[Math. Control, Systems and Signals, 13, 318-332, 2000]

Robust Properties of Risk–Sensitive Control*

Paul Dupuis† Matthew R. James‡ Ian Petersen§

Abstract

The purpose of the present paper is to precisely characterize and prove robustness properties of risk sensitive controllers. In particular, we establish a stochastic version of the small gain theorem. This theorem is expressed in terms of an inequality which bounds the average output power in terms of the input power. Since this inequality is closely related to the risk-sensitive criterion, our stochastic small gain theorem can be expressed in terms of the risk-sensitive criterion. This provides a concrete motivation for the use of the risk-sensitive criterion stochastic robustness.

Keywords. Risk sensitive stochastic control, robustness, small gain theorem.

1 Introduction

The development of risk–sensitive control has been intensive in recent years [BvS85, DM97, FHH97, FM95, HHM96, Jac73, JBE94, PMR96, Whi81], in part because of its connection with nonlinear robust (or H^∞) control [GD88, FM95, Jam92]. The latter topic has also been extensively developed, and was itself motivated by the successes of the linear H^∞ control formulation introduced by Zames [Zam81]. The original motivations for Zames work

*This research was supported in part by the National Science Foundation (NSF-DMS-9403820), the Army Research Office (DAAH4-96-I-0070) and the Australian Research Council.

†Lefschetz Center for Dynamical Systems, Division of Applied Mathematics, Brown University, Providence, R.I. 02912 USA. dupuis@cfm.brown.edu

‡Department of Engineering, Australian National University, Canberra, ACT 0200 Australia. Matthew.James@anu.edu.au +61 2 6249 4889

§Department of Electrical Engineering, Australian Defence Force Academy, Canberra, ACT 2600 Australia. irp@routh.ee.adfa.oz.au

were the following. Given a model for a true system, it is almost inevitable that there are dynamics that are left unmodeled, that assumptions regarding distributions are only approximate and even then may change with time, and so on. In such a situation, one may question the value of a control or regulation scheme that is designed under an idealized model. The goal of robust control was the development of design criteria that could deal with such model uncertainty. Optimality with respect to one of the standard cost criteria (e.g., L^2 norm of the output) was rejected, and instead the objective of the design criteria was good qualitative behavior (e.g., stability) across a wide range of system models. For example, a basic result within the theory is the *disturbance attenuation bound* [BB95], e.g., which guarantees a bound on a norm of the state of the system in terms of a corresponding norm on the disturbances that enter the system. This provides a sort of guaranteed performance that is *independent* of the statistical properties of the disturbances. This approach to control design is completely deterministic, with the disturbances modeled simply by unknown functions taking values in an appropriate space (e.g., an L^2 space). A second important result is the *small gain theorem* [Zam66, Vid93]. If controls are designed with respect to a nominal model and achieve a certain performance measure (the L^2 gain), then the small gain theorem characterizes a range of system models (which hopefully will include the "true" model) for which good performance is guaranteed. Such results are not possible with standard design criteria, and provide a quantitative measure of the additional robustness properties that are possible when robust criteria are used. There has been an enormous amount of work on robust control, including extensions to nonlinear systems and formulations in terms of differential games [BB95]. In the differential game formulation there are two players, one representing the disturbance entering the system which will attempt to degrade system performance, and the other representing the "true" control in the system.

A risk-sensitive criterion for stochastic optimal control was introduced by Jacobson [Jac73]. This performance criterion is defined as the average-of-exponential cost associated with a controlled process. In the context of linear-gaussian systems and quadratic cost, Jacobson solved for the optimal state feedback controller which is expressed in terms of a sign-indefinite Riccati equation, and further he established an equivalence with a deterministic dynamic game problem. Interest in the connections between risk–sensitive control and robust control starts with Glover and Doyle [GD88] and continues with many papers. A typical finite time risk–sensitive control problem is the following. The system dynamics are of the form

$$dx^\varepsilon = b(X^\varepsilon, u)dt + \varepsilon^{\frac{1}{2}}\sigma(X^\varepsilon)dw,\ X^\varepsilon(0) = x,$$

and with minimal cost function

$$V^\varepsilon(x) = \inf_u \varepsilon \log E_x \exp\left(\frac{\mu}{\varepsilon}\left[\int_0^T c(X^\varepsilon, u)dt\right]\right).$$

Here the infimum is over an appropriate class of controls, $\varepsilon \in (0, \infty)$, and μ is a parameter related to the L^2 gain in robust control. The connection to robust control is obtained in the large deviation limit $\varepsilon \to 0$. Under appropriate conditions, one can prove that $V^\varepsilon(x)$ converges to the value of a deterministic differential game, with one player in the game representing the original "true" control in the risk–sensitive problem, and the other playing the role of the disturbance. Analogous results have also been obtained for infinite time problems. Since the differential game can be viewed as a state space formulation of a certain robust control problem, this suggests that the controls designed on the basis of the risk–sensitive criteria should also exhibit certain robustness properties.

It is perhaps surprising that controls designed using one of the most sensitive limits in probability, namely a large deviation limit, should be expected to be robust. The standard argument is that since the risk–sensitive control is based on a very conservative exponential cost criteria, and since boundedness of such an exponential cost for any given control should indicate strong stabilizing properties, properties such as those associated to controls based on robust criteria should be expected. Interesting interpretations of the risk-sensitive criterion are presented in [vS96].

The purpose of the present paper is to precisely characterize and prove these robustness properties. In particular, we establish a stochastic version of the small gain theorem (§3). This theorem (Theorem 3.1) is expressed in terms of an inequality which bounds the average output power in terms of the input power. Since this inequality is closely related to the risk-sensitive criterion, our stochastic small gain theorem can be expressed in terms of the risk-sensitive criterion (Corollary 3.3). This provides a concrete motivation for the use of the risk-sensitive criterion stochastic robustness.

Before presenting the stochastic small gain theorem, we begin in §2 with an analysis of the risk-sensitive criterion. In particular, we point out the connection with the above-mentioned power gain inequality. Central to this is a duality relationship involving the risk-sensitive criterion and the relative entropy of probability measures associated with nominal and actual systems. In §4 we consider two extensions of our ideas. Robust control via the solution of a risk-sensitive optimal control problem is considered in §4.1, and in §4.2 our results are extended to allow for uncertainty in the diffusion coefficients.

2 Analysis of a Risk–Sensitive Cost

We consider stochastic differential equations (SDEs) of the form

$$G_0 \ : \ \begin{cases} dX_t &= \bar{b}(X_t)dt + \varepsilon^{\frac{1}{2}}\sigma(X_t)dw_t, \\ \\ Z_t &= h(X_t), \end{cases} \tag{2.1}$$

where X is $n-$dimensional and w is a $d-$dimensional standard Wiener process with $d \in \{1, ..., n\}$. A state feedback control may incorporated into the definition of $\bar{b}$, in which case we have $\bar{b}(X_t) = b(X_t, u(X_t))$ for an appropriate function $b(\cdot,\cdot)$. We will also consider a functional $Z_t = h(X_t)$, where h is a Borel measurable function. Z isolates the state variables whose behavior must be controlled, and this process will take values in $I\!R^m$. We assume that the SDE (2.1) has a unique strong solution on $[0, \infty)$ for each initial condition in $I\!R^n$, and denote the underlying probability space by $(\Omega, \mathcal{F}, P)$. The system (2.1), referred to as the *nominal system*, is illustrated in Figure 1.

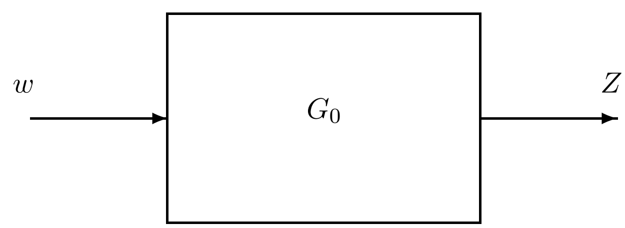

Figure 1: Nominal system

Let $\gamma > 0$ and $\varepsilon > 0$ be given. For $x \in I\!R^n$ and $T \in [0, \infty)$ we define

$$S(x, T) \doteq E_x \left[\exp \frac{1}{2\gamma^2\varepsilon} \int_0^T |Z_t|^2 dt \right], \tag{2.2}$$

where E_x denotes expectation conditioned on $X_0 = x$. This quantity is always well defined, although it may equal ∞.

Let $\mathcal{F}_t$ be a filtration satisfying the "usual conditions" of [KS88, p. 10], with the property that w is an $\mathcal{F}_t-$Wiener process. Let $\mathcal{V}_T$ denote the set of all $I\!R^d-$valued $\mathcal{F}_t-$progressively measurable processes v_t that are also square integrable:

$$\|v\|_T^2 \doteq E\left[\frac{1}{2}\int_0^T |v_t|^2 dt\right] < \infty.$$

We will also make use of finite power signals $v \in \mathcal{P}$ where

$$\|v\|_{power}^2 \doteq \limsup_{T\to\infty} \frac{1}{T}\|v\|_T^2.$$

and

$$\mathcal{P} \doteq \left\{ v_t \in \bigcap_{T>0} \mathcal{V}_T : \|v\|_{power} < \infty \right\}.$$

We have the following representation for $S(x,T)$: for each $x \in I\!R^n$ and $T \in [0,\infty)$,

$$\gamma^2 \varepsilon \log S(x,T) = \sup_{v\in\mathcal{V}_T} E_x\left[\frac{1}{2}\int_0^T \left(|\bar{Z}_t|^2 - \gamma^2|v_t|^2\right) dt\right], \tag{2.3}$$

where $\bar{Z}_t = h(\bar{X}_t)$, and $\bar{X}_t$ is the unique strong solution to the SDE

$$G \ : \ \begin{cases} d\bar{X}_t &= \bar{b}(\bar{X}_t)dt + \sigma(\bar{X}_t)v_t dt + \varepsilon^{\frac{1}{2}}\sigma(\bar{X}_t)dw_t, \\ \bar{Z}_t &= h(\bar{X}_t). \end{cases} \tag{2.4}$$

The representation (2.3) is discussed in the appendix. We shall refer to (2.4) as the *true system*, and view it as a *perturbation* of the nominal system, as illustrated in Figure 2.

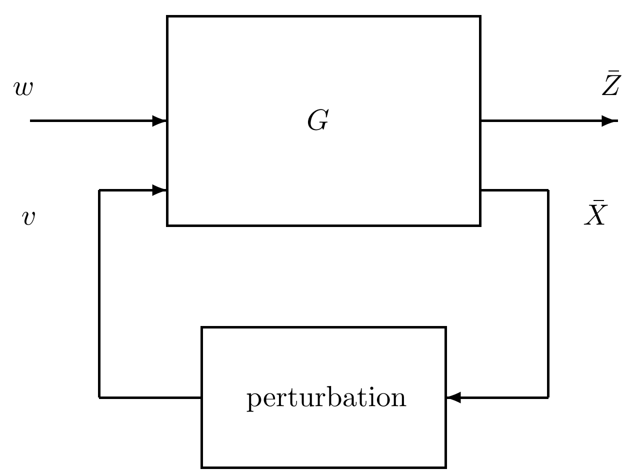

Figure 2: Actual system (a perturbation of the nominal).

The representation formula (2.3) gives

$$E_x\left[\frac{1}{2}\int_0^T |\bar{Z}_t|^2 dt\right] \le \gamma^2 E_x\left[\frac{1}{2}\int_0^T |v_t|^2 dt\right] + \gamma^2\varepsilon \log S(x,T) \tag{2.5}$$

for all $v \in \mathcal{V}_T$. Suppose the following bound on the large time asymptotics of the risk–sensitive cost holds:

$$\lambda^{\gamma,\varepsilon} \doteq \limsup_{T\to\infty} \frac{\gamma^2\varepsilon}{T} \log S(x,T) < \infty. \tag{2.6}$$

Finiteness of this cost in the present risk–sensitive context is closely related to a corresponding H_∞ bound in the context of deterministic robust control. Suppose we consider the small noise limit and assume $\limsup_{\epsilon\to 0}\lambda^{\gamma,\epsilon} = \lambda^\gamma < \infty$. It is shown in [FJ95, FM95] under this condition that there is finite $\beta^\gamma(x)$ so that for all $T > 0$

$$\frac{1}{2}\int_0^T |\bar{Z}_t|^2 dt \le \gamma^2 \frac{1}{2}\int_0^T |v_t|^2 dt + \lambda^\gamma T + \beta^\gamma(x), \tag{2.7}$$

where now the system is deterministic:

$$G \ : \ \begin{cases} \dot{\bar{X}}_t & = \bar{b}(\bar{X}_t) + \sigma(\bar{X}_t)v_t, \ \bar{X}_0 = x, \\ \\ \bar{Z}_t & = h(\bar{X}_t). \end{cases} \tag{2.8}$$

Inequality (2.7) is called a *power gain inequality*, and is studied in [DJ97]. The H_∞ or L_2 gain (in this deterministic context) corresponds to $\lambda^\gamma = 0$ (see [FJ95]); i.e.,

$$\frac{1}{2}\int_0^T |\bar{Z}_t|^2 dt \le \gamma^2 \frac{1}{2}\int_0^T |v_t|^2 dt + \beta^\gamma(x).$$

Returning to the stochastic case and using (2.5), (2.6) implies

$$\limsup_{T\to\infty}\frac{1}{T}E_x\left[\frac{1}{2}\int_0^T |\bar{Z}_t|^2 dt\right] \le \gamma^2 \limsup_{T\to\infty}\frac{1}{T}E_x\left[\frac{1}{2}\int_0^T |v_t|^2 dt\right] + \lambda^{\gamma,\epsilon} \tag{2.9}$$

for all $v \in \mathcal{P}$. In analogy with (2.7) we have for all $T > 0$

$$E_x\left[\frac{1}{2}\int_0^T |\bar{Z}_t|^2 dt\right] \le \gamma^2 E_x\left[\frac{1}{2}\int_0^T |v_t|^2 dt\right] + \lambda^{\gamma,\varepsilon}T + \beta^{\gamma,\epsilon}(x) \tag{2.10}$$

for some finite $\beta^{\gamma,\epsilon}(x)$ and all $v \in \mathcal{P}$. The following table summarizes the different possible interpretations of the bound (2.10):

λ	Interpretation
$\lambda^{\gamma,\epsilon} < \infty$	Stochastic finite power gain
$\lambda^{\gamma} < \infty$	Deterministic finite power gain
$\lambda^{\gamma} = 0$	Deterministic $\mathbf{L}_2$ gain

It is important to observe that even though a particular model is used in the definition of the risk–sensitive cost, the duality formula introduces the family of models (2.4). In the sequel, we will often interpret the model defined by (2.1) as an approximation to a more precise mathematical model of the physical process that is being controlled. One may wish to consider the surrogate (2.1) simply because an exact determination of the true model is impossible or very difficult. Alternatively, one may prefer (2.1) to a more detailed model because of analytical or computational motivations. Suppose that the "true" model takes the form (2.4) for some progressively measurable process $v_\cdot$. In this case the bound (2.5) is very attractive, in that it provides bounds for an ordinary (i.e., not risk– sensitive) cost for the true model (2.4) in terms of the risk–sensitive cost under the design model (2.1), plus a term that measures the size of the difference $v_\cdot$ between the design and the true models.

Several interesting choices are possible for v. What may be the simplest choice is the case where v is an exogenous noise (i.e., a process that is independent of the driving Wiener process) that perturbs the drift as indicated by (2.4). Another choice is $v_t \doteq b_1(\bar{X}_t)$. In this case, $\Delta b(\bar{X}_t) \doteq \sigma(\bar{X}_t) b_1(\bar{X}_t)$ can be interpreted as an error in the drift of the design model. However, in all cases it is apparent that the true model can differ from the design only in the drift term, and that the local quadratic variation (as a function of position) must be the same for all models. This is due to the presence of the relative entropy in the convex duality formula relating risk–sensitive and ordinary costs (see the Section 4.2 and [DE97]). In order for the bounds obtained from this formula to be meaningful, the relative entropy must be finite, and this in turn restricts the class of true models in the diffusion case to those obtained by perturbing the drift of the design model. Thus the class of comparable true models was determined by the choice of the design. We will return to this issue in §4.2, where we indicate how one can modify the design model if one wishes to accommodate modeling errors in quantities such as the quadratic variation.

In the next two sections the bound (2.10) will be used to identify robust properties of a control that makes the risk–sensitive cost $\lambda^{\gamma,\varepsilon}$ finite.

3 A Stochastic Small Gain Theorem

The *small gain theorem* is a basic and important result widely used in the analysis of feedback systems, see, e.g., [vdS96, Vid93]. In this section we give a stochastic generalization of this result. Referring to Figure 3, we are concerned with the stability of the closed loop system depicted, for a range of exogenous inputs $v^{(1)}$, $v^{(2)}$.

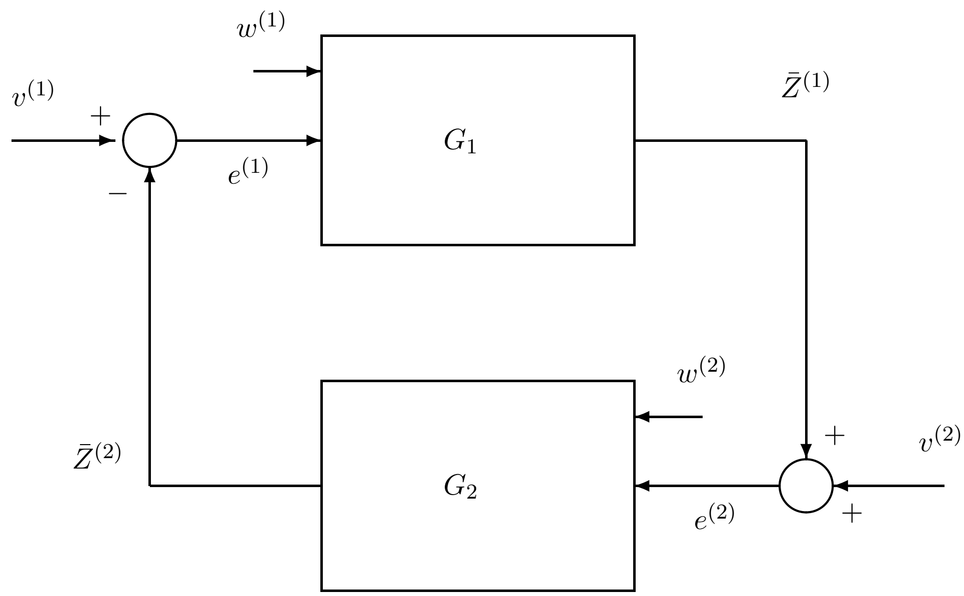

Figure 3: Small gain theorem.

By L^2_{loc} *stability* we mean that if $v^{(i)}$, $i = 1, 2$ are mean square integrable on $[0, T]$, then all signals in Figure 3 also are mean square integrable on $[0, T]$. By *power stability* we mean that if $v^{(i)}$, $i = 1, 2$ have finite power

$$\|v^{(i)}\|^2_{power} = \limsup_{T\to\infty} \frac{1}{T} E\left[\int_0^T |v^{(i)}(t)|^2 dt\right] < \infty, \tag{3.11}$$

then all signals in Figure 3 also have finite power.

The systems G_i, $i = 1, 2$ in Figure 3 are modelled by the dynamics (c.f. (2.4))

$$G_i \;:\; \begin{cases} d\bar{X}_t^{(i)} &= \bar{b}^{(i)}(\bar{X}_t^{(i)})dt + \sigma^{(i)}(\bar{X}_t^{(i)})e_t^{(i)}dt + \varepsilon^{\frac{1}{2}}\sigma^{(i)}(\bar{X}_t^{(i)})dw_t^{(i)}, \\ \bar{Z}_t^{(i)} &= h^{(i)}(\bar{X}_t^{(i)}). \end{cases} \tag{3.12}$$

We assume the existence of unique strong solutions to this system of equations.

Each system in isolation is assumed to satisfy a finite gain inequality (c.f. (2.10))

$$E_{x^{(i)}}\left[\frac{1}{2}\int_0^T |\bar{Z}_t^{(i)}|^2 dt\right] \leq \gamma_i^2 E_{x^{(i)}}\left[\frac{1}{2}\int_0^T |e_t^{(i)}|^2 dt\right] + \lambda^{(\gamma_i,\epsilon,i)}T + \beta^{(\gamma_i,\epsilon,i)}(x^{(i)}) \tag{3.13}$$

for all $e^{(i)} \in \mathcal{V}_T$, where $E_{x^{(i)}}$ denotes expectation given $\bar{X}_0^{(i)} = x^{(i)}$; this is analogous to the classical deterministic small gain theorem [Vid93, Theorem 6.6.1]. The interconnections in the closed-loop system Figure 3 are defined by the equations

$$\begin{aligned} e^{(1)} &= v^{(1)} - z^{(2)}, \\ e^{(2)} &= v^{(2)} + z^{(1)}. \end{aligned} \tag{3.14}$$

Theorem 3.1 *Consider the closed-loop system Figure 3 and assume that each system G_i satisfies the gain inequality (3.13), for all $T \geq 0$, $i = 1, 2$. Then the closed loop system is L^2_{loc} stable and power stable if*

$$\gamma_1\gamma_2 < 1. \tag{3.15}$$

Proof. The proof proceeds in the usual way, [Vid93, Theorem 6.6.1], except with modifications arising from our stochastic context.

Now from (3.13)

$$\begin{aligned} \|z^{(1)}\|_T^2 &\leq \gamma_1^2\|e^{(1)}\|_T^2 + \lambda^{(\gamma_1,\epsilon,1)}T + \beta^{(\gamma_1,\epsilon,1)}(x^{(1)}) \\ \|z^{(2)}\|_T^2 &\leq \gamma_2^2\|e^{(2)}\|_T^2 + \lambda^{(\gamma_2,\epsilon,2)}T + \beta^{(\gamma_2,\epsilon,2)}(x^{(2)}) \end{aligned} \tag{3.16}$$

and by (3.14)

$$\begin{aligned} \|e^{(1)}\|_T &\leq \|v^{(1)}\|_T + \|z^{(2)}\|_T \\ \|e^{(2)}\|_T &\leq \|v^{(2)}\|_T + \|z^{(1)}\|_T \end{aligned} \tag{3.17}$$

By assumption (3.15), there exists $\alpha > 0$ such that $(1+\alpha)^2\gamma_1^2\gamma_2^2 < 1$. Now for $a, b \geq 0$ we have $(a+b)^2 \leq (1+1/\alpha)a^2 + (1+\alpha)b^2$, and so (3.17) implies

$$\begin{aligned} \|e^{(1)}\|_T^2 &\leq (1+1/\alpha)\|v^{(1)}\|_T^2 + (1+\alpha)\|z^{(2)}\|_T^2 \\ \|e^{(2)}\|_T^2 &\leq (1+1/\alpha)\|v^{(2)}\|_T^2 + (1+\alpha)\|z^{(1)}\|_T^2 \end{aligned} \tag{3.18}$$

Substitution of (3.18) into (3.16) gives

$$\begin{aligned} \|z^{(1)}\|_T^2 &\leq \gamma_1^2(1+1/\alpha)\|v^{(1)}\|_T^2 + \gamma_1^2(1+\alpha)\|z^{(2)}\|_T^2 + \lambda^{(\gamma_1,\epsilon,1)}T \\ &\qquad +\beta^{(\gamma_1,\epsilon,1)}(x^{(1)}) \\ \|z^{(2)}\|_T^2 &\leq \gamma_2^2(1+1/\alpha)\|v^{(2)}\|_T^2 + \gamma_2^2(1+\alpha)\|z^{(1)}\|_T^2 + \lambda^{(\gamma_2,\epsilon,2)}T \\ &\qquad +\beta^{(\gamma_2,\epsilon,2)}(x^{(2)}), \end{aligned} \tag{3.19}$$

and solving gives

$$\begin{aligned} \|z^{(1)}\|_T^2 &\leq [1-(1+\alpha)^2\gamma_1^2\gamma_2^2]^{-1}\{\gamma_1^2(1+1/\alpha)\|v^{(1)}\|_T^2 \\ &\qquad + \gamma_1^2\gamma_2^2(1+1/\alpha)(1+\alpha)\|v^{(2)}\|_T^2 + \gamma_1^2(1+\alpha)\lambda^{(\gamma_2,\epsilon,2)}T \\ &\qquad +\gamma_1^2(1+\alpha)\beta^{(\gamma_2,\epsilon,2)}(x^{(2)}) + \lambda^{(\gamma_1,\epsilon,1)}T + \beta^{(\gamma_1,\epsilon,1)}(x^{(1)})\} \\ \|z^{(2)}\|_T^2 &\leq [1-(1+\alpha)^2\gamma_1^2\gamma_2^2]^{-1}\{\gamma_2^2(1+1/\alpha)\|v^{(2)}\|_T^2 \\ &\qquad + \gamma_1^2\gamma_2^2(1+1/\alpha)(1+\alpha)\|v^{(1)}\|_T^2 + \gamma_2^2(1+\alpha)\lambda^{(\gamma_1,\epsilon,1)}T \\ &\qquad +\gamma_2^2(1+\alpha)\beta^{(\gamma_1,\epsilon,1)}(x^{(1)}) + \lambda^{(\gamma_2,\epsilon,2)}T + \beta^{(\gamma_2,\epsilon,2)}(x^{(2)})\}. \end{aligned} \tag{3.20}$$

Inequalities (3.20) imply the L^2_{loc} stable property for each $T \geq 0$, and on dividing (3.20) by T and sending $T \to \infty$ we obtain power stability:

$$\begin{aligned} \|z^{(1)}\|^2_{power} &\leq [1-(1+\alpha)^2\gamma_1^2\gamma_2^2]^{-1}\{\gamma_1^2(1+1/\alpha)\|v^{(1)}\|^2_{power} \\ &\quad + \gamma_1^2\gamma_2^2(1+1/\alpha)(1+\alpha)\|v^{(2)}\|^2_{power} \\ &\quad + \gamma_1^2(1+\alpha)\lambda^{(\gamma_2,\epsilon,2)} + \lambda^{(\gamma_1,\epsilon,1)}\} \\ \|z^{(2)}\|^2_{power} &\leq [1-(1+\alpha)^2\gamma_1^2\gamma_2^2]^{-1}\{\gamma_2^2(1+1/\alpha)\|v^{(2)}\|^2_{power} \\ &\quad + \gamma_1^2\gamma_2^2(1+1/\alpha)(1+\alpha)\|v^{(1)}\|^2_{power} \\ &\quad + \gamma_2^2(1+\alpha)\lambda^{(\gamma_1,\epsilon,1)} + \lambda^{(\gamma_2,\epsilon,2)}\}. \end{aligned} \tag{3.21}$$

□

Remark 3.2 *Note that we get explicit bounds regarding the gain (3.20) and power gain (3.21).*

In view of (2.10), which shows that the gain inequality follows from the risk-sensitive criterion (2.6), we can restate the small gain theorem as follows.

Corollary 3.3 *Consider the closed-loop system Figure 3 and assume that each system G_i satisfies*

$$\lambda_1^{\gamma_1,\epsilon} < +\infty, \quad \lambda_2^{\gamma_2,\epsilon} < +\infty, \tag{3.22}$$

and $\gamma_1\gamma_2 < 1$. Then the closed-loop system is power stable.

We next apply the small gain theorem to obtain a *robust stability* result. The actual system is assumed to be

$$G \; : \; \begin{cases} d\bar{X}_t &= \bar{b}(\bar{X}_t)dt + \Delta\bar{b}(\bar{X}_t)dt + \varepsilon^{\frac{1}{2}}dw_t, \\ \bar{Z}_t &= \bar{X}_t \end{cases} \tag{3.23}$$

(i.e., (2.4) with $\sigma(x) = I$, $v_t = \Delta\bar{b}(\bar{X}_t)$, $h(x) = x$). The corresponding nominal system is

$$G_0 \ : \ \begin{cases} dX_t &= \bar{b}(X_t)dt + \varepsilon^{\frac{1}{2}} dw_t, \\ Z_t &= X_t \end{cases} \tag{3.24}$$

(i.e., (2.1) with $\sigma(x) = I$, $h(x) = x$). This is illustrated in Figure 4.

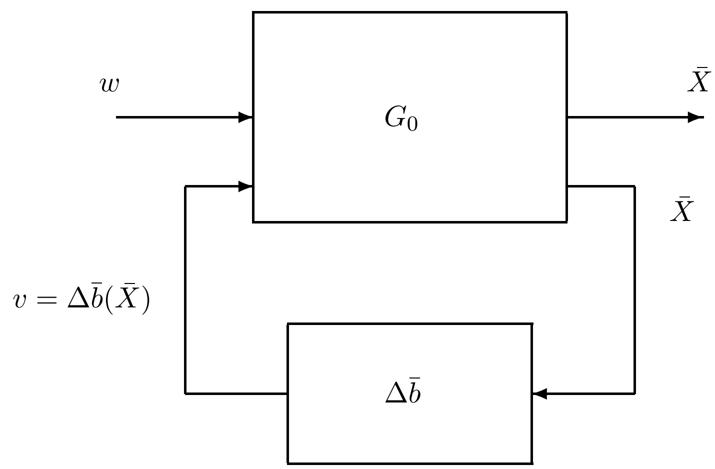

Figure 4: Robust stability

We assume that the perturbed drift satisfies

$$|\Delta\bar{b}(x)|^2 \leq C_{\Delta,1}|x|^2 + C_{\Delta,2}. \tag{3.25}$$

Theorem 3.4 *Consider the system (3.23) illustrated in Figure 4. Assume that*

$$\lambda^{\gamma,\epsilon} < +\infty, \quad \text{and} \quad \gamma^2 C_{\Delta,1} < 1.$$

Then the system (3.23) is power stable:

$$\|\bar{X}\|_{power} < +\infty.$$

PROOF. It follows from (2.10) that

$$\|\bar{X}\|_T^2 \leq \gamma^2 \|v\|_T^2 + \lambda^{\gamma,\epsilon} T + \beta^{\gamma,\epsilon}(x) \tag{3.26}$$

for all $T \geq 0$, where $v = \Delta\bar{b}(\bar{X})$. Substituting (3.25) into (3.26) we get

$$\|\bar{X}\|_T^2 \leq \gamma^2 C_{\Delta,1} \|\bar{X}\|_T^2 + \gamma^2 C_{\Delta,2} + \lambda^{\gamma,\epsilon} T + \beta^{\gamma,\epsilon}(x), \tag{3.27}$$

so that

$$\|\bar{X}\|_T^2 \leq [1 - \gamma^2 C_{\Delta,1}]^{-1} \{\gamma^2 C_{\Delta,2} T + \lambda^{\gamma,\epsilon} T + \beta^{\gamma,\epsilon}(x)\}. \tag{3.28}$$

Dividing inequality (3.28) by T and sending $T \to \infty$ we obtain

$$\|\bar{X}\|_{power}^2 \leq [1 - \gamma^2 C_{\Delta,1}]^{-1} \{\gamma^2 C_{\Delta,2} + \lambda^{\gamma,\epsilon}\}. \tag{3.29}$$

□

4 Extensions

Although we only consider the case of uncontrolled diffusion processes in Sections 2 and 3, many analogous results are also possible. In this section we consider briefly two such extensions. In §4.1 we discuss the case of controlled diffusions and the synthesis of controllers with the objective of obtaining robust stability as discussed above, and in §4.2 we consider cases where the diffusion coefficient can be uncertain.

4.1 Control

We consider the case of controlled diffusions with the same infinite horizon risk-sensitive cost. The controlled diffusion is equation (2.1), where $\bar{b}(X_t) = b(X_t, u(X_t))$, where we restrict to static state feedback controls $u : I\!R^n \to \mathcal{U}$, where $\mathcal{U}$ is the control space. In such a setting dependence of the risk-sensitive cost (2.2) on u will be denoted by $S(x, u, T)$. Consider problems for which one can establish large time asymptotics of the form

$$\lambda^{\gamma,\varepsilon} \doteq \limsup_{T\to\infty} \inf_u \frac{\gamma^2 \varepsilon}{T} \log S(x, u, T) < \infty. \tag{4.30}$$

If u^* is a minimizing (and admissible) feedback controller in (4.30), then the optimal closed-loop system with $\bar{b}(X_t) = b(X_t, u^*(X_t))$ will enjoy the robustness properties described in earlier sections. For details on obtaining the optimal controller, see, e.g., [FM95].

4.2 Noise Coefficient

A second interesting extension involves broadening the class of true models for which bounds can be obtained when a risk- sensitive cost is known to be finite. In this context, it is important to observe the key role that relative entropy plays in linking risk-sensitive and ordinary costs. Let $\mathcal{S}$ be a Polish space (a complete separable metric space) with the Borel $\sigma-$algebra, and let $\mathcal{P}(\mathcal{S})$ denote the set of probability measures defined on $\mathcal{S}$. Given $\theta \in \mathcal{P}(\mathcal{S})$ and $\nu \in \mathcal{P}(\mathcal{S})$, the *relative entropy of* ν *with respect to* θ is defined to be

$$R(\nu\|\theta) \doteq \int_{\mathcal{S}} \left(\log \frac{d\nu}{d\theta}\right) d\nu$$

if ν is absolutely continuous with respect to θ. Otherwise, we set $R(\nu\|\theta) \doteq \infty$.

If f is any non-negative function from $\mathcal{S}$ to the reals then (cf. [DE97])

$$\log \int_{\mathcal{S}} e^f d\theta = \sup_{\nu \in \mathcal{P}(\mathcal{S})} \left[\int_{\mathcal{S}} f d\nu - R(\nu\|\theta)\right].$$

With θ interpreted as the design model and ν as the true, it follows that whenever ν is absolutely continuous with respect to θ

$$\int_{\mathcal{S}} f d\nu \leq \log \int_{\mathcal{S}} e^f d\theta + R(\nu\|\theta). \tag{4.31}$$

Thus the absolute continuity of ν with respect to θ is a prerequisite if one wishes to obtain meaningful bounds on ordinary costs in terms of risk-sensitive costs and "distance" between models as measured by relative entropy. The requirement of absolute continuity is not purely technical, in that if one requires bounds for a sufficiently broad class of cost functions f (e.g., all bounded continuous functions), then finiteness of the relative entropy is also a necessary condition [DE97, Lemma 1.4.3]. Thus one possible approach to broadening the class of true systems is to restrict the class of possible cost functions further. An alternative, and the one we shall pursue here, is to modify the design model.

For the diffusion model discussed in Sections 2 and 3, the requirement of bounded relative entropy restricts the class of true models to those obtained by perturbing the drift in the SDE of the design model by a mean square integrable process. However, in certain problems one may wish to consider true models that are other perturbations of the design. For example, in many applications of diffusions the diffusion coefficient is difficult to estimate. If this is the case and if one wishes to continue to use risk-sensitive criteria to deal with this type of model uncertainty, then clearly the design model must be modified. In the rest of this section we will indicate how this can be done for one particular set of circumstances. The example will further illustrate the role that relative entropy plays, and it will also suggest how one might deal with other design models and desired perturbations.

As noted previously, two diffusion processes with different diffusion coefficients induce mutually singular measures. Because of this, one cannot obtain robustness with respect to errors in the diffusion coefficient via risk-sensitive costs if the design model is a diffusion. However, if one is willing to consider a more general class of design processes then it is possible to obtain bounds on a family of diffusion processes with different diffusion coefficients in terms of a single risk-sensitive cost. We will describe what appears to be the simplest possible example, and note that there are many variations and alternatives that will not be mentioned. In particular, we have not attempted to understand if the choice of design model is in any sense best, even though it is clear that this selection will determine the difficulty of evaluating the risk-sensitive cost. Our objective is to simply show that robustness with respect to these alternative types of modeling errors is indeed possible for *some* choice of design model.

For simplicity we restrict to the case of dimension one and fix $\varepsilon = 1$. Suppose one wishes to bound a performance measure of the diffusion processes

$$G \ : \ \begin{cases} d\bar{X}_t & = \bar{b}(\bar{X}_t)dt + \sigma dw_t, \\ \\ \bar{Z}_t & = h(\bar{X}_t), \end{cases} \tag{4.32}$$

Suppose that the values $\sigma \in \{a, A\}$ are to be considered, where $0 < a < A < \infty$. One can construct a process whose induced measure is absolutely continuous with respect to both diffusion processes, as well as a family of closely related processes, as follows. Let $\mathcal{F}_t$ be a filtration that satisfies the usual conditions, and suppose that w is an $\mathcal{F}_t-$Wiener process. Let Ξ be a continuous time Markov chain with states $\{a, A\}$, and suppose that the jump rates between these two states are $\eta_{a,A} \in (0, \infty)$ and $\eta_{A,a} \in (0, \infty)$. We suppose that the process Ξ is independent of the filtration $\mathcal{F}_t$, and define the process

$$G_0 \ : \ \begin{cases} dX_t & = \bar{b}(X_t)dt + \Xi dw_t, \\ \\ Z_t & = h(X_t), \end{cases} \tag{4.33}$$

Let $\mathcal{G}_t$ be the completion of the sigma algebra $\sigma(\mathcal{F}_t, \sigma(\Xi_s, 0 \le s \le t))$. Then w is a $\mathcal{G}_t-$Wiener process, and $\mathcal{G}_t$ satisfies the usual conditions. The pair (X_t, Ξ_t) form a jump Markov process that is progressively measurable with respect to $\mathcal{G}_t$.

For $x \in I\!R$, $\xi \in \{a, A\}$, and $T \in [0, \infty)$, we define

$$S(x, \xi, T) \doteq E_{x,\xi}\left[\exp \frac{1}{2\gamma^2}\int_0^T |Z_t|^2 dt\right],$$

where $E_{x,\xi}$ denotes expectation conditioned on $X_0 = x$ and $\Xi_0 = \xi$. As we will see, the introduction of this new "noise" will broaden the class of true models for which a bound of the form (4.31) is meaningful. The first step is to formulate the representation analogous to (2.3) for this type of process.

Because of the jump component Ξ_t, a representation for the processes (4.33) does not follow directly from the results proved in either [BD98] or [BD97]. However, the arguments there can be adapted with little difficulty to the simple jump diffusion (4.33), and one obtains the following representation. Define the function

$$\ell(a) \doteq \begin{cases} a \log a - a + 1 & \text{if } a \ge 0 \\ \infty & \text{else.} \end{cases}$$

Let $\mathcal{V}_T$ denote the set of all $I\!R-$valued $\mathcal{G}_t-$progressively measurable processes that are mean square integrable, and let $\mathcal{W}_T$ denote the set of all $I\!R-$valued $\mathcal{G}_t-$progressively measurable processes γ_t that satisfy

$$E\left[\int_0^T \ell(\gamma_t) dt\right] < \infty.$$

Given processes $\gamma_{a,A}$ and $\gamma_{A,a}$ in $\mathcal{W}_T$, let $\bar{\Xi}$ be the unique jump process with cadlag sample paths and state space $\{a, A\}$ such that

$$1_{\{\bar{\Xi}_t = a\}} - \int_0^t \left(-\gamma_{a,A} 1_{\{\bar{\Xi}_s = a\}} + \gamma_{A,a} 1_{\{\bar{\Xi}_s = A\}}\right) ds$$

and

$$1_{\{\bar{\Xi}_t = A\}} - \int_0^t \left(-\gamma_{A,a} 1_{\{\bar{\Xi}_s = A\}} + \gamma_{a,A} 1_{\{\bar{\Xi}_s = a\}}\right) ds$$

are $\mathcal{G}_t-$martingales. For any $v \in \mathcal{V}_T$ and $x \in I\!R$, define the processes $\bar{X}_t$ and $\bar{Z}_t$ by

$$G \ : \ \begin{cases} d\bar{X}_t & = \bar{b}(\bar{X}_t)dt + \Xi_t v_t dt + \Xi_t dw_t, \\ \\ \bar{Z}_t & = h(\bar{X}_t). \end{cases} \tag{4.34}$$

We then have the representation

$$\gamma^2 \log S(x,\xi,T) \\ = \sup_{v\in\mathcal{V}_T,\gamma_{a,A},\gamma_{A,a}\in\mathcal{W}_T} E_{x,\xi}\left[\int_0^T \left\{ \begin{array}{l} \frac{1}{2}\left(|\bar{Z}_t|^2 - \gamma^2|v_t|^2\right) \\ -\ell((\gamma_{a,A})_t/\eta_{a,A}) - \ell((\gamma_{A,a})_t/\eta_{A,a}) \end{array} \right\} dt\right].$$

We make use of the fact that $\ell(0) = 1$ and $\ell(1) = 0$. With the choices $(\gamma_{a,A})_t = \eta_{a,A}$ and $(\gamma_{a,A})_t = 0$, the process $\bar{\Xi}$ enters the state A (uniformly in the initial distribution) and stays there for all time. If $\bar{\Xi}(0) = A$ w.p.1 then $\bar{\Xi}$ remains in A for all time. Hence one obtains bounds of the form

$$\gamma^2 \log S(x,A,T) + T \geq E_{x,A}\left[\int_0^T \frac{1}{2}\left(|\bar{Z}_t|^2 - \gamma^2|v_t|^2\right) dt\right],$$

where $\bar{Z}$ will be associated with a diffusion process with diffusion coefficient A. In a similar fashion, one can also establish bounds on the performance of a diffusion process with coefficient a, and also a range of (non diffusion) models with quadratic variation between these two extremes.

Appendix

In this appendix, we discuss equation (2.3) as well as the general duality relation between relative entropy and exponential integrals.

The representation (2.3) is proved in [BD98] for the case where $\mathcal{F}_t$ is the completion of the filtration generated by the Wiener process w. It is straightforward to extend this to the general case where $\mathcal{F}_t$ is any filtration satisfying the usual conditions that makes w a Wiener process, and indeed [BD97] gives this form of the representation for Hilbert space valued Wiener processes. Equation (2.3) is obtained from the more general convex duality formula mentioned in Section 4.2:

$$\log \int_{\mathcal{S}} e^f \, d\theta = \sup_{\nu}\left[\int_{\mathcal{S}} f \, d\nu - R(\nu\|\theta)\right], \tag{0.35}$$

where $\mathcal{S}$ is a Polish space (a complete separable metric space) with the Borel $\sigma-$algebra, $\mathcal{P}(\mathcal{S})$ denoting the set of probability measures defined on $\mathcal{S}$, and

$$R(\nu\|\theta) \doteq \int_{\mathcal{S}} \left(\log \frac{d\nu}{d\theta}\right) d\nu$$

if ν is absolutely continuous with respect to θ and $R(\nu\|\theta) \doteq \infty$ otherwise. General properties of the relative entropy function (convexity, lower semi-continuity, etc.) are discussed in Section 1.4 of [DE97]. Apart from the scaling constants γ and ε, one can formally identify the logarithm of the risk sensitive cost $S(x,T)$ with the left hand side of (0.35), the integral of $|\bar{Z}_t|^2$ with the integral of f with respect to ν, and the integral of $|v_t|^2$ with the relative entropy of the measure induced by $\varepsilon^{\frac{1}{2}}w_{\cdot} + \int_0^{\cdot} v_t dt$ with respect to the measure induced by $\varepsilon^{\frac{1}{2}}w_{\cdot}$ (the identification is only formal because a single underlying probability space appears in (2.3)). Details of a rigorous derivation may be found in the references listed above.

References

[BB95] T. Basar and P. Bernhard. *H_∞ Optimal Control and Related Minimax Design Problems*. Birkhauser, Boston, 1995.

[BD97] A. Budhiraja and P. Dupuis. Representations and large deviations for Hilbert space valued diffusions. Technical report, Brown University, 1997.

[BD98] M. Boué and P. Dupuis. A variational representation for certain functionals of Brownian motion. *The Annals of Prob.*, 1998.

[BvS85] A. Bensoussan and J. H. van Schuppen. Optimal control of partially observable stochastic systems with an exponential-of-integral performance index. *SIAM Journal on Control and Optimization*, 23:599–613, 1985.

[DE97] P. Dupuis and R.S. Ellis. *A Weak Convergence Approach to the Theory of Large Deviations*. Wiley, New York, 1997.

[DJ97] P. Dower and M.R. James. Finite power gain analysis of nonlinear systems. *Int. J. Robust & Nonlinear Control*, 1997.

[DM97] P. Dupuis and W.M. McEneaney. Risk-sensitive and robust escape criteria. *SIAM J. on Control and Opt.*, 1997.

[FHH97] W.H. Fleming and D. Hernandez-Hernandez. Risk-sensitive control of finite state machines on an infinite horizon, I. to appear, SIAM J. Control Optim., 1997.

[FJ95] W.H. Fleming and M.R. James. The risk-sensitive index and the h_2 and H_∞ norms for nonlinear systems. *J. Mathematics of Control, Signals, and Systems*, 8:199–221, 1995. to appear.

[FM95] W.H. Fleming and W.M. McEneaney. Risk-sensitive control on an infinite time horizon. *SIAM J. Control and Optimization*, 33:1881–1915, 1995.

[GD88] K. Glover and J. Doyle. State space formulea for all stabilizing controllers that statisfy an H_∞ norm bound and relations to risk-sensitivity. *Systems and Control Letters*, 11:167–172, 1988.

[HHM96] D. Hernandez-Hernandez and S.I. Marcus. Risk-sensitive control of markov processes in countable state space. *Systems and Control Letters*, pages 147–155, 1996.

[Jac73] D.H. Jacobson. Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Trans. Aut. Control*, 18(2):124–131, 1973.

[Jam92] M.R. James. Asymptotic analysis of nonlinear stochastic risk-sensitive control and differential games. *Mathematics of Control, Signals and Systems*, 5(4):401–417, 1992.

[JBE94] M. R. James, J.S. Baras, and R.J. Elliott. Risk-sensitive control and dynamic games for partially observed discrete-time nonlinear systems. *IEEE Trans. Automatic Control*, 39(4):780–792, 1994.

[KS88] I. Karatzas and S. E. Shreve. *Brownian Motion and Stochastic Calculus*. Springer-Verlag, New York, 1988.

[PMR96] P. Dai Pra, L. Meneghini, and W.J. Runggaldier. Connections between stochastic control and dynamic games. *Math. Control, Signals and Systems*, 9(4):303–326, 1996.

[vdS96] A.J. van der Schaft. *L_2-Gain and Passivity Techniques in Nonlinear Control*. Springer Verlag, New York, 1996.

[Vid93] M. Vidyasagar. *Nonlinear Systems Analysis*. Prentice-Hall, Englewood Cliffs, 1993. Second Edition.

[vS96] J.H. van Schuppen. Risk-sensitivity of stochastic control problems. presentation at MTNS 96, St. Louis, 1996.

[Whi81] P. Whittle. Risk-sensitive linear quadratic Gaussian control. *Advances in Applied Probability*, 13:764–777, 1981.

[Zam66] G. Zames. On the input-output stability of time-varying nonlinear feedback systems. part i: Conditions derived using concepts of loop gain, conicity, and positivity. part ii: Conditions involving circles in the frequency plane and sector nonlinearities. *IEEE Trans. Aut. Control*, 11:228–238, 465–476, 1966.

[Zam81] G. Zames. Feedback and optimal sensitivity: Model reference transformation, multiplicative seminorms and approximate inverses. *IEEE Trans. Aut. Control*, 26:301–320, 1981.