Linear Controllers are Uniformly Optimal for the Witsenhausen Counterexample

Michael Rotkowitz[1,2]

IEEE Conference on Decision and Control, 2006

Abstract

In 1968, Witsenhausen introduced his celebrated counterexample, which illustrated that when an information pattern is nonclassical, the controllers which optimize an expected quadratic cost may be nonlinear. It is shown here that for the Witsenhausen Counterexample, when one instead considers the induced norm, then linear controllers are in fact optimal.

1 Introduction

A classical information pattern assumes that, at every time step, the controller can access not only information from that time, but from all preceding times as well. When this holds, and when the dynamics are linear, the cost quadratic, and the noise Gaussian, optimal controllers are linear.

The Witsenhausen Counterexample [6] showed that when a nonclassical information pattern exists, then nonlinear controllers may be optimal for the LQG norm.

In this paper, we will consider the same problem setup, but for the optimum of the norm induced by the 2-norm, sometimes called uniformly optimal control.

In the case of the L_1-norm, which is induced by a different norm, it was shown that [1] linear controllers are optimal over those which are differentiable in the origin, by a simple argument which would extend to most induced norms including ours. However, it was also shown that when one drops this differentiability assumption, and can optimize the controller for each possible direction of the input noise, then such a nonlinear controller may outperform all linear controllers [5], even for a centralized information pattern.

In the case of uniformly optimal control, it has been shown that linear controllers are optimal for the centralized case [2, 3], just as in LQG control. This paper shows that linear controllers are uniformly optimal as well for a problem with a nonclassical information pattern, and in particular, for the same problem which elucidated the possibility of nonlinear optimality for the LQG problem.

[1]Research School of Information Sciences and Engineering, The Australian National University, Canberra, Australia, (e-mail: rotkowitz@stanfordalumni.org)

[2]This work was primarily completed at the Royal Institute of Technology, Stockholm, Sweden, where it was supported by the European Commission through the Integrated Project RUNES and by the Swedish Strategic Research Foundation through an INGVAR grant.

2 Preliminaries

We introduce the problem setup, several basic definitions which will be used throughout the paper, and their immediate consequences.

2.1 Problem Formulation

Given noise

$$w = \begin{bmatrix} w_1 \\ w_2 \end{bmatrix}$$

and control laws

$$u_1 = \gamma_1(y_1) \qquad u_2 = \gamma_2(y_2)$$

the system then evolves as follows, as indicated in Figure 1

$$\begin{aligned} x_0 &= \sigma w_1 & v &= w_2 \\ y_1 &= x_0 & y_2 &= x_1 + v \\ u_1 &= \gamma_1(y_1) & u_2 &= \gamma_2(y_2) \\ x_1 &= x_0 + u_1 & x_2 &= x_1 - u_2 \end{aligned}$$

and we wish to keep the following variables small

$$\begin{aligned} z_1 &= \sqrt{k} u_1 \\ z_2 &= x_2 \end{aligned}$$

We can then put all of this together to obtain

$$z = \begin{bmatrix} \sqrt{k}\ \gamma_1(\sigma w_1) \\ \sigma w_1 + \gamma_1(\sigma w_1) - \gamma_2(\sigma w_1 + \gamma_1(\sigma w_1) + w_2) \end{bmatrix}$$

In the original problem [6], the noise was normally distributed

$$w \sim \mathcal{N}(0, I)$$

and we seeked γ_1, γ_2 to minimize

$$\mathbb{E}\ \|z\|_2$$

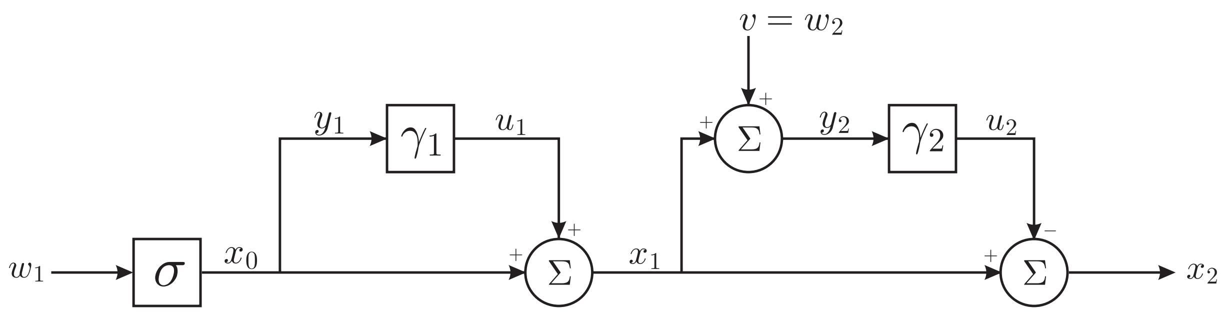

Figure 1: The Witsenhausen Counterexample

In this paper, we instead consider the induced norm, and seek γ_1, γ_2 to minimize

$$\sup_{w \neq 0} \frac{\|z\|_2}{\|w\|_2}$$

In order for the 2-norm to be consistent with Witsenhausen's objective function, the weighting on u_1 needs to be $\sqrt{k}$. Note when considering the references, however, that some, such as [4], use a weighting of k.

2.2 Polar Coordinates

We will consider polar coordinates

$$\begin{aligned} w_1 &= r\cos\theta \\ w_2 &= r\sin\theta \end{aligned}$$

such that the signals passed to the controllers become

$$\begin{aligned} y_1 &= \sigma r\cos\theta \\ y_2 &= \sigma r\cos\theta + \gamma_1(\sigma r\cos\theta) + r\sin\theta \end{aligned}$$

The gain for given control laws and a given input can be defined as

$$\begin{aligned} Q^2(\gamma_1,\gamma_2,r,\theta) &= \frac{\|z\|_2^2}{\|w\|_2^2} \\ &= \left(\frac{\sqrt{k}\ \gamma_1(\sigma r\cos\theta)}{r}\right)^2 \\ &\quad + \left(\sigma\cos\theta + \frac{\gamma_1(\sigma r\cos\theta)}{r}\right. \\ &\quad \left. - \frac{\gamma_2(\sigma r\cos\theta + \gamma_1(\sigma r\cos\theta) + r\sin\theta)}{r}\right)^2 \end{aligned}$$

We then define the cost for given control laws as

$$J(\gamma_1,\gamma_2) = \sup_{r>0}\ \sup_{\theta}\ Q(\gamma_1,\gamma_2,r,\theta)$$

and the optimal cost as

$$J^* = \inf_{\gamma_1}\ \inf_{\gamma_2}\ J(\gamma_1,\gamma_2)$$

Unless we state otherwise, a sup over θ is assumed to be over $|\theta| \leq \pi$.

We state a fairly obvious lemma and omit the proof.

Lemma 1. *$J(\gamma_1,\gamma_2)$ is finite only if $\gamma_1(0) = \gamma_2(0) = 0$.*

2.3 Definition of Piecewise Affine Control Laws

For any $a_-, a_+, b_-, b_+ \in \mathbb{R}$, we define piecewise affine control laws

$$\gamma_1(y_1) = \begin{cases} a_- y_1 & \text{if } y_1 \leq 0 \\ a_+ y_1 & \text{if } y_1 \geq 0 \end{cases}$$

$$\gamma_2(y_2) = \begin{cases} b_- y_2 & \text{if } y_2 \leq 0 \\ b_+ y_2 & \text{if } y_2 \geq 0 \end{cases}$$

If we let

$$a(\theta) = \begin{cases} a_- & \text{if } y_1 < 0 \\ a_+ & \text{if } y_1 \geq 0 \end{cases} \qquad b(\theta) = \begin{cases} b_- & \text{if } y_2 < 0 \\ b_+ & \text{if } y_2 \geq 0 \end{cases}$$

such that

$$\begin{aligned} \gamma_1(y_1(r,\theta)) &= a(\theta)y_1(r,\theta) \\ \gamma_2(y_2(r,\theta)) &= b(\theta)y_2(r,\theta) \end{aligned}$$

the gain for given control laws of this form and a given input can then be defined as

$$
\begin{aligned}
&Q^2_{pa}(a_-, a_+, b_-, b_+, r, \theta) \\
&\quad = \left(\frac{\sqrt{k}\ a(\theta)(\sigma r \cos\theta)}{r} \right)^2 \\
&\quad\quad + \left(\sigma \cos\theta + \frac{a(\theta)(\sigma r \cos\theta)}{r} \right. \\
&\quad\quad \left. - \frac{b(\theta)(\sigma r \cos\theta + a(\theta)(\sigma r \cos\theta) + r \sin\theta)}{r} \right)^2 \\
&\quad = \left(\sqrt{k}\ a(\theta) \sigma \cos\theta \right)^2 \\
&\quad\quad + (\sigma \cos\theta (1 + a(\theta))(1 - b(\theta)) + \sin\theta)^2
\end{aligned} \tag{1}
$$

Noting that this is independent of r, we hereafter refer to the gain as $Q_{pa}(a_-, a_+, b_-, b_+, \theta)$, and then define the cost for given control laws as

$$J_{pa}(a_-, a_+, b_-, b_+) = \sup_{\theta} Q_{pa}(a_-, a_+, b_-, b_+, \theta)$$

and the optimal cost as

$$J^*_{pa} = \inf_{a_-, a_+} \inf_{b_-, b_+} J_{pa}(a_-, a_+, b_-, b_+)$$

2.4 Definition of Linear Control Laws

We similarly define the subclass of linear control laws and their costs. For any $a, b \in \mathbb{R}$, we define

$$
\begin{aligned}
\gamma_1(y_1) &= a y_1 \\
\gamma_2(y_2) &= b y_2
\end{aligned}
$$

The gain for given control laws of this form and a given input can then be defined as

$$
\begin{aligned}
Q^2_l(a, b, \theta) &= \left(\sqrt{k}\ a \sigma \cos\theta \right)^2 \\
&\quad + (\sigma \cos\theta (1 + a)(1 - b) + \sin\theta)^2
\end{aligned} \tag{2}
$$

the cost for given control laws as

$$J_l(a, b) = \sup_{\theta} Q_l(a, b, \theta)$$

and the optimal cost as

$$J^*_l = \inf_a \inf_b J_l(a, b)$$

3 Linear Dominates Piecewise Affine

In this section, we show that the best piecewise affine control law is actually linear, and also find the optimal gains.

3.1 Optimal Values for Second Controller

Let θ^* be defined as the point at which $y_1 > 0$ and $y_2 = 0$. Thus

$$\cos(\theta^*) > 0 \quad \text{and} \quad \sigma(1 + a_+)\cos(\theta^*) + \sin(\theta^*) = 0$$

or

$$\theta^* = \arctan(-\sigma(1 + a_+)) \cap (-\pi/2, \pi/2)$$

We now focus on the half-circle $|\theta| \leq \pi/2$, from here until Section 3.4. The value of Q_{pa} on this half-circle is unaffected by a_-, and we thus refer to the gain as $Q_{pa}(a, b_-, b_+, \theta)$. We then define the worst-case cost on the half-circle for given control laws as

$$J_{hc}(a, b_-, b_+) = \sup_{|\theta| \leq \pi/2} Q_{pa}(a, b_-, b_+, \theta)$$

and the optimal such cost as

$$J^*_{hc} = \inf_a \inf_{b_-, b_+} J_{hc}(a, b_-, b_+)$$

At θ^*, the gain Q_{pa} is unaffected by our choice of b_-, b_+, and so it represents a lower bound on our worst case cost J_{hc}, an idea which is formalized at the end of this subsection. Therefore, if we can choose b_-, b_+ such that $Q_{pa}(\theta^*)$ is the maximum over all θ, then it would be an upper bound as well, and we will have found the optimal values of b_-, b_+.

This is only possible if the left and right derivatives with respect to θ are both non-positive. We can achieve $\frac{\partial}{\partial\theta} Q_{pa}(a, b_-, b_+, \theta)\big|_{\theta=\theta^*} = 0$ if $b_- = b_+ = b^*(a)$ where

$$b^*(a) = \frac{\sigma^2(ka^2 + (1 + a)^2)}{1 + \sigma^2(1 + a)^2} \tag{3}$$

For a given value of a, it remains to be checked whether this value of b leads to a peak or a trough at θ^*. Due to the periodicity of Q_{pa} we only need to show that the second derivative at this point is negative, and we will have the peak over all angles.

Since the value at θ^* is unaffected by our choice of b_-, b_+, we have

$$
\begin{aligned}
&Q_{pa}(a, b_-, b_+, \theta^*(a)) \\
&\quad\quad = Q_{pa}(a, b^*(a), b^*(a), \theta^*(a)) \quad \forall\ b_-, b_+
\end{aligned}
$$

which implies

$$J_{hc}(a, b_-, b_+) \geq Q_{pa}(a, b^*(a), b^*(a), \theta^*(a)) \quad \forall\ b_-, b_+$$

which then implies

$$\inf_{b_-, b_+} J_{hc}(a, b_-, b_+) \geq Q_{pa}(a, b^*(a), b^*(a), \theta^*(a)) \tag{4}$$

which further implies

$$J^*_{hc} \geq \inf_a Q_{pa}(a, b^*(a), b^*(a), \theta^*(a)) \tag{5}$$

3.2 Optimal Value for First Controller

In this subsection, for each value of $a \in \mathbb{R}$ we assume that $\theta = \theta^*(a)$ and that $b_- = b_+ = b^*(a)$ as defined above. We then seek the value of a which minimizes $Q_{pa}(a, b^*(a), b^*(a), \theta^*(a))$, the lower bound from (4), We define this lower bound as $Q_{lb}(a)$, and plugging (3) into (1) we find that

$$Q_{lb}^2(a) = \frac{\sigma^2(ka^2 + (1+a)^2)}{1 + \sigma^2(1+a)^2}$$

This function is plotted in Figure 2, along with its asymptotes.

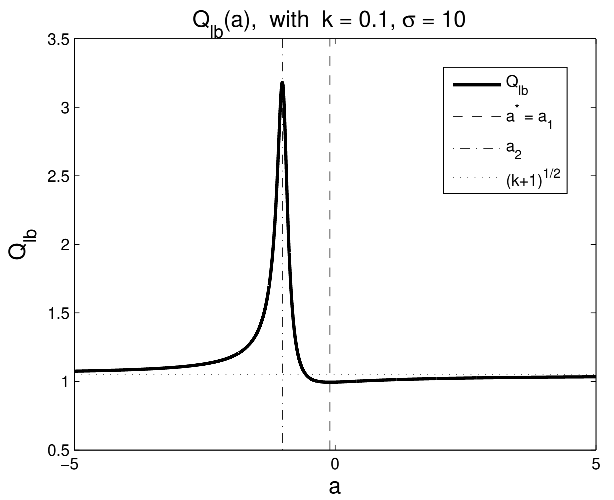

Figure 2: $Q_{lb}(a) = Q_{pa}(a, b^*(a), b^*(a), \theta^*(a))$

Letting $\alpha = k\sigma^2 + k + 1$ and $\beta = 2\sigma\sqrt{k}$, we find that $\frac{\partial}{\partial a} Q_{lb}^2(a) = 0$ iff $a \in \{a_1, a_2\}$ where

$$a_1 = 2(-\alpha + \sqrt{\alpha^2 - \beta^2})/\beta^2$$
$$a_2 = 2(-\alpha - \sqrt{\alpha^2 - \beta^2})/\beta^2$$

These points are also shown in Figure 2, and we see that the minimum occurs at a_1. We now prove that this is always the case.

Noting that

$$\alpha^2 - \beta^2 = (k\sigma^2 - 1)^2 + (2\sigma^2 + 1)k^2 + 2k > 0$$

and thus, that $a_1, a_2 \in \mathbb{R}$ and $a_1, a_2 < 0$, we can express the second derivative at these points as

$$\left.\frac{\partial^2 Q_{lb}^2(a)}{\partial a^2}\right|_{a \in \{a_1, a_2\}} = \frac{2\sigma^2(k\sigma^2 a^2 - 1)}{a(1 + \sigma^2(1+a)^2)^2}$$

Then,

$$k\sigma^2 a_1^2 - 1 = \left(\frac{2\sqrt{\alpha^2 - \beta^2}}{\beta^2}\right)\left(-\alpha + \sqrt{\alpha^2 - \beta^2}\right) < 0 \tag{6}$$

and

$$k\sigma^2 a_2^2 - 1 = \left(\frac{2\sqrt{\alpha^2 - \beta^2}}{\beta^2}\right)\left(\alpha + \sqrt{\alpha^2 - \beta^2}\right) > 0$$

which leads to

$$\frac{k\sigma^2 a_1^2 - 1}{a_1} = \sqrt{\alpha^2 - \beta^2} > 0$$

and

$$\frac{k\sigma^2 a_2^2 - 1}{a_2} = -\sqrt{\alpha^2 - \beta^2} < 0$$

Hence the second derivative is positive only for a_1, and $Q_{lb}(a)$ must achieve its infemum either at a_1 or at $\pm\infty$.

Letting $\gamma = k\sigma^2 - k - 1$,

$$Q_{lb}^2(a_1) = \frac{(2k+1)\alpha^2 - (k+1)\beta^2 + \gamma^2 + 2(\gamma - k\alpha)\sqrt{\alpha^2 - \beta^2}}{\alpha^2 + (k-1)\beta^2 + \gamma^2 + 2\alpha\sqrt{\alpha^2 - \beta^2}}$$

As $a \to \pm\infty$, $Q_{lb}^2(a) \to k + 1$.

$$(k+1) - Q_{lb}^2(a_1) = \frac{k\beta^2\sqrt{\alpha^2 - \beta^2}}{\alpha^2 + (k-1)\beta^2 + \gamma^2 + 2\alpha\sqrt{\alpha^2 - \beta^2}} \; > \; 0$$

and thus

$$\inf_a \; Q_{lb}(a) \; = \; Q_{lb}(a_1)$$

We hereafter refer to $a^* = a_1$. We have now shown

$$\inf_a Q_{pa}(a, b^*(a), b^*(a), \theta^*(a)) = Q_{pa}(a^*, b^*(a^*), b^*(a^*), \theta^*(a^*)) \tag{7}$$

and combining (5), (7), we have

$$J_{hc}^* \;\geq\; Q_{pa}(a^*, b^*(a^*), b^*(a^*), \theta^*(a^*)) \tag{8}$$

3.3 Optimal Cost on Half-Circle

We showed that for any a, choosing $b_- = b_+ = b^*(a)$ would yield a lower bound on the induced norm at $\theta^*(a)$, which would be equal to the induced norm if the second derivative was negative.

Now that we have the value of a^* for which this bound is minimized, we need only check the second derivative for that value.

$$\left.\frac{\partial^2 Q_{pa}^2(a, b^*, b^*, \theta)}{\partial \theta^2}\right|_{\theta=\theta^*} = \frac{2\sigma^2(ka^2 + (1+a)^2)}{1 + \sigma^2(1+a)^2}(k\sigma^2 a^2 - 1)$$

We've already seen in (6) that

$$k\sigma^2(a^*)^2 - 1 < 0$$

and thus,

$$\left.\frac{\partial^2 Q_{pa}^2(a, b^*, b^*, \theta)}{\partial\theta^2}\right|_{\theta=\theta^*} < 0$$

The first derivative was already forced to zero by our choice of b^*, and thus

$$J_{hc}(a^*, b^*(a^*), b^*(a^*)) \;=\; Q_{pa}(a^*, b^*(a^*), b^*(a^*), \theta^*(a^*))$$

This tells us that

$$\inf_{b_-,b_+} J_{hc}(a^*, b_-, b_+) \;\leq\; Q_{pa}(a^*, b^*(a^*), b^*(a^*), \theta^*(a^*)) \tag{9}$$

and combining (4), (9), we have

$$\inf_{b_-,b_+} J_{hc}(a^*, b_-, b_+) \;=\; Q_{pa}(a^*, b^*(a^*), b^*(a^*), \theta^*(a^*))$$

which then implies

$$J_{hc}^* \;\leq\; Q_{pa}(a^*, b^*(a^*), b^*(a^*), \theta^*(a^*)) \tag{10}$$

Lastly, combining (8), (10), we get

$$J_{hc}^* \;=\; Q_{pa}(a^*, b^*(a^*), b^*(a^*), \theta^*(a^*)) \tag{11}$$

3.4 Optimal Cost on Full Circle

In this subsection, we show that this optimum over the half-circle is also the optimum over the whole circle, and thus, that piecewise affine offers no advantage over linear.

$$\begin{aligned} J_{pa}^* \;&\geq\; J_{hc}^* \qquad (12)\\ &=\; Q_{pa}(a^*, a^*, b^*(a^*), b^*(a^*), \theta^*(a^*)) \end{aligned}$$

where the inequality holds because the supremum is taken over a subset in the latter expression, and the equality follows from (11).

Since $\cos(\theta+\pi) = -\cos\theta$ and $\sin(\theta+\pi) = -\sin\theta$ it follows from (1) that

$$Q_{pa}(a, a, b, b, \theta + n\pi) = Q_{pa}(a, a, b, b, \theta) \quad \forall\ n \in \mathbb{Z}$$

and thus, when the control laws are linear the cost may be determined on the half-circle, and so

$$\begin{aligned} &J_{pa}(a^*, a^*, b^*(a^*), b^*(a^*)) \\ &\quad= \sup_{\theta}\ Q_{pa}(a^*, a^*, b^*(a^*), b^*(a^*), \theta) \\ &\quad= \sup_{|\theta|\leq\frac{\pi}{2}}\ Q_{pa}(a^*, a^*, b^*(a^*), b^*(a^*), \theta) \\ &\quad= J_{hc}(a^*, a^*, b^*(a^*), b^*(a^*)) \\ &\quad= Q_{pa}(a^*, a^*, b^*(a^*), b^*(a^*), \theta^*(a^*)) \end{aligned}$$

which implies

$$J_{pa}^* \;\leq\; Q_{pa}(a^*, a^*, b^*(a^*), b^*(a^*), \theta^*(a^*)) \tag{13}$$

and thus, combining (12), (13)

$$J_{pa}^* \;=\; Q_{pa}(a^*, a^*, b^*(a^*), b^*(a^*), \theta^*(a^*))$$

In other words, the infemum is achieved where $a_- = a_+$ and $b_- = b_+$, and the optimal piecewise affine controller is linear.

The optimal linear controller must then result in the same optimal cost

$$J_l^* \;=\; Q_l(a^*, b^*(a^*), \theta^*(a^*)) \tag{14}$$

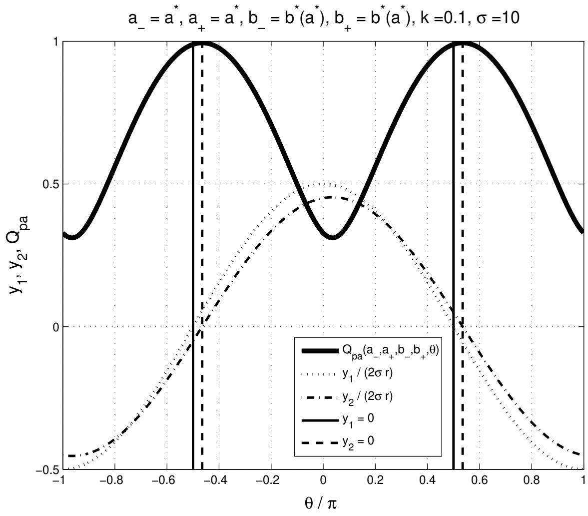

Figure 3: $y_1, y_2, Q_{pa}(a^*, a^*, b^*(a), b^*(a), \theta)$

The value of the gain with these optimal values is shown along the circle in Figure 3, and the signals passed to the controllers at each angle are shown as well. Note that the worst case cost indeed occurs where $y_2 = 0$, that is, at θ^*.

4 Linear Dominates All Nonlinear

If we had assumed that γ_1, γ_2 were merely left- and right-differentiable in the origin, then we would be done, since the induced norm would always be lower bounded by arbitrarily small inputs, and thus only piecewise affine controllers of the form considered in the previous section would have to be considered at all. But we need not make any such assumption, as shown in this section.

The following theorem states that the optimal linear controller achieves the same performance as the optimal overall controller.

Theorem 2.

$$J^* = J_l^*$$

Proof. Since the best controller is at least as good as the best linear controller, we clearly have

$$J^* \le J_l^* \tag{15}$$

Given any $\gamma_1, \gamma_2 : \mathbb{R} \to \mathbb{R}$, let

$$a = \gamma_1(1)$$

and then let

$$\begin{aligned} \theta_0 &= \arctan(-\sigma(1+a)) \\ &= \theta^*(a) \end{aligned}$$

and then choose

$$r_0 = \sec(\theta_0) \; / \; \sigma$$

such that

$$\begin{aligned} y_1(r_0, \theta_0) &= \sigma r_0 \cos(\theta_0) \\ &= 1 \end{aligned}$$

and

$$y_2(r_0, \theta_0) = 0$$

and thus, using this, Lemma 1, (2), and (14), respectively,

$$\begin{aligned} &Q^2(\gamma_1, \gamma_2, r_0, \theta_0)) \\ &= \left(\sqrt{k}\sigma a \cos(\theta_0)\right)^2 + \left(\sigma \cos\theta + \sqrt{k}\sigma a \cos(\theta_0) - \frac{\gamma_2(0)}{r_0}\right)^2 \\ &= \left(\sqrt{k}\sigma a \cos(\theta_0)\right)^2 + \left(\sigma \cos\theta + \sqrt{k}\sigma a \cos(\theta_0)\right)^2 \\ &= Q_l^2(a, b^*(a), \theta^*(a)) \\ &= J_l^2(a, b^*(a)) \end{aligned}$$

which yields

$$J(\gamma_1, \gamma_2) \ge J_l(a, b^*(a))$$

In other words, for arbitrary γ_1, γ_2 we can find a linear controller which performs at least as well, and thus

$$J^* \ge J_l^* \tag{16}$$

and the desired result follows from (15), (16). ∎

5 Discussion and Conclusions

We have shown that given the same setup as the Witsenhausen Counterexample, if our objective is to optimize the induced norm, then linear controllers are optimal. This was shown to hold true for arbitrary values of the constants. We showed in Section 3 that the best piecewise affine controller is linear, and specifically calculated those control laws. Then in Section 4, we showed that for any control laws, one can find linear controllers which perform just as well, and thus, that linear controllers are optimal.

This approach was partially redundant. It was not necessary to first show that the optimal controller is optimal amongst all piecewise affine controllers, since it is later shown that a linear controller dominates any given control law. However, it is not yet clear which approach will be most useful to generalize these results to other information structures, and so Section 3 is presented in its entirety.

This result raises the question of which, if not all, other information structures admit uniformly optimal controllers which are linear.

Acknowledgments

The author would like to thank C. Langbort for his excitement upon hearing this problem suggested, and thus for the motivation to revisit it and solve it.

References

[1] M. A. Dahleh and J. S. Shamma. Rejection of persistent bounded disturbances: Nonlinear controllers. *Systems and Control Letters*, 18:245–252, 1992.

[2] A. Feintuch and B. A. Francis. Uniformly optimal control of linear feedback systems. *Automatica*, 21(5):563–574, 1985.

[3] P. P. Khargonekar and K. R. Poola. Uniformly optimal control of linear time-invariant plants: Nonlinear time-varying controllers. *Systems and Control Letters*, 6:303–308, 1986.

[4] S. Mitter and A. Sahai. Information and control: Witsenhausen revisited. *Lecture Notes in Control and Information Sciences*, 241:281–293, 1999.

[5] A. A. Stoorvogel. Nonlinear l_1 optimal controllers for linear systems. *IEEE Transactions on Automatic Control*, 40(4):694–696, 1995.

[6] H. S. Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal of Control*, 6(1):131–147, 1968.