A Class of Near-Optimal Local Minima for Witsenhausen's Problem

Amir Ajorlou and Ali Jadbabaie

Laboratory for Information and Decision Systems
Institute for Data, Systems, and Society
Massachusetts Institute of Technology

58th IEEE Conference on Decision and Control
December 2019

Motivation

Witsenhausen's counterexample:

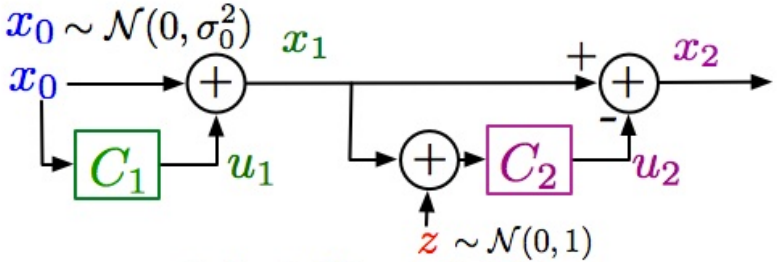

$\min\{k^2\mathbb{E}[u_1{}^2] + \mathbb{E}[x_2{}^2]\}$

- A simple two-stage LQG decentralized control problem.
- Optimal controller is nonlinear in regime $k^2\sigma_0^2 = 1$ for large σ_0.
- Asymmetry of information can lead to nonlinear solutions in multi-stage decentralized decision-making.
- Optimal controller still unknown, conjectured to be near piecewise-linear.

Motivation

A cherry-pick of related work:[1]

- Nonlinear designs largely outperform linear strategies (Mitter and Sahai (99)).
- Bounds on proximity to optimality using information theoretic techniques (Grover and Sahai (2013)).
- Heuristic techniques to approximate optimal solution (Li et al. (2009), Baglietto et al. (2001), Lee et al. (2001)).
- Little known on topological properties of optimal solution (continuous or not?, number of fixed points, etc.), or even local optima.

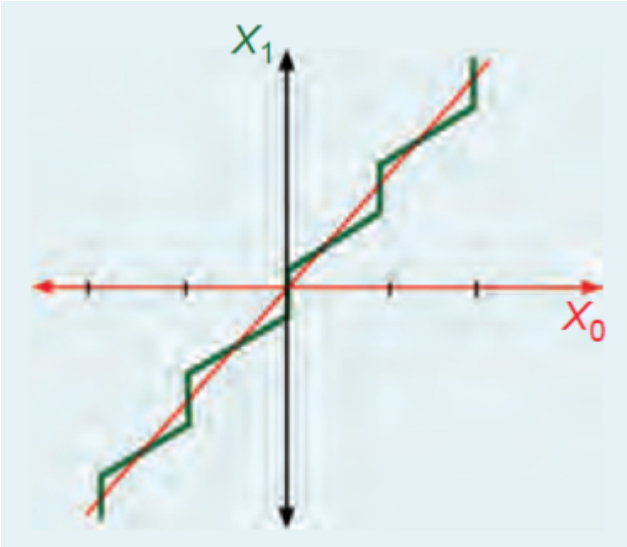

Slopey quantizer (Sahai (2010))

- Wu and Verdú (2011) using optimal transport theory: optimal controller is strictly increasing with a real analytic left inverse.
- This does not imply continuity of the optimal solution.

[1] There is plenty of elegant related work which could not be covered in this short talk.

Our Work

1. Pose the problem as a leader-follower coordination game:
   - $\theta \sim N(0, \sigma^2)$ observable to leader
   - Payoffs:

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

   - Follower observes $s = a_L + \delta$, where $\delta \sim N(0, 1)$.

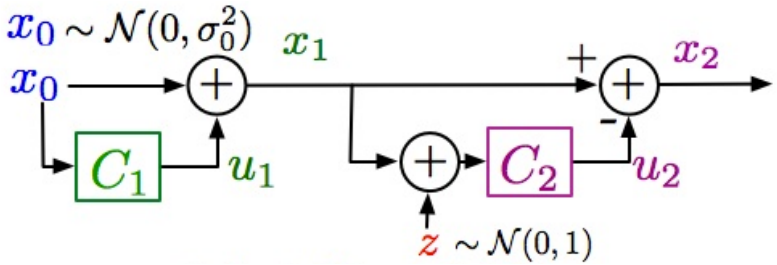

$$\min\{k^2 \mathbb{E}[u_1{}^2] + \mathbb{E}[x_2{}^2]\}$$
$$= \min\{k^2 \mathbb{E}[(x_0 - x_1)^2] + \mathbb{E}[(u_2 - x_1)^2]\}$$

2. Construct a class of near-piecewise-linear strategies for the leader, invariant under best response for large σ.
3. Show existence of equilibria with leader's strategy in this set.
4. Show existence of a collection of local minima of the same order as of the global optimal solution.

Invariant Set of Near-Piecewise-Linear Strategies

Perfect Bayesian equilibria of the game:

$$a_F^*(s) = \mathbb{E}_{\nu^*}[a_L^*|s] = \int_{\mathbb{R}} a_L \nu^*(a_L|s) da_L,$$

$$a_L^*(\theta) = \underset{a_L}{\operatorname{argmax}} -r_L(\theta - a_L)^2 - (1 - r_L) \int_{\mathbb{R}} (a_F^*(s) - a_L)^2 \phi(s - a_L) ds.$$

- $\nu^*(\cdot|s)$: the follower's belief on leader's action given s.

Partition $N(0, \sigma^2)$ into $2m + 1$ segments:

- $B_k^0 = [b_k^0, b_{k+1}^0)$ for $k \in \mathbb{N}_m$.
  $B_0^0 = (b_{-1}^0, b_1^0)$, $B_{-k}^0 = (b_{-k-1}^0, b_{-k}^0]$.
- Symmetric: $b_{-k}^0 = -b_k^0$. Unbounded tails: $b_{m+1}^0 = -b_{-m-1}^0 = +\infty$.

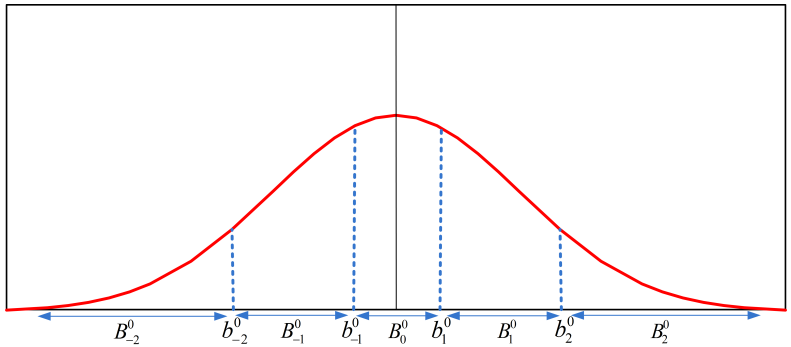

Invariant Set of Near-Piecewise-Linear Strategies

Perfect Bayesian equilibria of the game:

$$a_F^*(s) = \mathbb{E}_{\nu^*}[a_L^*|s] = \int_{\mathbb{R}} a_L \nu^*(a_L|s) da_L,$$

$$a_L^*(\theta) = \underset{a_L}{\operatorname{argmax}} -r_L(\theta - a_L)^2 - (1 - r_L) \int_{\mathbb{R}} (a_F^*(s) - a_L)^2 \phi(s - a_L) ds.$$

- $\nu^*(\cdot|s)$: the follower's belief on leader's action given s.

Partition $N(0, \sigma^2)$ into $2m + 1$ segments:

- $B_k^0 = [b_k^0, b_{k+1}^0)$ for $k \in \mathbb{N}_m$.
  $B_0^0 = (b_{-1}^0, b_1^0)$, $B_{-k}^0 = (b_{-k-1}^0, b_{-k}^0]$.
- Symmetric: $b_{-k}^0 = -b_k^0$. Unbounded tails: $b_{m+1}^0 = -b_{-m-1}^0 = +\infty$.
- Let centroid $c_k^0 = \mathbb{E}_{N(0,\sigma^2)}[\theta|\theta \in B_k^0]$.

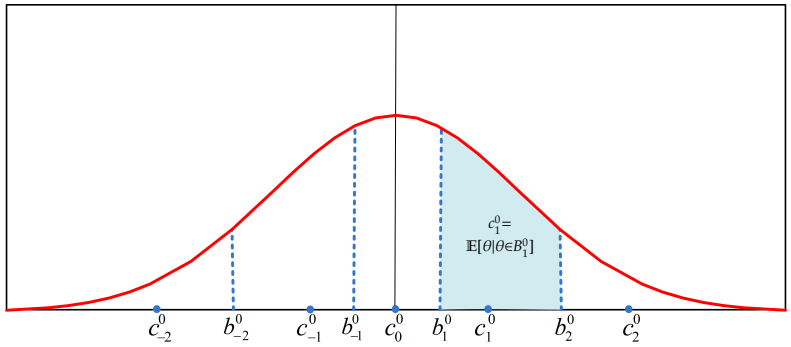

Invariant Set of Near-Piecewise-Linear Strategies

Perfect Bayesian equilibria of the game:

$$a_F^*(s) = \mathbb{E}_{\nu^*}[a_L^*|s] = \int_{\mathbb{R}} a_L \nu^*(a_L|s) da_L,$$

$$a_L^*(\theta) = \underset{a_L}{\operatorname{argmax}} -r_L(\theta - a_L)^2 - (1 - r_L) \int_{\mathbb{R}} (a_F^*(s) - a_L)^2 \phi(s - a_L) ds.$$

- $\nu^*(\cdot|s)$: the follower's belief on leader's action given s.

Partition $N(0, \sigma^2)$ into $2m + 1$ segments:

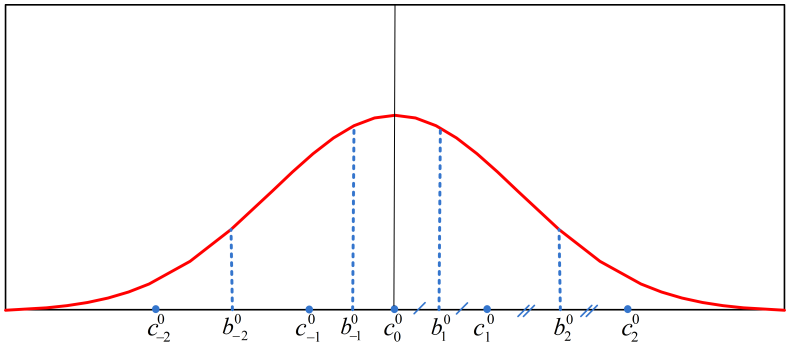

- $B_k^0 = [b_k^0, b_{k+1}^0)$ for $k \in \mathbb{N}_m$.
  $B_0^0 = (b_{-1}^0, b_1^0)$, $B_{-k}^0 = (b_{-k-1}^0, b_{-k}^0]$.
- Symmetric: $b_{-k}^0 = -b_k^0$. Unbounded tails: $b_{m+1}^0 = -b_{-m-1}^0 = +\infty$.
- Let centroid $c_k^0 = \mathbb{E}_{N(0,\sigma^2)}[\theta|\theta \in B_k^0]$. Then, $b_k^0 = \frac{c_{k-1}^0 + c_k^0}{2}$ (interval endpoints are equidistant from adjacent centroids).

Such a partition exists and is unique.

An Invariant Set of Near-Piecewise-Linear Strategies

$A_L^m(r_L, \sigma)$: a set of $(2m+1)$-segmented increasing odd functions.

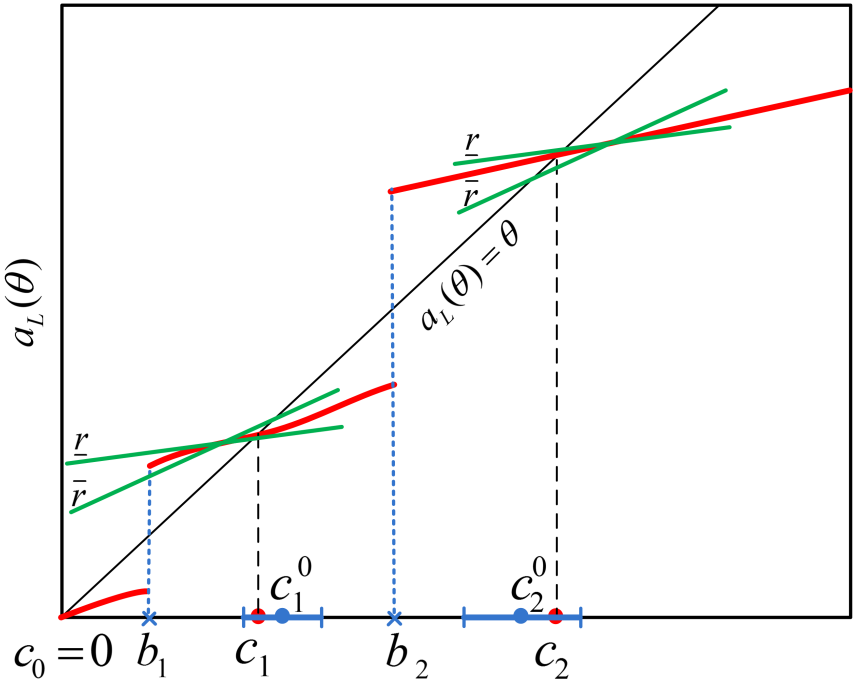

We consider strategies with the following properties:

- segments close to B_k^0
- a fixed point in each segment in a certain vicinity of c_k^0
- almost linear with a slope close to r_L over each segment

An Invariant Set of Near-Piecewise-Linear Strategies

$A_L^m(r_L, \sigma)$: a set of $(2m+1)$-segmented increasing odd functions.

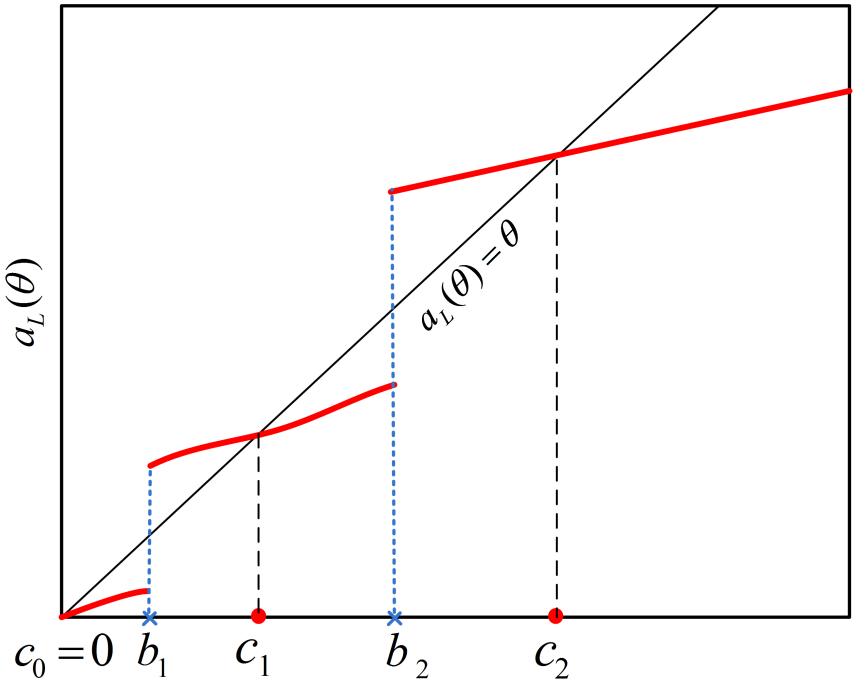

Property 1: For every $a_L(\theta) \in A_L^m(r_L, \sigma)$, there exist $2m+1$ segments $B_k = [b_k, b_{k+1})$, for $k \in \mathbb{N}_m$, $B_0 = (-b_1, b_1)$, and $B_{-k} = (b_{-k-1}, b_{-k}]$, with $b_{m+1} = -b_{-m-1} = +\infty$ such that:

- $a_L(\theta)$ is increasing and odd, and is smooth over each interval.
- $a_L(\theta)$ has a unique fixed point ($a_L(c_k) = c_k$) in each segment.

An Invariant Set of Near-Piecewise-Linear Strategies

$A_L^m(r_L, \sigma)$: a set of $(2m+1)$-segmented increasing odd functions.

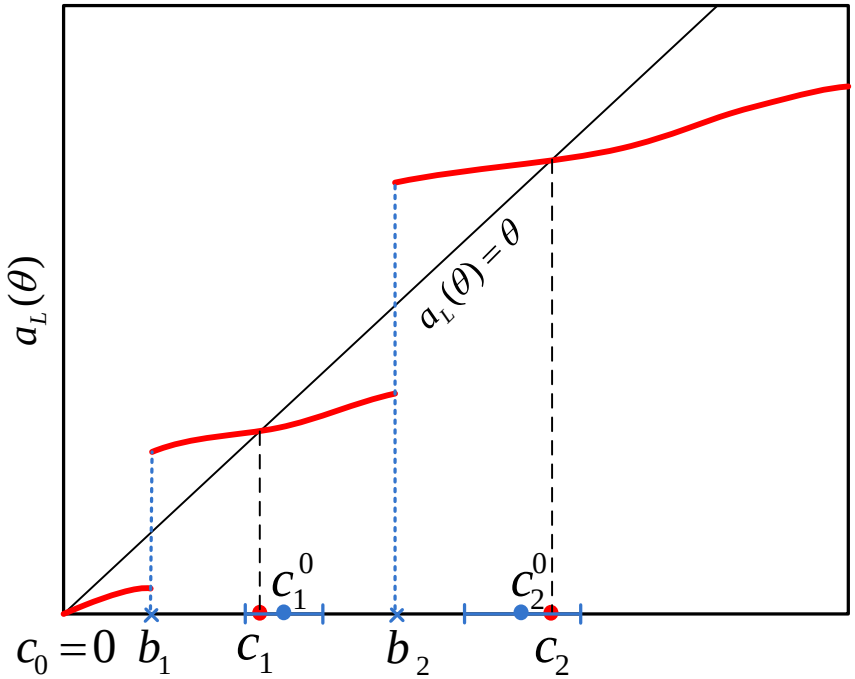

Property 2: For every $k \in \mathbb{N}_m$, $|b_k - \frac{c_{k-1}+c_k}{2}| \leq 0.1 r_L$. Moreover, $|c_k - c_k^0| \leq 2.9$.

Let $\bar{x}_k := x_k^0 + 3$ and $\underline{x}_k := x_k^0 - 3$, where quantizer's half-steps $x_k^0 = \frac{c_k^0 - c_{k-1}^0}{2}$. Then $\bar{x}_k$ and $\underline{x}_k$ bound lengths of half-intervals $[c_{k-1}, b_k]$ and $[b_k, c_{k+1}]$.

An Invariant Set of Near-Piecewise-Linear Strategies

$A_L^m(r_L, \sigma)$: a set of $(2m+1)$-segmented increasing odd functions.

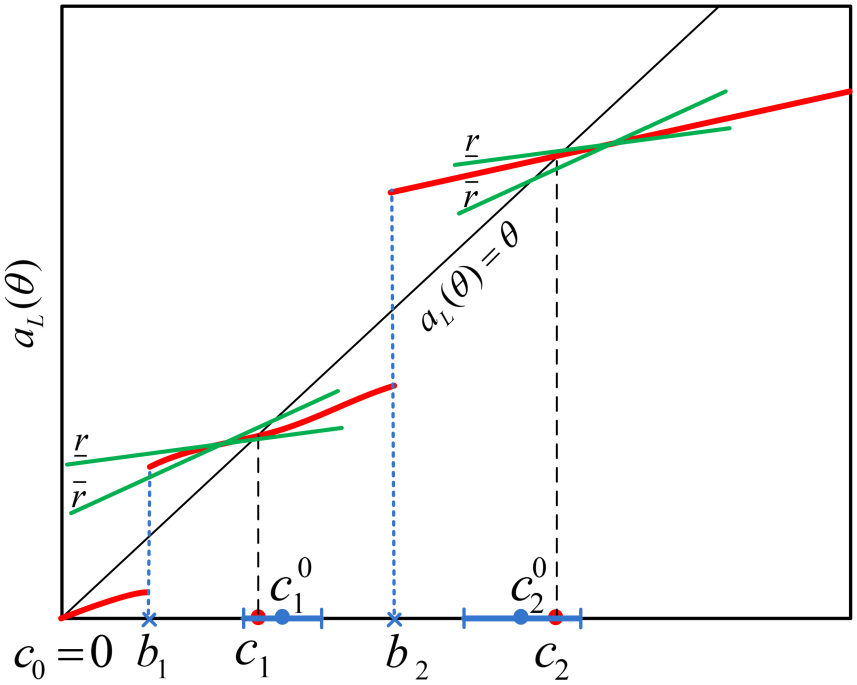

Property 3: Over inner intervals B_k $(k \neq -m, m)$, $\underline{r} \leq \frac{d}{d\theta} a_L(\theta) \leq \bar{r}$, where $\underline{r} = r_L(1 - 0.5 r_L^2 \sigma^2)$ and $\bar{r} = r_L(1 + 0.5 r_L^2 \sigma^2)$. Same bound over the tail if $b_m < \theta < c_m + \sqrt{e}\sigma\bar{x}_m$. For $\theta > c_m + \sqrt{e}\sigma\bar{x}_m$ we have $a_L(\theta) \leq c_m + 3r_L(\theta - c_m)$.

An Invariant Set of Near-Piecewise-Linear Strategies

$A_L^m(r_L, \sigma)$: Invariance under best response operator

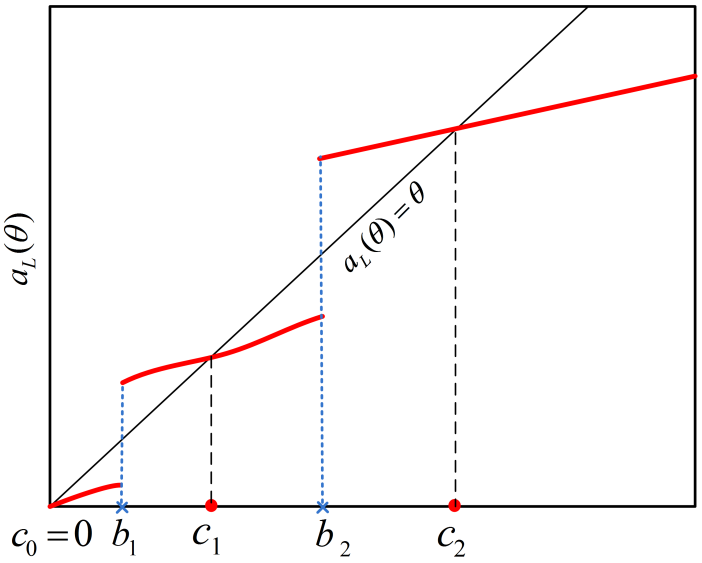

An Invariant Set of Near-Piecewise-Linear Strategies

$A_L^m(r_L, \sigma)$: Invariance under best response operator

$$a_L(\theta) \xrightarrow{\text{BR}} a_F(s)$$

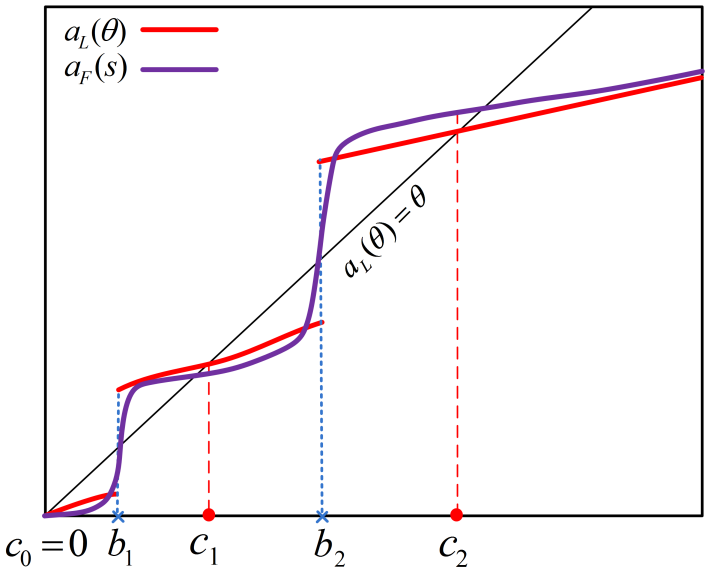

An Invariant Set of Near-Piecewise-Linear Strategies

$A_L^m(r_L, \sigma)$: Invariance under best response operator

$$a_L(\theta) \xrightarrow{\text{BR}} a_F(s) \xrightarrow{\text{BR}} \tilde{a}_L(\theta)$$

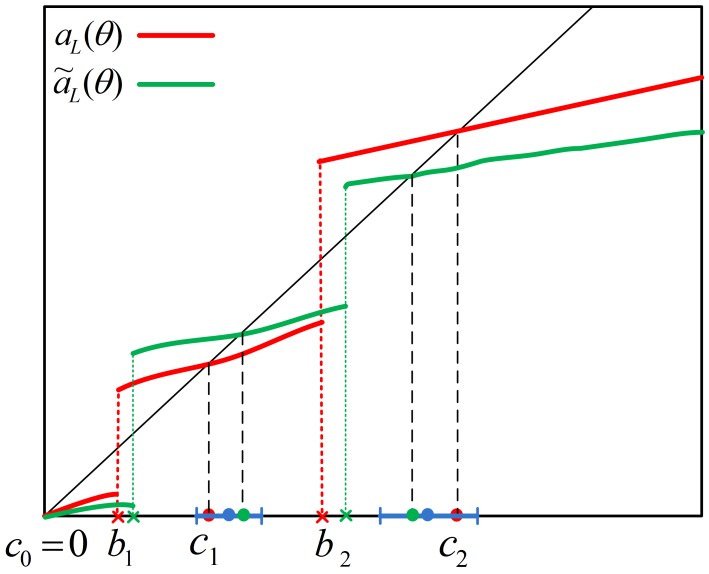

An Invariant Set of Near-Piecewise-Linear Strategies

$A_L^m(r_L, \sigma)$: Invariance under best response operator

$$a_L(\theta) \xrightarrow{\text{BR}} a_F(s) \xrightarrow{\text{BR}} \tilde{a}_L(\theta)$$

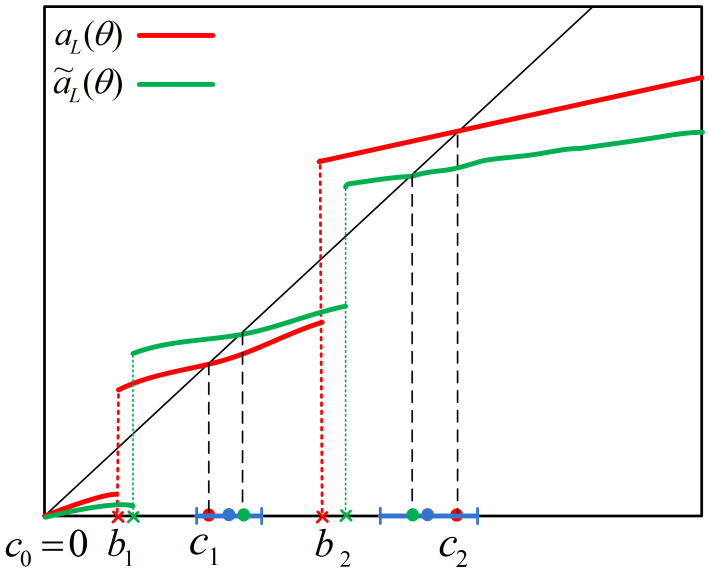

$A_L^m(r_L, \sigma)$ invariant for $m \in M(\sigma) = \{m \in \mathbb{N} | 2\sqrt{2\ln\sigma} + 4 < x_1^0 < 4\sqrt{\ln\sigma}\}$ and sufficiently large σ in the regime $\frac{1}{2} \leq r_L\sigma^2 \leq 1$.

- choice of $m \in M(\sigma) \Rightarrow$ *efficient strategies*.
- uses results from asymptotic quantization theory.

Nonlinear Equilibria

Theorem 1

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 for any

$$m \in M(\sigma) = \{m \in \mathbb{N} | 2\sqrt{2\ln\sigma} + 4 < x_1^0 < 4\sqrt{\ln\sigma}\},$$

and sufficiently large σ in the regime $\frac{1}{2} \leq r_L\sigma^2 \leq 1$. Moreover, the game has an equilibrium for which:

i) $a_L^*(\theta, r_L, \sigma) \in A_L^m(r_L, \sigma)$, *and*

ii) $(a_L^*(\theta, r_L, \sigma), a_F^*(s) = \mathbb{E}_\delta[a_L^*|s])$ *maximizes the expected payoff of the leader over all pair of strategies* $(a_L(\theta, r_L, \sigma), a_F(s) = \mathbb{E}_\delta[a_L|s])$ *where* $a_L(\theta, r_L, \sigma) \in A_L^m(r_L, \sigma)$.

Proof based on careful best response analysis, very involved!

Local Minima and Performance Guarantees

Expected disutility of the leader:

$$U(a_L, a_F) = r_L \int_{\mathbb{R}} (\theta - a_L(\theta))^2 \frac{\phi(\frac{\theta}{\sigma})}{\sigma} d\theta$$

$$+(1 - r_L) \int \int_{\mathbb{R}^2} (a_F(s) - a_L(\theta))^2 \phi(s - a_L(\theta)) \frac{\phi(\frac{\theta}{\sigma})}{\sigma} ds d\theta. \quad (1)$$

Lemma 1

Any pair of equilibrium strategies (a_L^*, a_F^*) characterized by Theorem 1, where $a_L^* \in A_L^m(r_L, \sigma)$ and $a_F^*(s) = \mathbb{E}_\delta[a_L^*|s]$ is a local minimum of U.

- U maps to the original Witsenhausen's cost.
- (a_L^*, a_F^*) equilibrium $\Rightarrow$ cost cannot be improved by changing (only) one of the strategies.
- Not ruling out a lower cost by simultaneously changing both.
- Idea of the proof:
  - The best response image of an infinitesimal variation in a_L^* also lies in A_L^m.
  - (a_L^*, a_F^*) is the minimizer of U over all pairs (a_L, a_F) with $a_L \in A_L^m$ (Theorem 1).

Performance Guarantees

- Analytically proved the existence of local minima for Witsenhausen's problem with a near-piecewise-linear strategy for the leader (or first controller).
- Next: evaluate their performance w.r.t optimal cost.
- Consider pair of strategies (a_L^0, a_F^0):
  - a_L^0: piecewise-linear strategy with segments B_k^0, constant slope r_L, and centroids c_k^0 as fixed points.
  - a_F^0: optimal $(2m+1)$-level MSE quantizer (i.e., constant value of c_k^0 over segment B_k^0).
- Any pair (a_L, a_F) with $a_L \in A_L^m$ can be used to upper-bound $U(a_L^*, a_F^*)$ from Theorem 1:

$$a_L^0 \in A_m^L \Rightarrow U(a_L^*, a_F^*) \leq U(a_L^0, a_F^0).$$

- $U(a_L^0, a_F^0)$ is easier to play with ...
- Exact asymptotics are known for expected loss of optimal $(2m+1)$-level MSE quantizer.

Performance Guarantees (cont'd)

$$U(a_L^0, a_F^0) = r_L(1-r_L)^2 D_L^0 + (1-r_L) D_F^0,$$

$$D_L^0 = \int_{\mathbb{R}} (\theta - a_F^0(\theta))^2 \frac{\phi(\frac{\theta}{\sigma})}{\sigma} d\theta,$$

$$D_F^0 = \int \int_{\mathbb{R}} (a_F^0(s) - a_L^0(\theta))^2 \phi(s - a_L^0(\theta)) \frac{\phi(\frac{\theta}{\sigma})}{\sigma} ds d\theta.$$

- D_F^0 can be upper-bounded as: $D_F^0 \leq 4\sqrt{\frac{2}{e}} \frac{(2-r_L)^2}{(1-r_L)^2} \phi(\frac{x_1^0}{\sqrt{2}}) + r_L^2 D_L^0.$
- D_L^0 is the expected loss of an optimal $(2m+1)$-level MSE quantizer for a source $\theta \sim N(0, \sigma^2)$.
- Exact asymptotics for D_L^0 from asymptotic quantization theory.
- for large m, $D_L^0 \approx \frac{c_\infty}{(2m+1)^2}$, where c_∞ is the *Panter-Dite constant* of a normal source:

$$c_\infty = \frac{1}{12} \left(\int_{\mathbb{R}} (\frac{\phi(\frac{\theta}{\sigma})}{\sigma})^{\frac{1}{3}} d\theta \right)^3 = \frac{\sqrt{3}\pi}{2} \sigma^2.$$

- Another exact asymptotic: $(2m+1)\frac{x_1^0}{\sigma} \approx \frac{\sqrt{6\pi}}{2} \Rightarrow D_L^0 \approx \frac{(x_1^0)^2}{\sqrt{3}}.$

Performance Guarantees (cont'd)

$$U(a_L^0, a_F^0) = r_L(1-r_L)^2 D_L^0 + (1-r_L) D_F^0,$$

$$D_L^0 = \int_{\mathbb{R}} (\theta - a_F^0(\theta))^2 \frac{\phi(\frac{\theta}{\sigma})}{\sigma} d\theta,$$

$$D_F^0 = \int \int_{\mathbb{R}} (a_F^0(s) - a_L^0(\theta))^2 \phi(s - a_L^0(\theta)) \frac{\phi(\frac{\theta}{\sigma})}{\sigma} ds d\theta.$$

- D_F^0 can be upper-bounded as: $D_F^0 \leq 4\sqrt{\frac{2}{e}} \frac{(2-r_L)^2}{(1-r_L)^2} \phi(\frac{x_1^0}{\sqrt{2}}) + r_L^2 D_L^0.$
- D_L^0 is the expected loss of an optimal $(2m+1)$-level MSE quantizer for a source $\theta \sim N(0, \sigma^2)$.
- Exact asymptotics for D_L^0 from asymptotic quantization theory.
- for large m, $D_L^0 \approx \frac{c_\infty}{(2m+1)^2}$, where c_∞ is the *Panter-Dite constant* of a normal source:

$$c_\infty = \frac{1}{12} \left(\int_{\mathbb{R}} (\frac{\phi(\frac{\theta}{\sigma})}{\sigma})^{\frac{1}{3}} d\theta \right)^3 = \frac{\sqrt{3}\pi}{2} \sigma^2.$$

- Another exact asymptotic: $(2m+1)\frac{x_1^0}{\sigma} \approx \frac{\sqrt{6\pi}}{2} \Rightarrow D_L^0 \approx \frac{(x_1^0)^2}{\sqrt{3}}.$

Performance Guarantees (cont'd)

Lemma 2

For the pair of equilibrium strategies (a_L^*, a_F^*) characterized by Theorem 1, where $a_L^* \in A_L^m(r_L, \sigma)$, $a_F^*(s) = \mathbb{E}_\delta[a_L^*|s]$, and $m \in M(\sigma)$ we have

$$\liminf_{\sigma\to\infty} \frac{\frac{r_L(1-r_L)(x_1^0)^2}{\sqrt{3}} + 4\sqrt{\frac{2}{e}}\frac{(2-r_L)^2}{(1-r_L)}\phi(\frac{x_1^0}{\sqrt{2}})}{U(a_L^*, a_F^*)} \geq 1.$$

- The above asymptotic upper bound on $U(a_L^*, a_F^*)$ is minimized when $x_1^0 \approx 2\sqrt{2\ln\sigma}$ for large σ, yielding a cost $\approx \frac{8r_L \ln\sigma}{\sqrt{3}}$.
- $M(\sigma) = \{m \in \mathbb{N} | 2\sqrt{2\ln\sigma} + 4 < x_1^0 < 4\sqrt{\ln\sigma}\} \Rightarrow$ cost (asymptotically) as low as $\frac{8r_L \ln\sigma}{\sqrt{3}}$ is achievable for some $m \in M(\sigma)$.
- To compare with the global optimum, we can use known lower bounds on the optimal cost of Witsenhausen's problem.
- Next lemma is an immediate result of the bounds derived in Grover et al. (2013).

Performance Guarantees (cont'd)

Lemma 3

Denote with $U^*(\sigma)$ the minimum value of the cost functional $U(a_L, a_F)$ in the regime $r_L\sigma^2 = 1$. Then, $\limsup_{\sigma\to\infty} \frac{\frac{\ln\sigma}{6\sigma^2}}{U^*(\sigma)} \leq 1$.

A quite loose but helpful lower bound!

Theorem 2

Any pair of equilibrium strategies (a_L^*, a_F^*) characterized by Theorem 1, where $a_L^* \in A_L^m(r_L, \sigma)$ and $a_F^*(s) = \mathbb{E}_\delta[a_L^*|s]$, is a local minimum of the cost functional U in (1). Moreover, the set $M(\sigma)$ is nonempty for sufficiently large values of σ and,

$$\liminf_{\sigma\to\infty} \frac{\frac{8r_L \ln\sigma}{\sqrt{3}}}{\min_{m\in M(\sigma)} U(a_L^*, a_F^*)} \geq 1.$$

In the regime $r_L\sigma^2 = 1$, all these local minima are within constant factor of the optimal cost, with at least one being less than 27.8 times away from the optimal value as $\sigma \to \infty$.

Concluding Remarks

- Witsenhausen's seminal counterexample: a simple two-stage LQG decentralized control problem, where optimal solution is nonlinear.
- Little known about the structure of global and local optima.
- Simulations suggest optimality of *near-piecewise-linear* strategies in large σ regime.
- We view the problem as a leader-follower game of incomplete information.
- We prove existence of *near-piecewise-linear* local optima with a cost at most a constant factor away from the optimal one, in large σ regime.

Concluding Remarks

- Still many fundamental open questions about the structural properties of local/global optima.
- It is well-known that optimal solution has a real analytic left inverse. Same can be shown for local optima.
- This does not imply continuity. Can the optimal solution be discontinuous?
- When viewed as an LQG leader-follower game, is there a phase transition in the emergence of nonlinear equilibria?
- Can nonlinear equilibria emerge when actions are strategic substitutes?
- The latter relates to the yet open conjecture on uniqueness of equilibrium in classical Kyle model in Finance.

Thank You!