The Evolution of Cooperation in an Iterated Survival Game

John Wakeley
joint work with Martin Nowak

Harvard University

CIRM Luminy
27 June 2018

Ideas behind this work

Cooperative behaviors might enhance survival.

Survival is a natural measure of payoff (or utility) in iterated evolutionary games.

Two-player games are like diploid selection models.

Image via justseeds.org, full text of Peter Kropotkin, Mutual Aid: A Factor of Evolution (1902) at Project Gutenberg at www.gutenberg.org/ebooks/4341

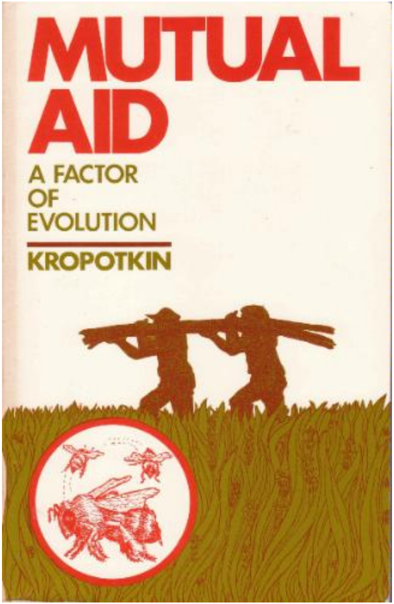

Mutual aid and mutual support (Kropotkin 1902)

Two aspects of animal life impressed me most during the journeys which I made in my youth to Eastern Siberia and Northern Manchuria. One of them was the extreme severity of the struggle for existence which most species of animals have to carry on against an inclement Nature. . . And the other was that even in those few spots where animal life teemed in abundance, I failed to find—although I was eagerly looking for it—that bitter struggle for the means of existence, among animals belonging to the same species, which was considered by most Darwinists (though not always by Darwin himself) as the dominant characteristic of struggle for life, and the main factor of evolution.

On the other hand, whenever I saw animal life in abundance. . . I saw Mutual Aid and Mutual Support carried on to an extent which made me suspect in it a feature of the greatest importance for the maintenance of life, the preservation of species, and its further evolution.

Preview of the iterated two-player survival game

The game:

- Two individuals are confronted with an unavoidable task, which we model as an iterated game.
- The game always includes n iterations.
- Surviving to the end of the game is the only payoff.
- If your partner dies, you must complete the task alone.

A specific question:

With selection via this game, under what circumstances will cooperative or helping behaviors evolve?

Background: Prisoner's Dilemma and iterated games

A represents cooperation, B defection.

Payoff matrix of Axelrod (1984), Rappoport and Chammah (1965):

		Partner A	Partner B
Individual	A	3	0
	B	5	1

Repeat (iterate) the game n times, and award cumulative payoffs:

Partner:	B	B	B	B	B	B	B	B	B	B
Individual:	B	B	B	B	B	B	B	B	B	B

Payoff: $1 + 1 + 1 + 1 + 1 + 1 + 1 + 1 + 1 + 1 = 10$

General symmetric two-player games (single iteration)

Two types of individuals: A and B (e.g. cooperator and defector)

Payoff matrix:

		Partner A	Partner B
Individual	A	a	b
	B	c	d

Individual payoffs (a, b, c and d) depend on

- the (focal) individual's type
- the type of its partner

Symmetry: the same matrix applies to both players.

More about symmetric two-player games, and cooperation

A is the more, and B the less cooperative type (assume $a > d$)

Payoff matrix:

		Partner A	Partner B
Individual	A	a	b
	B	c	d

Two key differences in payoff, depending on Partner

- $a - c$ is the advantage of A compared to B when Partner is A
- $b - d$ is the advantage of A compared to B when Partner is B

Three kinds of cooperative dilemmas, and no dilemma

	A	B
A	a	b
B	c	d

All possible two-player games fall into one of four categories, which will be exemplified here by

- Prisoner's Dilemma $(a - c < 0,\ b - d < 0)$
- Stag Hunt $(a - c > 0,\ b - d < 0)$
- Hawk-Dove game $(a - c < 0,\ b - d > 0)$
- Harmony Game $(a - c > 0,\ b - d > 0)$

The last term comes from De Jaegher and Hoyer (2016).

Heads up: later these will be referred to as PD, SH, HD, HG, in order of decreasing level of cooperative dilemma (N. 2012).

Evolutionary dynamics in infinite populations

	A	B
A	a	b
B	c	d

If x is the frequency of A, and pairs form at random in proportion to the frequencies of A and B, then in a continuous-time model

$$\frac{dx}{dt} = x(1-x)\Big((a-c)x + (b-d)(1-x)\Big)$$

and in a discrete-time model

$$\Delta x = x(1-x)\Big((a-c)x + (b-d)(1-x)\Big)/\bar{w}$$

in which $\bar{w}$ is the mean fitness. Note: no mutation.

- $(a-c)$ gives the direction of selection when A is common.
- $(b-d)$ gives the direction of selection when A is rare.

Three of four kinds of dynamics

	A	*B*
A	*a*	*b*
B	*c*	*d*

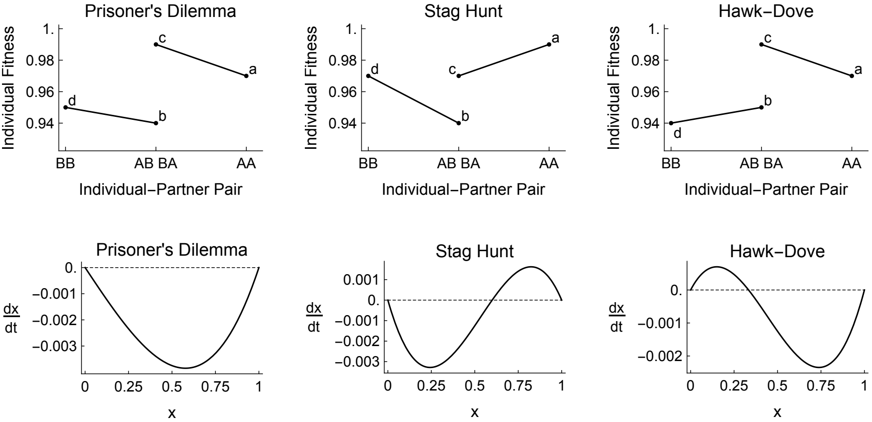

In diploid population genetics, these would be called: (negative) directional selection (PD), underdominance (SH) and overdominance (HD).

	A	B
A	$a(n)$	$b(n)$
B	$c(n)$	$d(n)$

How does repetition change the game?

- Now consider individual payoffs $a(n)$, $b(n)$, $c(n)$ and $d(n)$ for a game which is repeated n times.
- If payoffs accumulate additively as in the earlier example of the Prisoner's Dilemma, then

$$a(n) = na,\ b(n) = nb,\ c(n) = nc,\ d(n) = nd.$$

The signs of $a(n) - c(n)$ and $b(n) - d(n)$ will not change.

How does repetition change the game?

Example from literature: De Jaegher and Hoyer (2016), Model 2

- Two individuals protect a public good and their private goods.
- Individual's private good has a beneficial effect on partner.
- Public and private goods have different values for individuals.
- A number of attacks by predators occur.
- Individuals have $1/2$ chance of being the one who's attacked.
- Defectors can sustain 1 attack at a cost to the public good, but > 1 causes the public good to be lost completely.
- There is a fixed cost of being a cooperator.
- Cooperators are immune to attacks.

Under this ecological/behavioral model, the two key differences in payoff change non-linearly as the number of attacks increases.

Numerical example of Model 2 of De Jaegher and Hoyer (2016)

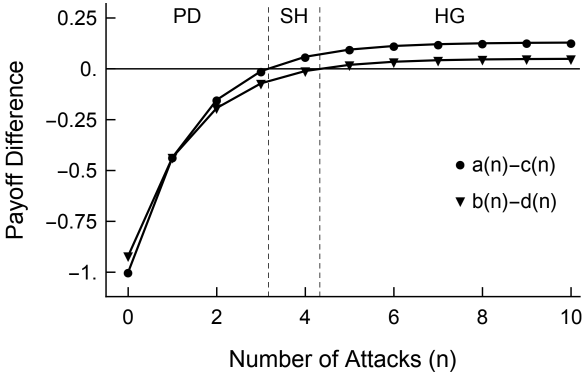

We studied this sort of phenomenon in our iterated survival game.

The iterated two-player survival game

	A	B	$\{\}$
A	a	b	a_0
B	c	d	d_0

Description, as before:

- Two individuals are confronted with an unavoidable task, which we model as an iterated game.
- The game always includes n iterations.
- Surviving to the end of the game is the only payoff.
- If your partner dies, you must complete the task alone.

For each iteration of the game:

- General 2×2 payoff matrix with *expected* payoffs a, b, c, d.
- These expected payoffs are probabilities of survival.
- Survival probabilities when alone: a_0 and d_0.
- A and B are fixed types, or pure, non-reactive strategies.

Flow diagram of the iterated two-player survival game

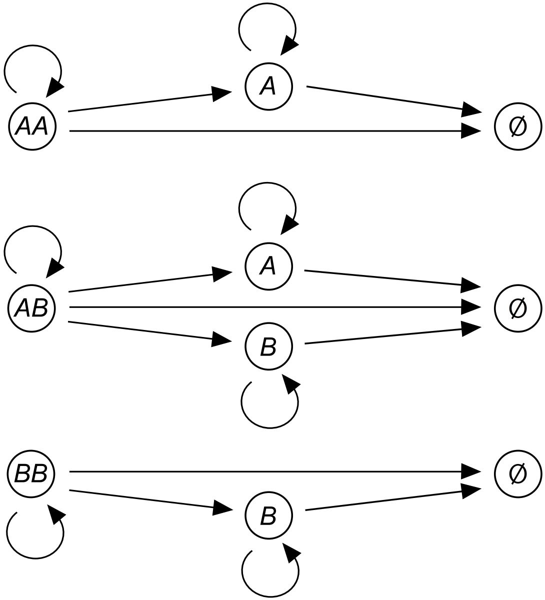

Stochastic process with six states:

- pair state AA
- pair state AB
- pair state BB
- loner state A
- loner state B
- empty state $\emptyset$

A six-state Markov model of the iterated two-player survival game

$$
\begin{array}{c} \\ AA \\ AB \\ BB \\ A \\ B \\ \emptyset \end{array}
\begin{array}{c}
\begin{array}{cccccc} AA & AB & BB & A & B & \emptyset \end{array} \\
\begin{pmatrix}
a^2 & 0 & 0 & 2a(1-a) & 0 & (1-a)^2 \\
0 & bc & 0 & b(1-c) & c(1-b) & (1-b)(1-c) \\
0 & 0 & d^2 & 0 & 2d(1-d) & (1-d)^2 \\
0 & 0 & 0 & a_0 & 0 & 1-a_0 \\
0 & 0 & 0 & 0 & d_0 & 1-d_0 \\
0 & 0 & 0 & 0 & 0 & 1
\end{pmatrix}
\end{array}
$$

Note: here the members of a pair are not labeled Individual and Partner.

Details about n-step transition probabilities

If the single-iteration matrix on the previous slide is called M, then M^n gives the transition probabilities for n iterations.

Games with a large number of iterations can be characterized by the magnitudes of the eigenvalues of M, which are

$$\lambda = \left\{1, a^2, bc, d^2, a_0, d_0\right\}.$$

Payoffs are always survival probabilities, and we assume

$$0 < a, b, c, d, a_0, d_0 < 1.$$

Ø is an absorbing state. In the limit $n \to \infty$, both players die.

Example n-step calculations

	A	B	$\{\}$
A	a	b	a_0
B	c	d	d_0

Probability that both members of an AA pair survive n iterations:

$$\text{Prob}(AA \to AA) = a^{2n}$$

Probability that A dies along the way but B survives n iterations:

$$\begin{aligned}\text{Prob}(AB \to B) &= \sum_{i=1}^{n} (bc)^{i-1} c(1-b) d_0^{n-i} \\ &= \left((bc)^n - d_0^n\right) \frac{c(1-b)}{bc - d_0}\end{aligned}$$

Key point: an individual's situation can change drastically.

The two key fitness differences in this model

	A	B
A	$a(n)$	$b(n)$
B	$c(n)$	$d(n)$

$$a(n) - c(n) = \frac{a - a_0}{a^2 - a_0} a^{2n} + \frac{a(1-a)}{a_0 - a^2} a_0^n - \frac{c - d_0}{bc - d_0} (bc)^n - \frac{c(1-b)}{d_0 - bc} d_0^n$$

$$b(n) - d(n) = \frac{b - a_0}{bc - a_0} (bc)^n + \frac{b(1-c)}{a_0 - bc} a_0^n - \frac{d - d_0}{d^2 - d_0} d^{2n} - \frac{d(1-d)}{d_0 - d^2} d_0^n$$

which may then be used, for example, in

$$\frac{dx}{dt} = x(1-x)\Big(\big[a(n) - c(n)\big]x + \big[b(n) - d(n)\big](1-x)\Big)$$

Prolonged survival games and cooperation

	A	B
A	$a(n)$	$b(n)$
B	$c(n)$	$d(n)$

Single-iteration games collapse to the usual two-player games:

$$a(1) - c(1) = a - c \quad \text{and} \quad b(1) - d(1) = b - d.$$

Very, very long games become neutral:

$$\lim_{n\to\infty} a(n) - c(n) = 0 \quad \text{and} \quad \lim_{n\to\infty} b(n) - d(n) = 0.$$

The logic of large-n conclusions: if $dx/dt < 0$ when $n = 1$ and the neutral limit is approached *from above*, then there exists a number of iterations above which A becomes favored.

Example: Prisoner's Dilemma, low survival

	A	B	$\{\}$
A	0.8	0.6	0.3
B	0.9	0.7	0.3

Parameters: $a = 0.8$, $b = 0.6$, $c = 0.9$, $d = 0.7$ and $a_0 = d_0 = 0.3$.

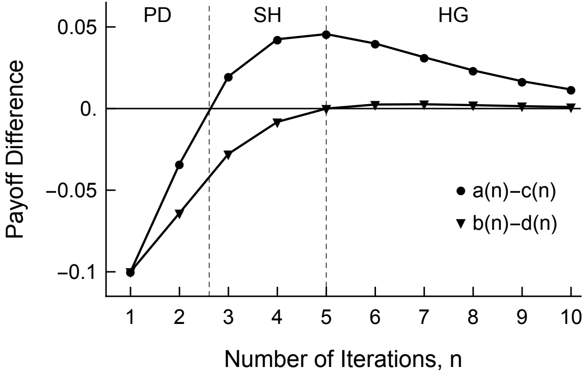

PD, high survival and different a, d spacings

	A	B	$\{\}$
A	a	b	a_0
B	c	d	d_0

$$c > a > d > b > a_0 = d_0$$

	A	B	$\{\}$
A	0.97	0.94	0.9
B	0.99	0.95	0.9

	A	B	$\{\}$
A	0.9733	0.94	0.9
B	0.99	0.9567	0.9

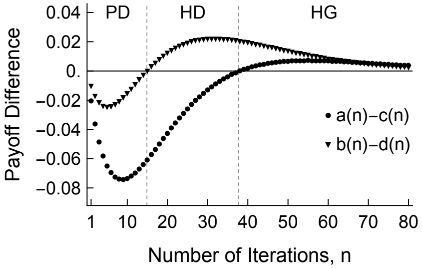

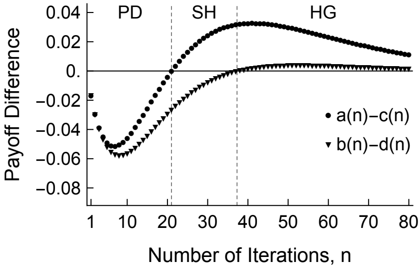

Same two PD cases depicted differently

	A	B	$\{\}$
A	a	b	a_0
B	c	d	d_0

	A	B	$\{\}$
A	0.97	0.94	0.9
B	0.99	0.95	0.9

	A	B	$\{\}$
A	0.9733	0.94	0.9
B	0.99	0.9567	0.9

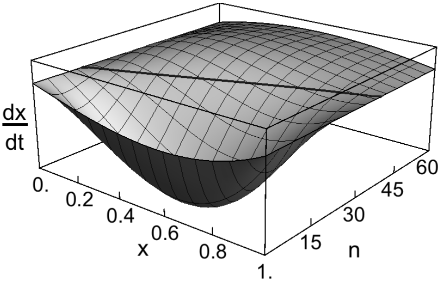

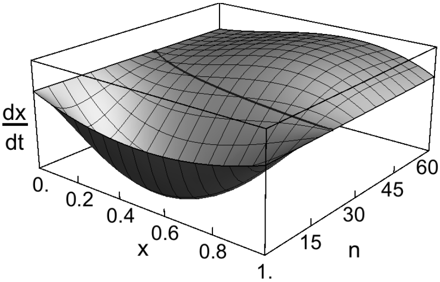

Example Stag Hunt and Hawk-Dove game

	A	B	$\{\}$
A	a	b	a_0
B	c	d	d_0

	A	B	$\{\}$
A	0.99	0.94	0.9
B	0.97	0.97	0.9

	A	B	$\{\}$
A	0.97	0.95	0.9
B	0.99	0.94	0.9

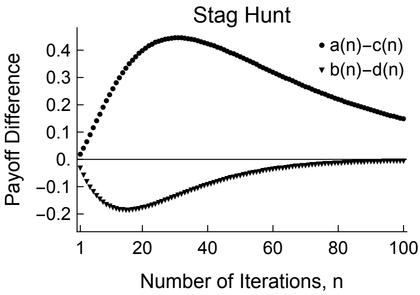

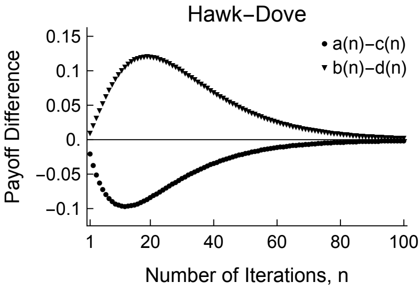

Large-n approximation for examples so far

	A	B
A	$a(n)$	$b(n)$
B	$c(n)$	$d(n)$

If AA survives better than AB and BB, and better than A and B, this means a^2 is the largest non-unit eigenvalue. Then for large n

$$a(n) - c(n) \approx \frac{a - a_0}{a^2 - a_0} a^{2n} > 0$$

$$b(n) - d(n) \approx 0$$

In this case, the cooperative type A eventually becomes favored:

$$\frac{dx}{dt} \approx x^2(1 - x)\frac{a - a_0}{a^2 - a_0} a^{2n} > 0 \quad \text{for all } x \in (0, 1)$$

when n is very large, regardless of the type of game at $n = 1$.

Other large-n approximations

	A	B
A	$a(n)$	$b(n)$
B	$c(n)$	$d(n)$

Depending on which of the non-unit eigenvalues $\{a^2, bc, d^2, a_0, d_0\}$ is the largest, a number of other large-n results can be obtained.

$$a(n) - c(n) = \frac{a - a_0}{a^2 - a_0} a^{2n} + \frac{a(1-a)}{a_0 - a^2} a_0^n - \frac{c - d_0}{bc - d_0} (bc)^n - \frac{c(1-b)}{d_0 - bc} d_0^n$$

$$b(n) - d(n) = \frac{b - a_0}{bc - a_0} (bc)^n + \frac{b(1-c)}{a_0 - bc} a_0^n - \frac{d - d_0}{d^2 - d_0} d^{2n} - \frac{d(1-d)}{d_0 - d^2} d_0^n$$

$$\frac{dx}{dt} = x(1-x)\Big(\big[a(n) - c(n)\big]x + \big[b(n) - d(n)\big](1-x)\Big)$$

Ex. game that doesn't enhance survival

	A	B	$\{\}$
A	0.93	0.91	0.97
B	0.95	0.90	0.97

$$a_0 = d_0 = 0.97, \text{ with } a = 0.93,\ b = 0.91,\ c = 0.95,\ d = 0.90$$

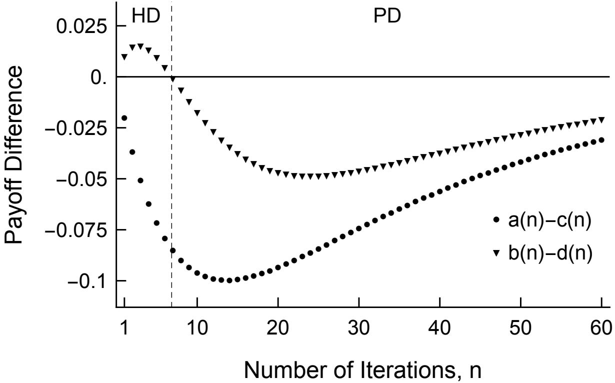

Finite populations

- No population is infinite. Use N for the population size.
- We hope dx/dt and Δx are of heuristic value for large N.
- But A can fix or be lost regardless of dx/dt or Δx.
- To decide whether A is favored or disfavored, can appeal to fixation probabilities or equilibrium frequencies.

For one of the main findings of this work—that the cooperative type A becomes uniformly favored ($dx/dt > 0$ for all $x \in (0, 1)$) when n is large—extension to finite N is fairly straightforward.

Modified Moran model with viabilities from the survival game

A population of N (haploid) individuals. In each time step:

- two individuals chosen without replacement to play the game;
- 0, 1, or 2 individuals do not survive;
- 0, 1, or 2 parents chosen with replacement;
- their offspring replace the individuals who died;
- offspring gets parent's type without modification.

Sampling randomly without replacement precludes self-interaction.

... with replacement gives everyone an equal chance to reproduce.

Transition probabilities for the number (K) of type A individuals

$$P_{K,K+2} = \frac{N-K}{N}\frac{N-K-1}{N-1}\text{Prob}(BB \to \emptyset)\frac{K}{N}\frac{K}{N}$$

$$P_{K,K+1} = \frac{N-K}{N}\frac{N-K-1}{N-1}\left[\text{Prob}(BB \to B)\frac{K}{N} + \text{Prob}(BB \to \emptyset)2\frac{K}{N}\frac{N-K}{N}\right] + 2\frac{K}{N}\frac{N-K}{N-1}\left[\text{Prob}(AB \to A)\frac{K}{N} + \text{Prob}(AB \to \emptyset)\frac{K}{N}\frac{K}{N}\right]$$

$$P_{K,K-1} = 2\frac{K}{N}\frac{N-K}{N-1}\left[\text{Prob}(AB \to B)\frac{N-K}{N} + \text{Prob}(AB \to \emptyset)\frac{N-K}{N}\frac{N-K}{N}\right] + \frac{K}{N}\frac{K-1}{N-1}\left[\text{Prob}(AA \to A)\frac{N-K}{N} + \text{Prob}(AA \to \emptyset)2\frac{K}{N}\frac{N-K}{N}\right]$$

$$P_{K,K-2} = \frac{K}{N}\frac{K-1}{N-1}\text{Prob}(AA \to \emptyset)\frac{N-K}{N}\frac{N-K}{N}$$

$$P_{K,K} = 1 - P_{K,K+2} - P_{K,K+1} - P_{K,K-1} - P_{K,K-2}$$

Expected value of ΔK over one time step in this Moran model

The expected change in the number of A individuals is given by

$$E[\Delta K] = 2P_{K,K+2} + P_{K,K+1} - P_{K,K-1} - 2P_{K,K-2}$$

which simplifies to

$$E[\Delta K] = 2\frac{K}{N}\frac{N-K}{N}\left(a(n)\frac{K-1}{N-1} - c(n)\frac{K}{N-1} + b(n)\frac{N-K}{N-1} - d(n)\frac{N-K-1}{N-1}\right)$$

with $a(n)$, $b(n)$, $c(n)$, $d(n)$ as before. Note that $x = K/N$.

Then, in the case of interest, when a^2 is the largest non-unit eigenvalue and $a(n) \gg b(n), c(n), d(n)$ for large n, we have $E[\Delta K] > 0$ for $K \geqslant 2$. We need to reexamine only the smallest frequency ($K = 1$ or $x = 1/N$).

Examination of the terminal frequency ($K = 1$ or $x = 1/N$)

When there is only one cooperator in the population:

- For the majority of Prisoner's Dilemmas and for many Hawk-Dove games, the next largest eigenvalue will be bc. In this case

$$E[\Delta K; K = 1] \approx 2 \frac{1}{N} \frac{N-1}{N} \left(\frac{b - a_0}{bc - a_0} - \frac{1}{N-1} \frac{c - d_0}{bc - d_0} \right) (bc)^n$$

  which will be > 0 unless N is very small.

- In the Stag Hunt, the second largest eigenvalue will be d^2. Here

$$E[\Delta K; K = 1] \approx -2 \frac{1}{N} \frac{N-2}{N} \frac{d - d_0}{d^2 - d_0} d^{2n}$$

  which will be < 0 unless $N = 2$.

Designing games directly in terms of survival

Any iterated evolutionary game can be re-parameterized in terms of fitness (here survival, or viability). To illustrate, we can design a game in the spirit of Model 2 of De Jaegher and Hoyer (2016).

Five parameters:

d_0	baseline individual survival
g	survival benefit of partnership
u	cost of cooperation
v	benefit of having an A partner
n	number of iterations (attacks)

Individuals survive an attack with probabilities that depend on their type and their partner's type:

A with A:	$a = d_0 + g - u + v$
A with B:	$b = d_0 + g - u$
B with A:	$c = d_0 + g + v$
B with B:	$d = d_0 + g$
A alone:	$a_0 = d_0 - u$
B alone:	d_0

Akin to De Jaegher and Hoyer (2016)

	A	B	$\{\}$
A	0.94	0.91	0.78
B	0.96	0.93	0.80

$d_0 = 0.8$, $g = 0.13$, $u = 0.02$, $v = 0.03$ give the matrix above, and

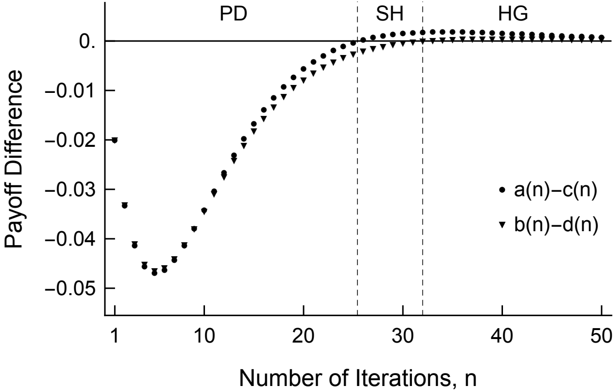

Findings for iterated two-player survival games

- When the game enhances survival, cooperative behaviors may become favored as the level of adversity *increases*.

  1. When the number of iterations, n, is large.

     This is true for all types of single-iteration games, including Prisoner's Dilemmas, Stag Hunts, and Hawk-Dove games.

  2. When probabilities of survival are small.

     Then cooperation becomes favored sooner, for smaller n.

  This holds even with loner survival probabilities $d_0 > a_0$.

- When the game does not enhance survival, such behaviors may be even more disfavored for large n than at $n = 1$.

Thank you!