

Convex Analysis in Stochastic Teams and Asymptotic Optimality of Finite Model Representations and Quantized Policies

Serdar Yüksel

Part of this is joint work with Naci Saldi and Tamás Linder

Queen's University, Canada

Department of Mathematics and Statistics

ACCESS-FORCES CPS Workshop; KTH, 2015

Stochastic Dynamic Team Problems

- Review of Information Structures in Decentralized Control
- Existence and Structural Properties
- Convexity Properties
- Approximation of Team Problems and Asymptotic Optimality of Finite Representations
- Witsenhausen's Counterexample: Non-convexity, Existence and Approximations

Witsenhausen's *Intrinsic Model*

A decentralized control system is called *sequential*, if there is a pre-defined order in which the decision makers (DMs) act. The model consists of:

- A collection of spaces $\{\Omega, \mathcal{F}, (\mathbb{U}^i, \mathcal{U}^i), (\mathbb{Y}^i, \mathcal{Y}^i), i \in \mathcal{N}\}$, specifying the system's control and measurement spaces which are assumed to be standard Borel. $N = |\mathcal{N}|$ is the number of control actions taken. Recall that a *standard Borel space* is a subset of a complete, separable and metric space.
- A *measurement constraint*: The \mathbb{Y}^i -valued observation variables are given by $y^i = \eta^i(\omega, \mathbf{u}^{-i})$, $\mathbf{u}^{-i} = \{u^k, k \leq i - 1\}$.
- A *design constraint*: $\underline{\gamma} = \{\gamma^1, \gamma^2, \dots, \gamma^N\}$: $u^i = \gamma^i(y^i)$, with $y^i = \eta^i(\omega, \mathbf{u}^{-i})$, and γ^i, η^i measurable functions. Let Γ^i denote the set of all admissible policies for DM i and $\Gamma = \prod_k \Gamma^k$.

Characterization of information structures

- A sequential team is *static*, if the information available at every decision maker is only affected by exogenous disturbances (Nature); that is no other decision maker can affect the information at any given decision maker.
- A sequential team problem is *dynamic* if the information available to at least one DM is affected by the action of at least one other DM.
- An IS $\{y^i, 1 \leq i \leq N\}$ is *classical* if y^i contains all of the information available to DM k for $k < i$.
- An IS is *quasi-classical* or *partially nested*, if whenever u^k , for some $k < i$, affects y^i , y^i contains y^k .
- An IS which is not partially nested is *nonclassical*.

Characterization of information structures

- A sequential team is *static*, if the information available at every decision maker is only affected by exogenous disturbances (Nature); that is no other decision maker can affect the information at any given decision maker.
- A sequential team problem is *dynamic* if the information available to at least one DM is affected by the action of at least one other DM.
- An IS $\{y^i, 1 \leq i \leq N\}$ is *classical* if y^i contains all of the information available to DM k for $k < i$.
- An IS is *quasi-classical* or *partially nested*, if whenever u^k , for some $k < i$, affects y^i , y^i contains y^k .
- An IS which is not partially nested is *nonclassical*.

Characterization of information structures

- A sequential team is *static*, if the information available at every decision maker is only affected by exogenous disturbances (Nature); that is no other decision maker can affect the information at any given decision maker.
- A sequential team problem is *dynamic* if the information available to at least one DM is affected by the action of at least one other DM.
- An IS $\{y^i, 1 \leq i \leq N\}$ is *classical* if y^i contains all of the information available to DM k for $k < i$.
- An IS is *quasi-classical* or *partially nested*, if whenever u^k , for some $k < i$, affects y^i , y^i contains y^k .
- An IS which is not partially nested is *nonclassical*.

Characterization of information structures

- A sequential team is *static*, if the information available at every decision maker is only affected by exogenous disturbances (Nature); that is no other decision maker can affect the information at any given decision maker.
- A sequential team problem is *dynamic* if the information available to at least one DM is affected by the action of at least one other DM.
- An IS $\{y^i, 1 \leq i \leq N\}$ is *classical* if y^i contains all of the information available to DM k for $k < i$.
- An IS is *quasi-classical* or *partially nested*, if whenever u^k , for some $k < i$, affects y^i , y^i contains y^k .
- An IS which is not partially nested is *nonclassical*.

Characterization of information structures

- A sequential team is *static*, if the information available at every decision maker is only affected by exogenous disturbances (Nature); that is no other decision maker can affect the information at any given decision maker.
- A sequential team problem is *dynamic* if the information available to at least one DM is affected by the action of at least one other DM.
- An IS $\{y^i, 1 \leq i \leq N\}$ is *classical* if y^i contains all of the information available to DM k for $k < i$.
- An IS is *quasi-classical* or *partially nested*, if whenever u^k , for some $k < i$, affects y^i , y^i contains y^k .
- An IS which is not partially nested is *nonclassical*.

Optimal Policies

Let $\underline{\gamma} = \{\gamma^1, \dots, \gamma^N\}$, and a cost function be defined as:

$$J(\underline{\gamma}) = E[c(\omega_0, \mathbf{u})],$$

for some non-negative loss (or cost) function $c : \Omega \times \prod_k \mathcal{U}^k \rightarrow \mathbb{R}$.

Definition

For a given stochastic team problem with a given information structure, $\{J; \Gamma^i, i \in \mathcal{N}\}$, a policy (strategy) N -tuple $\underline{\gamma}^ := (\gamma^{1*}, \dots, \gamma^{N*})$ is an optimal team decision rule if*

$$J(\underline{\gamma}^*) = \inf_{\underline{\gamma} \in \Gamma} J(\underline{\gamma}) =: J^*$$

Optimal Policies

Definition

An N -tuple of strategies $\underline{\gamma}^* := (\gamma^{1*}, \dots, \gamma^{N*})$ constitutes a person-by-person optimal (pbp optimal) solution if, for all $\beta \in \Gamma^i$ and all $i \in \mathcal{N}$, the following inequalities hold:

$$J^* := J(\underline{\gamma}^*) \leq J(\underline{\gamma}^{-i*}, \beta),$$

where

$$(\underline{\gamma}^{-i*}, \beta) := (\gamma^{1*}, \dots, \gamma^{i-1*}, \beta, \gamma^{i+1*}, \dots, \gamma^{N*}). \quad (1)$$

Witsenhausen's equivalent model and static reduction of sequential dynamic teams

Following Witsenhausen'88, we say that two information structures are equivalent if:

- (i) The policy spaces are isomorphic in the sense that policies under one information structure are realizable under the other information structure,
- (ii) the costs achieved under identical policies are identical almost surely and
- (iii) if there are constraints in the admissible policies, the isomorphism among the policy spaces preserves the constraint conditions.

Witsenhausen's equivalent model and static reduction of sequential dynamic teams

Following Witsenhausen'88, we say that two information structures are equivalent if:

- (i) The policy spaces are isomorphic in the sense that policies under one information structure are realizable under the other information structure,
- (ii) the costs achieved under identical policies are identical almost surely and
- (iii) if there are constraints in the admissible policies, the isomorphism among the policy spaces preserves the constraint conditions.

Witsenhausen's equivalent model and static reduction of sequential dynamic teams

Following Witsenhausen'88, we say that two information structures are equivalent if:

- (i) The policy spaces are isomorphic in the sense that policies under one information structure are realizable under the other information structure,
- (ii) the costs achieved under identical policies are identical almost surely and
- (iii) if there are constraints in the admissible policies, the isomorphism among the policy spaces preserves the constraint conditions.

Witsenhausen's equivalent model and static reduction of sequential dynamic teams

Following Witsenhausen'88, we say that two information structures are equivalent if:

- (i) The policy spaces are isomorphic in the sense that policies under one information structure are realizable under the other information structure,
- (ii) the costs achieved under identical policies are identical almost surely and
- (iii) if there are constraints in the admissible policies, the isomorphism among the policy spaces preserves the constraint conditions.

Witsenhausen's equivalent model and static reduction of sequential dynamic teams

Witsenhausen shows that a large class of sequential team problems admit an equivalent information structure which is static. This is called the *static reduction* of an information structure.

Earlier, for partially observed (or quasi-classical) information structures, a similar reduction was studied by Ho and Chu('72) in the context of LQG systems and a class of invertible non-linear systems.

An equivalence between sequential dynamics teams and their static reduction is as follows.

Witsenhausen's equivalent model and static reduction of sequential dynamic teams

Consider a dynamic team setting according to the intrinsic model where there are N time stages, and each DM observes, $y^k = \eta_k(\omega, u^1, u^2, \dots, u^{k-1})$, and the decisions are generated by $u^k = \gamma_k(y^k)$. The resulting cost under a given team policy is $J(\underline{\gamma}) = E[c(\omega, \mathbf{y}, \mathbf{u})]$, where $\mathbf{y} = \{y^k, k \in \mathcal{N}\}$.

This dynamic team can be converted to a static team provided for every $t \in \mathcal{N}$, there exists a function f_t for all S :

$$P(y^t \in S | \omega, u^1, \dots, u^{t-1}) = \int_S f_t(\omega, u^1, u^2, \dots, u^{t-1}, y^t) Q_t(dy^t).$$

Witsenhausen's equivalent model and static reduction of sequential dynamic teams

We can then write

$$P(d\omega, d\mathbf{y}) = P(d\omega) \prod_{t=1}^N f_t(\omega_0, u^1, u^2, \dots, u^{t-1}, y^t) Q_t(dy^t).$$

The cost function $J(\underline{\gamma})$ can then be written as

$$J(\underline{\gamma}) = \int P(d\omega) \prod_{t=1}^N (f_t(y_t, \omega_0, u^1, u^2, \dots, u^{t-1}, y^t) Q_t(dy^t)) c(\omega, \mathbf{y}, \mathbf{u}), \quad (2)$$

where now the measurement variables can be regarded as **independent** and by incorporating the $\{f_t\}$ terms into c , we can obtain an equivalent *static team* problem. Hence, the essential step is to appropriately adjust the probability space and the cost function.

Stochastic Dynamic Team Problems

- Review of Information Structures in Decentralized Control
- **Existence and Structural Properties**
- Convexity Properties
- Approximation of Team Problems and Asymptotic Optimality of Finite Representations
- Witsenhausen's Counterexample: Non-convexity, Existence and Approximations

Strategic Measures, Convexity Properties, and Optimal Solutions

We can view a measurable policy as a special case of randomized policies. This interpretation has many useful properties, one being the topological use of the space of probability measures.

For stochastic control problems, *strategic measures* (Schäl'75, Dynkin-Yushkevich'79, Feinberg'96) are defined as the set of probability measures induced by admissible control policies.

In the following, we discuss the case for stochastic team problems.

Strategic Measures, Convexity Properties, and Optimal Solutions

We can view a measurable policy as a special case of randomized policies. This interpretation has many useful properties, one being the topological use of the space of probability measures.

For stochastic control problems, *strategic measures* (Schäl'75, Dynkin-Yushkevich'79, Feinberg'96) are defined as the set of probability measures induced by admissible control policies.

In the following, we discuss the case for stochastic team problems.

Strategic Measures, Convexity Properties, and Optimal Solutions

Let $L_A(\mu)$ be the set of strategic measures induced by all admissible team policies with $(\omega, \mathbf{y}) \sim \mu$. In the following, $B = B^0 \times \prod_k B^k$ are used to denote all the Borel sets in $\Omega \times \prod_k \mathbb{U}^k$,

$$L_A(\mu) := \left\{ P \in \mathcal{P} \left(\Omega \times \prod_{k=1}^N (\mathbb{Y}^k \times \mathbb{U}^k) \right) : \right. \\ \left. P(B) = \int \mu(d\omega, d\mathbf{y}) \prod_k 1_{\{u^k = \gamma^k(y^k) \in B^k\}}, \gamma^k \in \Gamma^k \right\} \quad (3)$$

Strategic Measures, Convexity Properties, and Optimal Solutions

Let $L_R(\mu)$ be the set of strategic measures induced by all admissible team policies with $\omega, \mathbf{y} \sim \mu$ with individually randomized policies (that is, with independent randomizations):

$$L_R(\mu) := \left\{ P \in \mathcal{P} \left(\Omega \times \prod_{k=1}^N (\mathbb{Y}^k \times \mathbb{U}^k) \right) : P(B) = \int \mu(d\omega, d\mathbf{y}) \prod_k \Pi^k(du^k | y^k) \right\}$$

where Π^k takes place from the set of stochastic kernels from \mathbb{Y}^k to \mathbb{U}^k for each k .

Strategic Measures, Convexity Properties, and Optimal Solutions

Consider $\Upsilon = [0, 1]^N$ and let

$$L_C(\mu) := \left\{ P \in \mathcal{P} \left(\Omega \times \prod_{k=1}^N (\mathbb{Y}^k \times \mathbb{U}^k) \right) : P(\mathbf{B}) = \int \eta(dz) L_A(\mu, \underline{\gamma}(z)), \right. \\ \left. \eta \in \mathcal{P}(\Upsilon) \right\}$$

Here, $\underline{\gamma}(z)$ denotes a collection of team policies measurably parametrized by $z \in \Upsilon$.

Strategic Measures, Convexity Properties, and Optimal Solutions

Finally, let L_{CR} denote the set of strategic measures that are induced by some common randomness and arbitrary independent randomness.

$$L_{CR}(\mu) := \left\{ P \in \mathcal{P} \left(\Omega \times \prod_{k=1}^N (\mathbb{Y}^k \times \mathbb{U}^k) \right) : \right. \\ \left. P(B) = \int \eta(dz) \mu(d\omega, d\mathbf{y}) \prod_k \Pi^k(du^k | y^k, z) \right\}$$

Strategic Measures for Static Teams and Optimality of Deterministic Policies

Theorem

- (i) L_R has the following representation

$$L_R(\mu) = \left\{ P \in \mathcal{P} \left(\Omega \times \prod_{k=1}^N (\mathbb{Y}^k \times \mathbb{U}^k) \right) : P(B) = \int U(dz) L_A(\mu, \underline{\gamma}(z)), \right. \\ \left. U \in \mathcal{P}(\Upsilon), U(dw_1, \dots, dw_N) = \prod_s \eta_k(dw_k), \eta_k \in \mathcal{P}([0, 1]) \right\}, \quad (4)$$

that is U is constructed by the product of N independent random variables.

- (ii) $L_C(\mu)$ is convex. Its extreme points form $L_A(\mu)$. The sets L_R and L_{CR} are not convex for general sequential teams.
- $\inf_{\underline{\gamma} \in \Gamma} J(\underline{\gamma}) = \inf_{P \in L_A(\mu)} \int P(ds) c(s) = \inf_{P \in L_R(\mu)} \int P(ds) c(s) = \inf_{P \in L_C(\mu)} \int P(ds) c(s)$. **Deterministic policies are optimal among all.**

Strategic measures for dynamic teams

Theorem

- (i) $L_R(\mu)$ has the following representation so that for any $P \in L_R(\mu)$,

$$P(B) = \int U(dz)L_A(\mu, \underline{\gamma}(z))(B)$$
$$U(dv_1, \dots, dv_N) = \prod_s \eta_s(dv_s), \eta_s \in \mathcal{P}([0, 1])\}, (5)$$

where η_s is the Lebesgue measure on $[0, 1]$ and $\underline{\gamma}(z)$ is a collection of deterministic policies parametrized by z .

- (ii)

$$\inf_{\underline{\gamma} \in \Gamma} J(\underline{\gamma}) = \inf_{P \in L_A(\mu)} \int P(ds)c(s) = \inf_{P \in L_R(\mu)} \int P(ds)c(s)$$

In particular, deterministic policies are optimal among the randomized class.

Convexity of sets of Strategic Measures

Theorem

- *If the sequential team is not classical (and not necessarily non-classical), the set of strategic measures is not convex.*
- *If the information structure is classical, and if randomized policies are allowed so that, DM i has access to $y^k, u^k, k < i$ and y^i , then the set of strategic measures is convex.*

Existence of optimal team policies

Establishing the existence and structure of optimal policies is a challenging problem.

More specific setups and non-existence results have been studied in Witsenhausen'69, Wu-Verdu'11, Y.-Linder'12.

Considering the set of randomized strategic measures and convexification of these measures allow for placing a useful topology, that of weak convergence of probability measures, on the strategy spaces.

Existence of optimal team policies

Theorem

(i) Consider a static or dynamic team. Let the loss function c be lower semi-continuous in x, \mathbf{u} and $L_R(\mu)$ be a compact subset under weak topology. Then, there exists an optimal team policy. This policy is deterministic and induces a strategic measure in L_A .

(ii) Consider a static team or the static reduction of a dynamic team with c denoting the loss function. Let c be lower semi-continuous in x, \mathbf{u} and $L_C(\mu)$ be a compact subset under weak topology. Then, there exists an optimal team policy. This policy is deterministic and induces a strategic measure in L_A .

Sufficient conditions for existence of optimal policies

Theorem (Gupta, Y., Basar, Langbort'15)

Consider a static team where the action sets $\mathbb{U}^i, i \in \mathcal{N}$ are compact.

Furthermore, if the measurements satisfy:

$$P(d\mathbf{y}|\omega_0) = \prod_{i=1}^n Q^i(dy^i|\omega_0),$$

where $Q^i(dy^i|\omega_0) = \eta^i(y^i, \omega_0)\nu^i(dy^i)$ for some measure ν and continuous η that satisfy that for every $\epsilon > 0, \exists \delta$ such that if $d(a, b) < \delta$:

$|\eta^i(b, \omega_0) - \eta^i(a, \omega_0)| \leq \epsilon h^i(a, \omega_0)$, with $\sup_{\omega_0} \int h^i(a, \omega_0)\nu^i(dy^i) < \infty$, and $c(\omega_0, \mathbf{u})$ is continuous, then the set $L_R(\mu)$ is weakly compact and there exists an optimal team policy (which is deterministic and hence in $L_A(\mu)$).

Existence of optimal team policies: Proof Sketch

The existence result also applies to static reductions for sequential dynamic teams, and a class of teams with unbounded cost functions and non-compact action spaces.

The issue is the closedness property of the set of strategic measures achieved by independent randomization: A sequence of conditionally probability measures may converge to a limit which is not conditionally independent.

The proof builds on the fact that, conditioned on the channel properties, a weak limit of a sequence of joint probability measures that satisfies condition independent properties is also conditionally independent.

Example for a channel which satisfies the desired continuity properties is the additive Gaussian channel: $y^k = \omega_0 + v^k$.

Existence of optimal team policies: Proof Sketch

The existence result also applies to static reductions for sequential dynamic teams, and a class of teams with unbounded cost functions and non-compact action spaces.

The issue is the closedness property of the set of strategic measures achieved by independent randomization: A sequence of conditionally probability measures may converge to a limit which is not conditionally independent.

The proof builds on the fact that, conditioned on the channel properties, a weak limit of a sequence of joint probability measures that satisfies condition independent properties is also conditionally independent.

Example for a channel which satisfies the desired continuity properties is the additive Gaussian channel: $y^k = \omega_0 + v^k$.

Existence of optimal team policies: Proof Sketch

The existence result also applies to static reductions for sequential dynamic teams, and a class of teams with unbounded cost functions and non-compact action spaces.

The issue is the closedness property of the set of strategic measures achieved by independent randomization: A sequence of conditionally probability measures may converge to a limit which is not conditionally independent.

The proof builds on the fact that, conditioned on the channel properties, a weak limit of a sequence of joint probability measures that satisfies condition independent properties is also conditionally independent.

Example for a channel which satisfies the desired continuity properties is the additive Gaussian channel: $y^k = \omega_0 + v^k$.

Existence of optimal team policies: Proof Sketch

The existence result also applies to static reductions for sequential dynamic teams, and a class of teams with unbounded cost functions and non-compact action spaces.

The issue is the closedness property of the set of strategic measures achieved by independent randomization: A sequence of conditionally probability measures may converge to a limit which is not conditionally independent.

The proof builds on the fact that, conditioned on the channel properties, a weak limit of a sequence of joint probability measures that satisfies condition independent properties is also conditionally independent.

Example for a channel which satisfies the desired continuity properties is the additive Gaussian channel: $y^k = \omega_0 + v^k$.

Dynamic Teams and Witsenhausen's Counterexample


The existence result applies for Witsenhausen's counterexample.

It is known that classical LQG team problems admit solutions which are linear.

[Witsenhausen'68] showed that when there are measurability and information constraints leading to a non-classical information structure, this result is no longer true.

[Witsenhausen'68]: Even LQG problems may admits solutions which are non-linear.

Witsenhausen's Counterexample



$$y_0 = x_0, \quad u_0 = \mu_0(y_0), \quad x_1 = x_0 + u_0,$$

$$y_1 = x_1 + w_1, \quad u_1 = \mu_1(y_1), \quad x_2 = x_1 + u_1.$$

The goal is to minimize the expected performance index for some $k > 0$

$$Q_W(x, u_0, u_1) = k(u_0)^2 + x_2^2$$

Witsenhausen's Counterexample

This is the celebrated Witsenhausen's counterexample.

It is described by a linear system; all primitive variables are Gaussian.

Yet optimal team policy is non-linear [Witsenhausen'68].

Witsenhausen established that a solution exists ([Wu-Verdú'11] provided an alternative proof using Transport Theory), and established that an optimal policy is non-linear.

Witsenhausen's Counterexample

Suppose x and w_1 are two independent, zero-mean Gaussian random variables with variance σ^2 and 1. An equivalent representation is:

$$u_0 = \gamma_0(x), \quad u_1 = \gamma_1(u_0 + w).$$

$$Q_W(x, u_0, u_1) = k(u_0 - x)^2 + (u_1 - u_0)^2, \quad (6)$$




Figure: Flow of information in Witsenhausen's counterexample.

Witsenhausen's Counterexample

Now consider a different choice for Q :

$$Q_{\text{TC}}(x, u_0, u_1) = k(u_0)^2 + (u_1 - x)^2, \quad (7)$$

where again $k > 0$.




Figure: Flow of information in Witsenhausen's counterexample.

Witsenhausen's Counterexample

The version of this problem where the soft constraint is replaced by a hard power constraint, $E[(u_0)^2] \leq k$, is known as the *Gaussian Test Channel* (GTC).

In this context γ_0 is the *encoder* and γ_1 the *decoder*, where the latter's optimal choice is the conditional mean of x given y , that is $E[x|y]$.

The best encoder for the GTC can be shown to be linear (a scaled version of the source output, x), which in turn leads to a linear optimal decoder.

The approach here is through information theoretic arguments [Goblick'65][Berger'71].

Witsenhausen's Counterexample

Now, consider the more general version of (7):

$$Q_{\text{GTC}}(x, u_0, u_1) = k(u_0)^2 + (u_1 - x)^2 + b_0 u_0 x, \quad (8)$$

where b_0 is a scalar. In this case, an optimal solution is linear [Bansal-Bacsar'87]. The difference between (8) and Witsenhausen's problem is that Q in the former has a product term between the decision rules of the two agents while here it does not.




Figure: Flow of information in Witsenhausen's counterexample.

Witsenhausen's Counterexample

Hence, it is not only the nonclassical nature of the information structure but also the structure of the performance index that determines whether linear policies are optimal in these quadratic dynamic decision problems with Gaussian statistics and nonclassical information.

Furthermore, the noise distribution is also crucial: If the noise variables are discrete, it can be shown that the Witsenhausen's counterexample does not admit an optimal solution [Y.-Basar'13].

Witsenhausen's Counterexample

Hence, it is not only the nonclassical nature of the information structure but also the structure of the performance index that determines whether linear policies are optimal in these quadratic dynamic decision problems with Gaussian statistics and nonclassical information.

Furthermore, the noise distribution is also crucial: If the noise variables are discrete, it can be shown that the Witsenhausen's counterexample does not admit an optimal solution [Y.-Basar'13].

Witsenhausen's Counterexample

The static reduction of the Witsenhausen's counterexample is a two controller static team where the observations y^1 and y^2 of the two controllers are *independent* zero-mean Gaussian random variables.

The control laws γ^1 and γ^2 are to be chosen to minimize

$$J(\gamma^1, \gamma^2) = E[(y^1 + u^1 - u^2)^2 + (ku^1)^2 e^{(y^1 + u^1)(2y^2 - y^1 - u^1)/2}]$$

Theorem

The Witsenhausen's counterexample admits an optimal solution.

Also applies to: The LQG problem, the output feedback control problem, the relay channel problem etc.

Stochastic Dynamic Team Problems

- Review of Information Structures in Decentralized Control
- Existence and Structural Properties
- Convexity Properties
- Approximation of Team Problems and Asymptotic Optimality of Finite Representations
- Witsenhausen's Counterexample: Non-convexity, Existence and Approximations

Convexity of Static Team Problems

Definition

A (static or dynamic) team problem is convex on Γ if $J(\underline{\gamma}) < \infty$ for all $\underline{\gamma} \in \Gamma$ and for any $\alpha \in (0, 1), \underline{\gamma}_1, \underline{\gamma}_2 \in \Gamma$:

$$J(\alpha\underline{\gamma}_1 + (1 - \alpha)\underline{\gamma}_2) \leq \alpha J(\underline{\gamma}_1) + (1 - \alpha)J(\underline{\gamma}_2)$$

Theorem

Consider a static team. $J(\underline{\gamma})$ is convex if $c(\omega, \mathbf{u})$ is convex in \mathbf{u} provided that $J(\underline{\gamma}) < \infty$ for all $\underline{\gamma} \in \Gamma$.

Convexity of Static Team Problems

The *join* of two σ -fields over some set \mathbb{X} is the coarsest σ -field containing both. The *meet* of two σ -fields is the finest σ -field which is a subset of both. In addition, let \mathcal{F}_j be the *join* of the σ -field, that is $\mathcal{F}_j = \bigcup_k \mathcal{F}^k$.

Let \mathcal{F}^i be the σ -field generated by η^i and let $\mathcal{F}_c = \bigcap_k \mathcal{F}^k$ be the meet of these fields, this is termed as *common knowledge* by Aumann'76 for finite probability spaces.

Convexity of Static Team Problems

The *join* of two σ -fields over some set \mathbb{X} is the coarsest σ -field containing both. The *meet* of two σ -fields is the finest σ -field which is a subset of both. In addition, let \mathcal{F}_j be the *join* of the σ -field, that is $\mathcal{F}_j = \bigcup_k \mathcal{F}^k$.

Let \mathcal{F}^i be the σ -field generated by η^i and let $\mathcal{F}_c = \bigcap_k \mathcal{F}^k$ be the meet of these fields, this is termed as *common knowledge* by Aumann'76 for finite probability spaces.

Convexity of Static Team Problems

Theorem

(i) If a team problem is convex, then $E[c(\omega, \mathbf{u})|\mathcal{F}_c]$ is convex in u almost surely.

(ii) If

$$E[c(\omega, \mathbf{u})|\mathcal{F}_j]$$

is convex in u almost surely, then the team problem is convex on the set of team policies that satisfy $J(\underline{\gamma}) < \infty$.

A generalization of Radner and Krainak et. al.'s theorems

Theorem

Let $\{J; \Gamma^i, i \in \mathcal{I}\}$ be a static stochastic team problem where $U^i \equiv \mathbb{R}^{m_i}, i \in \mathcal{I}$, the loss function $E[L(\xi, \cdot) | \mathcal{F}_j]$ is convex and continuously differentiable in almost surely, and $J(\underline{\gamma})$ is bounded from below on \mathcal{G} . Let $\underline{\gamma}^*$ be a policy N -tuple with a finite cost ($J(\underline{\gamma}^*) < \infty$), and suppose that for every $\underline{\gamma} \in \mathcal{G}$ such that $J(\underline{\gamma}) < \infty$, the following holds:

$$\sum_{i \in \mathcal{I}} E\{\nabla_{u^i} c(\omega; \underline{\gamma}^*(\mathbf{y}))[\gamma^i(y^i) - \gamma^{i*}(y^i)]\} \geq 0, \quad (9)$$

Then, $\underline{\gamma}^*$ is a team-optimal policy, and it is unique if $E[c(\omega, \cdot) | \mathcal{F}_j]$ is strictly convex in almost surely.

A generalization of Radner and Krainak et. al.'s theorems

- (c.1) For all $\underline{\gamma} \in \Gamma$ such that $J(\underline{\gamma}) < \infty$, the following random variables have well-defined (finite) expectations (i.e., mean values)

$$\nabla_{u^i} c(\omega; \underline{\gamma}^*(\cdot))[\gamma^i(y^i) - \gamma^{i*}(y^i)], \quad i \in \mathcal{N}$$

- (c.2) Γ^i is a Hilbert space for each $i \in \mathcal{N}$, and $J(\underline{\gamma}) < \infty$ for all $\underline{\gamma} \in \Gamma$.
Furthermore,

$$E_{\xi|y^i} \{ \nabla_{u^i} c(\omega; \underline{\gamma}^*(\mathbf{y})) \} \in \Gamma^i, \quad i \in \mathcal{N}.$$

Theorem

Let $\{J; \Gamma^i, i \in \mathcal{N}\}$ be a static stochastic team problem which satisfies all the hypotheses of the previous theorem, but instead of (9), let either (c.1) or (c.2) be satisfied. Then, if $\underline{\gamma}^$ is a pbpo policy it is also team optimal. Such a policy is unique if $E[c(\omega; \cdot) | \mathcal{F}_j]$ is strictly convex in \mathbf{u} , a.s.*

Convexity of Sequential Dynamic Teams

The static reduction of a sequential dynamic team problem, if exists, is not unique. However, the following holds.

Theorem

A stochastic dynamic team problem with a static reduction is convex if and only if its static reduction is.

Non-convexity of Witsenhausen's Counterexample

Consider the celebrated Witsenhausen's counterexample: This is a dynamic non-classical team problem with y^1 and w^1 zero-mean independent Gaussian random variables with unit variance and $u^1 = \gamma^1(y^1)$, $u^2 = \gamma^2(u^1 + w^1)$ and the cost function $c(\omega, u^1, u^2) = k^2(y^1 - u^1)^2 + (u^1 - u^2)^2$ for some $k > 0$: The static reduction proceeds as follows, with $\eta(x) = \frac{1}{\sqrt{2\pi}}e^{-x^2/2}$:

$$\begin{aligned} & \int (k(u^1 - y^1)^2 + (u^1 - u^2)^2) Q(dy^1) \gamma^1(du^1|y^1) \gamma_1(du^2|y^2) P(dy^2|u^1) \\ &= \int (k(u^1 - y^1)^2 + (u^1 - u^2)^2) Q(dy^1) \gamma^1(du^1|y^1) \gamma_1(du^2|y^2) \eta(y^2 - u^1) dy^2 \\ &= \int \left((ku_0^2 + (u_0 - u_1)^2) \gamma^1(du^1|y^1) \gamma_1(du^2|y^2) \frac{\eta(y^2 - u^1) dy^2}{\eta(y^2)} \right) Q(dy^1) Q(dy^2) \end{aligned}$$

where Q denotes a Gaussian measure with zero mean and unit variance and η its density.

Quasi-classical information structures: Reduction through information equivalence

An *IS* is partially nested if an agent's information at a particular stage t can depend on the action of some other agent at some stage $t' \leq t$ only if she also has access to the information of that agent at stage t' .

Partially nested information structures include the cases where explicit information exchange in a decentralized system among decision makers is faster than information propagation through system dynamics.

Theorem

Consider a partially nested stochastic dynamic team with a convex cost function. The team problem is convex.

Convexity of Sequential Dynamic Teams

Ho and Chu established this result for the special setup involving the partially nested LQG teams. In this case, optimal policies are linear through an equivalence to static teams: Consider the following dynamic team with N DMs, with DM k having the following measurement

$$y^k = C^k \xi + \sum_{i:i \rightarrow k} D_{ik} u^i, \quad (10)$$

where ξ is an exogenous random variable picked by nature, and $i \rightarrow k$ denotes the precedence relation that the action of DM i affects the information of DM k and u^i is the action of DM i .

Quasi-classical information structures: Reduction through information equivalence

If the information structure is quasi-classical, then the information available to DM k , \mathcal{I}^k , can be represented with:

$$\mathcal{I}^k = \{y^k, \{\mathcal{I}^i, i \rightarrow k\}\}.$$

That is, DM k has access to the information available to all the signaling agents. Such an IS is equivalent to the IS $\mathcal{I}^k = \{\tilde{y}^k\}$, where \tilde{y}^k is a static measurement given by

$$\tilde{y}^k = \left\{ C^k \xi, \{C^i \xi, i \rightarrow k\} \right\}. \quad (11)$$

Such a conversion can be done provided that the policies adopted by the agents are deterministic.

Stochastic partial nestedness: A probabilistic definition of nestedness, its relation to convexity and signaling

When the information structure is non-classical or not quasi-classical, the decision makers may use their actions to communicate with each other. This phenomenon is known as signalling.

When signaling is present, the problem has a communications flavour and any communication problem is inherently non-convex.

It is known that quasi-classical information structures eliminate the incentive for signaling, since the future decision makers already have access to the information at the signaling decision maker.

Stochastic partial nestedness: A probabilistic definition of nestedness, its relation to convexity and signaling

In the following, we exhibit that the static reduction provides an effective method to identify when lack of a signaling incentive can be established and perhaps can lead to a more refined probability and information structure dependent characterization of *nestedness*, that generalizes partial nestedness.

Definition

The information structure of a sequential team problem is stochastically partially nested, if for an arbitrary cost function $c : \Omega \times \prod_k \mathbb{U}^k \rightarrow \mathbb{R}$ there exists a static reduction of this team which does not alter the loss function.

Stochastic partial nestedness: A probabilistic definition of nestedness, its relation to convexity and signaling

This definition implies the following result.

Lemma

Consider a sequential team problem with a stochastically partially nested information structure. If the cost function $c(\omega, \mathbf{u})$ is convex in \mathbf{u} , then the team problem is convex.

Proof. The static reduction of this team preserves convexity of the loss function, for an arbitrary convex loss function $l : \Omega \times \prod_k \mathcal{U}^k \rightarrow \mathbb{R}$. Thus, the reduced problem, and hence the original problem is convex. \diamond

Stochastic partial nestedness: A probabilistic definition of nestedness, its relation to convexity and signaling

Example

$$x_{t+1}^1 = a_1 x_t^1 + u_t^1 + w_t^1, \quad x_{t+1}^2 = a_2 x_t^2 + u_t^2 + w_t^2$$

$$x_{t+1}^3 = a_3 x_t^3 + u_t^1 + u_t^2 + w_t^3$$

$$y_t^1 = (x_t^1 + v_t^1, x_t^2 + v_t^2 + v_t^{21}, x_t^3 + v_t^{31})$$

$$y_t^2 = (x_t^1 + v_t^1 + v_t^{12}, x_t^2 + v_t^2, x_t^3 + v_t^{32}),$$

$$J = E \left[\sum_{t=0}^{T-1} \left((x_t^1)^2 + (x_t^2)^2 + \rho_1 (u_t^1)^2 + \rho_2 (u_t^2)^2 \right) \right],$$

with $\rho_1, \rho_2 > 0$. Measurements are: $I_t^i = \{y_t^i, I_{t-1}^i\}$, with $I_0^i = y_0^i$. This system is non-classical. But, an optimal team policy is linear.

Stochastic Dynamic Team Problems

- Review of Information Structures in Decentralized Control
- Existence and Structural Properties
- Convexity Properties
- Approximation of Team Problems and Asymptotic Optimality of Finite Representations
- Witsenhausen's Counterexample: Non-convexity, Existence and Approximations

Asymptotic Optimality of Finite Models in Stochastic Control

The approximation result builds on a sequence of recent studies on identifying conditions on *when a finite model can be used to construct approximately optimal policies* for a Markov Decision Problem with Borel state and action spaces [Saldi, Y., Linder'13,'14,'15].

Conditions on the transition kernels: Weak continuity, setwise continuity or total variation continuity

Conditions on cost functions: Lipschitz continuity.

Discounted cost vs. average cost: Recurrence conditions

It turns out that the results are applicable to team problems, leading to the following results.

Approximation of Team Problems and Optimal Solutions

Consider an N -agent static team problem in which DM i , $i = 1, \dots, N$, observes a random variable y^i and takes an action u^i .

Given any state realization x , the random variable y^i has a distribution $W^i(\cdot | x)$; that is, $W^i(\cdot | x)$ is a stochastic kernel on Y^i given X .

Approximation of Static Team Problems

The team cost function c is a non-negative function of the state, observations, and actions; that is, $c : \mathbf{X} \times \mathbf{Y} \times \mathbf{U} \rightarrow [0, \infty)$, where $\mathbf{Y} := \prod_{i=1}^N \mathbf{Y}^i$ and $\mathbf{U} := \prod_{i=1}^N \mathbf{U}^i$. For Agent i , the set of strategies Γ^i is given by

$$\Gamma^i := \{ \gamma^i : \mathbf{Y}^i \rightarrow \mathbf{U}^i, \gamma^i \text{ is measurable} \}.$$

Recall that $\Gamma = \prod_{i=1}^N \Gamma^i$. Then, the cost of the team $J : \Gamma \rightarrow [0, \infty)$ is given by

$$J(\underline{\gamma}) = \int_{\mathbf{X} \times \mathbf{Y}} c(x, \mathbf{y}, \mathbf{u}) \mathbb{P}(dx, d\mathbf{y}),$$

where $\mathbf{u} = \underline{\gamma}(\mathbf{y})$. Here, $\mathbb{P}(dx, d\mathbf{y}) := \mathbb{P}(dx) \prod_{i=1}^N W^i(dy^i|x)$ denotes the joint distribution of the state and observations. Therefore, we have

$$J^* = \inf_{\underline{\gamma} \in \Gamma} J(\underline{\gamma}^N).$$

Approximation of Static Team Problems

In this section, we impose the following assumptions.

Assumption

- (a) *The cost function c is bounded and continuous in \mathbf{u} .*
- (b) *For each i , \mathbf{U}^i is a convex subset of a locally convex vector space.*
- (c) *For each i , \mathbf{Y}^i is compact.*

Approximation of Static Team Problems

We first prove that the minimum cost achievable by continuous strategies is equal to the optimal cost J^* . To this end, for each i , we define

$$\Gamma_c^i := \{ \gamma^i \in \Gamma^i : \gamma^i \text{ is continuous} \} \text{ and } \Gamma_c := \prod_{i=1}^N \Gamma_c^i.$$

Proposition

We have

$$\inf_{\underline{\gamma} \in \Gamma_c} J(\underline{\gamma}) = J^*.$$

Approximation of Static Team Problems

Let d_i denote the metric on Y^i . Since Y^i is compact, one can find a finite set $Y^{n,i} := \{y_{i,1}, \dots, y_{i,i_n}\} \subset Y^i$ such that $Y^{n,i}$ is an $1/n$ -net in Y^i ; that is, for any $y \in Y^i$ we have

$$\min_{z \in Y^{n,i}} d_i(y, z) < \frac{1}{n}.$$

Define function $Q_{n,i}$ mapping Y^i to $Y^{n,i}$ by

$$Q_{n,i}(y) := \arg \min_{z \in Y^{n,i}} d_i(y, z).$$

For each n , $Q_{n,i}$ induces a partition $\{\mathcal{S}_{i,j}\}_{j=1}^{i_n}$ of Y^i given by

$$\mathcal{S}_{i,j} := \{y \in Y^i : Q_{n,i}(y) = y_{i,j}\}.$$

For any $\gamma^i \in \Gamma^i$, we let $\gamma^{n,i}$ denote the strategy $\gamma^i \circ Q_{n,i}$.

Approximation of Static Team Problems

Define

$$\Gamma^{n,i} := \{\gamma^i \in \Gamma^i : \gamma^i \text{ is constant on each } \mathcal{S}_{i,j}\}$$

and so, $\gamma^{n,i} \in \Gamma^{n,i}$ for each $\gamma^i \in \Gamma^i$. We let $\Gamma_n := \prod_{i=1}^N \Gamma^{n,i}$. The following theorem states that optimal policy $\underline{\gamma}^*$ can be approximated by policies in Γ_n .

Theorem

We have

$$\lim_{n \rightarrow \infty} \inf_{\underline{\gamma} \in \Gamma_n} J(\underline{\gamma}) = J^*.$$

Approximation of Static Team Problems

For each n , define stochastic kernels $W^{n,i}(\cdot | x)$ on $Y^{n,i}$ given X as follows:

$$W^{n,i}(\cdot | x) := \sum_{j=1}^{i_n} W(\mathcal{S}_{i,j}|x)\delta_{y_{i,j}}(\cdot).$$

Let $\Pi^{n,i} := \{\pi^i : Y^{n,i} \rightarrow \mathbf{U}^i, \pi^i \text{ measurable}\}$ and $\Pi_n := \prod_{i=1}^N \Pi^{n,i}$. Define $J_n : \Pi_n \rightarrow [0, \infty)$ as

$$J_n(\underline{\pi}) := \int_{X \times Y_n} c(x, \mathbf{y}, \mathbf{u}) \mathbb{P}_n(dx, d\mathbf{y}),$$

where $\underline{\pi} = (\pi^1, \dots, \pi^N)$, $\mathbf{u} = \underline{\pi}(\mathbf{y})$, $Y_n = \prod_{i=1}^N Y^{n,i}$, and $\mathbb{P}_n(dx, d\mathbf{y}) = \mathbb{P}(dx) \prod_{i=1}^N W^{n,i}(dy^i | x)$.

Approximation of Static Team Problems

Theorem

For any $\varepsilon > 0$, there exists a sufficiently large n such that the optimal (or almost optimal) policy $\underline{\pi}^ \in \Pi_n$ for the cost J_n is ε -optimal for the original team problem when $\underline{\pi}^* = (\pi^{1*}, \dots, \pi^{N*})$ is extended to Υ via $\gamma^i = \pi^{i*} \circ Q_{n,i}$.*

Approximation of Dynamic Team Problems

Theorem

Suppose that a static reduction exists, the cost is continuous, and $f_i(w_0, \mathbf{u}^{i-1}, y^i)$ is continuous in \mathbf{u}^{i-1} for $i = 1, \dots, N$. Then, the static reduction of the dynamic team model satisfies the existence conditions.

Observe that neither the Witsenhausen's counterexample nor the point-to-point communication problem satisfy the compactness condition. In the following, we discuss this important setting.

Stochastic Dynamic Team Problems

- Review of Information Structures in Decentralized Control
- Existence and Structural Properties
- Convexity Properties
- Approximation of Team Problems and Asymptotic Optimality of Finite Representations
- Witsenhausen's Counterexample: Non-convexity, Existence and Approximations

Approximation of Witsenhausen's Counterexample: Prior Work

- Lee-Lao-Ho, TAC'01
- Baglietto-Parisini-Zoppoli, TAC'01
- Li-Marden-Shamma, CDC'09
- Gnecco-Sanguinetti, INOC'09, OL'12
- Gnecco-Sanguinetti-Gaggero, SICON'12
- McEneaney-Han, Automatica'15

The following question has not been answered, to our knowledge?: Does there exist a computational scheme that would generate policies with costs arbitrarily close to optimum? What is the (optimal) value of the Witsenhausen counterexample?

Approximation of Witsenhausen's Counterexample

Recall that we have two agents. Agent 1 observes a zero mean and unit variance Gaussian random variable y^1 and decides its strategy u^1 , and Agent 2 observes $y^2 = u^1 + v$ and decides its strategy u^2 . Here, v is a zero mean and unit variance Gaussian noise independent of y^1 . The cost function of the team is given by

$$c(y^1, u^1, u^2) = l(u^1 - y^1)^2 + (u^2 - u^1)^2.$$

It was shown earlier that this problem can be reduced to a static team problem in which agents observe mutually independent zero mean and unit variance Gaussian random variables.

Approximation of Static Team Problems

Note that strategy spaces of the original problem and its static reduction are identical, and same strategies induce same team costs.

For any $k \in \mathbb{R}_+$, we let $K := [-k, k]$ and

$$\Gamma^{i,k} := \{\gamma^i \in \Gamma^i : \gamma^i(\mathbf{Y}^i) \subset K\},$$

where Γ^i denotes the strategy space of Agent i ; that is, the set of measurable functions from \mathbf{Y}^i to \mathbf{U}^i , where $\mathbf{Y}^i = \mathbf{U}^i = \mathbb{R}$ for $i = 1, 2$.

Approximation of Static Team Problems

Lemma

For any $\varepsilon > 0$, there exists $k \in \mathbb{R}_+$ such that

$$\inf_{(\gamma^1, \gamma^2) \in \Gamma^1 \times \Gamma^{2,k}} J(\gamma^1, \gamma^2) \leq \inf_{(\gamma^1, \gamma^2) \in \Gamma^1 \times \Gamma^2} J(\gamma^1, \gamma^2) + \varepsilon.$$

Recall that Γ_c^i denotes the set of continuous strategies of Agent i . Define $\Gamma_c^{i,k} := \Gamma^{i,k} \cap \Gamma_c^i$, for $i = 1, 2$.

Proposition

For any $k \in \mathbb{R}_+$, we have

$$\inf_{(\gamma^1, \gamma^2) \in \Gamma^1 \times \Gamma^{2,k}} J(\gamma^1, \gamma^2) = \inf_{(\gamma^1, \gamma^2) \in \Gamma_c^1 \times \Gamma_c^{2,k}} J(\gamma^1, \gamma^2).$$

Approximation of Witsenhausen's Counterexample

As a result, one can search for near optimal strategies for Witsenhausen's counterexample over the set $\Gamma_c^1 \times \Gamma_c^{2,k}$, for k sufficiently large.

We can show that any strategy in $\Gamma_c^1 \times \Gamma_c^{2,k}$ for arbitrary $k \in \mathbb{R}_+$ can be approximated with arbitrary precision by quantized strategies.

Fix any k . Let us choose $(\gamma^1, \gamma^2) \in \Gamma_c^1 \times \Gamma_c^{2,k}$ such that $J(\gamma^1, \gamma^2) < \infty$. Fix any $\delta > 0$. There exists $L = [-l, l]$ such that

$$\left| J(\gamma^1, \gamma^2) - \int_{L \times L} \tilde{c}(\gamma^1, y^1, \gamma^2, y^2) \mathbb{P}(dy^1) \mathbb{P}(dy^2) \right| < \frac{\delta}{2}. \quad (12)$$

Approximation of Witsenhausen's Counterexample

Let us quantize the interval L using a uniform quantizer, denoted as q , having $N(l)$ output levels; that is,

$$q : L \rightarrow \{y_1, \dots, y_{N(l)}\} \subset L$$

and

$$q^{-1}(y_j) = \left[y_j - \frac{\tau}{2}, y_j + \frac{\tau}{2} \right),$$

where $\tau = \frac{2L}{N(l)}$.

Define the quantized strategy $(\gamma^{1,q}, \gamma^{2,q})$ as follows:

$$\gamma^{i,q}(y^i) = \begin{cases} \gamma^i \circ q(y^i), & \text{if } y^i \in L \\ 0, & \text{otherwise.} \end{cases}$$

Approximation of Witsenhausen's Counterexample

To compute a near optimal policy for Witsenhausen's counterexample it is sufficient to choose a strategy based on the quantized observations

$(q(y^1), q(y^2))$ for sufficiently large $(l, N(l))$, where $q : L \rightarrow \{y_1, \dots, y_{N(l)}\}$ is extended to \mathbb{R} by mapping $\mathbb{R} \setminus L$ to 0.

In other words, for each $(l, N(l))$, let $Y_{l,N(l)}^i := \{0, y_1, \dots, y_{N(l)}\}$ (i.e., output levels of the extended q) and define probability measure $\mathbb{P}_{l,N(l)}$ on $Y_{l,N(l)}^i$ as

$$\mathbb{P}_{l,N(l)}(y_i) = \mathbb{P}(q^{-1}(y_i)). \quad (13)$$

Moreover, let $\Pi_{l,N(l)}^i := \{\pi^i : Y_{l,N(l)}^i \rightarrow \mathbf{U}_i, \pi^i \text{ measurable}\}$ and define

$$J_{l,N(l)}(\pi^1, \pi^2) := \sum_{j,i=0}^{N(l)} \tilde{c}(\pi^1, y_i, \pi^2, y_j) \mathbb{P}_{l,N(l)}(y_i) \mathbb{P}_{l,N(l)}(y_j).$$

Approximation of Witsenhausen's Counterexample: Asymptotic Optimality of Finite Representations

Theorem

For any $\varepsilon > 0$, there exists $(l, N(l))$ such that the optimal policy $(\pi^{1}, \pi^{2*}) \in \Pi_{l, N(l)}^1 \times \Pi_{l, N(l)}^2$ for the cost $J_{l, N(l)}$ is ε -optimal for Witsenhausen's counterexample when (π^{1*}, π^{2*}) is extended to $Y^1 \times Y^2$ via $\gamma^i = \pi^{i*} \circ q$, $i = 1, 2$. In particular, quantized policies are asymptotically optimal.*

In fact, the action space can also be quantized with an arbitrarily small loss: Thus, a numerical algorithm can be constructed so that a sequence of successively refined *finite models* can be obtained whose solution limit will lead to the value of Witsenhausen's counterexample.

Conclusion

- Review of Information Structures in Decentralized Control
- Existence and Structural Properties
- Convexity Properties
- Approximation of Team Problems and Asymptotic Optimality of Finite Representations
- Witsenhausen's Counterexample: Non-convexity, Existence and Approximations