Revision Protocols, Orders of Limits, and Stochastic Stability

William H. Sandholm

University of Wisconsin

This project:

1. Orders of Limits for Stationary Distributions, Stochastic Dominance, and Stochastic Stability (*TE* (2010))
2. Stochastic Imitative Game Dynamics with Committed Agents
3. Exponential Protocols and Large Population Potential Games

Evolutionary Game Theory

Population games model strategic interactions in which:

1. The number of agents is large.
2. Individual agents are small.
3. Agents are anonymous: an agent's payoff depends his own action and the distribution of others' actions.

Applications: network congestion, public goods with externalities, macro spillovers, decentralized control, cultural evolution, language formation, market behavior. . .

Evolutionary dynamics

Traditionally, predictions in game theory are based on equilibrium:
each agent's choice is optimal given the choices of the others.

These predictions rely on the assumption of equilibrium knowledge:
that agents correctly anticipate how other agents will behave.

In contexts with large numbers of agents, this assumption seems untenable.

We therefore consider an explicitly dynamic model of individual choice:

Each agent's behavior is described by a revision protocol, a myopic rule that tells him when and how to select a new strategy (Weibull (1995)).

A game, a revision protocol, and a population size define a Markov process $\{X_t^N\}$ on the set of population states.

Analyses follow one of two approaches, depending on the time span of interest:

medium-to-long run: analysis of ODE approximation.

very long run: analysis of stationary distribution.

Some objectives of evolutionary game theory

1. Define interesting classes of games and revision protocols
    - easy to describe
    - relevant to applications
    - have attractive theoretical properties

Some classes of games:

potential games (Monderer and Shapley (1996), Sandholm (2001))
(ex.: congestion games, common interest games, externality pricing games, decentralized control systems)

supermodular games (Topkis (1979))
(ex.: coordinated effort/investment games, arms races)

stable games (Hofbauer and Sandholm (2009))
(ex.: zero-sum games, wars of attrition, congestion w. cost asymmetries)

Some objectives of evolutionary game theory

1. Define interesting classes of games and revision protocols
   - easy to describe
   - relevant to applications
   - have attractive theoretical properties

Some classes of dynamics:

imitative (Taylor and Jonker (1978), Weibull (1995), Hofbauer (1995))

best response (Gilboa and Matsui (1991), Hofbauer (1995))

perturbed best response (Fudenberg and Levine (1998), Hofbauer and Sandholm (2002, 2007))

excess payoff (Brown and von Neumann (1950), Sandholm (2005))

pairwise comparison (Smith (1984), Sandholm (2010))

Some objectives of evolutionary game theory

2. Study the resulting dynamic processes
   2.1 Finite horizon, large population ⇒ approximation by ODE

Formal approximation theorem (Benaïm and Weibull (2003))

Analyses of dynamics:

Global convergence in potential/supermodular/stable games

Local stability of rest points / connections with ESS (Taylor and Jonker (1978), Ritzberger and Weibull (1995), Cressman (1997), Sandholm (2010))

Nonconvergence and its consequences

There are games in which all reasonable processes fail to converge to equilibrium. (Hofbauer and Swinkels (1996), Hart and Mas-Colell (2003))

There are games in which nearly all reasonable processes allow strictly dominated strategies to survive. (Hofbauer and Sandholm (2010))

Some objectives of evolutionary game theory

2. Study the resulting dynamic processes

   2.2 Infinite horizon, full support revision protocol
   ⇒ description using unique stationary distribution

Stochastic stability and equilibrium selection

(Foster and Young (1990), Kandori, Mailath, and Rob (1993), Young (1993))

Suppose that "error probabilities" are small (references above)
or that the population size is large (Binmore and Samuelson (1997)).

Then the stationary distribution typically is concentrated
on a single component of stochastically stable states.

Stochastic stability provides unique predictions of very long run behavior, even in games with multiple strict equilibria.

This talk considers stochastic stability analysis in two-strategy games:

1. under noisy best response protocols;
2. under imitative protocols.

Issues to be addressed:

- How does equilibrium selection depend on the choice of revision protocol?
- When is equilibrium selection independent of the limits (in η, in N) used to define stochastic stability?

Noisy best response protocols

The model:

$S = \{0, 1\}$ strategy set

N population size

$x \in \mathcal{X}^N = \{0, \frac{1}{N}, \frac{2}{N}, \ldots, 1\}$ fraction of agents playing strategy 1

$F^N : \mathcal{X}^N \to \mathbf{R}^2$ is the payoff function

$F_i^N(x)$ is the payoff to strategy i at state x.

We assume that F^N converges uniformly to some $F : [0, 1] \to \mathbf{R}^2$.

Agents are clever: decisions depend on the payoff difference

$\Delta F^N(x) = F_1^N(x) - F_0^N(x - \frac{1}{N})$.

The evolutionary process and its stationary distribution.

Agents receive revision opportunities according to independent, rate 1 Poisson processes.

When a current strategy i player receives a revision opportunity, he switches to strategy $j \neq i$ with probability $\rho^\eta(a) \in (0,1)$, where $a \in \mathbf{R}$ represents the current payoff advantage of j over i.

(The parameter $\eta > 0$ is called the noise level. More details shortly.)

These assumptions define an irreducible birth and death process $\{X_t^{N,\eta}\}_{t\geq 0}$ on $\mathcal{X}^N$ with stationary distribution

$$\frac{\mu_x^{N,\eta}}{\mu_0^{N,\eta}} = \prod_{j=1}^{Nx} \frac{\frac{N-j+1}{N} \cdot \rho^\eta\left(\Delta F^N(\frac{j}{N})\right)}{\frac{j}{N} \cdot \rho^\eta\left(-\Delta F^N(\frac{j-1}{N})\right)}.$$

How does $\mu^{N,\eta}$ behave as η and N approach their limits?

Revision protocols and their cost functions

$\rho^\eta(a)$ = the prob. of switching to a strategy with payoff advantage a.

We clearly want $\lim_{\eta \to 0} \rho^\eta(a) = \begin{cases} 1 & \text{if } a > 0, \\ 0 & \text{if } a < 0. \end{cases}$

Def. The cost of switching to a strategy with payoff disadvantage $d \in \mathbf{R}$ is

(1) $\quad c(d) = -\lim_{\eta \to 0} \eta \log \rho^\eta(-d). \qquad$ i.e., $\rho^\eta(-d) = \exp\left(-\eta^{-1}(c(d) + o(1))\right)$

Assumptions:

1. the limit in (1) exists for all $d \in \mathbf{R}$, with convergence uniform on compact intervals;
2. c is nondecreasing;
3. $c(d) = 0$ whenever $d < 0$;
4. $c(d) > 0$ whenever $d > 0$.

Examples

1. Best response with mutations (KMR (1993), Young (1993))

$$\rho^{\varepsilon}(a) = \begin{cases} 1-\varepsilon & \text{if } a > 0, \\ \varepsilon & \text{if } a \le 0. \end{cases}$$

$$\Rightarrow c(d) = 1 \text{ for } d \ge 0 \quad (\eta = \exp(-\varepsilon^{-1}))$$

2. Logit (Blume (1993, 1997))

$$\rho^{\eta}(a) = \frac{\exp(\eta^{-1}a)}{\exp(\eta^{-1}a)+1}$$

$$\Rightarrow c(d) = d \text{ for } d > 0$$

3. Probit (Myatt and Wallace (2003))

$$\rho^{\sigma^2}(a) = \mathbb{P}(\sigma Z + a > \sigma Z'), \quad Z, Z' \text{ i.i.d. standard normals}$$

$$\Rightarrow c(d) = \tfrac{1}{4}d^2 \text{ for } d > 0 \quad (\eta = \sigma^2)$$

Asymptotics of the stationary distribution

Define the relative cost function $\tilde{c}(d) = c(d) - c(-d)$ (nondecreasing, sign-preserving, odd).

Define the ordinal potential function

$$I(x) = \int_0^x \tilde{c}(\Delta F(y))\,\mathrm{d}y.$$

Note that $\operatorname{sgn}(I'(x)) = \operatorname{sgn}(F_1(x) - F_0(x))$.

Examples:

BRM: $I_{\text{sgn}}(x) = \int_0^x \operatorname{sgn}(\Delta F(y))\,\mathrm{d}y.$

logit: $I_1(x) = \int_0^x \Delta F(y)\,\mathrm{d}y$

probit: $I_2(x) = \int_0^x \frac{1}{4}\langle \Delta F(y)\rangle^2\,\mathrm{d}y,$ where $\langle a\rangle^2 = \operatorname{sgn}(a)\,a^2$

Shift I vertically so that its maximum is 0.

$$\Delta I(x) = I(x) - \max_{y\in[0,1]} I(y).$$

Theorem (informal):

For either order of limits in η and N, the exponential rates of decay of the stationary distribution weights $\mu_x^{N,\eta}$ are described by ΔI.

Theorem (formal):

The stationary distributions $\mu^{N,\eta}$ satisfy

i. $$\lim_{N\to\infty}\lim_{\eta\to 0}\max_{x\in\mathcal{X}^N}\left|\frac{\eta}{N}\log\mu_x^{N,\eta}-\Delta I(x)\right|=0 \quad \textit{and}$$

ii. $$\lim_{\eta\to 0}\lim_{N\to\infty}\max_{x\in\mathcal{X}^N}\left|\frac{\eta}{N}\log\mu_x^{N,\eta}-\Delta I(x)\right|=0.$$

Stochastic stability

Which states retain mass in $\mu^{N,\eta}$ when N is large and η is small?

We call x^* uniquely stochastically stable if for every open set $O \subseteq \mathbf{R}$ containing x^*,

$$\lim_{N\to\infty} \lim_{\eta\to 0} \mu^{N,\eta}(O) = \lim_{\eta\to 0} \lim_{N\to\infty} \mu^{N,\eta}(O) = 1.$$

Corollary: *If* $\operatorname{argmax}_{x\in[0,1]} I(x) = \{x^*\}$, *then* x^* *is uniquely stochastically stable.*

Let us focus on coordination games:

$\mathrm{sgn}(\Delta F(x)) = \mathrm{sgn}(x - x^*)$ for some $x^* \in (0,1)$.

$\Rightarrow$ three Nash equilibria: $e_0 = 0$, $e_1 = 1$, and x^*.

For any of our three protocols (BRM, logit, probit), it is easy to construct examples in which the chosen protocol selects a different equilibrium than the other two.

Example: Nonlinear Stag Hunt. $F_H(x) = h$ and $F_S(x) = sx^2$.

Nash equilibria: $e_H = 0$, $e_S = 1$, $x^* = \sqrt{h/s}$.

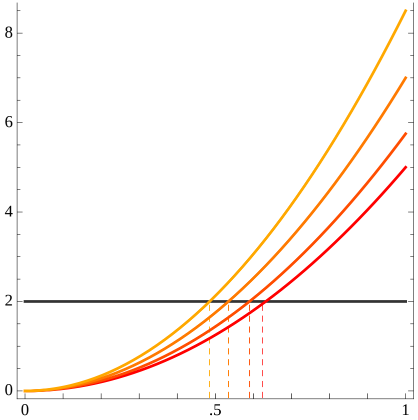

Figure: Payoffs and mixed equilibria when $h = 2$ and $s = 5, 5.75, 7$, and 8.5.

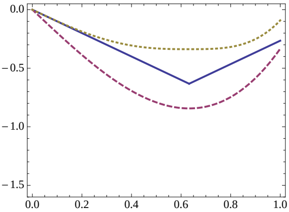

$h = 2,\ s = 5$

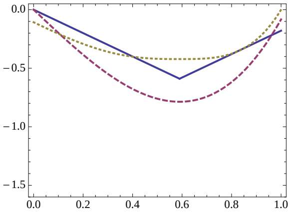

$h = 2,\ s = 5.75$

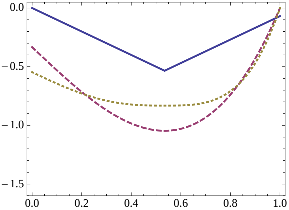

$h = 2,\ s = 7$

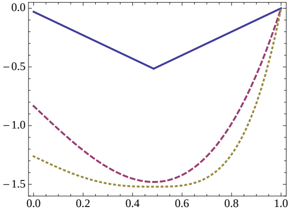

$h = 2,\ s = 8.5$

ΔI_{sgn} (blue), ΔI_1 (purple), and ΔI_2 (yellow).

Can we characterize the equilibria that are stochastically stable under every noisy best response protocol?

Define $\bar{G}_i(a) = \text{meas}(\{x \in [0,1] : F_i(x) - F_j(x) > a\}) = 1 - G_i(a)$.

In words: $\bar{G}_i(a)$ is the measure of the set of states at which the payoff to strategy i exceeds the payoff to strategy j by more than a.

Strategy i is stochastically dominant if

$\bar{G}_i(a) \geq \bar{G}_j(a)$ for all $a > 0$.

Aside: e_i has the bigger basin of attraction $\Leftrightarrow \bar{G}_i(0) \geq \bar{G}_j(0)$
In this case, strategy i is called risk dominant.

Theorem (informal):

A pure state is stochastically stable under every noisy best response protocol if and only if the corresponding strategy is stochastically dominant.

Theorem (formal):

(i) State e_i is weakly stochastically stable under every noisy best response protocol if and only if strategy i is stochastically dominant in F.

(ii) If strategy i is strictly stochastically dominant in F, then state e_i is uniquely stochastically stable under every noisy best response protocol.

Idea of the proof

Whether state 1 or state 0 is stochastically stable depends on whether $I(1)$ is greater than or less than $I(0) = 0$.

$$\begin{aligned} I(1) &= \int_0^1 \tilde{c}(\Delta F(y))\,\mathrm{d}y \\ &= \int_0^1 c(F_1(y) - F_0(y))\,\mathrm{d}y - \int_0^1 c(F_0(y) - F_1(y))\,\mathrm{d}y \\ &= \int_0^\infty c(a)\,\mathrm{d}G_1(a) - \int_0^\infty c(a)\,\mathrm{d}G_0(a). \end{aligned}$$

Thus the theorem reduces to a minor variation on the standard characterization of first-order stochastic dominance:

G_1 stochastically dominates G_0 if and only if $\int c\,\mathrm{d}G_1 \geq \int c\,\mathrm{d}G_0$ for every nondecreasing c.

Corollary: *Let F be a coordination game with linear payoffs.*

Then under any noisy best response protocol, state e_i is stochastically stable if and only if strategy i is risk dominant.

Proof: When F is linear, risk dominance and stochastic dominance are equivalent.

Imitation with committed agents

The model

Agents receive revision opportunities according to independent, rate 1 Poisson processes.

Suppose an i player receives a revision opportunity.

With probability $1 - \varepsilon$, he picks an opponent at random; if the opponent is playing $j \neq i$, the player switches with probability $r_{ij}(\pi)$. ($r_{01}(\cdot)$ and $r_{10}(\cdot)$ are bounded away from 0)

With the remaining probability ε, the player switches to j.

Binmore and Samuelson (1997) show that in this model, stochastic stability analysis can be sensitive to the order in which the limits in ε and N are taken.

Example

Committed agents

Suppose we introduce two additional agents:
  one always plays strategy 0,
  the other always plays strategy 1.

Now both strategies are always available for imitation
  (though sometimes just barely available when N is large).

In this case:

1. The evolutionary process is irreducible even without mutations.
2. (Theorem) If mutations are included, the two double limits agree.

Without mutations ($\varepsilon = 0$)

Define:

$$J(x) = \int_0^x \log \frac{r_{01}(F(y))}{r_{10}(F(y))}\, \mathrm{d}y;$$

$$\Delta J(x) = J(x) - \max_{y \in [0,1]} J(y).$$

Theorem:

As N grows large, the exponential rates of decay of the stationary distribution weights μ_x^N are described by $\Delta J(x)$.

Formally:

$$\lim_{N \to \infty} \max_{x \in \mathcal{X}^N} \left| \frac{1}{N} \log \mu_x^N - \Delta J(x) \right| = 0.$$

Mean dynamics and equilibrium selection: an example

Suppose N is large, and consider these two imitative protocols:

(S) $r_{ij}(\pi) = \lambda\left(\pi_j - m\right)$ imitation of success (Hofbauer (1995))

(D) $r_{ij}(\pi) = \lambda\left(M - \pi_i\right)$ imitation driven by dissatisfaction (Weibull (1995))

The medium-to-long run behavior of the Markov process $\{X_t^N\}$ is governed by a differential equation called the mean dynamic, which describes the expected motion of the process.

Under protocols (S) and (D) above, the mean dynamic is the replicator dynamic:

$$\dot{x} = p_x - q_x = \lambda x\,(1 - x)\,\Delta F(x).$$

Over finite time spans, behavior under (S) and (D) is indistinguishable.

What about very long run behavior?

Protocols (S) and (D) generate different rate of decay functions:

$$J_S(x) = \int_0^x \log \frac{F_1(y) - m}{F_0(y) - m}\, \mathrm{d}y;$$

$$J_D(x) = \int_0^x \log \frac{M - F_0(y)}{M - F_1(y)}\, \mathrm{d}y.$$

Now consider the coordination game

$$\begin{pmatrix} F_0(x) \\ F_1(x) \end{pmatrix} = \begin{pmatrix} 2 & 2 \\ 1 & 3 \end{pmatrix} \begin{pmatrix} 1 - x \\ x \end{pmatrix} = \begin{pmatrix} 2 \\ 1 + 2x \end{pmatrix}.$$

Its mixed equilibrium is $x^* = \frac{1}{2}$.

$$\begin{pmatrix} F_0(x) \\ F_1(x) \end{pmatrix} = \begin{pmatrix} 2 & 2 \\ 1 & 3 \end{pmatrix} \begin{pmatrix} 1-x \\ x \end{pmatrix} = \begin{pmatrix} 2 \\ 1+2x \end{pmatrix} \qquad x^* = \tfrac{1}{2}.$$

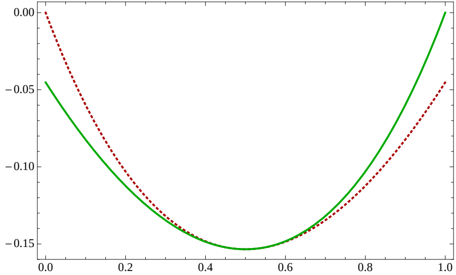

Figure: ΔJ_S (red) and ΔJ_D (green) when $m = 0$ and $M = 4$.

Protocol (S) selects the payoff dominated equilibrium,
while protocol (D) selects the payoff dominant equilibrium.

The same is true in any nearby game.

Intuition 1: Why are there distinct equilibrium selections?

Mean dynamics depend on differences in transition probabilities.

$$\dot{x} = p_x - q_x = x\,(1-x)\Big(r_{01}(F(x)) - r_{10}(F(x))\Big).$$

Stationary distributions depend on ratios of transition probabilities, and so on differences in logs of transition probabilities.

$$\frac{\mu_x^N}{\mu_0^N} = \prod_{j=1}^{Nx} \frac{p_{(j-1)/N}^N}{q_{j/N}^N}; \qquad J(x) = \int_0^x \Big(\log r_{01}(F(y)) - \log r_{10}(F(y))\Big)\,\mathrm{d}y.$$

Intuition 2: Why these particular selections?

Because of the concavity of $\log(\cdot)$ and Jensen's inequality.

Summary

We considered stochastic stability analysis in two-strategy games:

1. under noisy best response protocols;
2. under imitative protocols.

The issues we addressed:

- How does equilibrium selection depend on the choice of revision protocol?
- When is equilibrium selection independent of the limits (in η, in N) used to define stochastic stability?

Example: Imitation in a Hawk-Dove game

$$\begin{pmatrix} F_0(x) \\ F_1(x) \end{pmatrix} = \begin{pmatrix} 1 & 2 \\ 3 & 1 \end{pmatrix} \begin{pmatrix} 1-x \\ x \end{pmatrix} = \begin{pmatrix} 1+x \\ 3-2x \end{pmatrix} \quad \Rightarrow \quad x^* = \tfrac{2}{3}$$

Consider imitation of success: $r_{ij}(\pi) = \lambda\, \pi_j$

This generates an evolutionary process with transition probabilities

$$p_x^{N,\varepsilon} = (1-x) \cdot \Big((1-\varepsilon)\, \lambda \tfrac{Nx}{N-1} \pi_1 + \varepsilon \Big),$$
$$q_x^{N,\varepsilon} = x \cdot \Big((1-\varepsilon)\, \lambda \tfrac{N(1-x)}{N-1} \pi_0 + \varepsilon \Big).$$

When N is large, the mean dynamic of the process is

$$\begin{aligned} \dot{x} &\approx p_x^{N,\varepsilon} - q_x^{N,\varepsilon} \\ &\approx (1-\varepsilon)\,\lambda\, x\,(1-x)\,\Delta F(x) \;+\; 2\varepsilon\,(\tfrac{1}{2} - x). \end{aligned}$$

When $\varepsilon = 0$, this is the replicator dynamic. There is a stable rest point at x^* and unstable rest points at e_0 and e_1.

When $\varepsilon > 0$, the two boundary rest points disappear, leaving a globally stable rest point near x^*.

We argue next that in the Hawk-Dove game,

1. the prediction based on large population stochastic stability agrees with that of the mean dynamic;
2. the prediction based on small noise stochastic stability does not.

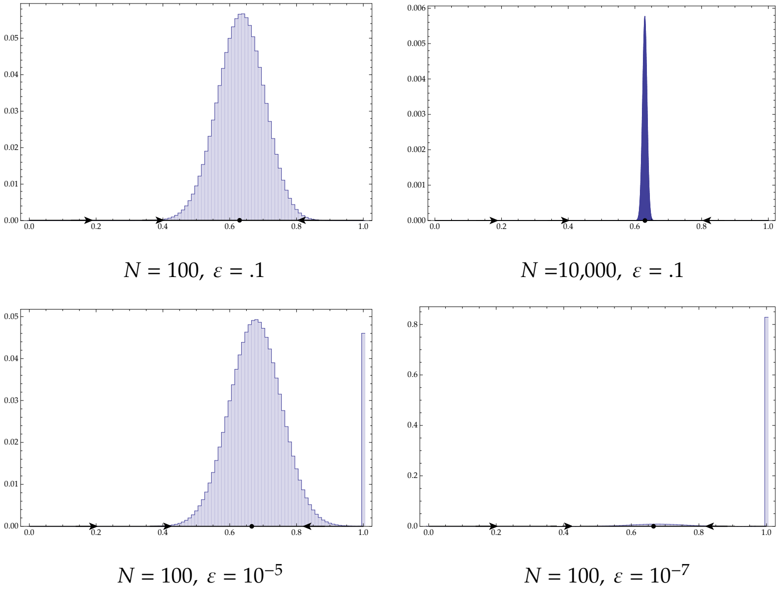
0.05
0.04
0.03
0.02
0.01
0.00
0.0
0.2
0.4
0.6
0.8
1.0
$N = 100, \ \varepsilon = .1$
0.006
0.005
0.004
0.003
0.002
0.001
0.000
$N = 10{,}000, \ \varepsilon = .1$
$N = 100, \ \varepsilon = 10^{-5}$
$N = 100, \ \varepsilon = 10^{-7}$

Back

A formal analysis (for all two-strategy games)

Define:

$$J(x) = \int_0^x \log \frac{r_{01}(F(y))}{r_{10}(F(y))}\, \mathrm{d}y;$$

$$\Delta J(x) = J(x) - \max_{y \in [0,1]} J(y).$$

Theorem: *Under imitative protocols with mutations, the stationary distributions $\mu^{N,\varepsilon}$ satisfy*

i. $\lim_{N \to \infty} \lim_{\varepsilon \to 0} \frac{1}{N} \log \frac{\mu_1^{N,\varepsilon}}{\mu_0^{N,\varepsilon}} = J(1)$,
while $\frac{\mu_x^{N,\varepsilon}}{\mu_0^{N,\varepsilon}}, \frac{\mu_x^{N,\varepsilon}}{\mu_1^{N,\varepsilon}} \in O(\varepsilon)$ *when* $x \in \mathcal{X}^N - \{0, 1\}$;

ii. $\lim_{\varepsilon \to 0} \lim_{N \to \infty} \max_{x \in \mathcal{X}^N} \left| \frac{1}{N} \log \mu_x^{N,\varepsilon} - \Delta J(x) \right| = 0.$

Exponential protocols and potential games

Potential games (Monderer and Shapley (1996), Sandholm (2001))

F is a potential game if it admits a potential function $f : \mathbf{R}^n \to \mathbf{R}$:

$$\nabla f(x) = F(x) \text{ for all } x \in X \quad (\text{i.e., } \frac{\partial f}{\partial x_i}(x) = F_i(x)).$$

Potential games are characterized by externality symmetry:

(ES) $DF(x)$ is symmetric for all $x \in X$.

(Other interesting classes of games:

stable games: $DF(x)$ negative semidefinite.

supermodular games: $\frac{\partial(F_{i+1} - F_i)}{\partial(e_{j+1} - e_j)}(x) \geq 0$.)

Examples

1. Random matching in common interest games

Normal form game A is a common interest game game if A is symmetric:

$$A_{ji} = A_{ij} \text{ for all } i, j \in S.$$

Since $F(x) = Ax$, we have $DF(x) = A$, so F is a potential game.

Its potential function is $f(x) = \frac{1}{2}x'Ax$.

2. Congestion games

$F_i(x) = -\sum_{\lambda \in \Lambda_i} c_\lambda(u_\lambda(x))$	payoff to path i
x_i	mass of players choosing path i
$u_\lambda(x) = \sum_{i:\, \lambda \in \Lambda_i} x_i$	utilization of link λ
$c_\lambda(u_\lambda)$	(increasing) cost of delay on link λ

F is a potential game with potential function $f(x) = -\sum_{\lambda \in \Lambda_i} \int_0^{u_\lambda(x)} c_\lambda(v)\,\mathrm{d}v$.

Example: A common interest game and its potential function

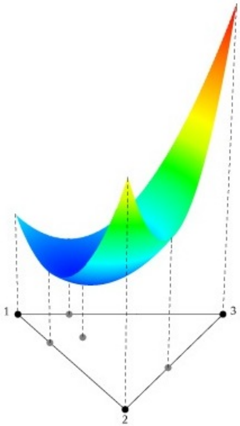

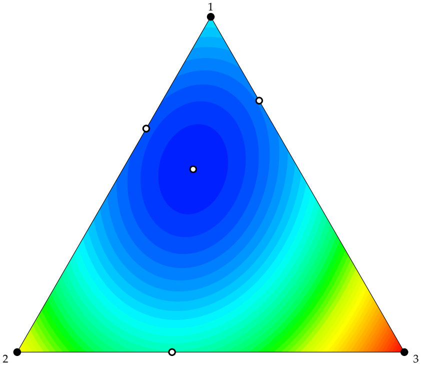

$$F(x) = \begin{pmatrix} F_1(x) \\ F_2(x) \\ F_3(x) \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 3 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}; \quad f(x) = \frac{1}{2}\left((x_1)^2 + 2(x_2)^2 + 3(x_3)^2\right).$$

We focus on behavior in finite population potential games.

F^N is a *finite population potential game* if there is a $f^N\colon \mathcal{X}^N \to \mathbf{R}$ such that

$$F^N_j(x + \tfrac{1}{N}(e_j - e_i)) - F^N_i(x) = f^N(x + \tfrac{1}{N}(e_j - e_i)) - f^N(x) \quad \forall x \in \mathcal{X}^N_i, i \in S.$$

ex. $F^N(x) = Ax, A$ symmetric

$\Rightarrow\; f^N(x) = \frac{1}{2}N\,x'Ax + \frac{1}{2}\sum_{k\in S} A_{kk}x_k.$

$\therefore$ as $F^N \to F$, $\frac{1}{N}f^N \to f$.

We also assume agents are clever: when making decisions, each agent accounts for his (size $\frac{1}{N}$) impact on the population state.

Exponential protocols

We call $\rho : \mathbf{R}^n \times X \to \mathbf{R}^{n\times n}$ a direct exponential protocol if

$$\rho_{ij}(\pi, x) = \frac{\exp(\eta^{-1}\psi(\pi_i, \pi_j))}{d_{ij}(\pi)},$$

where the functions $\psi : \mathbf{R}^2 \to \mathbf{R}$ and $d : \mathbf{R}^n \to (0,\infty)^{n\times n}$ satisfy

(*) $\quad \psi(\pi_i, \pi_j) - \psi(\pi_j, \pi_i) = \pi_j - \pi_i,$ and

(§) $\quad d_{ij}(\pi) = d_{ji}(\pi).$

(*) and (§) allow choice probs to depend on payoffs in a variety of ways:

pos. dependence on candidate payoff:	$\psi(\pi_i, \pi_j) = \pi_j$
neg. dependence on current payoff:	$\psi(\pi_i, \pi_j) = -\pi_i$
pos. dependence on payoff difference:	$\psi(\pi_i, \pi_j) = \frac{1}{2}(\pi_j - \pi_i)$
pos. dependence on pos. payoff difference:	$\psi(\pi_i, \pi_j) = [\pi_j - \pi_i]_+$
neg. dependence on neg. payoff difference:	$\psi(\pi_i, \pi_j) = -[\pi_j - \pi_i]_-$

Theorem: *The process $\{X_t^{N,\eta}\}$ is reversible with stationary distribution*

$$(\dagger) \qquad \mu_x^{N,\eta} = \frac{1}{K^N} \frac{N!}{\prod_{k \in S} (Nx_k)!} \exp(\eta^{-1} f^N(x)).$$

Recall that $f^N(x) \approx Nf(x)$.

$\Rightarrow$ stationary distribution weights are exponential in both η^{-1} and N.

The weight state x receives depends on:

1. the value of potential at x,
2. the weight on Nx in *multinomial*$(N; \frac{1}{n}, \ldots, \frac{1}{n})$.

If η is small and N is large, states near $\text{argmax}_x f(x)$ receive most of the weight in $\mu^{N,\eta}$.

(More precise descriptions of the asymptotics of $\mu^{N,\eta}$ are possible. . .)

Imitative exponential protocols

We call $\rho : \mathbf{R}^n \times X \to \mathbf{R}^{n\times n}$ an *imitative exponential protocol* if

$$\rho_{ij}(\pi, x) = x_j \, \frac{\exp(\eta^{-1}\psi(\pi_i, \pi_j))}{d_{ij}(\pi, x)},$$

where the functions $\psi : \mathbf{R}^2 \to \mathbf{R}$ and $d : \mathbf{R}^n \to (0, \infty)^{n\times n}$ satisfy

(*) $\psi(\pi_i, \pi_j) - \psi(\pi_j, \pi_i) = \pi_j - \pi_i$, and

(§) $d_{ij}(\pi, x) = d_{ji}(\pi, y)$.

Under a direct protocol, agents select candidate strategies "from a list".

Under an imitative protocol, agents select candidate strategies by observing opponents' current choices.

The Markov process generated by an imitative protocol is not irreducible. (For example, vertices are absorbing states).

To restore irreducibility, we introduce one committed agent for each strategy $i \in S$.

This changes transition probabilities by $O(\frac{1}{N})$.

Theorem: *The process $\{X_t^{N,\eta}\}$ is reversible with stationary distribution*

$$(\ddagger) \qquad \mu_x^{N,\eta} = \frac{1}{\kappa^N} \exp(\eta^{-1} f^N(x)).$$

Compare to the stationary distribution under direct protocols:

$$(\dagger) \qquad \mu_x^{N,\eta} = \frac{1}{K^N} \frac{N!}{\prod_{k \in S} (N x_k)!} \exp(\eta^{-1} f^N(x)).$$

Unlike in the case of direct choice, here there is no bias toward states with high multinomial weights.

Note: If f^N is constant, (‡) is the uniform distribution on $\mathcal{X}^N$, while (†) is a multinomial distribution.

Directions for future research

1. Equilibrium selection in the large N limit in nonreversible cases
2. Modifications to address waiting times
   - time varying noise levels (as in simulated annealing)
   - local interaction structures (lattices, general networks)
3. Modeling local interactions as interacting particle systems (see Blume (1993, 1995), Ellison (2000); also Nowak, Durrett, . . .)
4. Hydrodynamic limits (and others)
   - global interactions $\Rightarrow$ ODE limit
   - local interactions $\Rightarrow$ PDE limit