

Evolution in Finite Populations

Schaffer *J. Th. Bio.* 1988

Proposes a definition of ESS in finite populations that is intended to capture the idea of “resisting invasion by mutants.”

Recall the definition of ESS:

σ is an ESS if, for all $\sigma' \neq \sigma$, either

(a) $u(\sigma, \sigma) > u(\sigma', \sigma)$ or

(b) $u(\sigma, \sigma) = u(\sigma', \sigma)$ and $u(\sigma, \sigma') > u(\sigma', \sigma')$.

This is equivalent to requiring that

(a') σ is a NE: there is no other strategy σ' with $u(\sigma', \sigma) > u(\sigma, \sigma)$

and

(b') If $u(\sigma', \sigma) = u(\sigma, \sigma)$
(so σ' is an *alternate best reply* to σ)
then $u(\sigma', \sigma') < u(\sigma, \sigma')$

Schaffer takes this form as his starting point, and modifies (a') and (b') to reflect the finite population.

Specifically, he replaces (a') with the condition that when a single mutant is playing σ' and $N-1$ others are playing σ , the payoff to σ is higher than that to σ' . That is,

$$u(\sigma', \sigma) \leq (1 - 1/(N - 1))u(\sigma, \sigma) + (1/(N - 1))u(\sigma, \sigma')$$

(this is the version for 2-player games, the paper considers the general case of C players).

He calls this an “equilibrium condition;” it isn't a Nash equilibrium condition, but it does look like a condition for the share of σ to grow under a replicator-like dynamic.

It is also very close to the condition used to motivate ESS: Given a small invasion, the payoff of the ESS should exceed the payoff of the mutant, that is

$$(1 - \varepsilon)u(\sigma, \sigma) + \varepsilon u(\sigma, \sigma') > (1 - \varepsilon)u(\sigma', \sigma) + \varepsilon u(\sigma', \sigma')$$

the differences are the weak inequality and that a single mutant never has to play itself, so the payoff against itself doesn't matter.

It's clear that the ESS on this definition need not be Nash;

Strongest difference is for $N = 2$, where the finite ESS condition is

$$u(\sigma', \sigma) \leq u(\sigma, \sigma').$$

Only the cross-terms matter!

	<i>A</i>	<i>B</i>
<i>A</i>	3,3	1,2
<i>B</i>	2,1	0,0

Here *A* is a dominant strategy but it is not an ESS for $N=2$.

Can't say what -will- happen w/o an explicit dynamic.

Fudenberg, Nowak and Taylor [2004]:
Stochastic Evolution in Finite Populations

We consider a population of N agents playing a 2×2 symmetric game with payoff matrix

$$\begin{array}{cc} & A & B \\ \begin{array}{c} A \\ B \end{array} & \left(\begin{array}{cc} a, a & b, c \\ c, b & d, d \end{array} \right) \end{array}$$

Let i denote the number of individuals using strategy A .

The fitness of individuals using strategy A is

$$f_i = \frac{a(i-1) + b(N-i)}{N-1}$$

The fitness of individuals using strategy B is

$$g_i = \frac{ci + d(N-i-1)}{N-1}.$$

(These payoffs are deterministic- everyone plays everyone else.)

Because of the finite population, the two types of agent face a slightly different distribution of opponents' play.

“Frequency-dependent Moran process”

(standard Moran process is a model of genetic drift)

Each period, each agent “reproduces” at a rate proportional to its fitness, so that the number of A -offspring is if_i and the number of B -offspring is $(N - i)g_i$.

One offspring is chosen at random to enter the population, so that the probability of adding an A offspring is $\frac{if_i}{if_i + (N - i)g_i}$.

However, there is a probability μ_1 that an A -offspring is a “mutant” that plays B instead of A , and a probability $\mu_2 = k\mu_1$ that a B -offspring plays A instead of B .

After reproduction, one randomly chosen agent is removed from the population, so that the aggregate population size is constant; each old agent has probability $1/N$ of being removed.

Because the population size is constant, this process can be viewed as a model of imitation:

Each period one agent at random is selected to update, and the choice of a new strategy is influenced both by the prevailing payoffs and the prevailing popularity of the choices.

This is a “birth-death” process: the states are integers, and in each period the state can move by at most one step.

Let $P_{i,j}$ denote the probability of a transition from state i to state j . Then $P_{i,j} = 0$ if $|i - j| > 1$, and $P_{0,1} = \mu_2 = 1 - P_{0,0}$;

$$P_{i,i+1} = \Pr[\text{Add an A and remove a B}]$$

$$= \frac{if_i(1 - \mu_1) + (N - i)g_i\mu_2}{if_i + (N - i)g_i} \frac{N - i}{N}$$

$$P_{i,i-1} = \text{Probability} [\text{add a B and remove an A}]$$

$$P_{i,i} = 1 - P_{i,i+1} - P_{i,i-1} \quad \text{for } i = 1 \text{ to } N-1; \text{ and}$$

$$P_{N,N-1} = \mu_1 = 1 - P_{N,N}.$$

Let \hat{P} be the version of the process where the mutation rates are identically zero. In this process, states 0 and N are absorbing, and the others are transient.

We expect that the invariant distribution for small μ will be concentrated on these two endpoints.

Intuitively, after each mutation, the process will reach an absorbing state of the no-mutation process and remain there for a long time before the next mutation occurs. Moreover, we expect that the relative probabilities of these states will be determined by the probabilities of moving from one to the other when mutations are rare, so we need to solve the no-mutation process first.

$$\text{Let } x^*(k, N) \equiv \lim_{\mu_1 \rightarrow 0} x(\mu_1, k\mu_1, N).$$

Lemma 1: Under \hat{P} , the probability of absorption at state N, starting from state 1, is

$$\phi_{1,N} = \frac{1}{1 + \sum_{i=1}^{N-1} \prod_{j=1}^i \frac{g_j}{f_j}}, \text{ and the probability}$$

of absorption at 0, starting at N-1, is

$$\phi_{N-1,0} = 1 - \phi_{N-1,N} = \frac{\prod_{j=1}^{N-1} \frac{g_j}{f_j}}{1 + \sum_{i=1}^{N-1} \prod_{j=1}^i \frac{g_j}{f_j}}.$$

So the relative probabilities are determined by the product of the ratios of payoffs to the two strategies.

Reasons for this simple form:

1) The general formula involves the product of the ratios of the one-step transition probabilities.

2) Here

$$\frac{P_{i,i-1}}{P_{i,i+1}} = \frac{\left((N-j)g_i / ((N-j)g_i + jf_i) \right) (j/N)}{\left(jf_i / ((N-j)g_i + jf_i) \right) ((N-j)/N)} = \frac{g_i}{f_i}$$

Lemma 2: The weight that the limit distribution $x^*(k, N)$ assigns to state N is

$$\frac{k\phi_{1,N}}{k\phi_{1,N} + \phi_{N-1,0}} = \frac{k\prod_{j=1}^{N-1} f_j}{k\prod_{j=1}^{N-1} f_j + \prod_{j=1}^{N-1} g_j}.$$

Intuition: Consider the induced Markov chain on the two states 0 and N. (That is, ignore all periods where the state has some other value- such periods will be vanishingly rare as $\theta \rightarrow \infty$.)

The probability of a transition from state 0 to state N is $x_0\mu_2\phi_{1,N}$:

the probability that of state 0, times the probability of a mutation to A, followed by the probability that the no-mutation process is absorbed at all A starting from only 1 A.

The probability of a transition from state N to state 0 is $x_N\mu_1\phi_{N-1,0}$.

The displayed formula comes from equating these terms.

The paper gives a direct proof of lemma 2, but the induced-chain argument can be made rigorous. It holds in higher dimension (i.e for games with more strategies) and for more general adjustment processes.

The key assumptions are that

(1) the “vertices” (states where everyone plays the same action”) are absorbing in the non-mutation process,

(2) the probability of two or more simultaneous mutations is of lower order than the probability of one mutation, and

(3) that the “induced Markov chain” on the vertices (with transition probabilities derived from considering a single mutation and then following the no-mutation process) has a unique invariant distribution.

(Fudenberg and Imhof [2004], in preparation.)

Lemma 3: The long-run probability of state N , is more than $1/2$ exactly when $\phi_{N-1,0} < k\phi_{1,N}$, or equivalently when

$$\phi_{1,N} / \phi_{N-1,0} = \prod_{j=1}^{N-1} \frac{f_j}{g_j} \equiv \gamma > 1/k. \text{ Moreover,}$$

$x_N^*(k, N)$ converges to 1 or 0 as γ converges to ∞ or 0 respectively.

In papers such as Kandori, Mailath and Rob [1993] where the ratio of the mutation probabilities has no effect on the limit distribution so long as it is bounded away from 0 and infinity. This is because the radius (the number of mutations required to leave the basin of an absorbing state) of one of the absorbing states is smaller than the radius of the other one, so that regardless of N the limit distribution assigns positive probability to a single point. Here, the radius of each absorbing state is 1.

We use Lemma 3 to study how the limit distribution depends on the size of the population and on the parameters of the payoff matrix.

Plug in payoff functions into formula for γ :

$$\gamma = \frac{(b + (N - 2)a) \dots ((N - 2)b + a) ((N - 1)b)}{(c + (N - 2)d) \dots ((N - 2)c + d) \dots ((N - 1)c)}.$$

Note that when $N = 2$, $\gamma = b/c$:

As in the Schaffer's finite-population ESS, only the off-diagonal terms matter when $N = 2$.

Note that multiplying all of the f_i and g_i by a constant has no effect on γ , and so has no effect on the limit of the ergodic distributions.

But γ may change when a constant is added to all of the fitness functions.

Main result is a classification of the outcomes w.r.t. the payoff matrix- most of which come from a simple comparison of terms.

Theorem:

(a) If $b > c, a > d$, then $x_N^*(1, N) > 1/2$ for all N ,
and for any k ,

$$\lim_{N \rightarrow \infty} x_N^*(k, N) = \lim_{N \rightarrow \infty} \lim_{\theta \rightarrow \infty} x_N(k, \theta, N) = 1$$

(b) If $b > c, a < d$, then whether $x_N^*(k, N) > 1/2$
depends on k and the population size.

A sufficient condition for $x_N^*(k, N) > 1/2$ is
 $b - c > (N - 2)(d - a)$.

Moreover,

(b.1) In subcase $b > d > a > c$,
 $\lim_{N \rightarrow \infty} x_N^*(k, N) = 1$.

(b.2) In subcase $d > b > c > a$,
 $\lim_{N \rightarrow \infty} x_N^*(k, N) = 0$.

(b.3) In subcases $d > b > a > c$ and
 $d > a > b > c$, there are two pure-
strategy Nash equilibria, and
 $\lim_{N \rightarrow \infty} x_N^*(k, N)$ is 1 or 0

as $\int_0^1 \ln(b + (a - b)x)dx$ is greater or less
than $\int_0^1 \ln(c + (d - c)x)dx$.

The risk-dominant equilibrium need not
be selected.

Cases (c), (d) are symmetric.

Corollary: If the game has a strictly dominant strategy, the probability assigned this strategy by the limit distribution converges to 1 as N goes to infinity.

If the game has two strict Nash equilibria, then except for knife-edge cases there is an equilibrium to which the limit distribution assigns probability converging to 1 as N goes to infinity, but the risk-dominant equilibrium need not be selected.

Idea of Proof: Recall

$$\gamma = \frac{(b + (N - 2)a) \dots ((N - 2)b + a) ((N - 1)b)}{(c + (N - 2)d) \dots ((N - 2)c + d) \dots ((N - 1)c)}.$$

In case (a), the first term in the numerator exceeds the corresponding term in the denominator, as does the second, etc, so $\gamma > 1$, and lemma 3 implies that $x_N^*(1, N) > 1/2$. Moreover, the ratio of each pair of terms in γ is bounded away from 1, so it becomes infinite as N grows large

In case (b), if $b + (N - 2)a > c + (N - 2)d$, the pairwise comparison of terms again shows that $\gamma > 1$. The results for large N in subcases (b.1), and (b.2), also come from a comparison of corresponding terms.

The argument for large N in subcase (b.3) involves approximating γ by a ratio of integrals, using the idea that

$$\prod_{j=0}^{N-1} \left[a \frac{j}{N-1} + b \left(1 - \frac{j}{N-1} \right) \right] =$$

$$\exp \left(\sum_j \ln \left[a \frac{j}{N-1} + b \left(1 - \frac{j}{N-1} \right) \right] \right)$$

is approximately

$$\exp \left(N \int_0^1 \ln(a + (b-a)x) dx \right).$$

Because the selection is driven by the ratio of this expression to $\exp \left(N \int_0^1 \ln(c + (d-c)x) dx \right)$, the risk dominant equilibrium need not be selected.

To see why, consider a game where $d > b > a > c$, and $a + b = c + d \equiv 2\mu$ so that neither strategy is risk dominant. Then the two integrals are the expectations of the logarithm of two random variables with the same mean. The log is a concave function, so the expected value of the log is reduced by a mean-preserving spread, and (A, A) is selected because $b - a < d - c$.

Intuitively, when neither strategy is risk dominant they are equally fit at the point $i = N/2$, but the support of the long-run distribution depends on the transition probabilities at every state, and these are not determined by the value of the fitness functions at the midpoint.

Our model makes a somewhat odd prediction in the subcases of (a) and case (b) that correspond to “hawk-dove games,” meaning they are symmetric games with two asymmetric pure-strategy equilibria. Since we are working in a one-population model, the asymmetric equilibria in the hawk-dove case cannot arise. Moreover, in the limit of vanishingly small mutation rates, the process spends almost all of its time in the two homomorphous states.

But calculations show that when the mutation rate is high enough the process spends most of its time near the mixed-strategy equilibrium.

Fudenberg and Hojman [2004] show that the order of limits on N and ε doesn't matter provided that the only recurrent classes of the $\varepsilon = 0$ process are attractors for the deterministic continuous time dynamics obtained by sending N to infinity.

This is true for the case of two symmetric pure-strategy equilibria, but not for the hawk-dove case.

In hawk-dove, the interior mixed equilibrium is selected provided that $\varepsilon_N \geq e^{-\beta N}$ for some constant β .

Partial Intuition: Benaim and Weibull show that in the limit $N \rightarrow \infty$ the process is concentrated on the attractors of the deterministic process. When these contain the absorbing states of the unperturbed process we “expect” the two limits to yield distributions with the same support. In the one-dimensional case, the formulas for the relative probabilities involve the same ratio of integrals. (We don't think this extends to higher dimension.)

Possible extensions:

MxM coordination games, as in Kandori and Rob.

Multiplayer games, including “playing the field.” As in models of technology choice.

Multiple populations.

Other evolutionary processes such as “Wright-Fischer.”

Main question: which results hold for a wide range of evolutionary processes?