

Exact Analytic Solutions for Optimal Control Problems Under Multiplicative Noise [★]

Dongsung Huh, Terrence Sejnowski

*Univ. of California San Diego, La Jolla, CA 92093 USA,
Howard Hughes Medical Institute, Salk Institute, La Jolla, CA 92037
USA,
e-mail: [huh, terry]@salk.edu*

Abstract:

Control-dependent (multiplicative) noise makes it difficult to achieve optimal control because large control signals amplify noise. This paper considers a minimal (one-dimensional) system that includes multiplicative noise and solves the optimal control problem for arbitrary cost functions. In a limit when the control-cost approaches zero, this formulation becomes analytically solvable. The analysis reveals several important properties that are absent in traditional LQ analysis. In particular, multiplicative noise makes optimal solutions depend on the global features of cost functions.

Keywords: Optimal control theory, Stochastic optimal control, Control dependent noise, Generalized solutions of Hamilton-Jacobi equations

1. INTRODUCTION

For most physical systems with noisy dynamics, the noise statistics generally depends to some extent on the system state and the control signal rather than being strictly constant (additive noise). Often, such signal dependent noise has significant effects in control systems such as aerospace thrust control, digital feedback control, chemical reactors, analog mechanical actuators, biological muscles, and synaptic transmission of neural systems (see McLane [1971], Ibrahim [1985], van Wingerden and De Koning [1984], Wagenaar [1989], Bolotin [1984], Ruan and Choudhury [1993] and Harris and Wolpert [1998]). In these models, noise is often modeled as being amplified by the control signal in a multiplicative manner.

Most control problems with multiplicative noise have been analyzed on the basis of linear control model with quadratic cost functions on state and control signals (see Wonham [1967], Willems and Willems [1976], Phillis [1985], Kubrusly and Costa [1985], Ruan and Choudhury [1993], Yaz and Skelton [1994], Boyd et al. [1994], El Ghaoui [1995], Lu and Skelton [2000], Todorov [2005]).

Multiplicative noise imposes an additional challenge in optimal control because large control signals amplify noise. For convex state-costs, larger noise means increased average cost, which effectively penalizes large control (effective regularization). Therefore, convex state-costs promotes smaller, risk-averse control, which is in direct conflict with the original role of state-costs: promoting fast transition/stabilization toward a desired state.

This naturally leads to the following questions:

- (1) Is it better to use larger or smaller control signals when the state-cost gets steeper? ¹
- (2) What if the state-cost is not everywhere convex? Is it ideal to be risk-seeking and use larger control where the state-cost is locally concave? ²
- (3) Can the effective-regularization due to the state-cost completely substitute the control-cost? ³

These questions cannot be answered in the traditional LQ settings (linear-control, quadratic-cost).

In this paper, we investigate a minimal system that includes multiplicative noise and compute optimal solutions for arbitrary state-cost functions. Especially, we focus on a limiting case when the control-cost approaches zero, which reveals the trade-off between the state-cost and the effective regularization in a most clear manner.

Organization of the paper Section 2 describes the optimal control problem. An incomplete solution is given in section 3. Section 4 provides additional information necessary for completing the solution. Sections 5–9 describe complete solutions for various classes of state-costs. Section 11 generalizes the results to the cases when control-cost exists. Section 12 concludes the paper.

2. PROBLEM DESCRIPTION

2.1 Deterministic problem

Let us begin with describing a deterministic control problem. Consider a one-dimensional dynamical system

$$\dot{x} = f(x, u), \quad (1)$$

¹ Quick Answer: Smaller

² Quick Answer: No

³ Quick Answer: Depends on the state-cost.

[★] This work was supported by Howard Hughes Medical Institute.

which is assumed to have a unique inverse-dynamics

$$u = f^{-1}(x, \dot{x}), \quad (2)$$

where $x \in \mathbf{R}$ is the state, and $u \in \mathbf{R}$ is the control. Then, the cost function, which is typically represented as a function of x and u , can equally well be represented as a function of x and \dot{x} :

$$\mathcal{L}(x, \tau \dot{x}), \quad (3)$$

where τ — a constant with dimension of *time* — is shown explicitly. Such time constant is required for dimensional consistency within the cost function. τ can be regarded as a weight parameter for the control-cost, because it turns out to inversely scale the optimal state-transition-rate. (See section 11.)

2.2 Stochastic problem with multiplicative noise

In this paper, we consider a stochastic problem in which control-dependent (multiplicative) noise is introduced to the dynamics:

$$dx = f(x, u) \cdot (dt + \sigma \cdot d\xi), \quad (4)$$

where $\xi(t) \in \mathbf{R}$ is a standard Wiener process ($E[d\xi] = 0$, $E[d\xi^2] = dt$), and σ is the diffusion-rate constant — with dimension of $\sqrt{\text{time}}$. We call

$$\mu \equiv f(x, u) \quad (5)$$

the average state-transition-rate, or simply *drift*.

The goal is to minimize the expectation of total integrated cost (without temporal discount)

$$E_{\xi(\cdot)} \left[\int_0^\infty \mathcal{L}(x(t), \tau \mu(t)) dt \right]. \quad (6)$$

In this stochastic setting, the effect of control at one state tends to diffuse out to a range of nearby states. The extent of the diffusion is characterized by the ratio of time constants: $\beta \equiv \sigma^2/\tau$. In deterministic limit ($\beta \ll 1$), the control problem becomes highly localized. For larger β , the problem gets much harder to solve, as it requires consideration of longer-range effects of control. (See section 11.)

However, the interesting case is the limit $\beta \gg 1$ (or $\tau \rightarrow 0$), in which the characteristics of control-dependent noise is most emphasized. Ironically, this limit turns out to simplify the control problem, because the cost-function reduces to a pure state-cost (without control-cost):

$$Q(x) \equiv \mathcal{L}(x, 0). \quad (7)$$

In following sections, we analytically solve the control problem for arbitrary state-costs in the zero-control-cost limit ($\tau \rightarrow 0$), and extend the analysis to finite β cases in section 11.

3. HAMILTON-JACOBI-BELLMAN EQUATION

We assume the state-cost $Q(x)$ does not diverge to infinity for finite x and that it has a unique global minimum (which will be removed in section 9). Without loss of generality, we set the global minimum to be zero at $x = 0$:

$$\min_x Q(x) = Q(0) = 0.$$

Evidently, optimal drift at $x = 0$ will be zero, since it stops accumulation of any further cost. Therefore, the problem

is equivalent to a *first-exit* problem whose *exit-state* — where state-transition permanently ends — is $x = 0$.

Solving stochastic optimal control problems typically involves *cost-to-go*, the minimum total expected cost:

$$V(x) \equiv \min_{\mu(\cdot)} E_{\xi(\cdot)} \left[\int_0^\infty Q(\tilde{x}(t)) dt \right], \quad (8)$$

where $\tilde{x}(t)$ is a stochastic trajectory with initial state $\tilde{x}(0) = x$.

For stochastic first-exit problems, cost-to-go satisfies the following Hamilton-Jacobi-Bellman (HJB) equation

$$0 = Q(x) + \min_{\mu} \left[\mu(x) V_x(x) + \frac{1}{2} \sigma^2 \mu(x)^2 V_{xx}(x) \right], \quad (9)$$

where subscripts represent differentiations. This equation is non-trivial to solve because the minimum operation makes it a non-linear differential equation.

Once cost-to-go is obtained, optimal drift can be found in a closed form (from the minimum operation of (9)):

$$\mu^*(x) = -\frac{V_x(x)}{\sigma^2 V_{xx}(x)}, \quad (10)$$

and optimal control law follows from inverse-dynamics

$$u^*(x) = f^{-1}(x, \mu^*(x)). \quad (11)$$

Note that cost-to-go must be convex ($V_{xx} \geq 0$), since otherwise the minimum in (9) does not exist.

3.1 Solution of the nonlinear HJB equation

Substituting (10) to HJB (9), we obtain a nonlinear differential equation

$$Q = \frac{V_x^2}{2\sigma^2 V_{xx}}. \quad (12)$$

However, if we take the reciprocal of (12), this becomes a linear differential equation

$$\frac{1}{2\sigma^2 Q} = \frac{V_{xx}}{V_x^2} = -\frac{d}{dx} \left(\frac{1}{V_x} \right) \quad \text{for } x \neq 0, \quad (13)$$

and $1/V_x$ can be calculated by a simple integration

$$\frac{1}{V_x(x)} = \int_x^\infty \frac{dx'}{2\sigma^2 Q(x')} + c \quad \text{for } x > 0, \quad (14)$$

where c is an integration constant, yet undetermined.

Then, according to (10) and (12), optimal drift can be calculated as the ratio between cost and cost-to-go derivative:

$$\mu^*(x) = -\frac{2Q(x)}{V_x(x)} \quad (15)$$

$$= -\frac{Q(x)}{\sigma^2} \left(\int_x^\infty \frac{dx'}{Q(x')} + 2c \right) \quad x > 0. \quad (16)$$

Solution for $x < 0$ can be obtained in a similar manner. In this paper, however, we shall only consider solutions for $x > 0$ for simplicity. Readers may assume state-costs being symmetric in x , while $V_x(x)$ and $\mu^*(x)$ anti-symmetric.

The solution obtained above will be complete once the integration constant is determined. Initially, we guessed this could be done by further analyzing the HJB eq (9). Indeed a useful result can be derived from the convexity property implied in (9):

Lemma 1. c must be non-negative ($c \geq 0$).

Proof. Remind that cost-to-go must be convex. According to (14), however, $1/V_x$ is a monotonically decreasing function that approaches c as $x \rightarrow \infty$. For the convexity condition to be satisfied, V_x must be non-negative for $x > 0$. Therefore, $c \geq 0$. \square

Unfortunately, this is as far as one can go with HJB equation alone. The truth is, it does not contain sufficient information to uniquely determine the cost-to-go function. To fill in the missing information, other properties of cost-to-go must be investigated as well.

4. FUNDAMENTAL PROPERTIES OF $V(X)$

In this section, we enumerate several essential properties of cost-to-go that follow from the definition (8). Some are very simple, and do not require proofs.

Proposition 2. (Axes rescaling)

If two cost functions are related by

$$Q_2(x) = Q_1\left(\frac{x}{b}\right) \quad b > 0, \quad (17)$$

then the same holds for cost-to-go functions, $V_2(x) = V_1(x/b)$, and their derivative are

$$V_{2x}(x) = \frac{1}{b} V_{1x}\left(\frac{x}{b}\right). \quad (18)$$

Similarly, if two cost functions are related by multiplication by a positive constant

$$Q_2(x) = aQ_1(x) \quad a > 0. \quad (19)$$

then the same holds for cost-to-go functions, $V_2(x) = aV_1(x)$, and their derivative are

$$V_{2x}(x) = aV_{1x}(x). \quad (20)$$

These properties derives from the fact that a control problem does not change when the units change (*e.g.* from meters to feet). In other words, a control problem is invariant to axis re-scaling.

Lemma 3. If two cost functions are related by

$$Q_1(x) \leq Q_2(x) \quad \forall x > 0, \quad (21)$$

then the same holds for cost-to-go functions: , *i.e.*

$$V_1(x) \leq V_2(x) \quad \forall x > 0. \quad (22)$$

Proof. Given a drift function $\mu(x)$ and a cost function $Q_i(x)$, define $\tilde{V}_i(x; \mu)$ as the expected total cost:

$$\tilde{V}_i(x; \mu) \equiv E \left[\int_0^\infty Q_i(\tilde{x}(t)) dt \right] \quad \tilde{x}(0) = x. \quad (23)$$

Then, $\tilde{V}_1(x; \mu) \leq \tilde{V}_2(x; \mu)$, since

$$E \left[\int_0^\infty (Q_1(\tilde{x}(t)) - Q_2(\tilde{x}(t))) dt \right] \leq 0. \quad (24)$$

Then, $V_1(x) = \min_\mu \tilde{V}_1(x; \mu) \leq \tilde{V}_1(x; \mu_2^*) \leq \tilde{V}_2(x; \mu_2^*) = V_2(x)$, where μ_2^* is the optimal drift for Q_2 . \square

Proposition 4. (Shifting and zero-padding)

Consider a cost $Q_1(x)$ and its cost-to-go $V_1(x)$. Then, for the shifted and zero-padded cost

$$Q_2(x) = \begin{cases} 0 & 0 \leq x \leq x_0 \\ Q_1(x - x_0) & x_0 < x \end{cases} \quad (25)$$

the resulting cost-to-go gets shifted and zero-padded by the same amount — x_0 :

$$V_2(x) = \begin{cases} 0 & 0 \leq x \leq x_0 \\ V_1(x - x_0) & x_0 < x \end{cases} \quad (26)$$

Lemma 5. If $c > 0$, $V(x)$ is upper-bounded by x/c .

Proof. Since $V_x(x)$ is a monotonically increasing function and $\lim_{x \rightarrow \infty} V_x(x) = 1/c$, $V_x(x)$ is upper-bounded by $1/c$. And $V(0) = 0$. \square

5. MONOMIAL COST

For a (superlinear) monomial cost-function

$$Q_1(x) = x^{1+\epsilon} \quad \epsilon > 0, \quad x > 0, \quad (27)$$

cost-to-go derivative (14) is

$$\frac{1}{V_{1x}(x)} = \frac{1}{2\sigma^2\epsilon x^\epsilon} + c. \quad (28)$$

Here, we use the invariant property (proposition 2) to determine the integration constant.

Lemma 6. For monomial costs, $c = 0$.

Proof. Consider two positive constants a, b that are related by $a = b^{1+\epsilon}$. Then, a re-scaled cost $Q_2(x) \equiv aQ_1(x/b)$ is identical to the original cost $Q_1(x)$:

$$Q_2(x) = a(x/b)^{1+\epsilon} = x^{1+\epsilon}. \quad (29)$$

Then, it follows the cost-to-go's should also be identical: $V_2(x) = V_1(x)$. According to (18) and (20), however,

$$\begin{aligned} \frac{1}{V_{2x}(x)} &= \frac{b}{a} \left(\frac{1}{2\sigma^2\epsilon(x/b)^\epsilon} + c \right) \\ &= \frac{1}{V_{1x}(x)} + c \left(\frac{1}{b^\epsilon} - 1 \right). \end{aligned} \quad (30)$$

Since b is arbitrary, $\therefore c = 0$. \square

Therefore, we obtain following cost-to-go

$$V_1(x) = \frac{2\sigma^2\epsilon}{1+\epsilon} x^{1+\epsilon}, \quad (31)$$

and its derivative

$$V_{1x}(x) = 2\sigma^2\epsilon x^\epsilon, \quad (32)$$

Intriguingly, optimal drift is linear for a monomial cost of any (superlinear) power:

$$\mu_1^*(x) = -\frac{2Q_1(x)}{V_{1x}(x)} = -\frac{x}{\sigma^2\epsilon}. \quad (33)$$

Moreover, the slope of optimal drift is inversely proportional to the power ϵ . This is the first evidence that a steeper state-cost leads to a slower drift.

Note that in the limit $\epsilon \rightarrow 0$ (linear monomial), cost-to-go and its derivative (31,32) collapse to zero and optimal drift (33) diverges to infinity.

6. SUPERLINEAR COST

In this section, we generalize lemma 6 to a more general class of cost-functions.

Definition: A cost-function $Q_1(x)$ is *superlinear*, if $\exists a, \epsilon > 0, \exists x_0 \geq 0$ s.t.

$$Q_1(x) \geq Q_2(x) \quad \forall x > 0, \quad (34)$$

where $Q_2(x)$ is a shifted monomial

$$Q_2(x) = \begin{cases} 0 & 0 \leq x \leq x_0 \\ a(x - x_0)^{1+\epsilon} & x_0 < x \end{cases} \quad (35)$$

Lemma 7. For superlinear costs, $c = 0$.

Proof. If $c > 0$, $V_1(x)$ is upper-bounded by a linear function $-x/c$ (lemma 5). Since $Q_1(x) \geq Q_2(x)$, however, $V_1(x)$ is lower-bounded by $V_2(x)$ (lemma 3), which grows superlinearly — eq (31). $\therefore c = 0$ by contradiction. \square

With lemma 7, the results from monomial cost-functions readily generalize to "polynomial" cost-functions:

$$\frac{1}{Q(x)} = \sum_i \frac{a_i}{x^{1+\epsilon_i}} \quad \epsilon_i > 0, \quad (36)$$

which covers a much broader class of cost-functions. It is not necessary for all coefficients a_i to be positive, but $Q(x)$ must be positive for all $x > 0$. Cost-to-go derivative is

$$\frac{1}{V_x(x)} = \frac{1}{2\sigma^2} \sum_i \frac{a_i}{\epsilon_i x^{\epsilon_i}} \quad (37)$$

and optimal drift

$$\mu^*(x) = -\frac{\sum_i a_i / (\epsilon_i x^{\epsilon_i})}{\sigma^2 \sum_i a_i / x^{1+\epsilon_i}}. \quad (38)$$

More generally, the following compositionality theorem holds for all superlinear cost-functions.

Theorem 8. (Compositionality Rule)

If three superlinear cost functions are related by

$$\frac{1}{Q_3(x)} = \frac{a_1}{Q_1(x/b_1)} + \frac{a_2}{Q_2(x/b_2)}, \quad (39)$$

then cost-to-go derivative are related by

$$\frac{1}{V_{3x}(x)} = \frac{a_1 b_1}{V_{1x}(x/b_1)} + \frac{a_2 b_2}{V_{2x}(x/b_2)}. \quad (40)$$

Proof. Because (13) is linear and integration constants are identically zero for all three cost functions. \square

Of course, a cost function need not be represented in a "polynomial" form to be solved. Table 1 and Fig 1 shows some examples. Note that optimal drift functions are no longer linear. Instead, they reflect the shape of state-cost. Also note that non-convex state-costs do not lead to infinite, risk-seeking control. All superlinear state-cost leads to risk-averse control.

7. SUBLINEAR COST

At the end of section 5, we showed that for linear state-costs, cost-to-go collapses to zero and optimal drift diverges to infinity: a linear cost is not steep enough to provide sufficient self-regularization. The same is true for all sublinear state-costs.

Definition: A cost function $Q_1(x)$ is *sublinear*, if $\exists a > 0, \exists b \geq 0$ (yet finite), s.t.

$$Q_1(x) \leq ax + b \quad \forall x \geq 0. \quad (41)$$

(As one can see, this includes linear costs as well.)

Lemma 9. For sublinear costs, cost-to-go is zero.

Proof. Define $Q_2(x)$ as the shifted and zero-padded version of $Q_1(x)$ by $x_0 = b/a$ (as in (25)). Then according to proposition 4

$$Q_2(x) \leq ax \quad \forall x \geq 0. \quad (42)$$

Since cost-to-go for a linear cost is zero, $V_2(x)$ is upper-bounded by zero (lemma 3). $\therefore V_1(x) = V_2(x) = 0$. \square

The following example shows the process of cost-to-go collapsing as the state-cost becomes sublinear.

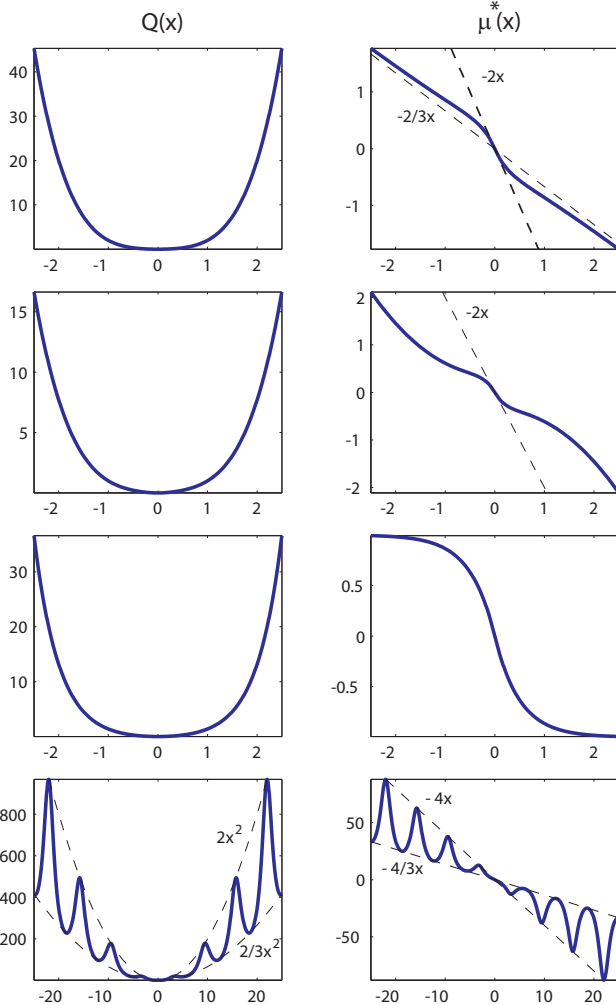


Fig. 1. Examples of superlinear costs. The cost functions and the optimal drift functions are shown in the same order as in Table 1. $\sigma^2 = 1$

$Q(x)$	$\frac{V_x(x)}{2\sigma^2}$	$\sigma^2 \mu^*(x)$
$x^2 + x^4$	$\frac{1}{\frac{1}{x} + \tan^{-1}(x) - \frac{\pi}{2}}$	$-\frac{(x+x^3)}{(1+x(\tan^{-1}(x) - \frac{\pi}{2}))^{-1}}$
$\frac{(x+x^3)^2}{1+3x^2}$	$x + x^3$	$-\frac{x+x^3}{1+3x^2}$
$\sinh(x)^2$	$\frac{\tanh(x)}{1 - \tanh(x)}$	$-\frac{1}{2}(1 - \exp(-2x))$
$\frac{(x + \sin(x)/2)^2}{1 + \cos(x)/2}$	$x + \sin(x)/2$	$-\frac{x + \sin(x)/2}{1 + \cos(x)/2}$

Table 1. Examples shown in Fig 1




Fig. 2. Saturating cost example in section 7.1: Plot of $Q(x)$ (43) and $V_x(x)$ derivative (45) for $a = 20, 60, 200, 10000$. Note that $V_x(x)$ is multiplied by $\sqrt{3a}$.

7.1 Example: Saturating cost

Consider the following cost

$$Q(x) = \frac{x^2}{(a+1) - a \left(\frac{\sqrt{3x/\sqrt{a}}}{\sinh(\sqrt{3x/\sqrt{a}})} \right)^2} \quad a \geq 0, \quad (43)$$

which is in fact composed of $Q_1(x) = x^2$ and $Q_2(x) = \sinh^2(x)$ as in (39). In the limit $a \rightarrow \infty$, this becomes a saturating function

$$Q(x) \approx x^2/(1+x^2). \quad (44)$$

According to theorem 8, the cost-to-go derivative is

$$V_x(x) = \frac{2\sigma^2 x}{(a+1) - a \left(\frac{\sqrt{3x/\sqrt{a}}}{\tanh(\sqrt{3x/\sqrt{a}})} - \frac{\sqrt{3x}}{\sqrt{a}} \right)}. \quad (45)$$

In the limit $a \rightarrow \infty$, this becomes

$$V_x(x) \approx 2\sigma^2 x / (1 + \sqrt{3ax}) \approx 0. \quad (46)$$

Notice that as a grows large, $V_x(x)$ becomes a step function with amplitude $1/\sqrt{3a}$. This is shown in Fig 2. Also note that although $Q(x)$ remains unchanged in the region $-2 \leq x \leq 2$, $V_x(x)$ in the same region goes through a drastic change as a increases. This is because the optimal solution depends not only on local features but on global shape of state-cost.

8. ALMOST-LINEAR COST

Previous sections showed that superlinear state-costs provide sufficient self-regularization, but sublinear (and lin-

ear) state-costs do not. The difference originates from the convergence property of the integration (14).

However, there exists a small class of cost-functions for which such distinction is impossible.

Definition: A cost $Q(x)$ is *almost-linear*, if $\exists \alpha, \beta > 0$ s.t.

$$\alpha x < Q(x) < \beta x^{1+\epsilon} \quad \forall x, \epsilon > 0 \quad (47)$$

(after appropriate shifting and zero-padding).

8.1 Examples

Consider two *almost-linear* cost-functions

$$Q_1(x) = (1+x) \cdot \log(1+x) \quad (48)$$

$$Q_2(x) = (1+x) \cdot \log^2(1+x). \quad (49)$$

Note that integration of $1/Q_1$ diverges

$$\int_x^\infty \frac{dx'}{Q_1(x')} = \left[\log(\log(1+x)) \right]_x^\infty = \infty. \quad (50)$$

Despite the subtle difference, however, integration of $1/Q_2$ converges and returns the following optimal solution,

$$\begin{aligned} V_x(x) &= 2\sigma^2 \log(1+x), \\ \mu(x) &= -\frac{1}{\sigma^2} (1+x) \cdot \log(1+x). \end{aligned} \quad (51)$$

where we assumed $c = 0$, yet without a formal proof.

Therefore, convergence property of *almost-linear* cost-functions can only be confirmed on case-by-case basis.

9. COST WITH MULTIPLE GLOBAL MINIMA

In this section, we extend our results to state-costs with multiple global minima. Let us begin with an example.

Consider the state-cost

$$Q(x) = (x^2 - 1)^2,$$

which is zero at $x = \pm 1$. Integrating $1/Q$ returns the following cost-to-go derivative

$$V_x(x) = \frac{8\sigma^2}{\frac{2x}{x^2-1} + \log\left(\frac{x-1}{x+1}\right)}, \quad (52)$$

which is a well-behaved real function for $|x| > 1$, but becomes complex for $-1 \leq x \leq 1$. In fact, this results follows from setting $c = 0$, which lacks justification for $-1 \leq x \leq 1$. A new principle is needed for state-space between global minima.

Lemma 10. If a state-cost has multiple global minima x_1, \dots, x_n s.t. $Q(x_i) = 0$, then

$$V(x) = 0 \quad x_1 \leq x \leq x_n \quad (53)$$

Proof. (1) $V(x_1) = V(x_n) = 0$, (2) $V(x) \geq 0, \forall x$, and (3) $V(x)$ must be convex everywhere. \square

This lemma implies absence of self-regulation between global minima. In this case, optimal drift diverges to infinity in the direction of the nearby global minimum.

10. UNIQUENESS OF NON-SMOOTH SOLUTIONS

All the examples shown up to now have smooth, differentiable cost-to-go solutions. However, solutions can be non-smooth for some cost functions.

10.1 Example

Consider the cost-function

$$Q(x) = x^2 + \sqrt{|x|}. \tag{54}$$

The corresponding cost-to-go derivative is

$$V_x(x) = \frac{3}{\log\left(\frac{x-\sqrt{x+1}}{x+2\sqrt{x+1}}\right) + 2\sqrt{3}\left(\pi/2 - \tan^{-1}\left(\frac{2\sqrt{x-1}}{\sqrt{3}}\right)\right)} \tag{55}$$

for $x > 0$. Notice that the cost-to-go derivative does not become zero in the limit $x \rightarrow 0^+$, but instead $V_x(x) \rightarrow \sqrt{27}/4\pi$. Since $V_x(x)$ is anti-symmetric in x , this solution is non-smooth at $x = 0$.

In general, one must be cautious when admitting non-smooth functions as potential solutions of a HJB equation, because this relaxed notion may introduce an infinite number of possible solutions. The problem of choosing a unique non-smooth solution is usually solved by introducing *viscosity condition*. For the control problems here considered, it can be shown that with all the conditions mentioned in section 4, uniqueness of solution is indeed guaranteed by allowing non-differentiability only at $x = 0$, and that these are indeed *viscosity solutions*.

11. GENERALIZATION TO FINITE β USING PERTURBATION THEORY

In this section, we extend the analysis to the cases when control-cost is non-zero. However, the analysis becomes very difficult when β is finite, and in fact impossible for most forms of control-cost.

Here we focus on quadratic control-cost. Consider

$$\mathcal{L}(x, \mu) = Q(x) + \frac{1}{2}(\tau\mu)^2. \tag{56}$$

for which the HJB equation is

$$0 = Q(x) + \min_{\mu} \left[\mu V_x(x) + \frac{1}{2}(\tau^2 + \sigma^2 V_{xx}(x)) \mu^2 \right]. \tag{57}$$

After explicit minimization, this simplifies to

$$\frac{1}{2Q} = \frac{\tau^2}{V_x^2} - \sigma^2 \frac{d}{dx} \left(\frac{1}{V_x} \right). \tag{58}$$

where optimal drift is

$$\mu^*(x) = -\frac{V_x(x)}{\tau^2 + \sigma^2 V_{xx}(x)} = -\frac{2Q(x)}{V_x(x)}. \tag{59}$$

Equation (58) is non-linear, and therefore it does not allow closed form solutions in general. If β is either very small or very large, however, the solution can be approximated as a perturbation series.

11.1 Perturbation Series I

Here we expand the solution near the deterministic limit. Equation (58) can be expressed as

$$\frac{1}{2Q} = U^2 - \beta \frac{dU}{dx}, \tag{60}$$

where $U \equiv \tau/V_x$ and $\beta \equiv \sigma^2/\tau$. We express U as a power series of β :

$$U(x) = U_0(x) + \beta U_1(x) + \beta^2 U_2(x) + \dots,$$

where U_0 represents the deterministic limit solution, and other U_k terms describe higher-order correction. When this is substituted in (60), we obtain

$$0 = \left(U_0^2 - \frac{1}{2Q} \right) + \beta \left(2U_0 U_1 - U_0' \right) + \beta^2 \left(U_1^2 + 2U_0 U_2 - U_1' \right) + \dots$$

Each term in parenthesis must vanish respectively, which leads to a series differential equations that can be solved iteratively. The solution is summarized in Table 2.

Note that the higher order terms introduce non-local features of state-cost, since U_i depends on the i^{th} order differentiation of $Q(x)$. Therefore, as β gets larger, the optimal solution will depend more on global features of state-cost.

11.2 Perturbation Series II

The solution can also be expanded near the zero-control-cost limit. We express (58) as

$$\frac{1}{2Q} = \alpha W^2 - \frac{dW}{dx}, \tag{61}$$

where $W \equiv \sigma^2/V_x$ and $\alpha \equiv 1/\beta^2 = \tau^2/\sigma^4$. and look for the solution of the form

$$W(x) = W_0(x) + \alpha W_1(x) + \alpha^2 W_2(x) + \dots$$

When this is substituted in (61), we find

$$0 = (W_0' + \frac{1}{2Q}) + \alpha(W_1' - W_0^2) + \alpha^2(W_2' - 2W_0 W_1) + \dots$$

which gives out a series of 1st-order differential equations that can be easily integrated iteratively (see Table 2).

Note that in this expansion, it is the lowest order term (W_0) that introduces the global-feature dependence of the solution, while the higher order terms gradually cancel out the global dependence. This is apparent in (61), since the αW^2 term acts as a non-linear filter that decays over distance.

Note that $c = 0$ is assumed for each integration, which remains to be proved.

Expansion I :	Expansion II :
$U(x) = \sum_{k=0}^{\infty} \left(\frac{\sigma^2}{\tau} \right)^k U_k(x)$	$W(x) = \sum_{k=0}^{\infty} \left(\frac{\tau}{\sigma^2} \right)^{2k} W_k(x)$
$\mu^*(x) = -2Q(x)U(x)/\tau$	$\mu^*(x) = -2Q(x)W(x)/\sigma^2$
$U_0 = \frac{1}{\sqrt{2Q}}$	$W_0(x) = \int_x^{\infty} \frac{1}{2Q(x')} dx'$
$U_1 = \frac{U_0'}{2U_0}$	$W_1(x) = \int_x^{\infty} -W_0^2(x') dx'$
$U_2 = \frac{U_1' - U_1^2}{2U_0}$	$W_2(x) = \int_x^{\infty} -2W_0 W_1 dx'$
$U_3 = \frac{U_2' - 2U_1 U_2}{2U_0}$	$W_3(x) = \int_x^{\infty} -(W_1^2 + 2W_0 W_2)$
$U_{k+1} = \frac{U_k' - \sum_{i,j} U_i U_j}{2U_0}$	$W_{k+1}(x) = \int_x^{\infty} -\sum_{i,j} W_i W_j dx'$
$1 \leq i, j \leq k, i + j = k + 1$	where $0 \leq i, j \leq k, i + j = k$

Table 2. Summary of Perturbation Series

Example: Quadratic State-Cost .

For the quadratic cost function,

$$\mathcal{L}(x, \mu) = x^2 + \frac{1}{2}(\tau\mu)^2,$$

the expansion terms are $W_0 = \frac{1}{2x}$, $W_1 = -\frac{1}{4x}$, $W_2 = \frac{1}{4x}$, $W_3 = -\frac{5}{16x}$, \dots , or in general

$$W_k = \frac{(-1)^k (2k-1)!!}{2x (k+1)!}.$$

This perturbation series converges to a closed form

$$\begin{aligned} W(x) &= \frac{1}{2x} \left(1 + \sum_{k=1}^{\infty} \frac{(2k-1)!!}{(k+1)!} (-\alpha)^k \right) \\ &= \frac{\sqrt{1+2\alpha}-1}{2\alpha x} \end{aligned} \quad (62)$$

which is indeed the solution of (61). Optimal drift is

$$\mu^*(x) = -\frac{2x}{\sigma^2 + \sqrt{\sigma^4 + 2\tau^2}}. \quad (63)$$

This solution is valid for all range of β .

12. CONCLUSION

This paper considered optimum control of systems with multiplicative noise. Our analysis is vastly different from the traditional linear-quadratic (LQ) analysis: Rather than analyzing a high dimensional system with a simple quadratic cost function, we analyzed a one dimensional system with complex, arbitrary cost functions. As a result, our analysis revealed exotic, even counter-intuitive, relationships between cost functions and the corresponding optimal feedback-control strategies.

Most importantly, we showed that multiplicative noise makes optimal solutions depend on the global features of cost functions, and that the extent of such global dependence is determined by β — the ratio of time constants. Other interesting discoveries include (1) feedback gain must decrease as the state-cost becomes steeper, and (2) optimality of risk-aversion vs risk-seeking strategies depends on superlinearity of state-cost and not convexity. Traditional LQ analysis cannot find these intriguing facts because the quadratic cost function is too simple.

Previously, it was suggested that a generalized, iterative-LQG (linear-quadratic-gaussian) method could approximately solve non-LQG problems with multiplicative noise (Todorov [2005]). However, it was not clear when such methods would fail. As mentioned above, the determining parameter is β . These methods would eventually fail as β gets sufficiently large, since LQG approximation only collects at most second order local information.

We also showed that analysis becomes simple in the large β limit. This limit occurs when the cost of control is negligible compared to the need for system stabilization. The iterative expansion method (expansion II) provides high accuracy solutions in such situations.

Unfortunately, the analytic approach presented here is limited to one-dimensional problems. However, many of the implications revealed by this analysis are expected to generalize to higher-dimensional problems and they could provide insights when solving practical problems.

REFERENCES

- V.V. Bolotin. *Random vibrations of elastic systems*. Springer, 1984. ISBN 9024729815.
- S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear matrix inequalities in system and control theory*. Society for Industrial Mathematics, 1994. ISBN 0898714850.
- L. El Ghaoui. State-feedback control of systems with multiplicative noise via linear matrix inequalities. *Syst. Control Lett.*, 24:223–228, February 1995. ISSN 0167-6911. doi: [http://dx.doi.org/10.1016/0167-6911\(94\)00045-W](http://dx.doi.org/10.1016/0167-6911(94)00045-W).
- C.M. Harris and D.M. Wolpert. Signal-dependent noise determines motor planning. *Nature*, 394(6695):780–784, 1998. ISSN 0028-0836.
- R.A. Ibrahim. *Parametric random vibration*. John Wiley & Sons, London, 1985.
- C. Kubrusly and O. Costa. Mean square stability conditions for discrete stochastic bilinear systems. *Automatic Control, IEEE Transactions on*, 30(11):1082–1087, 1985. ISSN 0018-9286.
- J. Lu and R.E. Skelton. Robust variance control for systems with finite-signal-to-noise uncertainty. *Automatica*, 36(4):511–525, 2000. ISSN 0005-1098.
- P. McLane. Optimal stochastic control of linear systems with state-and control-dependent disturbances. *Automatic Control, IEEE Transactions on*, 16(6):793–798, 1971. ISSN 0018-9286.
- Y. Phillis. Controller design of systems with multiplicative noise. *Automatic Control, IEEE Transactions on*, 30(10):1017–1019, 1985. ISSN 0018-9286.
- M. Ruan and A.K. Choudhury. Steady-state optimal controller with actuator noise variance linearly related to actuator signal variance. *Automatic Control, IEEE Transactions on*, 38(1):133–135, 1993. ISSN 0018-9286.
- E. Todorov. Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural computation*, 17(5):1084–1108, 2005. ISSN 0899-7667.
- A. van Wingerden and W. De Koning. The influence of finite word length on digital optimal control. *Automatic Control, IEEE Transactions on*, 29(5):385–391, 1984. ISSN 0018-9286.
- T.J. Wagenaar. Stability and stabilizability of chemical reactors modelled with stochastic parameters. *International Journal of Control*, 49(1):33–44, 1989. ISSN 0020-7179.
- J.L. Willems and J.C. Willems. Feedback stabilizability for stochastic systems with state and control dependent noise. *Automatica*, 12(3):277–283, 1976. ISSN 0005-1098.
- W.M. Wonham. Optimal stationary control of a linear systems with state-dependent noise. *SIAM Journal of Control*, 5:486–500, 1967.
- E. Yaz and R.E. Skelton. Parametrization of all linear compensators for discrete-time stochastic parameter systems. *Automatica*, 30(6):945–955, 1994. ISSN 0005-1098.