

EVOLUTIONARY GAME THEORY

Population dynamics, social learning, and conventions

Jörgen Weibull

Delhi Winter School, December 2017

1 Multi-population selection dynamics

- The replicator dynamic readily generalizes to arbitrary finite n -player games
- Domain now: finite normal-form games $G = \langle I, S, u \rangle$
 - $I = \{1, \dots, n\}$ the set of *player roles*
 - $S = \times_{i \in I} S_i$ the set of *pure-strategy profiles*, $s = (s_1, s_2, \dots, s_n)$
 - $u : S \rightarrow \mathbb{R}^n$ the *combined payoff function*, $u_i(s) \in \mathbb{R}$ being the payoff to the individual in player role $i \in I$
- A continuum "*player population*" of individuals for each player role, and all individuals play pure strategies

- Let $\Delta(S_i)$ be the mixed-strategy simplex for player role i
- Let $\square(S)$ be the polyhedron of mixed-strategy profiles,
$$\square(S) = \times_{i \in I} \Delta(S_i)$$
- Extend u from S to $\square(S)$ in the usual way, so that $u_i(x)$ is player i 's expected payoff when mixed-strategy profile $x = (x_1, \dots, x_n)$ is played

1.1 Examples

1. Taylor (1979):

$$\dot{x}_{ih} = [u_i(\mathbf{1}_{ih}, x_{-i}) - u_i(x)] \cdot x_{ih} \quad \forall i \in I, h \in S_i$$

(a) The *growth rate* of the subpopulation of h -strategists within player population i equals the absolute payoff advantage of its strategy:

$$g_{ih}(x) = u_i(\mathbf{1}_{ih}, x_{-i}) - u_i(x)$$


2. Maynard-Smith (1982) suggested instead population growth rates based on *relative* payoff advantage (assuming $u_i(x) > 0$ for all i and x):

$$\dot{x}_{ih} = \left[\frac{u_i(\mathbf{1}_{ih}, x_{-i})}{u_i(x)} - \mathbf{1} \right] \cdot x_{ih}$$


$$g_{ih}(x) = \frac{u_i(\mathbf{1}_{ih}, x_{-i})}{u_i(x)} - \mathbf{1}$$

1.2 Two-player games

Example 1.1 (Coordination) *In a 2×2 -coordination game each strict equilibrium is asymptotically stable and the mixed equilibrium is unstable:*




Example 1.2 (Hawk-Dove) *Again the interior NE is unstable while both strict equilibria are asymptotically stable. The long-run outcome is again “history dependent”. Generic polarization between player. Explain the difference from the single-population dynamics!*



Example 1.3 (Entry-deterrence) *Unique SPE but infinitely many other NE.*


Extensive form:



Normal form:

	<i>C</i>	<i>F</i>
<i>A</i>	1, 3	1, 3
<i>E</i>	2, 2	0, 0

Solution orbits of the Taylor replicator dynamic:



2 General selection dynamics

- Arbitrary finite games $G = \langle I, S, u \rangle$
- Generalize the dynamics to allow for a wide range of imitation behaviors and social learning:

$$\dot{x}_{ih} = g_{ih}(x) x_{ih}$$

where g is *regular* (Lipschitz continuous and $x_i \cdot g_i(x) \equiv 0$ for all i)

Definition 2.1 A growth-rate function g is **payoff monotonic** if

$$u_i(\mathbf{1}_{ih}, x_{-i}) > u_i(\mathbf{1}_{ik}, x_{-i}) \Rightarrow g_{ih}(x) > g_{ik}(x)$$

Definition 2.2 A growth-rate function g is **payoff positive** if

$$\begin{aligned} u_i(\mathbf{1}_{ih}, x_{-i}) > u_i(x) &\Rightarrow g_{ih}(x) > 0 \\ u_i(\mathbf{1}_{ih}, x_{-i}) < u_i(x) &\Rightarrow g_{ih}(x) < 0 \end{aligned}$$

- For any mixed-strategy profile x and any player i , let

$$B_i(x) = \{h \in S_i : u_i(\mathbf{1}_{ih}, x_{-i}) > u_i(x)\}$$

- these are the pure strategies that yield payoffs above average

Definition 2.3 *A growth-rate function g is weakly payoff-positive if:*

$$B_i(x) \neq \emptyset \Rightarrow g_{ih}(x) > 0 \text{ for some } h \in B_i(x)$$

Proposition 2.1 (Weibull, 1995) *For any regular weakly payoff-positive dynamic:*

(a) $x \in \square$ Lyapunov stable $\Rightarrow x \in \square^{NE}$

(b) $x^o \in \text{int}(\square) \wedge \lim_{t \rightarrow +\infty} \xi(x^o, t) = x^* \Rightarrow x^* \in \square^{NE}$

- Property (c) of the single-population replicator dynamics in symmetric two-player games can be showed to hold for arbitrary finite games if the growth rate has a certain monotonicity property with respect to payoffs:


Definition 2.4 (Hofbauer and Weibull, 1996) *A growth-rate function g is convex monotonic if:*

$$u_i(y_i, x_{-i}) > u_i(1_{ih}, x_{-i}) \Rightarrow y_i \cdot g_i(x) > g_{ih}(x)$$

Proposition 2.2 (Hofbauer and Weibull, 1996) *For any regular convex-monotonic selection dynamic: if $x^0 \in \text{int}(\square)$ and $k \in S_i$ is iteratively strictly dominated for player i , then*

$$\lim_{t \rightarrow +\infty} \xi_{ik}(t, x^0) = 0$$

Example 2.1 (outside-option game)



Three subgame-perfect equilibria, but only one, $s^* = (Ra, A)$, is compatible with “forward-induction”.

The purely reduced normal form:

	A	B
L	$2, v$	$2, v$
Ra	$3, 1$	$0, 0$
Rb	$0, 0$	$1, 3$

3 Stochastic evolution in finite populations

- We recall that any evolutionary process has two parts: *mutation* (creation of varieties) and *selection*
- However, all the above dynamics are pure selection dynamics. No mutations occur. Robustness against mutations was studied by considering dynamic stability (Lyapunov and asymptotic).
- Moreover, the above dynamics are all deterministic and treat each player population as a continuum.
- What can be said if the player-populations are *finite* and subject to *perpetually occurring random mutations* alongside selection?
- This is the research task we now tackle

3.1 Stochastic population processes

[Benaïm & Weibull, 2003]

- Consider any finite n -player game
- For each player role $i \in I$ there is a player population of constant and finite size N
- Individuals always play pure strategies (just as in the deterministic selection dynamics and in Nash's mass-action interpretation)
- A individual in player population i who plays a pure strategy $h \in S_i$ will be called an *h -strategist*

- At discrete times $t = 0, \delta, 2\delta, 3\delta, \dots$, for $\delta = 1/N$, one individual is randomly drawn for "strategy review"
- Equal probability for all individuals in the whole population to be drawn (i.i.d.)
- For each individual, the average time between two successive opportunities for strategy review is independent of N

Definition 3.1 A population state is a mixed-strategy profile $x = (x_1, \dots, x_n) \in \square(S)$ where each mixed strategy $x_i \in \Delta(S_i)$ is such that Nx_{ih} is a non-negative integer. Denote the (finite) set of population states $\Theta^N \subset \square(S)$.

Definition 3.2 At times $t \in D^N = \{0, 1/N, 2/N, \dots\}$, let the vector $X^N(t) = (X_1^N(t), \dots, X_n^N(t)) \in \Theta^N$ be the current population shares of h -strategists, for all pure strategies $h \in S_i$ and all player populations $i \in I$.

- We analyze population processes $\langle X^N(t) \rangle_{t \in D^N}$ that are (time homogeneous) Markov chains with transition probabilities

$$p_{ikh}^N(x) = \Pr \left[X_i^N(t + \frac{1}{N}) = x_i + \frac{1}{N} (\mathbf{1}_{ih} - \mathbf{1}_{ik}) \mid X^N(t) = x \right]$$

3.2 The mean-field equations

- The expected *net increase* in the population share of h -strategists in player population i , in population state x :

$$F_{ih}^N(x) = \sum_{k \neq h} p_{ikh}^N(x) - \sum_{k \neq h} p_{ihk}^N(x)$$

- Assume $F^N \rightarrow F$ as $N \rightarrow \infty$ (which we will see is the case in canonical examples)
- The *mean-field equations*, the flow approximation of the stochastic process (for N large) is

$$\dot{x}_{ih} = F_{ih}(x) \quad \forall i \in I, h \in S_i, x \in \square(S)$$

Q: Are the solution to these mean-field equations "good approximations" of the the movements of the stochastic population process?

3.3 Why care?

- Because if, for some stochastic population process, the associated mean-field F is of the form

$$F_{ih}(x) = g_{ih}(x) x_{ih}$$

for some growth-rate function g , then we can use the game-theoretic results for deterministic selection dynamics to obtain game-theoretic results for stochastic populationo processes

3.3.1 Example: aspiration and imitation (or Herbert Simon meets John Nash)

- Herbert Simon (1916-2001, pioneer in behavioral economics, coining the term "bounded rationality", and 1978 economics Nobel laureate) suggested that real economic agents do not maximize, they "satisfice" (seek to meet aspiration levels)
- Suppose individuals use the rule of thumb: "If I am dissatisfied with the performance of my current strategy, then I will imitate a randomly drawn individual in my own player population." [Gale, Binmore and Samuelson (1995), Björnerstedt and Weibull (1996), and Binmore and Samuelson (1997)]
- Suppose that the aspiration levels within each player population i are statistically independent and uniformly distributed on an interval $[a_i(x), b_i(x)]$ that contains the range of the payoff function π_i

- Then

$$F_{ih}(x) = \frac{1}{n} \frac{u_i(1_{ih}, x_{-i}) - u_i(x)}{b_i(x) - a_i(x)} \cdot x_{ih} .$$

- If $a_i(x) \equiv \alpha_i$ and $b_i(x) \equiv \beta_i$ for some $\alpha_i < \beta_i$, then this is but a player-specific time rescaling of the Taylor (1979) multi-population replicator dynamic.
- If instead $a_i(x) \equiv \alpha_i u_i(x)$ and $b_i(x) \equiv \beta_i u_i(x)$ for some $\alpha_i < \beta_i$ (and all payoffs are positive), then we obtain a player-specific time rescaling of the Maynard Smith (1982) multi-population replicator dynamics.
- If the mean field is a good approximation of these stochastic processes, then we have established that "Herbert Simon asymptotically meets John Nash": If a mean-field solution converges from some interior initial state, then the limit point is a Nash equilibrium, and then also

the stochastic population process will probabilistically converge to Nash equilibrium (in a precise sense)

3.4 The key lemma

- The answer to the question **Q** raised before is a qualified "yes". It hinges on the fact that the stochastic evolutionary process is exponentially well approximated by its mean field over bounded time intervals:

Proposition 3.1 (Benaïm & Weibull, 2003) *For every $T > 0$ there exist constants $c, K > 0$ such that for all $\varepsilon > 0$ and all N large enough:*

$$\Pr \left[\max_{0 \leq t \leq T} \|X^N(t) - \xi(t, x^o)\| \geq \varepsilon \mid X^N(0) = x_0 \right] \leq K e^{-c\varepsilon^2 N}$$

- In other words: the probability for any ε -deviation tends exponentially to zero as population size $N \rightarrow \infty$
- Using this, also asymptotic results can be obtain (see paper)

4 Social learning and stochastic stability

- Models of best replies to "recent history of play" with perpetual perturbations:
 - Young P. (1993a): "The evolution of conventions", *Econometrica*
 - Young P. (1993b): "An evolutionary model of bargaining", *Journal of Economic Theory*
 - Hurkens S. (1995): "Learning by forgetful players", *Games and Economic Behavior*
 - Young (1998): *Individual Strategy and Social Structure*. Princeton University Press.

4.1 Young's model in a nutshell

- Arbitrary finite game $G = \langle I, S, u \rangle$
- For each player role i a *finite* population of size N
- Each round $t = 0, 1, 2, \dots$ one individual is randomly drawn from each player population, and each of these n individuals draws a sample of size k (without replacement) from the $m \geq k$ last rounds of play, and these individuals play the game (see below)
- A *state* of this Markov chain is the m -history of recent play; the m most recent pure-strategy profiles
- A *successor history*, after any history, is obtained by adding the new strategy profile and deleting the oldest

- The *unperturbed process*: each individual called to play always plays a best reply against its k -sample from the current m -history
 - statistically independent samples across individuals in the same round and across periods
- The *perturbed process*:
 - with probability $1 - \varepsilon$: an individual called to play plays a best reply against its k -sample from the current m -history
 - with probability ε : the individual instead draws a pure strategy at random (uniformly)
- The perturbed process is *irreducible* (it can with positive probability get from any given state to any other given state in a finite number of rounds)

- If k/m is sufficiently small and m sufficiently large, then the perturbed process is also *aperiodic* (there exist no cycles that it cannot leave)
- A classical theorem in the theory of Markov chains is that any irreducible and aperiodic Markov chain with finite state space has a unique invariant probability distribution over its state space, a distribution to which it converges from any initial state
- Applied to the present class of perturbed Markov chains: There exists a unique invariant probability distribution μ^ε over m -histories of play, to which the process converges from any initial m -history.
- Next step: Take the limit of μ^ε as the probability for mistakes tends to zero, $\varepsilon \rightarrow 0$, this leads to $\mu^\varepsilon \rightarrow \mu^*$

- The so obtained limiting distribution μ^* is an invariant distribution under the unperturbed process (which may admit multiple invariant distributions)
- Any strategy profile that appears in a history in the support of the limiting distribution μ^* is called *stochastically stable*
- If a game has a unique stochastically stable strategy profile that repeats itself in every period, then it is a social *convention*

Proposition 4.1 (Young, 1993) *Let G be a 2×2 -coordination game. If $k/m \leq 1/2$ then the unperturbed process ($\varepsilon = 0$) converges from any initial state with probability one to one of the two strict equilibria. If $k/m \leq 1/2$ and m is sufficiently large, then the perturbed process ($\varepsilon > 0$) has a unique invariant distribution and, in the limit as $\varepsilon \rightarrow 0$, this places probability one on (repeated play of) the risk dominant equilibrium.*

- Risk dominance (Harsanyi & Selten, 1988)
- The mixed NE, although perfect, is never selected (in line with ESS and the replicator dynamic)

Example 4.1 Consider the coordination game

	a	b
a	4, 4	0, 2
b	2, 0	3, 3

The strict equilibrium (a, a) Pareto dominates the strict equilibrium (b, b) , but (b, b) risk dominates (a, a) .

- A generalization to arbitrary finite games:

Definition 4.1 (Young, 1998) *A finite normal-form game $G = (I, S, u)$ has non-degenerate best replies (NDBR) if, every pure strategy is either not a best reply to any mixed-strategy profile, or it is a best reply to a set of mixed-strategy profiles with non non-empty interior.*

- This is a generic property of finite normal-form games.

- For each player role i and any nonempty $T_i \subseteq S_i$, let $T = \times_{i \in N} T_i$ and let $\square(T) = \times_{i \in I} \Delta(T_i)$.
- Such block may be "closed under rational behavior" in a precise sense:

Definition 4.2 (Basu and Weibull, 1991) T is a **CURB** set if

$$\beta[\square(T)] \subseteq T.$$

- Examples: the entry-deterrence game, the outside-option game, coordination games. They all have (unique) social conventions.

Theorem 4.2 (Young, 1998) *Let G be a finite game with the NDBR property. The unperturbed process converges with probability one to a minimal CURB set if k/m is sufficiently small. In the limit as $\varepsilon \rightarrow 0$, the limit invariant distribution μ^* places unit probability on the strategy profiles on the minimal CURB set that have minimal stochastic potential.*

- Hence, set-valued conventions exist and can be mathematically identified