Nash Equilibrium to Social Optimum and Back: A Mean Field Perspective

René Carmona

Department of Operations Research & Financial Engineering
PACM
Princeton University

Van Eenam Lecture III, University of Michigan

Models of Competition: Nash Equilibria

- Say player i takes **action** α^i,
- Their cost J^i depends upon the actions $\alpha^1, \cdots, \alpha^N$ of **ALL** the players

$$J^i = J^i(\alpha^1, \cdots, \alpha^N)$$

- A **strategy profile** $(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$ is a **Nash equilibrium** if for every i and feasible action α^i

$$J^i(\hat{\alpha}^1, \cdots, \hat{\alpha}^{i-1}, \hat{\alpha}^i, \hat{\alpha}^{i+1}, \cdots, \hat{\alpha}^N) \le J^i(\hat{\alpha}^1, \cdots, \hat{\alpha}^{i-1}, \alpha^i, \hat{\alpha}^{i+1}, \cdots, \hat{\alpha}^N)$$

  whatever $i = 1, \cdots, N$ is !
- In other words,

  the system is in a Nash equilibrium if any player trying to deviate from their action cannot end up better off !

- Not traditional minimization (not the typical *steady state found in physics*)
  - **Lack of Uniqueness**: when they do exist, they are often in large numbers, often a continuum
  - Why should a system **settle** in a Nash equilibrium? Which one?
- These equilibria capture a notion of **stability**

Model of Cooperation

Nash Equilibria vs Social Optimality

If agents take actions $\alpha^1, \cdots, \alpha^N$, **Social Cost** is defined as:

$$J^{SC}(\alpha^1, \cdots, \alpha^N) = \frac{1}{N}\left[J^1(\alpha^1, \cdots, \alpha^N) + \cdots + J^N(\alpha^1, \cdots, \alpha^N)\right]$$

- If $(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$ is a Nash Equilibrium (NE)

  $$J^{SC}(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$$

  is the (average) cost to the population for settling in the Nash Equilibrium

Model of Cooperation

Nash Equilibria vs Social Optimality

If agents take actions $\alpha^1, \cdots, \alpha^N$, **Social Cost** is defined as:

$$J^{SC}(\alpha^1, \cdots, \alpha^N) = \frac{1}{N}\left[J^1(\alpha^1, \cdots, \alpha^N) + \cdots + J^N(\alpha^1, \cdots, \alpha^N)\right]$$

- If $(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$ is a Nash Equilibrium (NE)

  $$J^{SC}(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$$

  is the (average) cost to the population for settling in the Nash Equilibrium
- A **Central Planner** could **minimize** the social cost and find

  $$(\alpha^{1*}, \cdots, \alpha^{N*}) = \arg \inf_{(\alpha^1, \cdots, \alpha^N)} J^{SC}(\alpha^1, \cdots, \alpha^N)$$

  $J^{SC}(\alpha^{1*}, \cdots, \alpha^{N*})$ is the minimal social cost ! **Unfortunately**, it is **rarely** a Nash equilibrium

Model of Cooperation

Nash Equilibria vs Social Optimality

If agents take actions $\alpha^1, \cdots, \alpha^N$, **Social Cost** is defined as:

$$J^{SC}(\alpha^1, \cdots, \alpha^N) = \frac{1}{N}\left[J^1(\alpha^1, \cdots, \alpha^N) + \cdots + J^N(\alpha^1, \cdots, \alpha^N)\right]$$

- If $(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$ is a Nash Equilibrium (NE)

  $$J^{SC}(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$$

  is the (average) cost to the population for settling in the Nash Equilibrium
- A **Central Planner** could **minimize** the social cost and find

  $$(\alpha^{1*}, \cdots, \alpha^{N*}) = \arg \inf_{(\alpha^1, \cdots, \alpha^N)} J^{SC}(\alpha^1, \cdots, \alpha^N)$$

  $J^{SC}(\alpha^{1*}, \cdots, \alpha^{N*})$ is the minimal social cost ! **Unfortunately**, it is **rarely** a Nash equilibrium
- How **bad / suboptimal** can Nash equilibria be?

$$\textbf{PoS} = \frac{inf_{(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)\ NE} J^{SC}(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)}{J^{SC}(\alpha^{1*}, \cdots, \alpha^{N*})}$$

quantifying how **much worse** the best Nash equilibrium is

Price of Stability

$$\textbf{PoA} = \frac{sup_{(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)\ NE} J^{SC}(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)}{J^{SC}(\alpha^{1*}, \cdots, \alpha^{N*})}$$

quantifying how **much worse** the worst Nash equilibrium is

Price of Anarchy

Paradoxes and Price of Anarchy

- Game Theory is **replete with paradoxes** !
- e.g. **Braess's paradox** in selfish routing (static one-period deterministic game)
- How bad can a Nash Equilibrium (NE) be when compared to alternative solutions?
- Introduction of terminology **Price of Anarchy** (PoA) by **Koutsoupias-Papadimitriou**
- **Explicit PoA Bounds** for Selfish Routing Games (**T. Roughgarden - E. Tardos**

Goal of this part of the lecture:

compare Social Welfare for NE to what a Central Planer could achieve

PoA Bounds for Continuous Time Games

References:

- PoA for Deterministic Linear Quadratic N-player Games (**Basar - Zhu**) (2010)
- Related ideas in **M. Huang**'s presentation in Rome
- Efficiency in MFGs (**Balandat - Tomlin** (2013), **Cardaliaguet - Rainer** (2018))

Explicit computations for **LQ MFGs R.C. - Graves - Tan** (2017))

Compute smallest possible **Social Cost** per individual:

- Have a **Central Planner** find a *common feedback control*
- to **minimize** the Expected Cost per individual to the system

MFG Model

- Let the individuals take care of their optimizations
- Hope for a Nash Equilibrium
- Compute the Expected Cost (per individual) to the system

How much worse is the cost due to the NE?

R.C. - Delarue

R.C. - Dayanikli - Delarue - Lauriere

The Model The basic state controlled equation is:

$$dX_t = \alpha_t dt + \sigma dW_t$$

and for each fixed flow of probability measures $\boldsymbol{\mu} = (\mu_t)_{0 \le t \le T}$ we define the cost:

$$J^{\boldsymbol{\mu}}(\boldsymbol{\alpha}) = \mathbb{E}\Big[\int_0^T f(t, X_t, \mu_t, \alpha_t)dt + g(X_T, \mu_T)\Big]$$

Assume that the running cost function f is of the form:

$$f(t, x, \mu, \alpha) = \frac{1}{2}|\alpha|^2 + f_0(x, \mu).$$

Hamiltonian

$$H(t, x, \mu, y, \alpha) = \alpha \cdot y + \frac{1}{2}|\alpha|^2 + f_0(x, \mu)$$

Minimizer

$$\hat{\alpha}(t, x, \mu, y) = \arg\min H(t, x, \mu, y, \alpha) = -y$$

Assumption for Initial Model

(I) f_0 and g are **continuously differentiable** with respect to x, and **differentiable** with respect to μ (in the sense of ∂_μ).

(II) For any $(x, \mu) \in \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$, there exists a version of

$$v \mapsto \partial_\mu f_0(x, \mu)(v) \quad \text{resp.} \quad v \mapsto \partial_\mu g(x, \mu)(v)$$

such that the mapping

$$(x, \mu, v) \mapsto \partial_\mu f_0(x, \mu)(v) \quad \text{resp.} \quad (x, \mu, v) \mapsto \partial_\mu g(x, \mu)(v)$$

is continuous.

(III) $\partial_x f_0$ are $\partial_x g$ are **Lipschitz continuous**

(IV) $\partial_\mu f_0$ and $\partial_\mu g$ are **Lipschitz continuous** in the following sense.

$$\mathbb{E}\left[\left|\partial_\mu f_0(x', \mu')(X') - \partial_\mu f_0(x, \mu)(X)\right|^2\right] \leq L\left[|x' - x|^2 + \mathbb{E}\left[|X' - X|^2\right]\right]$$

$$\mathbb{E}\left[\left|\partial_\mu g(x', \mu')(X') - \partial_\mu g(x, \mu)(X)\right|^2\right] \leq L\left[|x' - x|^2 + \mathbb{E}\left[|X' - X|^2\right]\right]$$

for all $x, x' \in \mathbb{R}^d$, $\alpha, \alpha' \in \mathbb{R}^d$, $\mu, \mu' \in \mathcal{P}_2(\mathbb{R}^d)$, and any $\mathbb{R}^d$-valued random variables X and X' having μ and μ' as distributions.

(V) The functions f_0 and g are **convex** in (x, μ), convexity with respect to the measure argument being understood in the **displacement convexity** sense,

The Mean Field Game (MFG) System

In equilibrium, the state process $\mathbf{X} = (X_t)_{0 \le t \le T}$ and the adjoint process $\mathbf{Y} = (Y_t)_{0 \le t \le T}$ solve the following FBSDE of the McKean-Vlasov type:

$$\begin{cases} dX_t &= -Y_t dt + \sigma dW_t \\ dY_t &= -\partial_x f_0(X_t, \mathcal{L}(X_t))dt + Z_t dW_t, \\ Y_T &= \partial_x g(X_T, \mathcal{L}(X_T)), \end{cases}$$

The equilibrium strategy is given by

$$\alpha_t = -Y_t, \qquad 0 \le t \le T.$$

The equilibrium cost to an individual is

$$J^{\mathrm{MFG}}(\boldsymbol{\alpha}) = \mathbb{E}\Big[\int_0^T f(t, X_t, \mu_t, \alpha_t)dt + g(X_T, \mu_T)\Big]$$

with $\mu_t = \mathcal{L}(X_t)$ for each $t \ge 0$.

The Central Planner Problem

Minimize the McKean-Vlasov cost defined for each strategy $\boldsymbol{\alpha} = (\alpha_t)_{0 \le t \le T}$ as:

$$J^{\mathrm{MFC}}(\boldsymbol{\alpha}) = \mathbb{E}\Big[\int_0^T f(t, X_t, \mathcal{L}(X_t), \alpha_t)dt + g(X_T, \mu_T)\Big]$$

where $\mathbf{X}$ satisfies $dX_t = \alpha_t dt + \sigma dW_t$.

Social cost (per individual) defined as:

$$J^{\mathrm{MFC}} = \inf_{\boldsymbol{\alpha}} J^{\mathrm{MFC}}(\boldsymbol{\alpha})$$

Notice that:

$$J^{\mathrm{MFC}} \le J^{\mathrm{MFG}}(\boldsymbol{\alpha})$$

for all MFG equilibria $\boldsymbol{\alpha} = (\alpha_t)_{0 \le t \le T}$.

Convenient notation for the central planner optimization problem:

$$\boldsymbol{\alpha}^{\mathrm{MFC}} = \arg\inf_{\boldsymbol{\alpha}} J^{\mathrm{MFC}}(\boldsymbol{\alpha})$$

Accordingly:

$$\boldsymbol{\mu}^{\mathrm{MFC}} = (\mu_t^{\mathrm{MFC}})_{0 \le t \le T} \quad \text{with} \quad \mu_t^{\mathrm{MFC}} = \mathcal{L}(X_t^{\mathrm{MFC}}) \quad \text{and} \quad dX_t^{\mathrm{MFC}} = \alpha_t^{\mathrm{MFC}} dt + \sigma dW_t$$

The Central Planner Problem (cont.)

Because the (reduced) Hamiltonian is given by:

$$H(t, x, \mu, y, \alpha) = \alpha y + \frac{1}{2}|\alpha|^2 + f_0(x, \mu)$$

the solution is now given by the FBSDE

$$\begin{cases} dX_t &= -Y_t dt + \sigma dW_t \\ dY_t &= -\Big(\partial_x f_0(X_t, \mathcal{L}(X_t))dt + \tilde{\mathbb{E}}[\partial_\mu f_0(\tilde{X}_t, \mathcal{L}(X_t))(X_t)]\Big) dt + Z_t dW_t, \\ Y_T &= \partial_x g(X_T, \mathcal{L}(X_T)) + \tilde{\mathbb{E}}[\partial_\mu g(\tilde{X}_T, \mathcal{L}(X_T))(X_T)], \end{cases}$$

which is **different** from FBSDE giving the solution of the MFG problem !!!!!

Coercing Players to Change their Behaviors

Question: *Can we incentivize the individuals in a MFG by perturbing the running cost they incur in such a way that they end up behaving (in terms of their strategy and actual state) exactly as if they were adopting the optimal strategy identified by a central planner optimizing the original social cost?*

For $\lambda \in [0, 1]$ define

$$f_\lambda(x, \mu) = f_0(x, \mu) + \lambda \tilde{\mathbb{E}}[\frac{\delta f_0}{\delta m}(\tilde{X}, \mu)(x)]$$

and consider the MFG with **the same controlled state dynamics**, running cost function:

$$\frac{1}{2}|\alpha|^2 + f_\lambda(x, \mu)$$

and terminal cost

$$g_\lambda(x, \mu) = g(x, \mu) + \lambda \tilde{\mathbb{E}}[\frac{\delta g}{\delta m}(\tilde{X}, \mu)(x)]$$

so that the equilibrium state dynamics are given by the forward component of the solution of the FBSDE:

$$\begin{cases} dX_t &= -Y_t dt + \sigma dW_t \\ dY_t &= -\partial_x f_\lambda(X_t, \mathcal{L}(X_t))dt + Z_t dW_t, \\ Y_T &= \partial_x g_\lambda(X_T, \mathcal{L}(X_T)), \end{cases}$$

Since the L-derivative and the functional derivatives are related by

$$\partial_x \frac{\delta f}{\delta m}(\mu)(x) = \partial_\mu f(\mu)(x)$$

the solution of this MFG coincides for $\lambda = 1$ with the solution central planner MFC optimization problem

λ-INTERPOLATED MFG

For a generic flow $\boldsymbol{\mu} := (\mu_t)_{0\le t\le T}$ from $[0,T]$ to $\mathcal{P}_2(\mathbb{R}^d)$, the cost $J^{\lambda,\mathrm{MF}}(\boldsymbol{\alpha};\boldsymbol{\mu})$:

$$\begin{aligned} J^{\lambda,\mathrm{MF}}(\boldsymbol{\alpha};\boldsymbol{\mu}) :=& (1-\lambda)J(\boldsymbol{\alpha};\boldsymbol{\mu}) + \lambda J^{\mathrm{MFC}}(\boldsymbol{\alpha}) \\ =& \mathbb{E}\left[\frac{1}{2}\int_0^T |\alpha_t|^2 dt + \int_0^T \Big[(1-\lambda) f_0\big(X_t^{\boldsymbol{\alpha}},\mu_t\big) + \lambda f_0\big(X_t^{\boldsymbol{\alpha}},\mathcal{L}(X_t^{\boldsymbol{\alpha}})\big)\Big] dt\right] \\ &+ \mathbb{E}\Big[(1-\lambda) g\big(X_T^{\boldsymbol{\alpha}},\mu_T\big) + \lambda g\big(X_T^{\boldsymbol{\alpha}},\mathcal{L}(X_T^{\boldsymbol{\alpha}})\big)\Big]. \end{aligned} \tag{1}$$

DEFINITION

For a given $\lambda \in [0,1]$, we say that a (square-integrable) control process $\boldsymbol{\alpha}^{\boldsymbol{\lambda}}$ induces a λ-interpolated mean field equilibrium if $\boldsymbol{\alpha}^{\boldsymbol{\lambda}}$ solves the minimization problem

$$\inf_{\boldsymbol{\alpha}} J^{\lambda,\mathrm{MF}}(\boldsymbol{\alpha};\boldsymbol{\mu}^{\lambda}),$$

where $\boldsymbol{\mu}^{\lambda} := (\mu_t^{\lambda})_{0\le t\le T} = \mathcal{L}(X_t^{\lambda})$, for $t\in[0,T]$ where $\mathbf{X}^{\lambda}$ is the state process solving (**??**) controlled by $\boldsymbol{\alpha}^{\lambda}$.

Properties of λ-interpolated Equilibriums

1. if f_0 and g satisfy the **Lasry-Lions monotonicity condition**

$$\forall m, m' \in \mathcal{P}(\mathbb{R}^d), \quad \int_{\mathbb{R}^d} \Big(f_0(x, m') - f_0(x, m) \Big) d(m' - m)(x) \geq 0,$$

   and similarly for g, then

$$0 \leq \lambda < \lambda' \leq 1 \Longrightarrow J^{\lambda', \mathrm{MF}}(\boldsymbol{\alpha}^{\lambda'}; \boldsymbol{\mu}^{\lambda'}) \leq J^{\lambda, \mathrm{MF}}(\boldsymbol{\alpha}^{\lambda}; \boldsymbol{\mu}^{\lambda}).$$

2. f_0 and g are given as

$$f_0(x, \mu) = \int_{\mathbb{R}^d} \varphi_0(x - y) d\mu(y), \quad g(x, \mu) = \int_{\mathbb{R}^d} \psi(x - y) d\mu(y),$$

   for **even convex** functions φ_0 and ψ with **Lipschitz continuous derivatives**, then for any $\lambda \in [0, 1]$, there exists a unique λ-interpolated mean field equilibrium control $\boldsymbol{\alpha}^{\lambda}$ and the mapping

$$[0, 1] \ni \lambda \mapsto (X_t^{\lambda})_{0 \leq t \leq T}$$

   is continuous for the norm $\|\boldsymbol{X}\|_{\mathbb{S}_2} := \sup_{0 \leq t \leq T} \mathbb{E}[|X_t|^2]^{1/2}$.

INCENTIVIZING THE OPTIMAL SOCIAL COST PER INDIVIDUAL

Question: *Can we incentivize the players (still by perturbation of their cost functions) into a behavior which leads to the same equilibrium costs as those obtained under the rule of the central planner?*

Master equation for the value function of the **central planner** optimization problem:

$$
\begin{aligned}
&\partial_t V(t,x,\mu) + \frac{\sigma^2}{2}\Delta V(t,x,\mu) - \frac{1}{2}|\partial_x V(t,x,\mu)|^2 \\
&\quad + \frac{1}{2}\int_{\mathbb{R}^d}\int_{\mathbb{R}^d} \partial_\mu V(t,x',\mu)(x)\partial_\mu V(t,x'',\mu)(x)d\mu(x')d\mu(x'') + f_0(x,\mu) \\
&\quad + \int_{\mathbb{R}^d}\bigg[\Big(-\partial_x V(t,\tilde{x},\mu) - \int_{\mathbb{R}^d}\partial_\mu V(t,x',\mu)(\tilde{x})d\mu(x')\Big)\cdot \partial_\mu V(t,x,\mu)(\tilde{x}) \\
&\qquad\qquad + \frac{\sigma^2}{2}\mathrm{trace}\Big(\partial_v\partial_\mu V(t,x,\mu)(\tilde{x})\Big)\bigg] d\mu(\tilde{x}) = 0,
\end{aligned}
$$

with terminal condition $V(T,x,\mu) = g(x,\mu)$.

Master equation for a MFG with same controlled state equation and running cost function

$$\tilde{f}(t, x, \mu, \alpha) = \frac{1}{2}|\alpha|^2 + \tilde{f}_0(x, \mu)$$

given by:

$$\partial_t U(t, x, \mu) + \frac{\sigma^2}{2}\Delta U(t, x, \mu) - \frac{1}{2}|\partial_x U(t, x, \mu))|^2$$
$$- \int_{\mathbb{R}^d} \partial_x U(t, v, \mu) \cdot \partial_\mu U(t, x, \mu)(v) d\mu(v)$$
$$+ \frac{\sigma^2}{2} \int_{\mathbb{R}^d} \mathrm{trace}\left[\partial_v \partial_\mu U(t, x, \mu)(v)\right] d\mu(v) + \tilde{f}_0(x, \mu) = 0,$$

with terminal condition $U(T, x, \mu) = g(x, \mu)$.

So choosing:

$$\tilde{f}_0(x, \mu) = f_0(x, \mu) + \frac{1}{2}\int_{\mathbb{R}^d}\int_{\mathbb{R}^d} \partial_\mu V(t, x', \mu)(x) \cdot \partial_\mu V(t, x'', \mu)(x) d\mu(x') d\mu(x'')$$
$$- \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} \partial_\mu V(t, x', \mu)(\tilde{x}) \cdot \partial_\mu V(t, x, \mu)(\tilde{x}) d\mu(x') d\mu(\tilde{x})$$

does the trick since the master equations **are the same**.

PRICE OF INSTABILITY FOR A SOCIAL EQUILIBRIUM

Recall: *the social cost of any MFG (Nash) equilibrium is higher than the social cost (per individual) incurred when the individuals all agree to use the (common) strategy identified by a central planner*

I argue: *while less costly, the central planner solution is less stable*

Reasonable question: *By how much can an individual player's average cost be lowered by deviating unilaterally from the MFC optimal control* α^{MKV} *identified by the social planner?*

Single Player Deviation

Following the control identified by the social planner, an individual agent's cost is

$$J^* := J^{\mathrm{MKV}}(\boldsymbol{\alpha}^{\mathrm{MKV}}) = \mathbb{E}\Big[\int_0^T \Big(\frac{1}{2}|\alpha_t^{\mathrm{MKV}}|^2 + f_0(X_t^{\mathrm{MKV}}, \mu_t^{\mathrm{MKV}})\Big)\,dt + g(X_T^{\mathrm{MKV}}, \mu_T^{\mathrm{MKV}})\Big].$$

Single Player Deviation

Following the control identified by the social planner, an individual agent's cost is

$$J^* := J^{\mathrm{MKV}}(\boldsymbol{\alpha}^{\mathrm{MKV}}) = \mathbb{E}\Big[\int_0^T \Big(\frac{1}{2}|\alpha_t^{\mathrm{MKV}}|^2 + f_0(X_t^{\mathrm{MKV}}, \mu_t^{\mathrm{MKV}})\Big) dt + g(X_T^{\mathrm{MKV}}, \mu_T^{\mathrm{MKV}})\Big].$$

If allowed to deviate from this control, **still evolving in the same environment**, the **smallest cost** for the agent should be

$$\hat{J}_0 := J^{\boldsymbol{\mu}^{\mathrm{MKV}}}(\hat{\boldsymbol{\alpha}}) = \mathbb{E}\Big[\int_0^T \Big(\frac{1}{2}|\hat{\alpha}_t|^2 + f_0(\hat{X}_t, \mu_t^{\mathrm{MKV}})\Big) dt + g(\hat{X}_T, \mu_T^{\mathrm{MKV}})\Big]$$

where $d\hat{X}_t = \hat{\alpha}_t dt + \sigma dW_t$ and

$$\hat{\boldsymbol{\alpha}} = \arg\min_{\boldsymbol{\alpha}} \mathbb{E}\Big[\int_0^T \Big(\frac{1}{2}|\alpha_t|^2 + f_0(X_t^{\boldsymbol{\alpha}}, \mu_t^{\mathrm{MKV}})\Big) dt + g(X_T^{\boldsymbol{\alpha}}, \mu_T^{\mathrm{MKV}})\Big]$$

$$\text{s.t. } dX_t^{\boldsymbol{\alpha}} = \alpha_t dt + \sigma dW_t.$$

Single Player Deviation

Following the control identified by the social planner, an individual agent's cost is

$$J^* := J^{\text{MKV}}(\boldsymbol{\alpha}^{\text{MKV}}) = \mathbb{E}\Big[\int_0^T \Big(\frac{1}{2}|\alpha_t^{\text{MKV}}|^2 + f_0(X_t^{\text{MKV}}, \mu_t^{\text{MKV}})\Big)dt + g(X_T^{\text{MKV}}, \mu_T^{\text{MKV}})\Big].$$

If allowed to deviate from this control, **still evolving in the same environment**, the **smallest cost** for the agent should be

$$\hat{J}_0 := J^{\mu^{\text{MKV}}}(\hat{\boldsymbol{\alpha}}) = \mathbb{E}\Big[\int_0^T \Big(\frac{1}{2}|\hat{\alpha}_t|^2 + f_0(\hat{X}_t, \mu_t^{\text{MKV}})\Big)dt + g(\hat{X}_T, \mu_T^{\text{MKV}})\Big]$$

where $d\hat{X}_t = \hat{\alpha}_t dt + \sigma dW_t$ and

$$\hat{\boldsymbol{\alpha}} = \arg\min_{\boldsymbol{\alpha}} \mathbb{E}\Big[\int_0^T \Big(\frac{1}{2}|\alpha_t|^2 + f_0(X_t^{\boldsymbol{\alpha}}, \mu_t^{\text{MKV}})\Big)dt + g(X_T^{\boldsymbol{\alpha}}, \mu_T^{\text{MKV}})\Big] \quad \text{s.t. } dX_t^{\boldsymbol{\alpha}} = \alpha_t dt + \sigma dW_t. \tag{2}$$

Notice this is a **classical control problem**!

PoI ≥ 0 by construction

PoI of a Social Optimum

Definition

The **Price of Instability (PoI)** is defined as the quantity:

$$\mathrm{PoI} = J^* - \hat{J}_0 \tag{3}$$

where

- J^* is the cost of the mean field control problem
- $\hat{J}_0$, is the cost of the optimal control in the classical control problem in the social planner environment

PoI of a Social Optimum

Definition

The **Price of Instability (PoI)** is defined as the quantity:

$$\text{PoI} = J^* - \hat{J}_0 \tag{3}$$

where

- J^* is the cost of the mean field control problem
- $\hat{J}_0$, is the cost of the optimal control in the classical control problem in the social planner environment

- Does not involve possible Nash equilibriums of the system.

PoI of a Social Optimum

Definition

The **Price of Instability (PoI)** is defined as the quantity:

$$\mathrm{PoI} = J^* - \hat{J}_0 \tag{3}$$

where

- J^* is the cost of the mean field control problem
- $\hat{J}_0$, is the cost of the optimal control in the classical control problem in the social planner environment

- Does not involve possible Nash equilibriums of the system.
- If $\mathrm{PoI} = 0$, $\boldsymbol{\alpha}^{\mathrm{MKV}}$ is an MFG equilibrium control. In this case, $\mathrm{PoS} = 1$,
- Furthermore, if the MFG equilibrium is unique, then we also have $\mathrm{PoA} = 1$.

Some Properties of PoI

Assume $\boldsymbol{\alpha}^{\mathrm{MKV}}$ is given by a **bounded feedback function** that is Lipschitz continuous in x.

1. If $\mathrm{PoI} = 0$, then it must hold,

$$\int_{\mathbb{R}^d} \partial_\mu f_0\left(x, \mu_t^{\mathrm{MKV}}\right)(y) d\mu_t^{\mathrm{MKV}}(x) = 0, \qquad y \in \mathbb{R}^d, t \in [0, T]$$

   and

$$\int_{\mathbb{R}^d} \partial_\mu g\left(x, \mu_T^{\mathrm{MKV}}\right)(y) d\mu_T^{\mathrm{MKV}}(x) = 0, \qquad y \in \mathbb{R}^d.$$

2. If f_0 and g are twice continuously differentiable, then

$$\boxed{\mathrm{PoI} \geq \frac{1}{4C} \mathbb{E} \int_0^T |Y_t|^2 dt,}$$

   where C is a constant which depends only on the model's coefficients and

$$Y_t = \mathbb{E}\left\{ \widetilde{\mathbb{E}}\left[\partial_\mu g\left(\widetilde{X}_T^{\mathrm{MKV}}, \mu_T^{\mathrm{MKV}}\right)\left(X_T^{\mathrm{MKV}}\right) + \int_t^T \partial_\mu f_0\left(\widetilde{X}_s^{\mathrm{MKV}}, \mu_s^{\mathrm{MKV}}\right)\left(X_s^{\mathrm{MKV}}\right) ds \right] \Big| \mathcal{F}_t \right\}.$$

Numerics for LQ Models

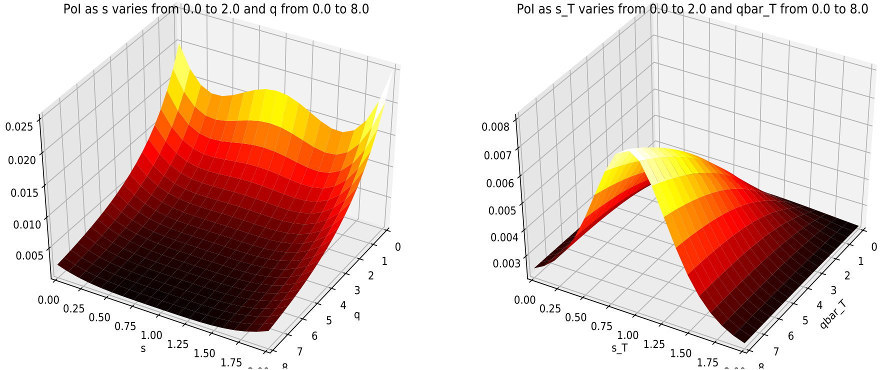

Price of Instability when s varies from $s = 0.0$ to $s = 2.0$ and q varies from 0.0 to 8.0 (left) and when s_T varies from $s_T = 0.0$ to $s_T = 2.0$ and $\bar{q}_T$ varies from 0.0 to 8.0. (right) all the other parameters are fixed to $s = 1.0$, $q = 1.0$, $\bar{q} = 1.0$, $q_T = 1.0$.

Thank You