PROBABILISTIC ANALYSIS OF MFGS
IV. GAMES OF TIMING AND FINITE STATE SPACE MEAN FIELD GAMES

René Carmona

Department of Operations Research & Financial Engineering
PACM
Princeton University

Minerva Lectures, Columbia U. October 28, 2016

Economic Models of Illiquidity & Bank Runs

- **Bryant** & **Dyamond-Dybvig**
  - deterministic, static, *undesirable* equilibrium
- **Morris-Shin** & **Rochet-Vives**
  - still static, investors' *private* (noisy) signals
- **He-Xiong**
  - **dynamic** continuous time model, **perfect observation**
  - exogenous randomness for **staggered** debt maturities
  - investors choose **to roll** or **not to roll**

O. Gossner's lecture: first game of timing

- diffusion model for the value of assets of the bank
- investors have private noisy signals
- investors choose a time to withdraw funds

M. Nutz Toy model for MFG game of timing with a continuum of players

Continuous Time Bank Run Model

Inspired by **Gossner**'s lecture

- N depositors
- Amount of each individual (initial & final) deposit $D_0^i = 1/N$
- Current interest rate r
- Depositors promised return $\bar{r} > r$
- Y_t = value of the assets of the bank at time t,
- Y_t Itô process, $Y_0 \geq 1$
- $L(y)$ liquidation value of bank assets if $Y = y$
- Bank has a credit line of size $L(Y_t)$ at time t at rate $\bar{r}$
- Bank uses credit line each time a depositor runs (withdraws his deposit)

Bank Run Model (cont.)

- Assets mature at time T, no transaction after that
- If $Y_T \geq 1$ every one is paid in full
- If $Y_T < 1$ **exogenous default**
- **Endogenous default** at time t if depositors try to withdraw **more** than $L(Y_t)$

Bank Run Model (cont.)

Each depositor $i \in \{1, \cdots, N\}$

- has access to a **private signal** X_t^i at time t

$$dX_t^i = dY_t + \sigma dW_t^i, \qquad i = 1, \cdots, N$$

- **chooses a time** $\tau^i \in \mathcal{S}^{X^i}$ at which to **TRY** to withdraw his deposit
- collects **return** $\bar{r}$ until time τ^i
- tries to **maximize**

$$J^i(\tau^1, \cdots, \tau^N) = \mathbb{E}\Big[g(\tau^i, Y_{\tau^i})\Big]$$

where

- $g(t, Y_t) = e^{-rt\wedge\tau}(L(Y_t) - N_t/N)^+ \wedge \frac{1}{N}$
- N_t number of withdrawals before t
- $\tau = \inf\{t;\ L(Y_t) < N_t/N\}$

Bank Run Model: case of full information

Assume

- $\sigma = 0$, i.e. Y_t is **public knowledge** !
- the function $y \hookrightarrow L(y)$ is also public knowledge
- $\tau^i \in \mathcal{S}^Y$

In **ANY** equilibrium

$$\tau^i = \inf\{t;\ L(Y_t) \leq 1\}$$

- Depositors withdraw at the **same time** (**run on the bank**)
- Each depositor gets his deposit back (**no one gets hurt!**)

Highly Unrealistic

Depositors should **wait longer** because of **noisy private signals**

Games of Timing

N players, states (observations / private signals) X_t^i at time t

$$dX_t^i = dY_t + \sigma dW_t^i$$

Y_t common unobserved signal (Itô process)

$$dY_t = \mu_t dt + \sigma_t dW_t^0$$

Each player maximizes

$$J^i(\tau^1, \cdots, \tau^N) = \mathbb{E}\Big[g(\tau^i, X_{\tau^i}, Y_{\tau^i}, \overline{\mu}^N([0, \tau^i])\Big]$$

where

- each τ^i is a $\mathcal{F}^{X^i}$ stopping time
- $\overline{\mu} = \frac{1}{N}\sum_{i=1}^N \delta_{\tau^i}$ empirical distribution of the τ^i's
- $g(t, x, y, p)$ is the reward to a player for
  - exercising his timing decision at time t when
  - his private signal is $X_t^I = x$,
  - the unobserved signal is $Y_t = y$,
  - the proportion of players who already exercised their right is p.

Abstract MFG Formulation

Recall

$$\begin{cases} dY_t = b_t dt + \sigma_t dW_t^0 \\ dX_t = dY_t + \sigma dW_t, \end{cases}$$

More generally:

1. The **states** of the players are given by a single measurable function

   $$X : \mathcal{C}([0,T]) \times \mathcal{C}([0,T]) \mapsto \mathcal{C}([0,T])$$

   progressively measurable $X(w^0, w)_t$ depends only upon $w^0_{[0,t]}$ and $w_{[0,t]}$,
2. $X^i = X(W^0, W^i)$ state process for player i
3. **Reward / cost function** F on $\mathcal{C}([0,T]) \times \mathcal{C}([0,T]) \times \mathcal{P}([0,T]) \times [0,T]$ **progressively measurable** $F(w^0, w, \mu, t)$ depends only upon $w^0_{[0,t]}$, $w_{[0,t]}$, and $\mu([0,s])$ for $0 \le s \le t$.

Approximate Nash Equilibria

Definition
If $\epsilon > 0$, a set $(\tau^{1,*}, \cdots, \tau^{N,*})$ of stopping time $\tau^{i,*} \in \mathcal{S}_{X^i}$ is said to be an ϵ-Nash equilibrium if for every $i \in \{1, \cdots, N\}$ and $\tau \in \mathcal{S}_{X^i}$ we have:

$$\mathbb{E}[F(W^0, W^i, \overline{\mu}^{N,-i}, \tau^{i,*})] \geq \mathbb{E}[F(W^0, W^i, \overline{\mu}^{N,-i}, \tau)] - \epsilon,$$

$\overline{\mu}^{N,-i}$ denoting the empirical distribution of $(\tau^{1,*}, \cdots, \tau^{i-1,*}, \tau^{i+1,*}, \cdots, \tau^{N,*})$.

Weak Characterization

the set of weak limits as $N \to \infty$ of ϵ_N - Nash equilibria when $\epsilon_N \searrow 0$ coincide with the set of weak solutions of the MFG equilibrium problem

Formulation of the MFG of Timing Problem

$$J(\mu,\tau) = \mathbb{E}[F(W^0, W, \mu, \tau)]$$

Definition
A stopping time $\tau^* \in \mathcal{S}_X$ is said to be a strong MFG equilibrium if for every $\tau \in \mathcal{S}_X$ we have:

$$J(\mu,\tau^*) \geq J(\mu,\tau)$$

with $\mu = \mathcal{L}(\tau^*|W^0)$.

MFG of Timing Problem

1. *Best Response Optimization*: for each random environment $\boldsymbol{\mu}$ solve

$$\hat{\theta} \in \arg \sup_{\theta \in \mathcal{S}_X, \theta \leq T} J(\mu,\theta);$$

2. *Fixed-Point Step*: find $\boldsymbol{\mu}$ so that

$$\forall t \in [0,T],\ \mu(W^0,[0,t]) = \mathbb{P}[\hat{\theta} \leq t | W^0].$$

Assumptions

(C) For each fixed $(w^0, w) \in \mathcal{C}([0,T]) \times \mathcal{C}([0,T])$, $(\mu, t) \mapsto F(w^0, w, \mu, t)$ is continuous.

(SC) For each fixed $(w^0, w, \mu) \in \mathcal{C}([0,T]) \times \mathcal{C}([0,T]) \times \mathcal{P}([0,T])$, $t \mapsto F(w^0, w, \mu, t)$ is upper semicontinuous.

(ID) For any progressively measurable random environments $\mu, \mu' : \mathcal{C}([0,T]) \mapsto \mathcal{P}([0,T])$ s.t. $\mu(w^0) \le \mu'(w^0)$ a.s.

$$M_t = F(W^0, W, \mu'(W^0), t) - F(W^0, W, \mu(W), t)$$

is a sub-martingale.

(ID) holds when F has **increasing differences** $t \le t'$ and $\mu \le \mu'$ imply:

$$F(w^0, w, \mu', t') - F(w^0, w, \mu', t) \ge F(w^0, w, \mu, t') - F(w^0, w, \mu, t).$$

(ID) $\Longrightarrow$ the expected reward J has also increasing differences

$$J(\mu', \tau') - J(\mu', \tau) \ge J(\mu, \tau') - J(\mu, \tau)$$

Major Disappointment: if$F(w^0, w, \mu, t) = G(\mu[0,t])$ for some real-valued continuous function G on $[0,1]$ which we assume to be differentiable on $(0,1)$. If F satisfies assumption **(ID)**, then G is constant.

Fixed Point Results on Order Lattices

Recall: A partially ordered set $(\mathcal{S}, \leq)$ is said to be a lattice if:

$$x \vee y = \inf\{z \in \mathcal{S};\ z \geq x, z \geq y\} \in \mathcal{S}$$

and (1)

$$x \wedge y = \sup\{z \in \mathcal{S};\ z \leq x, z \leq y\} \in \mathcal{S},$$

for all $x, y \in \mathcal{S}$. A lattice $(\mathcal{S}, \leq)$ is said to be complete if every subset $S \subset \mathcal{S}$ has a greatest lower bound $\inf S$ and a least upper bound $\sup S$, with the convention that $\inf \emptyset = \sup \mathcal{S}$ and $\sup \emptyset = \inf \mathcal{S}$.

Example The set $\mathcal{S}$ of stopping times of a right continuous filtration $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$

Fact 1: If $\mathcal{S}$ is a complete lattice and $\Phi : \mathcal{S} \ni x \mapsto \Phi(x) \in \mathcal{S}$ is order preserving in the sense that $\Phi(x) \leq \Phi(y)$ whenever $x, y \in \mathcal{S}$ are such that $x \leq y$, the set of fixed points of Φ is a non-empty complete lattice.

Another definition A real valued function f on a lattice $(\mathcal{S}, \leq)$ is said to be supermodular if for all $x, y \in \mathcal{S}$

$$f(x \vee y) + f(x \wedge y) \geq f(x) + f(y). \tag{2}$$

Existence of Strong Equilibria

Under assumptions **(SC)** and **(ID)** there exists a strong equilibrium. Moreover, if continuity **(C)** is assumed instead of semicontinuity **(SC)**, then there exist strong equilibria τ^* and θ^* such that for any strong equilibrium τ we have $\theta^* \leq \tau \leq \tau^*$ a.s.

MFGs in Finite State Spaces or Graphs

Equilibrium Pbs with Finitely Many States

Finite State Space $E = \{1, \cdots, M\}$ in lieu of $\mathbb{R}^d$

Motivation

- Vaccination Models: **Laguzet - Turinici**
- Computer network security (Botnet attacks) **Kolokolstov-Bensoussan**

Papers

- MFGs on Finite State Spaces **Gomes-Mohr-Souza**
- MFGs on Graphs **Guéant**
- MFGs with Major and Minor Players **R.C.-P.Wang**

In both cases

- Mean Field Interactions
  (dynamics and costs depend upon proportion of individuals in a given state)

Continuous Time, Finite State Dynamics

SDEs replaced by **Continuous Time Stochastic Processes** in **finite state space** E

For **convenience** give up on **open loop problems** use controls in **feedback form** so **markovian dynamics**

Distribution given by a **Q-matrix** $q_t = [q_t(x, x')]_{x,x' \in E}$:

$$q_t(x, x') = \textit{rate of jumping from state } x \textit{ to } x' \textit{ at time } t.$$

$$\mathbb{P}[X_{t+\Delta t} = x' | X_t = x] = q_t(x, x')\Delta t + o(\Delta t).$$

Properties of Q-matrices

$$\begin{cases} q_t(x, x') \geq 0 & \textit{if } x' \neq x \\ q_t(x, x) = -\sum_{x' \neq x} q_t(x, x') \end{cases}$$

Finite State Mean Field Game: Data

Jump rates

$$[0,T] \times E \times E \times \mathcal{P}(E) \times A \ni (t, x, x', \mu, \alpha) \hookrightarrow \lambda_t(x, x', \mu, \alpha)$$

Q-matrix

$$q_t(x, x') = \lambda_t(x, x', \mu, \alpha)$$

Costs

- Running cost function

$$[0,T] \times E \times \mathcal{P}(E) \times A \ni (t, x, \mu, \alpha) \hookrightarrow f(t, x, \mu, \alpha)$$

- terminal cost function

$$E \times \mathcal{P}(E) \ni (x, \mu) \mapsto g(x, \mu)$$

Remark,

If $\mu \in \mathcal{P}(E)$, $\mu = (\mu(\{x\})_{x \in E}$ finite dimensional **probability simplex!**

Finite State Mean Field Games

Hamiltonian

$$H(t,x,\mu,h,\alpha) = \sum_{x' \in E} \lambda_t(x,x',\mu,\alpha)h(x') \; + f(t,x,\mu,\alpha).$$

Hamiltonian minimizer

$$\hat{\alpha}(t,x,\mu,h) = \arg\inf_{\alpha \in A} H(t,x,\mu,h,\alpha),$$

Minimized Hamiltonian

$$H^*(t,x,\mu,h) = \inf_{\alpha \in A} H(t,x,\mu,h,\alpha) = H(t,x,\mu,h,\hat{\alpha}(t,x,\mu,h)).$$

HJB Equation

$$\partial_t u^{\mu}(t,x) + H^*(t,x,\mu_t,u^{\mu}(t,\cdot)) = 0, \qquad 0 \le t \le T, \; x \in E,$$

with terminal condition $u^{\mu}(T,x) = g(x,\mu_T)$.

The Master Equation Equation

$$\partial_t U + H^*(t,x,\mu,U(t,\cdot,\mu)) + \sum_{x'\in E} h^*(t,\mu,U(t,\cdot,\mu))(x')\frac{\partial U(t,x,\mu)}{\partial\mu(\{x'\})} = 0,$$

where the $\mathbb{R}^E$-valued function h^* is defined on $[0,T]\times\mathcal{P}(E)\times\mathbb{R}^E$ by:

$$\begin{aligned} h^*(t,\mu,u) &= \int_E \lambda_t\big(x,\,\cdot\,,\mu,\hat\alpha(t,x,\mu,u)\big)\, d\mu(x) \\ &= \sum_{x\in E}\lambda_t\big(x,\,\cdot\,,\mu,\hat\alpha(t,x,\mu,u)\big)\mu(\{x\}). \end{aligned}$$

System of Ordinary Differential Equations (**ODEs**)

A Cyber Security Model

- N computers in a network (**minor players**)
- One hacker / attacker (**major player**)
- Action of major player affect minor player states (even when $N >> 1$)
- Major player feels only μ_t^N the empirical distribution of the minor players' states

Finite State Space: each computer is in one of 4 states

- protected & infected
- protected & susceptible to be infected
- unprotected & infected
- unprotected & susceptible to be infected

Continuous time Markov chain in $E = \{DI, DS, UI, US\}$

Each **player's action** is intended to affect the **rates of change** from one state to another to minimize **expected costs**

$$J(\boldsymbol{\alpha}^0, \boldsymbol{\alpha}) = \mathbb{E}\left[\int_0^T (k_D \mathbf{1}_D + k_I \mathbf{1}_I)(X_t)dt\right]$$

$$J^0(\boldsymbol{\alpha}^0, \boldsymbol{\alpha}) = \mathbb{E}\left[\int_0^T (-f_0(\mu_t) + k_H \phi^0(\mu_t))dt\right]$$

Minor Players Transition Rates

$$\lambda_t(\cdot,\cdot,\mu,v_{\mathrm{H}},0) =$$

$$\begin{array}{c|cccc} & \mathrm{DI} & \mathrm{DS} & \mathrm{UI} & \mathrm{US} \\ \mathrm{DI} & \cdots & q_{\mathrm{rec}}^{\mathrm{D}} & 0 & 0 \\ \mathrm{DS} & v_{\mathrm{H}} q_{\mathrm{inf}}^{\mathrm{D}} + \beta_{\mathrm{DD}}\mu(\{\mathrm{DI}\}) + \beta_{\mathrm{UD}}\mu(\{\mathrm{UI}\}) & \cdots & 0 & 0 \\ \mathrm{UI} & 0 & 0 & \cdots & q_{\mathrm{rec}}^{\mathrm{U}} \\ \mathrm{US} & 0 & 0 & v_{\mathrm{H}} q_{\mathrm{inf}}^{\mathrm{U}} + \beta_{\mathrm{UU}}\mu(\{\mathrm{UI}\}) + \beta_{\mathrm{DU}}\mu(\{\mathrm{DI}\}) & \cdots \end{array}$$

$$\lambda_t(\cdot,\cdot,\mu,v_{\mathrm{H}},1) =$$

$$\begin{array}{c|cccc} & \mathrm{DI} & \mathrm{DS} & \mathrm{UI} & \mathrm{US} \\ \mathrm{DI} & \cdots & q^{\mathrm{D}_{\mathrm{rec}}} & \lambda & 0 \\ \mathrm{DS} & v_{\mathrm{H}} q_{\mathrm{inf}}^{\mathrm{D}} + \beta_{\mathrm{DD}}\mu(\{\mathrm{DI}\}) + \beta_{\mathrm{UD}}\mu(\{\mathrm{UI}\}) & \cdots & 0 & \lambda \\ \mathrm{UI} & \lambda & 0 & \cdots & q_{\mathrm{rec}}^{\mathrm{U}} \\ \mathrm{US} & 0 & \lambda & v_{\mathrm{H}} q_{\mathrm{inf}}^{\mathrm{U}} + \beta_{\mathrm{UU}}\mu(\{\mathrm{UI}\}) + \beta_{\mathrm{DU}}\mu(\{\mathrm{DI}\}) & \cdots \end{array}$$

FINITE PLAYERS MFGS

One major player and N minor players

- X_t^0 state of major player: $X_t^0 \in E^0 = \{1, 2, \ldots, d^0\}$
- X_t^j state of major player:: $X_t^j \in E = \{1, 2, \ldots, d\} \quad j = 1, \cdots, N$

At time $t \leq T$, the major player...

- can observe its own states X_t^0 and the empirical distribution μ_t^N of the minor player's states.
- chooses a control of the form $\alpha^0(t, X_t^0, \mu_t^N)$.

Each minor player...

- can observe its own states X_t^j, the state X_t^0 of the major player, and the empirical distribution μ_t^N.
- chooses a control of the form $\alpha(t, X_t^j, X_t^0, \mu_t^N)$.

The system evolves as a Continuous-Time Markov Chain

- The transition rate matrix of each player depends on his own states, major player's state and μ_t^N.
- The change of states are conditionally independent among the players.

Jump Rates of the System

- Minor players' jump rates:

$$\begin{aligned}[0, T] \times E \times E \times E^0 \times A^0 \times \mathcal{P}(E) \times A &\to \mathbb{R} \\ (t, x, x', x^0, \alpha^0, \mu, \alpha) &\to q(t, x, x', x^0, \alpha^0, \mu, \alpha)\end{aligned}$$

- Major player's jump rate:

$$\begin{aligned}[0, T] \times E^0 \times E^0 \times \mathcal{P}(E) \times A^0 &\to \mathbb{R} \\ (t, x^0, x'^0, \mu, \alpha^0) &\to q^0(t, x^0, x'^0, \mu, \alpha^0)\end{aligned}$$

- Major player's control and state impact EVERY player in the game.
- We assume that q and q^0 satisfies:

$$q(t, x, x', x^0, \alpha^0, \mu, \alpha) \geq 0, \quad q^0(t, x^0, x'^0, \mu, \alpha^0) \geq 0$$

$$q(t, x, x, x^0, \alpha^0, \mu, \alpha) = -\sum_{x' \neq x} q(t, x, x', x^0, \alpha^0, \mu, \alpha)$$

$$q^0(t, x^0, x'^0, \mu, \alpha^0) = -\sum_{x'^0 \neq x^0} q^0(t, x^0, x'^0, \mu, \alpha^0).$$

Jump Rates of the System

- The changes of states are conditionally independent among the players:

$$\mathbb{P}[X^0_{t+\Delta t} = j^0, X^1_{t+\Delta t} = j^1, \ldots, X^N_{t+\Delta t} = j^N | X^0_t = i^0, X^1_t = i^1, \ldots, X^N_t = i^N]$$

$$:= [\mathbb{1}_{i^0=j^0} + q^0(t, i^0, j^0, \alpha(t, i^0, \mu^N_t), \mu^N_t)\Delta t + o(\Delta t)]$$
$$\times \prod_{n=1}^{N} [\mathbb{1}_{i^n=j^n} + q(t, i^n, j^n, \beta^n(t, i^n, i^0, \mu^N_t), i^0, \alpha(t, i^0, \mu^N_t), \mu^N_t)\Delta t + o(\Delta t)]$$

- This is equivalent to define the transition rate matrix Q^N for the Markov Chain $(X^0_t, X^1_t, \ldots, X^N_t)$ with $M^0 \times M^N$ states.

  Here is how: we just retain the first order terms by expending the RHS of the above equality.

- Q^N is a HUGE sparse matrix as N grows.

Payoff and Symmetric Nash Equilibrium

Fix a finite horizon T.

- Major player's payoff:

$$J^{N,0}(\alpha, \beta^1, \ldots, \beta^N) := \mathbb{E}\left[\int_0^T f^0(t, \alpha(t, X_t^0, \mu_t^N), X_t^0, \mu_t^N)dt + g^0(X_T^0, \mu_T^N)\right]$$

- Minor player's payoff:

$$J^{N,n}(\alpha, \beta^1, \ldots, \beta^N) := \mathbb{E}[\int_0^T f(t, \beta^n(t, X_t^n, X_t^0, \mu_t^N), X_t^n, \alpha(t, X_t^0, \mu_t^N), X_t^0, \mu_t^N)dt$$
$$+ g(X_T^n, X_T^0, \mu_T^N)]$$

Our objective is to search for the Symmetric Nash Equilibrium.

i.e. to find $\alpha^* \in \mathbb{A}^0$ and $\beta^* \in \mathbb{A}$ such that for all $\alpha \in \mathbb{A}^0$ and $\beta \in \mathbb{A}$:

$$J^{N,0}(\alpha^*, \beta^*, \ldots, \beta^*) \geq J^{N,0}(\alpha, \beta^*, \ldots, \beta^*)$$
$$J^{N}(\alpha^*, \beta^*, \ldots, \beta^*) \geq J^{N}(\alpha^*, \beta^*, \ldots, \beta, \ldots \beta^*)$$

Formulation of the Mean Field Game

Why do we use MFG?

- $N-$player Game is difficult: as number of players grows, the dimension of the transition rate matrix of the system increases exponentially.
- Use MFG paradigm: consider the limit case where the number of minor player N tends to infinity.
- Propagation of Chaos: hope that the solution of the limit case provides an approximative equilibrium for N-player game when N is large.

Perks of MFG:

- The empirical distribution of the minor players' states has a tractable form of infinitesimal generator.
- Deviation of a SINGLE minor player's strategy has NO impact on the distribution of minor player's states.

Strategy of Solution

We employ a fixed point argument based on the **controls of the players**:

Step 1 (Major Player's Problem)

- Fix an admissible strategy $\mathbb{A} \ni \bar{\beta} = \bar{\beta}(t, X_t^n, X_t^0, \mu_t)$ for the minor players.
- Given that all the minor players use the strategy $\bar{\beta}$, solve for the optimal control of the major player $\alpha^*(\bar{\beta})$.

Step 2 (Representative Minor Player's Problem)

- Fix an admissible strategy $\mathbb{A}^0 \ni \bar{\alpha} = \bar{\alpha}(t, X_t^0, \mu_t)$ for the major player and a Markov strategy $\bar{\beta} = \bar{\beta}(t, X_t^n, X_t^0, \mu_t)$ for the minor players.
- Consider a population of minor players using strategy $\bar{\beta}$ and a major player using strategy $\bar{\alpha}$. Denote $\mu_t(\bar{\alpha}, \bar{\beta})$ the corresponding distribution of the population of minor players.
- Consider an additional minor player facing the major player $\bar{\alpha}$, and the distribution $\mu_t(\bar{\alpha}, \bar{\beta})$.
- Solve for his optimal control $\beta^*(\bar{\alpha}, \bar{\beta})$.

Step 3 (Fixed Point Argument)

Search for the fixed point $[\bar{\alpha}, \bar{\beta}] = [\alpha^*(\bar{\beta}), \beta^*(\bar{\alpha}, \bar{\beta})]$.