Mean Field Game Analysis of Tournaments

Yuchong Zhang [1]
(Joint works with M. Nutz[1], E. Bayraktar[2] and J. Cvitanic[3])

[1]Columbia University

[2]University of Michigan

[3]Caltech

Advances in Stochastic Analysis for Risk Modeling
CIRM, November 16, 2017

Outline

Motivation

- Rank-based criteria is everywhere:
  - sports, R&D, admission, ... any "competition"
  - people's utility often depends on relative performance
- May have a mixture of absolute and relative performance criteria.
- Goal: analyze large population competition of "timing"
  - stochastic control rather than stopping problems
  - want a model with tractability
  - describe equilibrium behavior
  - tournament design
- MFG where agents interact through the ranking of the hitting times

Mean Field Games

- Introduced by Lasry–Lions and Huang–Malhamé–Caines, 2006
- Nash equilibria for $N \to \infty$ players
- Interaction through empirical distribution ν of the private states
- Typical setting: each player controls a diffusion with some reward and cost-of-effort which depend on ν

- Coupled system: HJB and Kolmogorov PDEs, or FBSDEs
- Only Linear-Quadratic control can be solved explicitly

- Cardaliaguet, Carmona, Delarue, Fouque, Lacker, Yam, ...
- Mean field games of timing: Carmona–Delarue–Lacker, Nutz
- Toy example of Lasry–Lions, “When Does the Meeting Start?”
- Principal-agent problem: Élie, Mastrolia & Possaimai

General strategy for finding a Nash equilibrium

1. **Best-response step**:
   Solve a standard stochastic control problem for a representative player, given the strategy or performance of all other players.

2. **Fixed-point step**:
   Find a fixed point of the best-response mapping.

Outline

The Poissonian model

- A continuum of agents $i \in (I, \mathcal{I}, \mu)$; atomless
- Each agent controls intensity λ of her independent Poisson process
- Quadratic instantaneous cost $c\lambda_t^2\, dt$
- Goal is reached if process jumps (once)

Agents...

- are ranked according to their completion times
- are paid a rank-based reward $R : [0, 1] \to \mathbb{R}_+$ (decreasing, piecewise Lipschitz and left-continuous at $r = 1$)
- maximize expected reward minus cost

Interaction:

- Agents observe $\rho(t) = \mu\{i : \tau^i \leq t\}$ = proportion of agents completed by time t, and choose feedback control $\lambda(\rho(t))$
- Reward $R(\rho(\tau_i))$ and cost $c(\rho(t))$ depend on rank/proportion

Solving the MFG

- The representative player's problem:

$$v(r) = \sup_{\lambda \in \Lambda} E\left[R(\rho(\tau_\lambda)) - \int_0^{\tau_\lambda} c(\rho(t))\lambda(\rho(t))^2\, dt \,\middle|\, \rho(0) = r \right]$$

where $\Lambda =$set of feedback controls (piecewise Lipschitz),

$$\rho(t) = \int_0^t \bar{\lambda}(\rho(s))(1 - \rho(s))\, ds$$

- The fixed-point problem:

$$\bar{\lambda} \mapsto \rho \mapsto \text{optimal } \lambda$$

  - Agent space $(I, \mathcal{I}, \mu)$ is atomless
  - Hence this is Nash: no single agent influences ρ

Existence and uniqueness

Theorem

There exists a unique (a.e.) equilibrium optimal control $\lambda^* \in \Lambda$,

$$\lambda^*(r) = \frac{R(r) - \frac{1}{2\sqrt{1-r}} \int_r^1 \frac{R(y)}{\sqrt{1-y}}\,dy}{2c(r)}, \quad r \in [0,1)$$

In equilibrium, the value function of any agent before completion is

$$v(r) = \frac{1}{2\sqrt{1-r}} \int_r^1 \frac{R(y)}{\sqrt{1-y}}\,dy, \quad r \in [0,1)$$

- Value of a freshly started game: $V = v(0) = \frac{1}{2}\int_0^1 \frac{R(y)}{\sqrt{1-y}}dy$.
- V is independent of c: higher cost $\Rightarrow$ smaller optimal effort $\Rightarrow$ state ρ is slowed down$\Rightarrow$ same reward

Example: no cut-off

Reward R Equilibrium effort λ^*

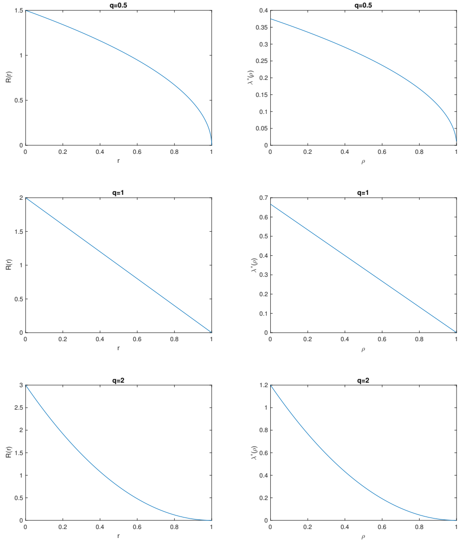

Example: cut-off at $\alpha = 60\%$

Reward R Equilibrium effort λ^*

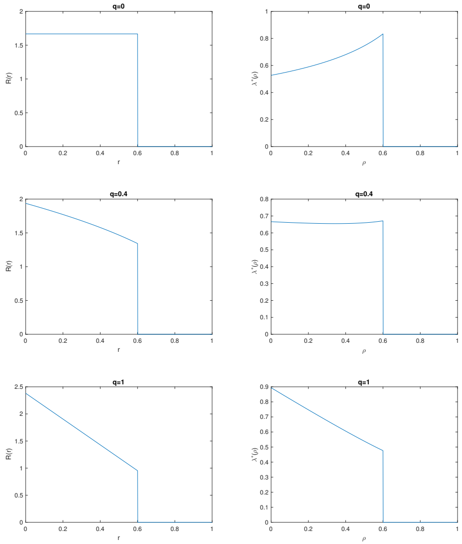

The principal's problem

- Given R, there exists a unique, deterministic equilibrium state ρ
- Time until α-fraction of the population has completed:

$$T_\alpha(R) = \inf\{t \geq 0 : \rho(t) \geq \alpha\} \in (0, \infty]$$

- Given $\alpha \in (0,1)$ and budget $B > 0$,

$$\text{minimize } T_\alpha(R) \quad \text{subject to} \quad \int_0^1 R(r)\,dr \leq B$$

- What reward scheme $R^* \geq 0$ can attain $T_\alpha^* = \inf_R T_\alpha(R)$?
- A constrained calculus of variation problem

Theorem

Suppose $\frac{c(r)(1-r)}{2-r}$ is decreasing. Given a fixed reward budget $B > 0$ and $\alpha \in (0,1)$, the optimal non-negative rank-based reward scheme to minimize $T_\alpha(R)$ is

$$R^*(r) = \frac{B}{C}\left\{\sqrt{\frac{c(r)}{2-r}} + \frac{1}{2}\int_r^\alpha \frac{1}{1-s}\sqrt{\frac{c(s)}{2-s}}ds\right\}\mathbf{1}_{[0,\alpha]}(r),$$

and the minimum α-completion time is

$$T_\alpha^* = \frac{4C^2}{B}, \quad \text{where} \quad C := \frac{1}{2}\int_0^\alpha \frac{\sqrt{c(r)(2-r)}}{1-r}dr.$$

The corresponding equilibrium effort is

$$\lambda^*(r) = \frac{B}{2C}\frac{1}{\sqrt{(2-r)c(r)}}\mathbf{1}_{[0,\alpha]}(r)$$

Plots for $\alpha = 25\%, 50\%, 75\%$, with c constant

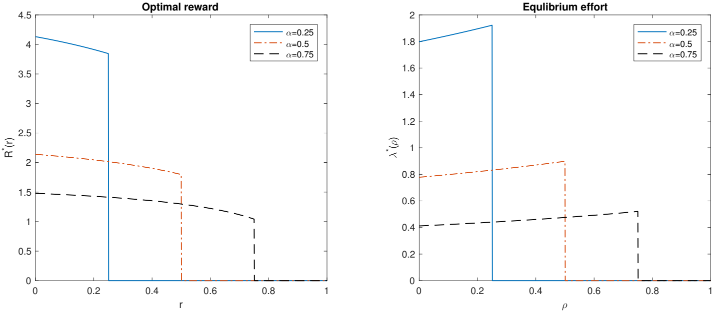

Outline

The diffusion model

- Reward function: $R(t,r) = 1_{\{t \le T\}} H(r) + 1_{\{t>T\}} 0$.
- Assume agents are homogeneous: (x_0, c, σ).
- Optimization problem for a representative player:

$$\sup_a E\left[R_\mu(\tau_\mu) - \int_0^{\tau_\mu \wedge T} c a_t^2 dt \right]$$

where $R_\mu(t) = R(t, F_\mu(t)) = 1_{\{t \le T\}} H(F_\mu(t))$,

$$dX_t = -a_t dt + \sigma dB_t, \;\; X_0 = x_0,$$
$$\tau_\mu = \inf\{t \ge 0 : X_t = 0\}.$$

- Nash equilibrium: fixed points of $\mu \mapsto \mathcal{L}(\tau_\mu^*)$.

Finding the best-response

- Given μ, the best-response can be computed by solving the HJB equation for the value function:

$$v_t + \sup_{a \geq 0} \left\{ -a v_x + \frac{1}{2}\sigma^2 v_{xx} - c a^2 \right\} = 0,$$

$$v(t,0) = H(F_\mu(t)), \quad v(T,x) = 0.$$

- The nonlinear HJB equation can be linearized using the Cole-Hopf transformation $u = \exp(\frac{v}{2c\sigma^2})$. u has stochastic representation:

$$u(t,x) = E\left[\exp\left(\frac{R_\mu(t + \tau^\circ_{x/\sigma})}{2c\sigma^2}\right)\right],$$

where τ°_m is the BM first passage time to level m.

- The optimal feedback control $a^* = -\sigma^2 u_x / u$.

Explicit Nash equilibrium

Proposition. The fixed point equation (for μ) is :

$$f_\mu(t) = \frac{u(t,0;\mu)}{u(0,x_0;\mu)} f_{\tau^\circ_{x_0/\sigma}}(t).$$

Theorem ($T = \infty$)

There is a unique equilibrium completion time distribution μ given in terms of its quantile function $r \mapsto T^\mu_r$ by

$$T^\mu_r = F^{-1}_{\tau^\circ_{x_0/\sigma}} \left(\frac{\int_0^r \exp\left(-\frac{H(z)}{2c\sigma^2}\right) dz}{\int_0^1 \exp\left(-\frac{H(z)}{2c\sigma^2}\right) dz} \right).$$

Moreover, the value of the game is given by

$$V_\infty = -2c\sigma^2 \ln \left(\int_0^1 \exp\left(-\frac{H(z)}{2c\sigma^2}\right) dz \right).$$

Explicit Nash equilibrium

Proposition. The fixed point equation (for μ) is :

$$f_\mu(t) = \frac{u(t,0;\mu)}{u(0,x_0;\mu)} f_{\tau^\circ_{x_0/\sigma}}(t).$$

Theorem ($T = \infty$)

There is a unique equilibrium completion time distribution μ given in terms of its quantile function $r \mapsto T_r^\mu$ by

$$T_r^\mu = F^{-1}_{\tau^\circ_{x_0/\sigma}}\left(\frac{\int_0^r \exp\left(-\frac{H(z)}{2c\sigma^2}\right)dz}{\int_0^1 \exp\left(-\frac{H(z)}{2c\sigma^2}\right)dz}\right).$$

Moreover, the value of the game is given by

$$V_\infty = -2c\sigma^2 \ln\left(\int_0^1 \exp\left(-\frac{H(z)}{2c\sigma^2}\right)dz\right).$$

Theorem ($T < \infty$)

There is a unique equilibrium completion time distribution μ given in terms of its quantile function $r \mapsto T_r^\mu$ by

$$T_r^\mu = F^{-1}_{\tau^\circ_{x_0/\sigma}}\left(\frac{1 - F_{\tau^\circ_{x_0/\sigma}}(T)}{1 - F_\mu(T)} \int_0^r \exp\left(\frac{-H(z)}{2c\sigma^2}\right) dz\right), \quad r \in [0, F_\mu(T)],$$

where the equilibrium terminal completion rate $F_\mu(T) \in (0,1)$ is the unique solution of

$$F_{\tau^\circ_{x_0/\sigma}}(T) = \frac{1 - F_{\tau^\circ_{x_0/\sigma}}(T)}{1 - F_\mu(T)} \int_0^{F_\mu(T)} \exp\left(\frac{-H(z)}{2c\sigma^2}\right) dz.$$

Moreover, the value of the game is given by

$$V = 2c\sigma^2 \ln\left(\frac{1 - F_{\tau^\circ_{x_0/\sigma}}(T)}{1 - F_\mu(T)}\right).$$

Dependence on the cost parameter

- (One-stage) Poissonian: game value is independent of c.
- Diffusion: game value is increasing in c when $T = \infty$, and non-monotone in c when $T < \infty$.

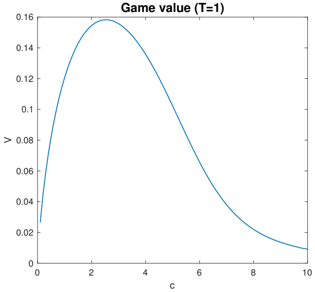

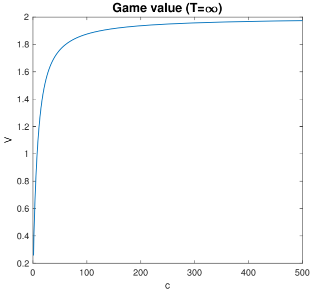

$$x_0 = 1,\ \sigma = 0.25,\ H(r) = 6(1-r)^2.$$

- Efficiency does not necessarily make people happier.

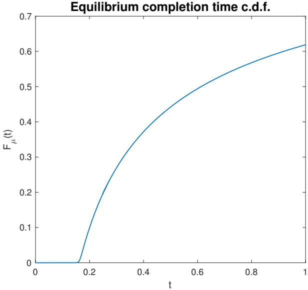

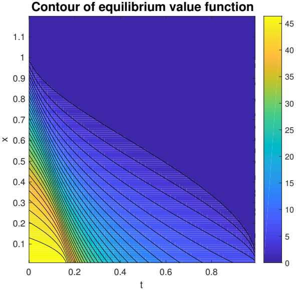

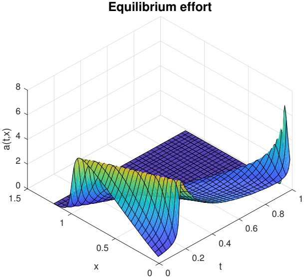

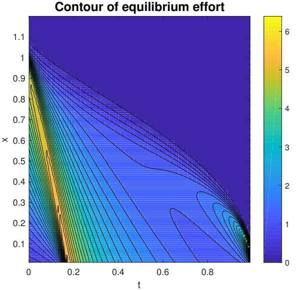

$$x_0 = c = 1,\ \sigma = 0.25,\ H(r) = 6(1-r)^2.$$

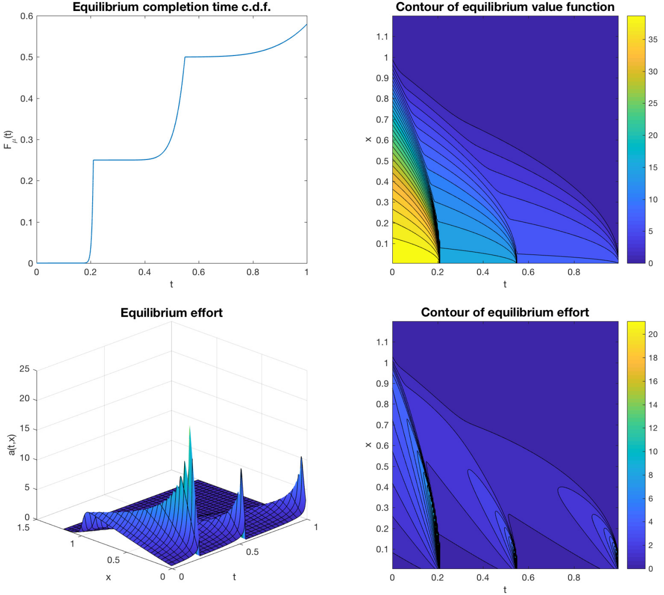

$$x_0 = c = 1,\ \sigma = 0.25,\ H = 5 \cdot 1_{[0,0.25)} + 2 \cdot 1_{[0.25,0.5)} + 1_{[0.5,1]}.$$

Minimizing the α-quantile

Objective: Given limited reward budget K, what (non-negative) reward scheme minimizes the time it takes $\alpha \in (0,1)$ fraction of the population to complete their projects in equilibrium?

Theorem: regardless of the tournament horizon, the unique (up to a.e. equivalence) optimal reward scheme is the uniform scheme with cutoff rank α:

$$H^*(r) = \frac{K}{\alpha} 1_{[0,\alpha]}(r).$$

- cf. Poissonian competition.

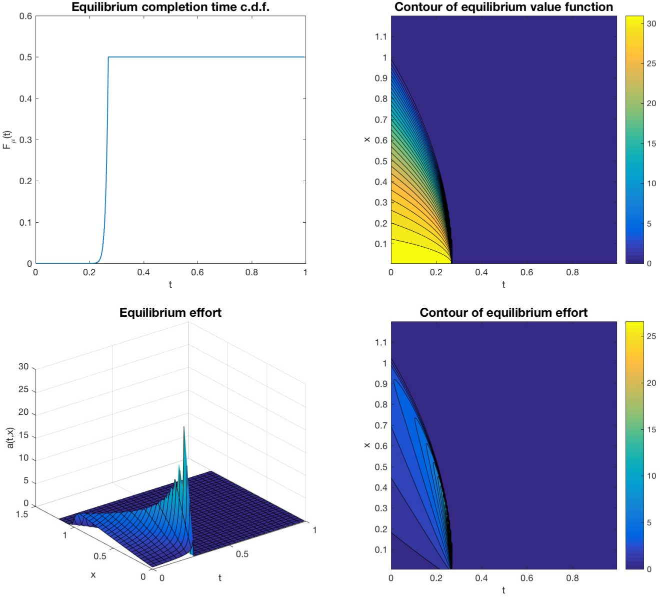

$$x_0 = c = 1,\ \sigma = 0.25,\ K = 2,\ \alpha = 0.5,\ T = 1.$$

Maximizing profit

Objective: Suppose each project completed at time t generates a profit $g(t)$ for the principal. What reward scheme maximizes the aggregate net profit $E[g(\tau) - R_\mu(\tau)]$, $\tau \sim \mu$ for the principal?

Theorem: Suppose $T = \infty$. Given a non-constant decreasing profit function $g \in C_b(\mathbb{R}_+)$, an optimal reward scheme $H^* \geq 0$ is given by

$$H^*(r) = g(F_{\mu^*}^{-1}(r) \wedge t_b^*) - g(t_b^*),$$

where t_b^* = argmax of a one-dimension static optimization problem, and f_{μ^*} has an explicit formula involving t_b^*.

- In equilibrium, the agents receive $g(t \wedge t_b^*) - g(t_b^*)$ for finishing at time t. It is optimal for the principal to align his interest with that of the agents.
- Optimizing over the set of feasible μ.

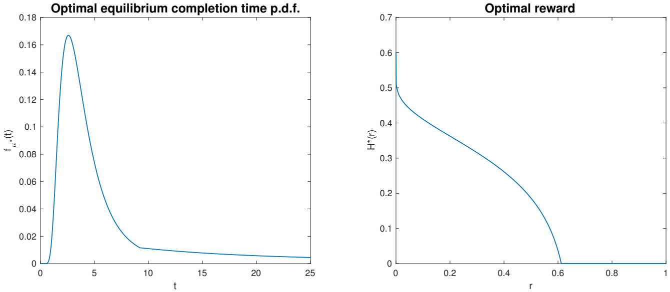

$g(t) = x_0 e^{-0.1t}$, $x_0 = c = 1$, $\sigma = 0.25$. The participation reward constraint starts to be binding at rank 0.61.

Outline

Conclusion

- Two tractable mean field models of competition
- Unique equilibrium in semi-closed form
  - Poissonian: game value is independent of c
  - Diffusion: game value is increasing in c (when $T = \infty$)
- Principal-agent problem explicitly solvable via calculus of variation
  - Poissonian: quantile-minimizing scheme is not uniform
  - Diffusion: quantile-minimizing scheme is uniform
- Open questions: multi-stage Poissonian model, heterogeneous agents, time-consistent formulation of the principal's problem, un-fixed pie, multiple tournaments, ...

Main references

E. Bayraktar, J. Cvitanić, and Y. Zhang.
Mean field game analysis of tournaments.
In preparation.

E. Bayraktar and Y. Zhang.
A rank-based mean field game in the strong formulation.
Electron. Commun. Probab., 21, paper no. 72:1–12, 2016.

R. Carmona and F. Delarue.
Probabilistic analysis of mean-field games.
SIAM J. Control Optim., 51(4):2705–2734, 2013.

M. Nutz.
A mean field game of optimal stopping.
Preprint, arXiv:1605.09112v1, 2016.

M. Nutz and Y. Zhang.
A mean field competition.
Preprint, arXiv:1708.01308v1, 2017.

Thank you for your attention!