

Diversity in Evolutionary Dynamics

Leonard J. Schulman

Caltech

January 2025, ITCS

Joint work with
Yuval Rabani (Hebrew University)
Alistair Sinclair (UC Berkeley)

Why is there diversity in nature?



Coco Island, Costa Rica. Image: travel.earth

Replicator dynamics

Mathematical framework for population genetics rests upon the notion of a “fitness function” $u(x)$ of genotypes $x \in \{0, 1\}^n$.

p_x = fraction of population with genotype x

$\hat{u} = \sum_x p_x u(x)$ average fitness

Replicator Dynamics:

$$\dot{p}_x = p_x(u(x) - \hat{u})$$

[Chastain et al., 2014]: this dynamic is the same as the *multiplicative weights update* algorithm, a strategy formulated for regret-minimization in the theory of repeated games [Hannan, 1957].

[Mehta et al., 2015] pointed out that this made a curious prediction for population genetics: the most fit genotype will eventually overtake all others. Species becomes a monoculture.

We do not see monocultures in nature, however.

Why not monocultures?

Why the discrepancy between reality and the theoretical prediction?

There are a number of reasonable explanations (not mutually exclusive).

- Mutations [Mehta et al., 2017]
- Speciation
- The mathematical assumptions are too far from reality
- “Eventually” is longer than geologic time

However, there is also an explanation requiring no apologies for, or extensions to, the model.

[Ehrlich and Raven, 1964]: “It is apparent that reciprocal selective responses have been greatly underrated as a factor in the origination of organic diversity.”

Motivated *Red Queen Hypothesis* [Van Valen, 1973]: co-evolution with parasites spurs continual genetic change.

Frequency-dependent fitness

Frequency-dependent fitness is a broader theme in population genetics, first expressed clearly in:

[Lewontin, 1958]: "It is becoming apparent that the adaptive values of genotypes in populations are functions, and often complex functions, of the gene frequency. This is not surprising since the environment of an organism includes the relative frequencies of other sorts of organisms in the population."

This led naturally to the work of Ehrlich and Raven and other work on co-evolution. However the work both in population genetics and in TCS, left unanswered the following questions:

- 1 Does competition between species indeed preclude monocultures, in replicator equation-driven dynamics?
- 2 How might the answer to this question depend upon game-theoretic conditions on the inter-species interaction?

Answers for a special case were given in [Piliouras and Schulman, 2018]; the general case is quite different and is addressed in the present paper.

Weak Selection regime

As in most of the literature, we assume:

- 1 Haploid species with binary alleles (each organism has genotype in $\{0, 1\}^n$).
- 2 Reproduction is sexual (\Rightarrow genes mix) and panmictic (mating uar)

The model has two important and non-obvious aspects.

(a) **Weak selection regime** Gene mixing (from sex) occurs more quickly than natural selection. So the distribution p_x over genotypes is always a *product distribution* $p_x = \prod_1^n p_{i,x_i}$ (where $x_i \in \{0, 1\}$ and $p_{i,0} + p_{i,1} = 1$). Replicator dynamics separate into one equation per gene.

(b) **Two species locked in a zero-sum game** (think predator-prey or host-parasite).

Species A has n genes and population distribution p_x for $x \in \{0, 1\}^n$.

Species B has m genes and population distribution q_y for $y \in \{0, 1\}^m$.

Side comment: something like our results likely survive in non-zero-sum games but this remains unexplored.

Game-theoretic context: Team Games

Given: two-player zero-sum game in normal form, $2^n \times 2^m$ payoff matrix u .

Form a two-team game: Player i of Team A controls bit x_i of the row label $x = (x_1, \dots, x_n)$. Player j of Team B controls bit y_j of the column label $y = (y_1, \dots, y_m)$.

Randomization at different players is independent. Player i chooses an action x_i according to an agreed-on distribution $(p_{i,0}, p_{i,1})$, independently of all other player choices. Similarly for team B.

Payoff to team A is $u(x, y)$, payoff to team B is $-u(x, y)$.

Duality Gap

Two-player zero-sum games have, famously, no duality gap (von Neumann). However, this no longer holds for two-team zero-sum games. Each team is weakened by its inability to coordinate randomness among players. (Second-mover advantage.)

The existence of such a gap was first observed by [von Stengel and Koller, 1997] and the gap was quantified by [Schulman and Vazirani, 2019].

A species is a team of genes; gene actions are alleles

Due to Weak Selection, each species' strategy, at any time, is a product distribution. Thus a species plays as a *team*. The available actions to a gene are its alleles. The same payoff is shared by all genes.

The replicator equation for species A separates into n single-coordinate replicator equations, one for each $p_{i,0}$. Similarly for species B .

Due to co-evolution, the fitness of an allele in species A depends upon the distribution q of species B . And vice versa.

Coupled replicator equations

$$\begin{aligned}\dot{p}_{i,b} &= p_{i,b}(\widehat{u}_{i,b}(p, q) - \widehat{u}(p, q)) \\ \dot{q}_{j,b} &= q_{j,b}(-\widehat{u}_{j,b}(p, q) + \widehat{u}(p, q))\end{aligned}\tag{1}$$

As a measure of diversity we'll be interested in the system entropy:

$$H(p) + H(q) = \left(\sum_i H(\{p_{i,0}, p_{i,1}\})\right) + \left(\sum_j H(\{q_{j,0}, q_{j,1}\})\right)$$

Previous work: binary phenotypes

Some inter-species interactions are governed chiefly by one feature. E.g., *flower depth* vs. *pollinator tongue length*.



Sword-billed hummingbird visiting a flower of *Datura sanguinea*

Image: Libor Vaicenbacher,

photographylife.com

Model: binary functions $f : \{0, 1\}^n \rightarrow \{0, 1\}$,
 $g : \{0, 1\}^m \rightarrow \{0, 1\}$, payoff in interaction (x, y)
can be expressed as $u(f(x), g(y))$.

[Piliouras and Schulman, 2018] This creates effective energy-conserving Hamiltonian dynamics in 2D; thanks to Poincaré-Bendixson theorem, system is periodic. \Rightarrow Theorem:

Entropic Property A

For any initial populations
 $(p(0), q(0)) \in (0, 1)^{n+m}$ (i.e., all genotypes are initially represented),

$$\liminf_{t \rightarrow +\infty} (H(p(t)) + H(q(t))) > 0 \quad (2)$$

General two-team zero-sum games $u(x, y)$

Unlike the binary-phenotype case, dynamics are *not* periodic. Our first result is negative: namely, we show that sexual reproduction does not guarantee Property A, i.e., the maintenance of diversity at all times:

Theorem 1

There exist replicator dynamics with no weak pure Nash equilibrium for which Property A fails. In fact, inequality (2) holds only on a set of initial conditions of measure zero.


Our concrete example used in the proof of Theorem 1 consists of two species, with just $n = m = 2$ genes, each species having only three phenotypes. Sharp contrast to the result for binary phenotypes.

Method: there is a potential function which is continually decreasing, and is 0 in corners.

Exemplifying game:

$$\begin{bmatrix} 1 & \frac{1}{2} & \frac{1}{2} & -1 \\ -\frac{1}{2} & 0 & 0 & -\frac{1}{2} \\ -\frac{1}{2} & 0 & 0 & -\frac{1}{2} \\ -1 & \frac{1}{2} & \frac{1}{2} & 1 \end{bmatrix}$$

Genome bit distributions (in each species) exhibit “spiraling-out.”
On each near-visit to a corner, system entropy increasingly small.



Our main result is a complementary positive statement, which says that in any non-degenerate game, diversity *is* maintained in the following weaker, “infinitely often” sense.

Entropic Property B

For any initial populations $(p(0), q(0)) \in (0, 1)^{n+m}$, during an infinite span of time the entropy is uniformly bounded away from 0. Formally:

$$\exists \varepsilon > 0 \text{ s.t. } \int_0^\infty \max\{0, H(p(t)) + H(q(t)) - \varepsilon\} dt = \infty.$$

We also use the following even weaker entropic property:

Entropic Property C

For any initial populations $(p(0), q(0)) \in (0, 1)^{n+m}$,

$$\int_0^\infty (H(p(t)) + H(q(t))) dt = \infty.$$

The following theorem summarizes the maintenance of diversity in general replicator dynamics, according to types of Nash equilibria the game has:

Theorem 2

The following results hold for general replicator dynamics:

- (i) If the game has no pure Nash equilibrium, then Property **B** holds.*
- (ii) If the assumption in (i) is weakened to assume only that the game has no strict pure Nash equilibrium, then Property **C** holds.*
- (iii) If the game has a strict pure Nash equilibrium, then Property **C** (and therefore also Property **B**) fails on a set of initial populations $(p(0), q(0)) \in (0, 1)^{n+m}$ of positive measure.*

General method for (i) and (ii): show that even though trajectory cannot get stuck in corners. Occasionally escapes (if only to a new corner) and this creates a high-entropy time interval.

In summary, our results refute the supposition that sexual reproduction ensures diversity at all times, but affirm a weaker assertion that extended periods of high diversity are necessarily a recurrent event.

Bibliography I

- E. Chastain, A. Livnat, C. Papadimitriou, and U. Vazirani. Algorithms, games, and evolution. *Proc. National Academy of Sciences*, 2014. doi: 10.1073/pnas.1406556111.
- P. R. Ehrlich and P. H. Raven. Butterflies and plants: A study in coevolution. *Evolution*, 18(4):586–608, Dec. 1964.
- J. Hannan. Approximation to Bayes risk in repeated play. *Contributions to the Theory of Games*, 3(2):97–139, 1957.
- R. C. Lewontin. A general method for investigating the equilibrium of gene frequency in a population. *Genetics*, 43(3):419–434, 1958. doi: 10.1093/genetics/43.3.419.
- R. Mehta, I. Panageas, and G. Piliouras. Natural selection as an inhibitor of genetic diversity: Multiplicative weights updates algorithm and a conjecture of haploid genetics [working paper abstract]. In *Proc. 7th ITCS*, page 73. ACM, 2015. doi: 10.1145/2688073.2688118.
- R. Mehta, I. Panageas, G. Piliouras, P. Tetali, and V. V. Vazirani. Mutation, sexual reproduction and survival in dynamic environments. In *Proc. 8th ITCS*, pages 16:1–16:29, 2017. doi: 10.4230/LIPIcs.ITCS.2017.16.
- G. Piliouras and L.J. Schulman. Learning dynamics and the co-evolution of competing sexual species. In *Proc. 9th ITCS*, pages 59:1–3, 2018. Full version available at <https://arxiv.org/abs/1711.06879>.
- L. J. Schulman and U. V. Vazirani. The duality gap for two-team zero-sum games. *Games and Economic Behavior*, 115(C): 336–345, 2019. doi: 10.1016/j.geb.2019.03.011.
- L. Van Valen. A new evolutionary law. *Evolutionary Theory*, 1:1–30, 1973.
- B. von Stengel and D. Koller. Team-maxmin equilibria. *Games and Economic Behavior*, 21(1):309 – 321, 1997. doi: 10.1006/game.1997.0527.