostile noise environments, the risk-sensitive controller performs better.

To conclude, it would be fair to say that the risk-sensitive suboptimal dual controller is expected to perform better than its risk-neutral counterpart in uncertain noise situations. But there is no general rule so far as to how to choose a suitable value or a suitable range of values of θ_1, for which the risk-sensitive controller will perform better.

4. CONCLUSIONS

Dual aspects of the risk-sensitive control have been studied in this paper. A dynamic programming equation solution to the optimal risk-sensitive dual control problem has been given. For the case of cost indices in terms of outputs rather than states, this dynamic programming equation shows the control and probing aspects of the risk-sensitive controller in a conveniently separated form. The difficulty involved in solving this equation even numerically calls for suboptimal risk-sensitive dual control strategies. One such strategy has been considered by extending the single-step risk-sensitive dual control cost criterion. Also, risk-sensitive cautious control has been studied for a SISO, minimum phase, ARX model. The suboptimal dual controller has been derived for the same model. Simulation studies carried out for the special case of an integrator with a time-varying gain show that the suboptimal risk-sensitive dual controller is more robust to uncertain noise environments than its risk-neutral counterpart.

ACKNOWLEDGEMENTS

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.

REFERENCES

1. Fel'dbaum, A. A., *Optimal Control Systems*, Academic Press, New York, 1965.
2. Åström, K. J. and B. Wittenmark, *Adaptive Control*, Addison-Wesley, New York, 1989.
3. Kumar, P. R. and P. Varaiya, *Stochastic Systems: Estimation, Identification and Adaptive Control*, Prentice Hall, New Jersey, 1986.
4. Åström, K. J., 'Optimal control of Markov processes with incomplete state information I', *J. Math. Anal. Appl.*, **10**, 174–205 (1965).
5. Åström, K. J., 'Optimal control of Markov processes with incomplete state information II', *J. Math. Anal. Appl.*, **26**, 403–406 (1969).
6. Sternby, J., 'A simple dual control problem with an analytical solution', *IEEE Trans. Automat. Contr.*, **21**, 840–844 (1976).
7. Åström, K. J. and B. Wittenmark, 'Problems of identification and control', *J. Math. Anal. Appl.*, **34**, 90–113 (1971).
8. Wieslander, J. and B. Wittenmark, 'An approach to adaptive control using real-time identification', *Automatica*, **7**, 211–217 (1971).
9. Jacobs, O. L. R. and J. W. Patchell, 'Caution and probing in stochastic control', *Int. J. Control*, **16**, 189–199 (1972).
10. Alster, J. and P. R. Bélanger, 'A technique for dual adaptive control', *Automatica*, **10**, 627–634 (1974).
11. Mosca, E., S. Rocchi and G. Zappa, 'A new dual active control algorithm', *Preprints 17th IEEE Conf. on Decision and Control*, San Diego, CA, IEEE Piscataway, NJ, 1978, pp. 509–512.
12. Wittenmark, B. and C. Elevitch, 'An adaptive control algorithm with dual features', *Preprints 7th IFAC Symposium on Identification and System parameter Estimation*, York, U.K., 1985, Pergamon Press Ltd, Oxford, U.K., 1985, pp. 587–592.
13. Sternby, J., '*Topics in dual control*', PhD Thesis TFRT-1012, Department of Automatic Control, Lund Institute of Technology, Lund, Sweden, 1977.
14. Milito, R., C. S. Padilla, R. A. Padilla and D. Cadroin, 'An innovations approach to dual control', *IEEE Trans. Automat. Contr.*, **27**, 132–137 (1982).
15. James, M. R., J. S. Baras and R. J. Elliott, 'Risk-sensitive control and dynamic games for partially observed discrete-time systems". *IEEE Trans. on Automatic Control*, **39**, 780–792 (1994).
16. Collings, I. B., M. R. James and J. B. Moore, 'An information-state approach to risk-sensitive tracking problems' (summary), *Journal of Mathematical Systems, Estimation, and Control*, **6**, 343–346 (1996).
17. Dey, S. and J. B. Moore, 'Risk-sensitive filtering and smoothing via reference probability methods', *IEEE Trans. on Automatic Control*, (1997), in press.
18. Dey, S. and J. B. Moore, 'Risk-sensitive filtering and smoothing for hidden Markov models', *Systems and Control Letters*, **25**, 361–366 (1995).
19. Moore, J. B., R. J. Elliott and S. Dey, 'Risk-sensitive generalizations of minimum variance estimation and control', *Proc. of the IFAC Symposium on Nonlinear Control System Design (NOLCOS)*, Tahoe City, CA, June 1995, Elsevier Science Ltd, Oxford, U.K., 1995, pp. 465–470.
20. Dey, S. and J. B. Moore, 'Risk-sensitive dual control', *Proc. of 34th IEEE Conf. on Decision and Control*, New Orleans, Louisiana, December 1995, IEEE, Piscataway, New Jersey, pp. 1042–1047.
21. Anderson, B. D. O. and J. B. Moore, *Optimal Filtering*, Prentice Hall, New Jersey, 1979.

