An Introduction to Evolutionary Games

Ross Cressman
Department of Mathematics
Wilfrid Laurier, Waterloo, Canada

http://www.wlu.ca/ wwwmath/faculty/cressman
/rcressman.research.html

Evolution of Cooperation

Hamilton's Rule
Kin Selection

Evolutionary game theory began in the 1970s when John Maynard Smith introduced the intuitive idea of an evolutionarily stable strategy (ESS) to predict the eventual behaviour of individuals (i.e. their strategy choice) without analyzing complex dynamical systems of evolution that may ultimately depend on many factors such as genetics and/or mating systems.

Definition: (Maynard Smith, 1982) An ESS is a strategy such that, if all members of a population adopt it, then no mutant strategy could invade the population under the influence of natural selection.

He goes on to say that, to make this definition precise, assumptions must be made about the evolving population.

Assume we have a mixed population consisting of mostly p^* individuals with a few individuals using strategy p. That is, the strategy distribution in the population is

$$(1-\epsilon)p^* + \epsilon p$$

where $\epsilon > 0$ is the small frequency of $p-$users in the population.

Let the fitness (i.e. reproductive success) of an individual using strategy q in this mixed population be

$$\pi(q, (1-\epsilon)p^* + \epsilon p).$$

Assume we have a mixed population consisting of mostly p^* individuals with a few individuals using strategy p. That is, the strategy distribution in the population is

$$(1-\epsilon)p^* + \epsilon p$$

where $\epsilon > 0$ is the small frequency of $p-$users in the population.

Let the fitness (i.e. reproductive success) of an individual using strategy q in this mixed population be

$$\pi(q, (1-\epsilon)p^* + \epsilon p).$$

Then one interpretation of Maynard Smith's requirement for p^* to be an ESS is that, for all $p \neq p^*$,

$$\pi(p, (1-\epsilon)p^* + \epsilon p) < \pi(p^*, (1-\epsilon)p^* + \epsilon p).$$

for all $\epsilon > 0$ sufficiently small.

At an ESS, for all $p \neq p^*$,

$$\pi(p,(1-\epsilon)p^* + \epsilon p) < \pi(p^*,(1-\epsilon)p^* + \epsilon p).$$

for all $\epsilon > 0$ sufficiently small.

By continuity as $\epsilon \to 0$, we have, for all $p \neq p^*$,

$$\pi(p,p^*) \leq \pi(p^*,p^*).$$

That is, we have our first result; namely, if p^* is an ESS, then it must be a (symmetric) Nash equilibrium (NE) for the game where $\pi(q,p^*)$ is the payoff to a player using strategy q when interacting with strategy p^*.

If $\pi(q,(1-\epsilon)p^* + \epsilon p)$ is linear in $\epsilon > 0$ and p^* is an ESS, then we also have the requirement that

$$\pi(p,p) < \pi(p^*,p).$$

whenever $\pi(p,p^*) = \pi(p^*,p^*)$.

Matrix Games (Symmetric normal form games with finitely many strategies)

Pure strategies: $e_1, e_2, ..., e_m$
Mixed strategies: $p = (p_1, ..., p_m) \in \Delta^m$ where

$$\Delta^m \equiv \{(p_1, ..., p_m) \mid \sum_{i=1}^{m} p_i, p_i \geq 0\}$$

Note: Pure strategies are unit vectors in $\mathbf{R}^m$.
Payoffs: If $\pi(e_i, e_j) = A_{ij}$, then

$$\pi(p, q) = \sum_{i,j=1}^{m} p_i A_{ij} q_j \equiv p \cdot Aq.$$

Theorem (Maynard Smith, 1982):

$p^* \in \Delta^m$ is an ESS of a matrix game iff
(a) $\pi(p, p^*) \leq \pi(p^*, p^*)$ for all $p \in \Delta^m$
(NE condition)
(b) If $\pi(p, p^*) = \pi(p^*, p^*)$ and $p \neq p^*$,
then $\pi(p, p) < \pi(p^*, p)$ (stability condition).

Theorem (Hofbauer & Sigmund, 1998):
$p^* \in \Delta^m$ is an ESS of a matrix game iff

$$\pi(p^*, p) > \pi(p, p)$$

for all $p \in \Delta^m$ sufficiently close (but not equal) to p^*.

Note 1: This characterization of ESS is also valid when $\pi(p, q)$ is of the form $\sum_{i=1}^{m} p_i f_i(q)$ and $f_i(q)$ is nonlinear in the components of q. The model is then called a frequency-dependent population game.

Note 2: A p^* that satisfies this inequality is also called locally superior (Weibull, 1995).

Replicator Equation

Some Assumptions:
1. Each individual receives a payoff from one random pairwise interaction per unit time.
2. Payoffs translate directly into fitness (i.e. reproductive success).
3. Offspring are clones of their only parent.

If $n_i(t)$ is the number of individuals using strategy e_i at time t, then for matrix games,

$$\dot{n}_i = n_i \pi(e_i, p).$$

Here $N \equiv n_1 + ... + n_m$ is the total population size and $p_i \equiv n_i/N$ is the frequency of strategy i in the population. From calculus,

$$\dot{p}_i = p_i(\pi(e_i, p) - \pi(p, p))$$

This is called the **replicator equation**.

Theorem (Taylor & Jonker, 1978):

If $p^* \in \Delta^m$ is an ESS of a matrix game, then it is locally asymptotically stable (l.a.s.) for the replicator equation.

Notes: 1. Taylor & Jonker proved this using linearization techniques but the best proof is based on the above characterization of an ESS as locally superior and works for frequency-dependent population games as well. This proof also shows that a matrix-game ESS p^* in the interior of Δ^m (that is, a completely mixed strategy) is globally asymptotically stable for the replicator equation since it is globally superior.

Theorem (Taylor & Jonker, 1978):

If $p^* \in \Delta^m$ is an ESS of a matrix game, then it is locally asymptotically stable (l.a.s.) for the replicator equation.

Notes: 1. Taylor & Jonker proved this using linearization techniques but the best proof is based on the above characterization of an ESS as locally superior and works for frequency-dependent population games as well. This proof also shows that a matrix-game ESS p^* in the interior of Δ^m (that is, a completely mixed strategy) is globally asymptotically stable for the replicator equation since it is globally superior.

2. The converse is not true. However, if individuals can play mixed strategies, then $p^* \in \Delta^m$ is an ESS iff it is l.a.s. for all mixed-strategy replicator equations. In fact, it could be argued that the mixed-strategy replicator equation with only two strategies p^* and p is Maynard Smith's original intuition of an uninvadable strategy.

Hawk-Dove Game

$$\begin{array}{ccc} & H & D \\ H & \frac{V}{2}-C & V \\ D & 0 & \frac{V}{2} \end{array}$$

(a) If $\frac{V}{2} \geq C$, then Hawk is the only ESS.

(b) If $\frac{V}{2} < C$, the unique ESS is $\frac{1}{C}(\frac{V}{2}, C - \frac{V}{2})$.

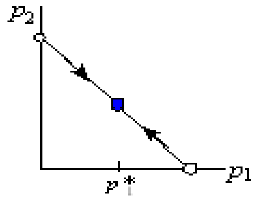

(c) Hawk-Dove

The replicator equation evolves to the ESS.

Hawk-Dove Game

$$\begin{array}{ccc} & H & D \\ \begin{array}{c} H \\ D \end{array} & \left[\begin{array}{c} \frac{V}{2} - C \\ 0 \end{array}\right. & \left.\begin{array}{c} V \\ \frac{V}{2} \end{array}\right] \end{array}$$

(a) If $\frac{V}{2} \geq C$, then Hawk is the only ESS.

(b) If $\frac{V}{2} < C$, the unique ESS is $\frac{1}{C}(\frac{V}{2}, C - \frac{V}{2})$.

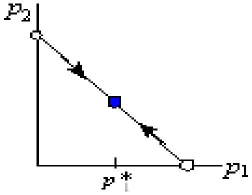

(c) Hawk-Dove

The replicator equation evolves to the ESS.

The expected payoff at equilibrium is $\frac{V}{2} - \frac{V^2}{4C}$.

The highest expected payoff $\frac{V}{2}$ occurs when the population is monomorphic at Dove.

In particular, evolutionary games are NOT models of group selection.

Doubly Symmetric Games (Weibull, 1995)
Partnership Games (Hofbauer & Sigmund, 1998)
Potential Games (Sandholm, 2011)

Theorem. Suppose that the $m \times m$ payoff matrix A is symmetric. Then $p^* \in \Delta^m$ is an ESS iff it is a strict local maximum of the average payoff function $p \cdot Ap$. Every such matrix game has at least one ESSet; namely,

$$\{p \in \Delta^m \mid p \in \arg\max\{q \cdot Aq \mid q \in \Delta^m\}\}.$$

For almost all symmetric payoff matrices, every ESSet consists of a finite number of ESS's.

For these games, almost all trajectories of the replicator equation converge to an ESS. That is, for almost all (i.e. up to a set of measure zero) points in Δ^m, the trajectory with this initial point converges to some ESS.

Proof. This is a corollary of the Fundamental Theorem of Natural Selection (at a single locus with multiple alleles) since these games are equivalent to viability selection in these circumstances where A_{ij} is the survival probability of genotype A_iA_j.

Example of an ESSet: Let $A \equiv \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 0 & 1 & 1 \end{bmatrix}$.

Then $E = \{p \in \Delta^m \mid p_1 = 0 \text{ or } p_3 = 0\}$ since A is symmetric and $p \cdot Ap \leq 1$ with equality iff $p \in E$.

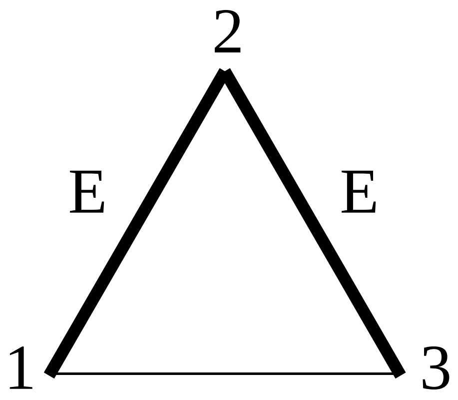

Folk Theorem of Evolutionary Games

For deterministic evolutionary dynamics

(a) A stable rest point is a NE.

(b) Any convergent trajectory evolves to a NE.

(c) A strict NE is asymptotically stable.

The theorem is true for the replicator equation applied to frequency-dependent population games with the understanding that part (b) applies for trajectories starting with all strategies present.

Unfortunately, many evolutionary games (even matrix games) have no stable rest points, strict NE, ESS or convergent trajectories, especially when there are a large number of strategies. This issue already appears for three-strategy games (e.g. Rock-Scissors-Paper games with cyclic dominance).

However, there are still many situations where the folk theorem applies as illustrated briefly by the following examples that I am particularly interested in.

Unfortunately, many evolutionary games (even matrix games) have no stable rest points, strict NE, ESS or convergent trajectories, especially when there are a large number of strategies. This issue already appears for three-strategy games (e.g. Rock-Scissors-Paper games with cyclic dominance).

However, there are still many situations where the folk theorem applies as illustrated briefly by the following examples that I am particularly interested in.

The n-stage war of attrition

The two players now have a finite sequence of interactions where their strategy choice at later interactions depend on earlier choices.
Here, players choose between Leave (L) and Remain (R) at each stage until some player chooses L or stage n is reached.

This is a symmetric simultaneity game whose extensive form has a large number of subgames.

The Three-Stage War of Attrition

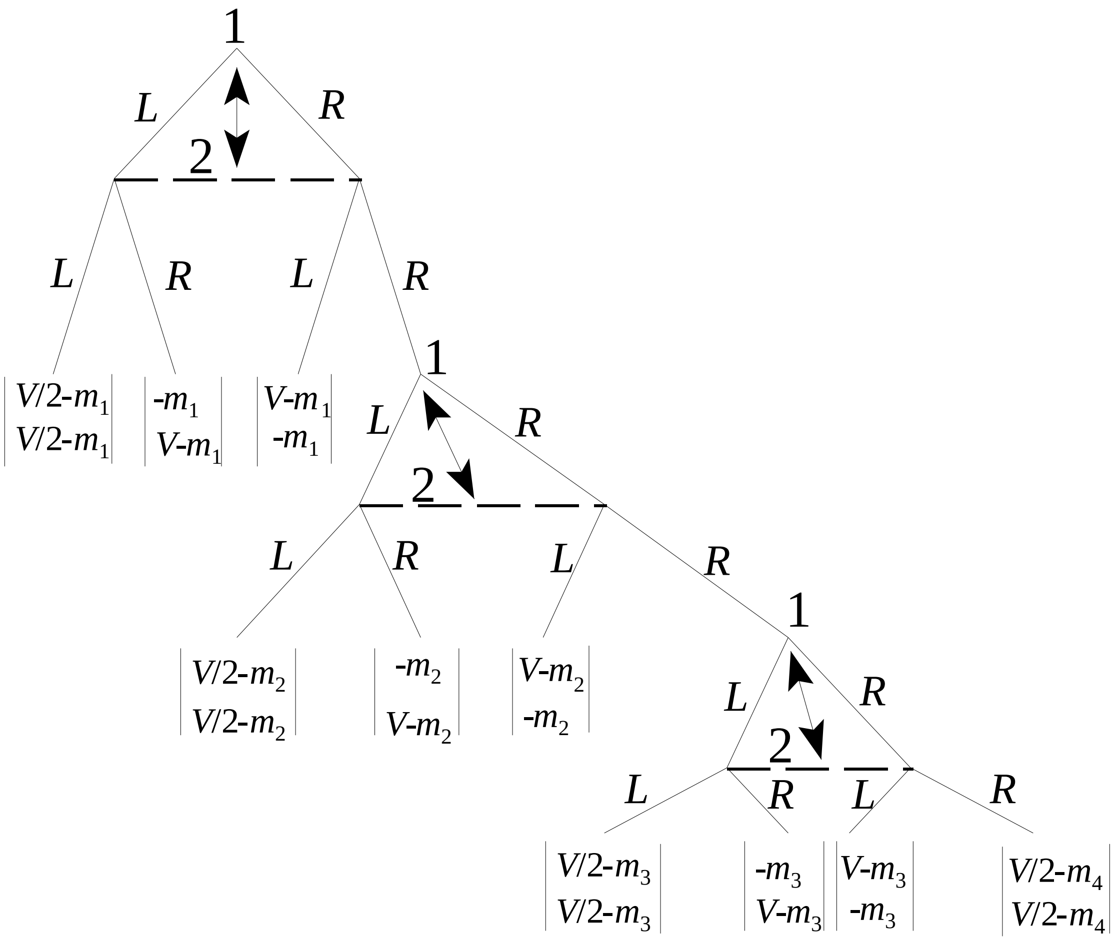

Theorem (Cressman, 2003):

Consider the reduced-strategy normal form of the N-Stage War of Attrition where the cost of waiting to stage i, m_i, is an increasing function of i. The game has $N+1$ pure strategies and a unique ESS $p^* \in \Delta^{(N+1)}$.

This ESS is pervasive (i.e. it reaches every player information set) and is found by backward induction. It is globally asymptotically stable for the N-dimensional replicator equation. The expected payoff at the ESS is less than that obtained by both players leaving immediately (i.e. less than $\frac{V}{2} - m_1$).

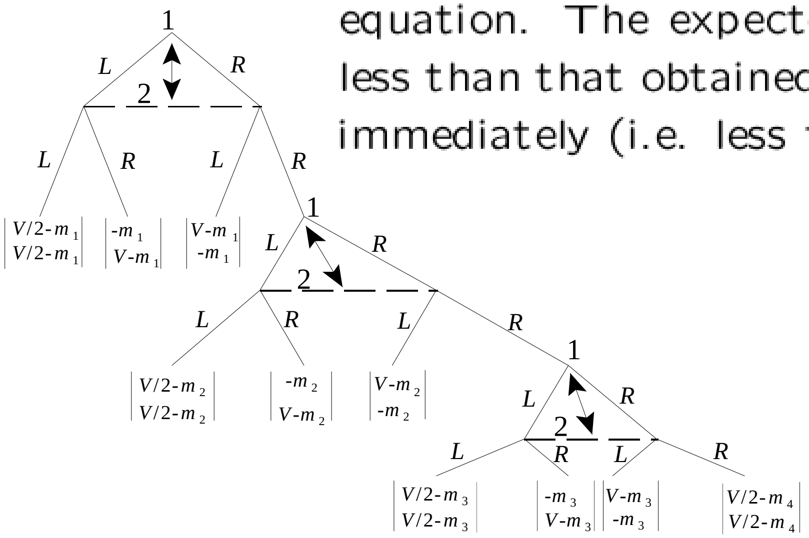

3. Theorem (Cressman, 2003)
Suppose Γ is a symmetric simultaneity game.

(a) If p^* is an asymptotically stable NE of the standard normal form of Γ under the replicator equation, then p^* is
(i) pervasive (i.e. p^* against p^* reaches all subgames) and
(ii) p^* is subgame perfect (i.e. p^* induces a NE in every subgame).

(b) If Γ has no moves by nature, then a pervasive NE p^* of the reduced-strategy normal form of Γ is asymptotically stable if and only if p^* is given by backward induction applied to the asymptotically stable pervasive NE of the subgames of Γ and their truncations.

Note: In part (b), we cannot replace "asymptotically stable" with "ESS". A well-known two-stage counterexample is given by van Damme (1991) (see also Selten, 1988).

A Symmetric Signaling Game

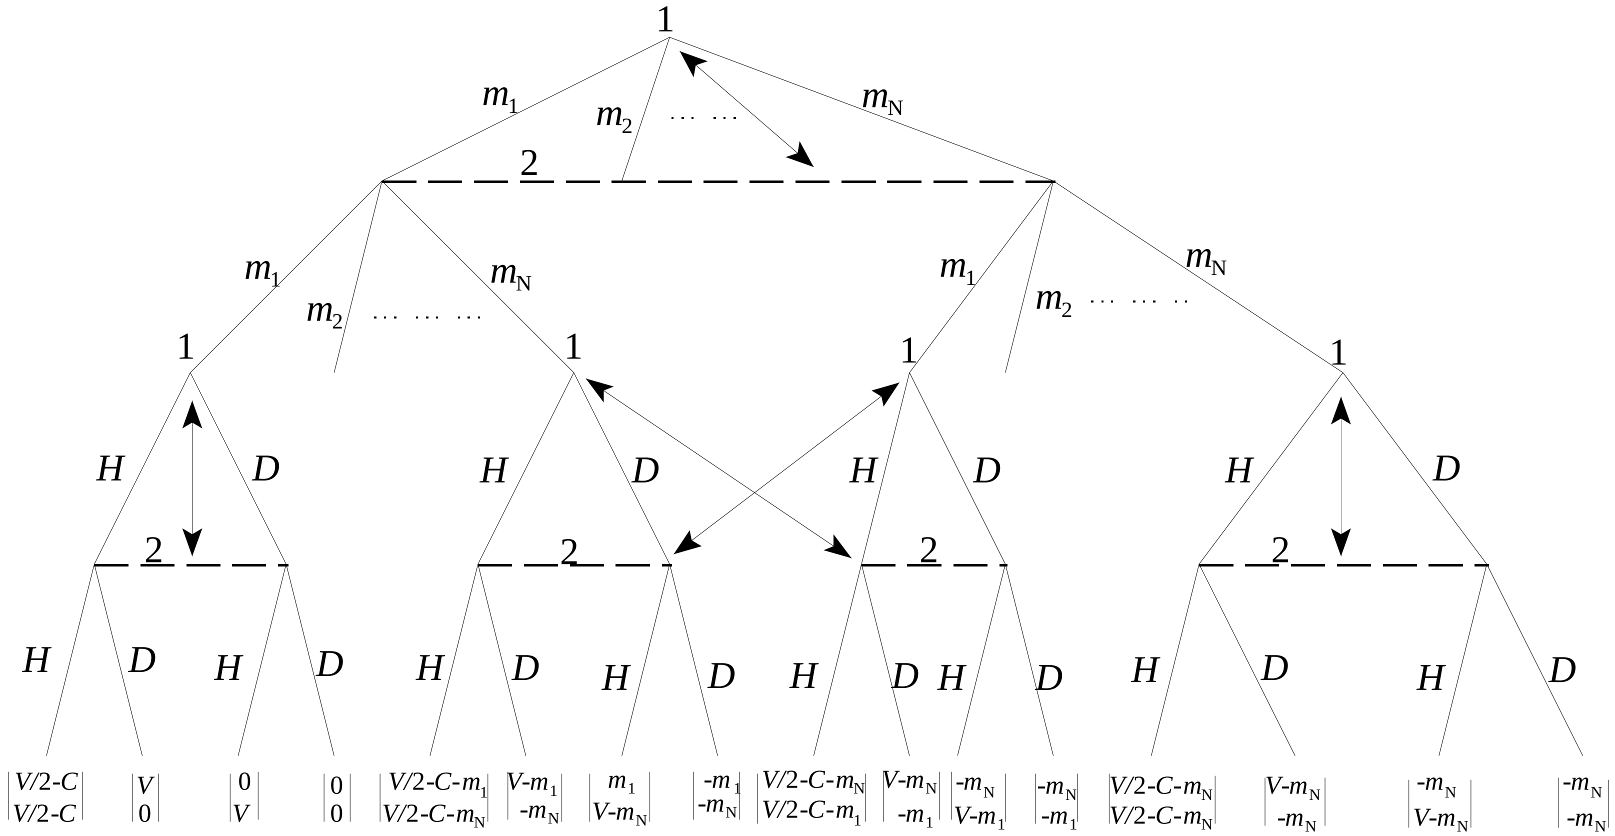

Single-Species Habitat Selection Game

1. Individuals choose which of H habitats (patches) to inhabit. Pure strategies: $e_1, e_2, ..., e_H$.
2. Payoff F_i in patch i depends only on the population size in i.
3. Assume F_i decreases as size increases.

When total population size is fixed at N, the payoff in i only depends on frequency p_i in patch i. Let

$$\pi(e_i, p) \equiv F_i(Np_i).$$

This is an example of a frequency-dependent population game.

As an example, the logistic effect is

$$F_i(Np_i) \equiv r_i \left(1 - \frac{Np_i}{K_i}\right).$$

Ideal Free Distribution (Fretwell & Lucas, 1969):
Species will distribute itself so that
(i) all occupied patches have the same payoff
(ii) the payoff in any unoccupied patch would be no greater than in any occupied patch.

In game-theoretic terms, the IFD is a NE of the habitat selection game.

Theorem (Fretwell & Lucas, 1969)
There exists a unique IFD p^* for each fixed population size N.

Theorem (Cressman & Krivan, 2006)
The IFD p^* is an ESS and it is globally asymptotically stable (under the replicator equation).

Theorem (Cressman & Krivan, 2006)
The IFD p^* is an ESS and it is globally asymptotically stable (under the replicator equation).

Proof. It is enough to show that $\pi(p^*, p) > \pi(p, p)$ for all $p \neq p^*$.

$$\begin{aligned}\pi(p^* - p, p) &= \sum_{i=1}^{H}(p_i^* - p_i)F_i(Np_i) \\ &\geq \sum_{i=1}^{H}(p_i^* - p_i)(F_i(Np_i) - F_i(Np_i^*)) \text{ (1)} \\ &> 0 \text{ (2)}\end{aligned}$$

(1) $F_j(Np_j^*) \leq F_i(Np_i^*)$ if $p_j^* = 0$ and $p_i^* \neq 0$

(2) $(p_i^* - p_i)(F_i(Np_i) - F_i(Np_i^*)) > 0$ if $p_i \neq p_i^*$.

Adaptive Movement

n_i is the population size (density) in patch i.

$$\dot{n}_i = \sum_{j=1}^{H} I_{ij}(n)n_j - \sum_{j=1}^{H} I_{ji}(n)n_i$$

where $I_{ij}(n)$ is the migration rate from patch j to patch i.

Note: This is not the replicator equation if animals can migrate to an unoccupied patch when this patch has higher payoff.

Theorem (Cressman & Krivan, 2006)
Suppose (i) animals never migrate to a patch with lower payoff and
(ii) some animals always migrate to a patch with the highest payoff.
Then the above migration dynamics evolves to the IFD corresponding to the total population size $N = \sum_{j=1}^{H} n_j$. That is, n_i evolves to Np_i^*.

Dynamics satisfying properties (i) and (ii) are called "better response" dynamics and include the best response dynamics. For general frequency-dependent population games, an ESS may not be stable under all better response dynamics.

Frequency and density dependent habitat selection game:

$$\dot{n}_i = n_i F_i(Np_i) + r\left[\sum_{j=1}^{H} I_{ij}(n)n_j - \sum_{j=1}^{H} I_{ji}(n)n_i\right]$$

where r is a positive parameter giving the time scale between population dynamics and migration (i.e. frequency) dynamics.

In this polymorphic model with a fixed time scale r, the density in each patch i evolves to the patch carrying capacity K_i.
This corresponds to the IFD, $p_i^* = K_i/N^*$, at the population equilibrium $N^* = K_i + ...K_H$.

Separation of time scales

Fast migration dynamics:
For fixed N, the frequency distribution p converges to the IFD $p^*(N)$. At this IFD, all individuals have the same fitness; namely,

$$\overline{F}(N) \equiv \sum p_i^*(N) F_i(p_i^*(N)N).$$

Note that $\overline{F}(N) = F_i(p_i^*(N)N)$ for each occupied patch (i.e. for $p_i^*(N) > 0$).

The slow population dynamics is then

$$\dot{N} = \sum \dot{n}_i = N\overline{F}(N)$$

as p instantaneously tracks the IFD.

Question: Does N evolve to the population equilibrium $N^* = K_i + ...K_H$?

Fast population dynamics:
For fixed p, the density of the monomorphic population evolves according to

$$\dot{N} = \sum \dot{n}_i = N\overline{F}(N) \quad with \quad \overline{F}(N) \equiv \sum p_i F_i(p_i N).$$

Suppose that N converges to a stable equilibrium $N^*(p)$ for each fixed p.

The slow frequency dynamics is then

$$\dot{p}_i = \sum_{j=1}^{H} I_{ij}(N^*(p)p)p_j - \sum_{j=1}^{H} I_{ji}(N^*(p)p)p_i$$

as N instantaneously tracks the $N^*(p)$.

Question: Does this system evolve to carrying capacity in each patch?

Theorem

For any better response evolutionary strategy dynamics and any choice of time scale r (including a complete separation of time scales), the frequency and density habitat selection game for a single species evolves to the IFD at the population density equilibrium.

That is, the system converges to the carrying capacity in each patch.

Theorem

For any better response evolutionary strategy dynamics and any choice of time scale r (including a complete separation of time scales), the frequency and density habitat selection game for a single species evolves to the IFD at the population density equilibrium.

That is, the system converges to the carrying capacity in each patch.

This result is not true for two (or more) species habitat selection games (e.g. for predator-prey systems or for competitive species) with adaptive movement.

For example, there are two-habitat predator-prey systems that reach a stable equilibrium under fast migration dynamics but become unstable when the time scales are reversed.

For example, there are two-habitat predator-prey systems that reach a stable equilibrium under fast migration dynamics but become unstable when the time scales are reversed.

That is, time scales are important.

The above discussion raises a basic question of how population dynamics interact with behavioral dynamics. In particular, what is the effect of different time scales for these processes?

Time scales also play an important role for games with continuous strategy spaces.
For example, the assumptions underlying the canonical equation of adaptive dynamics typically include population dynamics that operate on a fast time scale compared to evolutionary dynamics.

Adaptive dynamics for a one-dimensional trait space S.

S is a subinterval of real numbers.
$x \in S$ are the pure strategies.

$\pi(y, x; N)$ is the payoff to an individual using pure strategy y if the resident population is monomorphic with trait x and density N.

Assume (i) The resident population dynamics (i.e. $\dot{N} = N\pi(x, x; N)$) is on a fast time scale, (ii) For fixed x, population dynamics evolves to a stable positive equilibrium $N^*(x)$.

Define $\pi(y, x) \equiv \pi(y, x; N^*(x))$. In particular, $\pi(x, x) = 0$.

The canonical equation of adaptive dynamics models the evolution of this monomorphism x through trait substitution by successful invasion by nearby mutant traits y.

Specifically, at x in the interior of S,

$$\dot{x} = k(x)\frac{\partial \pi(y,x)}{\partial y}\,|_{y=x} \tag{1}$$

where $k(x)$ is a positive function related to the rate of mutation.

Theorem. Suppose that x^* is a rest point of (1) in the interior of S (i.e. $\pi_1(x^*, x^*) = 0$).

(a) x^* is asymptotically stable (also called convergence stable) if $\pi_{11} + \pi_{12} < 0$.
If x^* is convergence stable, then $\pi_{11} + \pi_{12} \leq 0$.

Here the second order partial derivatives are all evaluated at (x^*, x^*).

Theorem. Suppose that x^* is a rest point of (1) in the interior of S.
(a) x^* is convergence stable if $\pi_{11} + \pi_{12} < 0$.
If x^* is convergence stable, then $\pi_{11} + \pi_{12} \leq 0$.
(b) If $\pi_{11} < 0$, then x^* is a neighborhood strict NE. Conversely, if x^* is a NE, then $\pi_{11} \leq 0$.

Notes: 1. A strict NE x^* need not be convergence stable (c.f. Folk Theorem).

2. A convergence stable x^* need not be a (strict) NE.
In this case, a $y \in S$ can invade the dimorphism evenly split between $x^* + \varepsilon$ and $x^* - \varepsilon$ if and only if $\mid y - x^* \mid > \varepsilon$.
This forms the basis of an initial evolutionary branching.

3. A continuously stable strategy (CSS) x^* is a strict NE that is convergence stable.

Questions: 1 What happens if the population density time scale is not separated from the evolutionary time scale? One approach is the Darwinian dynamics of Vincent and Brown (2005).

2. What happens if the resident population is polymorphic. One approach is the replicator equation generalized to an infinite-dimensional dynamics on the set of probability measures over S (Cressman and Hofbauer, 2005).

Multi-dimensional trait space

S is a subinterval of $\mathbf{R}^n$ with x^* in its interior. The canonical equation of adaptive dynamics is now

$$\frac{dx}{dt} = k(x)C(x)\nabla_1 \pi(y,x)\mid_{y=x} . \qquad (2)$$

$\nabla_1 \pi(y,x)\mid_{y=x}$ is the gradient vector with i^{th} component $\frac{\partial \pi(y,x)}{\partial y_i}\mid_{y=x}$ and
$C(x)$ is an $n \times n$ covariance matrix
(i.e. positive definite and symmetric) modeling the mutation process in different directions.

Theorem. (Leimar, 2009) Suppose x^* in the interior of S is a rest point of (2).

(a) x^* is asymptotically stable for any choice of continuous covariance matrix $C(x)$ (i.e. x^* is (strongly) convergence stable) if $A + B$ is negative definite where

$$A_{ij} \equiv \frac{\partial^2 \pi(y,x^*)}{\partial y_j \partial y_i}\mid_{y=x^*} ; B_{ij} \equiv \frac{\partial^2 \pi(y,x)}{\partial y_i \partial x_j}\mid_{x=y=x^*} .$$

Theorem. (Leimar, 2009) Suppose x^* in the interior of S is a rest point of (2).

(a) x^* is convergence stable if $A+B$ is negative definite.

Conversely, if x^* is convergence stable, then $A+B$ is negative semi-definite.

(b) If A is negative definite, then x^* is a neighborhood strict NE. Conversely, if x^* is a NE, then A is negative semi-definite.

(c) If A and $A+B$ are negative definite, then x^* is a neighborhood CSS and neighborhood half-superior.

Conversely, if x^* is a neighborhood CSS or neighborhood half-superior, then A and $A+B$ are negative semi-definite.

Definition. (Cressman, 2009) Suppose the strategy space S of a symmetric game is a subset of $\mathbf{R}^n$ and $0 \leq p^* < 1$ is fixed. Strategy $x^* \in S$ is *neighborhood p^*-superior* if

$$\pi(x^*, P) > \pi(P, P)$$

for all probability measures $P \in \Delta(S)$ with $1 > P(\{x^*\}) \geq p^*$ and the support of P sufficiently close to x^*.

It is *neighborhood superior* if $p^* = 0$.
It is *neighborhood half-superior* if $p^* = \frac{1}{2}$.

Notes. 1. Half-superiority is the same as CSS by the above theorem (if borderline cases are excluded).
2. Neighborhood superiority is the same as convergence of the replicator equation.
It is also related to the neighborhood invader strategy (NIS) of Apaloo (1997).
3. It is the superiority concept that generalizes most easily to asymmetric games (such as two-species games) with continuous strategy spaces (Cressman, 2010).

Monomorphism versus Polymorphism

Time Scales: Behavioral; Ecological;
Evolutionary. Predator versus prey.

Two Species and Multi Species

Non Equilibrium Behavior

Invadable and Uninvadable
Replacement

Monomorphism versus Polymorphism

Time Scales: Behavioral; Ecological;
Evolutionary. Predator versus prey.

Two Species and Multi Species

Non Equilibrium Behavior

Invadable and Uninvadable
Replacement

Finite and Stochastic Effects

Spatial Effects

Evolution of Cooperation

Kin Selection

References

John Maynard Smith, 1982. Evolution and the Theory of Games, Cambridge University Press.

Jorgen Weibull, 1995. Evolutionary Game Theory, MIT Press

Josef Hofbauer and Karl Sigmund, 1998. Evolutionary Games and Population Dynamics, Cambridge University Press.

Ross Cressman, 2003. Evolutionary Dynamics and Extensive Form Games, MIT Press.

Tom Vincent and Joel Brown, 2005. Evolutionary Game Theory, Natural Selection, and Darwinian Dynamics, Cambridge University Press.

Fabio Dercole and Sergio Rinaldi, 2008. Analysis of Evolutionary Processes: The Adaptive Dynamics Approach and Its Applications, Princeton University Press

Merci