

MIT Open Access Articles

*Slopey quantizers are locally optimal
for Witsenhausen's counterexample*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Ajorlou, Amir, and Jadbabaie, Ali. "Slopy Quantizers Are Locally Optimal for Witsenhausen's Counterexample." 2016 IEEE 55th Conference on Decision and Control (CDC), December 12-14 2016, Las Vegas, Nevada, USA, Institute of Electrical and Electronics Engineers (IEEE), December 2016 © 2016 Institute of Electrical and Electronics Engineers (IEEE)

As Published: <http://dx.doi.org/10.1109/CDC.2016.7798805>

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <http://hdl.handle.net/1721.1/111985>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Slopy Quantizers Are Locally Optimal for Witsenhausen’s Counterexample

Amir Ajorlou[†] and Ali Jadbabaie[†]

Abstract—We study the perfect Bayesian equilibria of a leader-follower game of incomplete information. The follower makes a noisy observation of the leader’s action (who moves first) and chooses an action minimizing her expected deviation from the leader’s action. Knowing this, leader who observes the realization of the state, chooses an action that minimizes her distance to the state of the world and the ex-ante expected deviation from the follower’s action. We show the existence of what we call “near piecewise-linear equilibria” when there is strong complementarity between the leader and the follower and the precision of the prior is poor. As a major consequence of this result, we prove local optimality of a class of *slopy* quantization strategies which had been suspected of being the optimal solution in the past, based on numerical evidence for Witsenhausen’s counterexample.

Index Terms—Decentralized control, optimal stochastic control, incomplete information games, perfect Bayesian equilibrium.

I. INTRODUCTION

In his seminal work [1], Witsenhausen constructed a simple two-stage linear-quadratic-Gaussian (LQG) decentralized control problem where the optimal controller happens to be nonlinear. This example showed for the first time that linear quadratic Gaussian team problems can have nonlinear solutions. By resorting to Witsenhausen’s counterexample, [2] produced an example showing that the standard decentralized static output optimal control problem of linear deterministic systems could also admit optimal nonlinear solutions. For nearly half a century, this counterexample has been a subject of intense research across multiple communities ([3]–[8]).

The endogenous information structure of Witsenhausen’s counterexample, where the signal observed in the second stage is a noisy version of the control action in the first stage, gives rise to a nonclassical information structure. While the problem looks deceptively simple and quadratic on first look, it is actually a very complicated, nonconvex, functional optimization problem. This counterexample has shed light on intricacies of optimal decisions in stochastic team optimization problems with similar information structure. Naturally, this problem has given rise to a large body of literature. For example, [9] provides a variant of Witsenhausen’s counterexample with discrete primitive random variables and finite support, where no optimal solution exists. Another interesting variant, with the same information structure but different cost function, is the Gaussian test channel ([4], [10]) where the linear strategies can be shown to be optimal. Also, [11] shows that if the objective function is

changed to a worst case induced norm, the linear controllers dominate nonlinear policies. While [1] proves the existence of an optimal solution using tools from real and functional analysis, other works such as ([6], [8]) suggest *lifting* the problem to an equivalent optimization problem over the space of probability measures and then employing tools from the optimal transport theory [12].

Although the optimal strategy and optimal cost for Witsenhausen’s counterexample are still unknown, it can be shown that carefully designed nonlinear strategies can largely outperform the linear strategies (see, e.g., the multi-point quantization strategies proposed by [5]). This result, in particular, implies the fragility of the comparative statics and policies solely derived based on the linear strategies in problems with similar setting. A relevant line of research is to provide error bounds on the proximity to optimality for approximate solutions. [13]–[15] use information theoretic techniques and vector versions of the original problem to provide such bounds. There are also several works aiming to approximate the optimal solution. [16]–[19] employ different heuristic approaches, all confirming an almost piecewise-linear form for the optimal controller. However, a complete optimality proof for such strategies has been elusive.

In this paper, we view Witsenhausen’s problem as a leader-follower coordination game in which the action of the leader is corrupted by an additive noise, before reaching the follower. The leader aims to coordinate with the follower while staying close to the observed state, recognizing that her action is not observed perfectly. As a result, she needs to signal the follower in a manner that can be decoded efficiently. More than a mere academic counterexample, the above setup could model a scenario where coordination happens across generations and the insights of the leader who is from a different generation is corrupted/lost by the time the message reaches the future generations. If the leader can internalize the fact that her actions will not be observed perfectly, how should she act to make sure coordination occur? When the leader cares far more about coordination with the follower than staying “on the message”, the near piecewise-linear equilibrium strategy of the leader coarsens the observation in well-spaced intervals, rather than merely broadcasting a linearly scaled version of the observed state as the linear strategy would suggest. In this line, we prove the existence of a superior equilibrium strategy which coarsens the message via a “slopy quantizer”, making both the leader and follower better off compared to the linear equilibrium strategies.

To this end, we analyze the perfect Bayesian equilibria of this game and show that strong complementarity between

[†]Institute for Data, Systems, and Society, Massachusetts Institute of Technology (MIT), Cambridge, MA 02139, USA. E-mail: {ajorlou, jadbabai}@mit.edu. This work was supported by ARO MURI W911NF-12-1-0509.

the leader and the follower combined with a prior with poor enough precision can give rise to nonlinear equilibria, and in particular, equilibria in form of slopey quantizers. To the best of our knowledge, this work is the first providing an analytical rationale for the optimality of slopey quantization strategies for Witsenhausen’s counterexample.

The main idea behind the proof is to carefully construct a class of what we call *near piecewise-linear* strategies for the leader that stays invariant under the best response operator. By a near piecewise-linear strategy, we mean a piecewise strategy where the changes in the derivative in each segment are very small, making the strategy in each segment almost linear. Each strategy has a fixed number of segments, with leader’s action changing very slowly within each segment. As a result, leader’s actions stay very close to fixed points of the strategy in each segment. These fixed points are the values of the state which the leader does not distort. Therefore, well-spaced fixed points (combined with some appropriate bounds on the relative prior of the state of the world in different segments) reveal the leader’s actions to the follower with high probability, making the “signal” easily decodable. As a consequence, we can characterize the best response of the follower to leader’s strategy. Using this characterization, we show that the best response of the leader to follower’s strategy also varies very little, essentially remaining near piecewise-linear as well.

A key challenge in deriving the invariance property for this set of strategies for the leader is to bound and control the displacement in the fixed points and endpoints of the segments of leader’s strategy under the best response. A key observation here is that the fixed points of the leader’s best responses are *local minimizers* of the expected deviation of the leader’s action from the follower. This insight allows us to show that the fixed points of the leader’s best response lie in a tight neighborhood of the fixed points of the follower’s strategy. We then show that the fixed points of the follower’s strategy in turn lie in a vicinity of a convex combination of the leader’s fixed points and the expected value of the state of the world within each segment. Combining the two, we can derive an approximate dynamics for the displacement in the fixed points and endpoints of the segments in leader’s strategy under the best response. Using this approximate dynamics, we then characterize an invariant set of fixed points and interval endpoints for leader’s strategy, which we can then use in order to prove the existence of a near piecewise-linear equilibrium strategy for the leader.

II. MODEL

The game consists of a leader L and a follower F . Before the agents act, the state of the world θ is drawn from a normal distribution with zero mean and variance σ^2 . The leader can observe the realization of θ and acts first. The payoff of the leader is given as

$$u_L = -r_L(\theta - a_L)^2 - (1 - r_L)(a_F - a_L)^2, \quad (1)$$

where a_F is the action of the follower and $0 < r_L < 1$. The follower makes a private, noisy observation of the leader’s

action, $s = a_L + \delta$ where $\delta \sim N(0, 1)$. The payoff of the follower is

$$u_F = -(a_L - a_F)^2. \quad (2)$$

We consider the perfect Bayesian equilibria of the game and show that they reduce to the Bayes Nash equilibria due to the Gaussian noise in the observation. Denote with $a_L^*(\theta)$ and $a_F^*(s)$ the equilibrium strategies, and with $\nu^*(\cdot|s)$ the follower’s belief about leader’s action given s . Due to the normal noise in the observation, $\nu^*(\cdot|s)$ is fully determined by $a_L^*(\theta)$ and the prior as there are no off-equilibrium-path information sets. Equilibrium strategies should thus satisfy

$$\begin{aligned} a_F^*(s) &= \mathbb{E}_{\nu^*}[a_L^*|s] = \int_{-\infty}^{\infty} a_L \nu^*(a_L|s) da_L, \\ a_L^*(\theta) &= \operatorname{argmax}_{a_L} -r_L(\theta - a_L)^2 \\ &\quad - (1 - r_L) \int_{-\infty}^{\infty} (a_F^*(s) - a_L)^2 \phi(s - a_L) ds, \end{aligned} \quad (3)$$

where $\phi(\cdot)$ denotes the PDF of the standard normal distribution. We can easily characterize the linear equilibria of the game, following [1].

Theorem 1: Linear Bayes strategies of the leader and follower are of the form $a_F^*(s) = \mu s$ and $a_L^*(\theta) = \lambda \theta$, where $\mu = \frac{t^2}{1+t^2}$ and $\lambda = \frac{t}{\sigma}$, and t is a real root of the equation

$$\frac{r_L}{1 - r_L}(\sigma - t) = \frac{t}{(1 + t^2)^2}. \quad (4)$$

Our main objective in this paper is to show the existence of an equilibrium with a near piecewise-linear strategy for the leader in the regime where there is strong complementarity between the leader and the follower (when r_L is small) and the prior’s precision is poor (or large σ). To this end, and motivated mainly by [1], we focus on regime $\frac{1}{\sigma} \leq r_L \sigma^2 \leq 1$, and aim to prove the existence of such an equilibrium for sufficiently large values of σ (and hence small r_L).

III. NONLINEAR EQUILIBRIA

Our approach for proving the existence of an equilibrium with a near piecewise-linear strategy for the leader is to identify a set of such strategies for the leader that is invariant under the best response operator. We characterize such a set in the next section.

A. An Invariant Set of Near Piecewise-Linear Strategies for the Leader

Given $m \in \mathbb{N}$, consider a partition of the normal distribution $N(0, \sigma^2)$ into $2m + 1$ segments $\cup_{k=-m}^m B_k^0$, with $B_k^0 = [b_k^0, b_{k+1}^0]$ for $k \in \mathbb{N}_m$, $B_0^0 = (b_{-1}^0, b_1^0)$, and $B_{-k}^0 = (b_{-k-1}^0, b_{-k}^0]$, with $b_{-k}^0 = -b_k^0$ and $b_{m+1}^0 = -b_{-m-1}^0 = +\infty$. Denote with c_k^0 the expected value of $\theta \sim N(0, \sigma^2)$ in segment B_k^0 , that is, $c_k^0 = \mathbb{E}_{N(0, \sigma^2)}[\theta | \theta \in B_k^0]$. Clearly, $c_0^0 = 0$ and $c_{-k}^0 = -c_k^0$ for $k \in \mathbb{N}_m$.

We are in particular interested in a partition where the interval endpoints b_k^0 are the midpoints of $[c_{k-1}^0, c_k^0]$, i.e.,

¹Proofs are not included in this manuscript due to space limitations. See [20] for the proofs.

$b_k^0 = \frac{c_{k-1}^0 + c_k^0}{2}$ for $k \in \mathbb{N}_m$. We can show that such a partition exists and is unique. Next, we construct a set of $(2m+1)$ -segmented increasing odd functions, denoted by $A_L^m(r_L, \sigma)$ satisfying the following properties.

Property 1: For every $a_L(\theta) \in A_L^m(r_L, \sigma)$, there exist $2m+1$ segments $B_k = [b_k, b_{k+1}]$, for $k \in \mathbb{N}_m$, $B_0 = (-b_1, b_1)$, and $B_{-k} = (b_{-k-1}, b_{-k}]$, with $b_{m+1} = -b_{-m-1} = +\infty$ such that:

- $a_L(\theta)$ is increasing and odd (i.e., $a_L(-\theta) = -a_L(\theta)$), and is smooth over each interval.
- $a_L(\theta)$ has a unique fixed point in each segment. That is, for each interval B_k , $(-m \leq k \leq m)$, there exists a unique $c_k \in B_k$ such that $a_L(c_k) = c_k$, with $c_0 = 0$.

We also impose a constraint on the slope of $a_L(\theta)$ in each interval, keeping the slope very close to r_L , as well as a linear bound on $a_L(\theta)$ in the tail. More precisely, we impose the following property:

Property 2: For every $-m < k < m$ and $\theta \in B_k$, $r \leq \frac{d}{d\theta} a_L(\theta) \leq \bar{r}$, where $r = r_L(1 - 0.5r_L^2\sigma^2)$ and $\bar{r} = r_L(1 + 2.5r_L^2\sigma^2)$. For the tail interval B_m , $r \leq \frac{d}{d\theta} a_L(\theta) \leq \bar{r}$ for $b_m < \theta < c_m + \sigma^2$. For $\theta > c_m + \sigma^2$ we have $a_L(\theta) \leq c_m + 5r_L(\theta - c_m + \sigma)^2$.

The larger σ , the closer \bar{r} and r to r_L . For instance, choosing $\sigma \geq 16$ ensures $\bar{r} < 1.01r_L$ and $r > 0.998r_L$. Finally, we impose the constraint that interval endpoints b_k remain close to midpoints of $[c_{k-1}, c_k]$ and that fixed points c_k remain close to c_k^0 's.

Property 3: For every $k \in \mathbb{N}_m$,

$$|b_k - \frac{c_{k-1} + c_k}{2}| \leq 0.1r_L\sigma. \quad (5)$$

Moreover,

$$|c_k - c_k^0| \leq \frac{2k(2m-k)\zeta}{r_L}, \quad (6)$$

where $\zeta = 0.44r_L^2\sigma^2 + (2m+3.4)r_L^2\sigma$.

Our proof strategy is then to show that for any $m \in \mathbb{N}$, there exists $\sigma_m > 0$ such that, for every $\sigma \geq \sigma_m$ in regime $\frac{1}{\sigma} \leq r_L\sigma^2 \leq 1$, the set of strategies $A_L^m(r_L, \sigma)$ characterized by Property 1-3 is invariant under the best response operator.

B. Best Response Analysis

The first step in verifying the invariance of $A_L^m(r_L, \sigma)$ is to characterize the best response of the follower $a_F(s)$ to the leader's strategy $a_L(\theta) \in A_L^m(r_L, \sigma)$. We can then use these properties to find the updated best response of the leader to $a_F(s)$, denoted by $\tilde{a}(\theta)$ and enforce its inclusion in $A_L^m(r_L, \sigma)$. Choosing σ sufficiently large, we can ensure several useful properties for $A_L^m(r_L, \sigma)$ that can facilitate this process. We state these properties in the lemma below.

Lemma 1: There exists $\sigma_m \geq \max(16, 7m^2)$ such that for any $\sigma \geq \sigma_m$ in regime $\frac{1}{\sigma} \leq r_L\sigma^2 \leq 1$ we have the following:

- The lengths of the half intervals in any strategy $a_L(\theta) \in A_L^m(r_L, \sigma)$ are upper-bounded by σ . That is, $|c_{k-1} - b_k| \leq \sigma$ and $|c_k - b_k| \leq \sigma$ for $k \in \mathbb{N}_m$.

²We only state the properties (and in many cases the analysis) only for $\theta \geq 0$. The counterpart for $\theta \leq 0$ is immediate since the function is odd.

- There exists a $C > 3 \ln \sigma$ such that the lengths of the half intervals in any strategy $a_L(\theta) \in A_L^m(r_L, \sigma)$ are lower-bounded by C . That is, $|c_{k-1} - b_k| \geq C$ and $|c_k - b_k| \geq C$ for $k \in \mathbb{N}_m$. Moreover, $c_m \geq \sigma$ and $b_m \geq 0.5\sigma$.
- The lengths of the first half (second half) of the intervals in any strategy $a_L(\theta) \in A_L^m(r_L, \sigma)$ are in increasing order. That is, $|c_k - b_k| < |c_{k+1} - b_{k+1}|$ and $|c_{k-1} - b_k| < |c_k - b_{k+1}|$, for $k \in \mathbb{N}_{m-1}$. Moreover, $b_2 - b_1 > 2b_1$ ensuring an increasing order on the lengths of the intervals as well.
- Let e_k , $0 \leq k \leq m$, be the expected value of θ in segment B_k , i.e., $e_k = \mathbb{E}_{N(0, \sigma^2)}[\theta | \theta \in B_k]$. Then, the distances from e_k to the endpoints b_k and b_{k+1} also satisfies the lower and upper bounds C and σ . Moreover, $e_m \geq 0.5\sigma$.

By choosing $\sigma \geq \sigma_m$, we can exploit the properties stated for $A_L^m(r_L, \sigma)$ in the above lemma. In what follows, we use these properties to prove the invariance of the set of strategies $A_L^m(r_L, \sigma)$ under the best response. The follower's best response to the strategy of the leader $a_L(\theta)$ is the expected action of the leader given the observation $s = a_L + \delta$ and can be written as

$$a_F(s) = \mathbb{E}_\delta[a_L | s] = \frac{\int_{-\infty}^{\infty} a_L(\theta) \phi(s - a_L(\theta)) \phi(\frac{\theta}{\sigma}) d\theta}{\int_{-\infty}^{\infty} \phi(s - a_L(\theta)) \phi(\frac{\theta}{\sigma}) d\theta}. \quad (7)$$

Using this, we can easily show that $a_F(s)$ is analytic and increasing, with $\frac{d}{ds} a_F(s) = \text{Var}[a_L | s]$ (see [1] for a proof).

In order to characterize $a_F(s)$, we start by estimating the expected action of the leader and its variance conditioned on the interval to which θ belongs. Actions of the leader in interval B_k ($k \neq \pm m$) are well-concentrated around c_k . In fact $|a_L(\theta) - c_k| \leq \bar{r}\sigma$ for $\theta \in B_k$, from which the lemma below immediately follows.

Lemma 2: For $0 \leq k < m$, $|\mathbb{E}[a_L(\theta) | s, \theta \in B_k] - c_k| \leq \bar{r}\sigma$ and $\text{Var}[a_L(\theta) | s, \theta \in B_k] \leq \bar{r}^2\sigma^2$.

The analysis is a bit involved in the tail, since for $\theta > c_m$ the leader's actions are not in a bounded vicinity of c_m anymore. However, we can derive several useful properties for the tail as well.

Lemma 3: Consider a tail observation by the leader (i.e., $\theta \in B_m$). Then,

$$\mathbb{E}[a_L(\theta) | s, \theta \in B_m] - c_m \leq \bar{r}\sigma, \quad (8)$$

for $s \leq c_m + \sigma$, and

$$\mathbb{E}[a_L(\theta) | s, \theta \in B_m] - c_m \leq 5r_L\sigma(s - c_m + 1 + \frac{0.01}{\sigma^2}), \quad (9)$$

for $s > c_m + \sigma$. Also, $\mathbb{E}[a_L(\theta) | s, \theta \in B_m] - c_m \geq -\bar{r}\sigma$. As for the variance,

$$\text{Var}[a_L(\theta) | s, \theta \in B_m] \leq \begin{cases} \frac{1}{3}, & \text{for } s < c_{m-1} \\ 0.64\bar{r}^2\sigma^2, & \text{for } c_{m-1} \leq s \leq c_m + \sigma \\ 8.8r_L^2\sigma^2(s - c_m)^2, & \text{for } s > c_m + \sigma. \end{cases} \quad (10)$$

Let the signal observed by the follower be between c_k and c_{k+1} , i.e., $s = c_k + \delta$ with $0 \leq \delta \leq c_{k+1} - c_k$. Then, we claim

that the follower's posterior on θ given s has a negligible probability out of the neighboring intervals $B_k \cup B_{k+1}$. We first derive the following property for the relative prior of θ in B_k 's by relating it to e_k 's and using the bounds on the distance from e_k to the endpoints of the intervals, as well as the increasing order of the lengths of the intervals given by Lemma 1.

Lemma 4: For any $-m \leq k_1, k_2 \leq m$, we have

$$\frac{\text{Prob}[\theta \in B_{k_1}]}{\text{Prob}[\theta \in B_{k_2}]} \leq e^{\frac{(c_{k_2} - c_{k_1})^2}{64}}. \quad (11)$$

This lemma implies that the posterior on θ is more affected by the likelihoods rather than relative priors. Using this, we can bound the posterior of θ outside the neighboring intervals to s (i.e., out of $B_k \cup B_{k+1}$).

Lemma 5: Let the observed signal by the follower be $s = c_k + \delta$, where $0 \leq \delta \leq c_{k+1} - c_k$. Then, for any $r \geq 1$,

$$\frac{\text{Prob}[\theta \in B_{k-r}|s]}{\text{Prob}[\theta \in B_k|s]} \leq e^{-\frac{245(c_k - c_{k-r})^2}{512}}. \quad (12)$$

Similarly,

$$\frac{\text{Prob}[\theta \in B_{k+r+1}|s]}{\text{Prob}[\theta \in B_{k+1}|s]} \leq e^{-\frac{245(c_{k+r+1} - c_{k+1})^2}{512}}. \quad (13)$$

Using this lemma and the fact that the fixed points c_k are well-spaced, we can show that the effect of the intervals other than B_k and B_{k+1} on $a_F(s)$ are negligible. In order to characterize the follower's best response $a_F(s)$, we then need to focus only on the segments adjacent to the observed signal, and in particular figure out the weight of each of these two neighboring intervals in the follower's posterior on θ . We do this in the following lemma.

Lemma 6: Define

$$m_{k+1} = \frac{c_k + c_{k+1}}{2} + \frac{1}{\Delta_{k+1}} \ln \left(\frac{\text{Prob}[\theta \in B_k]}{\text{Prob}[\theta \in B_{k+1}]} \right), \quad (14)$$

where $\Delta_{k+1} = c_{k+1} - c_k$. Also, write the signal observed by the follower as $s = m_{k+1} + \delta$. Then, for $0 \leq k < m-1$,

$$e^{-\Delta_{k+1}(\delta + \bar{r}\sigma) - \frac{\bar{r}^2 \sigma^2}{2}} \leq \frac{\text{Prob}[\theta \in B_k|s]}{\text{Prob}[\theta \in B_{k+1}|s]} \leq e^{-\Delta_{k+1}(\delta - \bar{r}\sigma) + \frac{\bar{r}^2 \sigma^2}{2}}. \quad (15)$$

For the case involving the tail segment B_m ,

$$e^{-\Delta_m(\delta + \bar{r}\sigma) - \frac{\bar{r}^2 \sigma^2}{2}} \leq \frac{\text{Prob}[\theta \in B_{m-1}|s]}{\text{Prob}[\theta \in B_m|s]} \leq 1.168 e^{-\Delta_m(\delta - \bar{r}\sigma) + \frac{\bar{r}^2 \sigma^2}{2}}. \quad (16)$$

It is worth mentioning that m_{k+1} defined in the above lemma is quite close to the midpoint of c_k and c_{k+1} . In fact, it follows from Lemma 4 that $|m_{k+1} - \frac{c_k + c_{k+1}}{2}| < \frac{\Delta_{k+1}}{64}$. We can now characterize the best response of the follower $a_F(s)$ to the leader's strategy $a_L(\theta) \in A_L^m(r_L, \sigma)$ up to the first order.

Lemma 7: Let $s = m_{k+1} + \delta$, with $c_k \leq s \leq c_{k+1}$. Then

$$\begin{aligned} a_F(s) &\geq c_k + \frac{\Delta_{k+1}}{1 + 1.168 e^{-\Delta_{k+1}(\delta - \bar{r}\sigma) + \frac{\bar{r}^2 \sigma^2}{2}}} - 1.01 \bar{r} \sigma \\ a_F(s) &\leq c_k + \frac{\Delta_{k+1}}{1 + e^{-\Delta_{k+1}(\delta + \bar{r}\sigma) - \frac{\bar{r}^2 \sigma^2}{2}}} + 1.01 \bar{r} \sigma. \end{aligned} \quad (17)$$

Also,

$$0 \leq \frac{d}{ds} a_F(s) \leq 1.17 e^{-\Delta_{k+1}(|\delta| - \bar{r}\sigma)} (\Delta_{k+1} + 2\bar{r}\sigma)^2 + 1.01 \bar{r}^2 \sigma^2. \quad (18)$$

Corollary 1: A useful consequence of Lemma 7 is that

$$\begin{aligned} a_F(s) &\geq c_{k+1} - 1.17 e^{-\Delta_{k+1}(\delta - \bar{r}\sigma)} \Delta_{k+1} - 1.01 \bar{r} \sigma \\ a_F(s) &\leq c_k + 1.17 e^{\Delta_{k+1}(\delta + \bar{r}\sigma)} \Delta_{k+1} + 1.01 \bar{r} \sigma, \end{aligned} \quad (19)$$

where $s = m_{k+1} + \delta$, with $c_k \leq s \leq c_{k+1}$.

Note that the exponential terms in the above bounds vanish quite fast for large C and $|\delta|$. For small $|\delta|$, another useful upper bound on the derivative of $a_F(s)$ is

$$\frac{d}{ds} a_F(s) \leq \frac{1}{4} (\Delta_{k+1} + 2\bar{r}\sigma)^2 + 1.01 \bar{r}^2 \sigma^2. \quad (20)$$

Corollary 2: Let $s = m_{k+1} + \delta$, with $c_k \leq s \leq c_{k+1}$. Then,

$$\begin{aligned} c_k - 1.5 \bar{r} \sigma &\leq a_F(s) \leq c_k + 1.5 \bar{r} \sigma \quad \text{for } \delta < -0.65 \\ c_{k+1} - 1.5 \bar{r} \sigma &\leq a_F(s) \leq c_{k+1} + 1.5 \bar{r} \sigma \quad \text{for } \delta > 0.65. \end{aligned} \quad (21)$$

Roughly speaking, the above corollary says that, if the observed signal by the follower is far enough from the midpoint of c_k and c_{k+1} , then the optimal action of the follower is well-concentrated around c_k or c_{k+1} (whichever that is closer), and changes very slowly according to Lemma 7. However, $a_F(s)$ may have very high variations for s close to m_{k+1} as can be seen from Lemma 7.

The following lemma characterizes $a_F(s)$ when follower makes a tail observation.

Lemma 8: Let $s = c_m + \delta$, where $\delta > 0$. Then,

- i) for $\delta \leq \sigma$, $c_m - 1.01 \bar{r} \sigma \leq a_F(s) \leq c_m + \bar{r} \sigma$, and $0 \leq \frac{d}{ds} a_F(s) \leq 0.65 \bar{r}^2 \sigma^2$.
- ii) for $\delta > \sigma$, $c_m - 1.01 \bar{r} \sigma \leq a_F(s) \leq c_m + 5r_L \sigma (\delta + 1 + \frac{0.01}{\sigma^2})$, and $0 \leq \frac{d}{ds} a_F(s) \leq 9r_L^2 \sigma^2 \delta^2$.

Lemma 7 and 8 provide the first order characteristics of the best response of the follower to a leader's strategy $a_L(\theta) \in A_L^m(r_L, \sigma)$. We are now ready to analyze the leader's best response $\tilde{a}_L(\theta)$ to $a_F(s)$ and see if it stays in $A_L^m(r_L, \sigma)$. We have $\tilde{a}_L(\theta) = \text{argmax}_{a_L} \tilde{u}_L(\theta, a_L)$, where

$$\begin{aligned} \tilde{u}_L(\theta, a_L) &= -r_L(\theta - a_L)^2 \\ &\quad - (1 - r_L) \int_{-\infty}^{\infty} (a_F(s) - a_L)^2 \phi(s - a_L) ds. \end{aligned} \quad (22)$$

Lemma 9: Consider $\theta \in [c_k, c_{k+1}]$, $0 \leq k < m$. Then, there exists a unique $\tilde{b}_{k+1} \in [c_k, c_{k+1}]$ such that

$$\begin{aligned} |\tilde{a}_L(\theta) - c_k| &< 5\bar{r}\sigma \quad \text{for } \theta < \tilde{b}_{k+1}, \\ |\tilde{a}_L(\theta) - c_{k+1}| &< 5\bar{r}\sigma \quad \text{for } \theta \geq \tilde{b}_{k+1}. \end{aligned} \quad (23)$$

The points \tilde{b}_{k+1} determine the segments of the best response strategy $\tilde{a}_L(\theta)$. We can bound the derivative of $\tilde{a}_L(\theta)$ over these segments by incorporating Lemma 7 and Corollary 1 and 2 into the above bound.

Lemma 10: Consider $\theta \in [c_k, c_{k+1}]$, $0 \leq k < m$, with $\theta \neq \tilde{b}_{k+1}$. Then,

$$\begin{aligned} \frac{d}{d\theta} \tilde{a}_L(\theta) &\geq \frac{r_L}{r_L + (1 - r_L)(1 + 0.3\bar{r}^2\sigma^2)} \\ \frac{d}{d\theta} \tilde{a}_L(\theta) &\leq \frac{r_L}{r_L + (1 - r_L)(1 - 2.4\bar{r}^2\sigma^2)}. \end{aligned} \quad (24)$$

Using this lemma and the values $r = r_L(1 - 0.5r_L^2\sigma^2)$ and $\bar{r} = r_L(1 + 2.5r_L^2\sigma^2)$, we can easily verify that $r \leq \frac{d}{d\theta} \tilde{a}_L(\theta) \leq \bar{r}$. This means that Property 2 is preserved by the best response for $\theta \in [-c_m, c_m]$. We study the tail case later in Lemma 12. Next, we characterize the fixed points of the best response strategy $\tilde{a}_L(\theta)$.

Lemma 11: Define

$$\tilde{J}_L(a_L) = \int_{-\infty}^{\infty} (a_F(s) - a_L)^2 \phi(s - a_L) ds. \quad (25)$$

Then, $\tilde{J}_L(a_L)$ is strongly convex over $[c_k - 5\bar{r}\sigma, c_k + 5\bar{r}\sigma]$, with $\frac{d^2}{da_L^2} \tilde{J}_L(a_L) \geq 2(1 - 2.4\bar{r}^2\sigma^2)$. Let \tilde{c}_k be the unique solution of

$$\tilde{c}_k = \underset{a_L \in [c_k - 5\bar{r}\sigma, c_k + 5\bar{r}\sigma]}{\operatorname{argmin}} \tilde{J}_L(a_L). \quad (26)$$

Then, $\tilde{a}_L(\tilde{c}_k) = \tilde{c}_k$.

The above lemma implies that Property 1 is also preserved under the best response. Next lemma describes the tail properties of $\tilde{a}_L(\theta)$.

Lemma 12: If $\tilde{b}_m < \theta < \tilde{c}_m + \sigma^2$, then $r \leq \frac{d}{d\theta} \tilde{a}_L(\theta) \leq \bar{r}$. For $\theta > \tilde{c}_m + \sigma^2$, we have

$$\tilde{a}_L(\theta) \leq \tilde{c}_m + 5r_L(\theta + \sigma - \tilde{c}_m). \quad (27)$$

Now, in order to verify that the updated strategy $\tilde{a}_L(\theta)$ satisfies Property 3, we need to bound the displacements in the fixed points \tilde{c}_k and endpoints \tilde{b}_k .

Lemma 13: For the endpoints of the intervals corresponding to $\tilde{a}_L(\theta)$, we have

$$\left| \tilde{b}_{k+1} - \frac{\tilde{c}_k + \tilde{c}_{k+1}}{2} \right| \leq 0.1r_L\sigma. \quad (28)$$

Bounding the displacement in \tilde{c}_k has multiple steps: it involves relating the fixed point of the leader's best response $\tilde{a}_L(\theta)$ in interval \tilde{B}_k to the fixed point of $a_F(s)$ in B_k (i.e., s_k), followed by estimating s_k in terms of c_k and e_k (recall that $e_k = \mathbb{E}_{N(0, \sigma^2)}[\theta | \theta \in B_k]$, i.e., the expected value of θ over B_k). Finally we bound the displacement in e_k with the displacement of the interval endpoints using several properties of truncated normal distribution.

Lemma 14: Let s_k be the fixed point of $a_F(s)$ in the interval $[c_k - 5\bar{r}\sigma, c_k + 5\bar{r}\sigma]$, i.e., $a_F(s_k) = s_k$. Then,

$$|\tilde{c}_k - s_k| < 0.44r_L^2\sigma^2. \quad (29)$$

Lemma 15: s_k can be located based on c_k and e_k as

$$|s_k - (1 - r_L)c_k - r_L e_k| < (2m + 3.3)r_L^2\sigma. \quad (30)$$

Using Lemma 14 and 15, we can reach at

$$|\tilde{c}_k - (1 - r_L)c_k - r_L e_k| < 0.44r_L^2\sigma^2 + (2m + 3.3)r_L^2\sigma. \quad (31)$$

We can now use (31) and Lemma 13 to verify that Property 3 is also preserved by the best response, completing the proof

of the invariance of $A_L^m(r_L, \sigma)$ for $\sigma \geq \sigma_m$ given by Lemma 1.

Theorem 2: Consider the regime $\frac{1}{\sigma} \leq r_L\sigma^2 \leq 1$ and a given $m \in \mathbb{N}$. Then, there exists $\sigma_m > 0$ such that for any $\sigma \geq \sigma_m$, the set of $(2m+1)$ -segmented strategies $A_L^m(r_L, \sigma)$ for the leader, characterized by Property 1-3, is invariant under the best response. Moreover, the game described in Section II has an equilibrium for which $a_L^*(\theta, r_L, \sigma) \in A_L^m(r_L, \sigma)$.

Given $m \in \mathbb{N}$, the maximum deviation of fixed points c_k from their counterparts c_k^0 for strategies in $A_L^m(r_L, \sigma)$ is upper-bounded by $2m^2(0.44r_L\sigma^2 + (2m+3.4)r_L\sigma)$, according to (6). This upper bound does not grow unboundedly with σ in the regime $\frac{1}{\sigma} \leq r_L\sigma^2 \leq 1$. Furthermore, if $\sigma \rightarrow +\infty$ along a path where $r_L\sigma^2 \rightarrow 0$, this upper bound approaches zero, hence asymptotically identifying the equilibrium strategy.

Proposition 1: Let $a_L^0(\theta, r_L, \sigma)$ be the $(2m+1)$ -segmented piecewise-linear strategy with B_k^0 's as segments and c_k^0 's as fixed points and fixed common slope r_L , that is, $a_L^0(\theta, r_L, \sigma) = c_k^0 + r_L(\theta - c_k^0)$ for $\theta \in B_k^0$. Then,

$$\lim_{\substack{r_L\sigma^2 \rightarrow 0 \\ \sigma \rightarrow +\infty}} \mathbb{E}_{N(0, \sigma^2)}[|a_L^*(\theta, r_L, \sigma) - a_L^0(\theta, r_L, \sigma)|] \rightarrow 0. \quad (32)$$

IV. NUMERICAL EXAMPLES

In [1], Witsenhausen proves the existence of a nonlinear controller which outperforms the optimal linear strategy for a simple two-stage LQG decentralized control problem. The equivalent regime under our setup is $k^2\sigma^2 = 1$, where $k^2 = \frac{r_L}{1-r_L}$. Witsenhausen shows that for large enough values of σ the optimal controller in this regime is nonlinear. It has been a long-standing conjecture that the optimal controller is a near piecewise-linear controller ([14]). Theorem 2 proves the existence of an equilibrium for the game of Section II with a near piecewise-linear strategy for the leader for sufficiently large σ in the regime $\frac{1}{\sigma} \leq r_L\sigma^2 \leq 1$; This implies the existence of a slope quantized local optimum³ for Witsenhausen's counterexample in this regime (it is easy to verify that $k^2\sigma^2 = 1$ falls in this regime, since $k^2\sigma^2 = 1$ yields $\frac{1}{\sigma} < r_L\sigma^2 = 1 - r_L < 1$ for $\sigma > 2$). The aim of this section is to illustrate Theorem 2 for the special cases of $m = 1, 2$ by specifying explicit values for σ_m which guarantee the existence of 3-segmented and 5-segmented near piecewise-linear equilibria, respectively.

Proposition 2 (3-Segmented Equilibria): Suppose that $\frac{1}{\sigma} \leq r_L\sigma^2 \leq 1$ and $\sigma \geq 16$. Then, the game described in Section II has an equilibrium with $a_L^*(\theta, r_L, \sigma) \in A_L^1(r_L, \sigma)$; $a_L^*(\theta, r_L, \sigma)$ is a 3-segmented near piecewise-linear strategy possessing Property 1-3.

The strip containing this nonlinear strategy for the leader for the case $\sigma = 16$ and $r_L\sigma^2 = 1 - r_L$ is depicted in Figure 1 (the plot only shows the region $\theta \geq 0$; $a_L(-\theta) = -a_L(\theta)$).

³This is indeed stronger than a local optimum since, at an equilibrium, fixing one player's strategy the deviation in the other player's strategy does not have to be local.




Fig. 1. The strip containing the nonlinear equilibrium strategy for the leader for the case $\sigma = 16$ and $r_L\sigma^2 = 1 - r_L$.




Fig. 2. The strip containing the nonlinear equilibrium strategy for the leader for the case $\sigma = 60$ and $r_L\sigma^2 = 1 - r_L$.

Proposition 3 (5-Segmented Equilibria): Suppose that $\frac{1}{\sigma} \leq r_L\sigma^2 \leq 1$ and $\sigma \geq 60$.⁴ Then, the game described in Section II has an equilibrium with $a_L^*(\theta, r_L, \sigma) \in A_L^2(r_L, \sigma)$; $a_L^*(\theta, r_L, \sigma)$ is a 5-segmented near piecewise-linear strategy possessing Property 1-3.

The strip containing this nonlinear strategy for the leader for the case $\sigma = 60$ and $r_L\sigma^2 = 1 - r_L$ is depicted in Figure 2.

V. CONCLUSIONS

We studied Witsenhausen's counterexample in a leader-follower game setup where the follower makes noisy observations from the leader's action and aims to choose her action as close as possible to that of the leader. Leader who moves first and can see the realization of the state of the world chooses her action to minimize her ex-ante distance from the follower's action as well as the state of the world. We showed the existence of nonlinear perfect Bayesian equilibria in the regime where there is strong complementarity between

the leader and the follower when the prior has very poor precision. Our results provide the first analytical proof for the local optimality of near piecewise-linear controllers for the well-known Witsenhausen's counterexample, where the optimal controller is conjectured to be a slopey quantizer.

REFERENCES

- [1] H. S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM J. Control*, vol. 6, pp. 131–147, 1968.
- [2] T. Başar, "On the optimality of nonlinear designs in the control of linear decentralized systems," *IEEE Transactions on Automatic Control*, vol. 21, p. 797, 1976.
- [3] Y. Ho and T. Chang, "Another look at the nonclassical information structure problem," *IEEE Transactions on Automatic Control*, vol. 25, pp. 537–540, 1980.
- [4] R. Bansal and T. Başar, "Stochastic teams with nonclassical information revisited: When is an affine law optimal?," *IEEE Transactions on Automatic Control*, vol. 32, pp. 554–559, 1987.
- [5] S. Mitter and A. Sahai, "Information and control: Witsenhausen revisited," *Lecture Notes in Control and Information Sciences*, vol. 241, pp. 281–293, 1999.
- [6] Y. Wu and S. Verdú, "Witsenhausen's counterexample: A view from optimal transport theory," in *Proceedings of the 50th IEEE Conf. Decision Control (CDC)*, pp. 5732–5737, 2011.
- [7] P. Grover, A. B. Wagner, and A. Sahai, "Information embedding and the triple role of control," *IEEE Transactions On Information Theory*, vol. 61, pp. 1539–1549, 2015.
- [8] A. Gupta, S. Yüksel, T. Başar, and C. Langbort, "On the existence of optimal policies for a class of static and sequential dynamic teams," *SIAM J. Control Optim.*, vol. 53, pp. 1681–1712, 2015.
- [9] S. Yüksel and T. Başar, "Stochastic networked control systems: Stabilization and optimization under information constraints," *Birkhäuser*, 2013.
- [10] T. Başar, "Variations on the theme of the Witsenhausen counterexample," in *Proceedings of the 47th IEEE Conf. Decision Control (CDC)*, pp. 1614–1619, 2008.
- [11] M. Rotkowitz, "Linear controllers are uniformly optimal for the Witsenhausen counterexample," in *Proceedings of the 45th IEEE Conf. Decision Control (CDC)*, pp. 553–558, 2006.
- [12] C. Villani, *Optimal Transport: Old and New*. Springer-Verlag, 2009.
- [13] P. Grover and A. Sahai, "A vector version of Witsenhausen's counterexample: Towards convergence of control, communication and computation," in *Proceedings of the 47th IEEE Conf. Decision Control (CDC)*, pp. 1636–1640, 2008.
- [14] P. Grover and A. Sahai, "Vector Witsenhausen counterexample as assisted interference suppression," *Int. J. Syst., Control Commun (IJSCC) Spec. iss. Inf Process. Decision Making in Distributed Control Syst.*, vol. 2, pp. 197–237, 2010.
- [15] P. Grover, S. Y. Park, and A. Sahai, "Approximately optimal solutions to the finite-dimensional Witsenhausen counterexample," *IEEE Transactions on Automatic Control*, vol. 58, pp. 2189–2204, 2013.
- [16] N. Li, J. R. Marden, and J. S. Shamma, "Learning approaches to the Witsenhausen counterexample from a view of potential games," in *Proceedings of the 48th IEEE Conf. Decision Control (CDC)*, 2009.
- [17] M. Baglietto, T. Parisini, and R. Zoppoli, "Numerical solutions to the Witsenhausen counterexample by approximating networks," *IEEE Transactions on Automatic Control*, vol. 46, pp. 1471–1477, 2001.
- [18] J. T. Lee, E. Lau, and Y.-C. L. Ho, "The Witsenhausen counterexample: A hierarchical search approach for nonconvex optimization problems," *IEEE Transactions on Automatic Control*, vol. 46, pp. 382–397, 2001.
- [19] M. Mehmetoglu, E. Akyol, and K. Rose, "A deterministic annealing approach to Witsenhausen's counterexample," in *ISIT*, 2014.
- [20] A. Ajorlou and A. Jadbabaie, "Slopey quantizers are locally optimal for witsenhausen's counterexample," available at: <http://arxiv.org/abs/1604.03806>, 2016.

⁴It is to be noted that smaller values for σ_m may work here (and similarly in Proposition 2). These values are obtained using the bounds derived for generic m . Of course, given a specific value of m , these bounds can be tightened much further which may result in smaller values for σ_m .