Special Topics Seminar

Affine Laws and Learning Approaches for Witsenhausen Counterexample

Hajir Roozbehani

Dec 7, 2011

Outline

- Optimal Control Problems
  - Affine Laws
  - Separation Principle
  - Information Structure
- Team Decision Problems
  - Witsenhausen Counterexample
  - Sub-optimality of Affine Laws
  - Quantized Control
- Learning Approach

Linear Systems

Discrete Time Representation

In a classical multistage stochastic control problem, the dynamics are

$$
\begin{aligned}
x(t+1) &= Fx(t) + Gu(t) + w(t) \\
y(t) &= Hx(t) + v(t),
\end{aligned}
$$

where $v(t)$ and $y(t)$ are independent sequences of random variables and $u(t) = \gamma(y(t))$ is the control law (or decision rule). A cost function

$$J(\gamma, x(0))$$

is to be minimized.

Linear Systems

Discrete Time Representation

In a classical multistage stochastic control problem, the dynamics are

$$x(t+1) = Fx(t) + Gu(t) + w(t)$$
$$y(t) = Hx(t) + v(t),$$

where $v(t)$ and $y(t)$ are independent sequences of random variables and $u(t) = \gamma(y(t))$ is the control law (or decision rule). A cost function

$$J(\gamma, x(0))$$

is to be minimized.

Success Stories with Affine Laws

LQR

Consider a linear dynamical system

$$x(t+1) = Fx(t) + Gu(t), x(t) \in \mathbb{R}^n, u(t) \in \mathbb{R}^m$$

with complete information and the task of finding a pair $(x(t), u(t))$ that minimizes the functional

$$J(u(t)) = \sum_{t=0}^{T} [x(t)^{'} Qx(t) + u(t)^{'} Ru(t)],$$

subject to the described dynamical constraints and for $Q > 0, R > 0$. This is a convex optimization problem with an affine solution:

$$u^{*}(t) = -R^{-1} B^{'} P(t) x(t),$$

where $P(t)$ is to be found by solving algebraic Riccati equations.

Certainty Equivalence

LQR

Consider a linear dynamical system

$$x(t+1) = Fx(t) + Gu(t) + w(t), x(t) \in \mathbb{R}^n, u(t) \in \mathbb{R}^m$$

with imperfect information and the task of finding a law $u(t) = \gamma(x(t))$ that minimizes the functional

$$J(u(t)) = \sum_{t=0}^{T} \mathbb{E}[x(t)' Qx(t) + u(t)' Ru(t)],$$

subject to the described dynamical constraints and for $Q > 0, R > 0$. This is a convex optimization problem with an affine solution:

$$u^*(t) = -R^{-1} B' P(t)x(t),$$

where $P(t)$ is to be found by solving algebraic Riccati equations.

Classical vs Optimal Control

- Beyond its optimality properties, affinity enables us to make tight connections between classical and modern control.
- The steady state approximation $P(t) = P$ of LQR amounts to the classical proportional controller $u = -Kx$.

Figure: Hendrik Wade Bode and Rudolf Kalman

Optimal Filter

LQG

Now consider the problem of estimating the state of a dynamical system that evolves at the presence of noise

$$\begin{aligned} x(t+1) &= Fx(t) + Gu(t) + w(t) \\ y(t) &= Hx(t) + v(t), \end{aligned}$$

where $w(t)$ and $v(t)$ are independent stochastic processes.

- What is $\mathbb{E}[x(t)|\mathcal{F}_{Y(t)}]$? Kalman gave the answer: this is the dual of LQR that we just saw.
- Why is this important?
- How about the optimal smoother $\mathbb{E}[x(0)|\mathcal{F}_{Y(t)}]$?

Optimal Smoother

Linear Systems

Assume that the goal is to design a causal control

$$\gamma : y \to u$$
$$\pi : (x_0, u, w) \to y$$

that gives the best estimate of (uncertain) initial conditions of the system. Let $\mathcal{F}_t(\gamma(.))$ denote the filtration generated by control law $\gamma(.)$. For linear systems:

$$\text{var}(\mathbb{E}[x_0|\mathcal{F}_{Y_t}(u(t))]) = \text{var}(\mathbb{E}[x_0|\mathcal{F}_{Y_t}(0)])$$

(there is no reward for amplifying small perturbations)

Separation Principle

- The solution to all mentioned problems is linear when dealing with linear systems
- How about a problem that involves both estimation and control? i.e.,

$$\text{minimize} \quad \mathbb{E}[J(\gamma(y_t))].$$

subject to

$$\begin{aligned} x(t+1) &= Fx(t) + Gu(t) + w(t) \\ y(t) &= Hx(t) + v(t). \end{aligned}$$

Under some mild assumptions a composition of optimal control and optimal estimator is optimal

$$\begin{aligned} u^* &= -K(t)\hat{x}(t) \\ \hat{x} &= -L(t)y(t) \end{aligned}$$

(known as separation principle)

Role of Linearity in Separation Principle

- Fails for simplest forms of nonlinearity

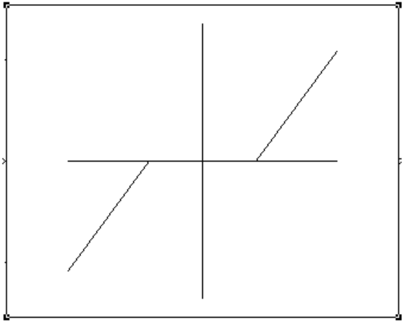

Information Structure

Let us think about the information required to implement an affine law in linear systems. Recall

$$x_{t+1} = Fx_t + Gu_t + w_t$$
$$y_t = Hx_t + v_t.$$

How does $y(t)$ depends on $u(\tau)$ for $\tau \leq t$? This is a convolution sum

$$y_t = \sum_{k=1}^{t} HF^k Gu_k = \sum_{k=1}^{t} D_k u_k$$

When the world is random

$$y_t = H\eta_t + \sum_{k=1}^{t} D_k u_k,$$

with $\eta_t = (x_0, w_1, w_2, ..., w_t, v_1, v_2, ..., v_t)'$.

- precedence $\Rightarrow$ dynamics are coupled ($D_k \neq 0$ for some k).

$$y_t = H\eta_t + \sum_{k=1}^{t} D_k u_k$$

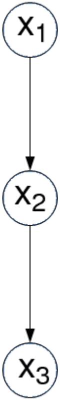

- perfect recall $\Rightarrow \eta_s \subset \eta_t \iff s \leq t$.

$$y_t = H\eta_t + \sum_{k=1}^{t} D_k u_k$$

Classical Structure

- perfect recall $\Rightarrow \eta_s \subset \eta_t \iff s \leq t$.
- precedence+ perfect recall $\Rightarrow$ classical structure [2].

$$y_t = H\eta_t + \sum_{k=1}^{t} D_k u_k$$

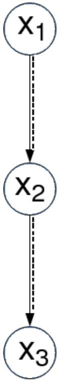

Classical Structure

- perfect recall $\Rightarrow \eta_s \subset \eta_t \iff s \leq t$.
- precedence+ perfect recall $\Rightarrow$ classical structure.
- equivalent to observing only external randomness.

$$y_t = H\eta_t$$

how does this contribute to separation?

Connection between Information Structure and Separation

- The fact that the information set can be reduced to $\{H\eta_t\}$ implies the separation (one cannot squeeze more information by changing the observation path!)
- This is mainly due to the fact that control depends in a deterministic fashion to randomness in external world.
- Main property that allows separation: use all of control to minimize the cost without having to worry how to gain more information!
- Rigorously proving the separation theorem, and classifying systems for which it holds is an unresolved matter in stochastic control [1].

Information Structure (Partially Nested)

- Same holds for partially nested structure [2](followers have perfect recall).

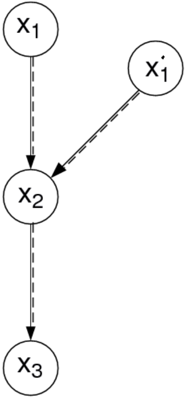

Figure: Adapted from [2]

Team Decision Problems

Recap on Success Stories

- The class of affine laws gives us strong results for dealing with various problems: optimal controller/filter/smoother/etc.
- But the success story had an end!

Decentralized control

- Are affine laws optimal when the information structure is non-classical?
- Conjectured to be true for almost a decade. Witsenhausen proved wrong [6].

Witsenhausen Counterexample

A classical example that shows affine laws are not optimal in decentralized control problems.

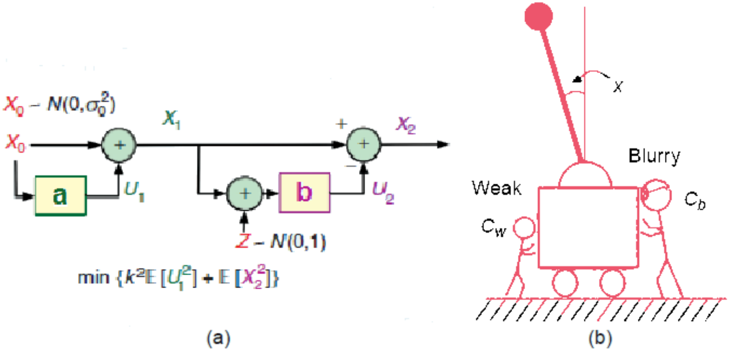

Figure: Adapted from [5]

Witsenhausen Counterexample

A classical example that shows affine laws are not optimal in decentralized control problems.

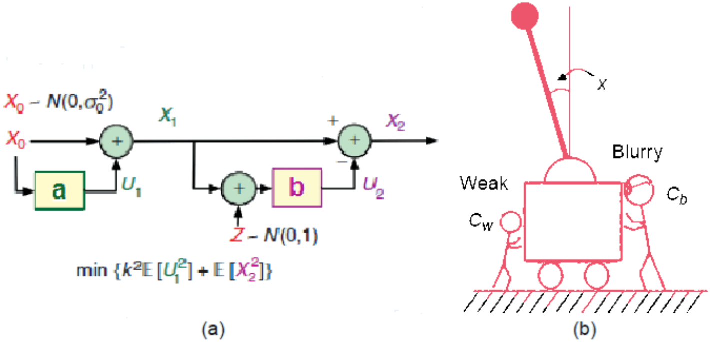

Figure: Adapted from [5]

- Without the noise on the communication channel, the problem is easy! (optimal cost zero).

Witsenhausen Counterexample

A classical example that shows affine laws are not optimal in decentralized control problems.

- We will see by an example why the change of information structure makes the problem non-convex

Witsenhausen Counterexample

A classical example that shows affine laws are not optimal in decentralized control problems.

- We will see by an example why the change of information structure makes the problem non-convex
- In essence, when one forgets the past, the estimation equality becomes control dependent. This is because control can vary the extent to which the forgotten data can be recovered (control has dual functionalities).

Witsenhausen Counterexample

A classical example that shows affine laws are not optimal in decentralized control problems.

- We will see by an example why the change of information structure makes the problem non-convex
- In essence, when one forgets the past, the estimation equality becomes control dependent. This is because control can vary the extent to which the forgotten data can be recovered (control has dual functionalities).
- Thus, the main difficulty is to find the first stage control (Witsenhausen characterized the optimal second stage control as a function of the first stage control [6]).

Witsenhausen Counterexample

A two stage problem ("encoder/decoder"):

- first stage: $x_1 = x_0 + u_1$ and $y_1 = x_0$, $x0 \sim N(0, \sigma^2)$
- second stage: $x_2 = x_1 - u_2$ and $y_2 = x_1 + w$, $w \sim N(0, 1)$

Note the non-classical structure $y_2 = \{x_1 + w\}$ as opposed to the classical $y_2 = \{x_0, x_1 + w\}$. The cost is

$$\mathbb{E}[ku_1^2 + x_2^2],$$

where k is a design parameter. Look for feedback laws $u_1 = \gamma(y_1), u_2 = \gamma(y_2)$ that minimize the cost.

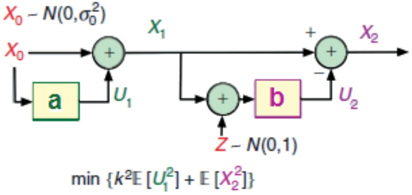

Optimal Affine Law

- The second stage is an estimation problem since $x_2 = x_1 - u_2$.
- Let $u_2 = by_2$ and $u_1 = ay_1$. What is the best estimate of x_1?

$$u_2 = \mathbb{E}[x_1 | y_2] = \frac{\mathbb{E}x_2 y_2}{\mathbb{E}y_2^2} = \frac{(1+a)^2\sigma^2}{(1+a)^2\sigma^2 + 1} y$$

- The expected cost is

$$k^2 a^2 \sigma^2 + \frac{(1+a)^2\sigma^2}{(1+a)^2\sigma^2 + 1} y.$$

Let $t = \sigma(1+a)$ and minimize w.r.t t to find the optimal gain as the fixed point of

$$\sigma - \frac{t}{k^2(1+t^2)^2}$$

Where Convexity Fails?

- The second stage is an estimation problem since $x_2 = x_1 - u_2$.
- Let $u_2 = by_2$ and $u_1 = ay_1$. What is the best estimate of x_1?

$$u_2 = \mathbb{E}[x_1|y_2] = \frac{\mathbb{E}x_2y_2}{\mathbb{E}y_2^2} = \frac{(1+a)^2\sigma^2}{(1+a)^2\sigma^2+1}y$$

- The expected cost is

$$k^2a^2\sigma^2 + \frac{(1+a)^2\sigma^2}{(1+a)^2\sigma^2+1}y.$$

Let $t = \sigma(1+a)$ and minimize w.r.t t to find the optimal gain as the fixed point of

$$\sigma - \frac{t}{k^2(1+t^2)^2}$$

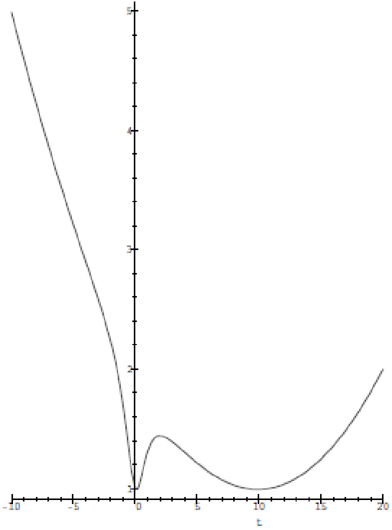

Figure: Expected Cost vs t [4]. Note the local minima!

Nonlinear Controllers

- For $k = 0.1$ and $\sigma = 10$, the expected cost of the optimal affine controller is $0.99 > 0$.
- Witsenhausen suggested a control law for u_1

$$u_1 = -x_0 + \sigma \text{sgn}(x_0),$$

  and a nonlinear control law for u_2

$$u_2 = \sigma \tanh(\sigma y_2).$$

- First stage control gives a binary output (tanh(.) is the MMSE).
- This gives an expected cost of 0.404. How bad can this ratio be?

Quantized Controllers

- Mitter and Sahai [4] proposed 1-bit quantized controllers

$$\gamma(y_1) = -y_1 + \sigma \operatorname{sgn}(y_1)$$
$$\gamma(y_2) = \sigma \operatorname{sgn}(y_2)$$

- The decoding error (proportional to the second stage cost) dies off with $e^{-\sigma^2/2}$.
- Can find limiting values of k and σ for which the expected cost of quantized to linear controller is zero.

Learning Approaches

Properties of Optimal Control

Consider a reformulation of the problem as shown.

- Let $x_1 = x_0 + \gamma(x_0) = f(x_0)$ and $x_2 = f(x_0) - g(f(x) + w)$
- The cost is then given by

$$\mathbb{E}[k^2(x - f(x))^2 + (f(x) - g(f(x) + w))^2]$$

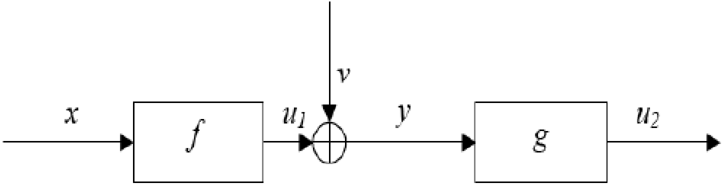

Figure: Witsenhausen Counterexample [3]

Learning Approaches

Properties of Optimal Control

- $f(x)$ is an odd function
- For a given $f(x)$

$$g_f^*(y) = \mathbb{E}[f(x)|y] = \frac{\mathbb{E}_x[f(x)\phi(y - f(x))]}{\mathbb{E}_x[\phi(y - f(x))]}$$

- The cost becomes

$$J(f) = k^2\mathbb{E}[(x - f(x))^2] + 1 - I(D_f),$$

where $I(D_f)$ is the fisher information of random variable y

$$I(D_f) = \int (\frac{d}{dy} D_f(y))^2 \frac{dy}{D_f(y)}$$

with density

$$D_f(y) = \int \phi(y - f(x))\phi(x; 0, \delta^2)dx$$

Learning Approaches

Where Convexity Fails?

- $f(x)$ is an odd function
- For a given $f(x)$

$$g_f^*(y) = \mathbb{E}[f(x)|y] = \frac{\mathbb{E}_x[f(x)\phi(y - f(x))]}{\mathbb{E}_x[\phi(y - f(x))]}$$

- The cost becomes

$$J(f) = k^2\mathbb{E}[(x - f(x))^2] + 1 - I(D_f),$$

where $I(D_f)$ is the fisher information of random variable y

$$I(D_f) = \int (\frac{d}{dy} D_f(y))^2 \frac{dy}{D_f(y)}$$

with density

$$D_f(y) = \int \phi(y - f(x))\phi(x; 0, \delta^2)dx$$

Learning Approaches

Where Convexity Fails?

- $f(x)$ is an odd function
- For a given $f(x)$

$$g_f^*(y) = \mathbb{E}[f(x)|y] = \frac{\mathbb{E}_x[f(x)\phi(y - f(x))]}{\mathbb{E}_x[\phi(y - f(x))]}$$

- The cost becomes

$$J(f) = k^2\mathbb{E}[(x - f(x))^2] + 1 - I(D_f),$$

where $I(D_f)$ is the fisher information of random variable y

$$I(D_f) = \int (\frac{d}{dy} D_f(y))^2 \frac{dy}{D_f(y)}$$

with density

$$D_f(y) = \int \phi(y - f(x))\phi(x; 0, \delta^2)dx$$

- Problem decomposes into convex + non-convex terms

Learning Approaches

Where Convexity Fails?

- $f(x)$ is an odd function
- For a given $f(x)$

$$g_f^*(y) = \mathbb{E}[f(x)|y] = \frac{\mathbb{E}_x[f(x)\phi(y - f(x))]}{\mathbb{E}_x[\phi(y - f(x))]}$$

- The cost becomes

$$J(f) = k^2\mathbb{E}[(x - f(x))^2] + 1 - I(D_f),$$

where $I(D_f)$ is the fisher information of random variable y

$$I(D_f) = \int (\frac{d}{dy} D_f(y))^2 \frac{dy}{D_f(y)}$$

with density

$$D_f(y) = \int \phi(y - f(x))\phi(x; 0, \delta^2) dx$$

- Other convex+non-convex decompositions: quadratic Wasserstein distance+MMSE [7].

Learning Approaches

Properties of Optimal Control

- The new formulation allows us to see why the non-classical problem is not convex ($-I(D_f)$ is concave).
- Cost of stage two can be written as $1 - I(D_f)$. Intuitively, this penalizes how hard it is at step 2 to decode the signal sent at step 1.
- Maximizing the Fisher information amounts to properly separating signals for a given noise level (does not matter if odd or even).
- Optimal control $f(x)$ is symmetric: stage one cost is symmetric (asking for symmetric $f(x)$) and stage two cost does not care!

Step Functions for $f(x)$

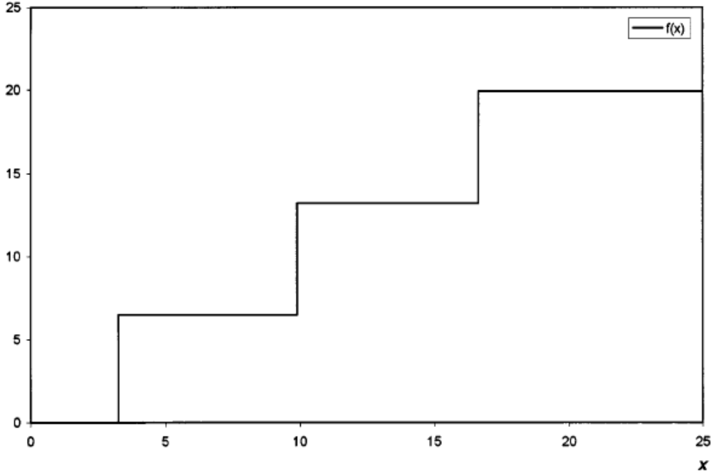

Figure: 3.5 step functions

Major Techniques to Solve the WHC

Solution of $f(x)$ [a]	Total cost J [b]
Optimal affine solution [22] (1968)	0.961852
1-step; by Witsenhausen [22] (1968)	0.404253
1-step; by Bansal and Basar [1] (1987)	0.365015
2-step; by Deng and Ho [4] (1999)	0.190
25-step; by Ho and Lee [12] (2000)	0.1717
2.5-main-step; by Baglietto *et al.* [2] (2001)	0.1701
3.5-main-step; by Lee *et al.* [13] (2001)	0.1673132
3.5-main-step; by our work in this paper (2009)	0.1670790

Figure: Some benchmark statistics [3]

Learning Approach to the WHC

- Divide $f(x)$ into intervals
- Can compute $g_f^*(y)$

$$g_f^*(y) = \frac{-\sum_{i=1}^{n+1} q_i a_i \phi(y + a_i) + \sum_{i=1}^{n+1}}{q_i a_i \phi(y - a_i)},$$

where

$$q_i = \int_{b_{i-1}}^{b_i} \phi(s, 0, \delta^2) ds.$$

- Similarly, for a choice of $(a_1, ..., a_n)$, one can compute the expected cost

$$f(x) = \begin{cases} a_1 & (0 = b_0 \leq x < b_1) \\ a_2 & (b_1 \leq x < b_2) \\ \vdots & \vdots \\ a_n & (b_{n-1} \leq x < b_n) \\ a_{n+1} & (b_n \leq x < b_{n+1} = \infty) \\ -f(-x) & x < 0 \end{cases}$$

Figure: Quantized controller

Learning Approach to the WHC

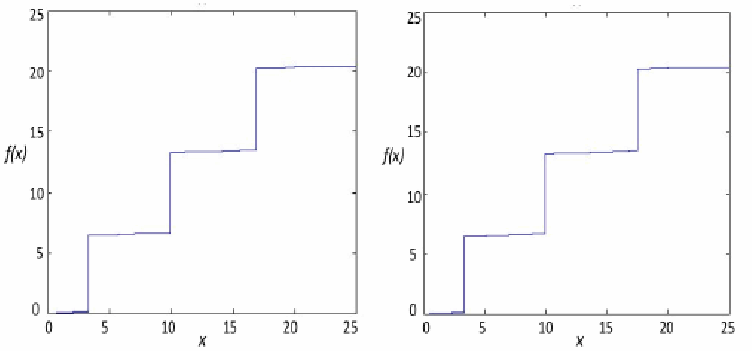

Figure: Optimized quantized control [3].

Learning Approach to the WHC

- Players: intervals $[b_{-i}, b_i)$, $i = \{1, ..., n+1\}$
- Decisions: value $a_i \in \{a | a = a_{max} \frac{\delta}{m} k, k = 0, ..., m\}$.
- Utility function: $U = -J$ (to be maximized).
- Use joint fictitious play with inertia, i.e.,

$$a_i^*(t) = \arg\max \frac{1}{t} \sum_{s=1}^{t} U_s(a_i, a_{-i}(s))$$

with probability $1 - \epsilon$. and

$$a_i^*(t) = a_i^*(t-1)$$

otherwise.

Learning Approach to the WHC

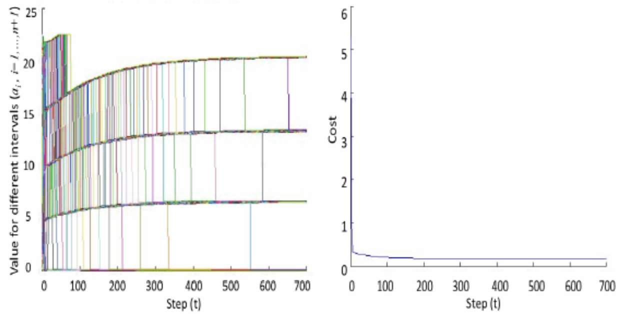

Fig. 2. Convergence processes of a_i and the total cost J in the case where $n = 400$, $m = 450$: each colored plot in the left figure shows one interval's decision value $f(x)$ at each step of iteration. There are totally 400 colored plots in the left figure since we discretize $f(x)$ into 400 steps. The right one shows the convergence of the total cost J.

Figure: Convergence to 3.5 (tilted) step functions [3]

References

[1] Tryphon Georgiou and Anders Lindquist. The separation principle in stochastic control. *arXiv:1103.3005v2*, 2011.

[2] Yu-Chi Ho and K'ai-Ching Chu. Team decision theory and information structures in optimal control problems–part i. *Automatic Control, IEEE Transactions on*, 17(1):15 – 22, feb 1972.

[3] Na Li, Jason R. Marden, and Jeff S. Shamma. Learning approaches to the Witsenhausen counterexample from a view of potential games. In *CDC*, pages 157–162, 2009.

[4] Sanjoy Mitter and Anant Sahai. Information and control: Witsenhausen revisited. In *Learning, control and hybrid systems (Bangalore, 1998)*, volume 241 of *Lecture Notes in Control and Inform. Sci.*, pages 281–293. Springer, London, 1999.

[5] Anant Sahai and Pulkit Grover. Demystifying the Witsenhausen counterexample. *IEEE Control Syst. Mag.*, 30(6):20–24, 2010.

[6] H. S. Witsenhausen. A counterexample in stochastic optimum control. *SIAM J. Control*, 6:131–147, 1968.

[7] Yihong Wu and Sergio Verdú. Witsenhausen's counterexample: A view from optimal transport theory. In *CDC*, 2011.