Multi-Agent Systems and Evolution

Day 2: Evolutionary Game Theory

Elias Fernández Domingos

Outline of the course

- Day 1: Introduction to Game Theory
- **Day 2: Evolutionary Game Theory**
- Day 3: Games on Networks
- Day 4: Practical challenges and connecting theory to Behavioural Experiments
- Day 5: Final remarks and Project presentations

Day 2: Evolutionary Game Theory

1. Evolutionary Stability
2. Infinite Populations
3. Finite Populations
4. Tutorial: how to reproduce an EGT paper

When you think about Game Theory...

games

economy

politics

However, there are many other strategic interactions, and many of them occur in Large populations!

However, there are many other strategic interactions, and many of them occur in Large populations!

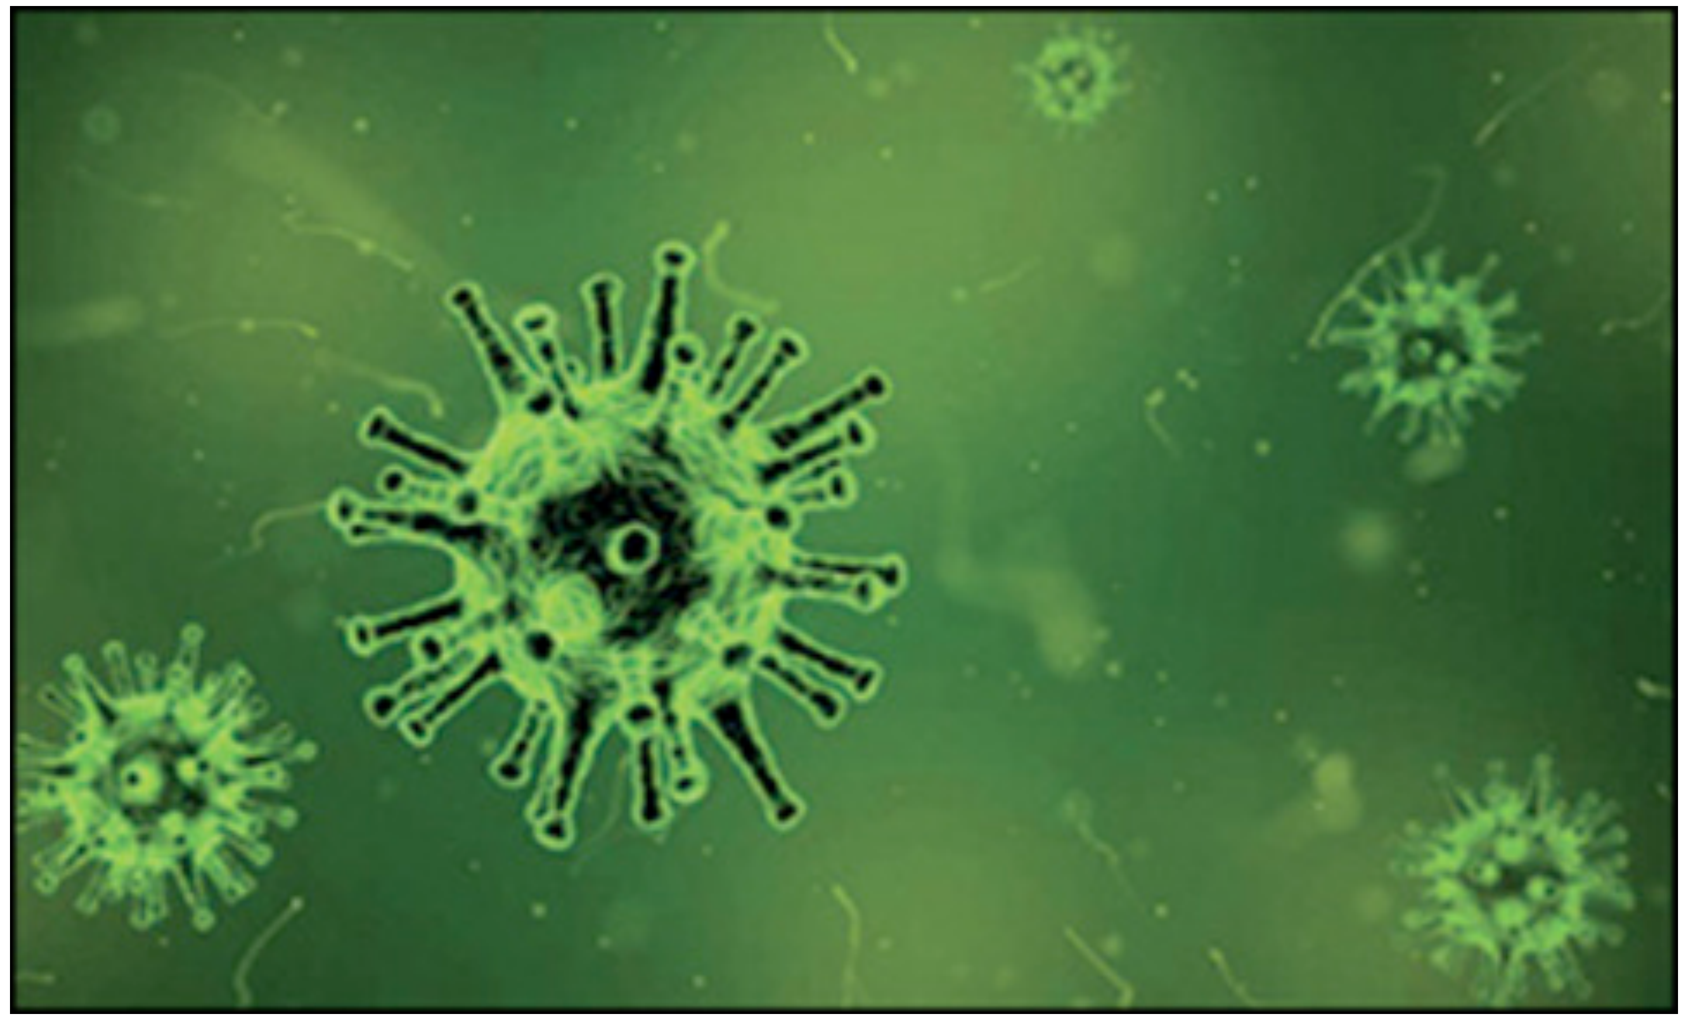

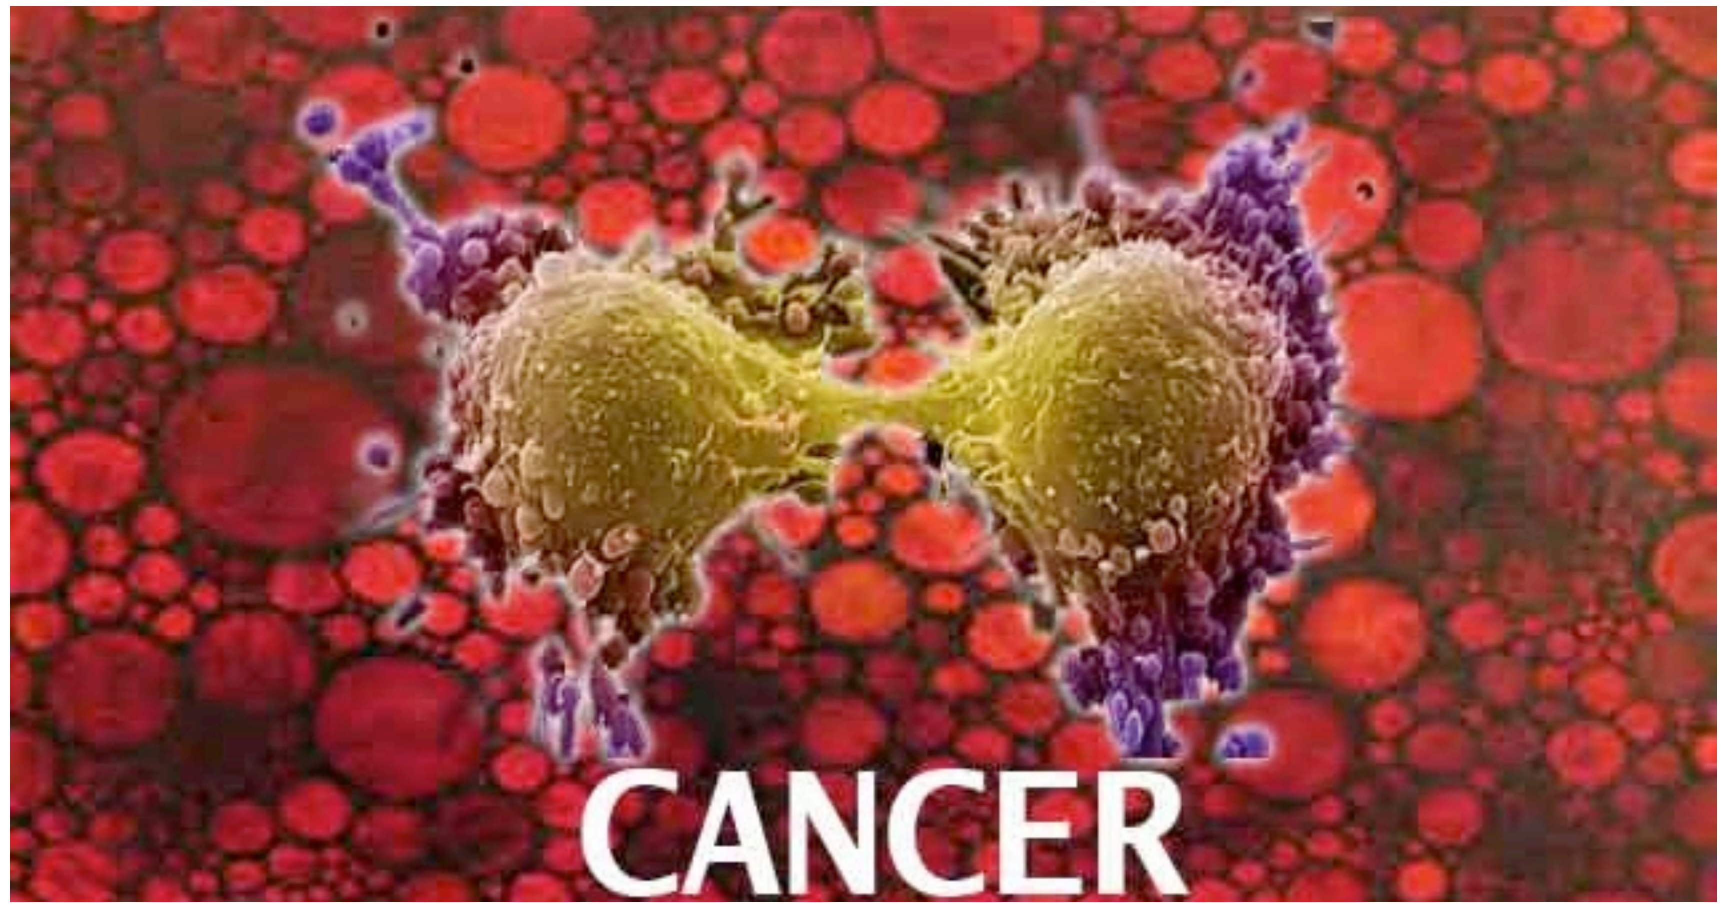

However, there are many other strategic interactions, and many of them occur in Large populations!

A change in perspective from individual to population

Complex individuals - Smaller populations

Simpler individuals - Larger populations

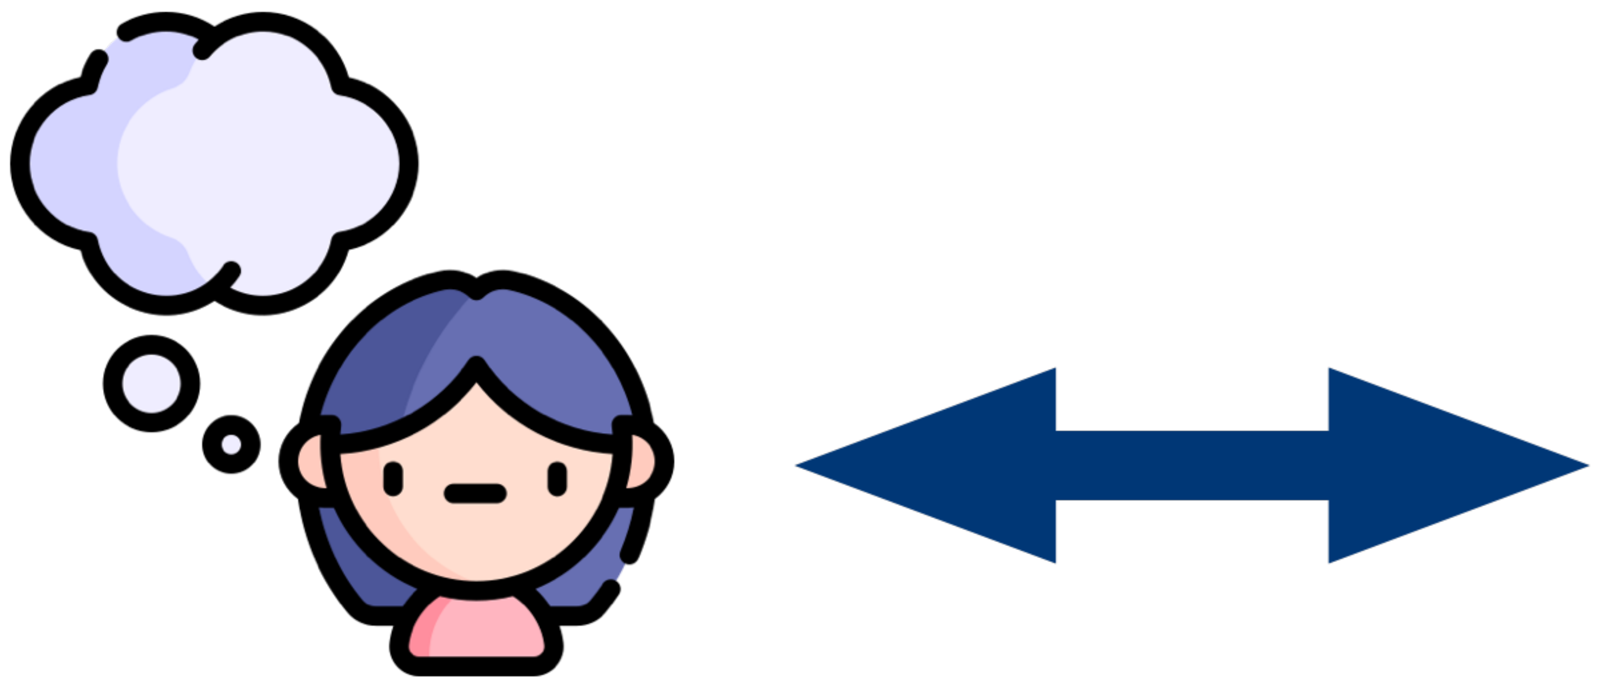

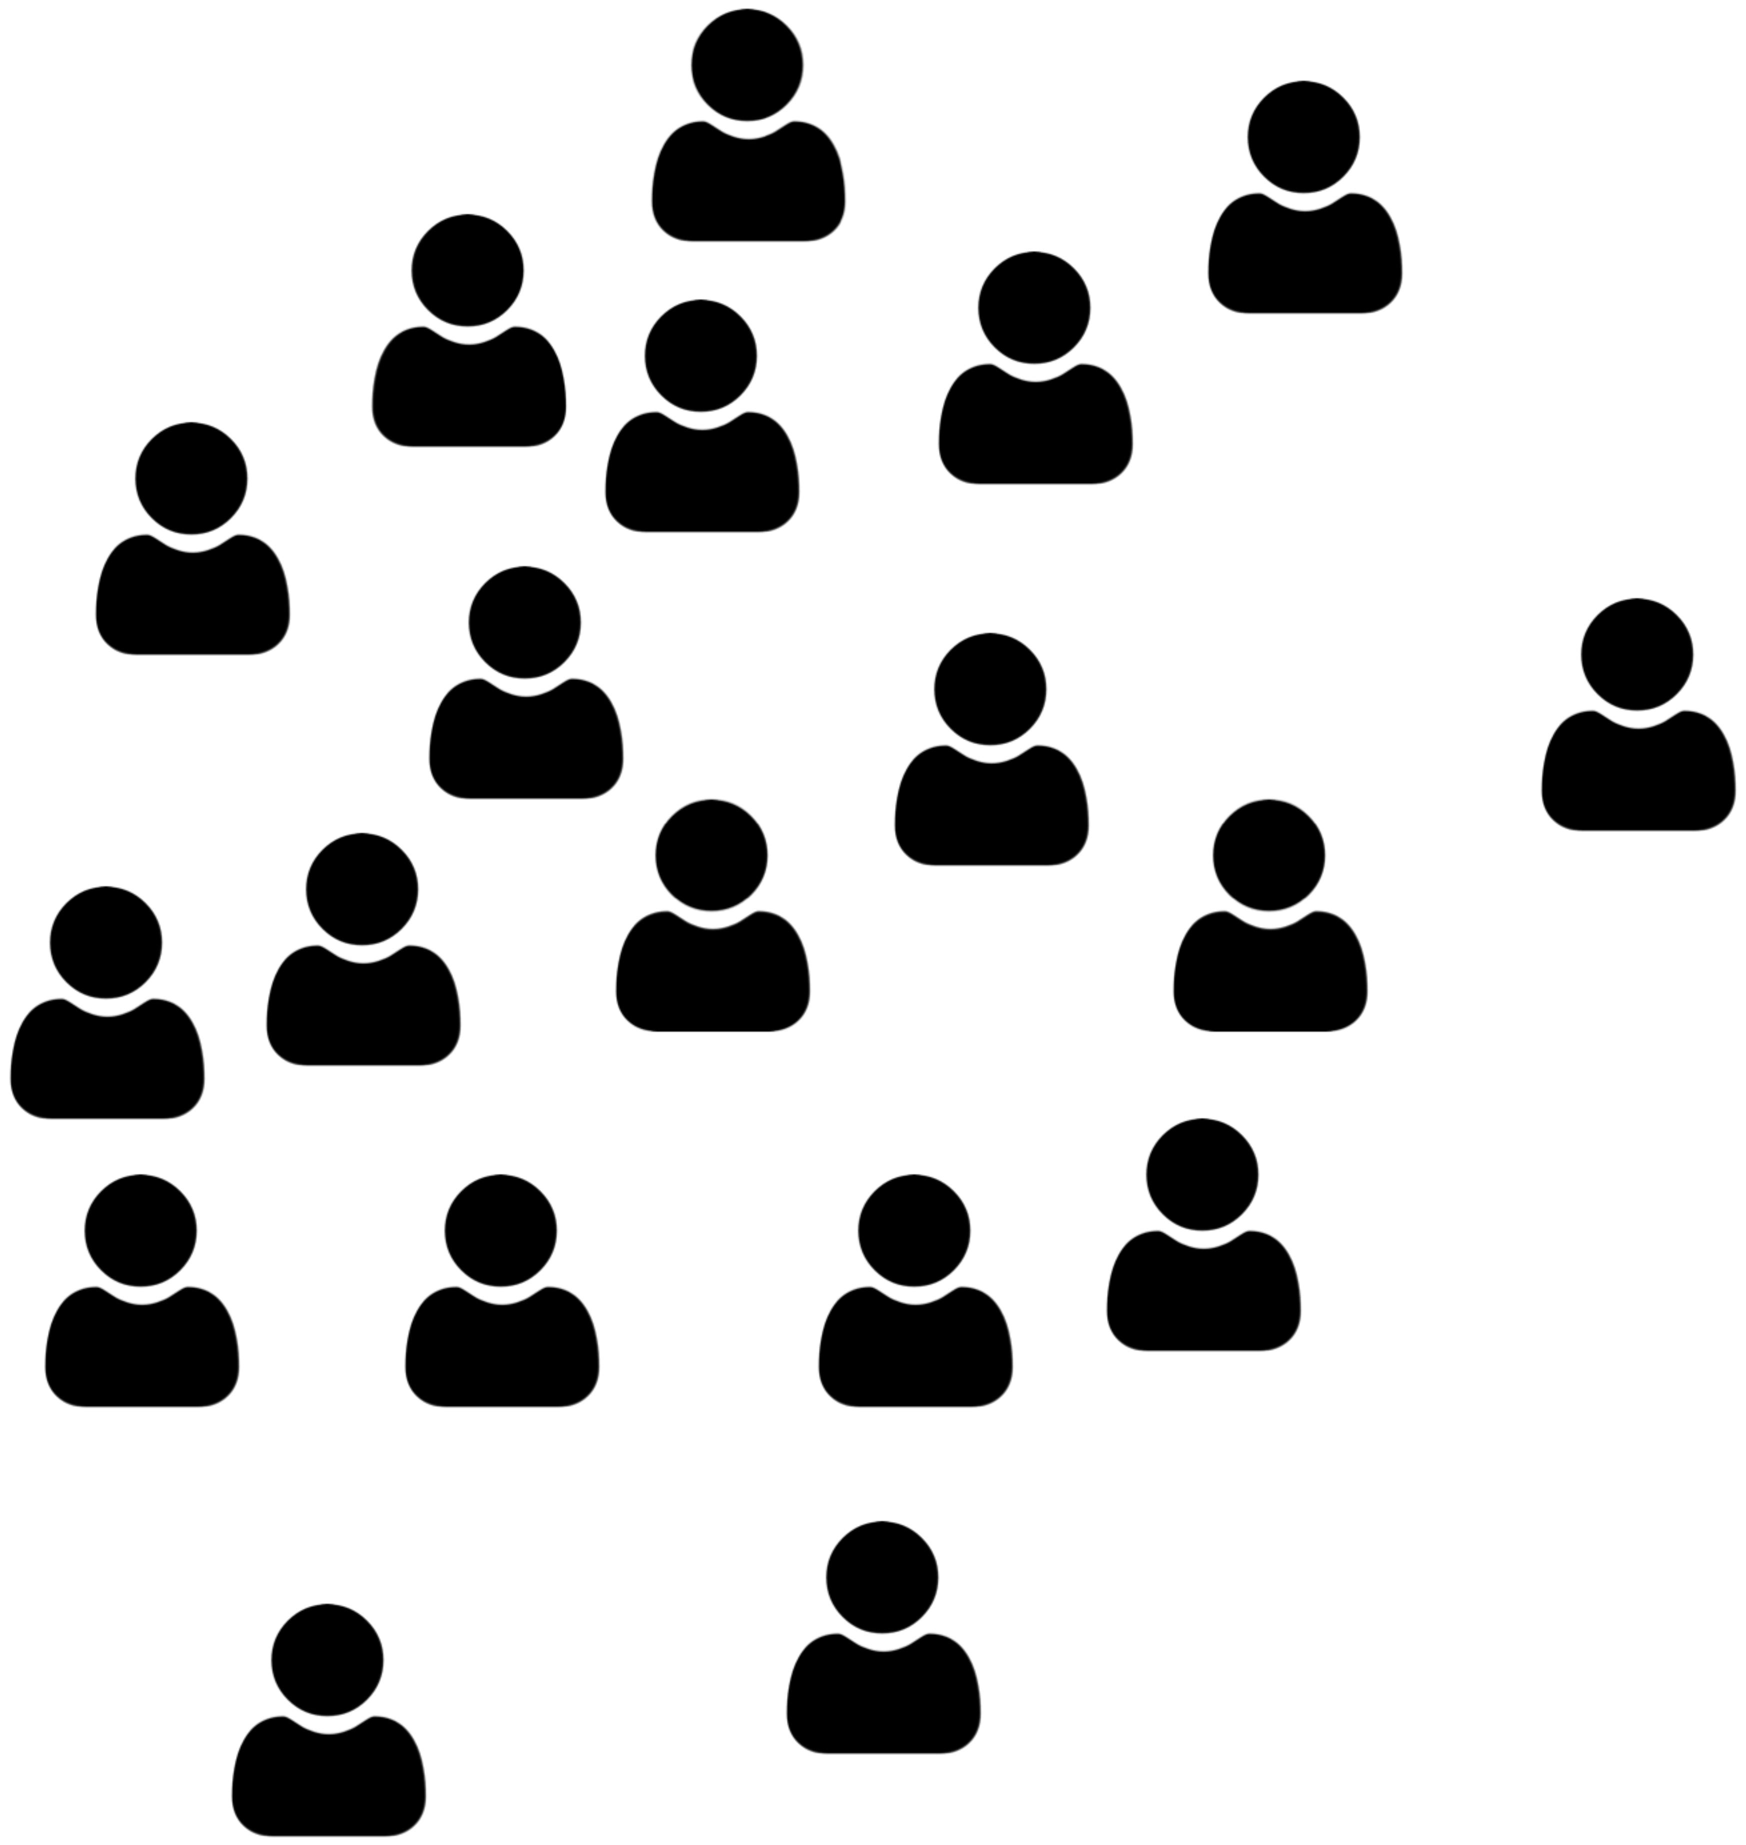

The complexity in EGT often comes from emergent behaviour due to the interactions of many individuals in a population.

https://blogs.bl.uk/untoldlives/2020/03/john-maynard-smith-evolutionary-biology-and-the-logic-of-animal-conflict.html

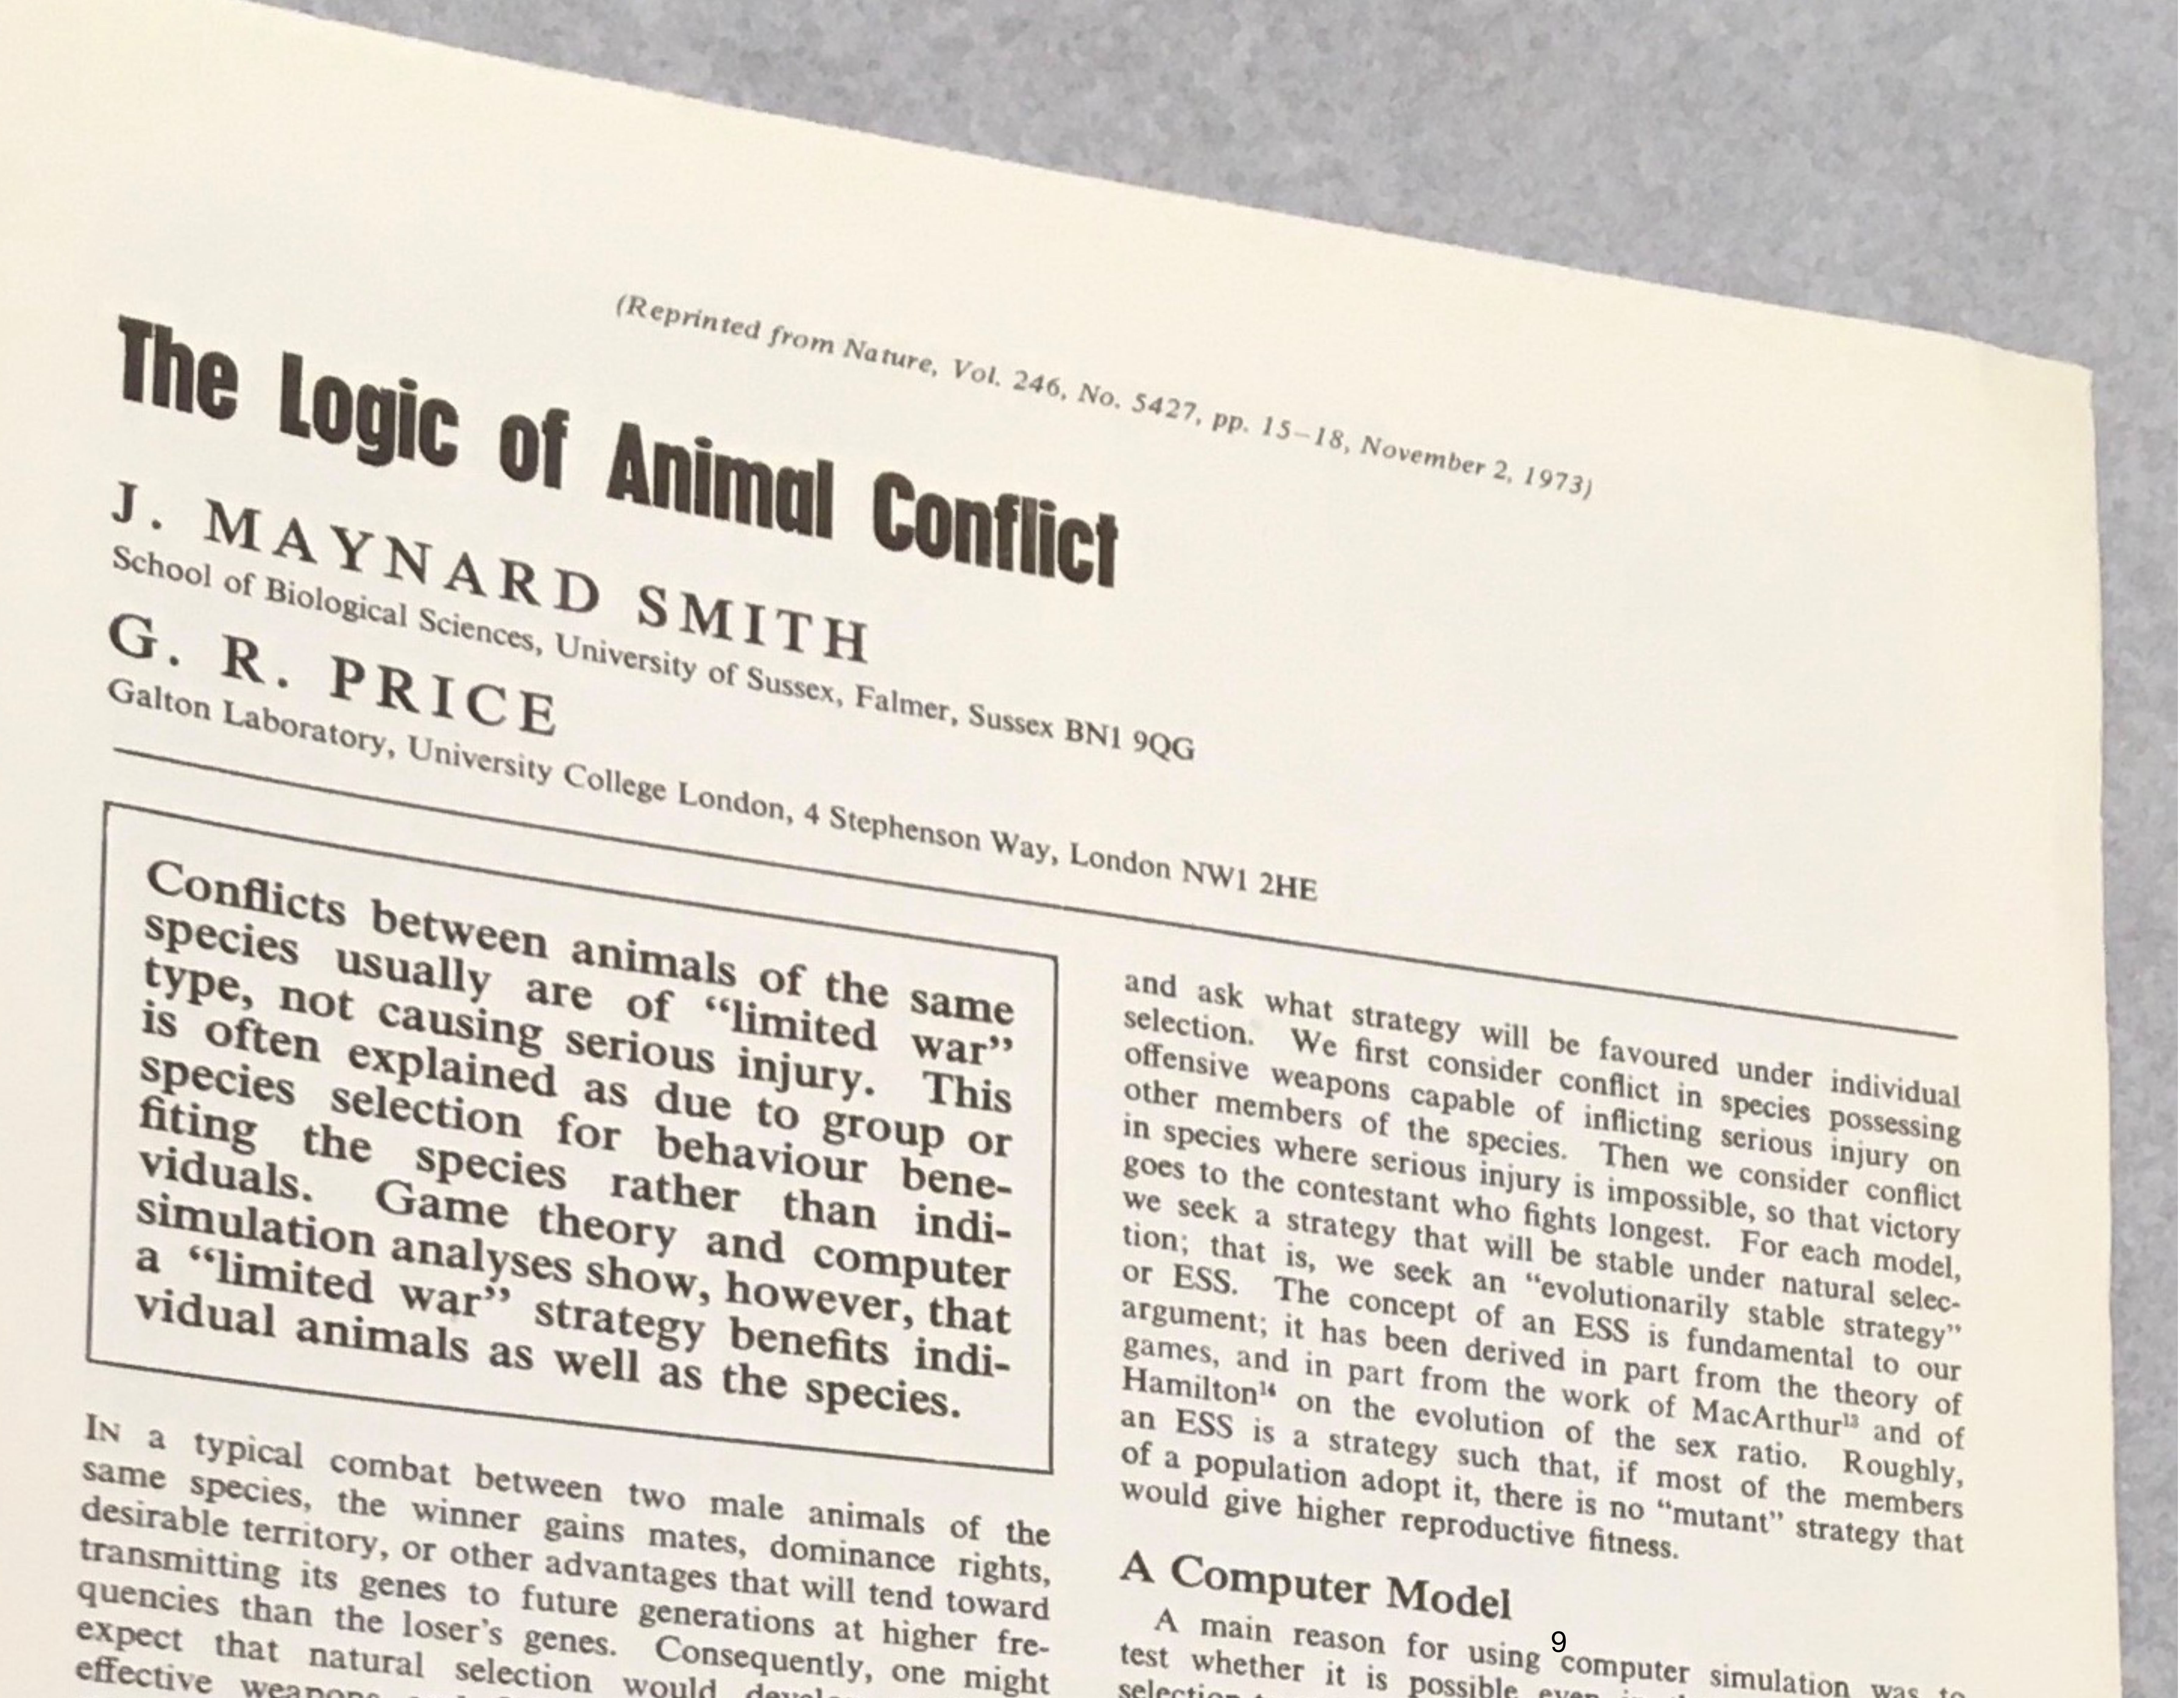

(Reprinted from Nature, Vol. 246, No. 5427, pp. 15–18, November 2, 1973)

The Logic of Animal Conflict

J. MAYNARD SMITH
School of Biological Sciences, University of Sussex, Falmer, Sussex BN1 9QG
G. R. PRICE
Galton Laboratory, University College London, 4 Stephenson Way, London NW1 2HE

Conflicts between animals of the same species usually are of "limited war" type, not causing serious injury. This is often explained as due to group or species selection for behaviour benefiting the species rather than individuals. Game theory and computer simulation analyses show, however, that a "limited war" strategy benefits individual animals as well as the species.

In a typical combat between two male animals of the same species, the winner gains mates, dominance rights, desirable territory, or other advantages that will tend toward transmitting its genes to future generations at higher frequencies than the loser's genes. Consequently, one might expect that natural selection would develop ... effective weapons ...

and ask what strategy will be favoured under individual selection. We first consider conflict in species possessing offensive weapons capable of inflicting serious injury on other members of the species. Then we consider conflict in species where serious injury is impossible, so that victory goes to the contestant who fights longest. For each model, we seek a strategy that will be stable under natural selection; that is, we seek an "evolutionarily stable strategy" or ESS. The concept of an ESS is fundamental to our argument; it has been derived in part from the theory of games, and in part from the work of MacArthur[13] and of Hamilton[14] on the evolution of the sex ratio. Roughly, an ESS is a strategy such that, if most of the members of a population adopt it, there is no "mutant" strategy that would give higher reproductive fitness.

A Computer Model

A main reason for using computer simulation was to test whether it is possible even ...

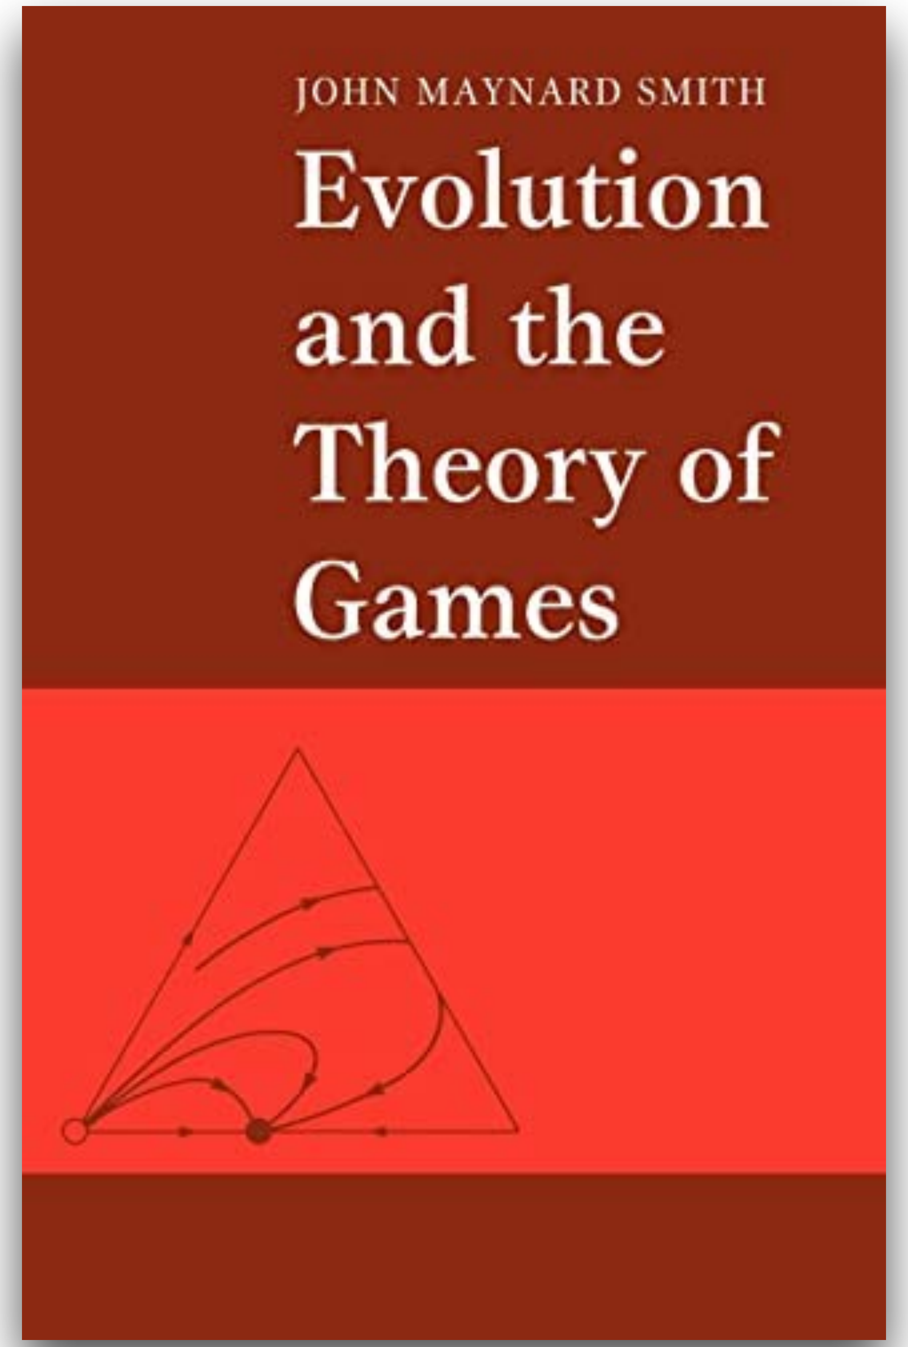

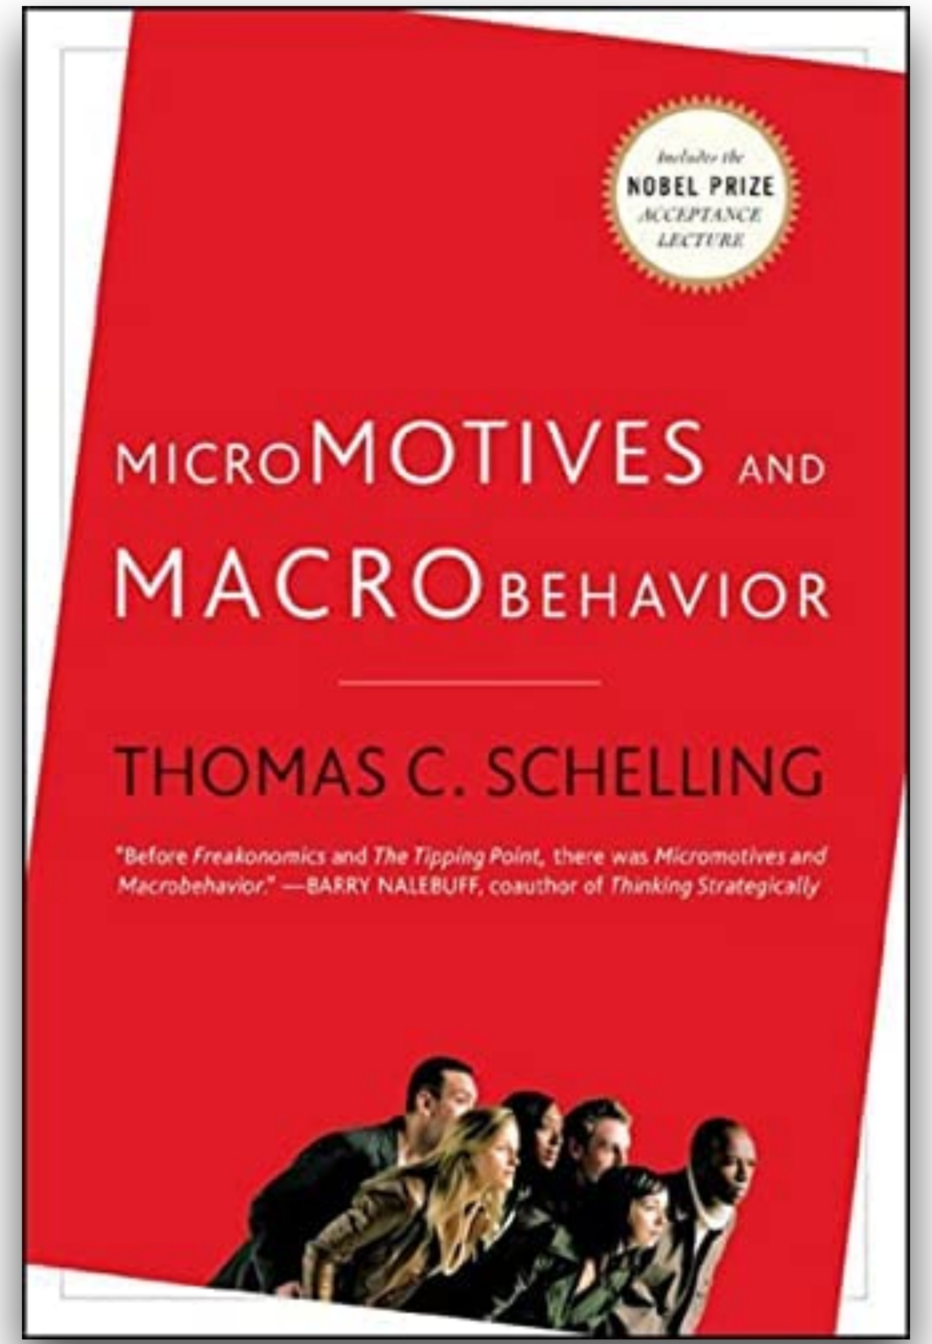

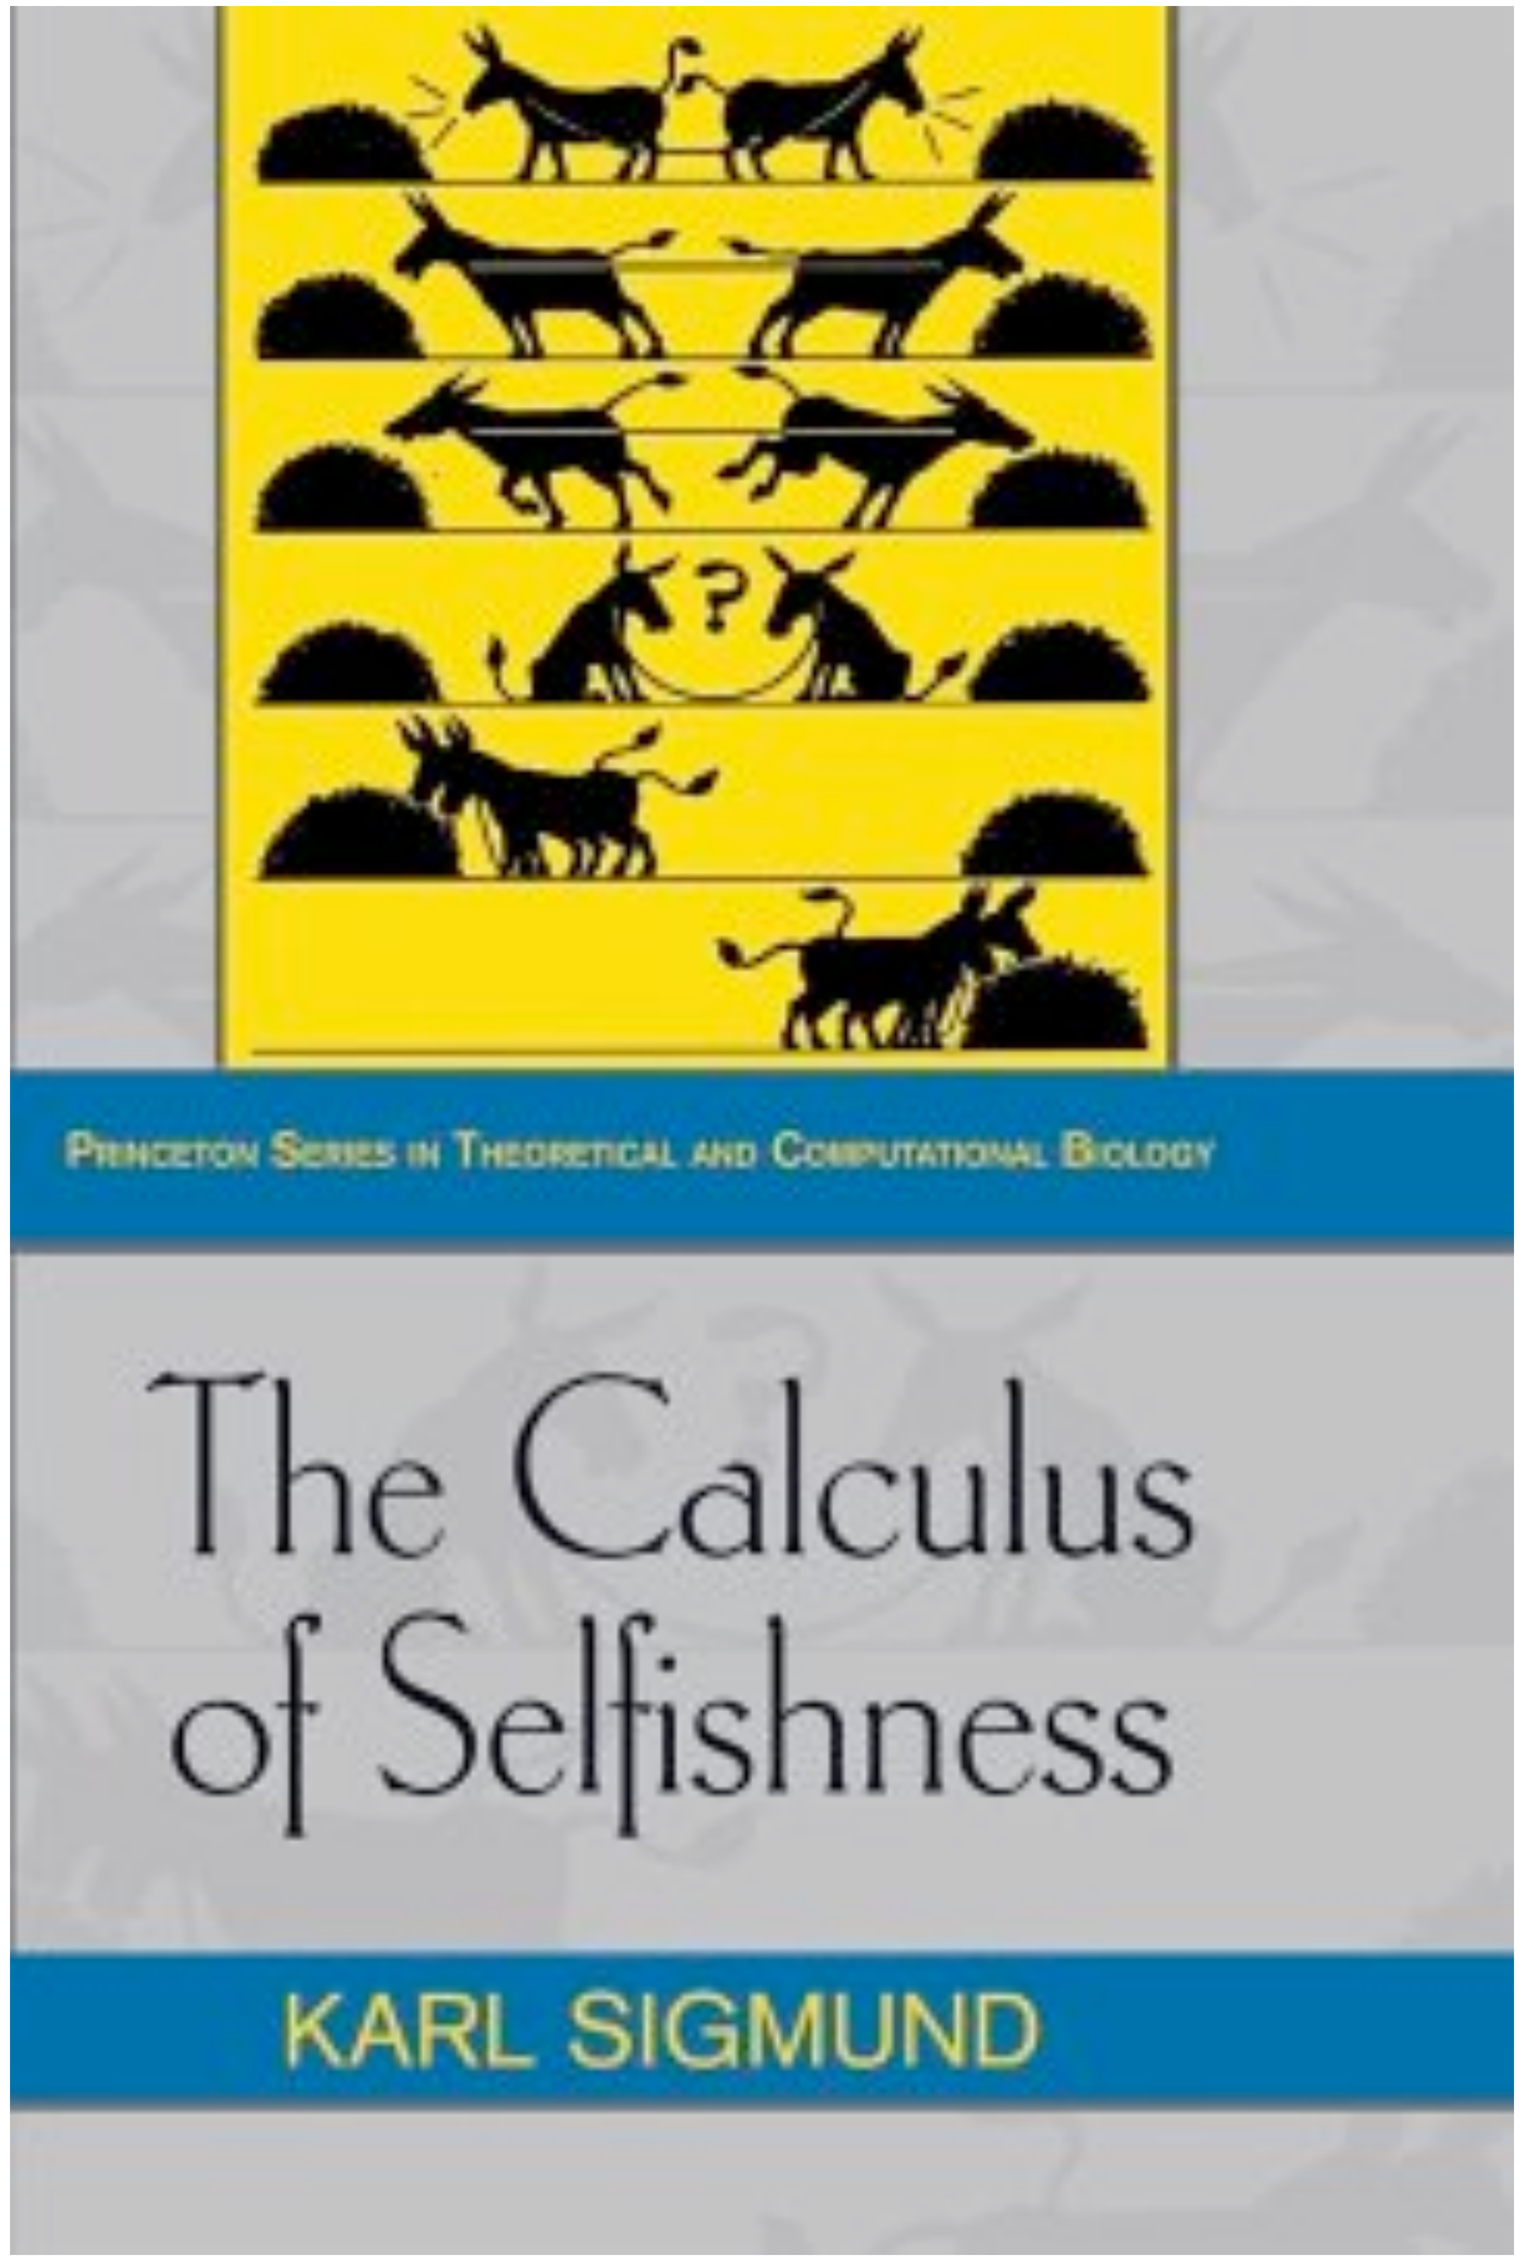
Princeton Series in Theoretical and Computational Biology
The Calculus of Selfishness
KARL SIGMUND

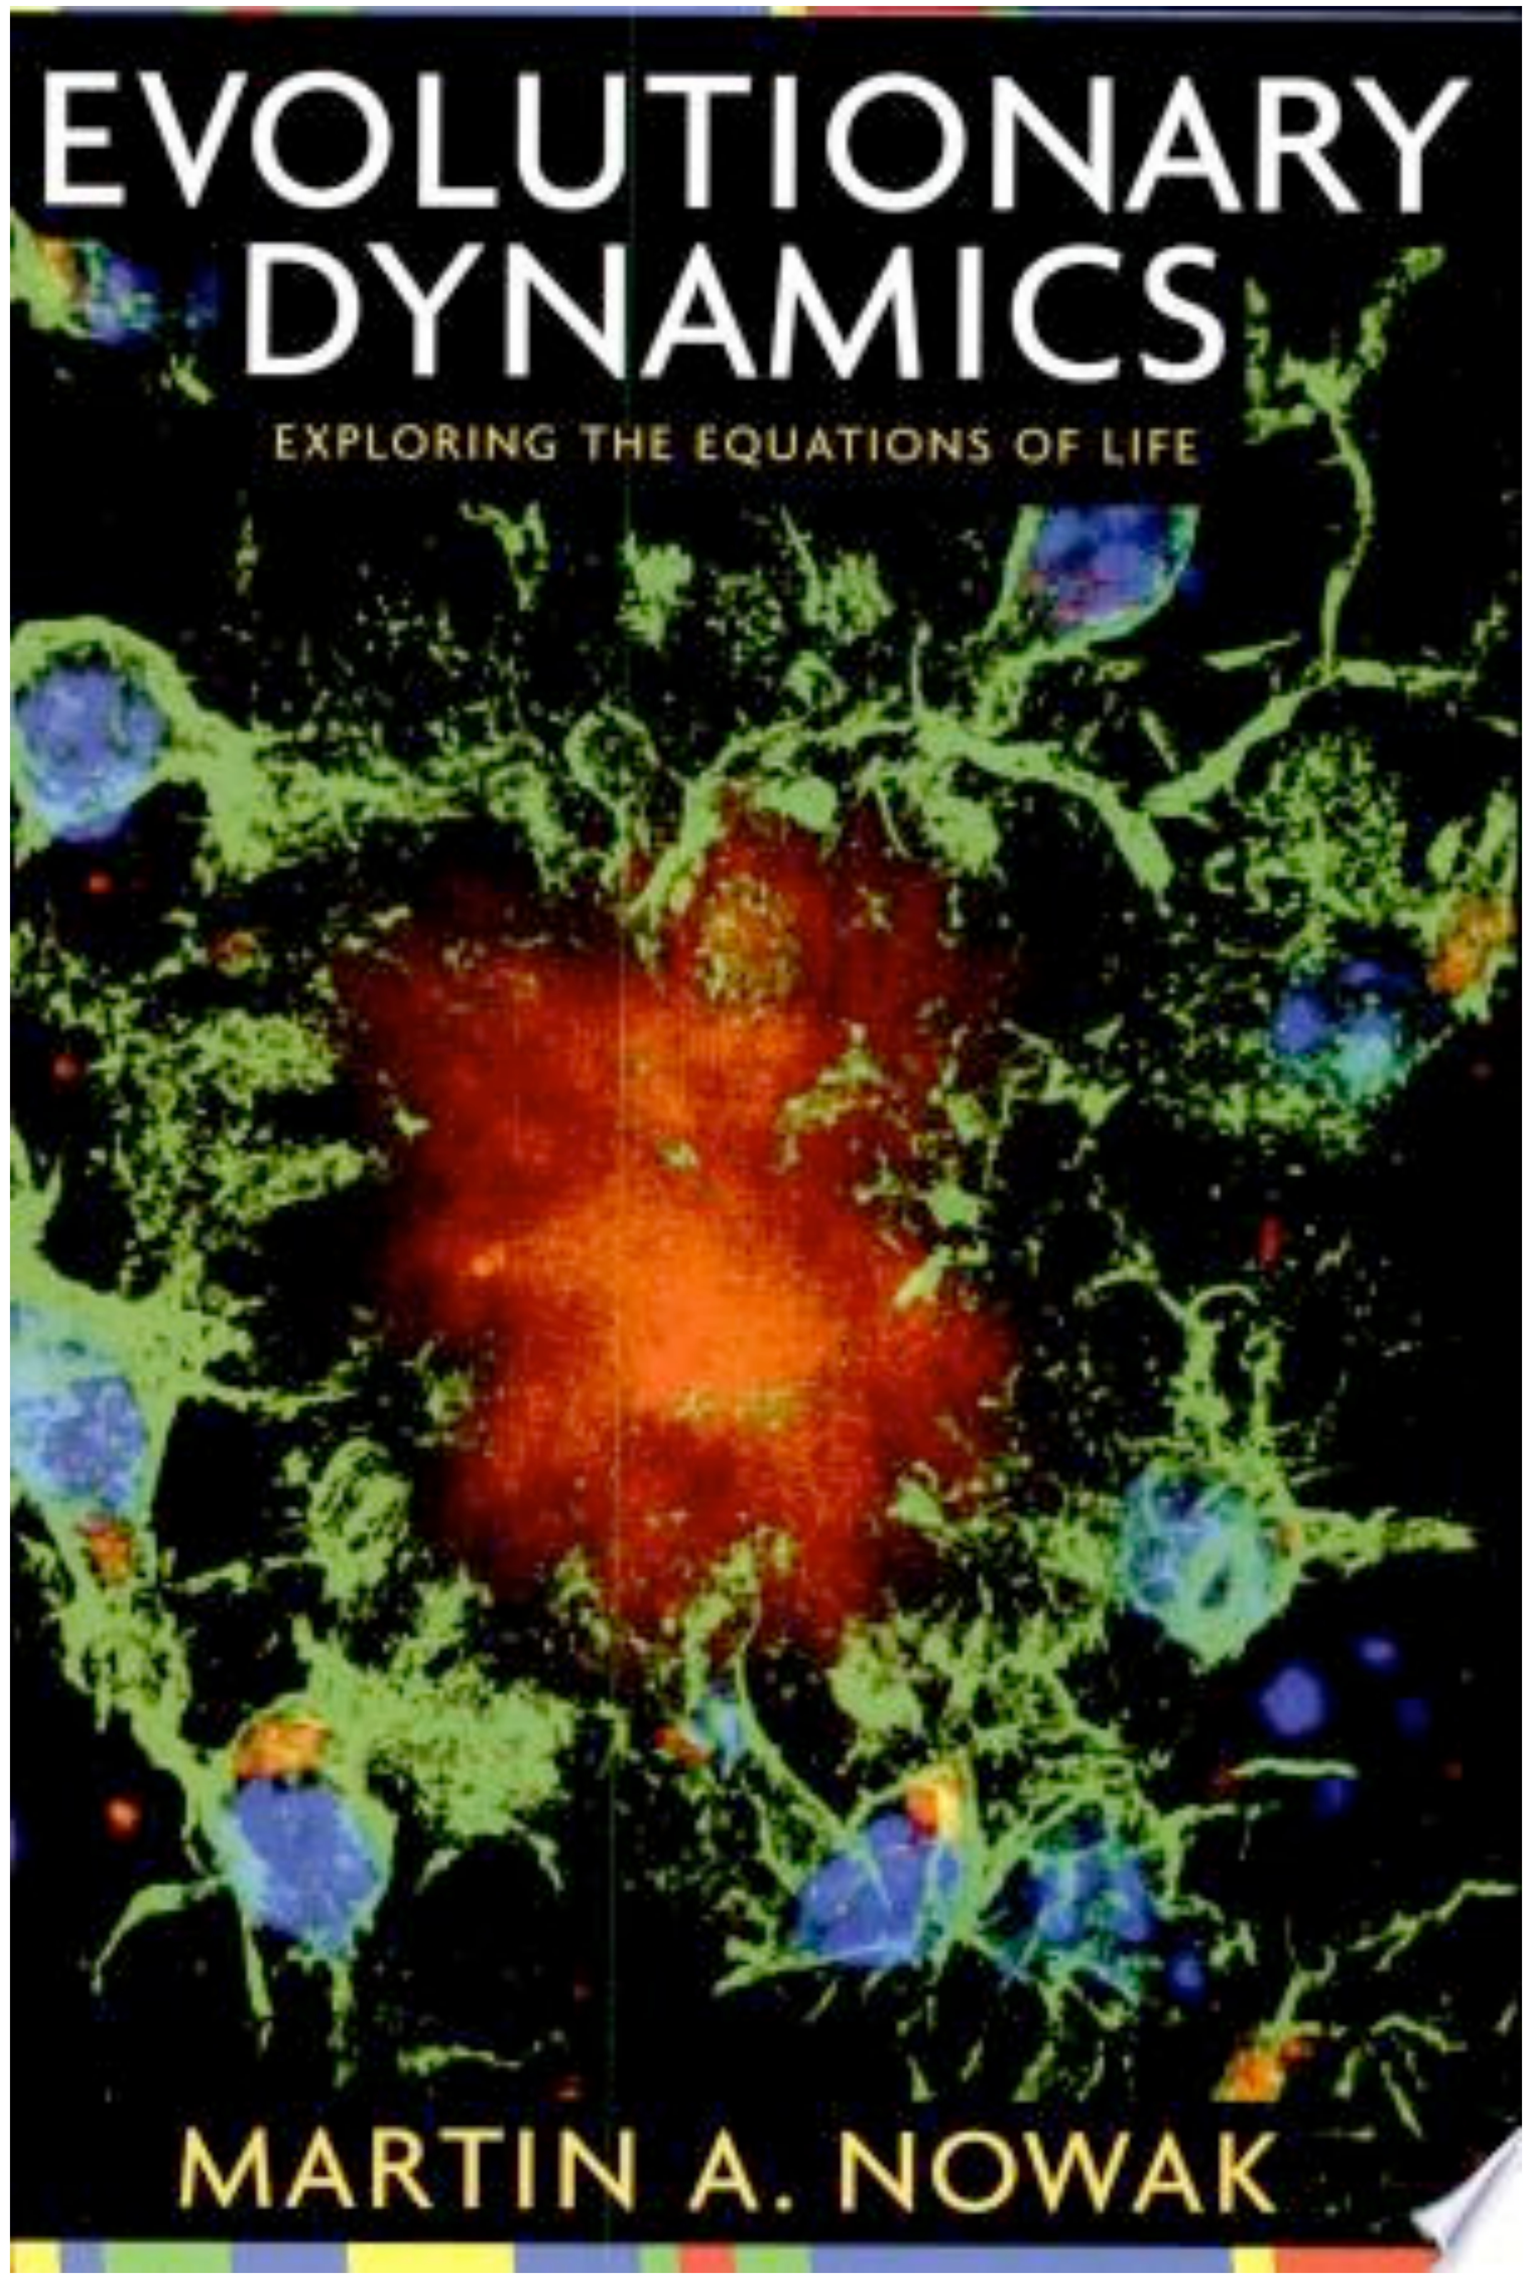
EVOLUTIONARY DYNAMICS
EXPLORING THE EQUATIONS OF LIFE
MARTIN A. NOWAK

Part 1: Evolutionary Stable Strategies (ESS)

Evolutionary Stable Strategy (ESS)

An **Evolutionary Stable Strategy** is an strategy that, if adopted by all individuals of a population, cannot be invaded by alternative or mutant strategies

A strategy S is evolutionary stable if it follows the following 2 conditions for all strategies $T \neq S$:

1. $\Pi(S, S) > \Pi(T, S)$ or
2. $\Pi(S, S) = \Pi(T, S)$ and $\Pi(S, T) > \Pi(T, T)$

Can C invade a population of D

Prisoners Dilemma, T>R, P>S

Greed and fear

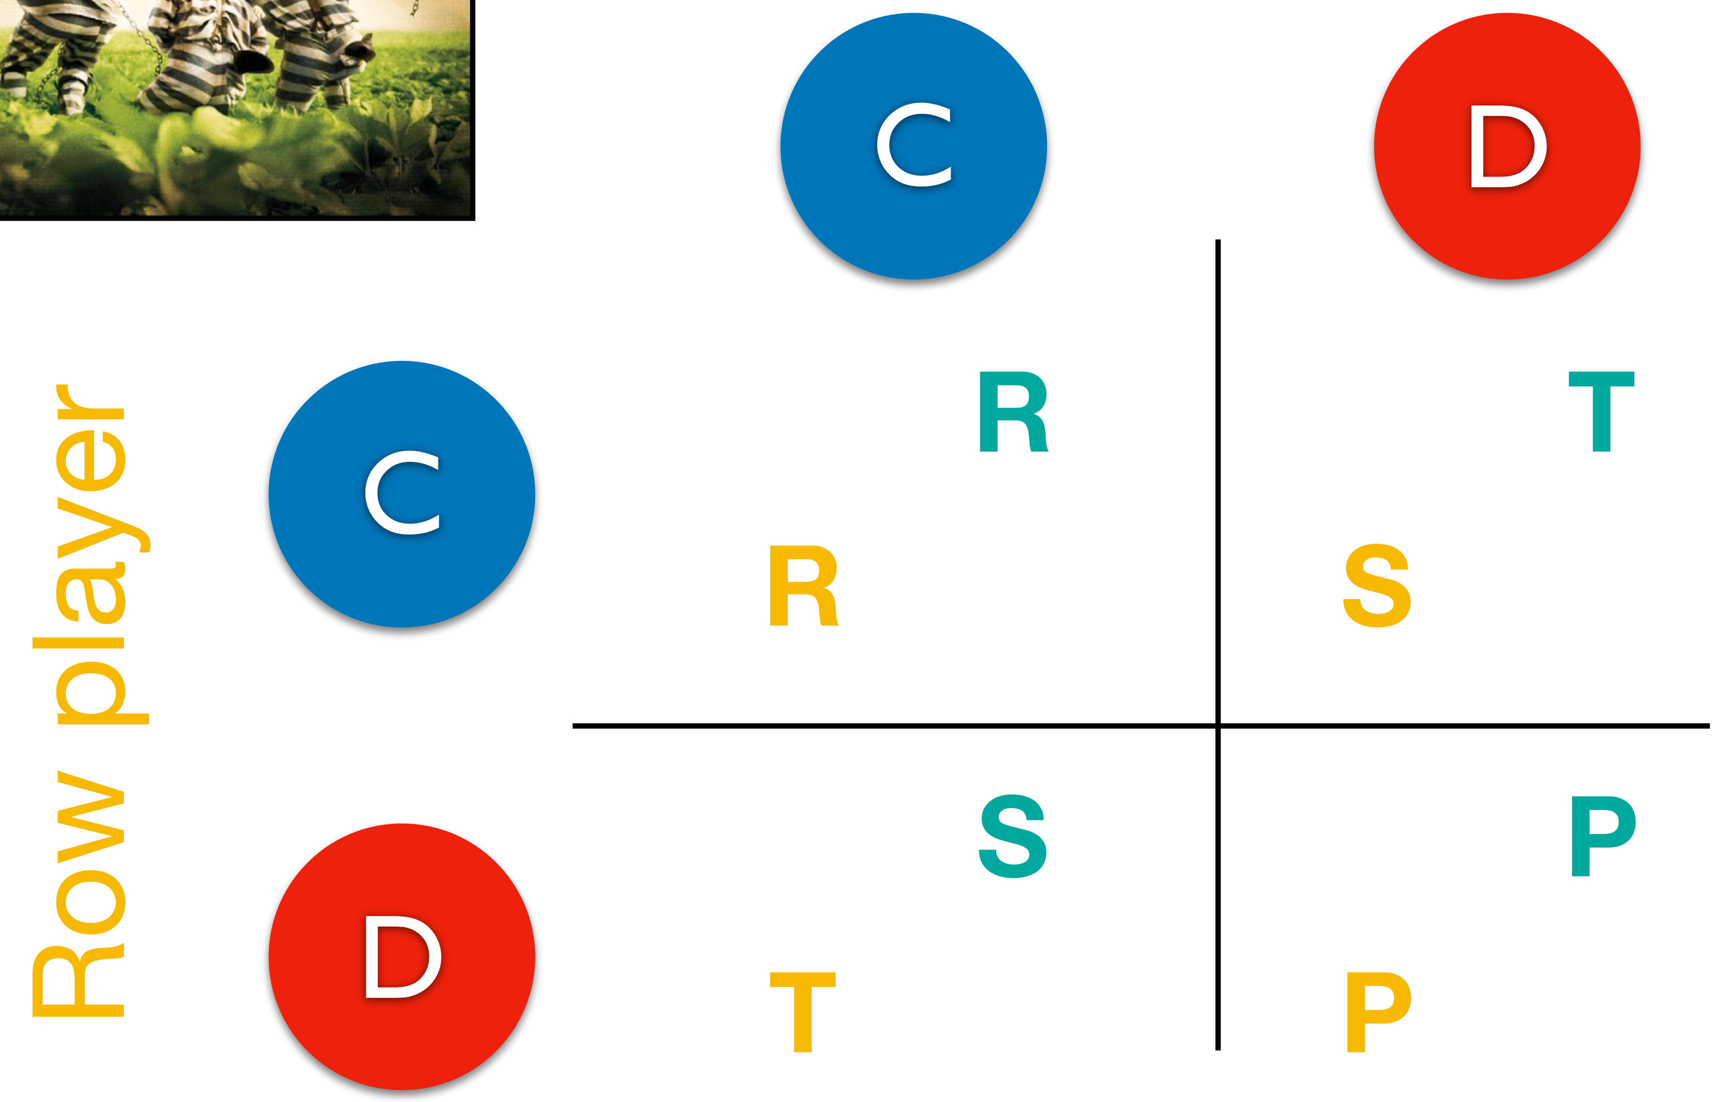

C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

Assume an **infinite population** of $(1-\epsilon)$ D players and ϵ C players

Success of ϵ fraction of C in a population with $(1-\epsilon)$ D players: $S(1-\epsilon)+R\epsilon$

Success of $(1-\epsilon)$ fraction of D in a population with ϵ C players: $P(1-\epsilon)+T\epsilon$

C players can take over the population when $S(1-\epsilon)+R\epsilon > P(1-\epsilon)+T\epsilon$

This happens when either $S>P$ or when $S=P, R>T$

If C cannot invade, D is an **Evolutionary Stable Strategy (ESS)**

…and inversely, can D invade a population of C

Prisoners Dilemma, T>R, P>S

Greed and fear

Column player

Row player	C	D
C	R, R	S, T
D	T, S	P, P

C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

Assume an **infinite population** of $(1-\epsilon)$ C players and ϵ D players

Success of ϵ fraction of D in a population with $(1-\epsilon)$ C players: $T(1-\epsilon)+P\epsilon$

Success of $(1-\epsilon)$ fraction of C in a population with ϵ D players: $R(1-\epsilon)+S\epsilon$

D players can take over the population when $T(1-\epsilon)+P\epsilon > R(1-\epsilon)+S\epsilon$

This happens when either $T>R$ or when $T=R, P>S$

If D cannot invade, C is an **Evolutionary Stable Strategy (ESS)**

	C	D
C	3 3	7 0
D	0 7	1 1

	C	D
C	7 7	3 0
D	0 3	1 1

	C	D
C	3 3	7 1
D	1 7	0 0

C No since $S < P$

D Yes since $T > R$

	C	D
C	3 / 3	7 / 0
D	0 / 7	1 / 1

ESS

	C	D
C	7 / 7	3 / 0
D	0 / 3	1 / 1

	C	D
C	3 / 3	7 / 1
D	1 / 7	0 / 0

D Is an **ESS**

	C	D
C	3, 3	0, 7
D	7, 0	1, 1

ESS

D Is an **ESS**

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

C No since $S < P$

D No since $T < R$

	C	D
C	3, 3	1, 7
D	7, 1	0, 0

	C	D
C	3, 3	0, 7
D	7, 0	1, 1 (ESS)

D Is an **ESS**

	C	D
C	7, 7 (ESS)	0, 3
D	3, 0	1, 1 (ESS)

C Is an **ESS**

D Is an **ESS**

	C	D
C	3, 3	1, 7
D	7, 1	0, 0

	C	D
C	3, 3	0, 7
D	7, 0	1, 1 (ESS)

ESS

D Is an **ESS**

	C	D
C	7, 7 (ESS)	0, 3
D	3, 0	1, 1 (ESS)

ESS

C Is an **ESS**

D Is an **ESS**

	C	D
C	3, 3	1, 7
D	7, 1	0, 0

C Yes since $S > P$

D Yes since $T > R$

	C	D
C	3, 3	0, 7
D	7, 0	1, 1

ESS

D Is an **ESS**

ESS

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

C Is an **ESS**

ESS

D Is an **ESS**

	C	D
C	3, 3	1, 7
D	7, 1	0, 0

No **ESS** ?

What about (1/5,4/5)?

Connection between NE and ESS

Given a symmetric two-player normal-form game $G = (\{1,2\}, A, u)$ and a mixed strategy s, If s is an ESS then (s, s) is a NE of the game G.

Given a symmetric two-player normal-form game $G = (\{1,2\}, A, u)$ and a mixed strategy s, If (s, s) is a **strict symmetric** NE then s is an ESS

Part 2: Infinite Populations

The success of a species depends both on its **fitness** and its **numbers**.

The replicator equation

Avg. fitness of the population

Frequency of strategy i

$$\dot{x}_i = x_i \left[f(x_i) - \sum_{j=1}^{n} x_j f(x_j) \right]$$

Fitness of strategy i

The replicator equation

$$\dot{x}_i = x_i \left[f(x_i) - \sum_{j=1}^{n} x_j f(x_j) \right]$$

increase

$$f_i(x) > \sum_{j=1}^{n} x_j f_j(x)$$

decrease

$$f_i(x) < \sum_{j=1}^{n} x_j f_j(x)$$

$$f_i(x) = \sum_{j=1}^{n} x_j f_j(x)$$

Equilibrium? Only if this is true $\forall i$

Important: in the following slides we assume the population participates in a 2-player symmetric game.

Matrix form of expected payoffs

Expected payoff of type of a type i in a population with state x

$$(A\boldsymbol{x})_i = \sum_{j=1}^{m} a_{ij}x_j$$

Average payoff in the population

$$\boldsymbol{x}^T A\boldsymbol{x} = \sum_i x_i(A\boldsymbol{x})_i = \sum_{i,j} a_{ij}x_ix_j$$

Nash equilibrium (again...)

$$G(x_i) = \dot{x}_i = x_i[(Ax)_i - x^T Ax]$$

Where the matrix A is a payoff matrix with element A_{ij} representing the fitness of strategy i over strategy j.

A symmetric $n \times n$ game has a symmetric NE $\boldsymbol{x}$ if

$$\boldsymbol{z}^T Ax \leq \boldsymbol{x}^T Ax \ \forall \boldsymbol{z} \in \Delta_n$$

Evolutionary stable state (ESS)

$$G(x_i) = \dot{x}_i = x_i[(Ax)_i - x^T Ax]$$

Where the matrix A is a payoff matrix with element A_{ij} representing the fitness of strategy i over strategy j.

Every symmetric NE is a rest point of the replicator equation, however, not every rest point of the replicator equation is an NE.

Evolutionary stable state (ESS)

$$G(x_i) = \dot{x}_i = x_i[(Ax)_i - x^T Ax]$$

Where the matrix A is a payoff matrix with element A_{ij} representing the fitness of strategy i over strategy j.

For all $\hat{\boldsymbol{x}} \neq \boldsymbol{x}$ in some neighbourhood of x (the perturbed state):

Equilibrium condition

$$\hat{\boldsymbol{x}}^T A\boldsymbol{x} \leq \boldsymbol{x}^T A\boldsymbol{x}$$

Stability condition

if $\hat{\boldsymbol{x}}^T A\boldsymbol{x} = \boldsymbol{x}^T A\boldsymbol{x}$ then $\hat{\boldsymbol{x}}^T A\hat{\boldsymbol{x}} < \boldsymbol{x}^T A\hat{\boldsymbol{x}}$

Evolutionary stable state (ESS)

$$G(x_i) = \dot{x}_i = x_i[(Ax)_i - x^T Ax]$$

Where the matrix A is a payoff matrix with element A_{ij} representing the fitness of strategy i over strategy j.

State x is **evolutionary stable** if for all $\hat{\boldsymbol{x}} \neq \boldsymbol{x}$ in some neighbourhood of x (the perturbed state), then

$$\boldsymbol{x}^T A \hat{\boldsymbol{x}} > \hat{\boldsymbol{x}}^T A \hat{\boldsymbol{x}}$$

Evolutionary stable state (ESS)

$$G(x_i) = \dot{x}_i = x_i[(Ax)_i - x^T Ax]$$

Where the matrix A is a payoff matrix with element A_{ij} representing the fitness of strategy i over strategy j.

State x is **evolutionary stable** if for all $\hat{\boldsymbol{x}} \neq \boldsymbol{x}$ in some neighbourhood of x (the perturbed state), then

$$\boldsymbol{x}^T A \hat{\boldsymbol{x}} > \hat{\boldsymbol{x}}^T A \hat{\boldsymbol{x}}$$

Evolutionary stable state (ESS)

$$G(x_i) = \dot{x}_i = x_i[(Ax)_i - x^T Ax]$$

Where the matrix A is a payoff matrix with element A_{ij} representing the fitness of strategy i over strategy j.

Evolutionary stable state extends the concept of ESS to mixed strategies through dynamic stability. That is, a population configuration (state) is stable if, after an infinitesimal change in the population (e.g., the introduction of a mutant), it converges to that state.

Payoff dominance and Risk dominance

A NE is considered **payoff dominant** if it is **Pareto superior** (all other NE provide less payoff to at least one player) to all other NE in the game.

A NE is considered **risk dominant** if it is perceived as "less risky" than all other NE. Risk perception here means that it maximises the expected payoff given the uncertainty about what the opponent(s) might do. It can also be seen as **the equilibria with the largest basin of attraction**.

Payoff dominance and Risk dominance

A strategy S **risk dominates** a strategy T if the expected payoff for a player i choosing S is bigger than the expected payoff of choosing T, that is $\Pi_i(p \mid S) > \Pi_i(p \mid T)$, where p is the is the risk factor of the pure NE (S, S), that is, the probability that an opponent will choose strategy S.

Social dilemma space

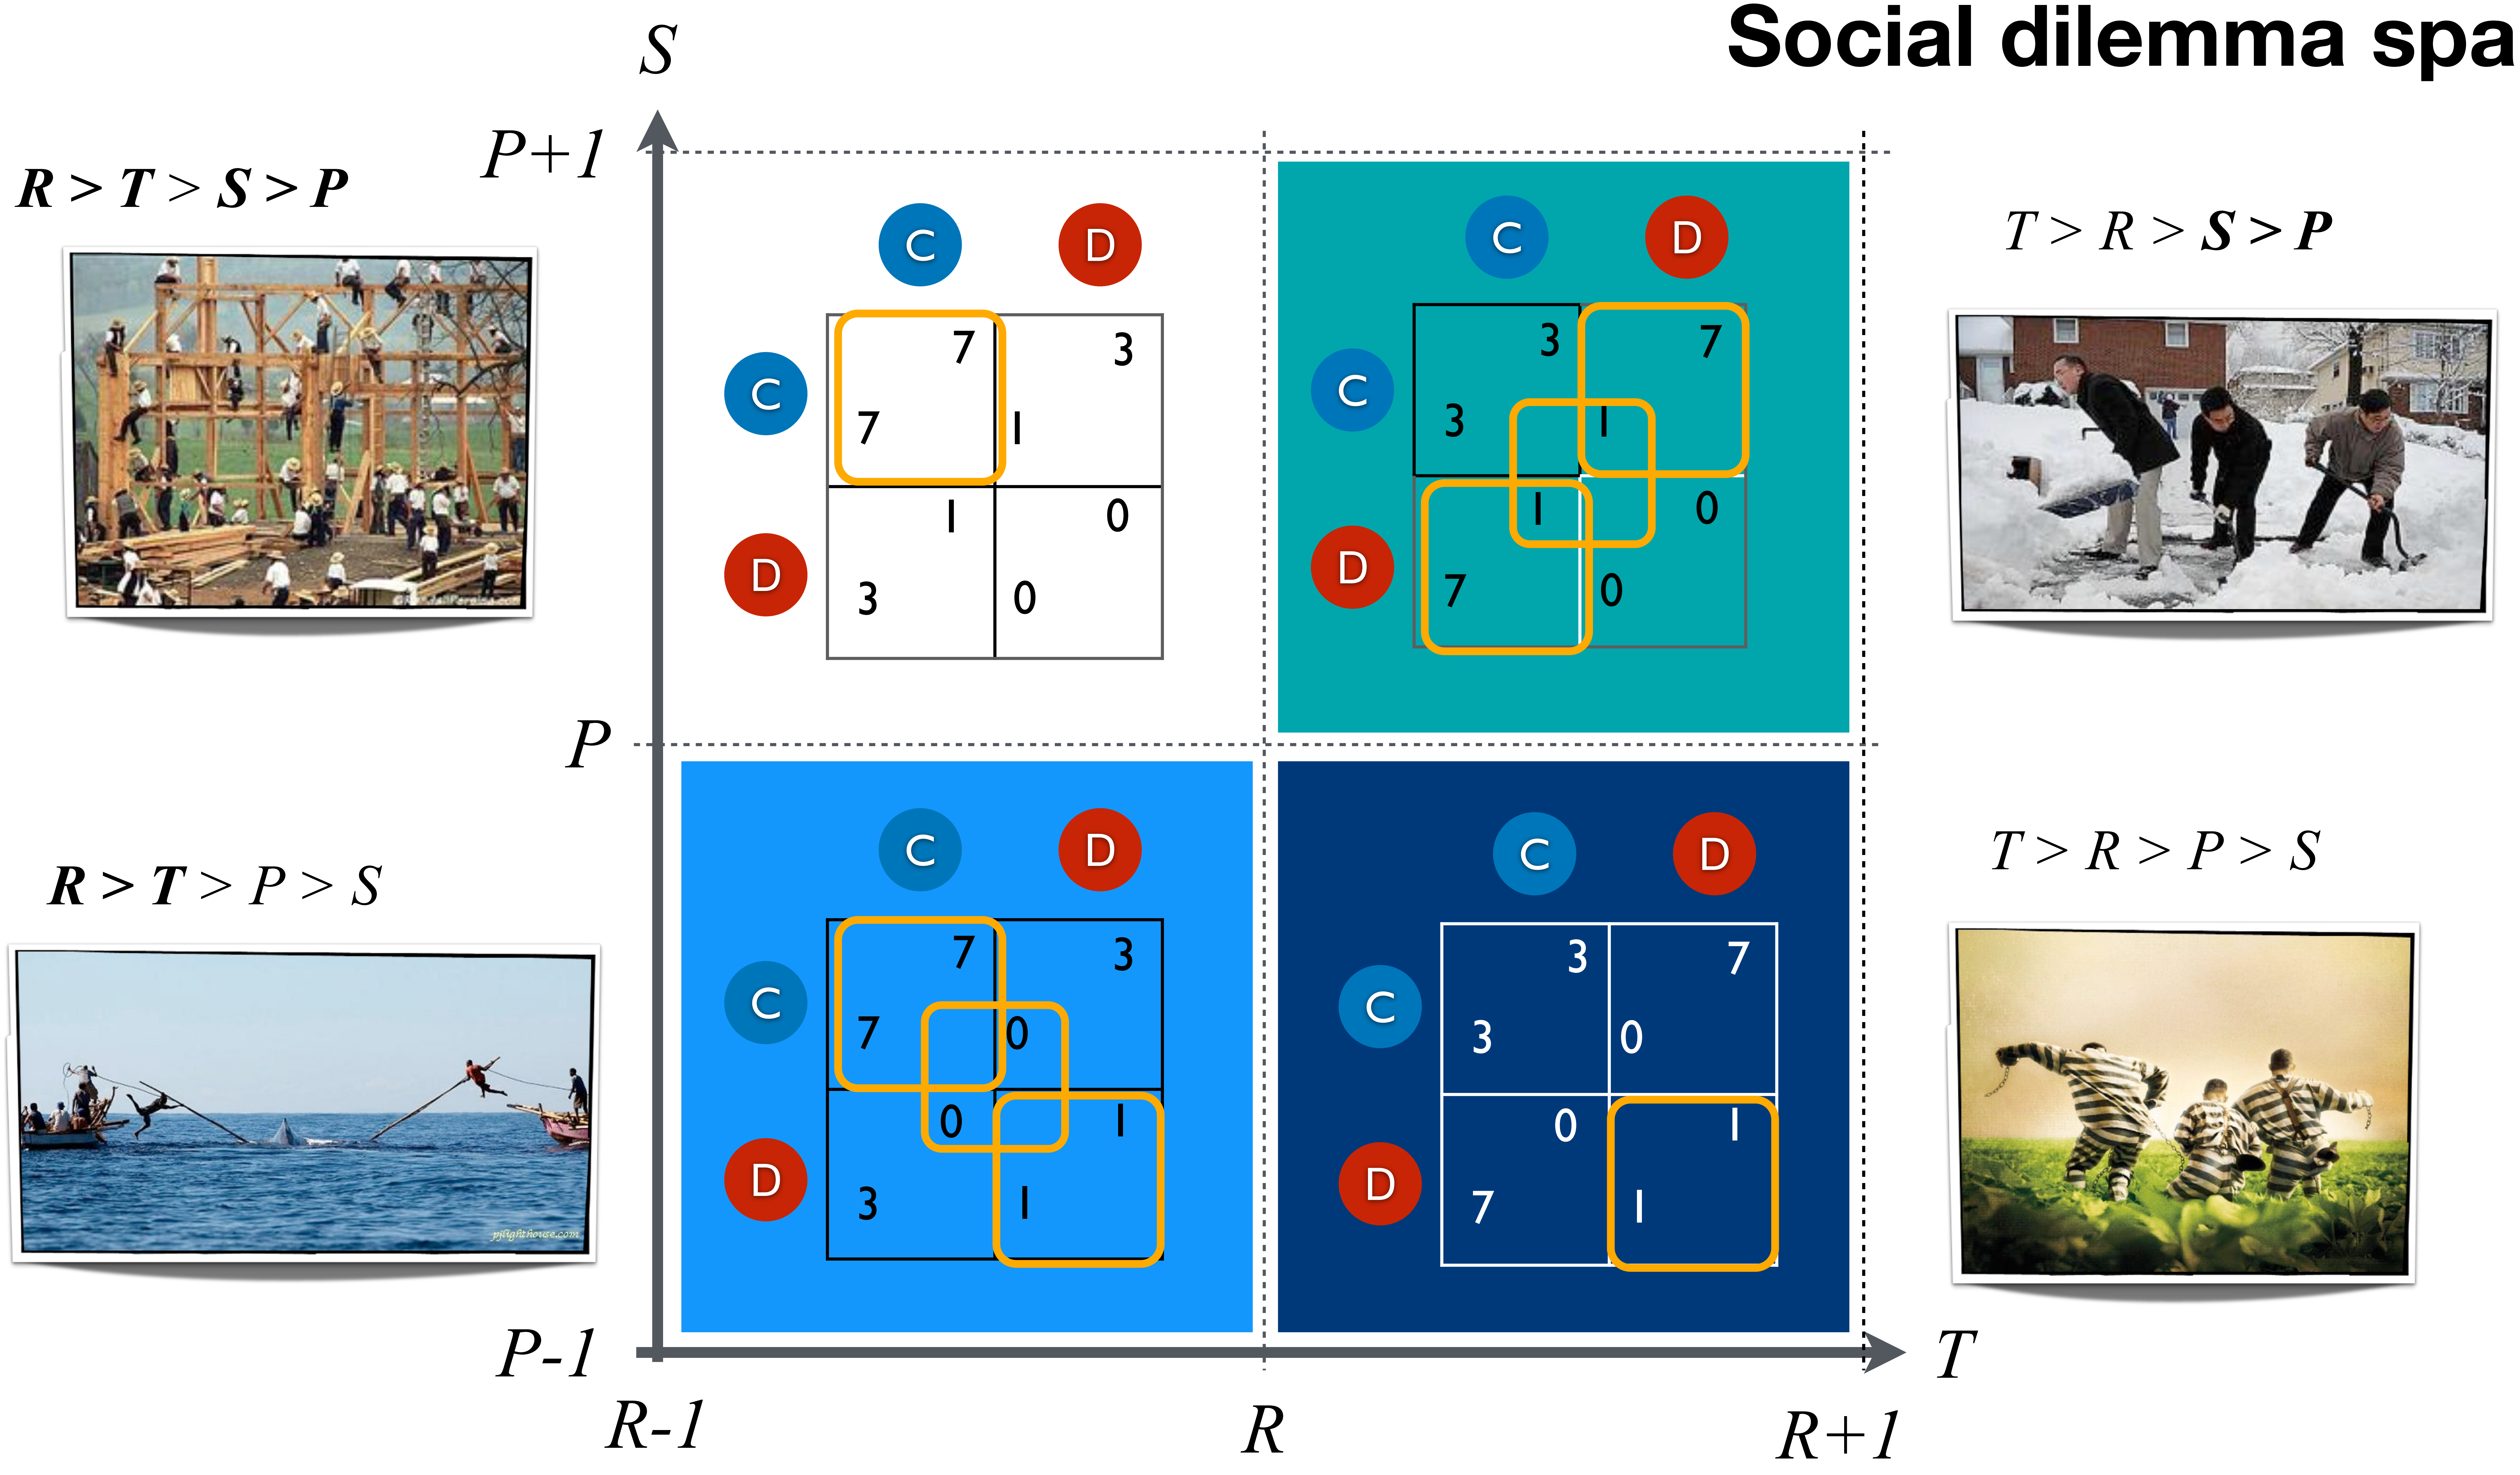

Replicator equation results for all social dilemmas

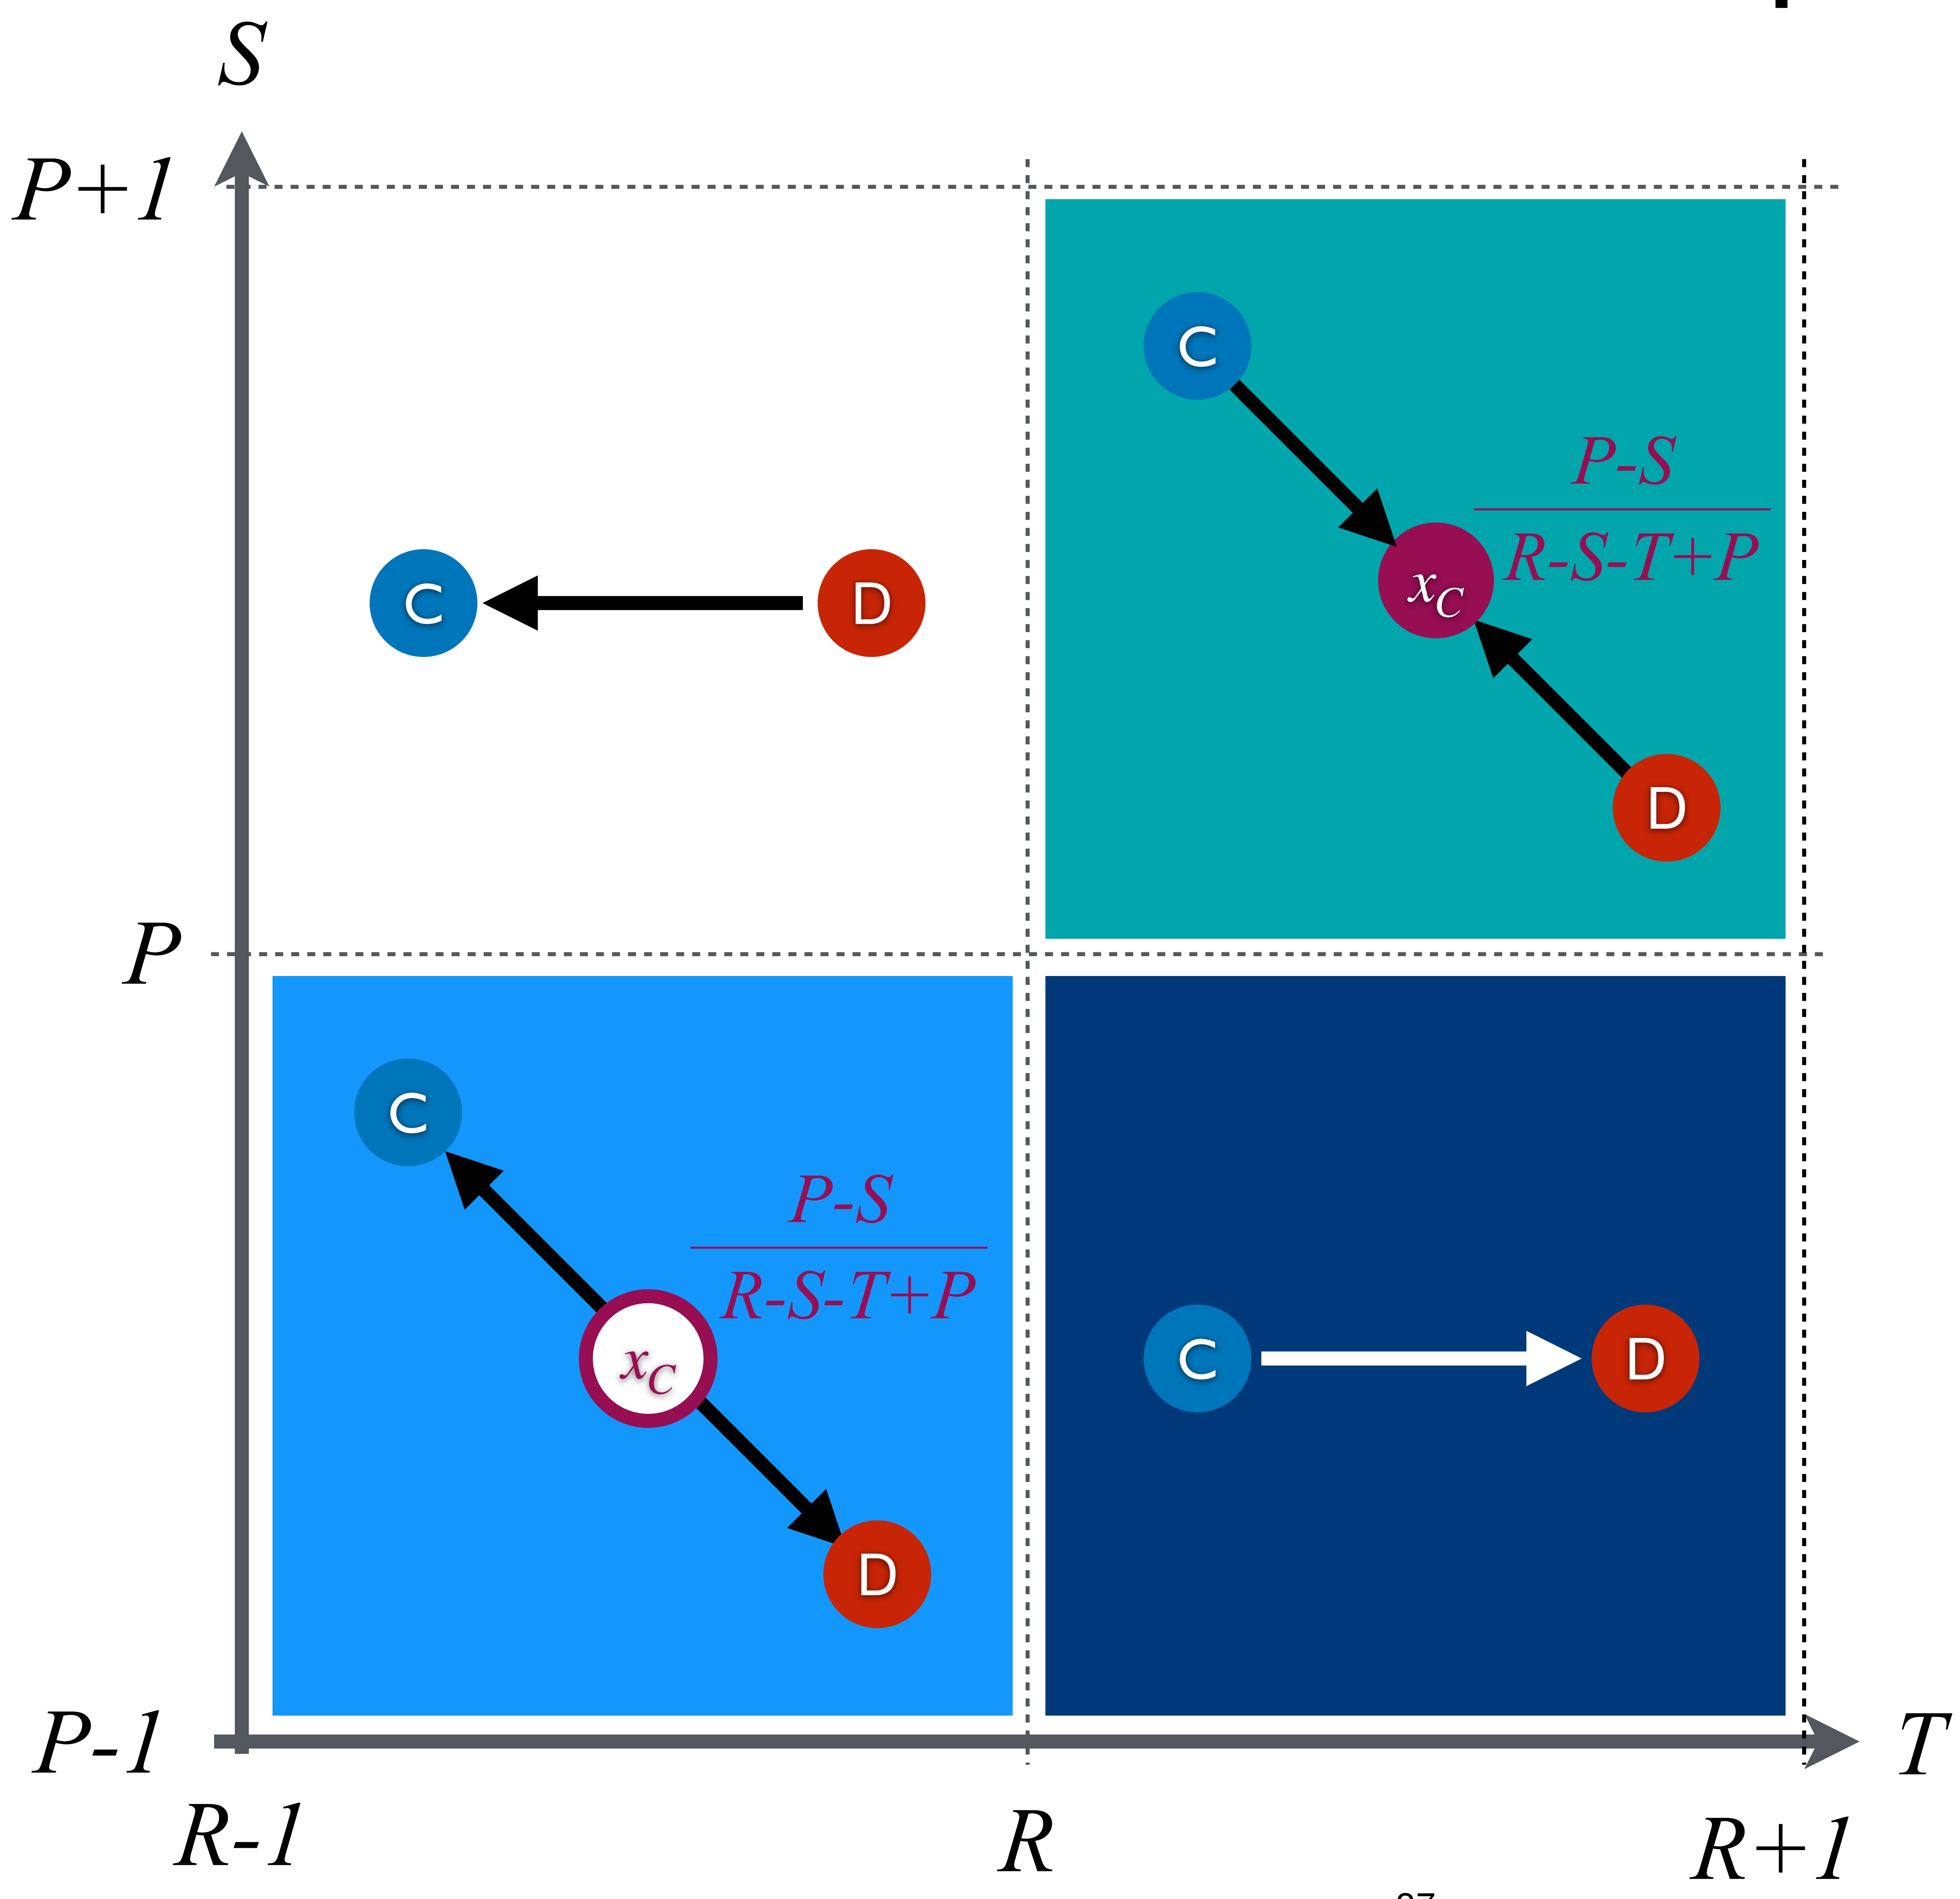

Example: The Hawk-Dove game

The Hawk-Dove game

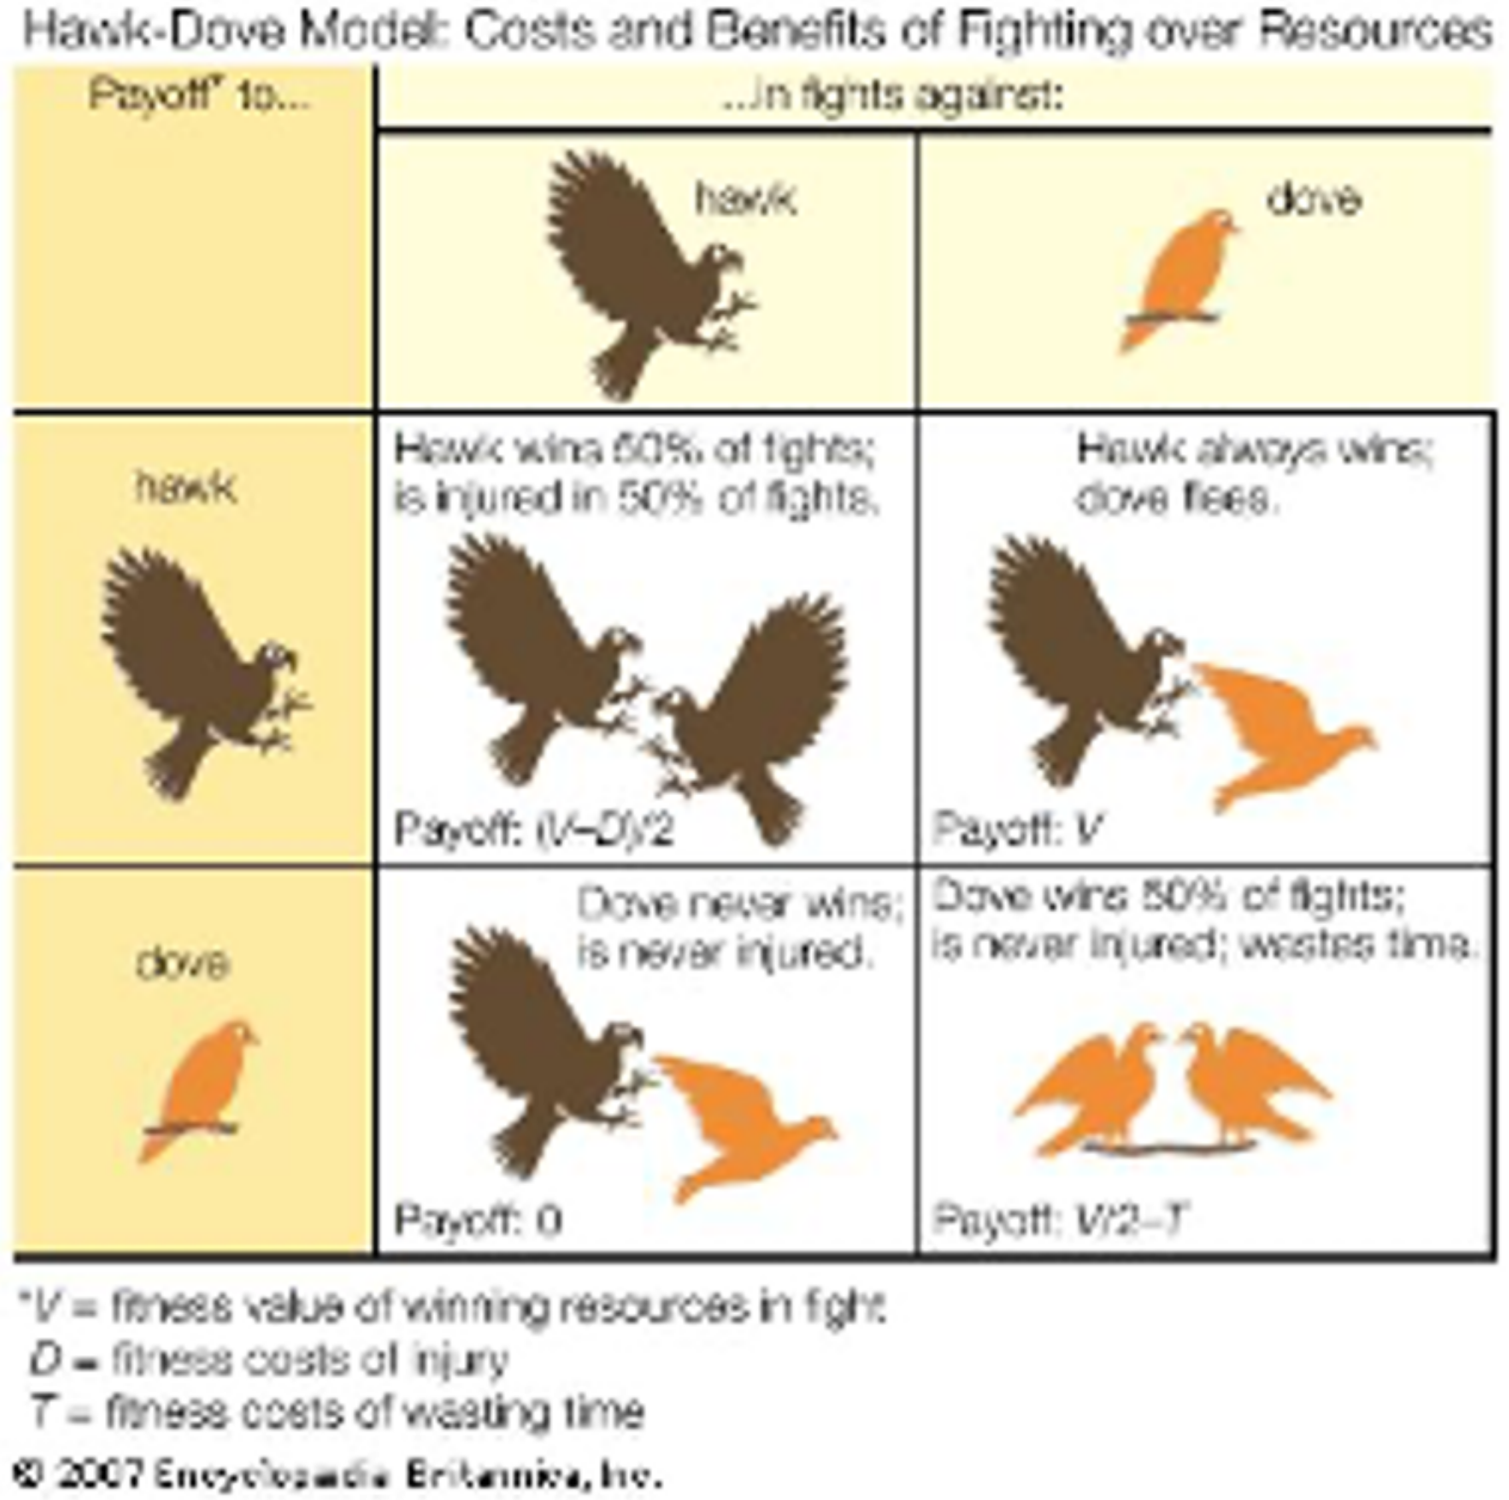

The hawk-Dove game is a coordination game formulated by John Maynard Smith and Georg Price. The aim of the game was to understand the resolution of conflicts by fighting in the animal kingdom. The game consists of two players, which each have the choice between two possible actions; either they take time to display (dove) before fighting or they can escalate immediately and fight (hawk). When both players escalate (hawk), they have a 50% risk of being injured (-D/2) and 50% of wining (V/2). When a dove fights another dove she also wins 50% of the time (V/2) but only after a period of mutual displays to show of strength (-T). Hawks always win against doves, resulting in a benefit for one (V) and not for the other (0).

Example: The Hawk-Dove game - Defection -dominant

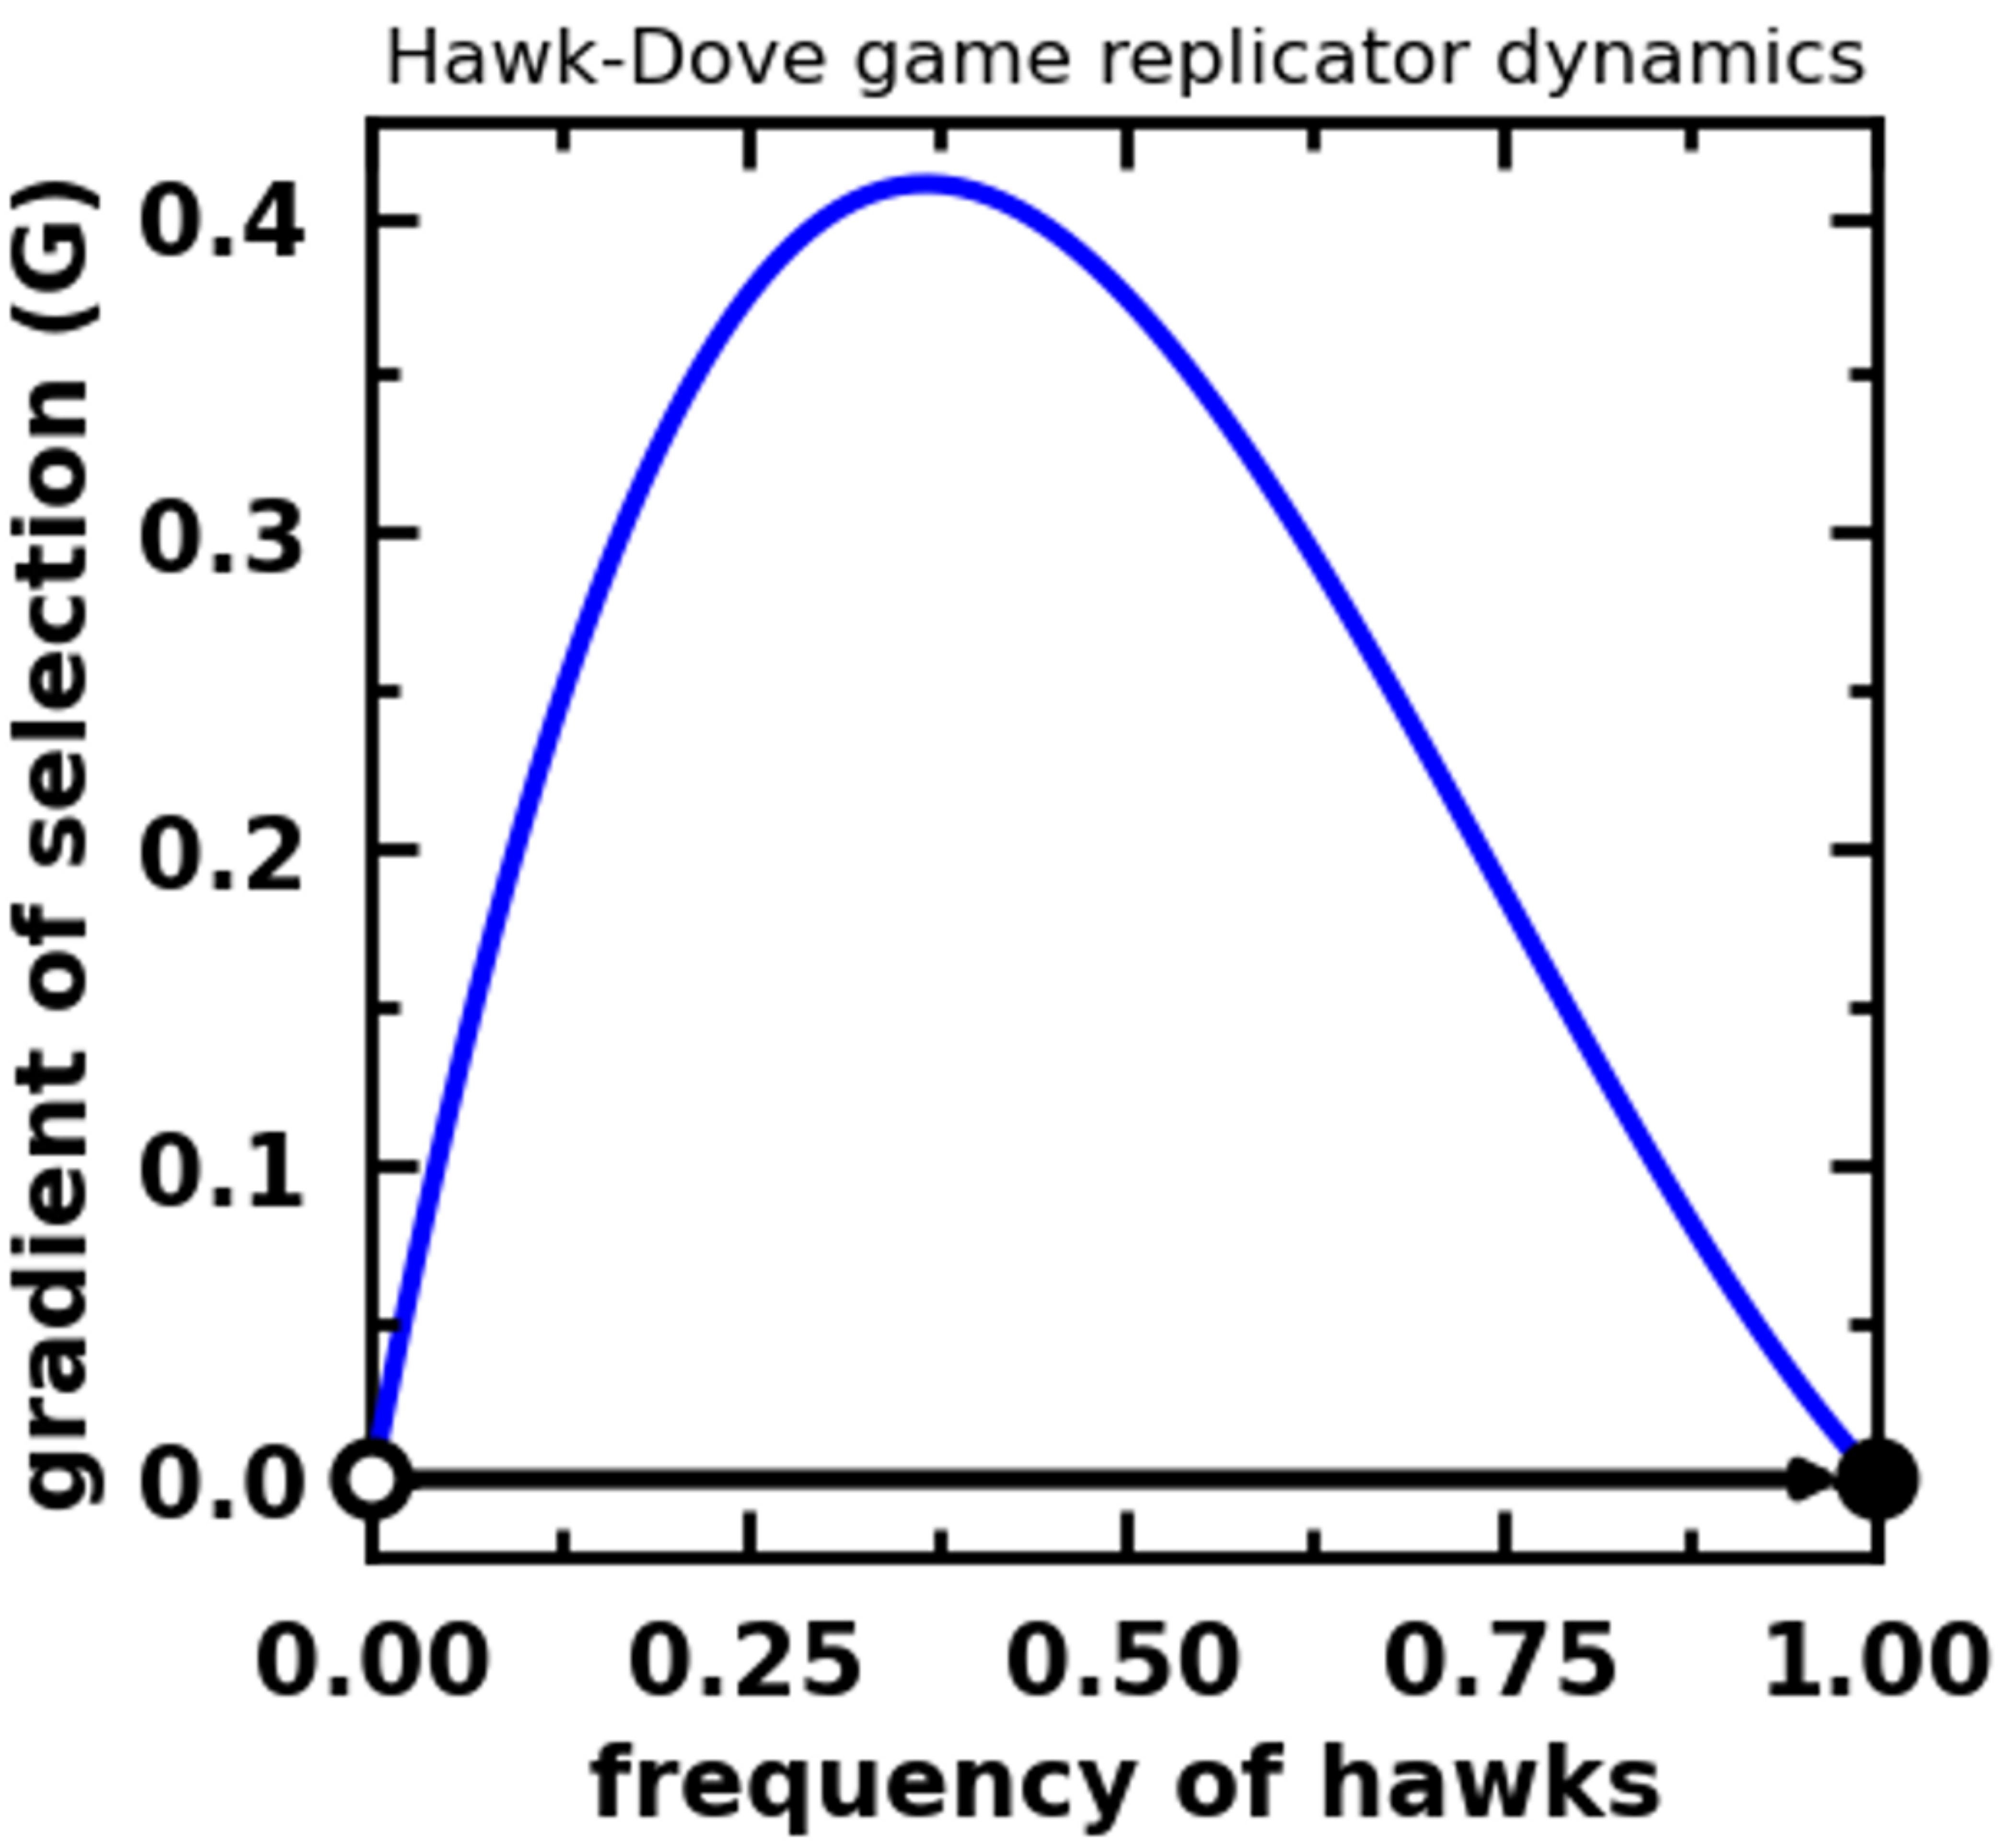

$$V = 3,\ D = 2,\ T = 1\ (V > D)$$

$$\begin{pmatrix} 0.5,\ 0.5 & 3,\ 0 \\ 0,\ 3 & 0.5,\ 0.5 \end{pmatrix}$$

Example: The Hawk-Dove game - Anti-coordination

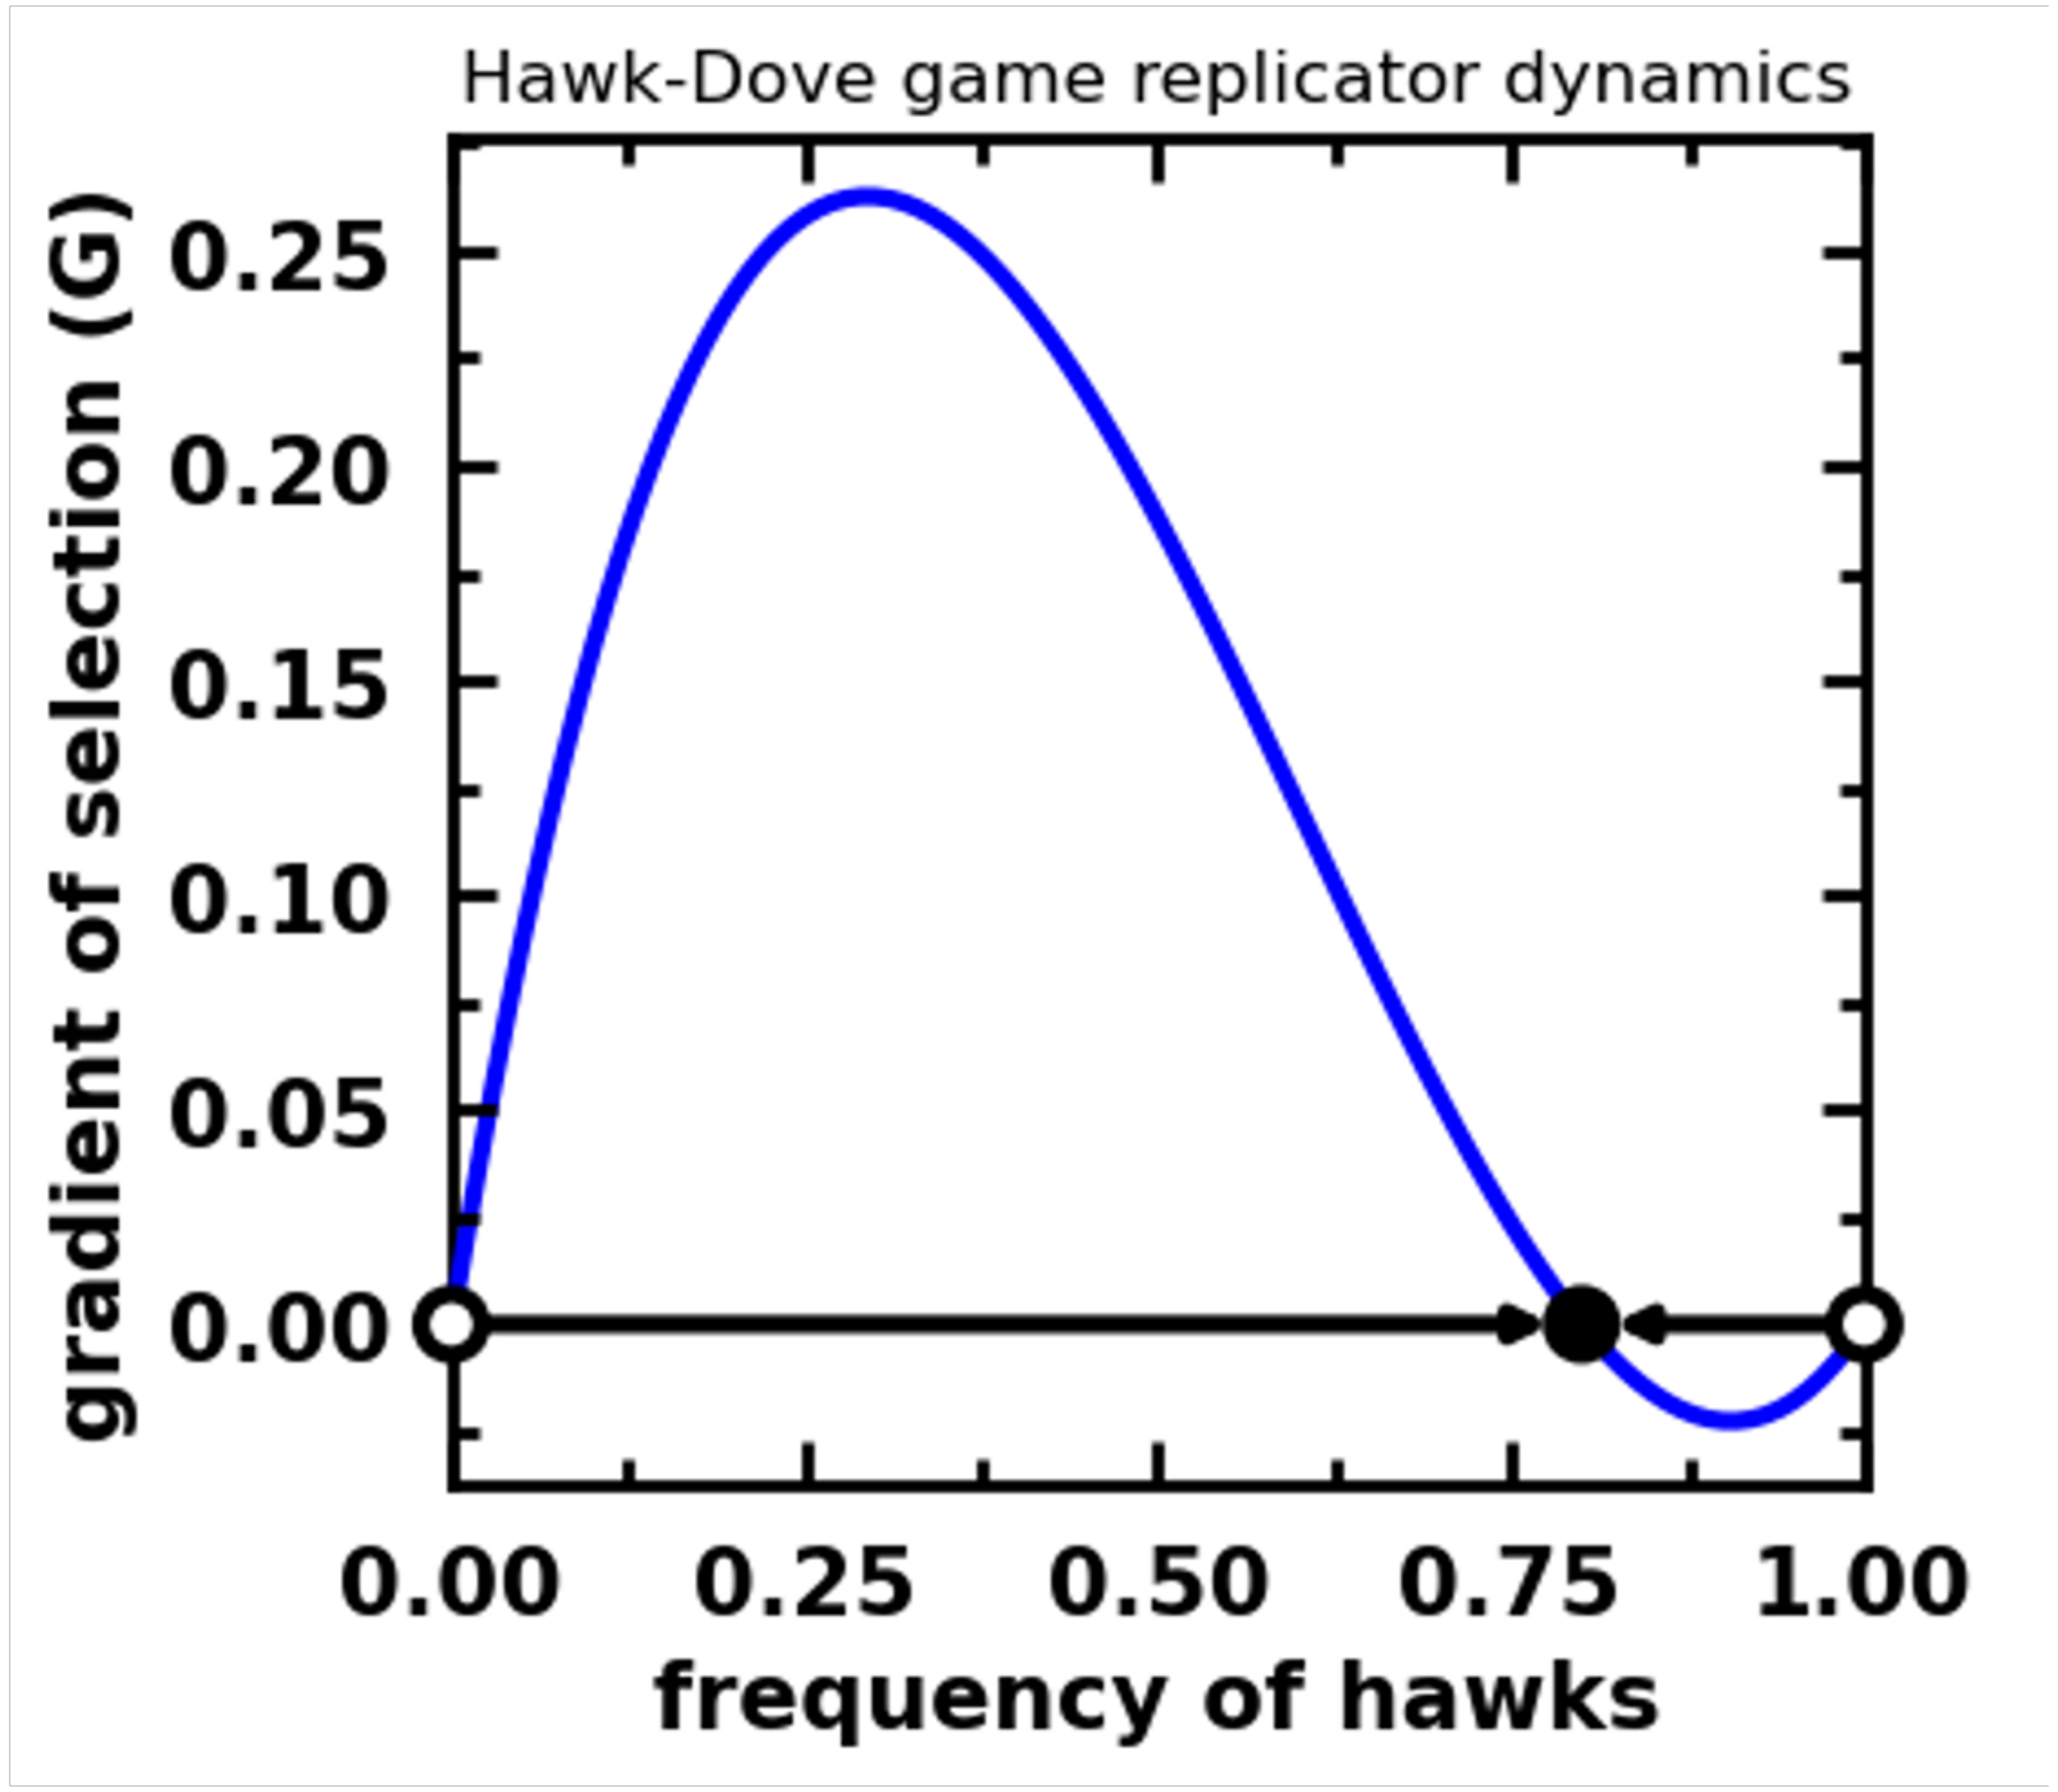

$$V = 2,\ D = 3,\ T = 1\ (V < D)$$

$$\begin{pmatrix} -0.5,\ -0.5 & 2,\ 0 \\ 0,\ 2 & 0,\ 0 \end{pmatrix}$$

Example: The Hawk-Dove game - Coordination

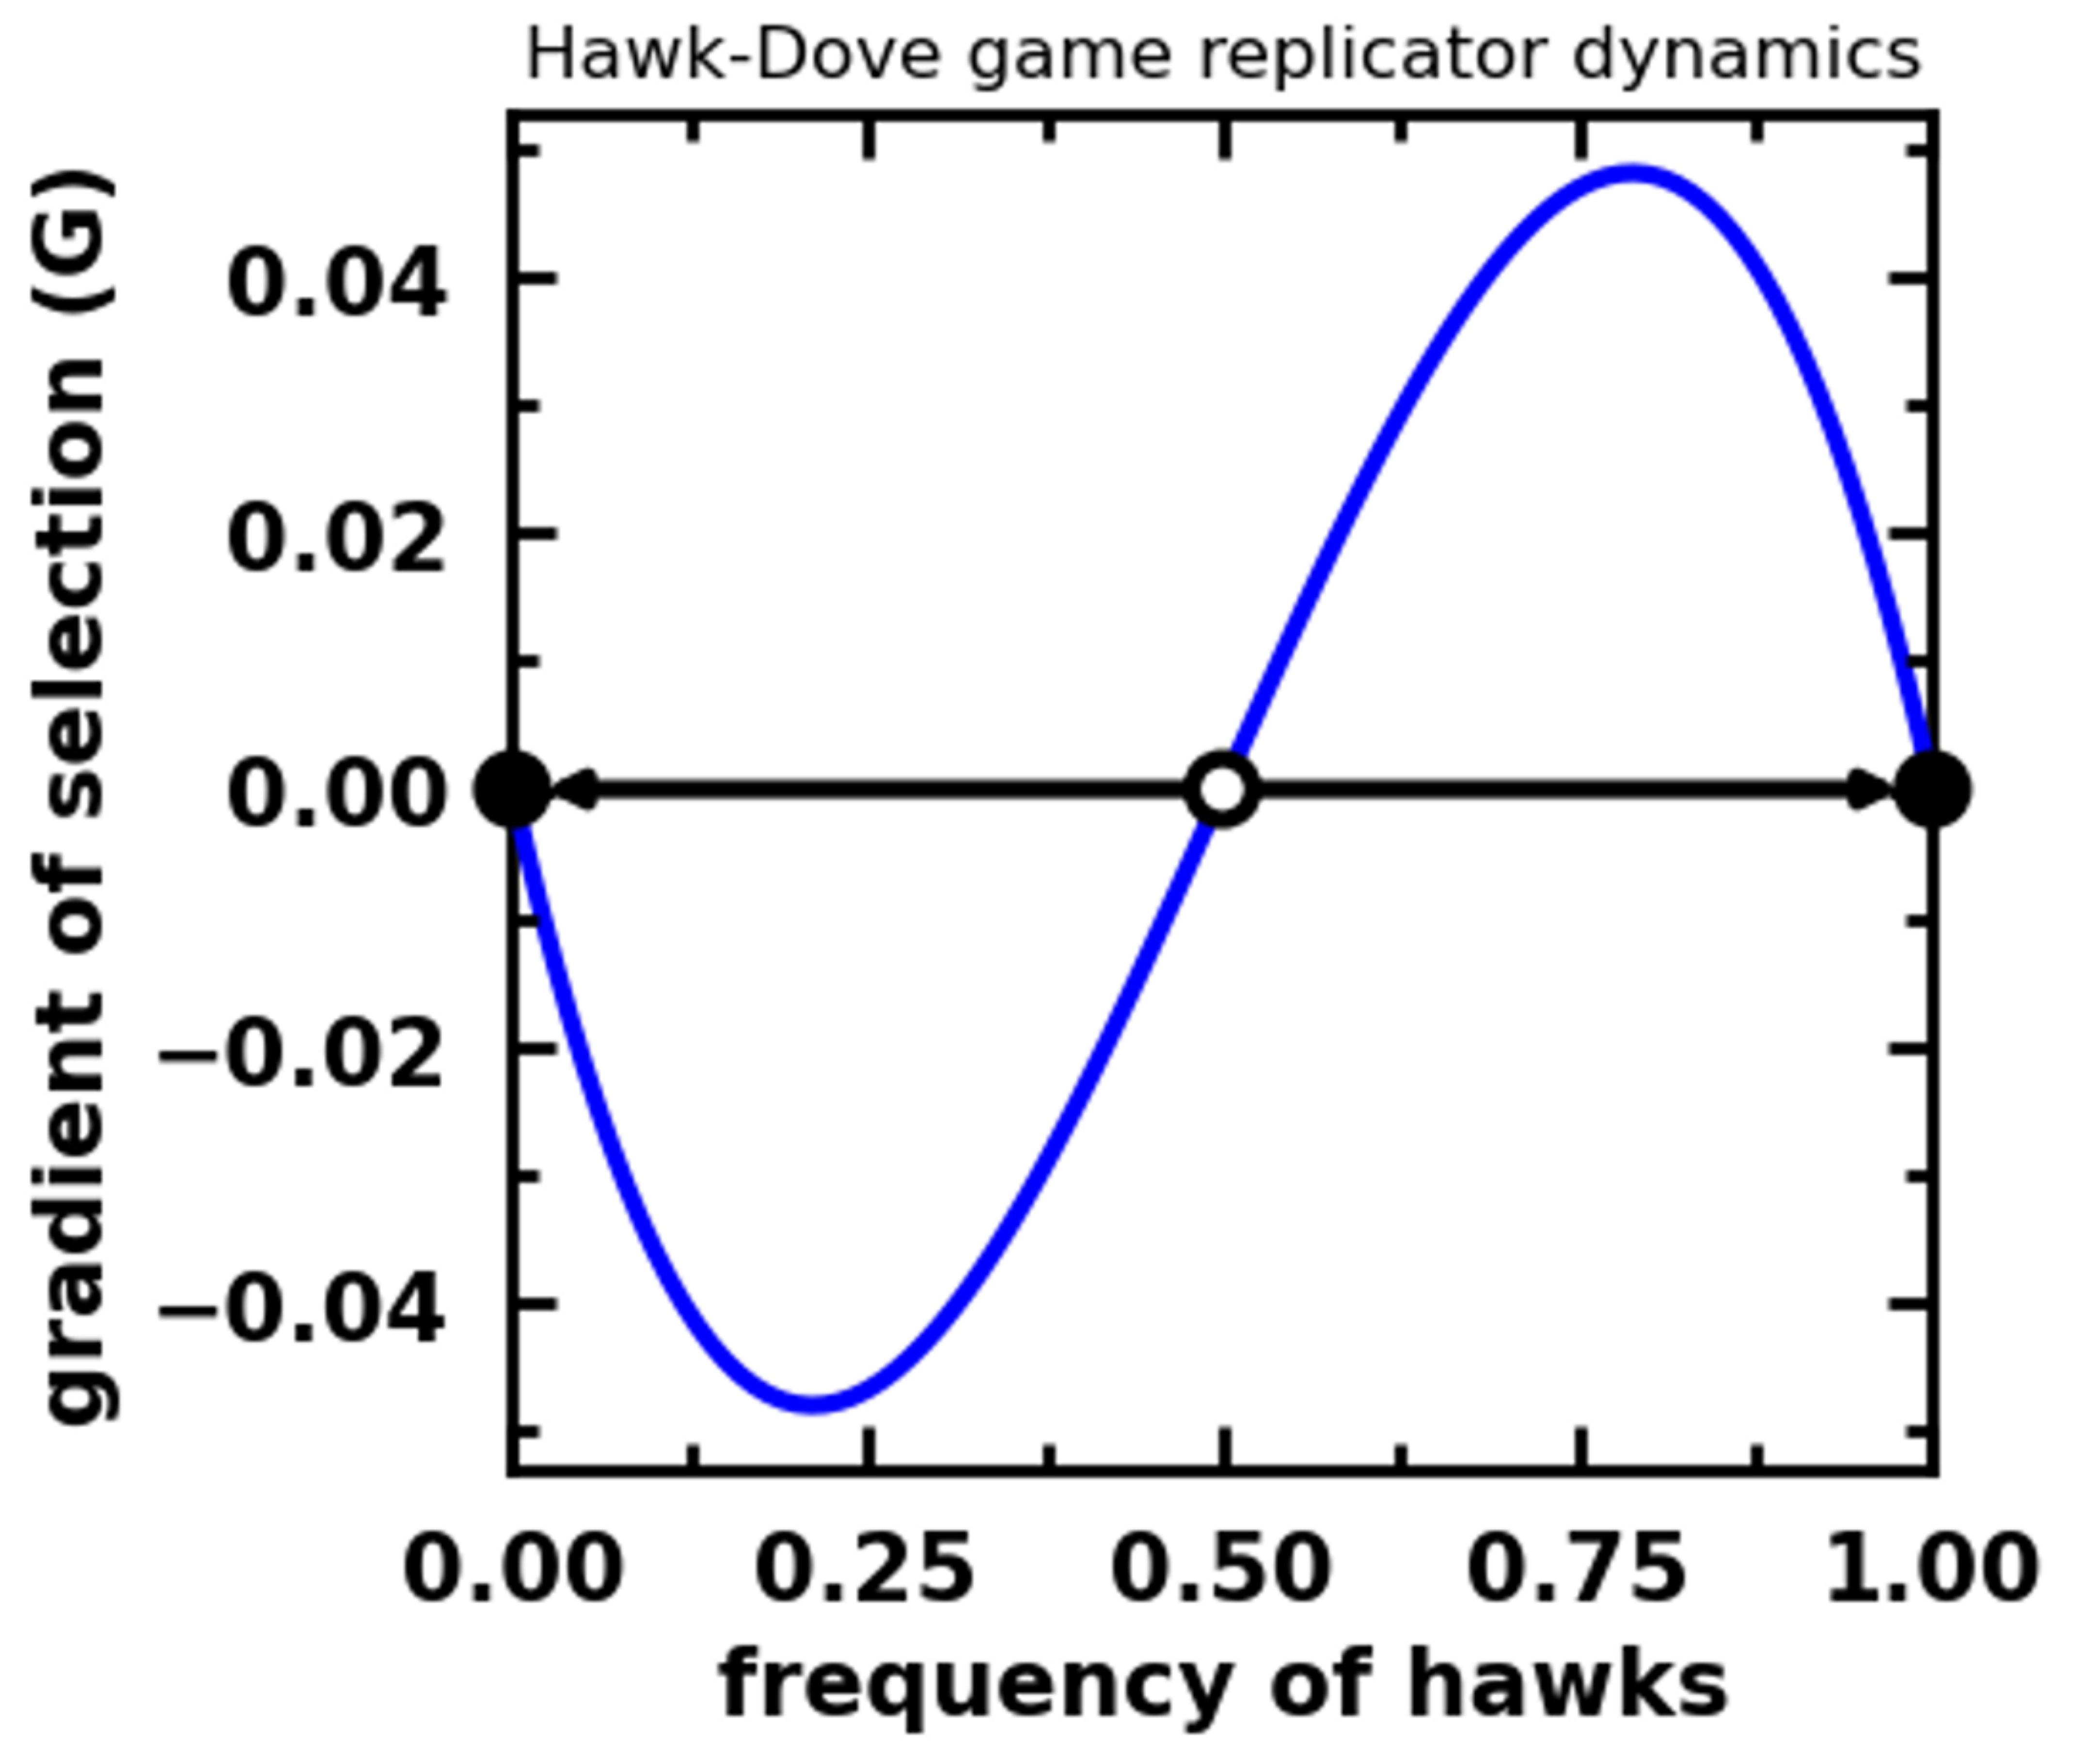

$$V = 3,\ D = 2,\ T = -2\ (V > D)$$

$$\begin{pmatrix} 0.5,\ 0.5 & 3,\ 0 \\ 0,\ 3 & 3.5,\ 3.5 \end{pmatrix}$$

Example: The Hawk-Dove game - Cooperation-dominant

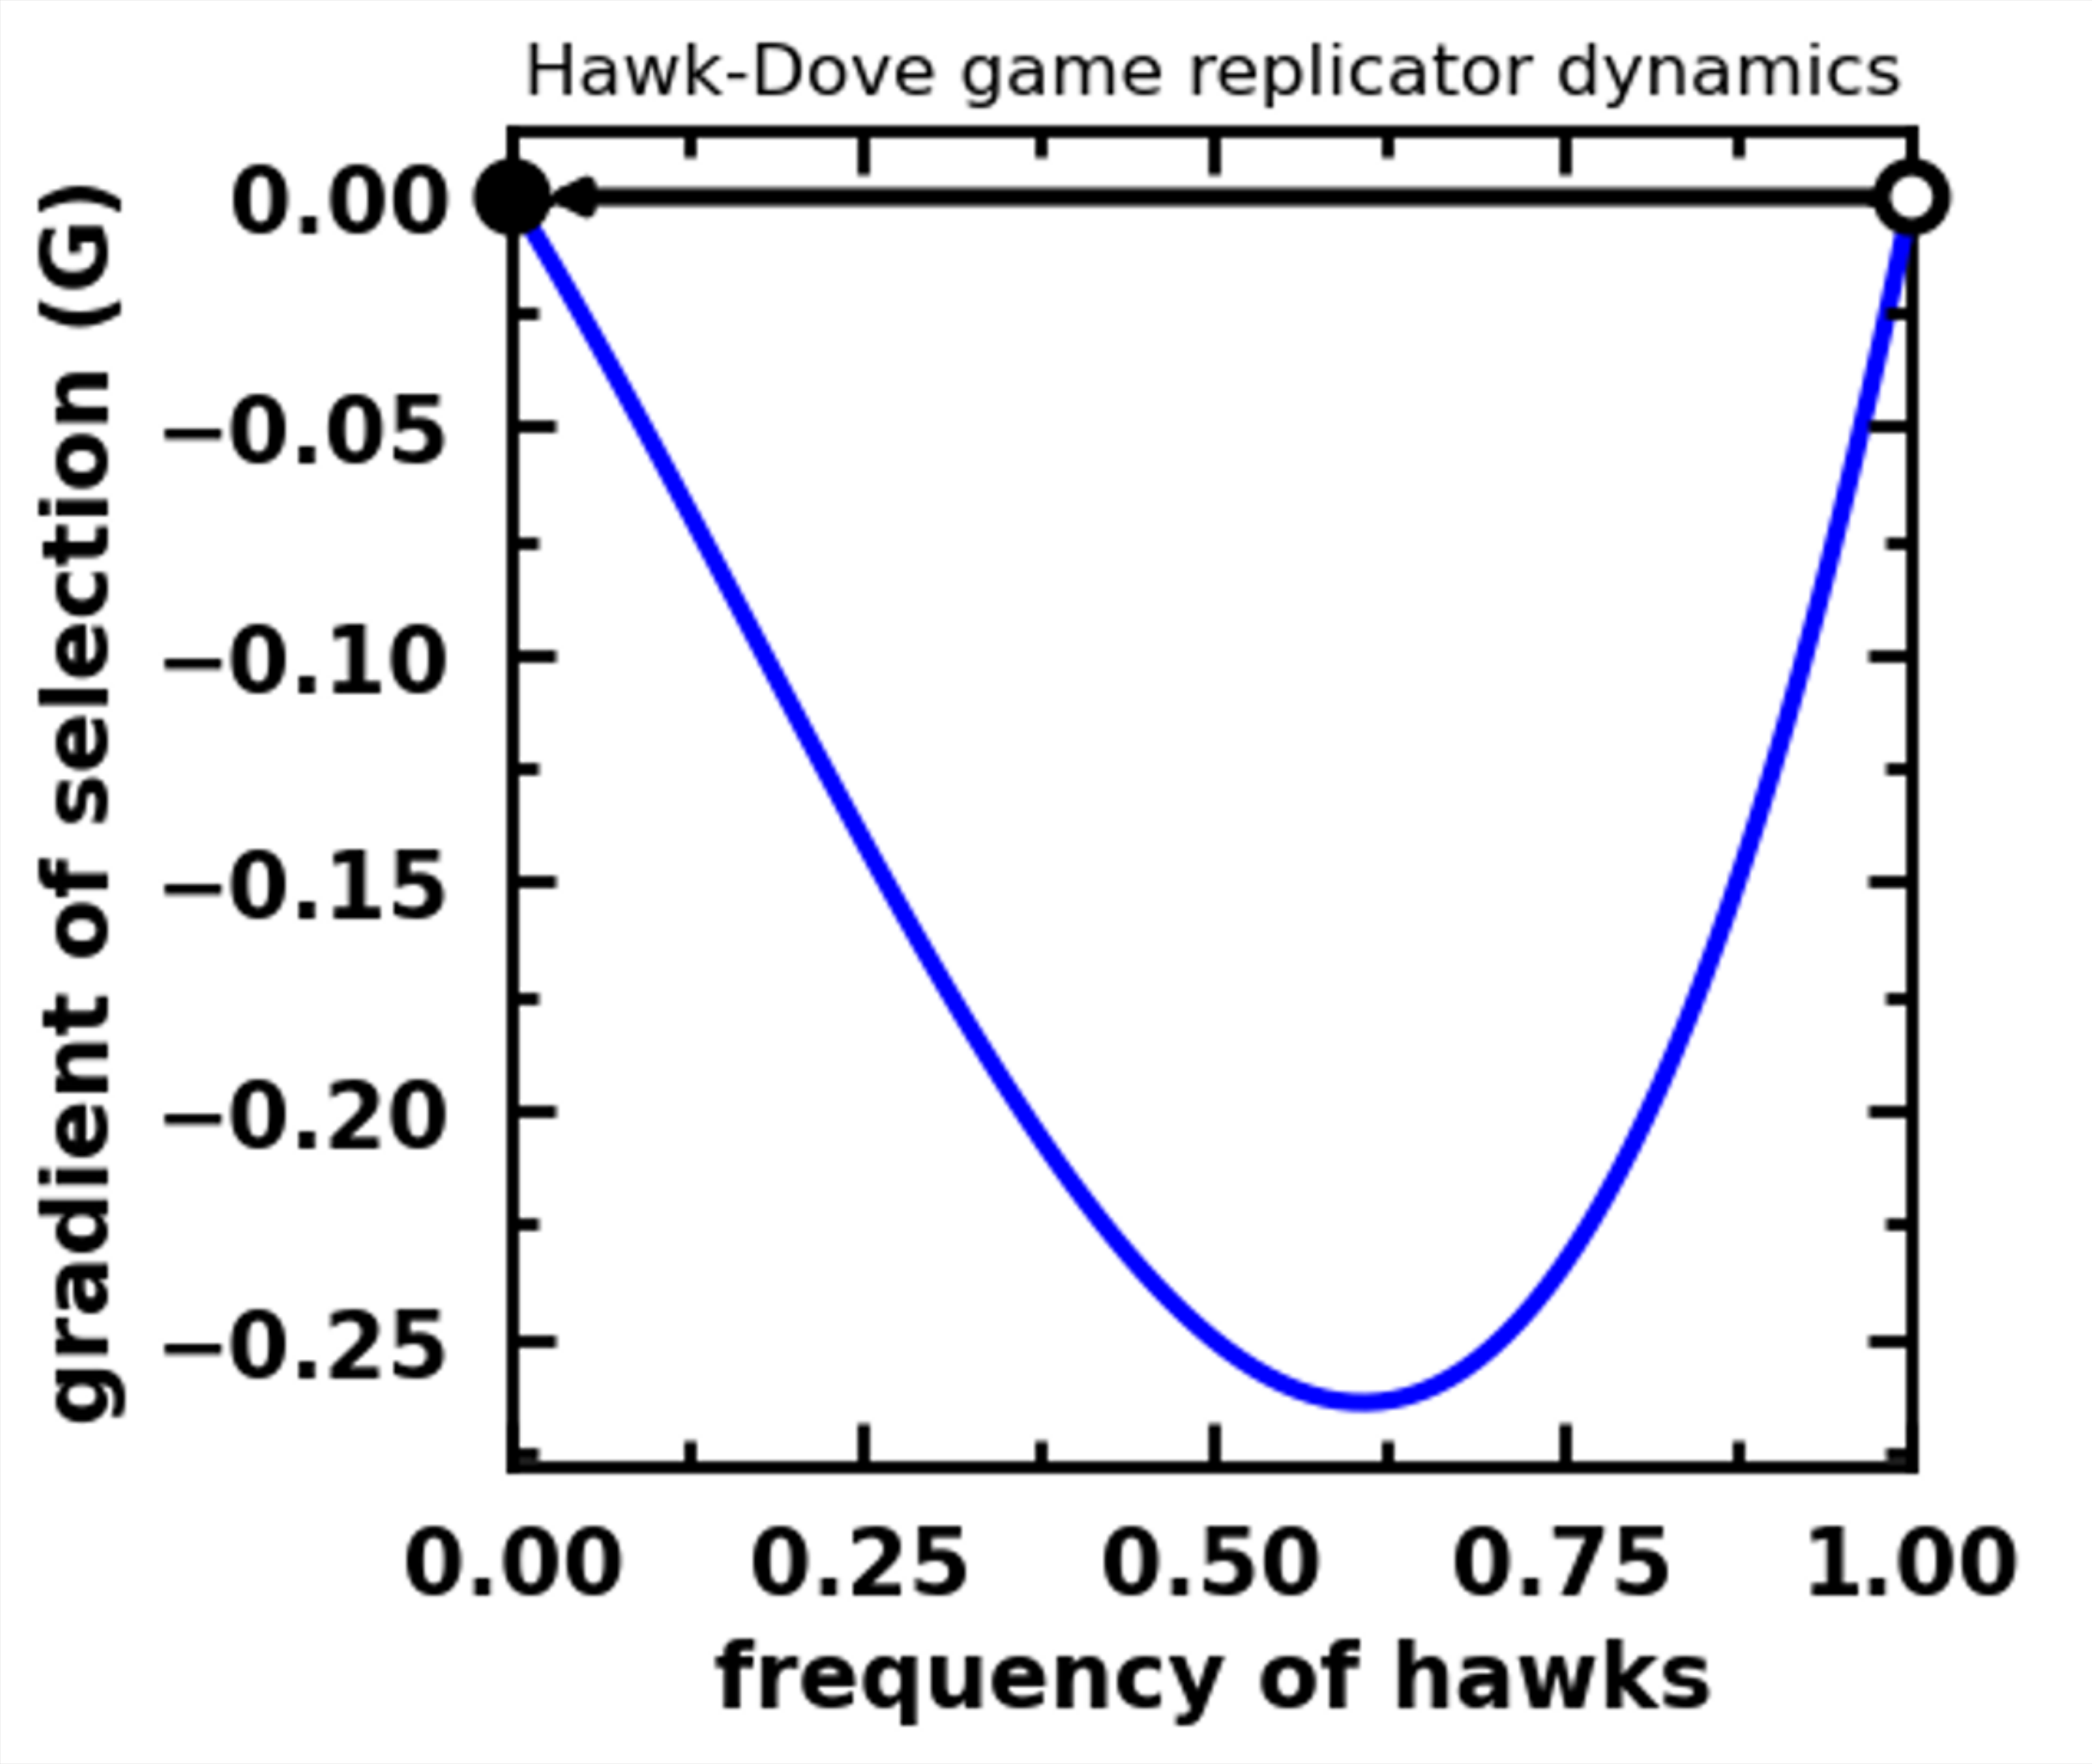

$$V = -1,\ D = 2,\ T = 0\ (V > D)$$

$$\begin{pmatrix} -1.5,\ -1.5 & -1,\ 0 \\ 0,\ -1 & -0.5,\ -0.5 \end{pmatrix}$$

Rock-Paper-Scissors

$$\begin{pmatrix} 0.5,\ 0.5 & 0,\ 1 & 1,\ 0 \\ 1,\ 0 & 0.5,\ 0.5 & 0,\ 1 \\ 0,\ 1 & 1,\ 0 & 0.5,\ 0.5 \end{pmatrix}$$

Rock-Paper-Scissors

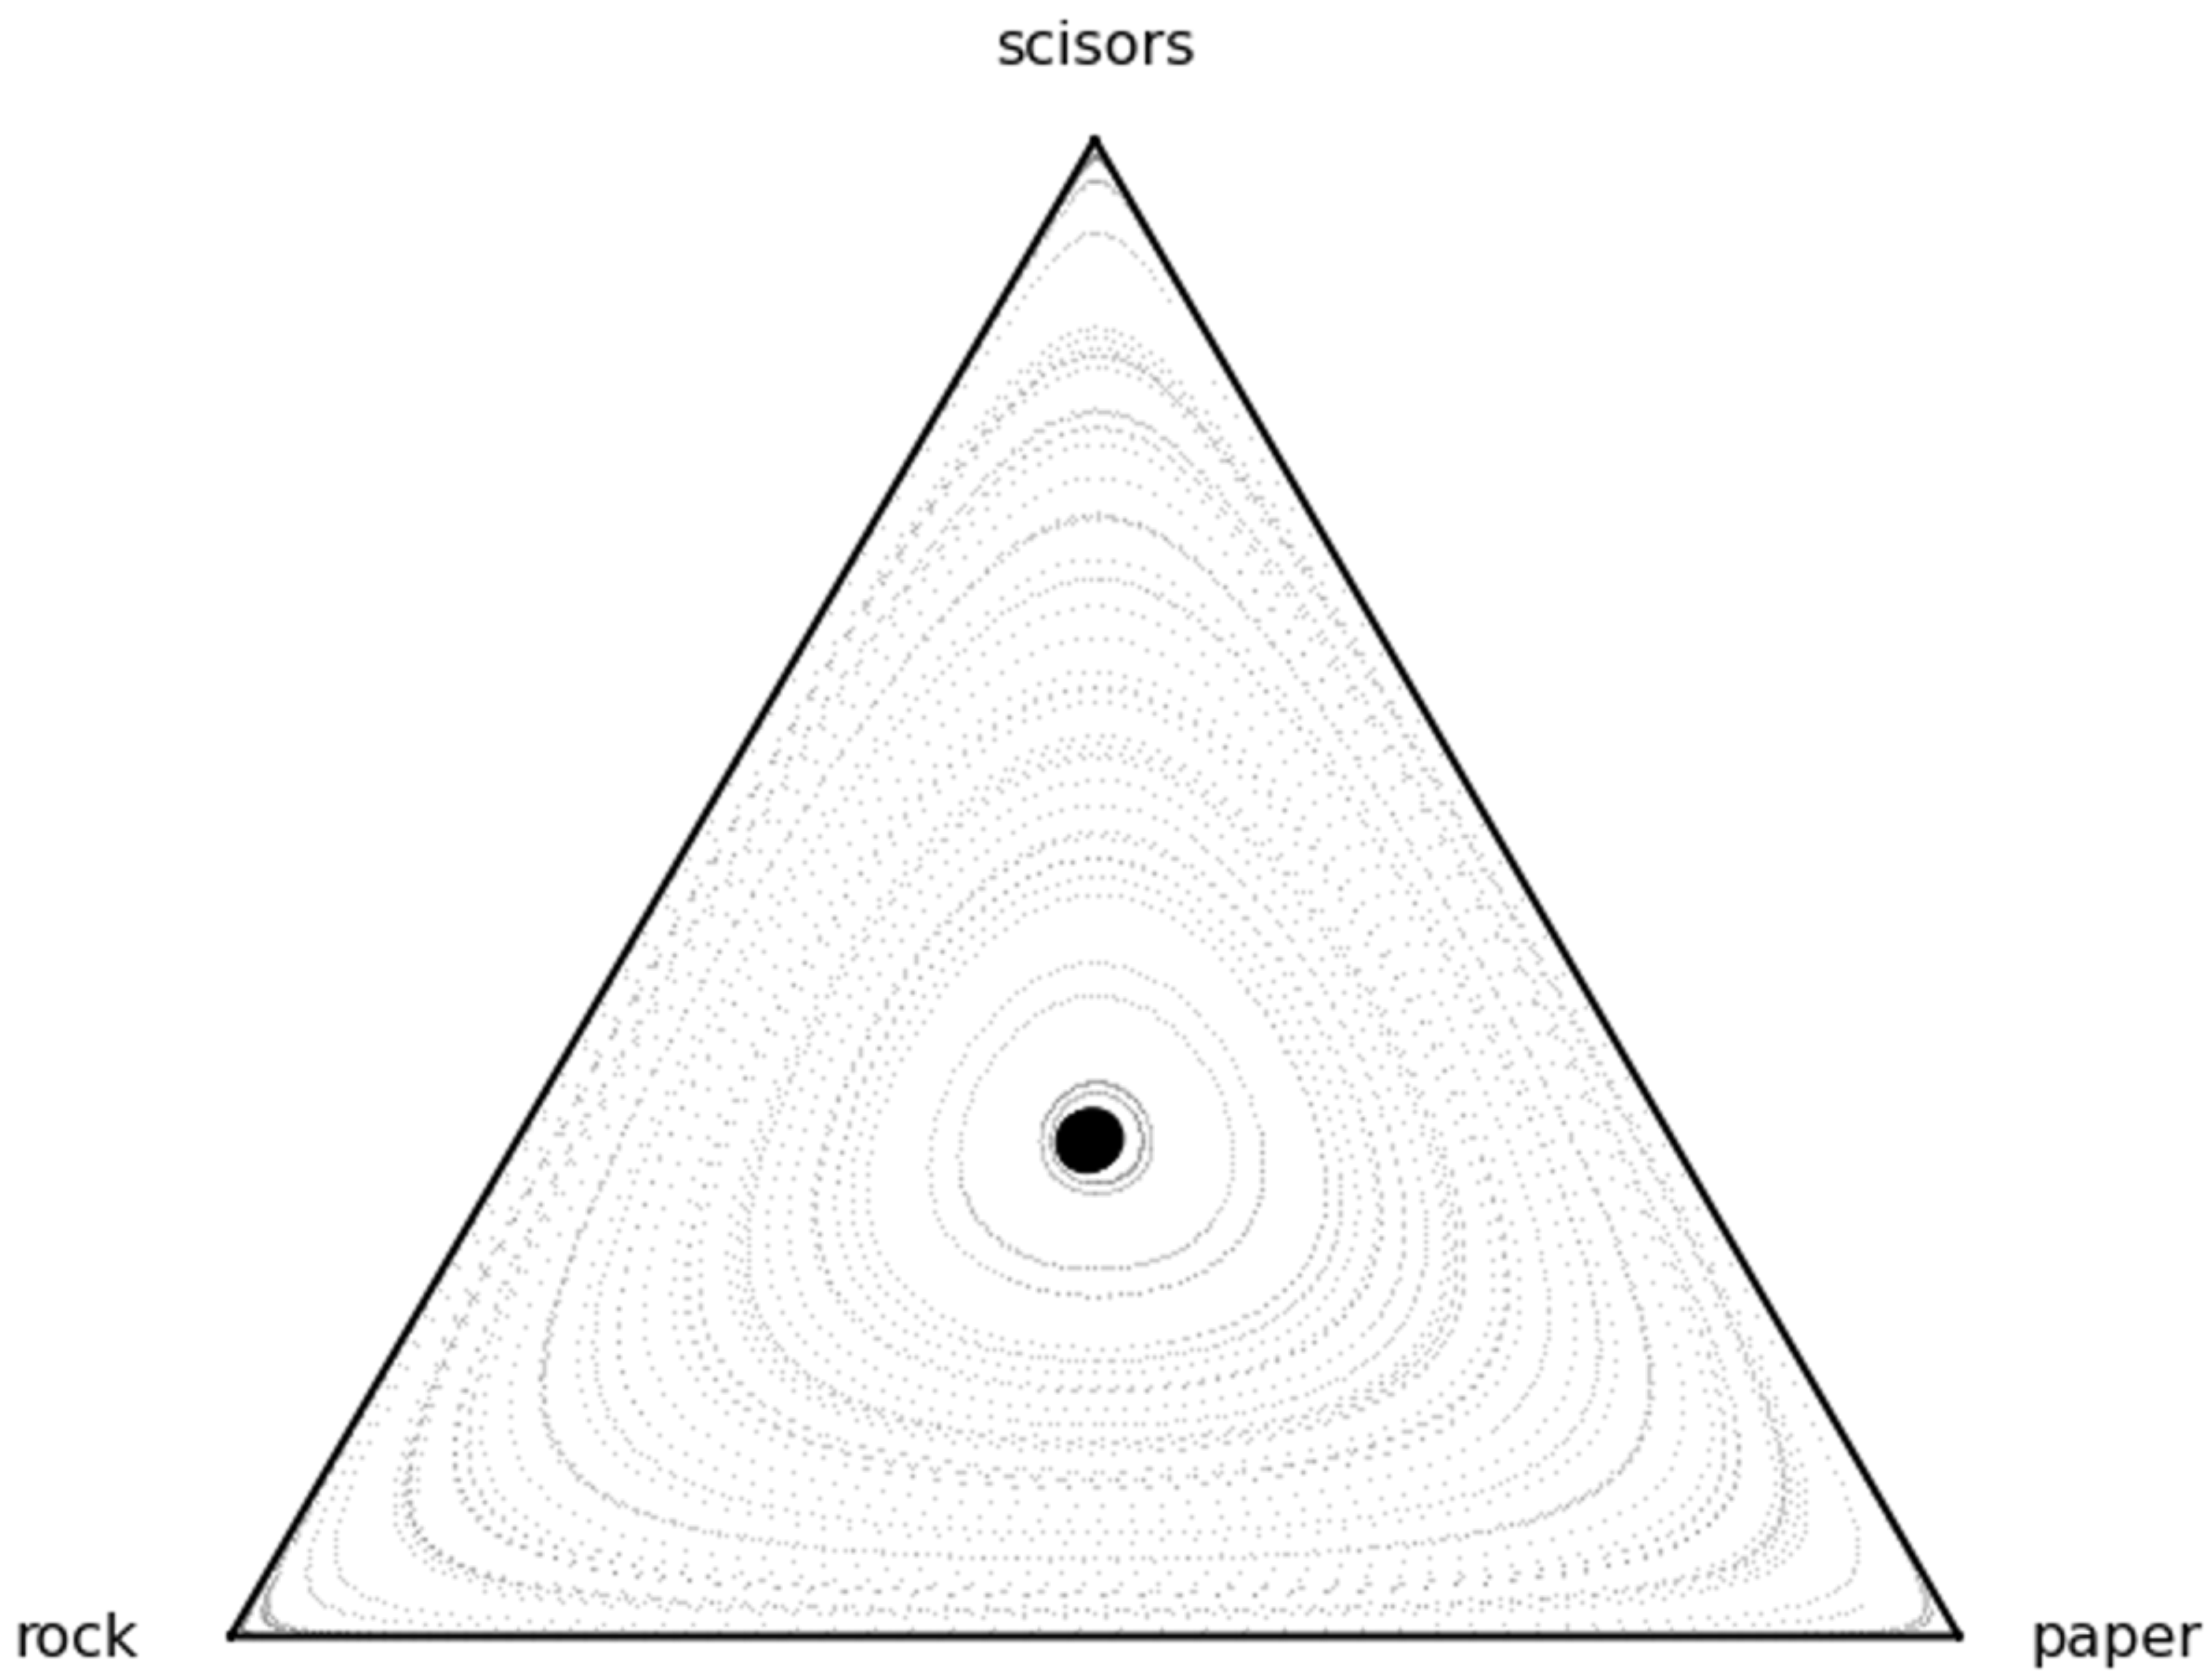

$$\begin{pmatrix} 0.5,\ 0.5 & 0,\ 1 & 1,\ 0 \\ 1,\ 0 & 0.5,\ 0.5 & 0,\ 1 \\ 0,\ 1 & 1,\ 0 & 0.5,\ 0.5 \end{pmatrix}$$

Hawk-dove-human Game

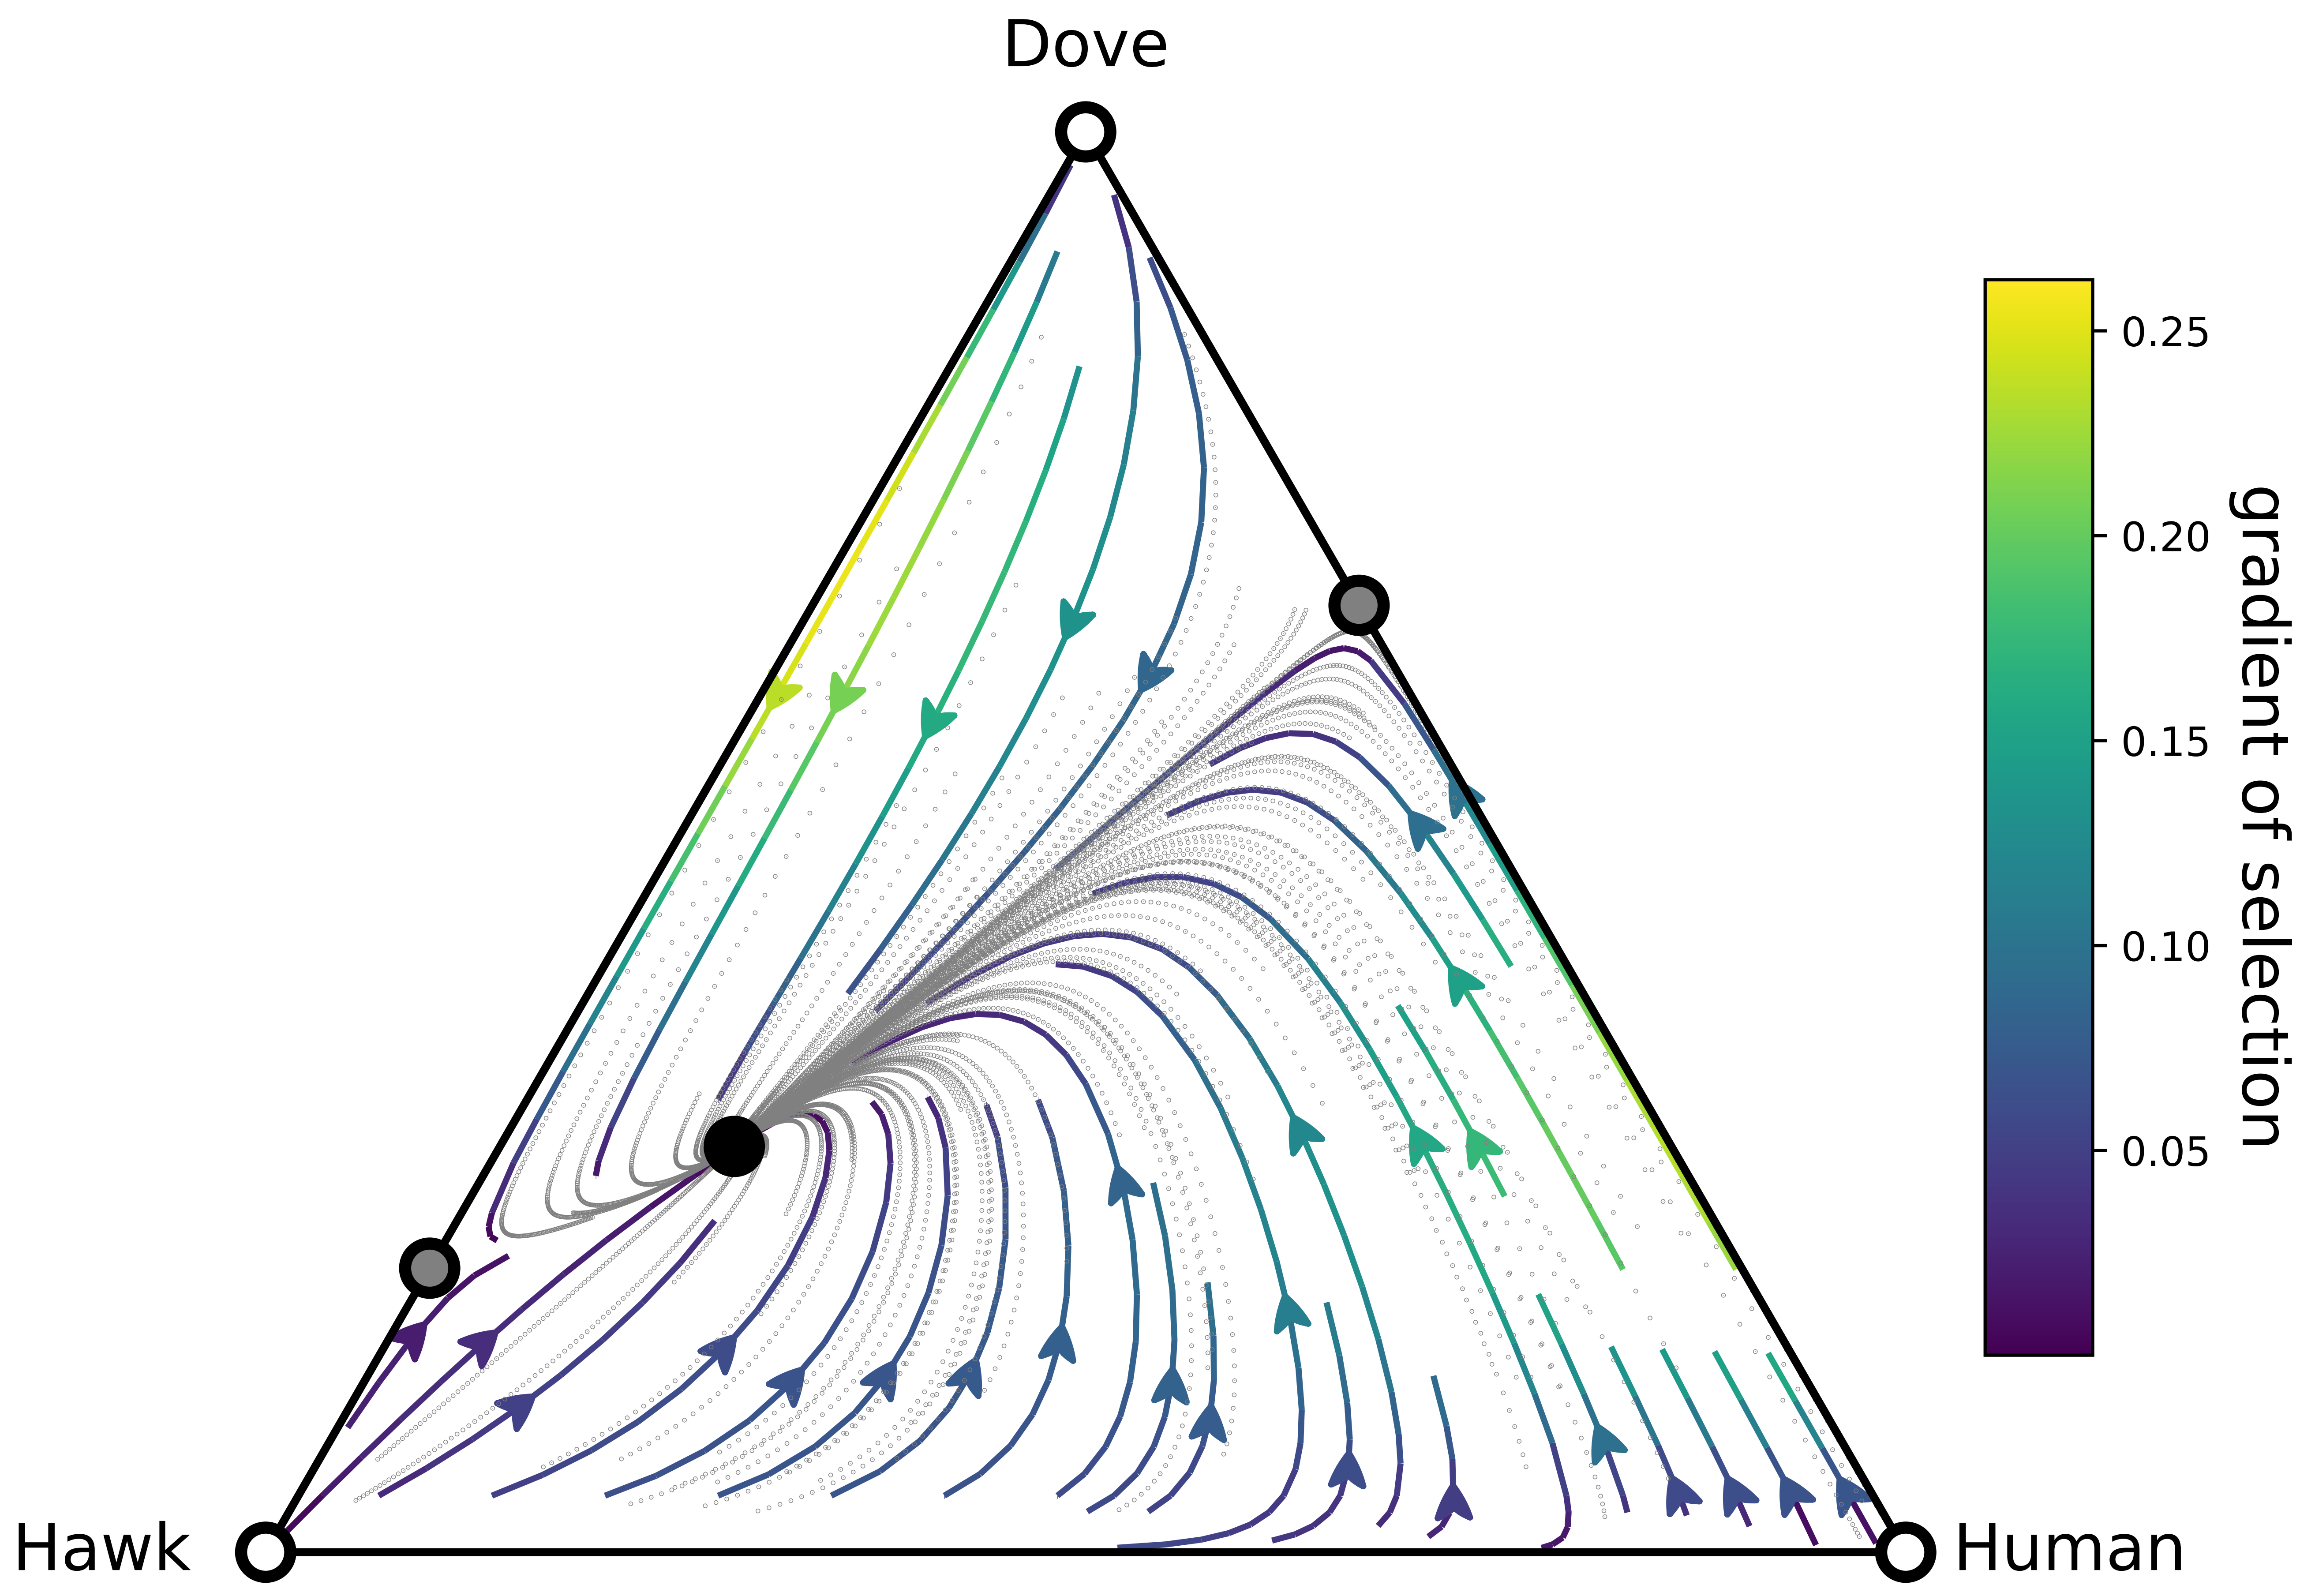

$$\begin{pmatrix} -0.5 & 2 & 0 \\ 0 & 0 & 2 \\ 0 & 1 & 0 \end{pmatrix}$$

Other dynamics

Lotka-Volterra (Predator-prey)

Replicator-mutator equation

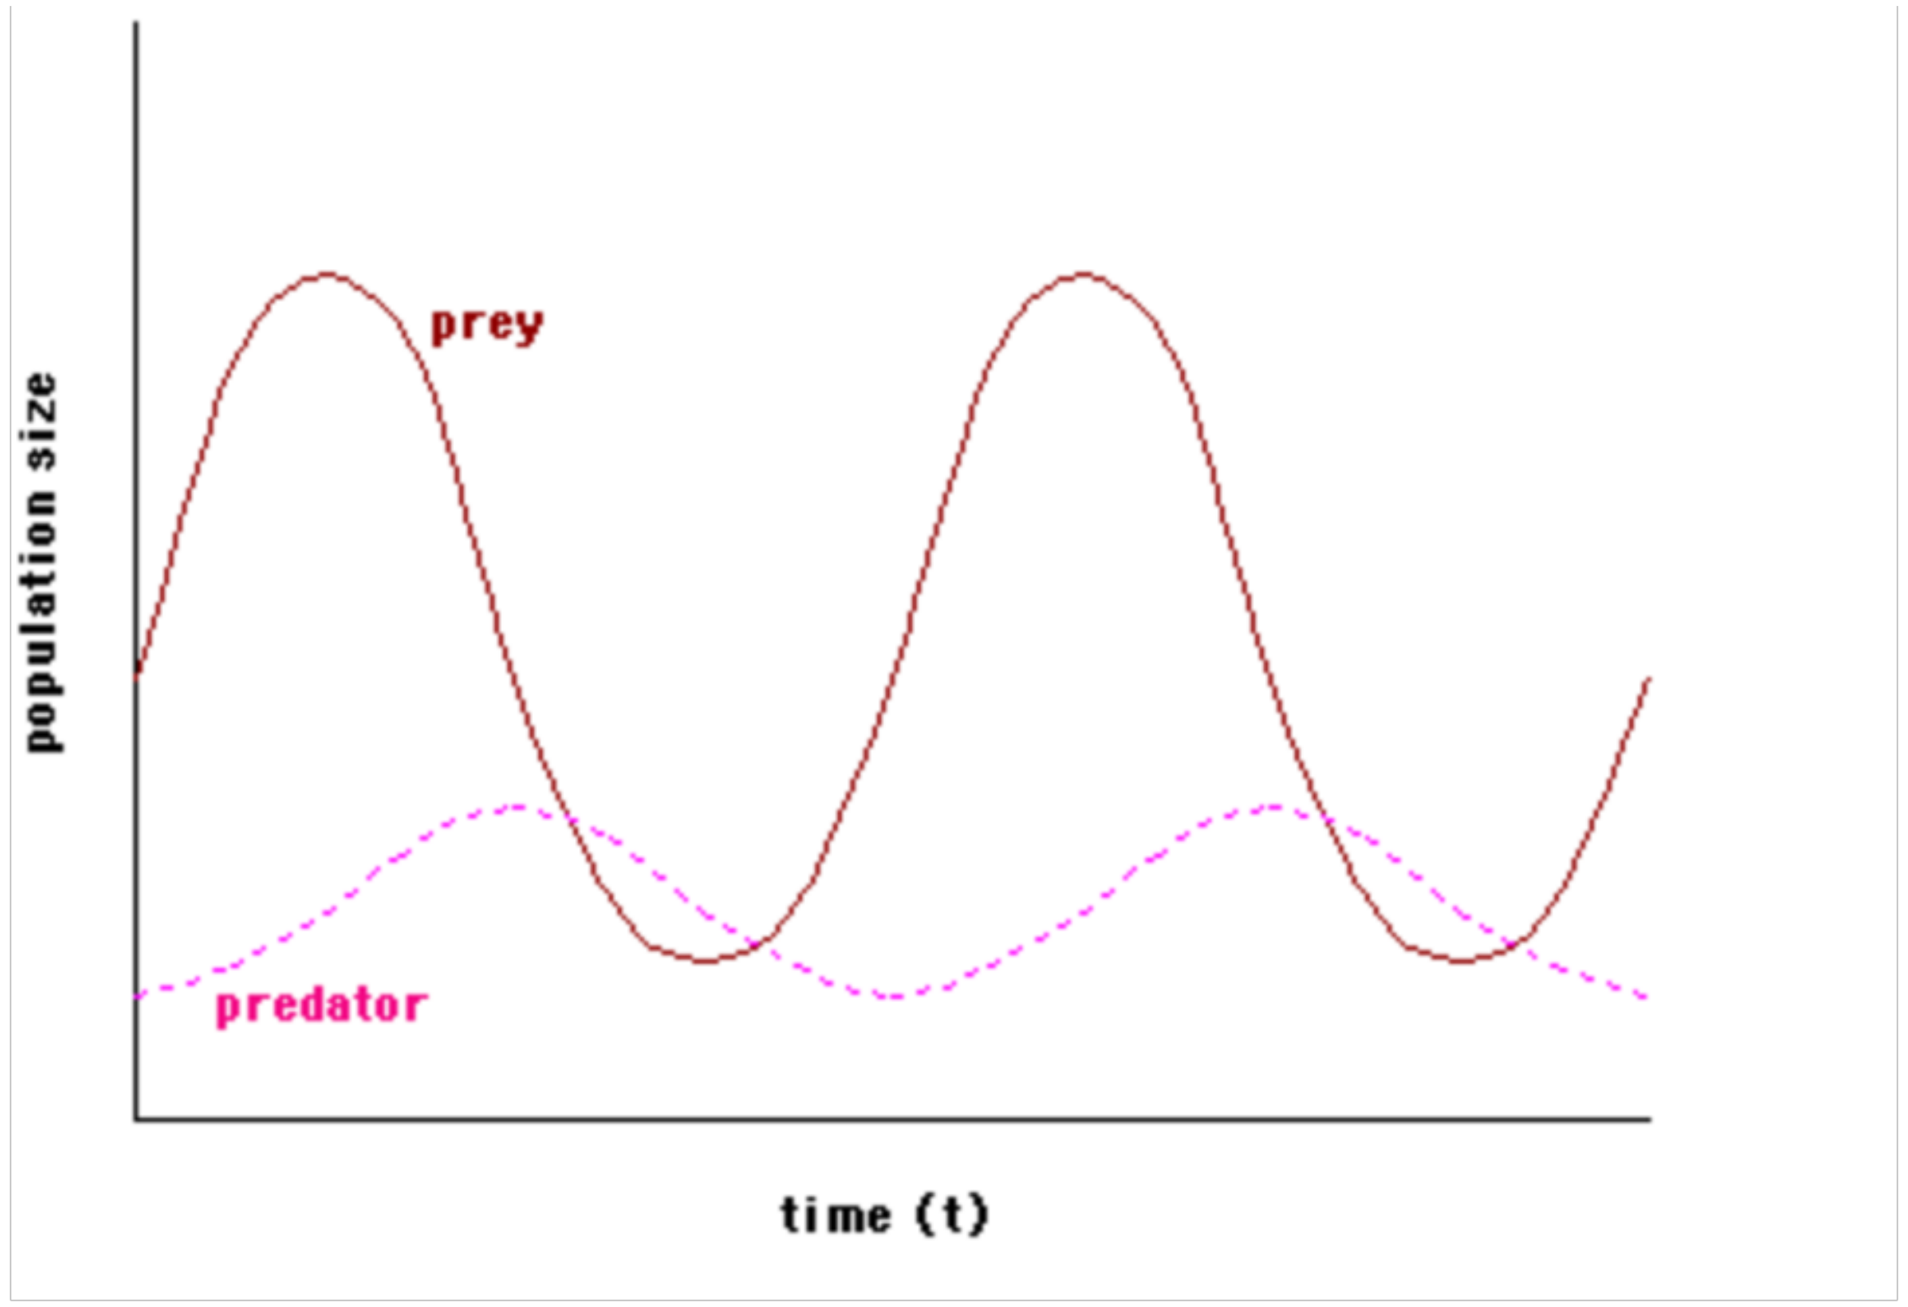

See http://www.tiem.utk.edu/~gross/bioed/bealsmodules/predator-prey.html

Asymmetric Games

	L	R
T	0, 0	-1, 1
B	X, -1	-2, -2

$$\dot{x}_i = x_i[(Ay)_i - x^T Ay]$$

$$\dot{y}_i = y_i[(Bx)_i - y^T Bx]$$

Payoffs A

	L	R
T	0	-1
B	X	-2

Payoffs B

	T	B
L	0	-1
R	1	-2

Asymmetric Games (see notebook!)

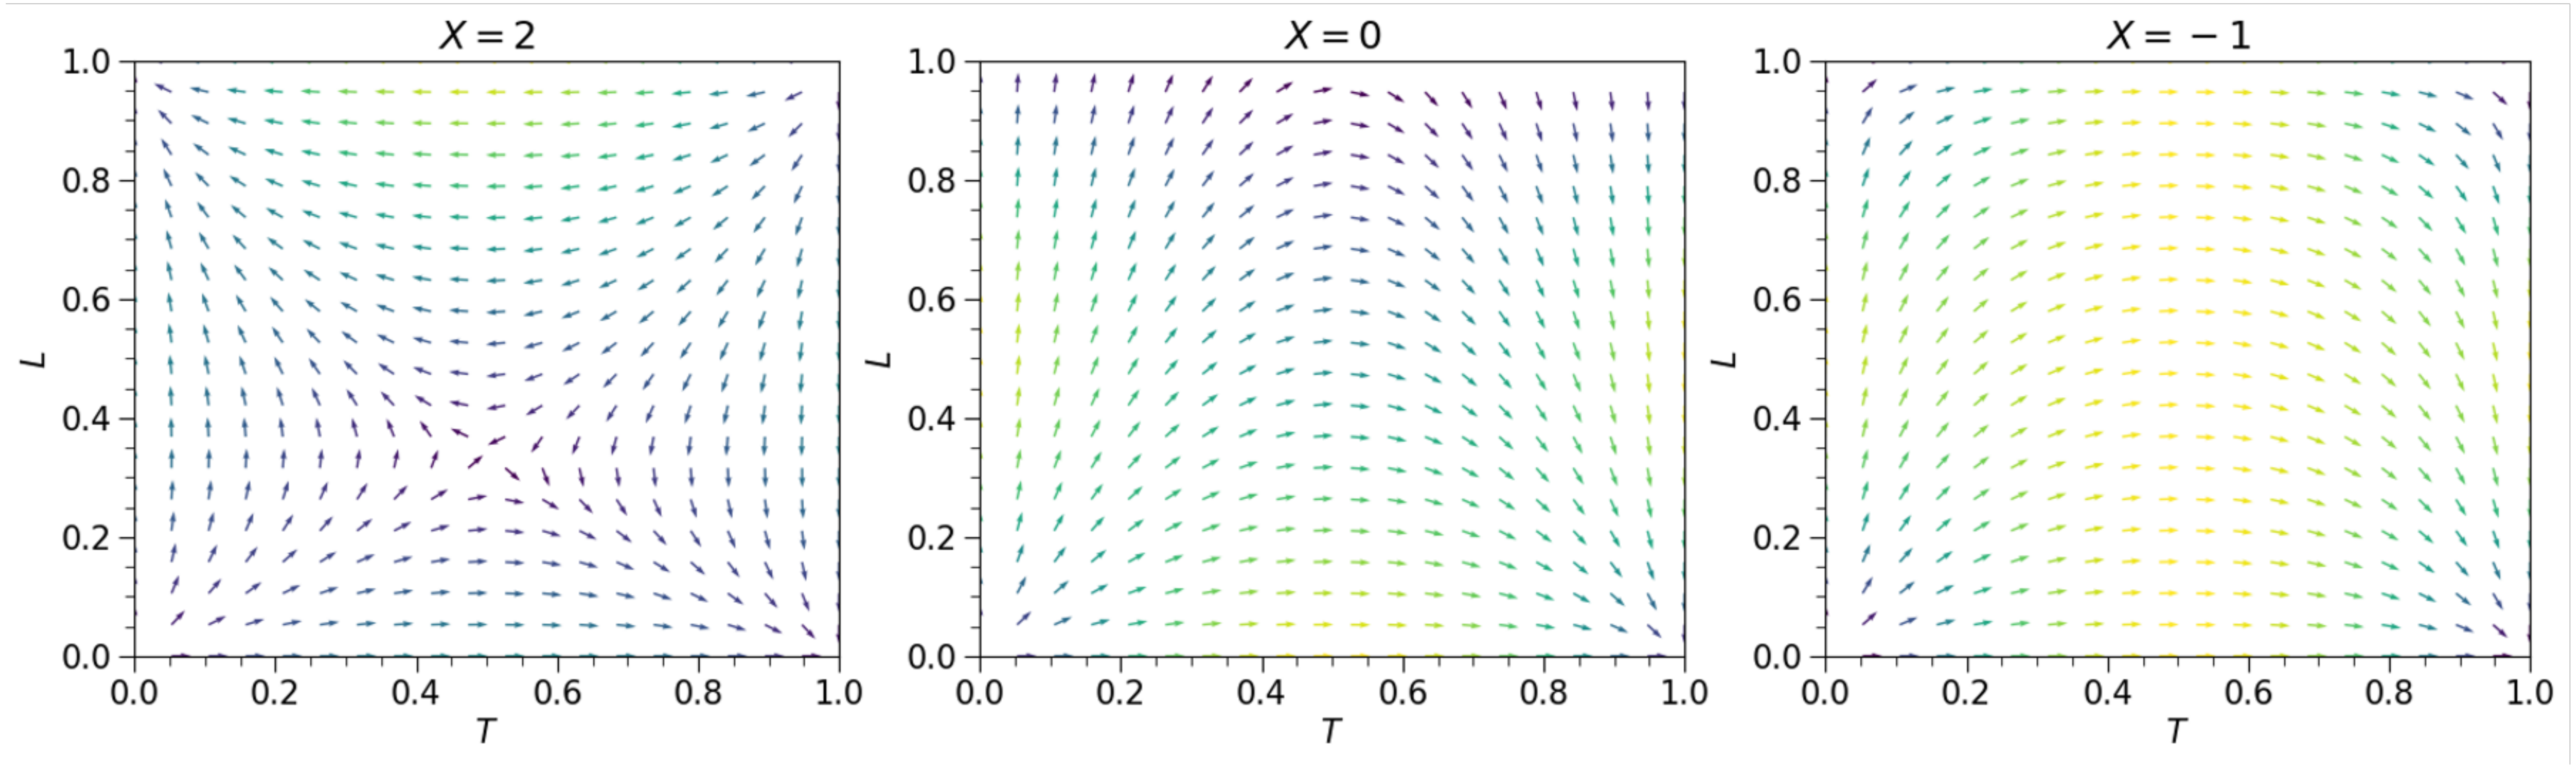

Part 3: Finite Populations

Evolution of trust

https://ncase.me/trust/

Evolutionary dynamics in Finite Populations

Evolutionary dynamics (Φ)

Moran process

Wright-Fisher process

(Agent-based simulation)

Adopted from Arne Traulsen

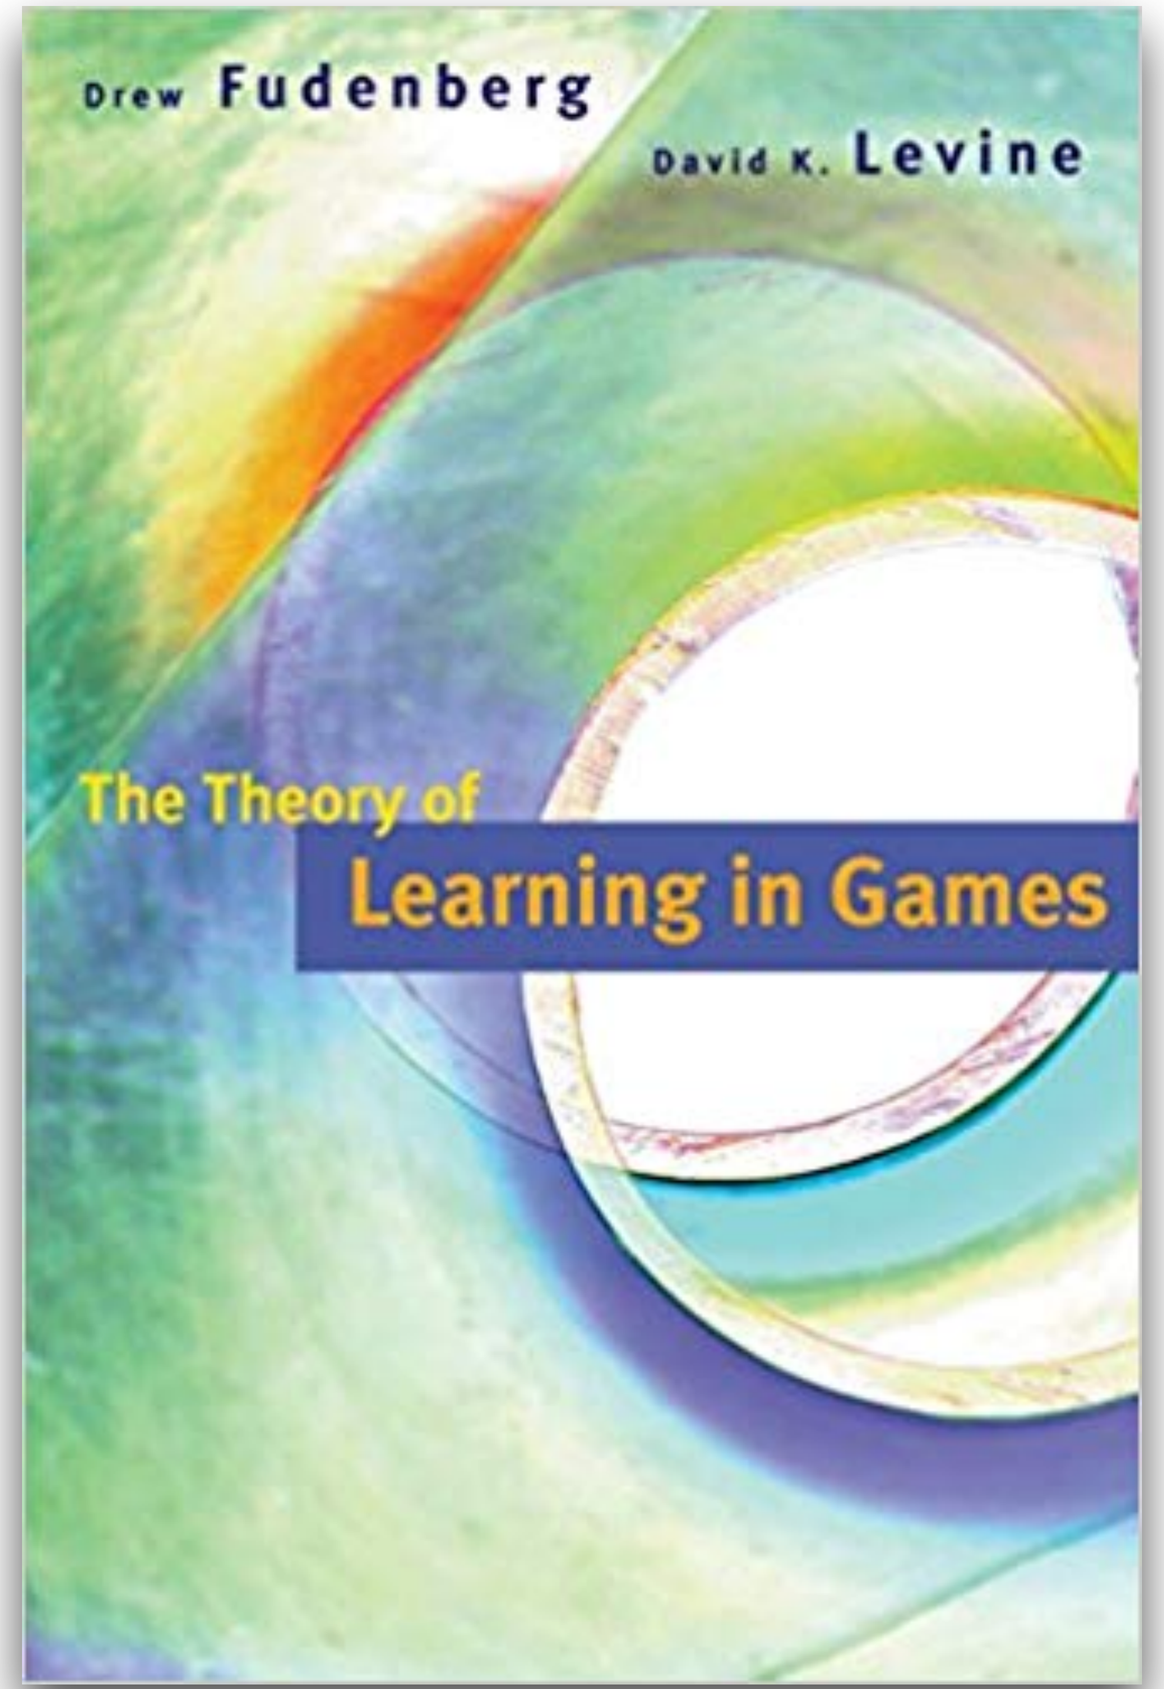

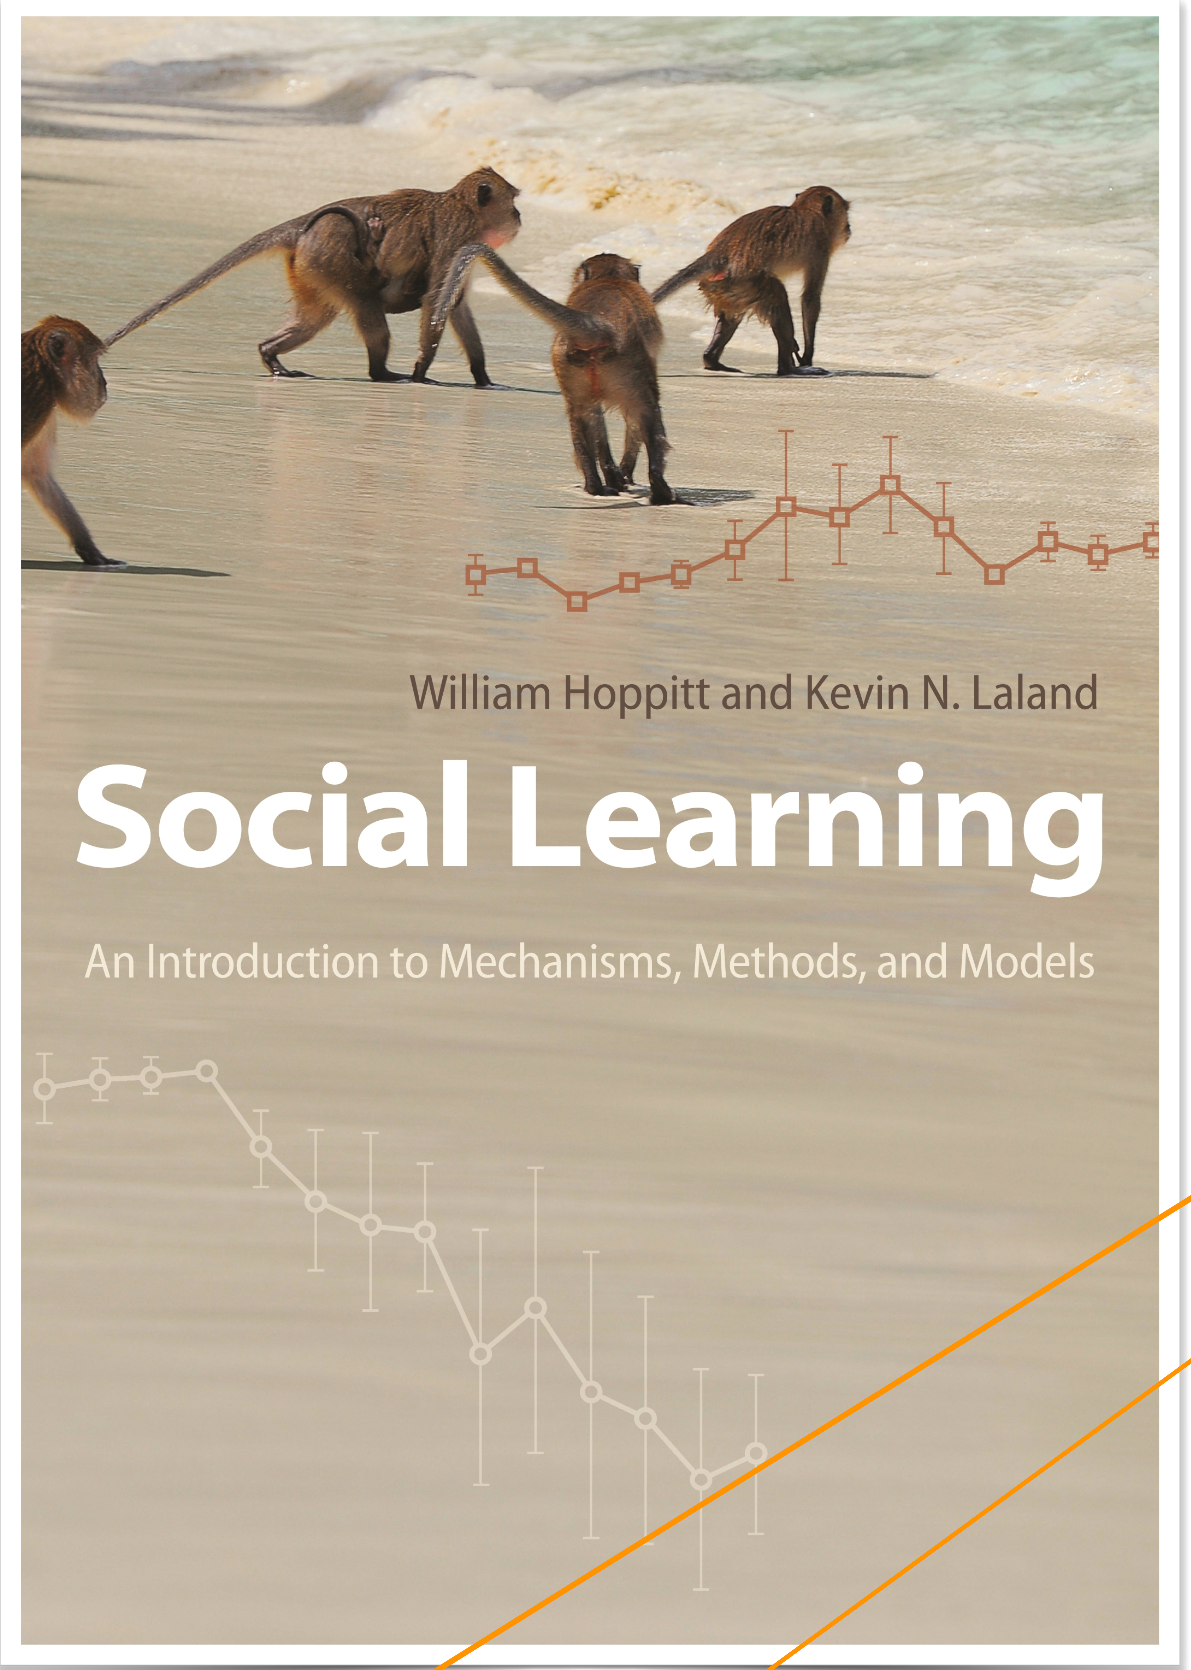

New behaviour is **acquired by observation/imitation**

Social learning is learning that is facilitated by observation , or interaction with, another individual or its products

Reinforcement learning (*Bush-Mosteller, Mach-Flach, Roth-Erev, …, learning automata, Q-learning*) models provide **individual learning by experience**

Evolutionary dynamics models use **learning by imitating the best**

A multi-agent model of social learning

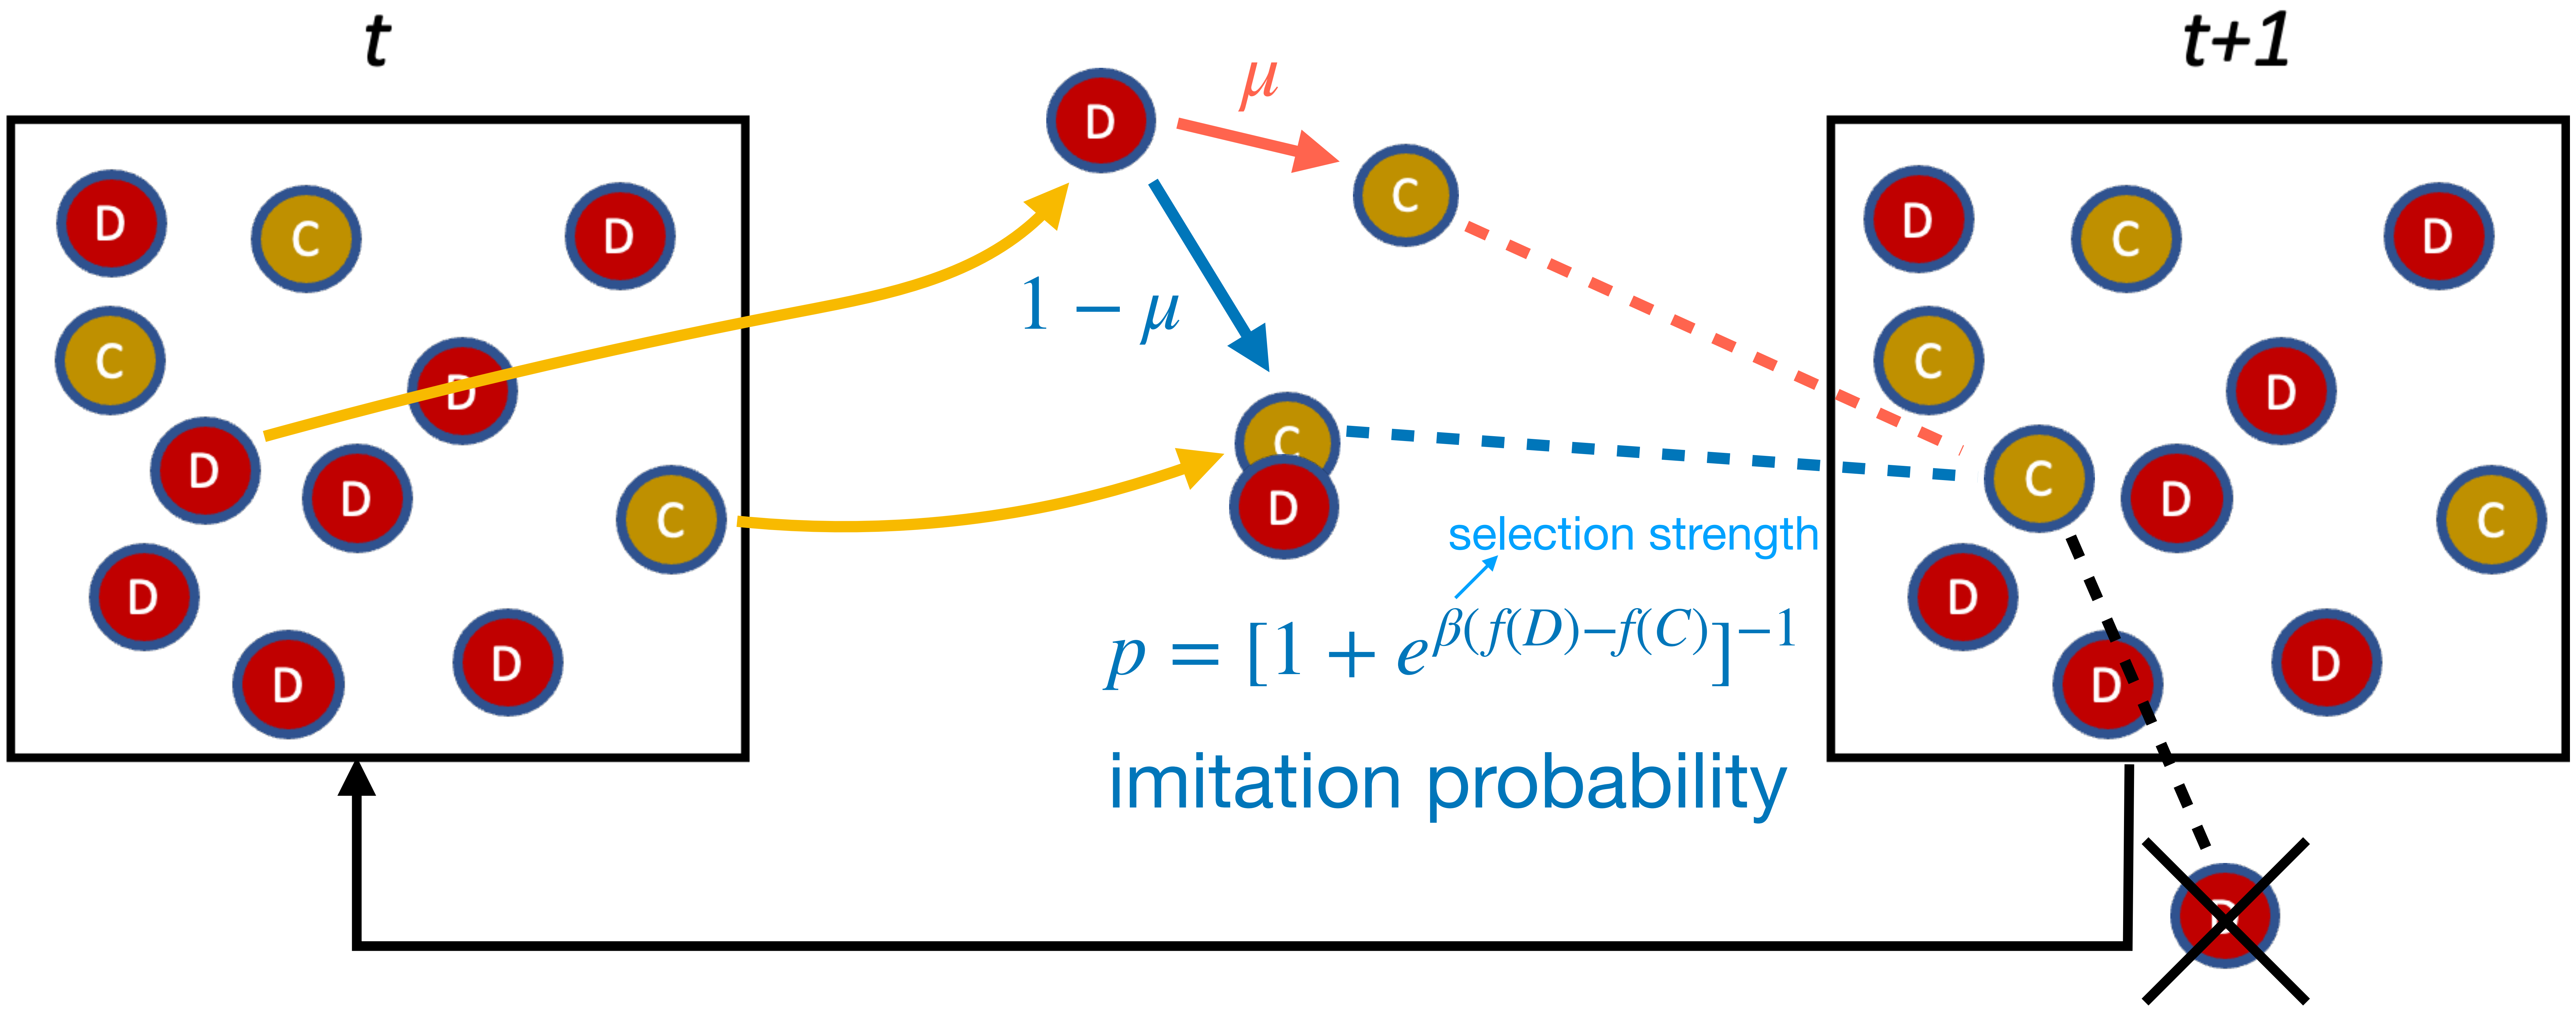

A multi-agent model of social learning

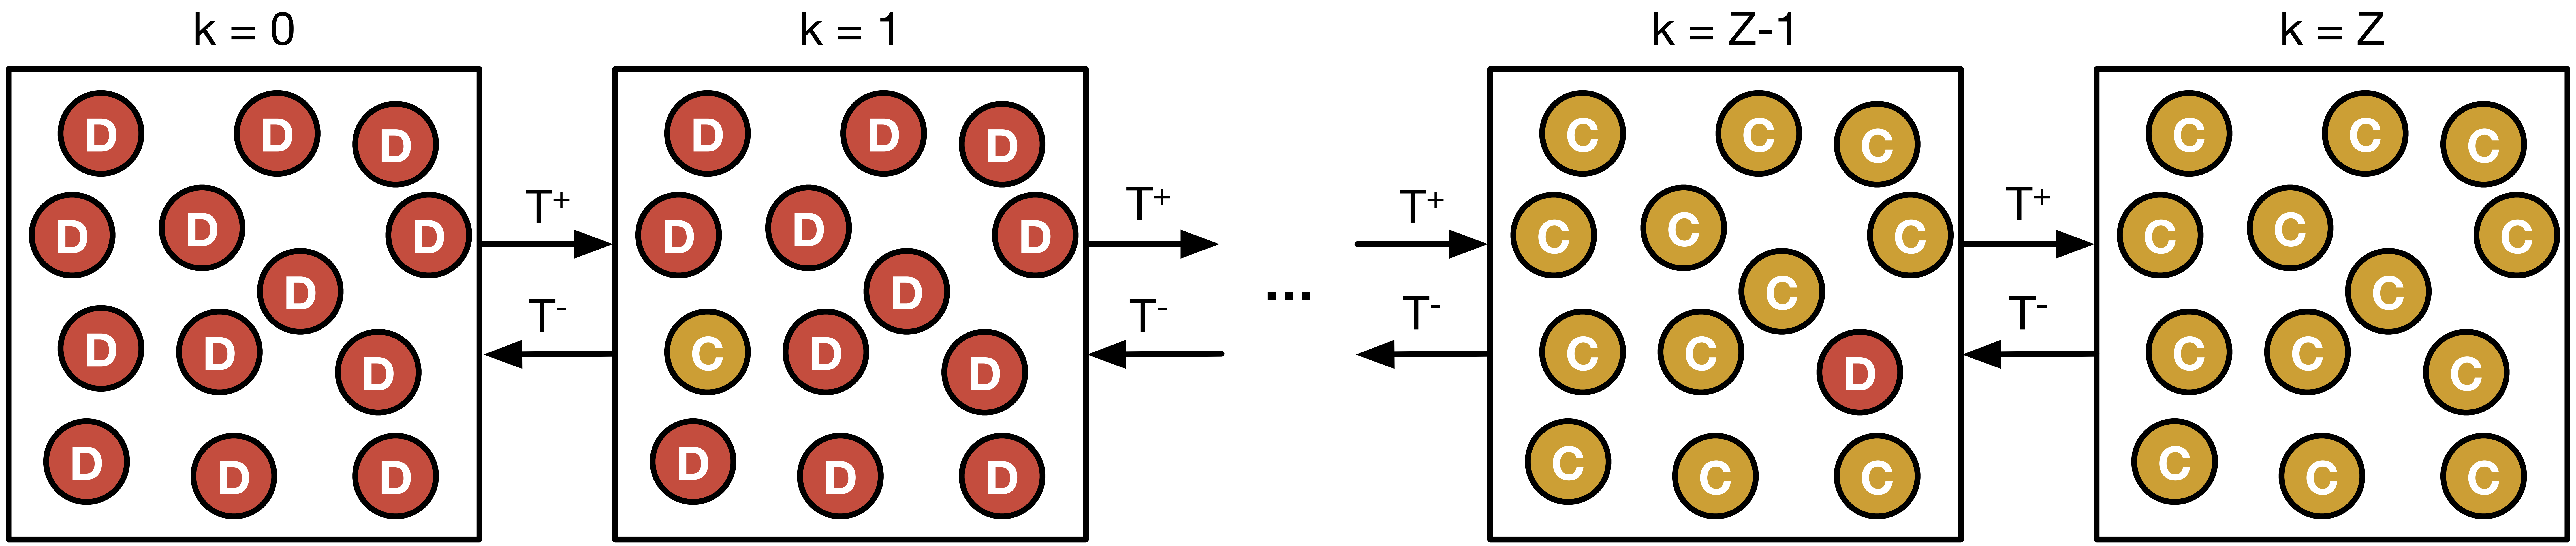

This process can be described by a **Markov Chain**

A multi-agent model of social learning

Fermi function

$$p \equiv [1 + e^{\beta(f_i(k) - f_j(k))}]^{-1}$$

Assuming 2 strategies C and D in a population of size Z, k C players and $Z - k$ D players

$$T^{+} = (1 - \mu)\frac{Z - k}{Z}\frac{k}{Z}[1 + e^{-\beta(f_C - f_D}]^{-1} + \mu\frac{Z - k}{Z}$$

$$T^{-} = (1 - \mu)\frac{k}{Z}\frac{Z - k}{Z}[1 + e^{\beta(f_C - f_D}]^{-1} + \mu\frac{k}{Z}$$

A multi-agent model of social learning

Fermi function

$$p \equiv [1 + e^{\beta(f_i(k) - f_j(k))}]^{-1}$$

Assuming 2 strategies C and D in a population of size Z, k C players and $Z - k$ D players

$$T^{+} = (1-\mu)\frac{Z-k}{Z}\frac{k}{Z}[1 + e^{-\beta(f_C - f_D}]^{-1} + \mu\frac{Z-k}{Z}$$

$$T^{-} = (1-\mu)\frac{k}{Z}\frac{Z-k}{Z}[1 + e^{\beta(f_C - f_D}]^{-1} + \mu\frac{k}{Z}$$

probability that the imitation process occurs and individuals adopting different strategies are selected

probability of mutating to strategy D

A multi-agent model of social learning

To calculate the fitness we now need to sample without replacement!

For 2-player games, we have:

$$f_D(k) = \frac{k-1}{Z-1}\Pi(D, D) + \frac{Z-k}{Z-1}\Pi(D, C)$$

$$f_C(k) = \frac{k}{Z-1}\Pi(C, D) + \frac{Z-k-1}{Z-1}\Pi(C, C)$$

A multi-agent model of social learning

To calculate the fitness we now need to sample without replacement!

For n-player games, where N is the size of the group, we have (hypergeometric sampling):

$$f_D(k) = \binom{Z-1}{N-1}^{-1} \sum_{j=0}^{N-1} \binom{k-1}{j} \binom{Z-k-1}{N-j-1} \Pi_D(j)$$

$$f_D(k) = \binom{Z-1}{N-1}^{-1} \sum_{j=0}^{N-1} \binom{k-1}{j} \binom{Z-k}{N-j-1} \Pi_C(j+1)$$

A multi-agent model of social learning

With this we can define the transition matrix T that maps the probabilities of transitioning from a state with k Ds to an adjacent state with $k+1$ or $k-1$ Ds:

$$
\begin{aligned}
T_{i+1,i} &= T^{-} \\
T_{i,i+1} &= T^{+} \\
T_{1,i} &= 1 - T^{+} - T^{-}
\end{aligned}
$$

Important indicators

Gradient of selection (most likely path for the evolution):

$$G(x) = T^{+}(k) - T^{-}(k) = (1-\mu)\frac{k}{Z}\frac{Z-k}{Z}\tan h\left(\frac{\beta}{2}[f_C(k) - f_D(k)]\right) + \mu$$

Stationary distribution (the time spent at each state)

The stationary distribution can be computed as the left eigenvector associated with the eigenvalue 1 of the transition matrix.

Vasconcelos, V. V., Santos, F. C. & Pacheco, J. M. A bottom-up institutional approach to cooperative governance of risky commons. *Nature Climate Change* **3**, 797–801 (2013).

Evolutionary dynamics

Evolutionary dynamics (Φ)

Moran process
Wright-Fisher process

Replication dynamics

$\dot{x}_i = x_i (\pi_i - <\pi>)$

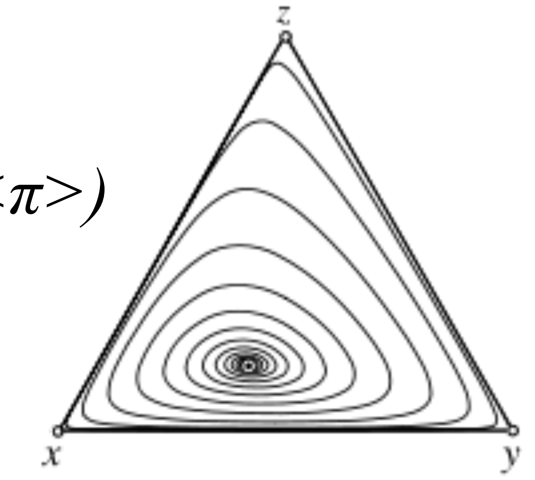

N→∞

(Nonlinear) dynamical systems

Agent-based simulation

Adopted from Arne Traulsen

Examples: Hawk-Dove: effect of Z

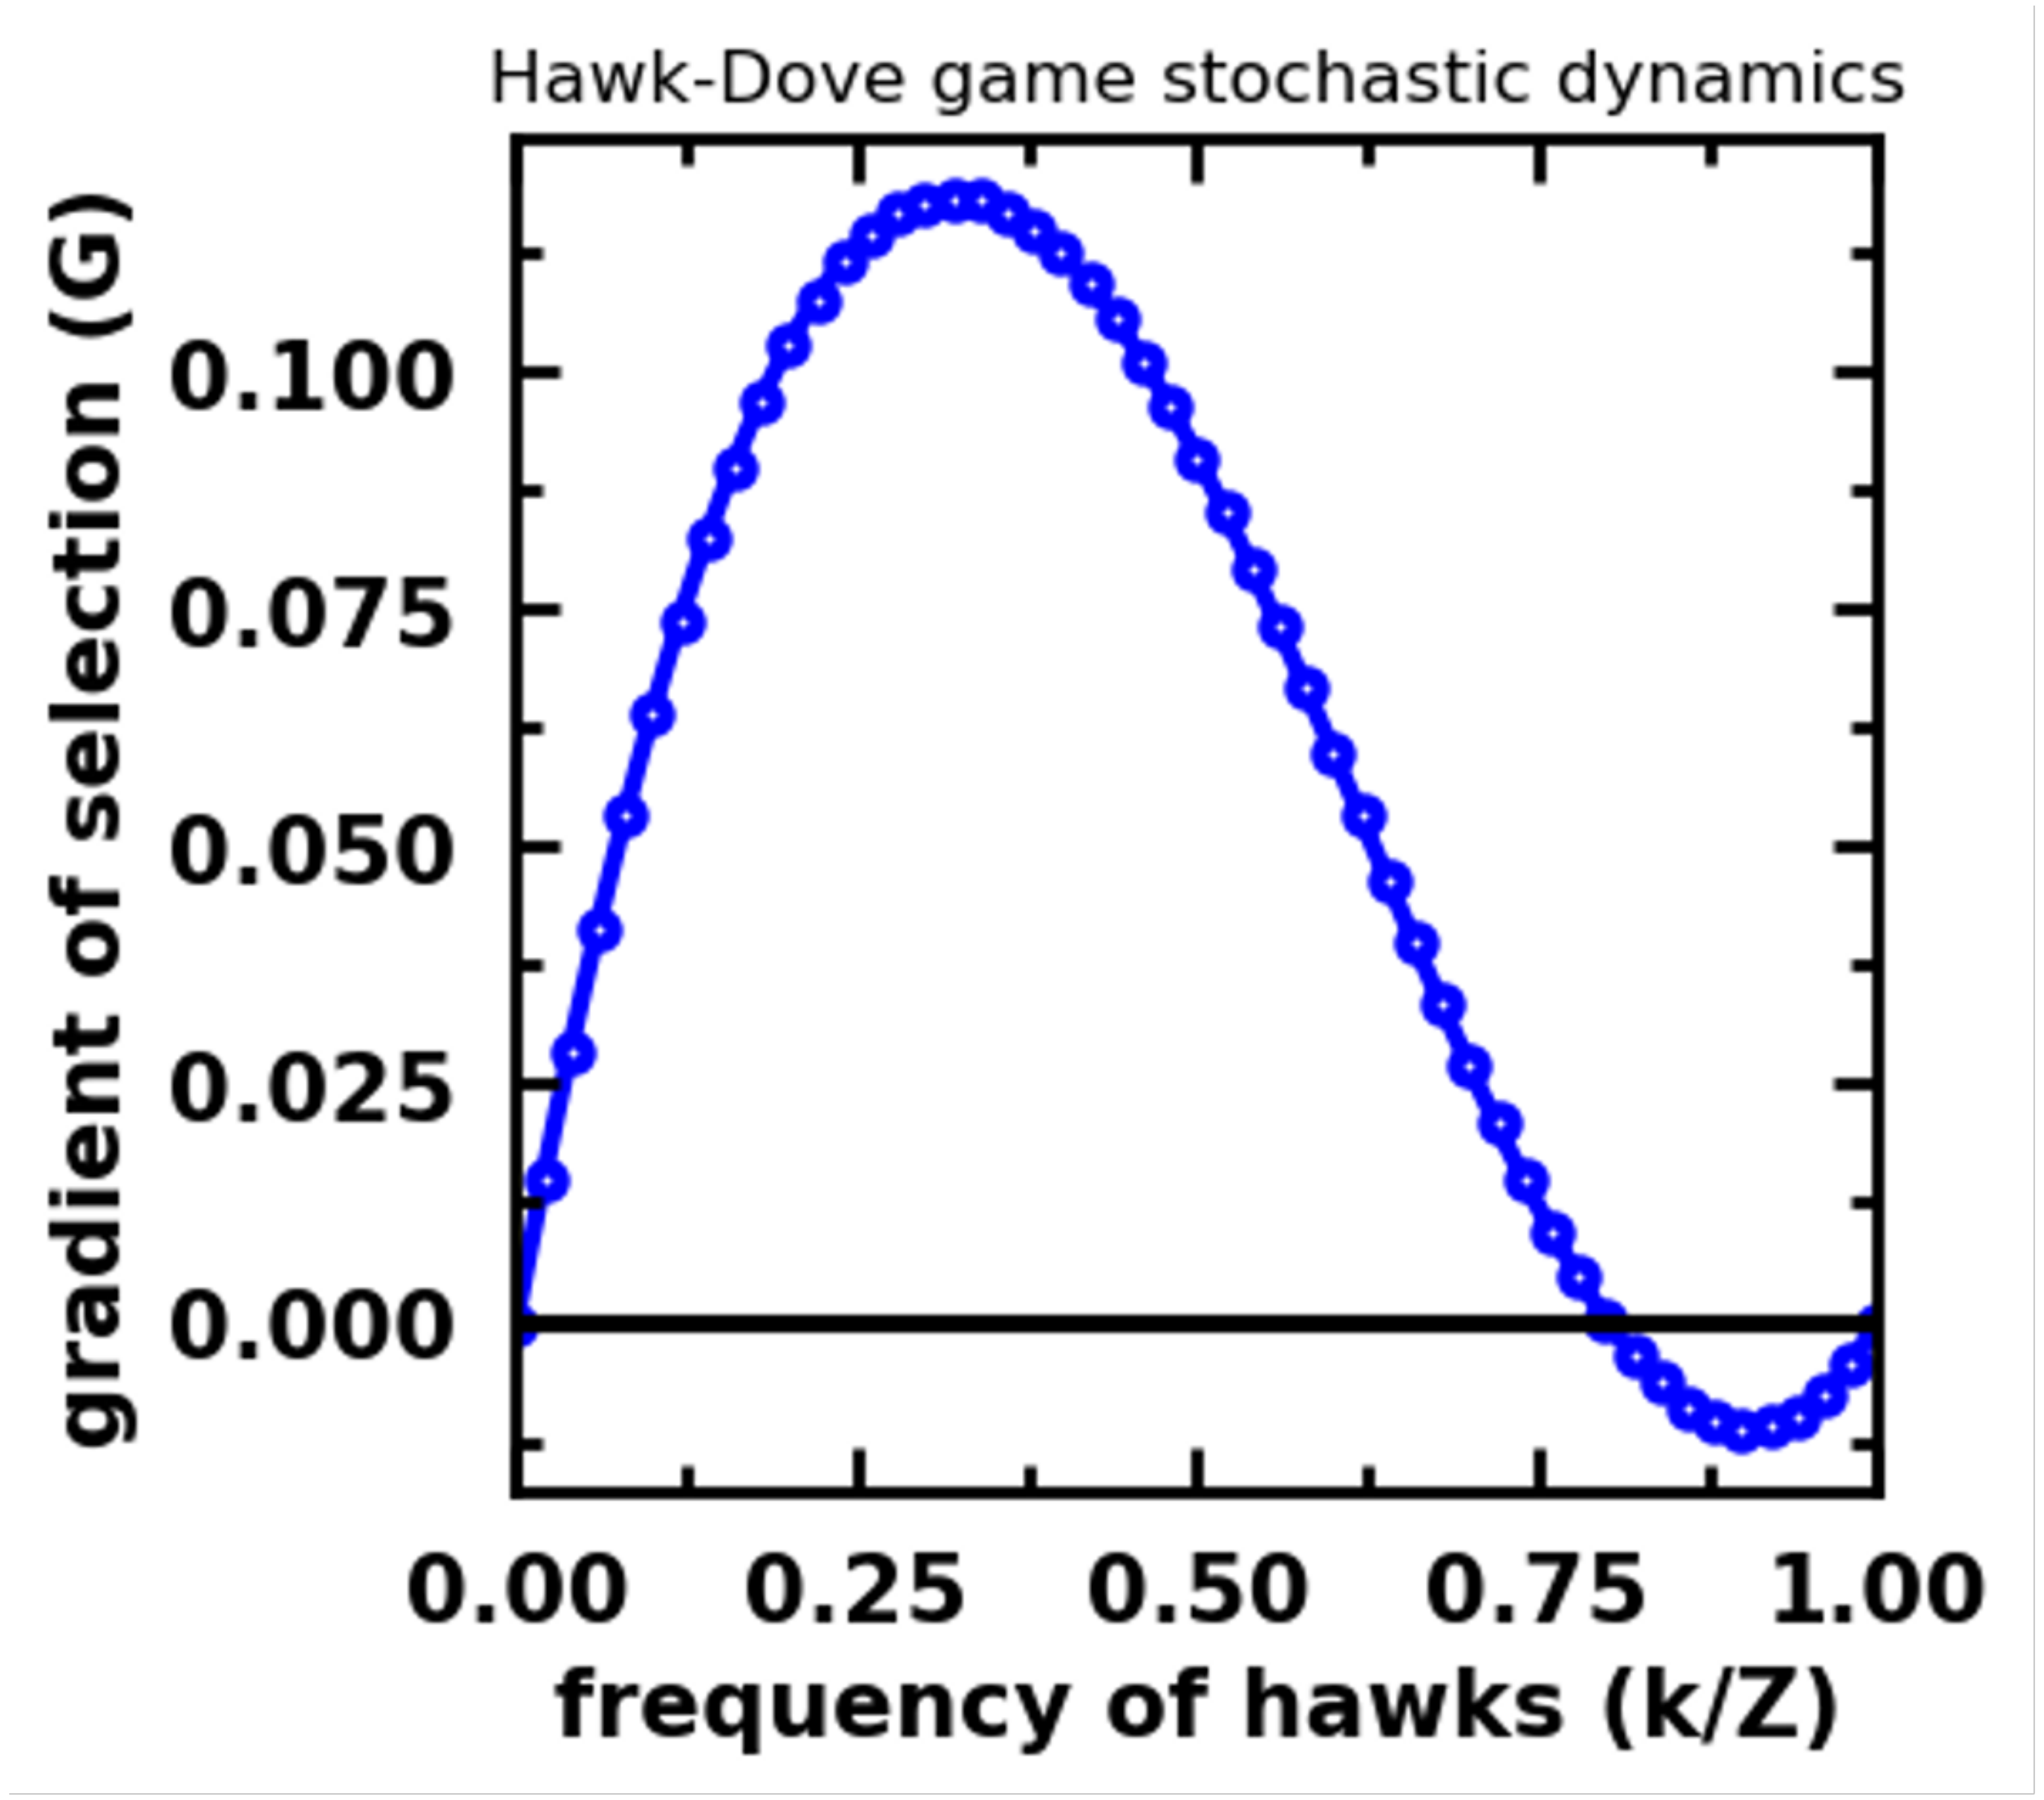

$V = 2,\ D = 3,\ T = 1\ (V < D)$

$$\begin{pmatrix} -0.5, \ -0.5 & 2, \ 0 \\ 0, \ 2 & 0, \ 0 \end{pmatrix}$$

$Z = 100$

$\beta = 1$

$\mu = 0$

$x_i \equiv [k_i/Z]$

Examples: Hawk-Dove: effect of Z

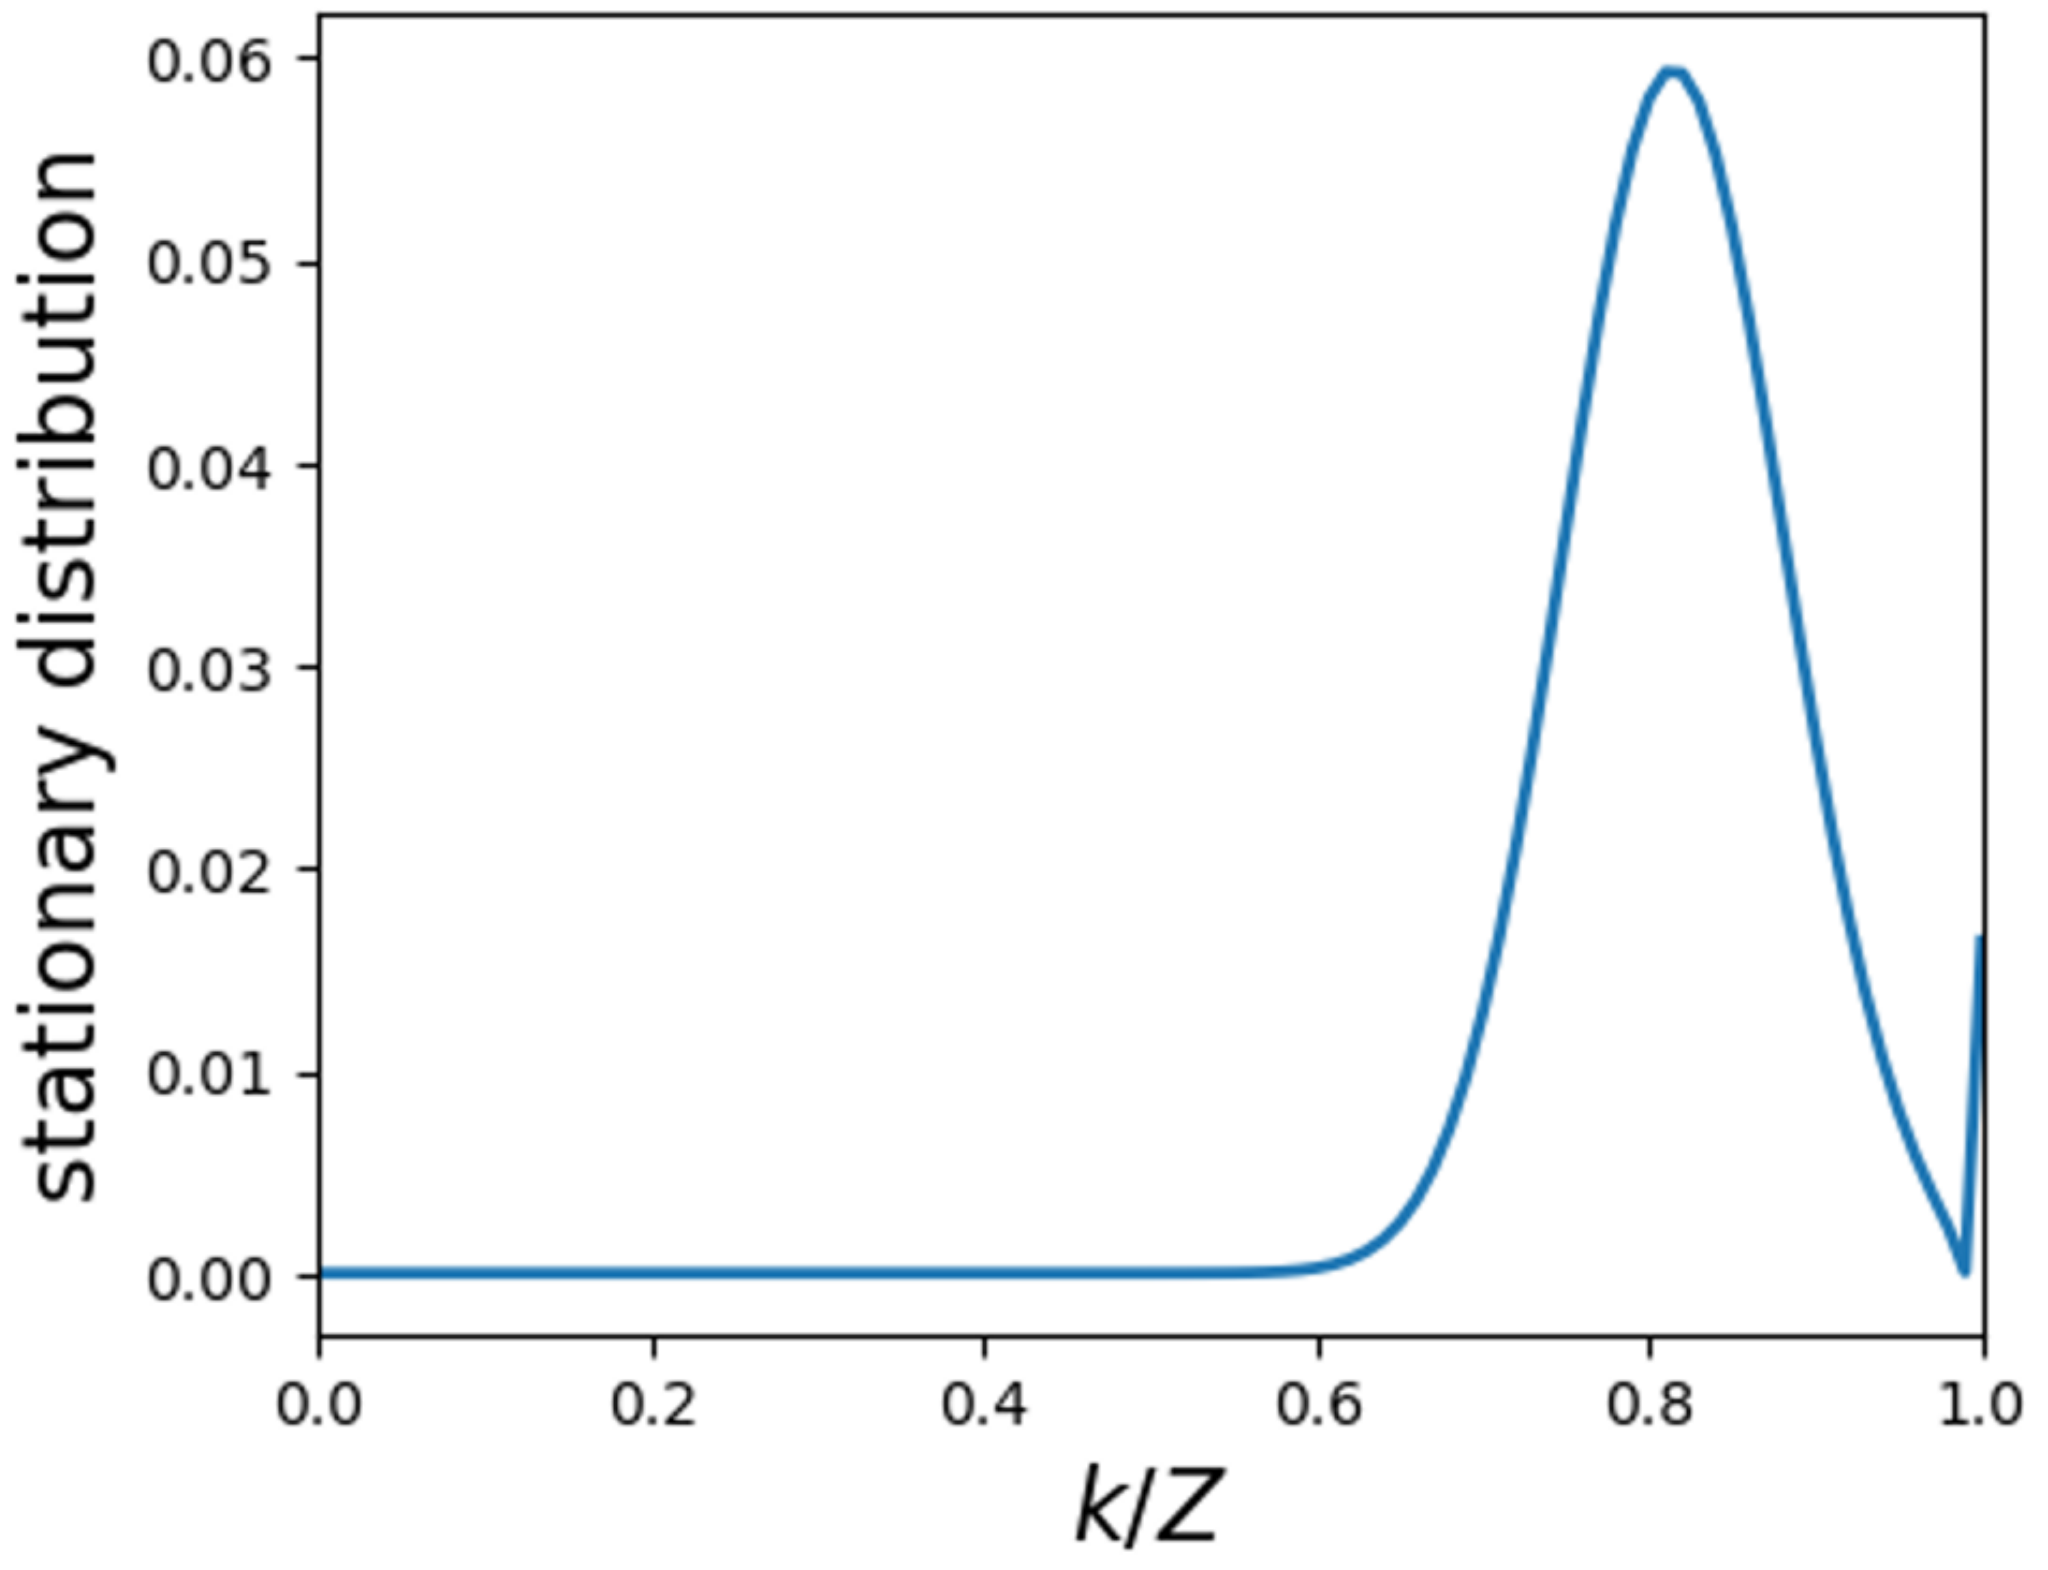

$V = 2,\ D = 3,\ T = 1\ (V < D)$

$$\begin{pmatrix} -0.5,\ -0.5 & 2,\ 0 \\ 0,\ 2 & 0,\ 0 \end{pmatrix}$$

$Z = 100$

$\beta = 1$

$\mu = 1e - 3$

$x_i \equiv [k_i/Z]$

Examples: Hawk-Dove: effect of Z

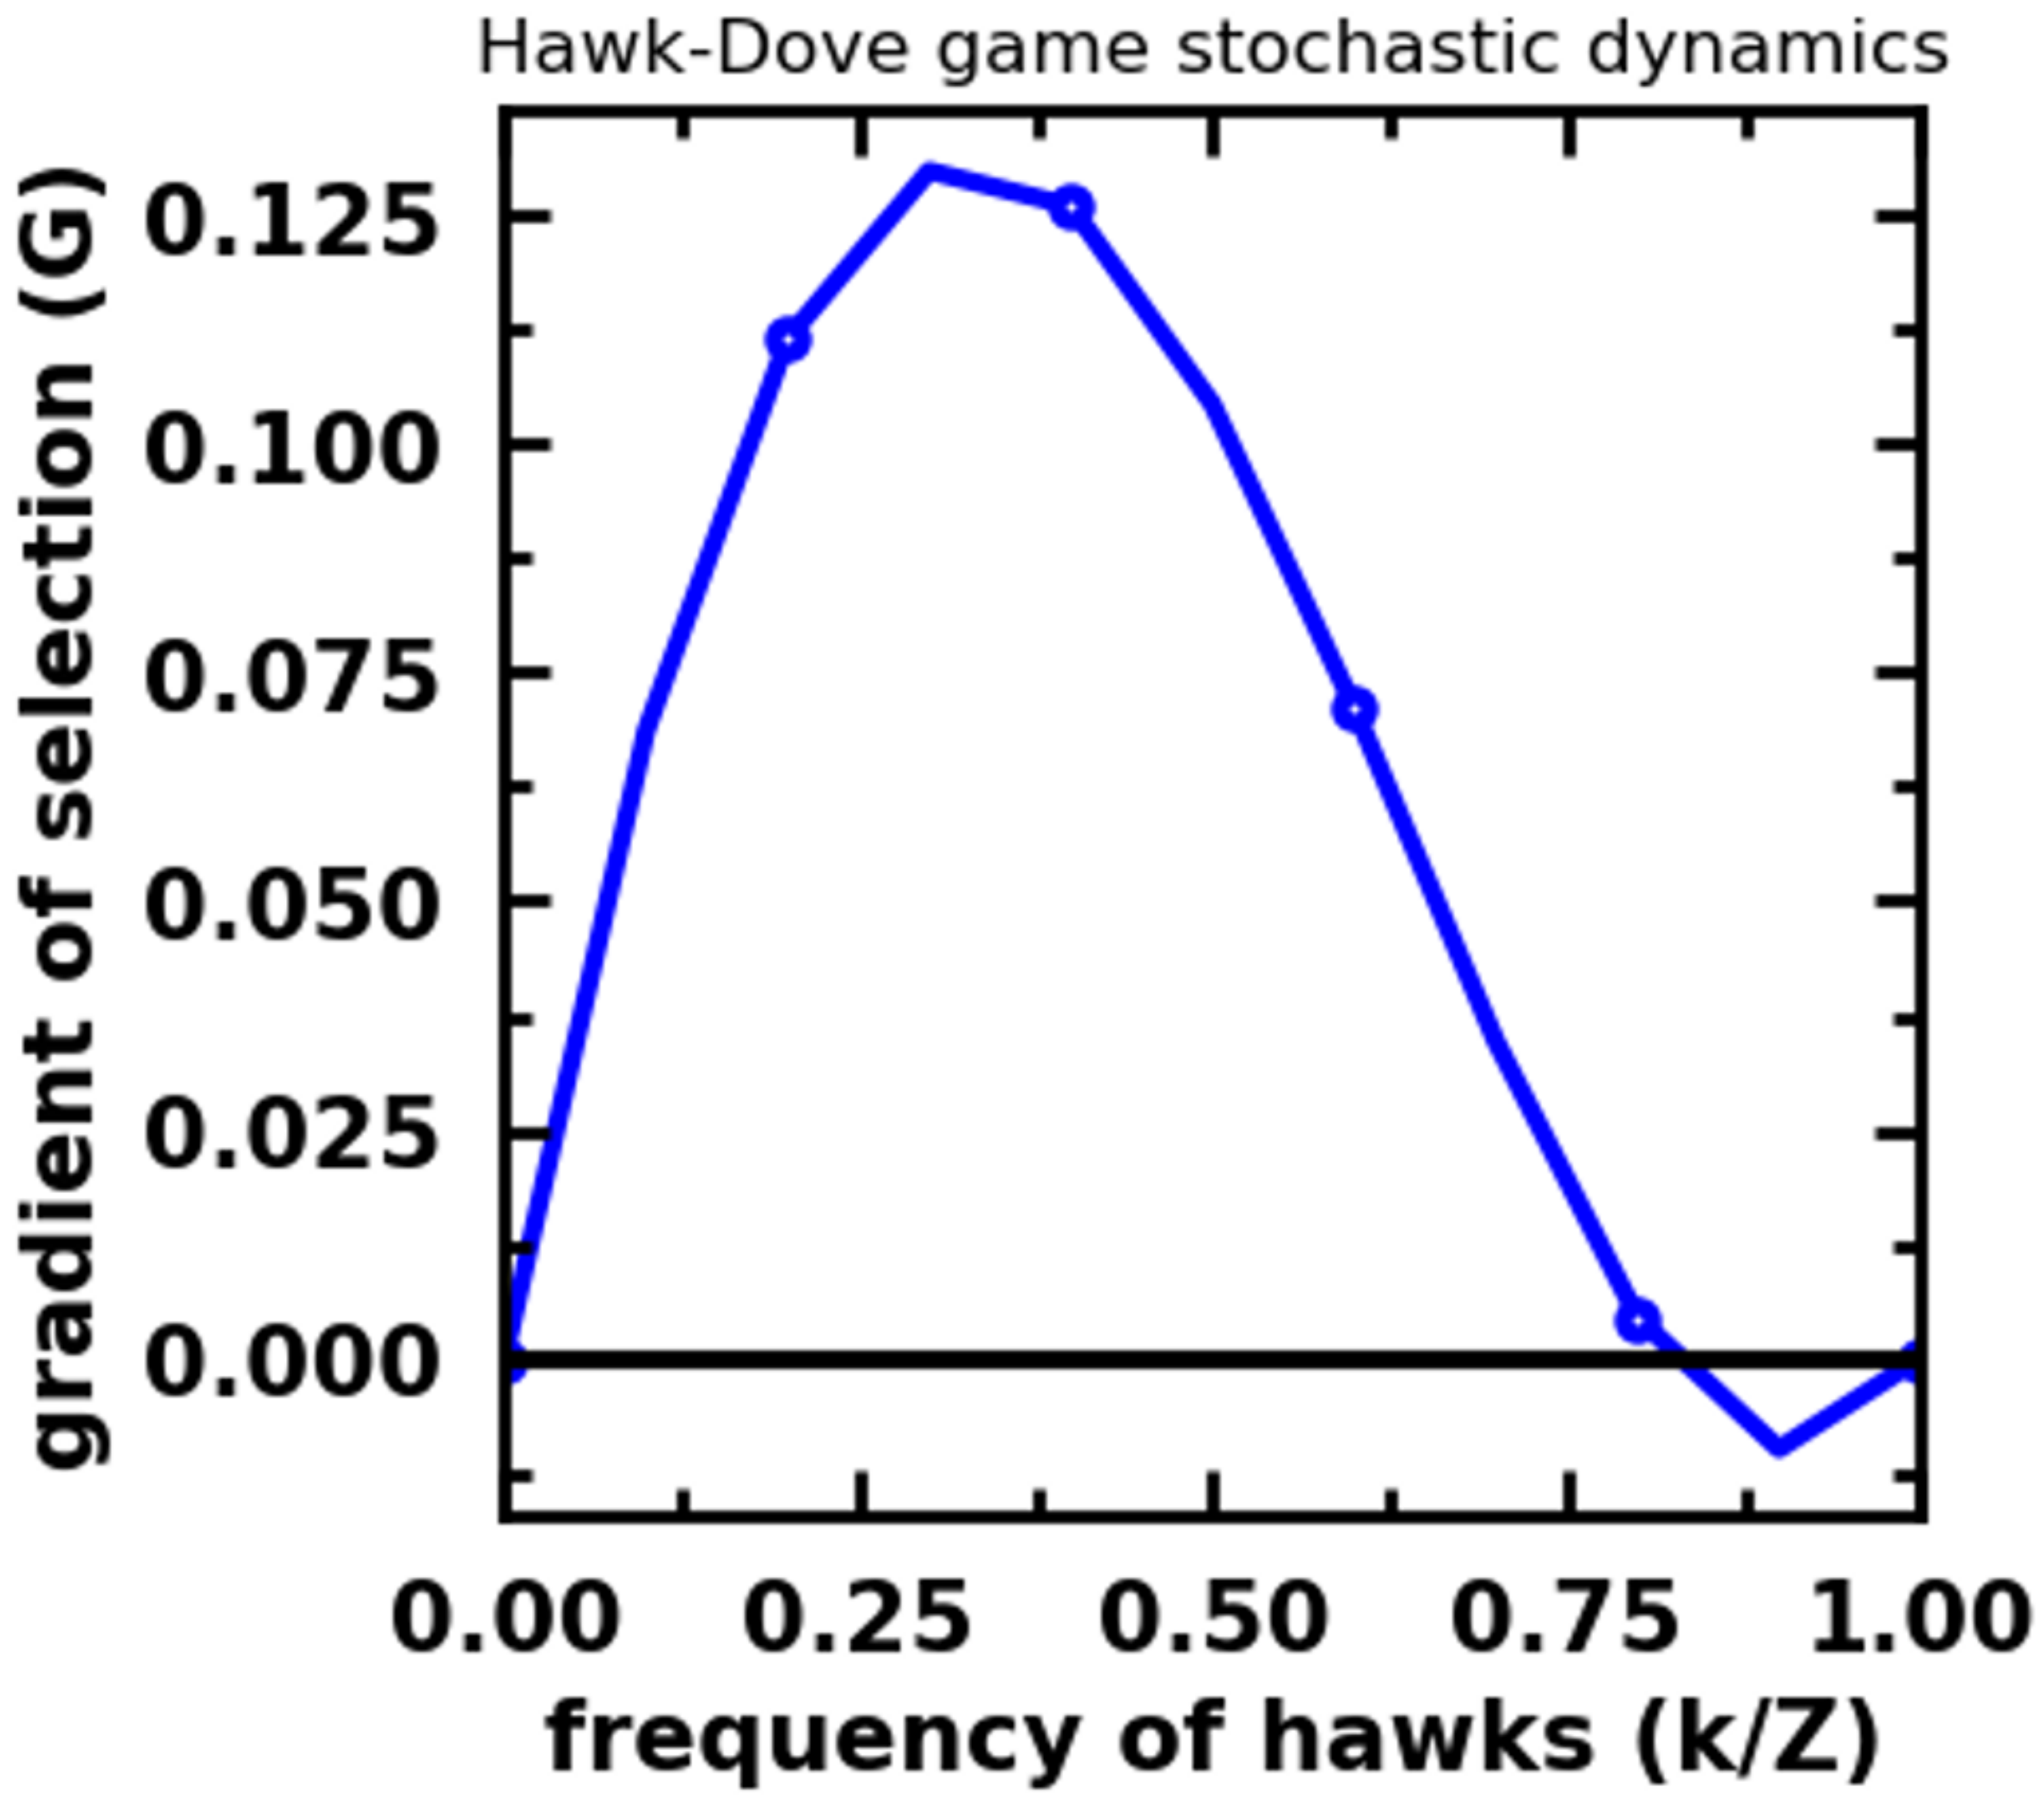

$V = 2,\ D = 3,\ T = 1\ (V < D)$

$$\begin{pmatrix} -0.5, \ -0.5 & 2, \ 0 \\ 0, \ 2 & 0, \ 0 \end{pmatrix}$$

$Z = 10$

$\beta = 1$

$\mu = 0$

$x_i \equiv [k_i/Z]$

Examples: Hawk-Dove: effect of Z

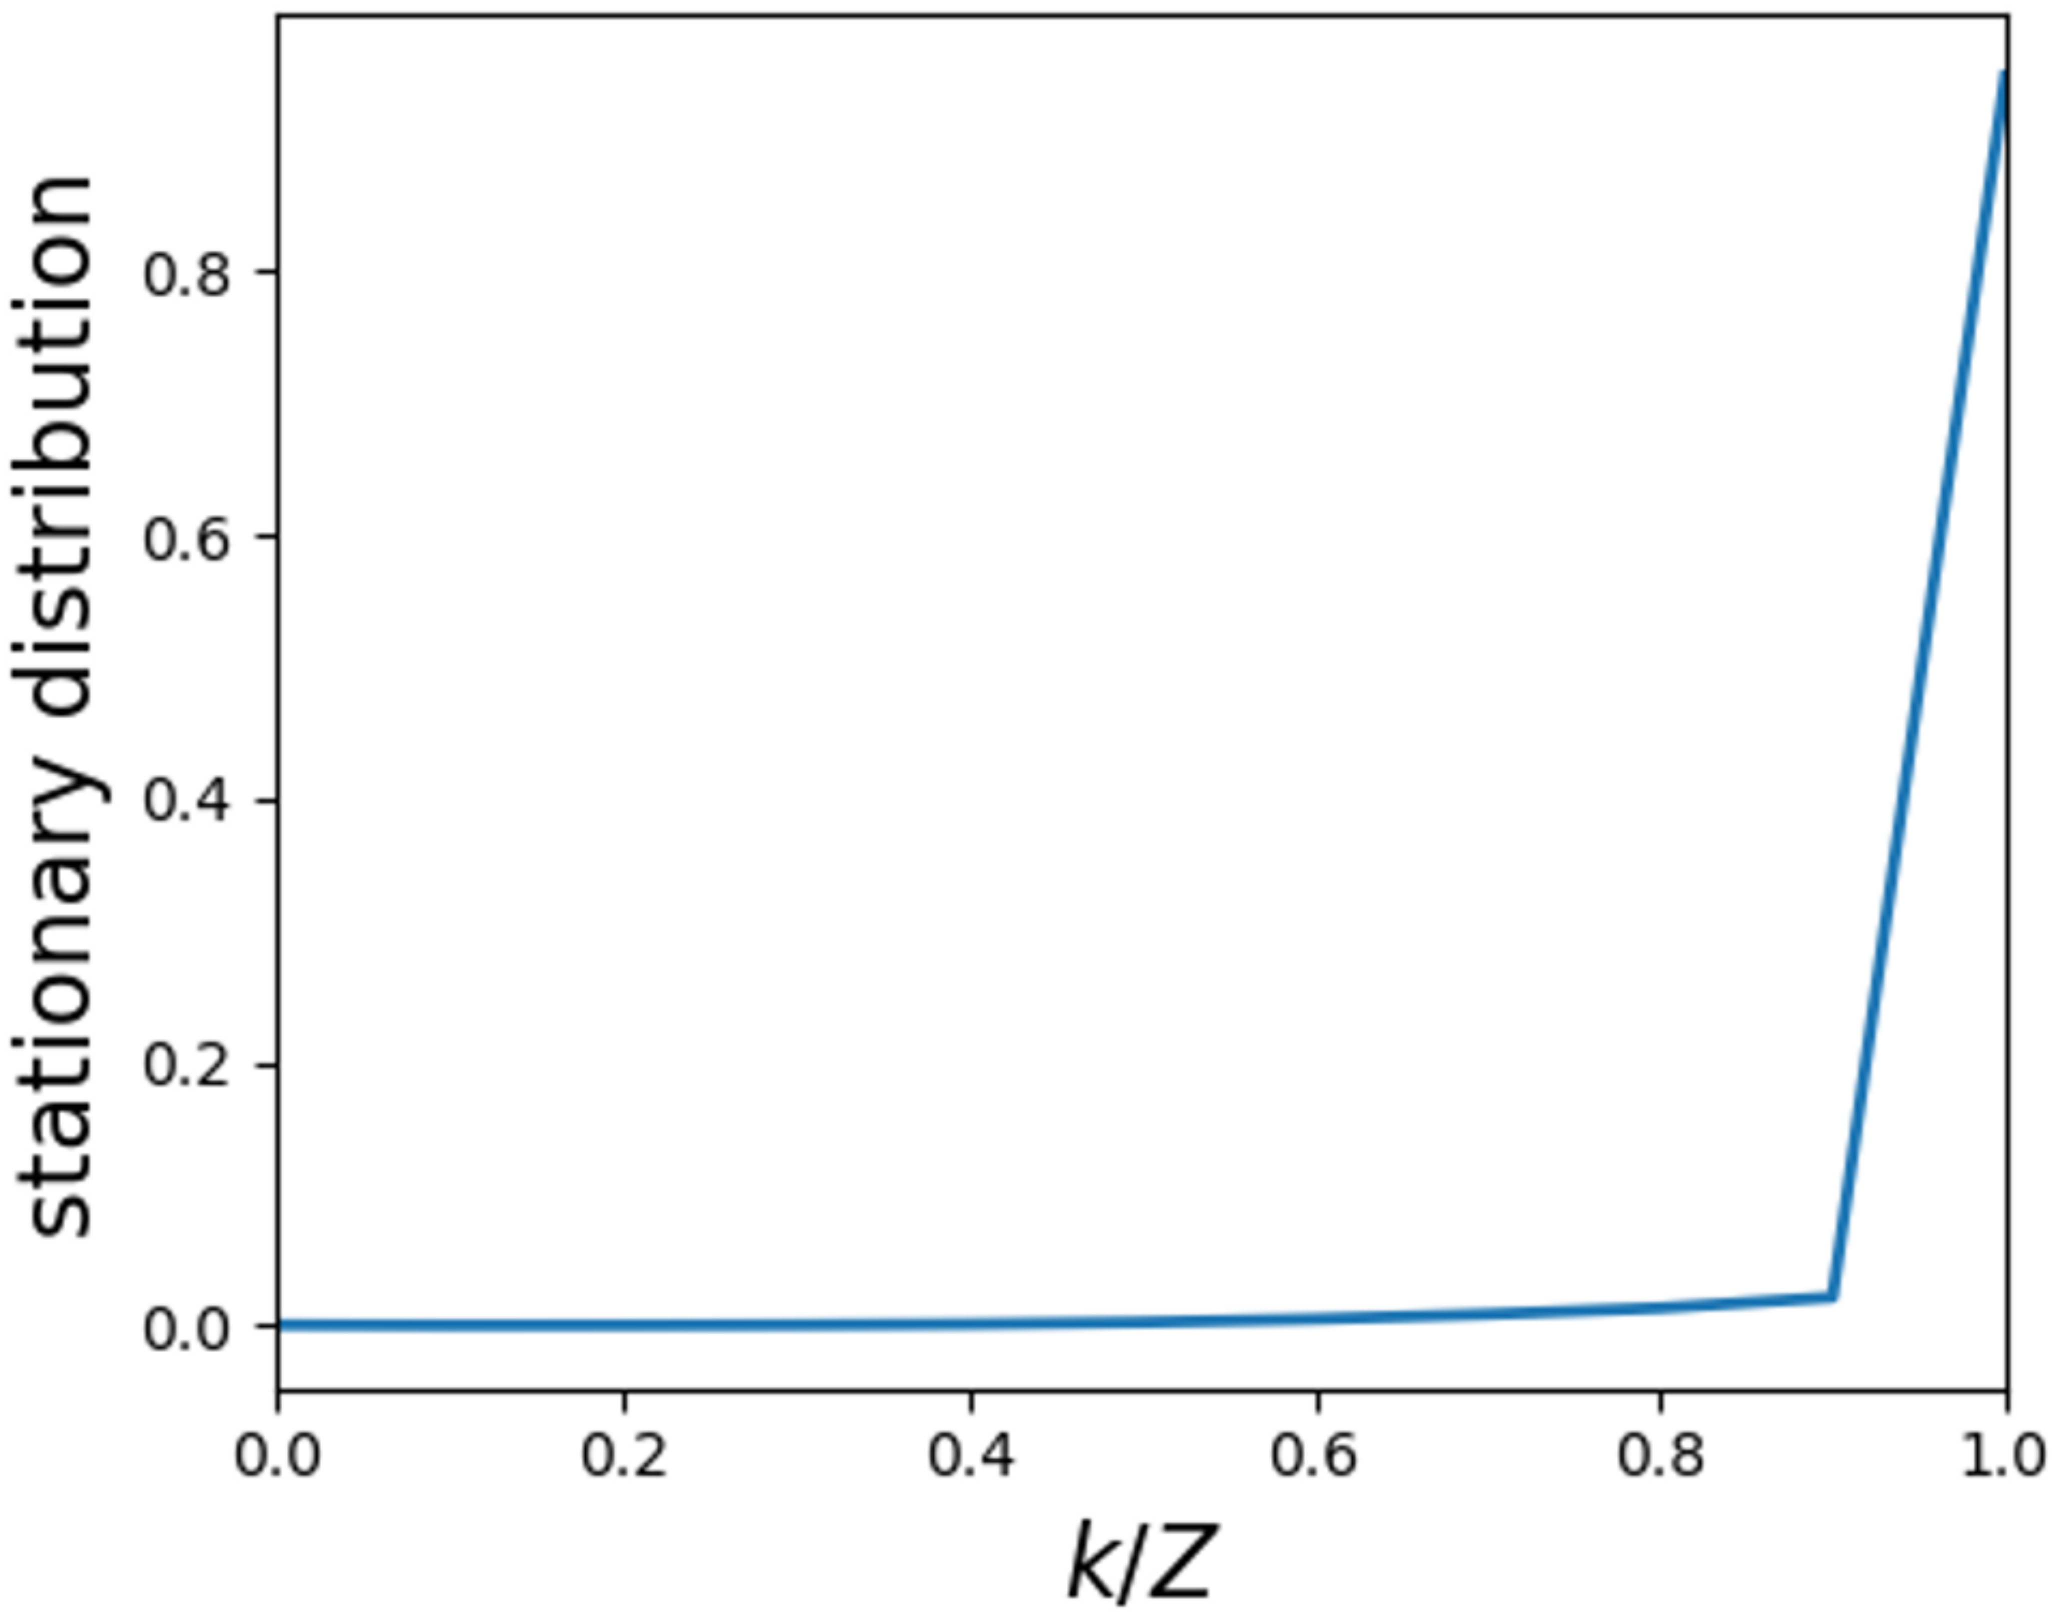

$V = 2,\ D = 3,\ T = 1\ (V < D)$

$$\begin{pmatrix} -0.5,\ -0.5 & 2,\ 0 \\ 0,\ 2 & 0,\ 0 \end{pmatrix}$$

$Z = 10$

$\beta = 1$

$\mu = 1e - 3$

$x_i \equiv [k_i/Z]$

Examples: Hawk-Dove: effect of β

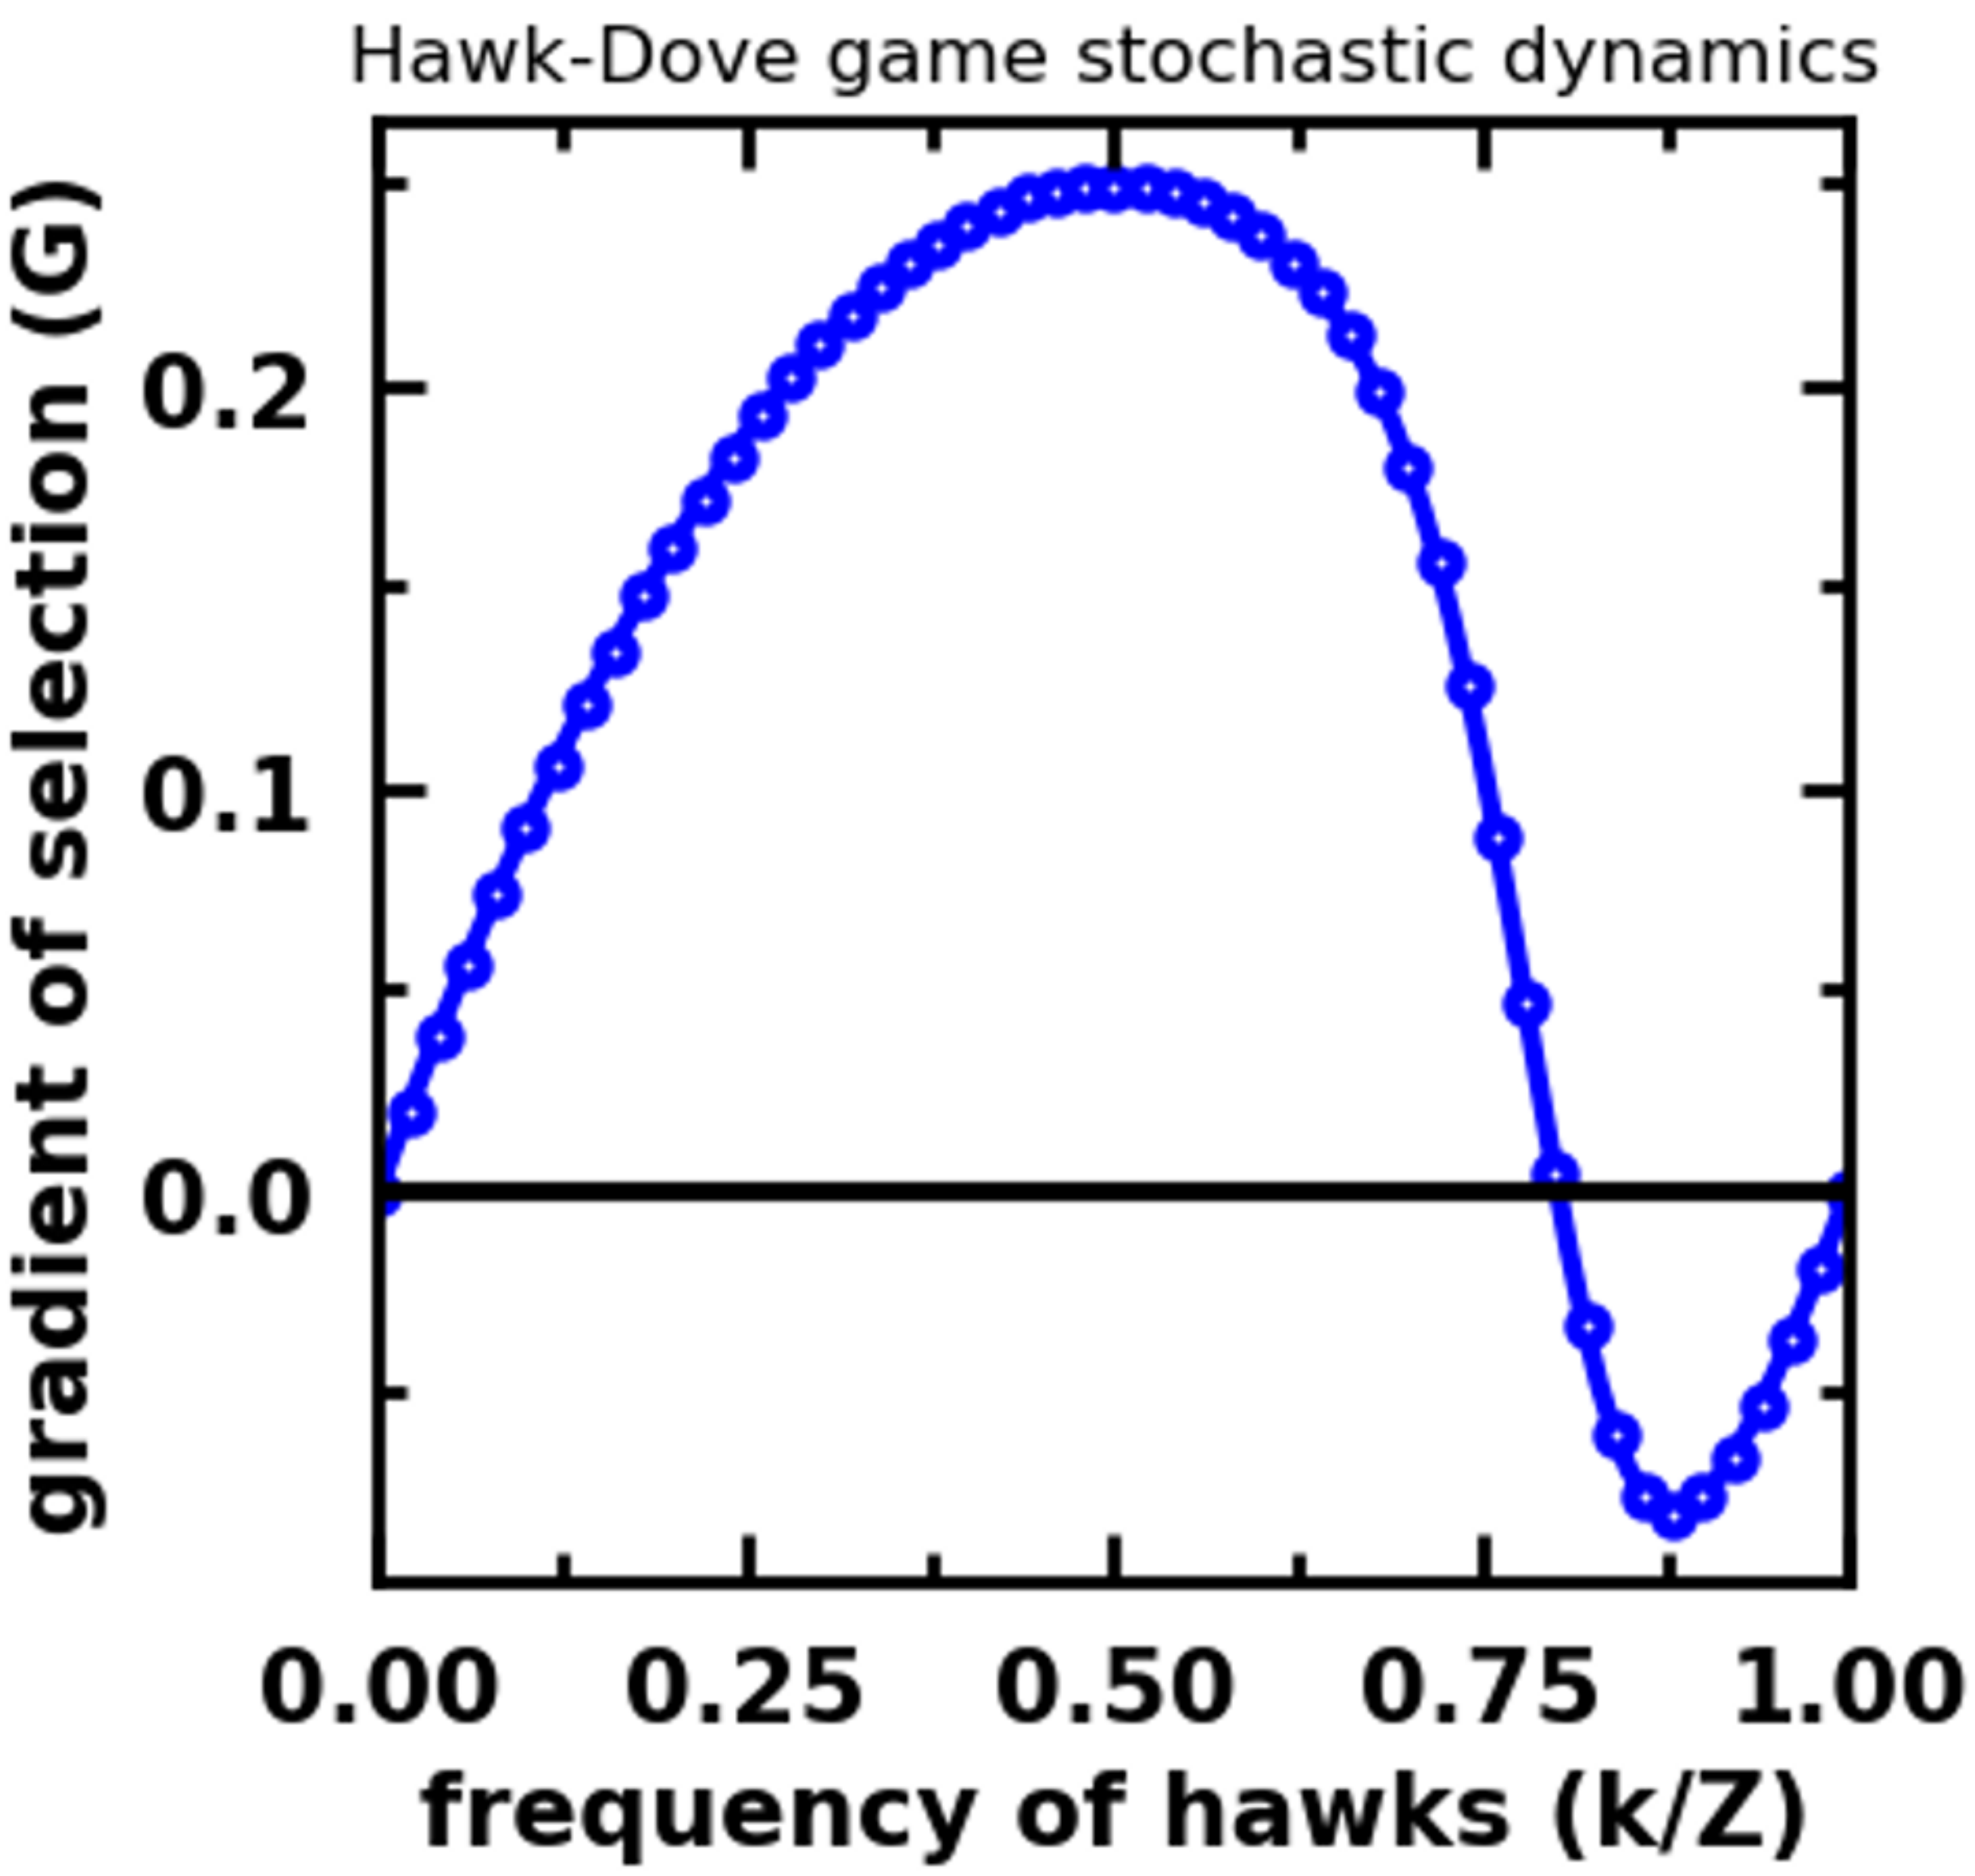

$V = 2,\ D = 3,\ T = 1\ (V < D)$

$$\begin{pmatrix} -0.5,\ -0.5 & 2,\ 0 \\ 0,\ 2 & 0,\ 0 \end{pmatrix}$$

$Z = 100$

$\beta = 10$

$\mu = 0$

$x_i \equiv [k_i/Z]$

Examples: Hawk-Dove: effect of β

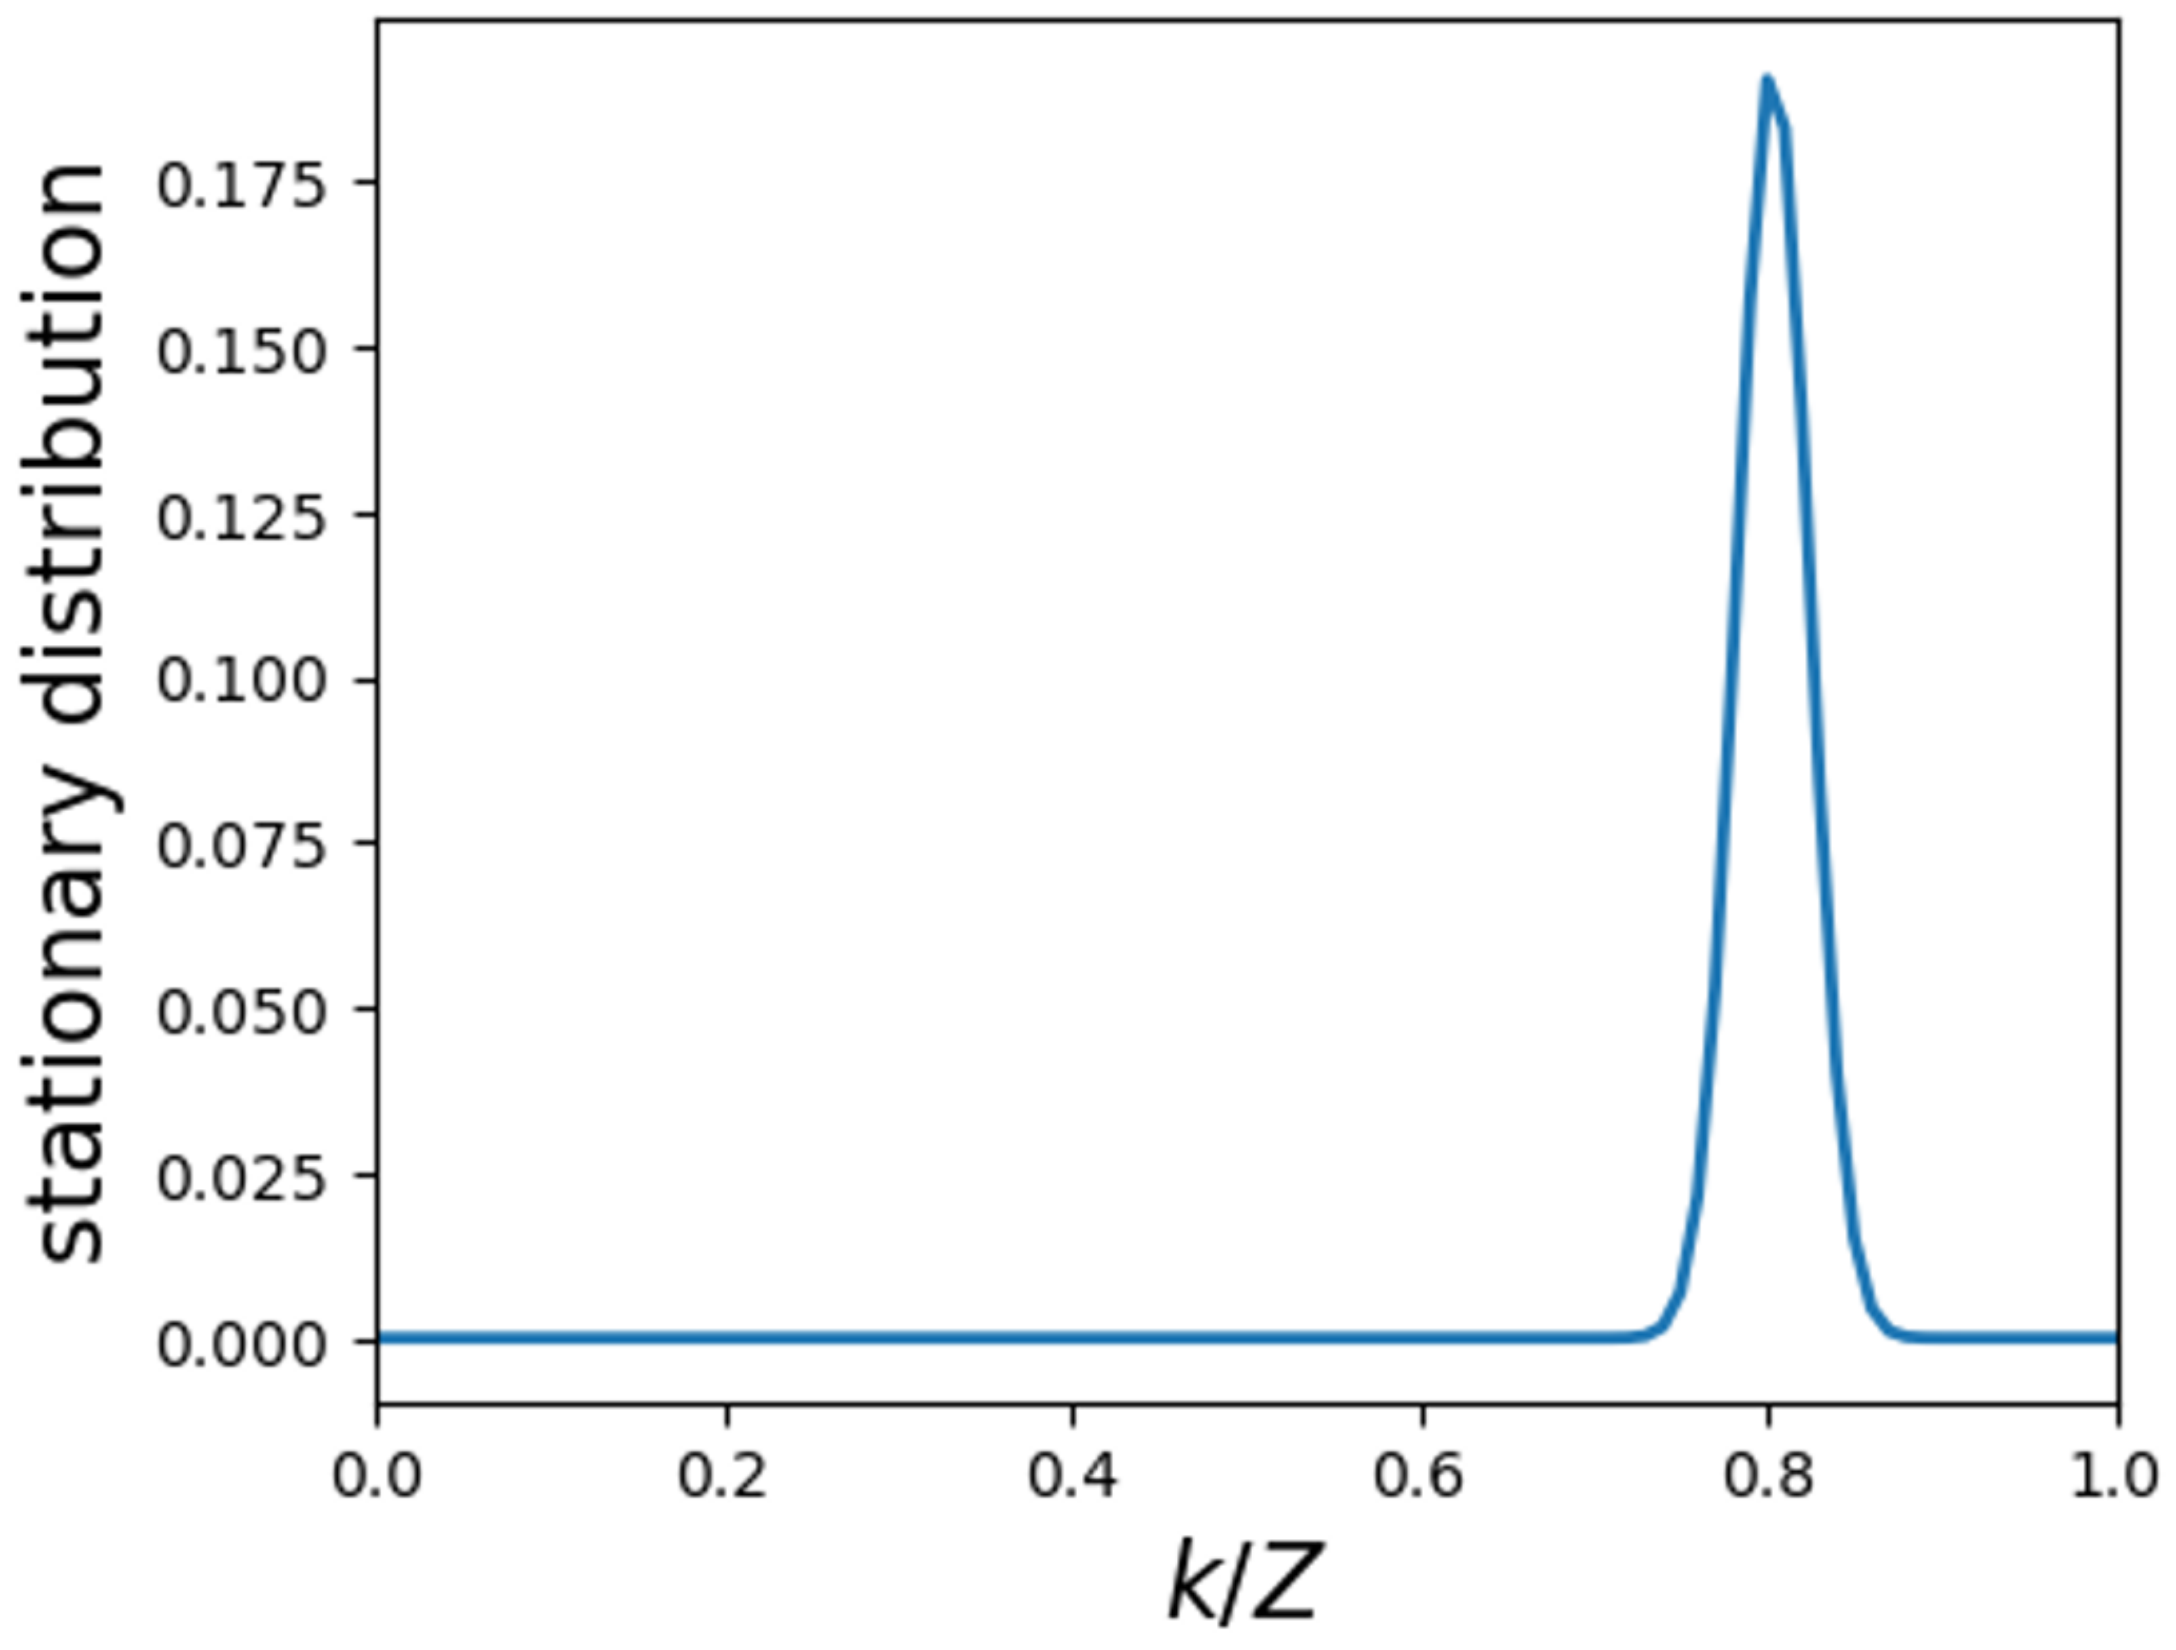

$V = 2,\ D = 3,\ T = 1\ (V < D)$

$$\begin{pmatrix} -0.5,\ -0.5 & 2,\ 0 \\ 0,\ 2 & 0,\ 0 \end{pmatrix}$$

$Z = 100$

$\beta = 10$

$\mu = 1e-3$

$x_i \equiv [k_i/Z]$

Examples: Hawk-Dove: effect of β

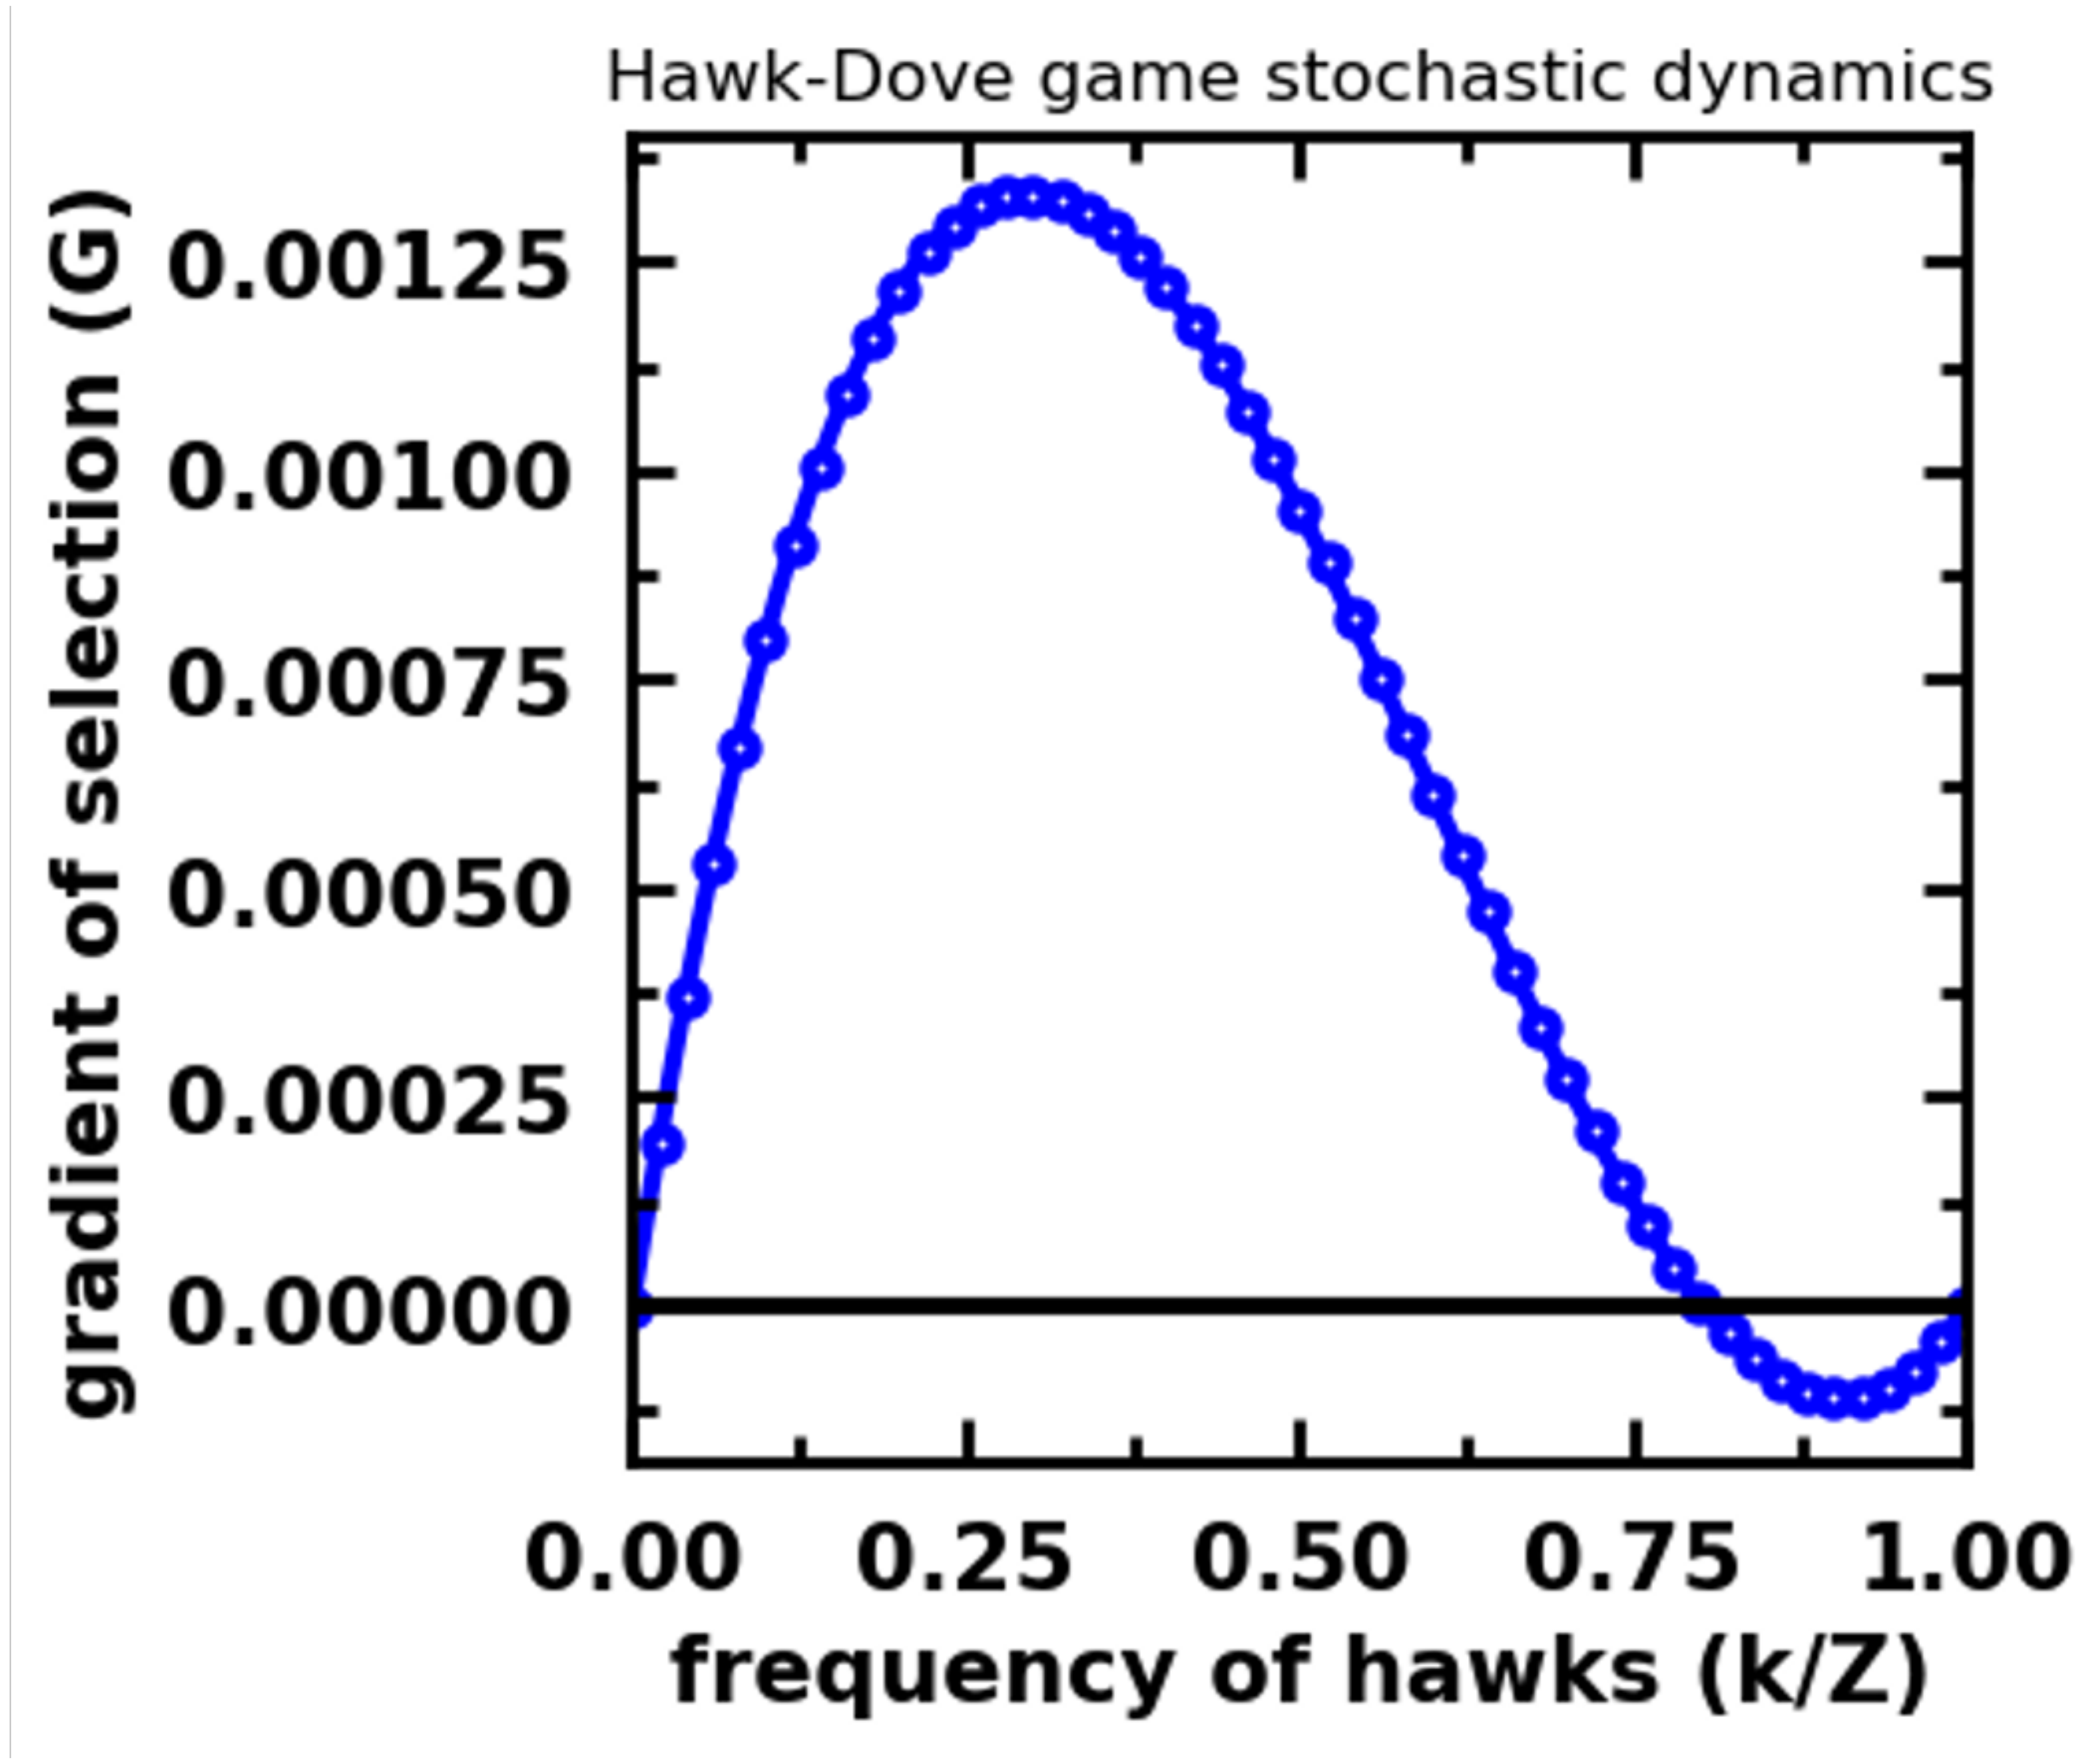

$V = 2,\ D = 3,\ T = 1\ (V < D)$

$$\begin{pmatrix} -0.5, & -0.5 & 2, & 0 \\ & 0, 2 & 0, & 0 \end{pmatrix}$$

$Z = 100$

$\beta = 1e - 2$

$\mu = 0$

$x_i \equiv [k_i/Z]$

Examples: Hawk-Dove: effect of μ

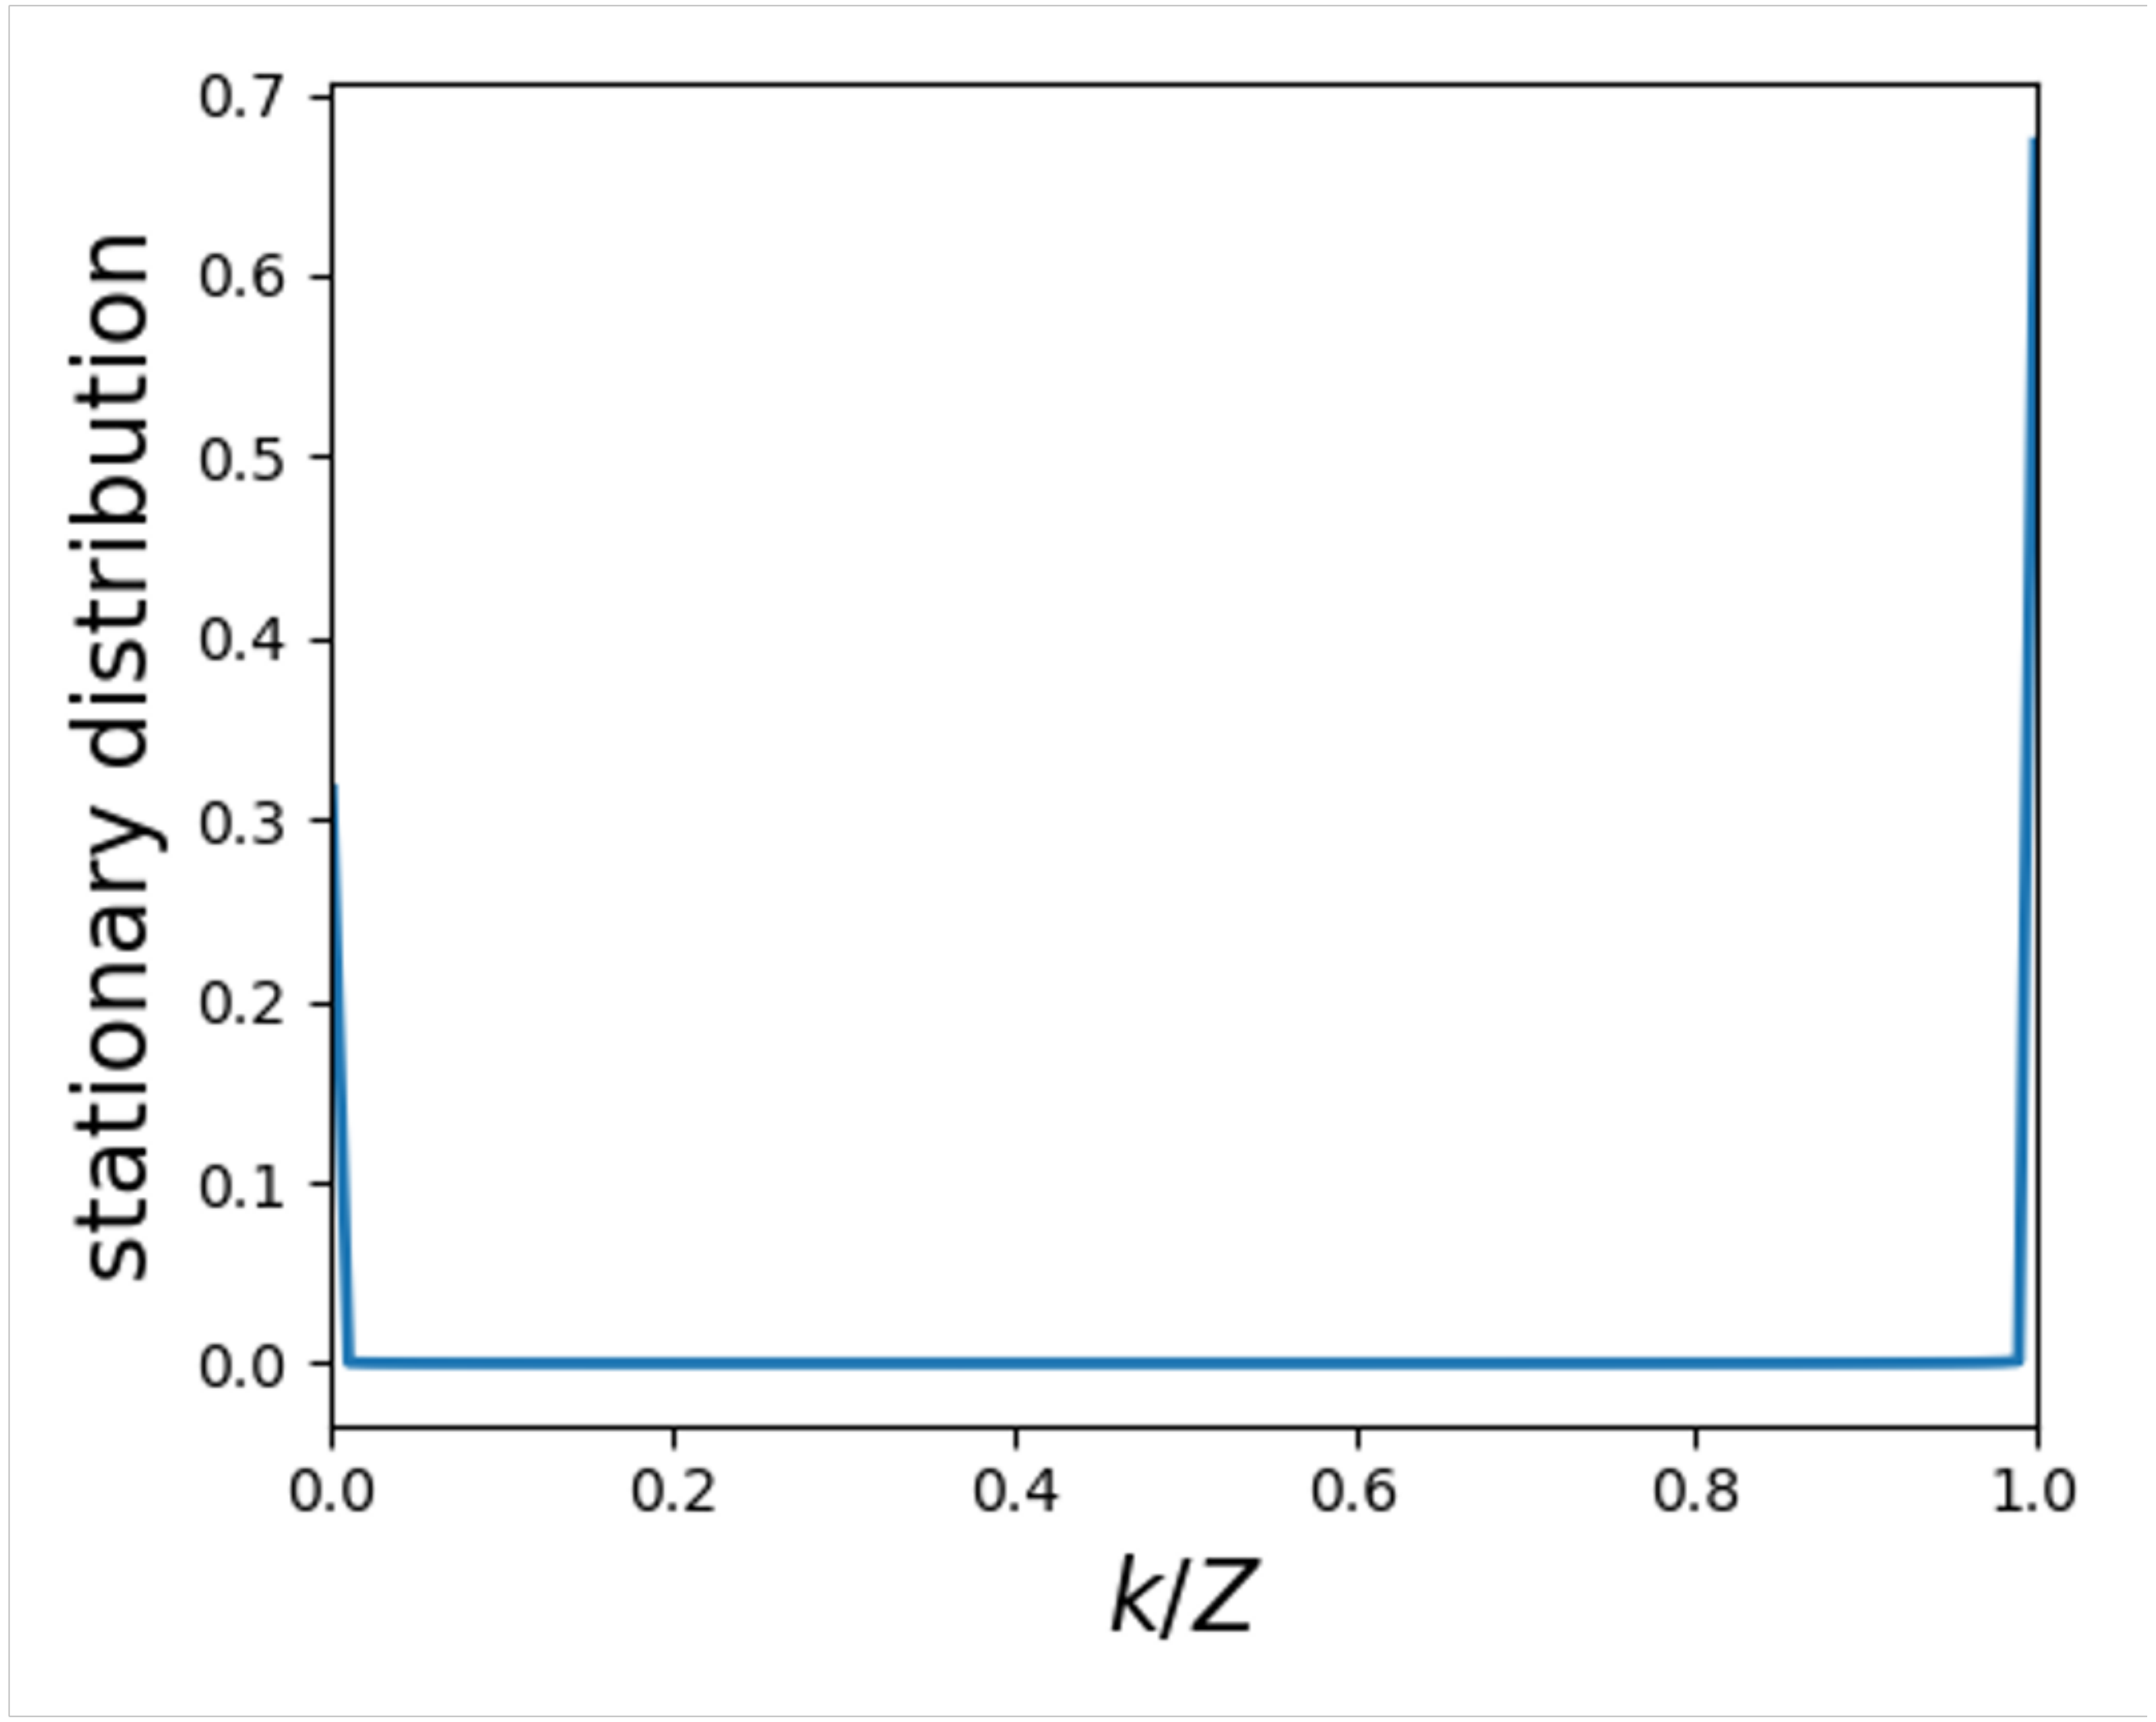

$V = 2,\ D = 3,\ T = 1\ (V < D)$

$$\begin{pmatrix} -0.5,\ -0.5 & 2,\ 0 \\ 0,\ 2 & 0,\ 0 \end{pmatrix}$$

$Z = 100$

$\beta = 1$

$\mu = 1e - 5$

$x_i \equiv [k_i/Z]$

Examples: Hawk-Dove: effect of μ

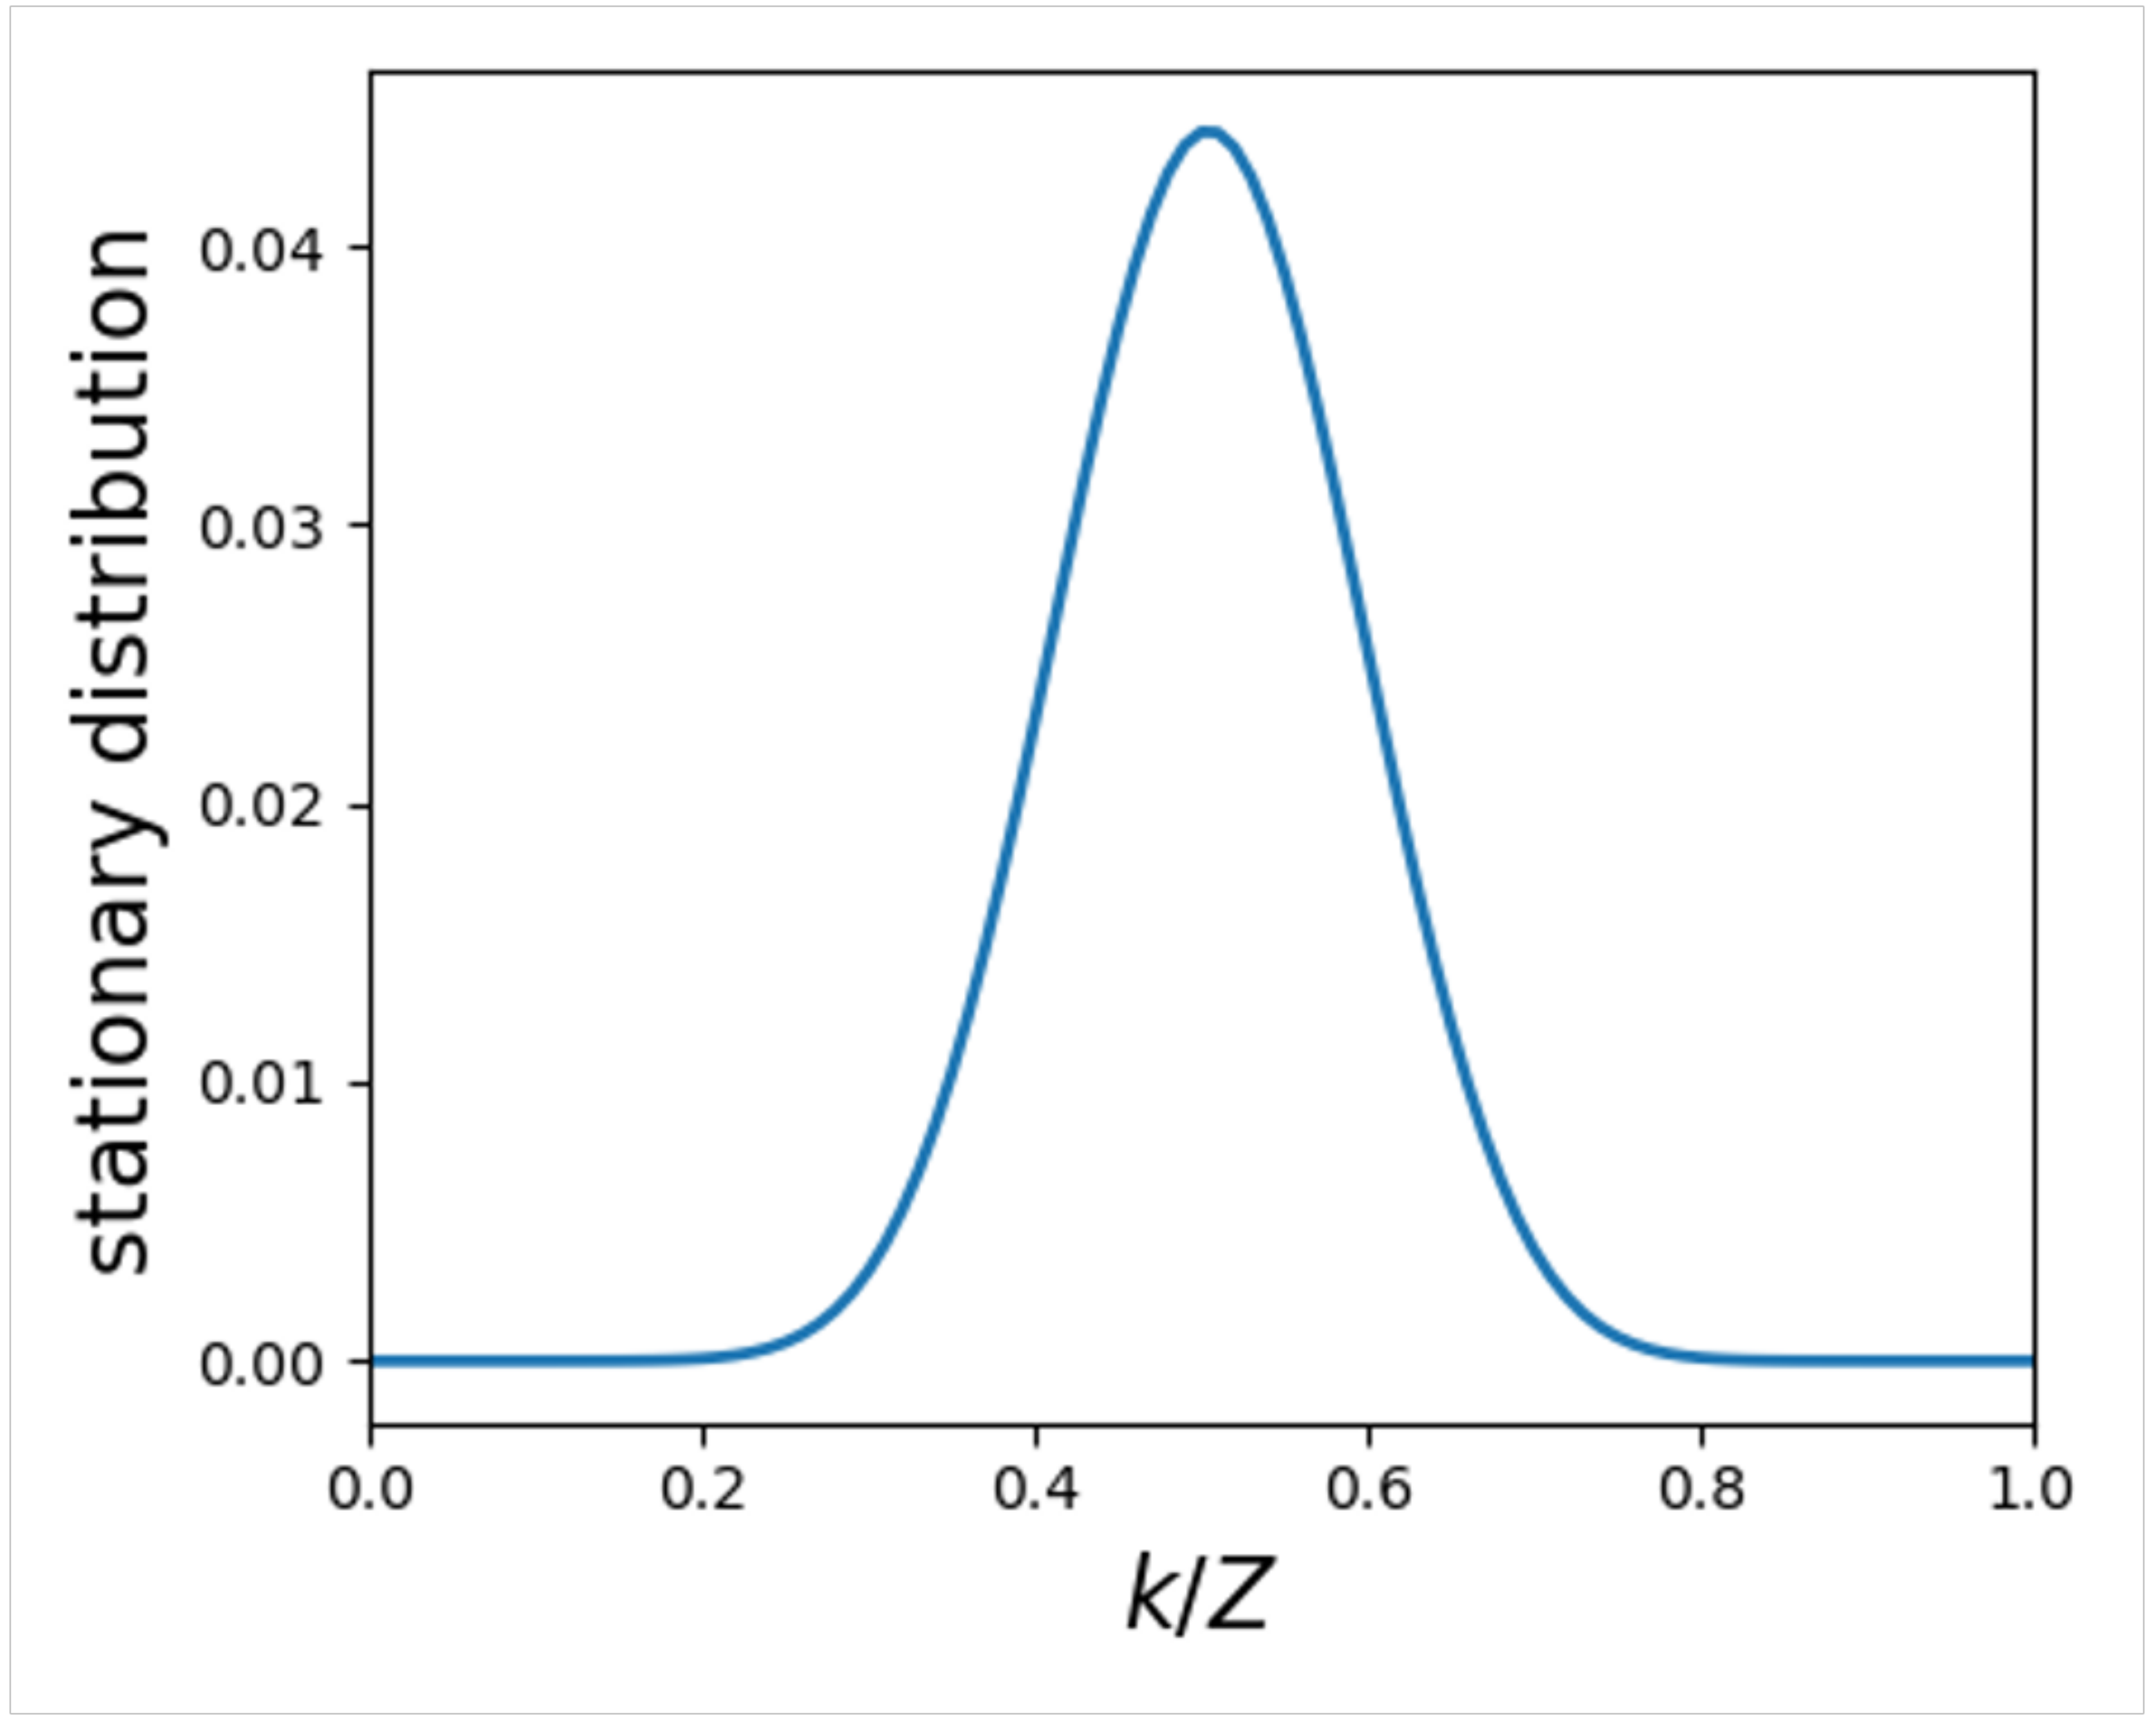

$V = 2,\ D = 3,\ T = 1\ (V < D)$

$$\begin{pmatrix} -0.5,\ -0.5 & 2,\ 0 \\ 0,\ 2 & 0,\ 0 \end{pmatrix}$$

$Z = 100$

$\beta = 1$

$\mu = 1e - 1$

$x_i \equiv [k_i / Z]$

Examples

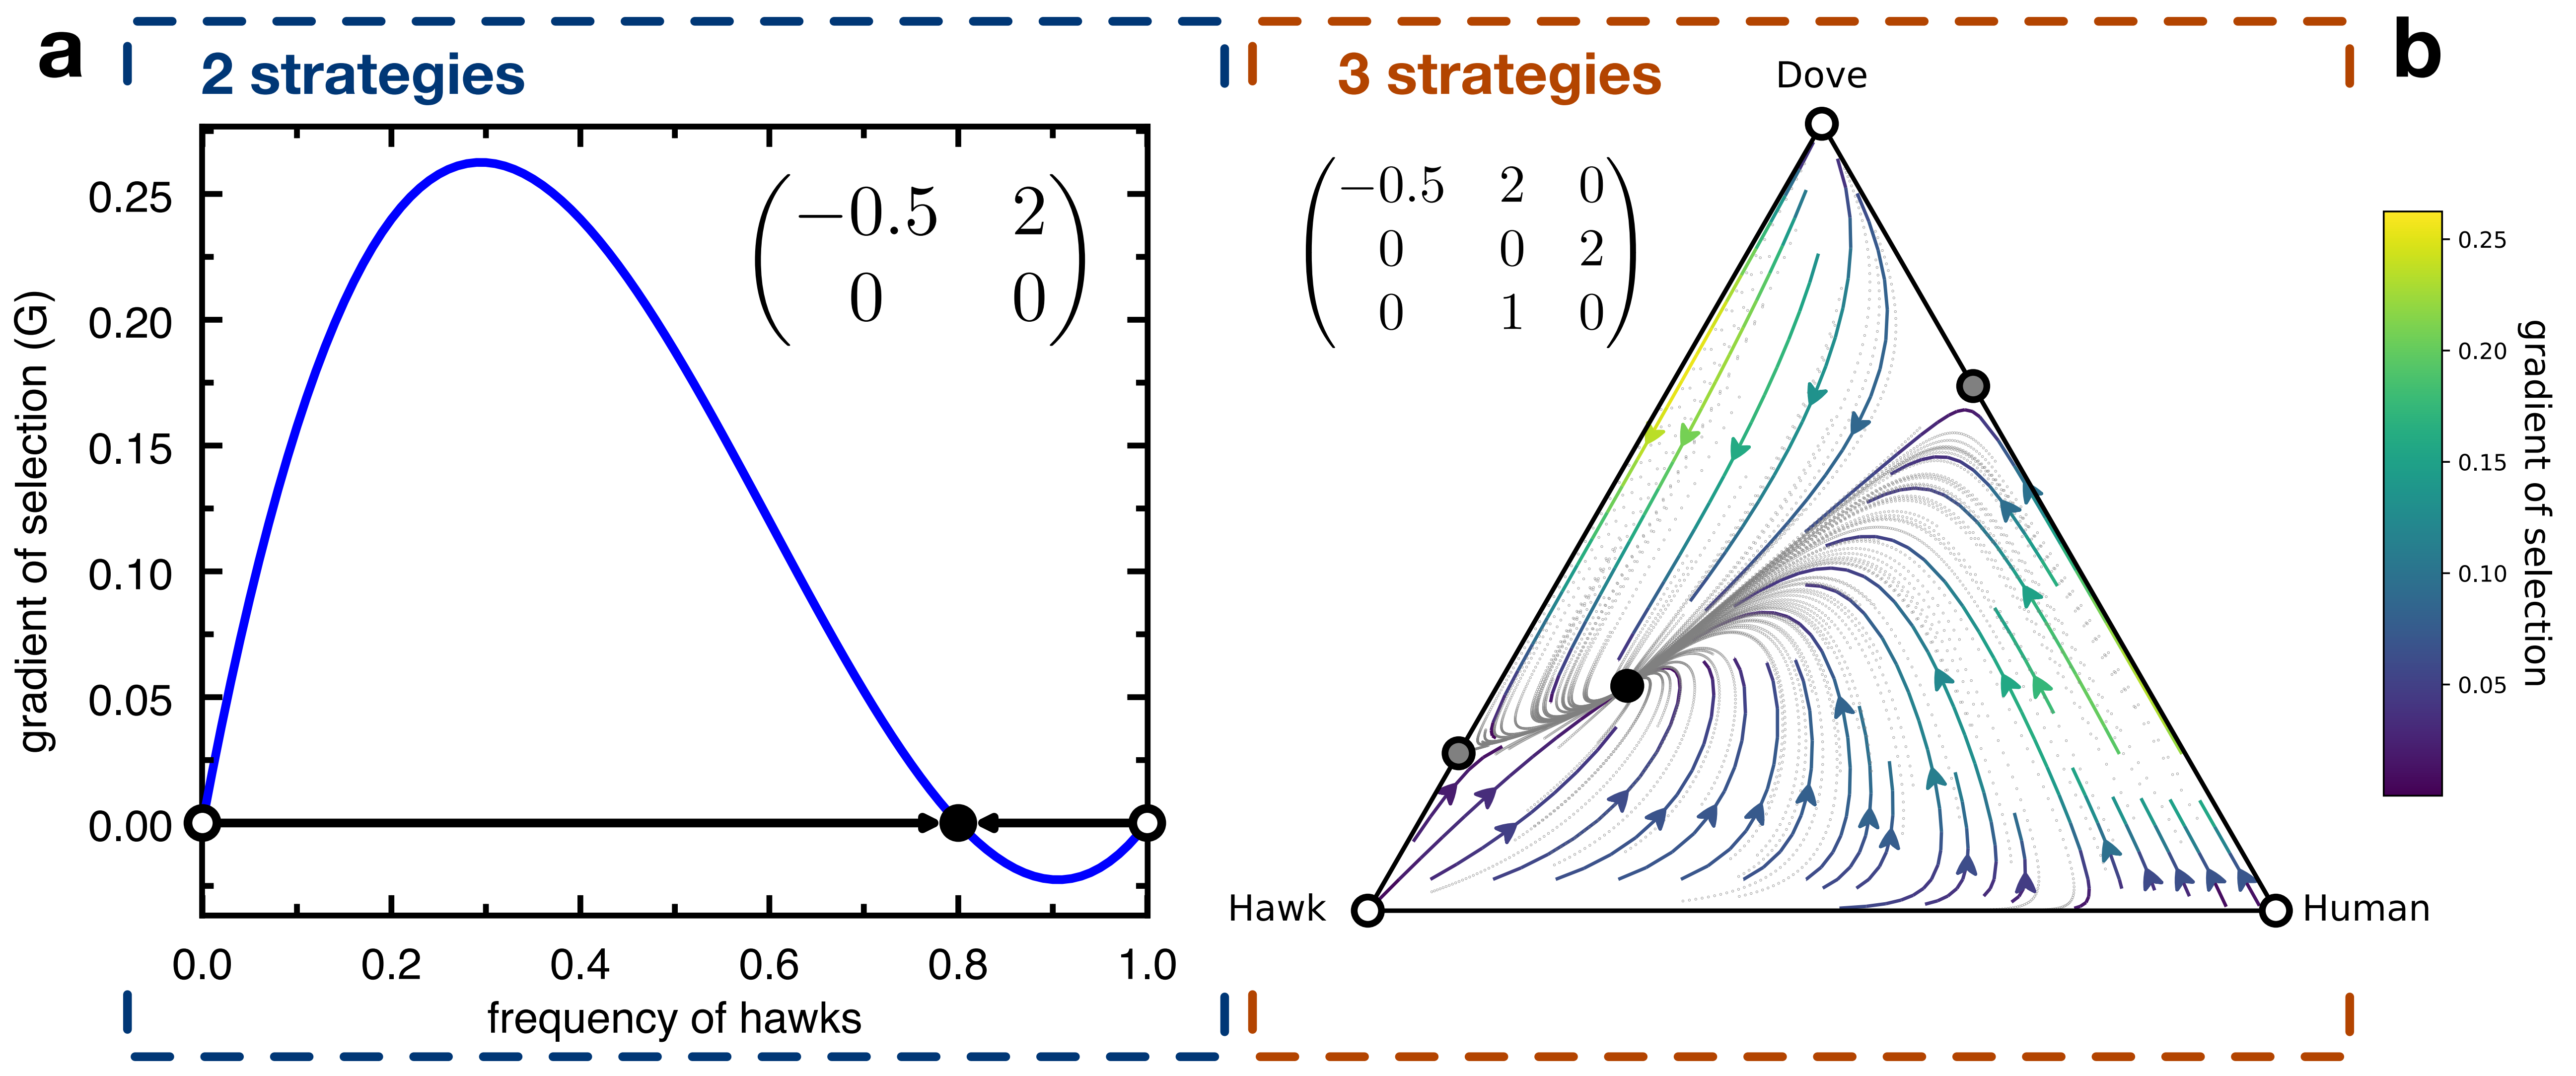

a
2 strategies
$\begin{pmatrix} -0.5 & 2 \\ 0 & 0 \end{pmatrix}$
gradient of selection (G)
frequency of hawks
b
3 strategies
$\begin{pmatrix} -0.5 & 2 & 0 \\ 0 & 0 & 2 \\ 0 & 1 & 0 \end{pmatrix}$
Dove
Hawk
Human
gradient of selection

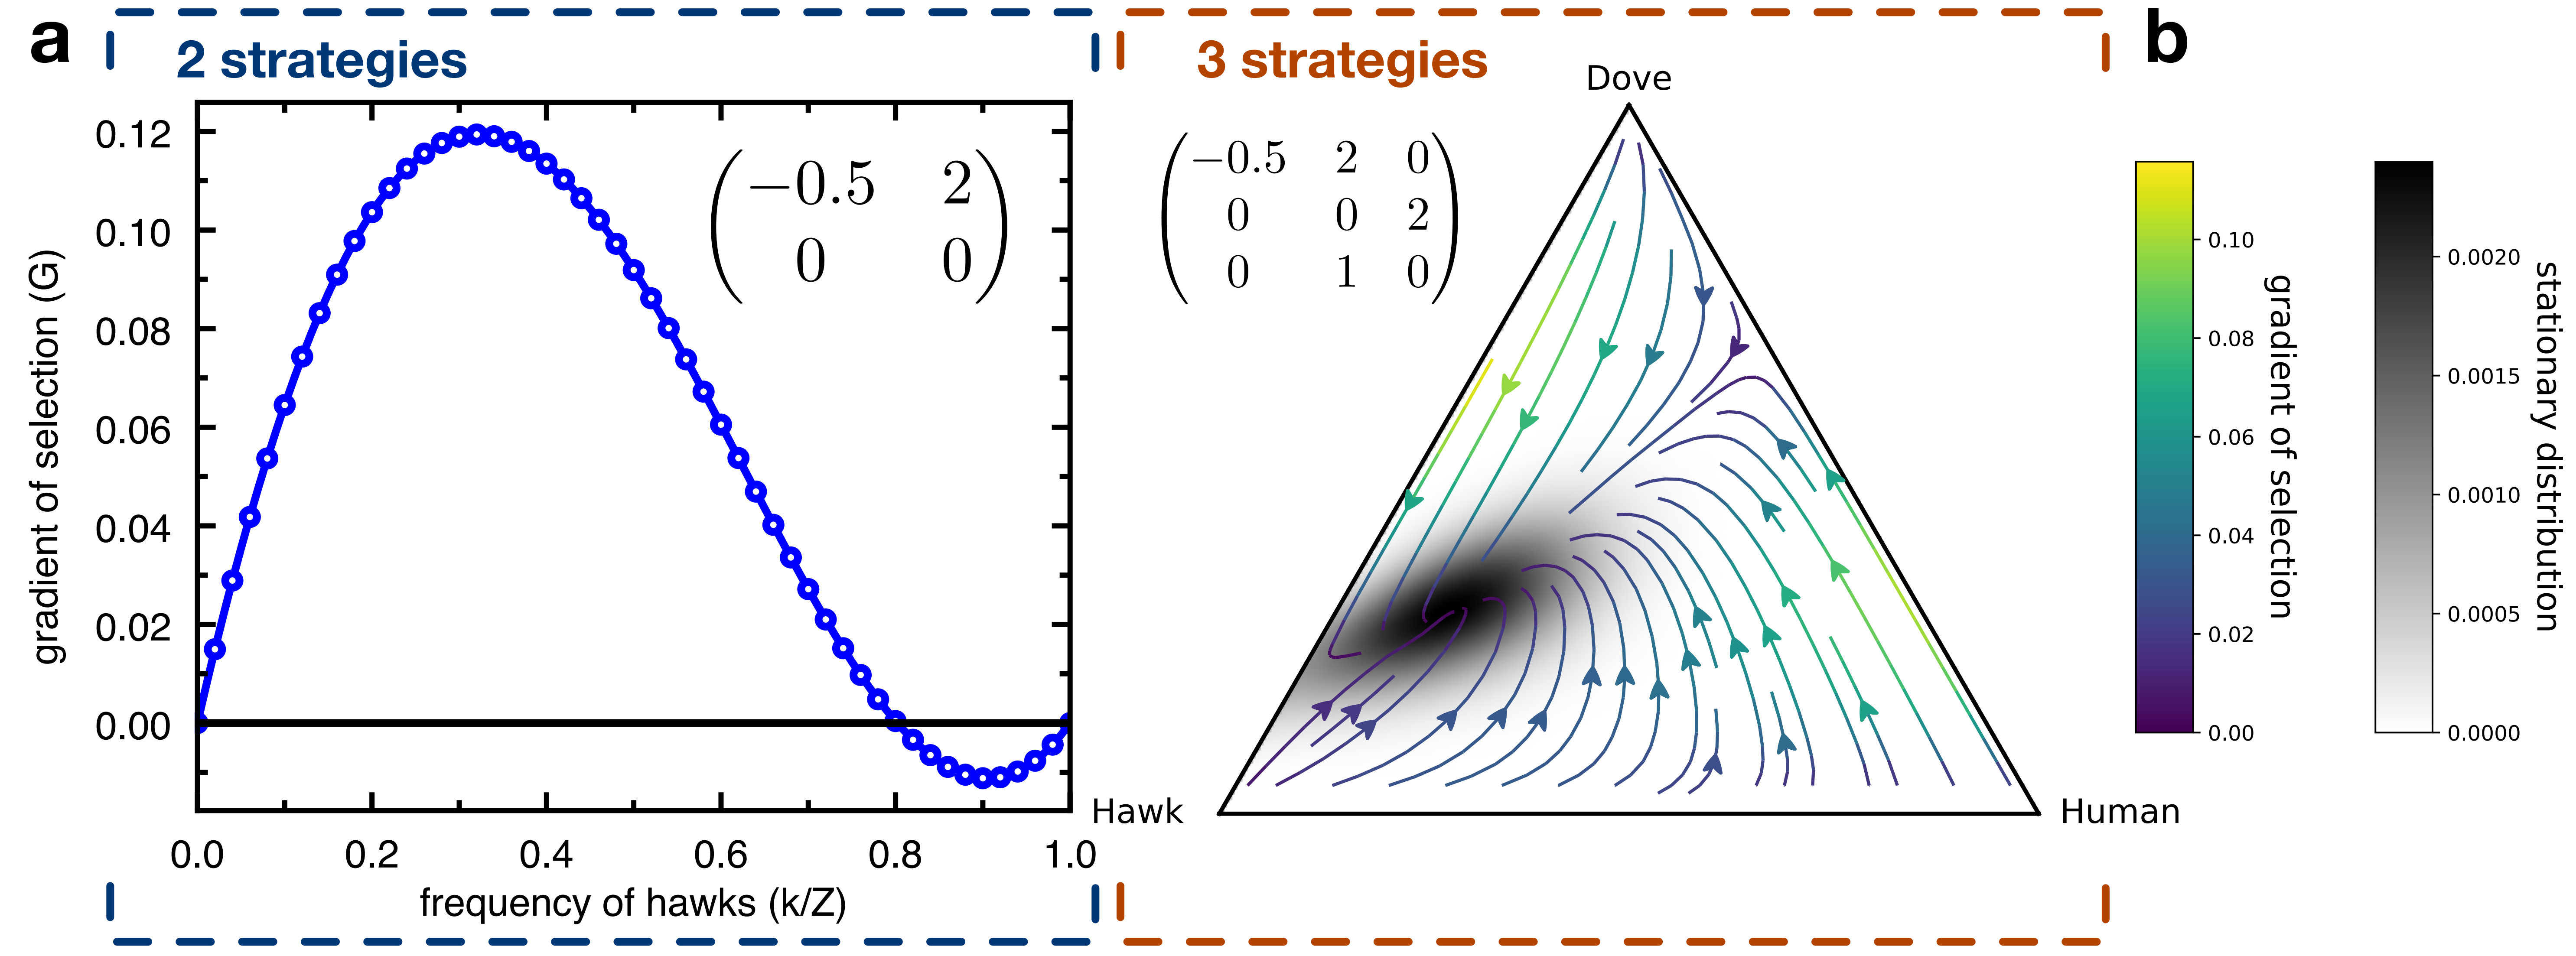

a
2 strategies
$\begin{pmatrix} -0.5 & 2 \\ 0 & 0 \end{pmatrix}$
gradient of selection (G)
frequency of hawks (k/Z)
b
3 strategies
$\begin{pmatrix} -0.5 & 2 & 0 \\ 0 & 0 & 2 \\ 0 & 1 & 0 \end{pmatrix}$
Dove
Hawk
Human
gradient of selection
stationary distribution

Fixation probabilities

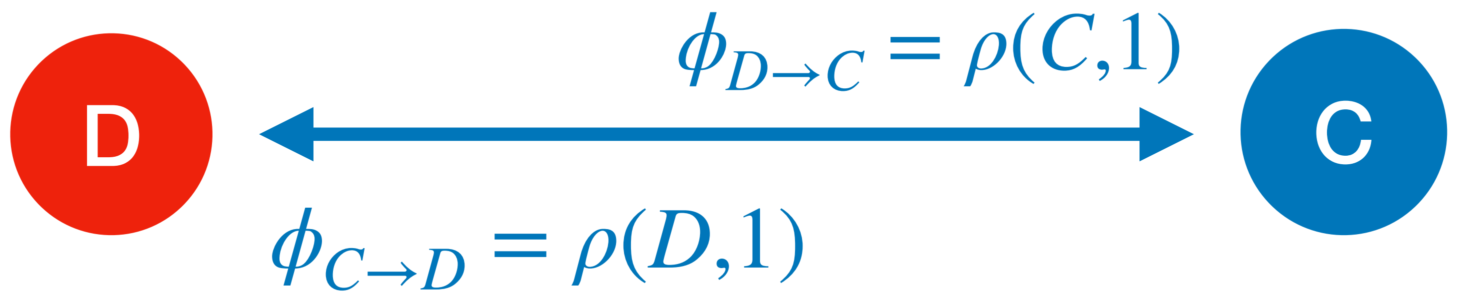

We want $\rho(C,1)$, the probability that one C can take over a population of D players, which is given by

$$\rho(1,C) = T^{-}(1)\rho(C,0) + T^{+}(1)\rho(C,2) + (1 - T^{-}(1) - T^{+}(1))\rho(C,1)$$

$$\rho(1,C) = \left(\sum_{k=0}^{Z-1} \prod_{i=1}^{k} \frac{T^{-}(i)}{T^{+}(i)} \right)^{-1}$$

Fixation probabilities, evolutionary robustness and risk dominance

Fixation probability (of a single mutant)

$$\rho_{ij} = \left(1 + \sum_{m=1}^{Z-1} \prod_{k=1}^{m} \frac{T^{-}(k)}{T^{+}(k)}\right)^{-1}$$

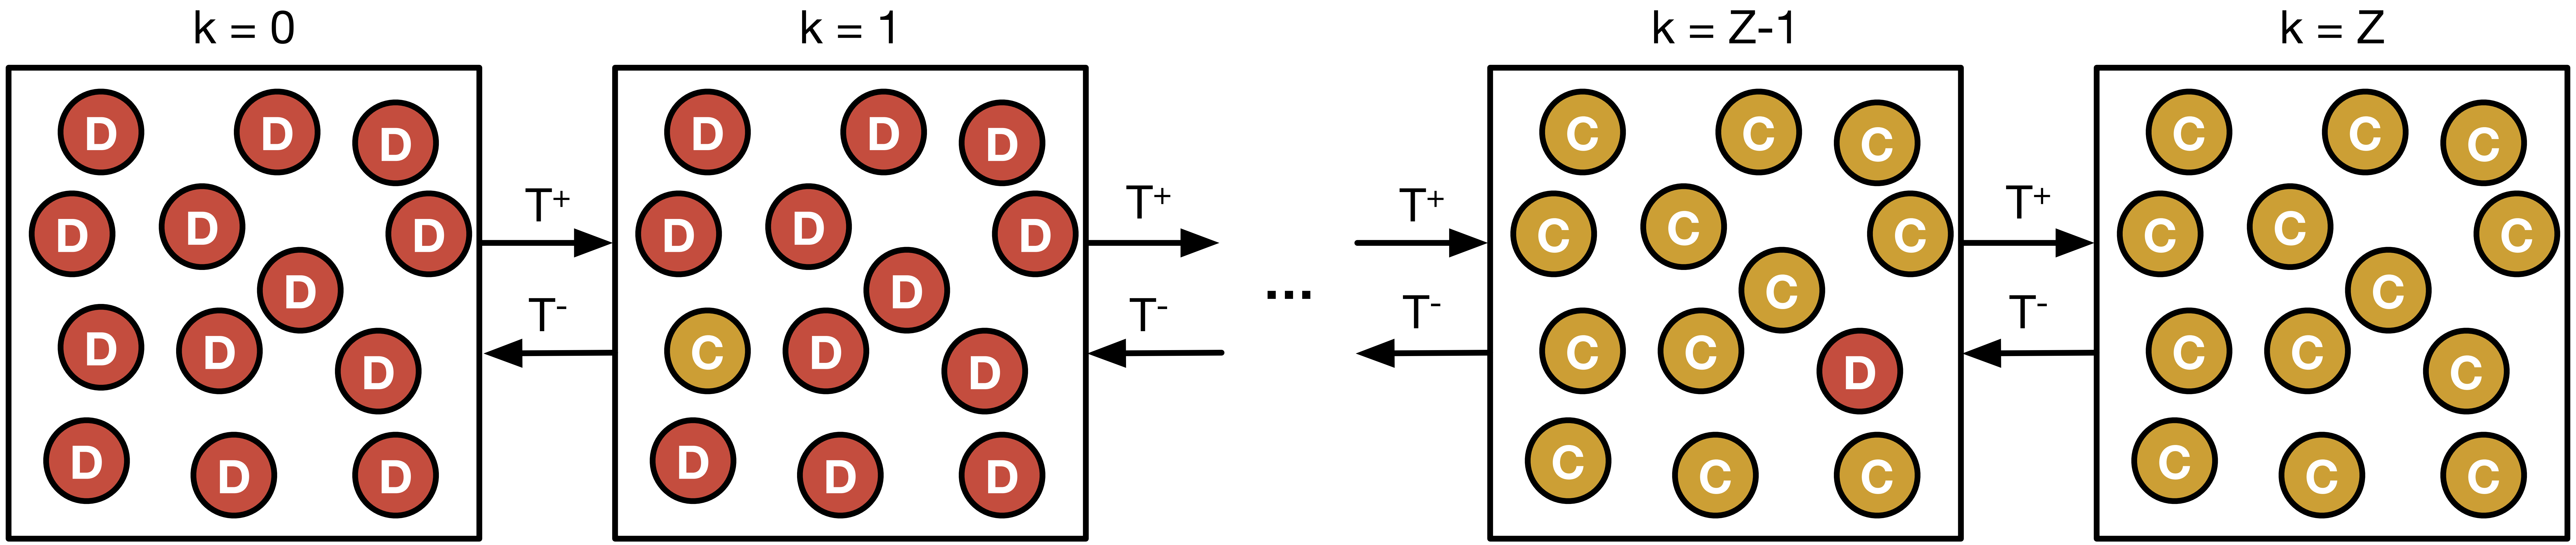

Fixation probabilities and evolutionary robustness

Neutral drift

"*In a finite, homogeneous population of size Z, a newly introduced **neutral mutation** (i.e., a mutation that does not change the payoff to either player) will eventually replace the entire population with probability* $\rho = \frac{1}{Z}$."

Survival of the fittest

A **deleterious mutation**, which is opposed by selection, will fix with probability $\rho < \frac{1}{Z}$

An **advantageous mutation**, which is favored by selection, will fix with probability $\rho > \frac{1}{Z}$.

Stewart, A. J., & Plotkin, J. B. (2013). From extortion to generosity, evolution in the iterated prisoner's dilemma. Proceedings of the National Academy of Sciences, 110(38), 15348-15353.

Solution concepts ; Evolutionary Robustness

Prisoners Dilemma, T>R, P>S

Greed and fear

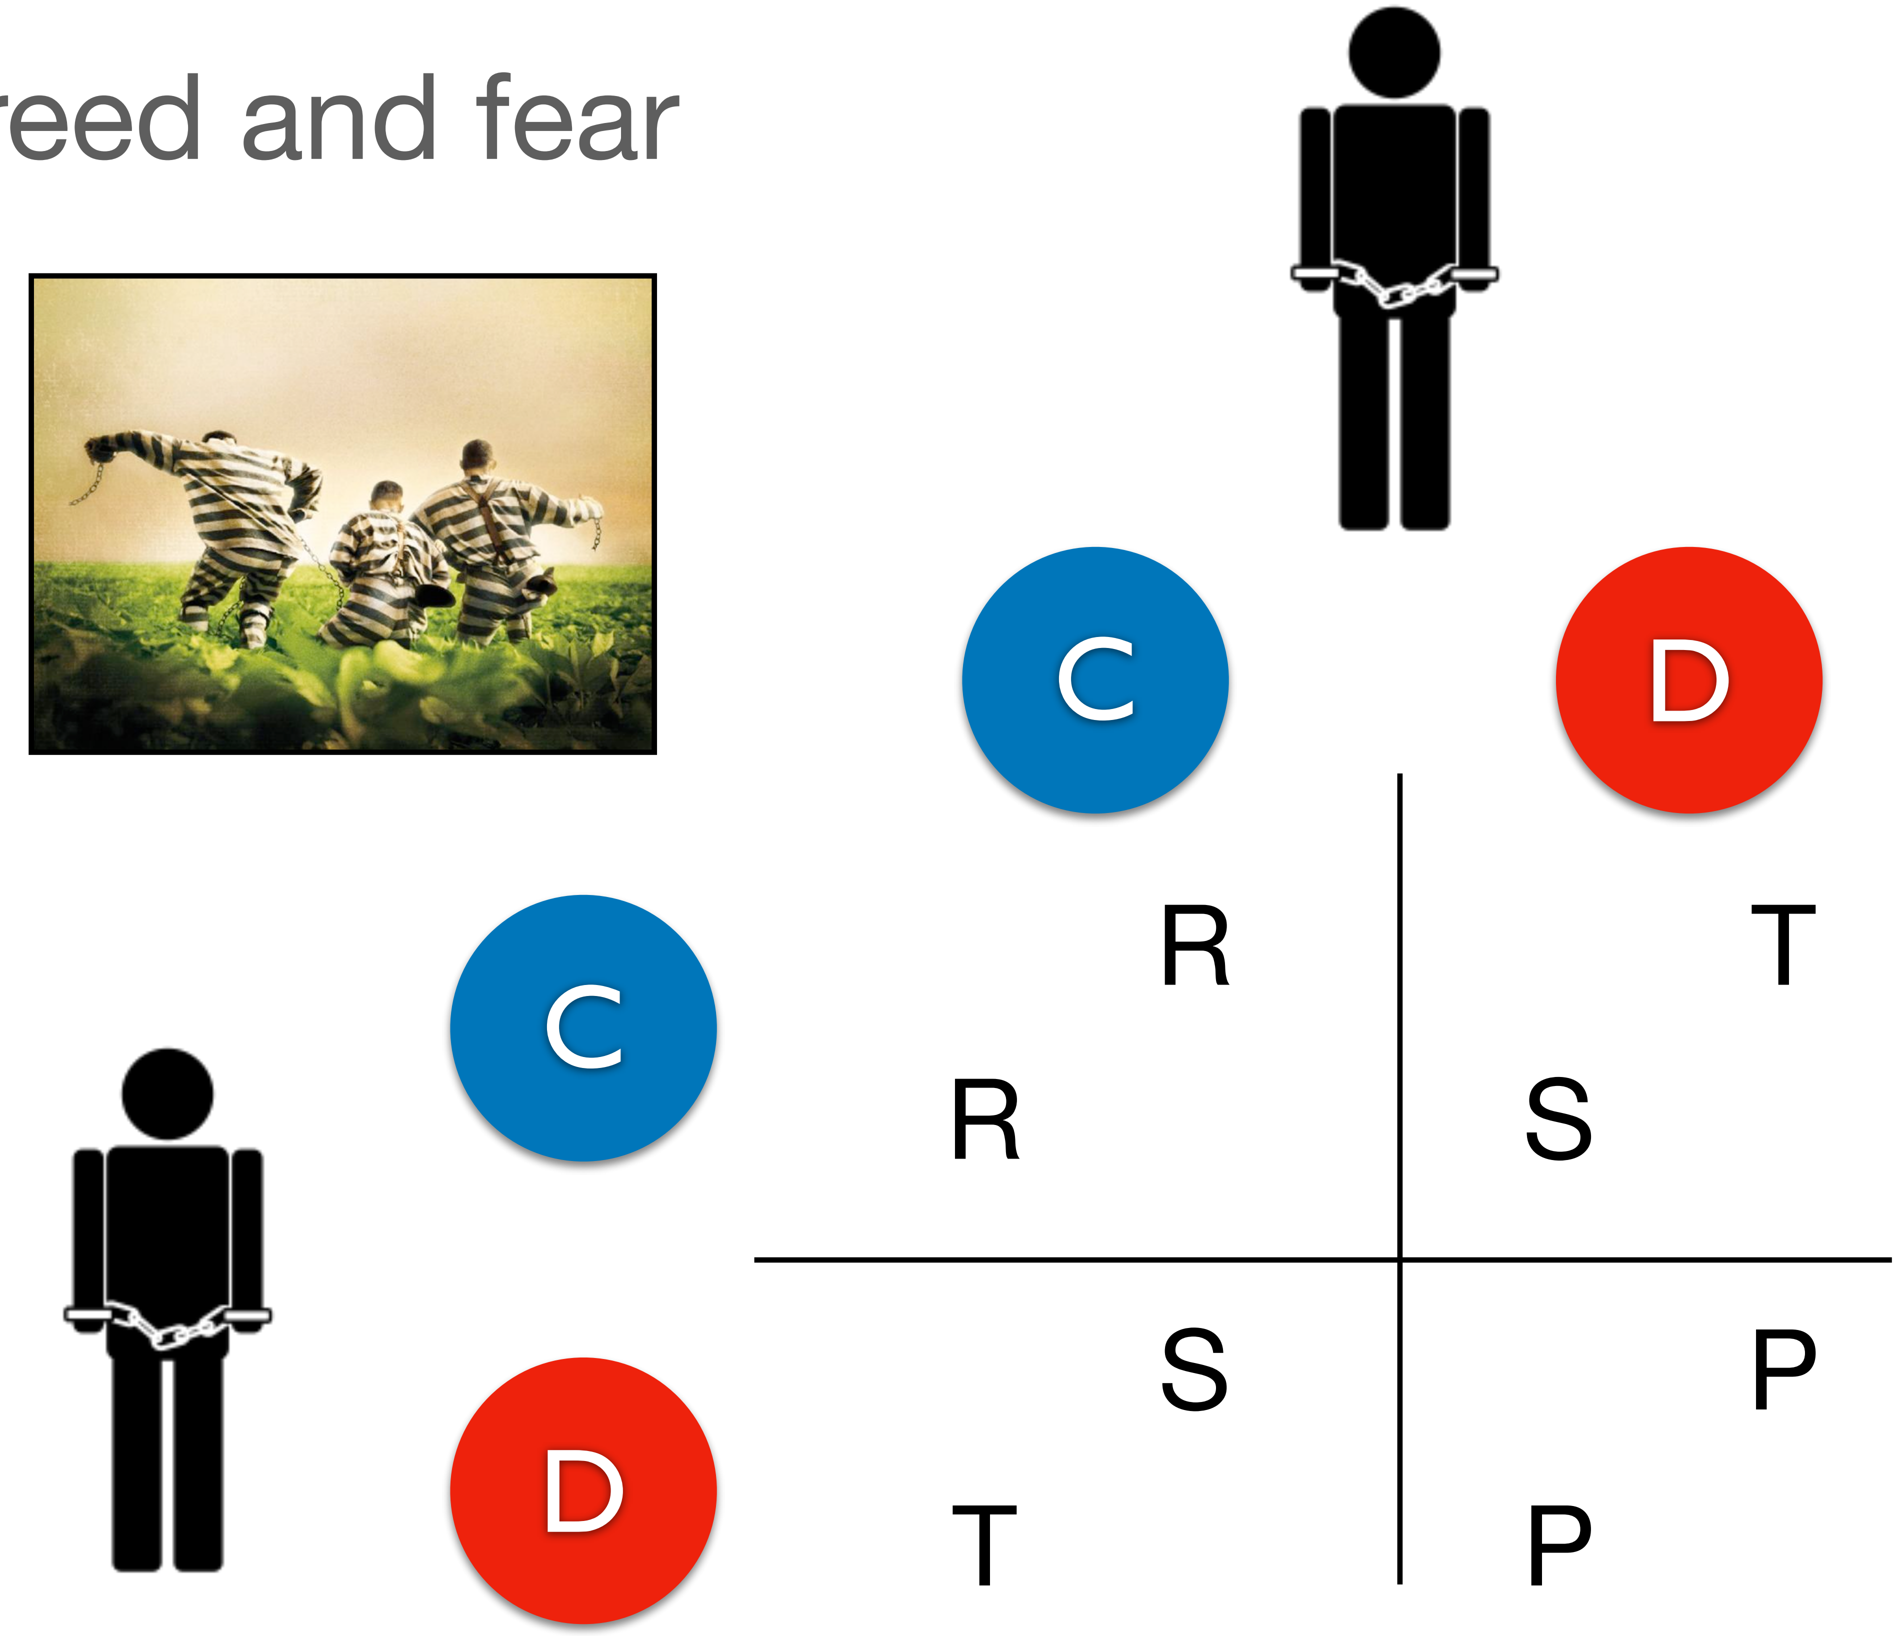

Remember a ***neutral mutation*** *can replace the entire population with probability* $\rho = \frac{1}{Z}$.

A strategy s^* is **Evolutionary Robust** against a mutant strategy s' if the latter has a fixation probability of $\rho \leq \frac{1}{Z}$ in a population with s^*

In terms of the reduced Markov chain, there are no edges leaving the state s^* with $\rho > \frac{1}{Z}$

In the limit of $Z \rightarrow \infty$ the condition reduces to the ESS condition

Stewart, A. J., & Plotkin, J. B. (2013). From extortion to generosity, evolution in the iterated prisoner's dilemma. Proceedings of the National Academy of Sciences, 110(38), 15348-15353.

But, as this can quickly become computationally intractable!

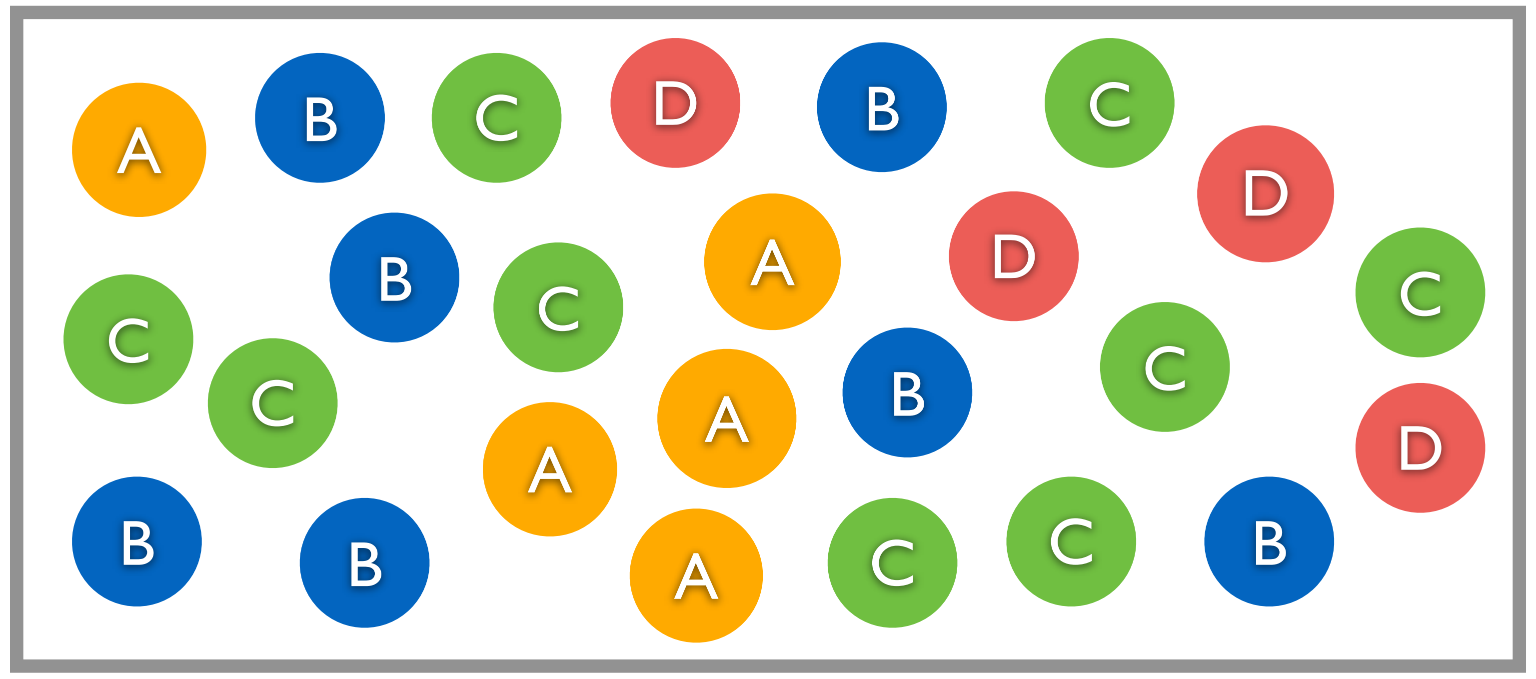

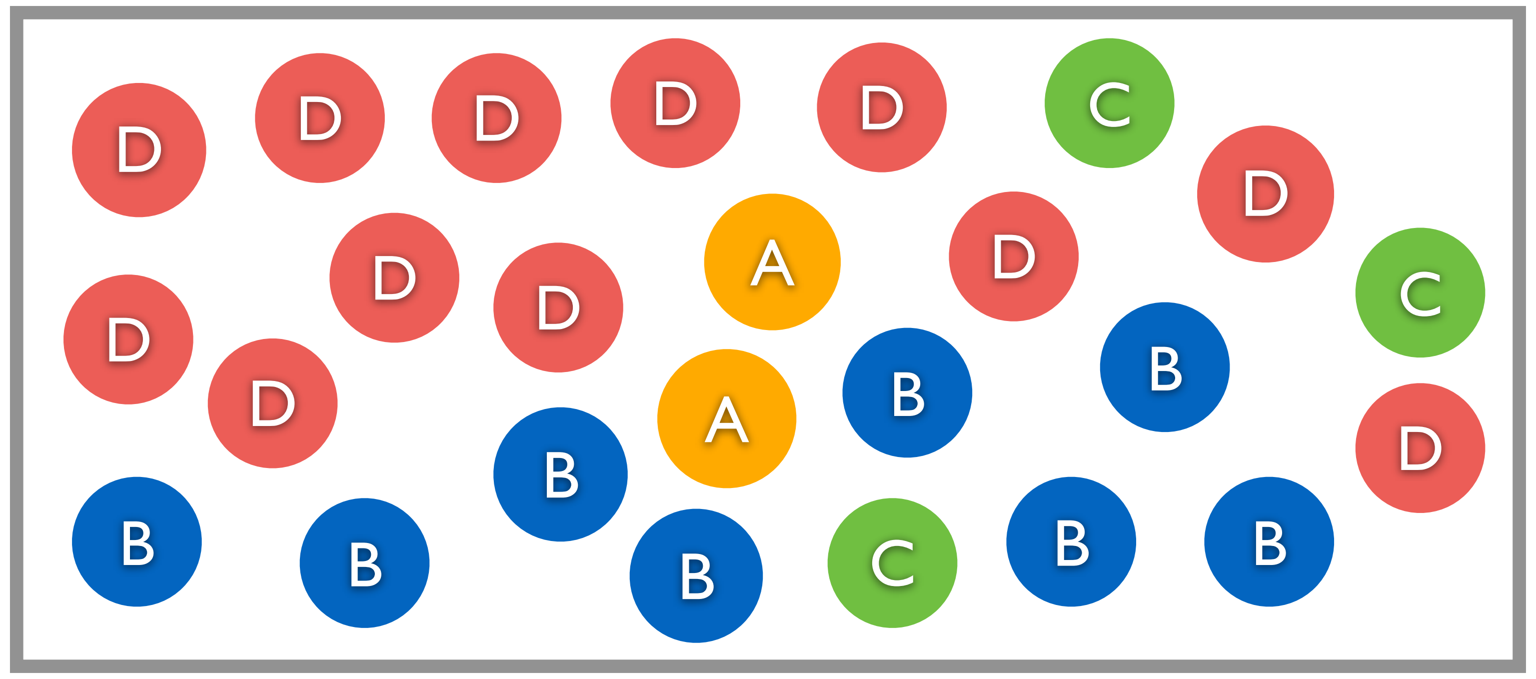

Evolutionary dynamics (Φ)

Replication dynamics

$\dot{x}_i = x_i\,(\pi_i - \langle\pi\rangle)$

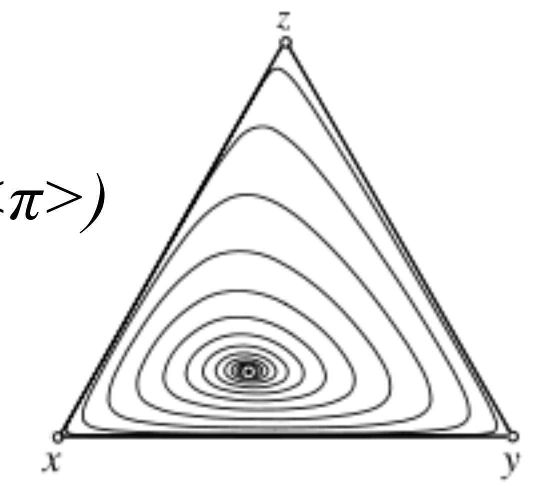

(Nonlinear) dynamical systems

N→∞

Moran process
Wright-Fisher process

Agent-based simulation

N finite

μ << 1

(stochastic)
Small mutation approximation

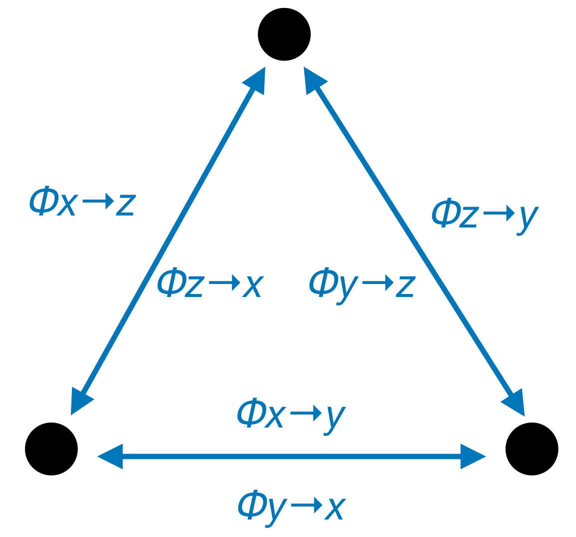

(Reduced) Markov chains

Adopted from Arne Traulsen

Small mutation limit (SML) and dominance between strategies

Under the assumption that **mutations are rare ($\mu \to 0$)** we always end up in a **monomorphic state**

This allows us to simplify the transition probabilities to

$$T^{\pm}(k) = \frac{k}{Z}\frac{Z-k}{Z}[1 + e^{\mp\beta(f_a - f_b)}]^{-1}$$

And to consider instead the **reduced Markov chain** which only **contains the vertices of the simplex**, so that the system is now characterised by the following transition probabilities:

$$T_{i,j} = \rho_{ij}/(n_s - 1) \qquad T_{i,i} = 1 - \sum_{\forall j} T_{i,j}$$

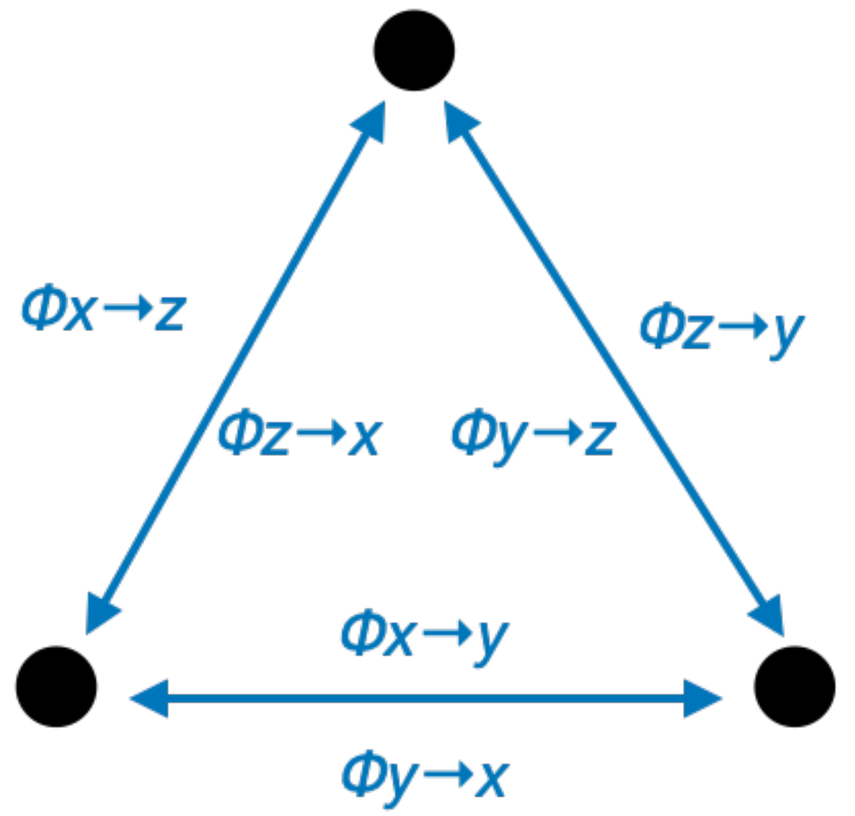

Example: Rock-Paper-Scissors

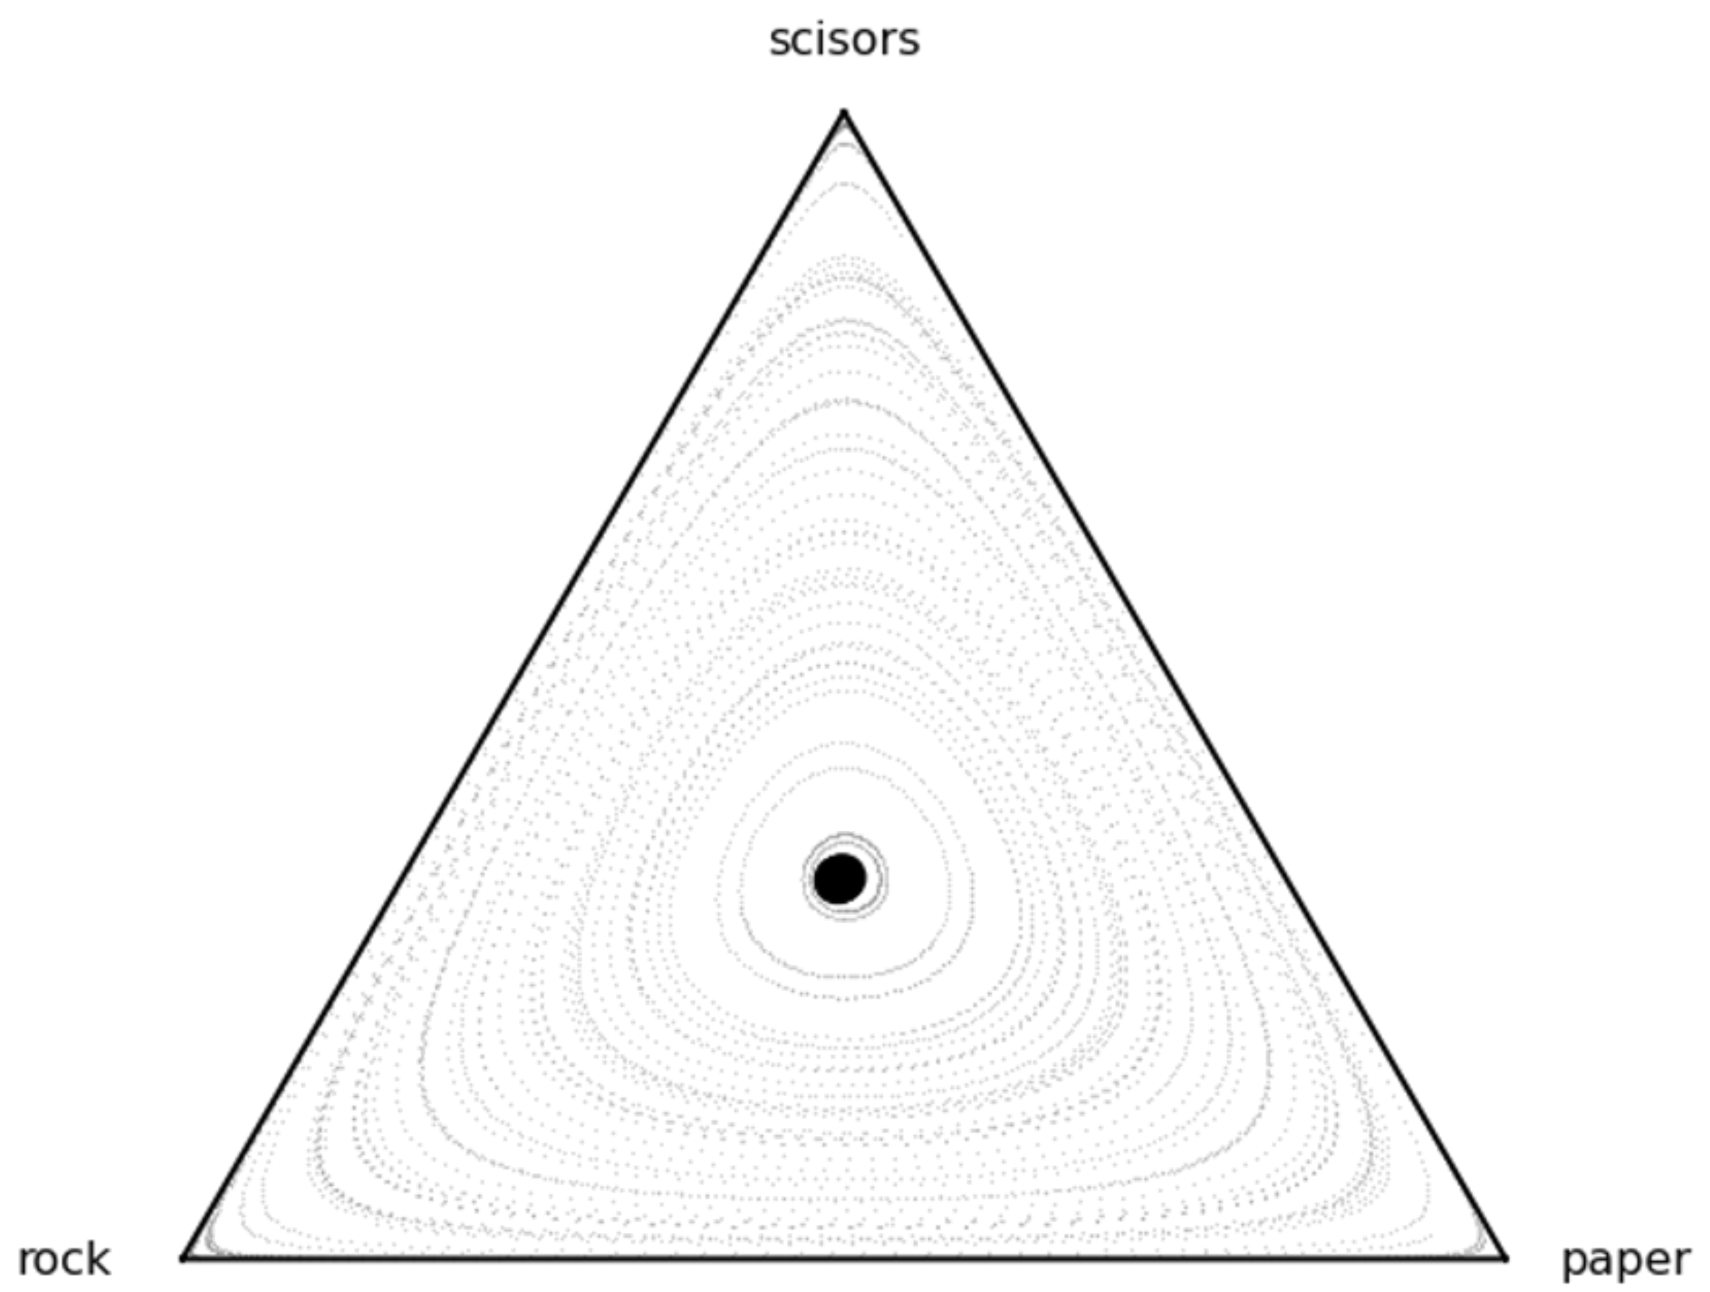

$$\begin{pmatrix} 0.5, 0.5 & 0, 1 & 1, 0 \\ 1, 0 & 0.5, 0.5 & 0, 1 \\ 0, 1 & 1, 0 & 0.5, 0.5 \end{pmatrix}$$

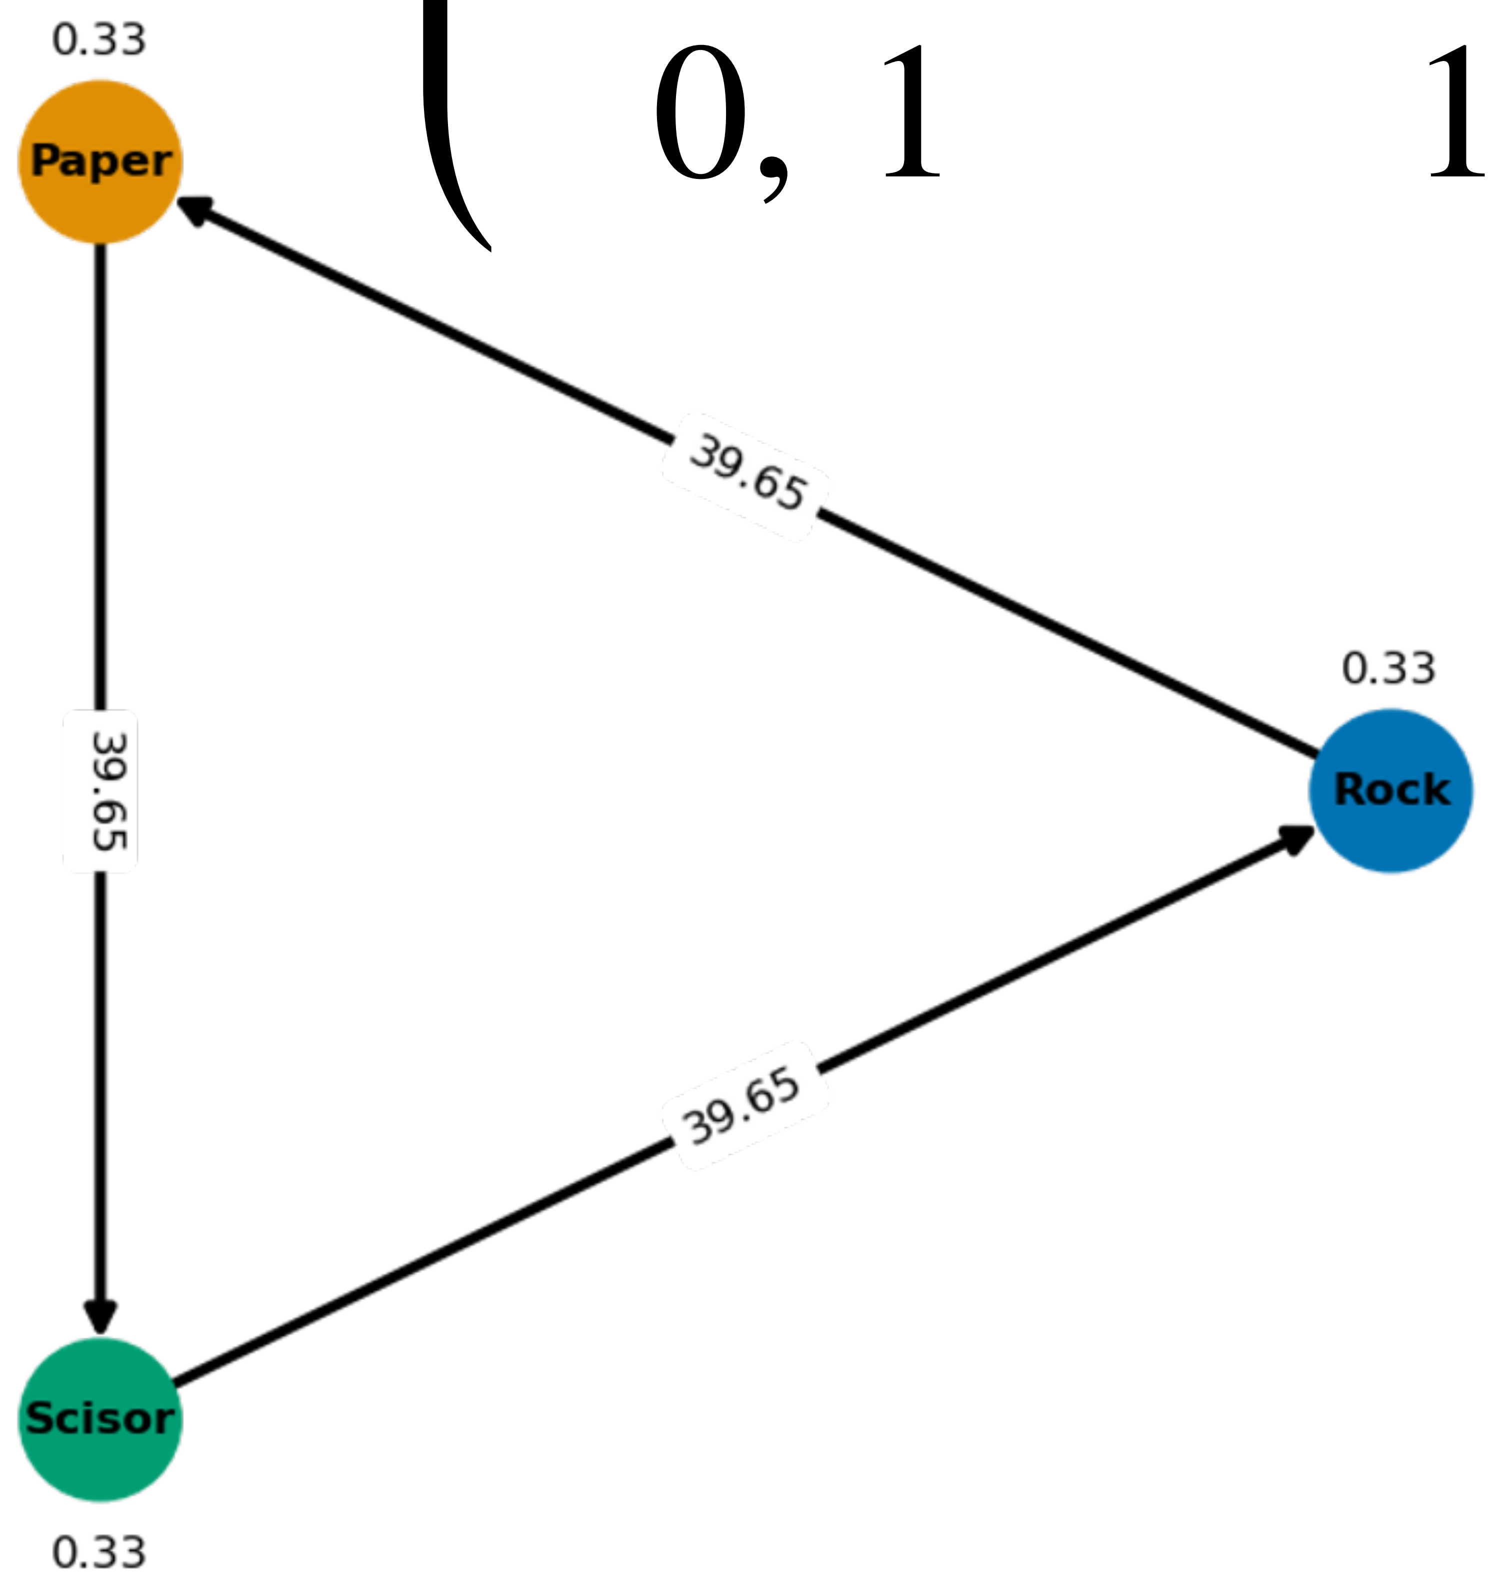

EGTtools demo

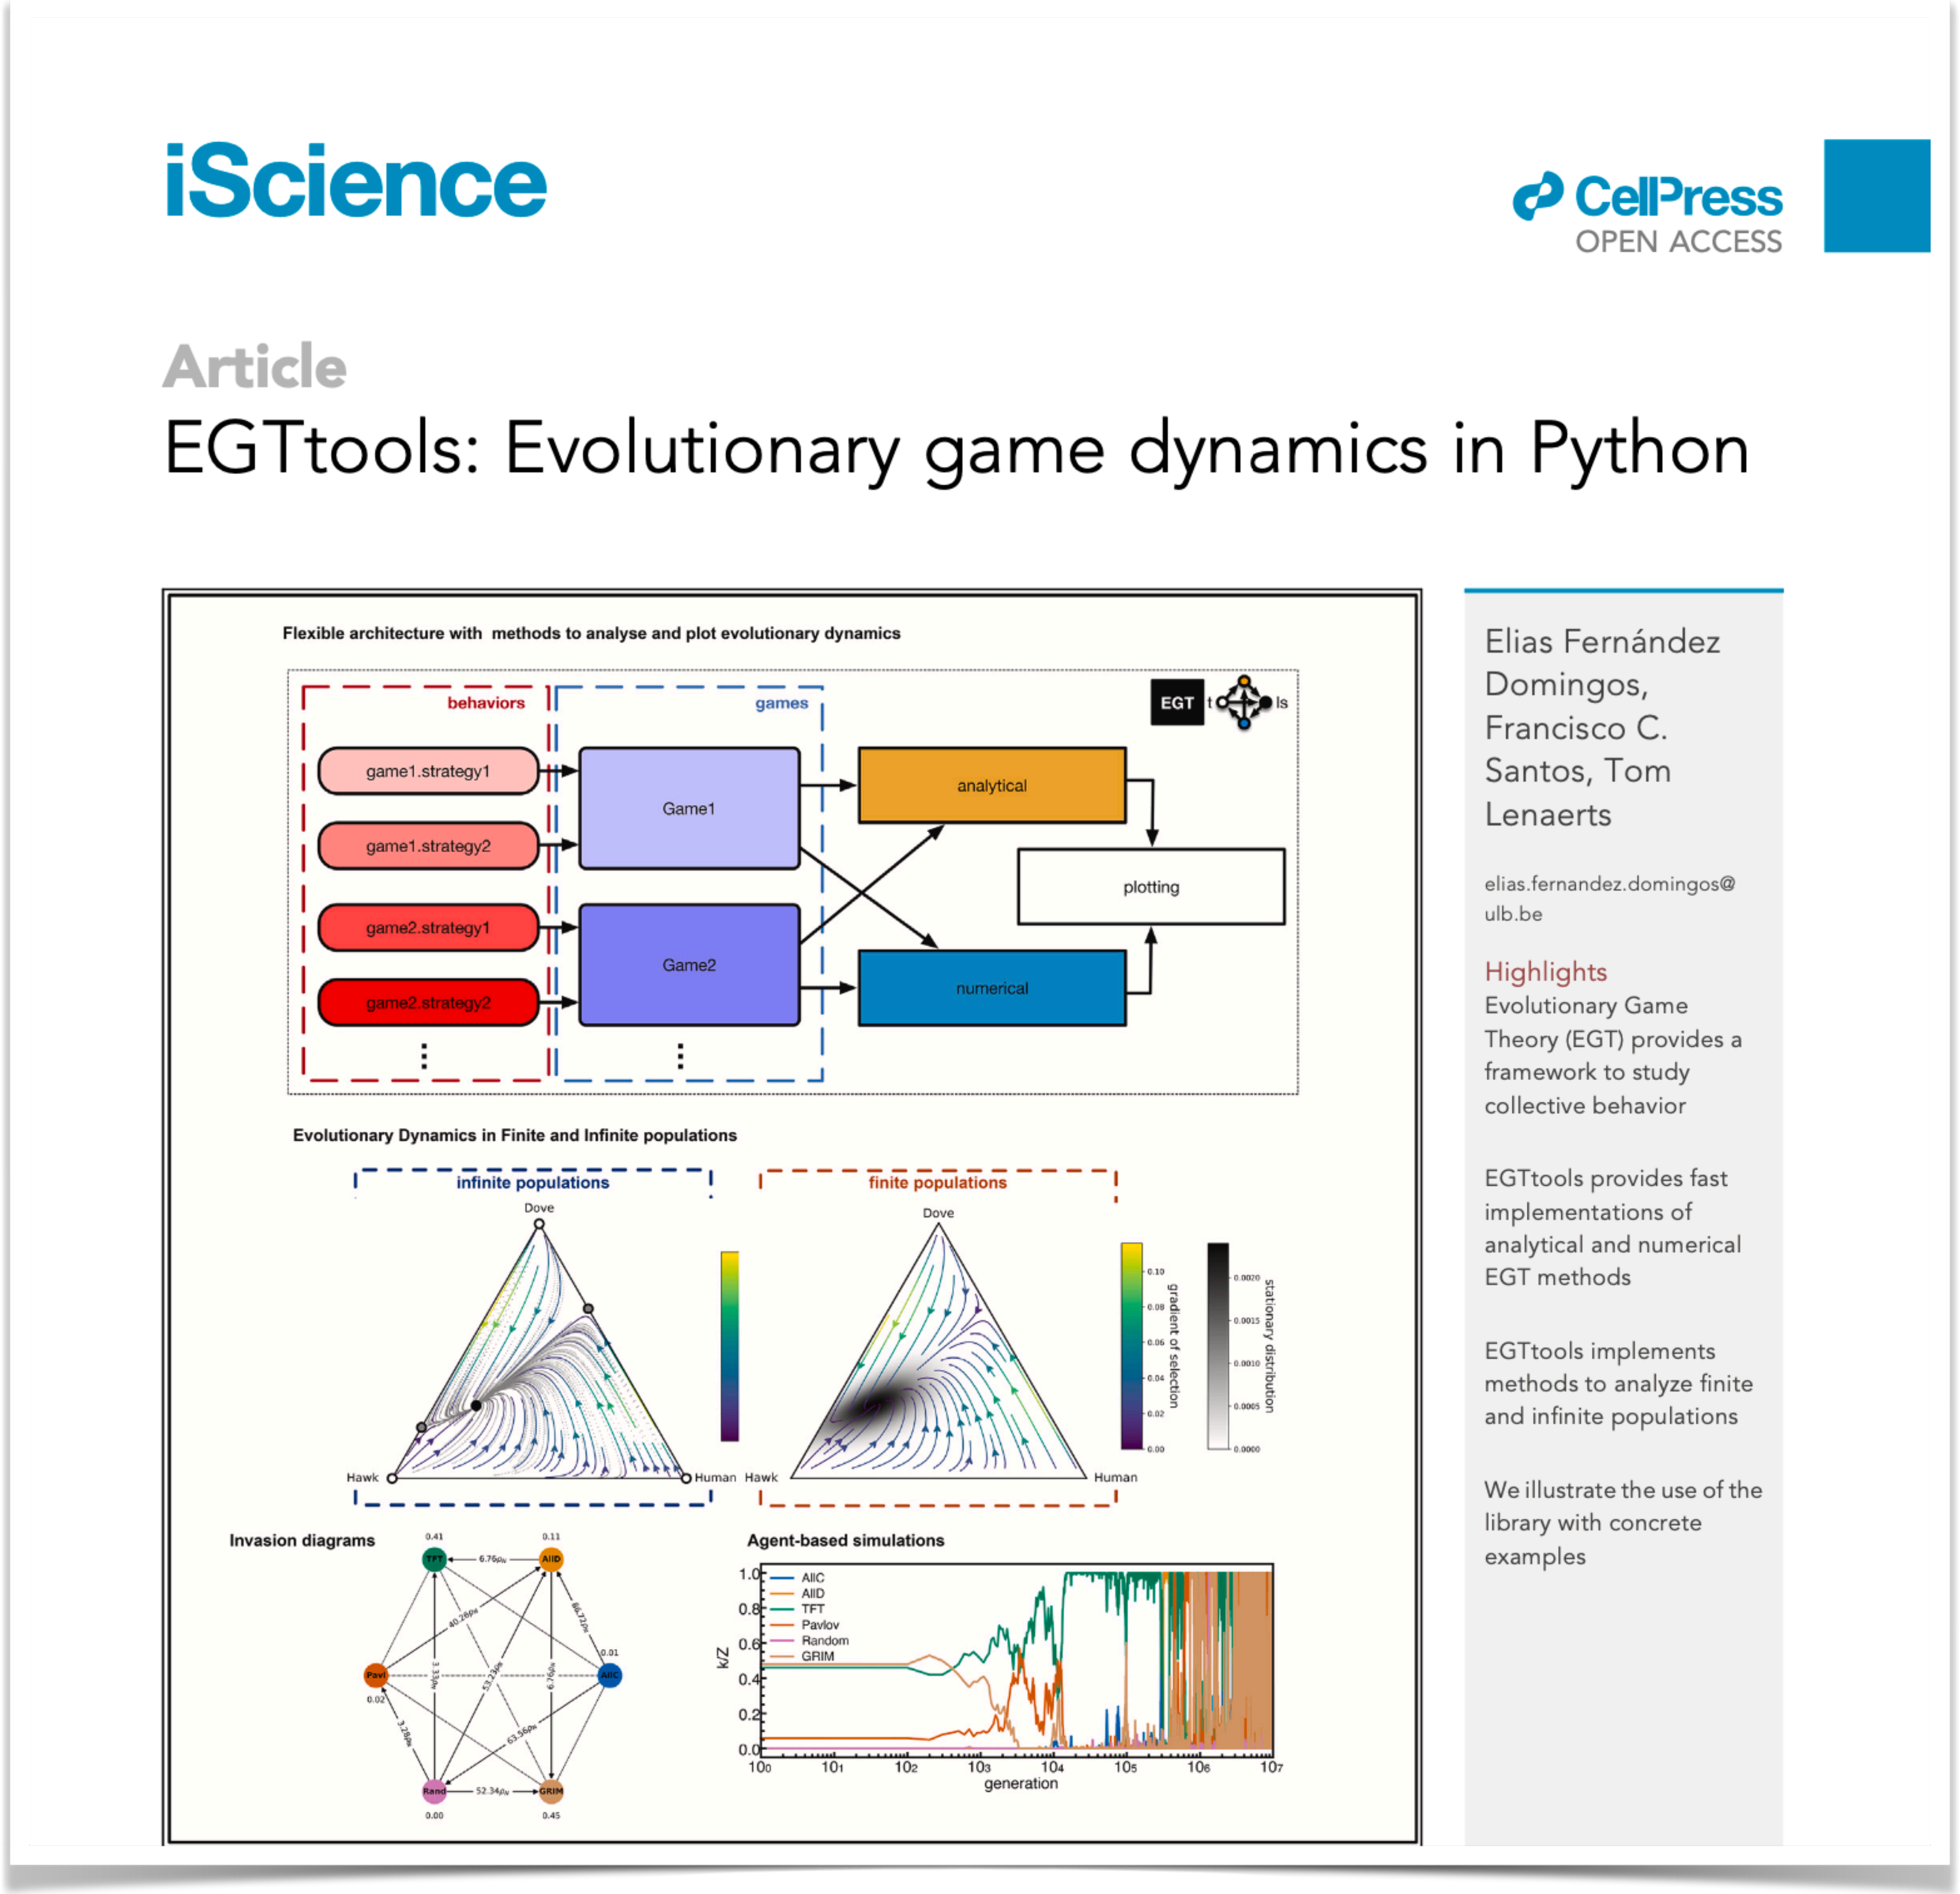
iScience
CellPress OPEN ACCESS

Article

EGTtools: Evolutionary game dynamics in Python

Elias Fernández Domingos, Francisco C. Santos, Tom Lenaerts

elias.fernandez.domingos@ulb.be

Highlights

Evolutionary Game Theory (EGT) provides a framework to study collective behavior

EGTtools provides fast implementations of analytical and numerical EGT methods

EGTtools implements methods to analyze finite and infinite populations

We illustrate the use of the library with concrete examples

Domingos, E. F., Santos, F. C., & Lenaerts, T. (2023). EGTtools: Evolutionary game dynamics in Python. Iscience, 26(4): 106419 https://doi.org/10.1016/j.isci.2023.106419

https://github.com/Socrats/EGTTools

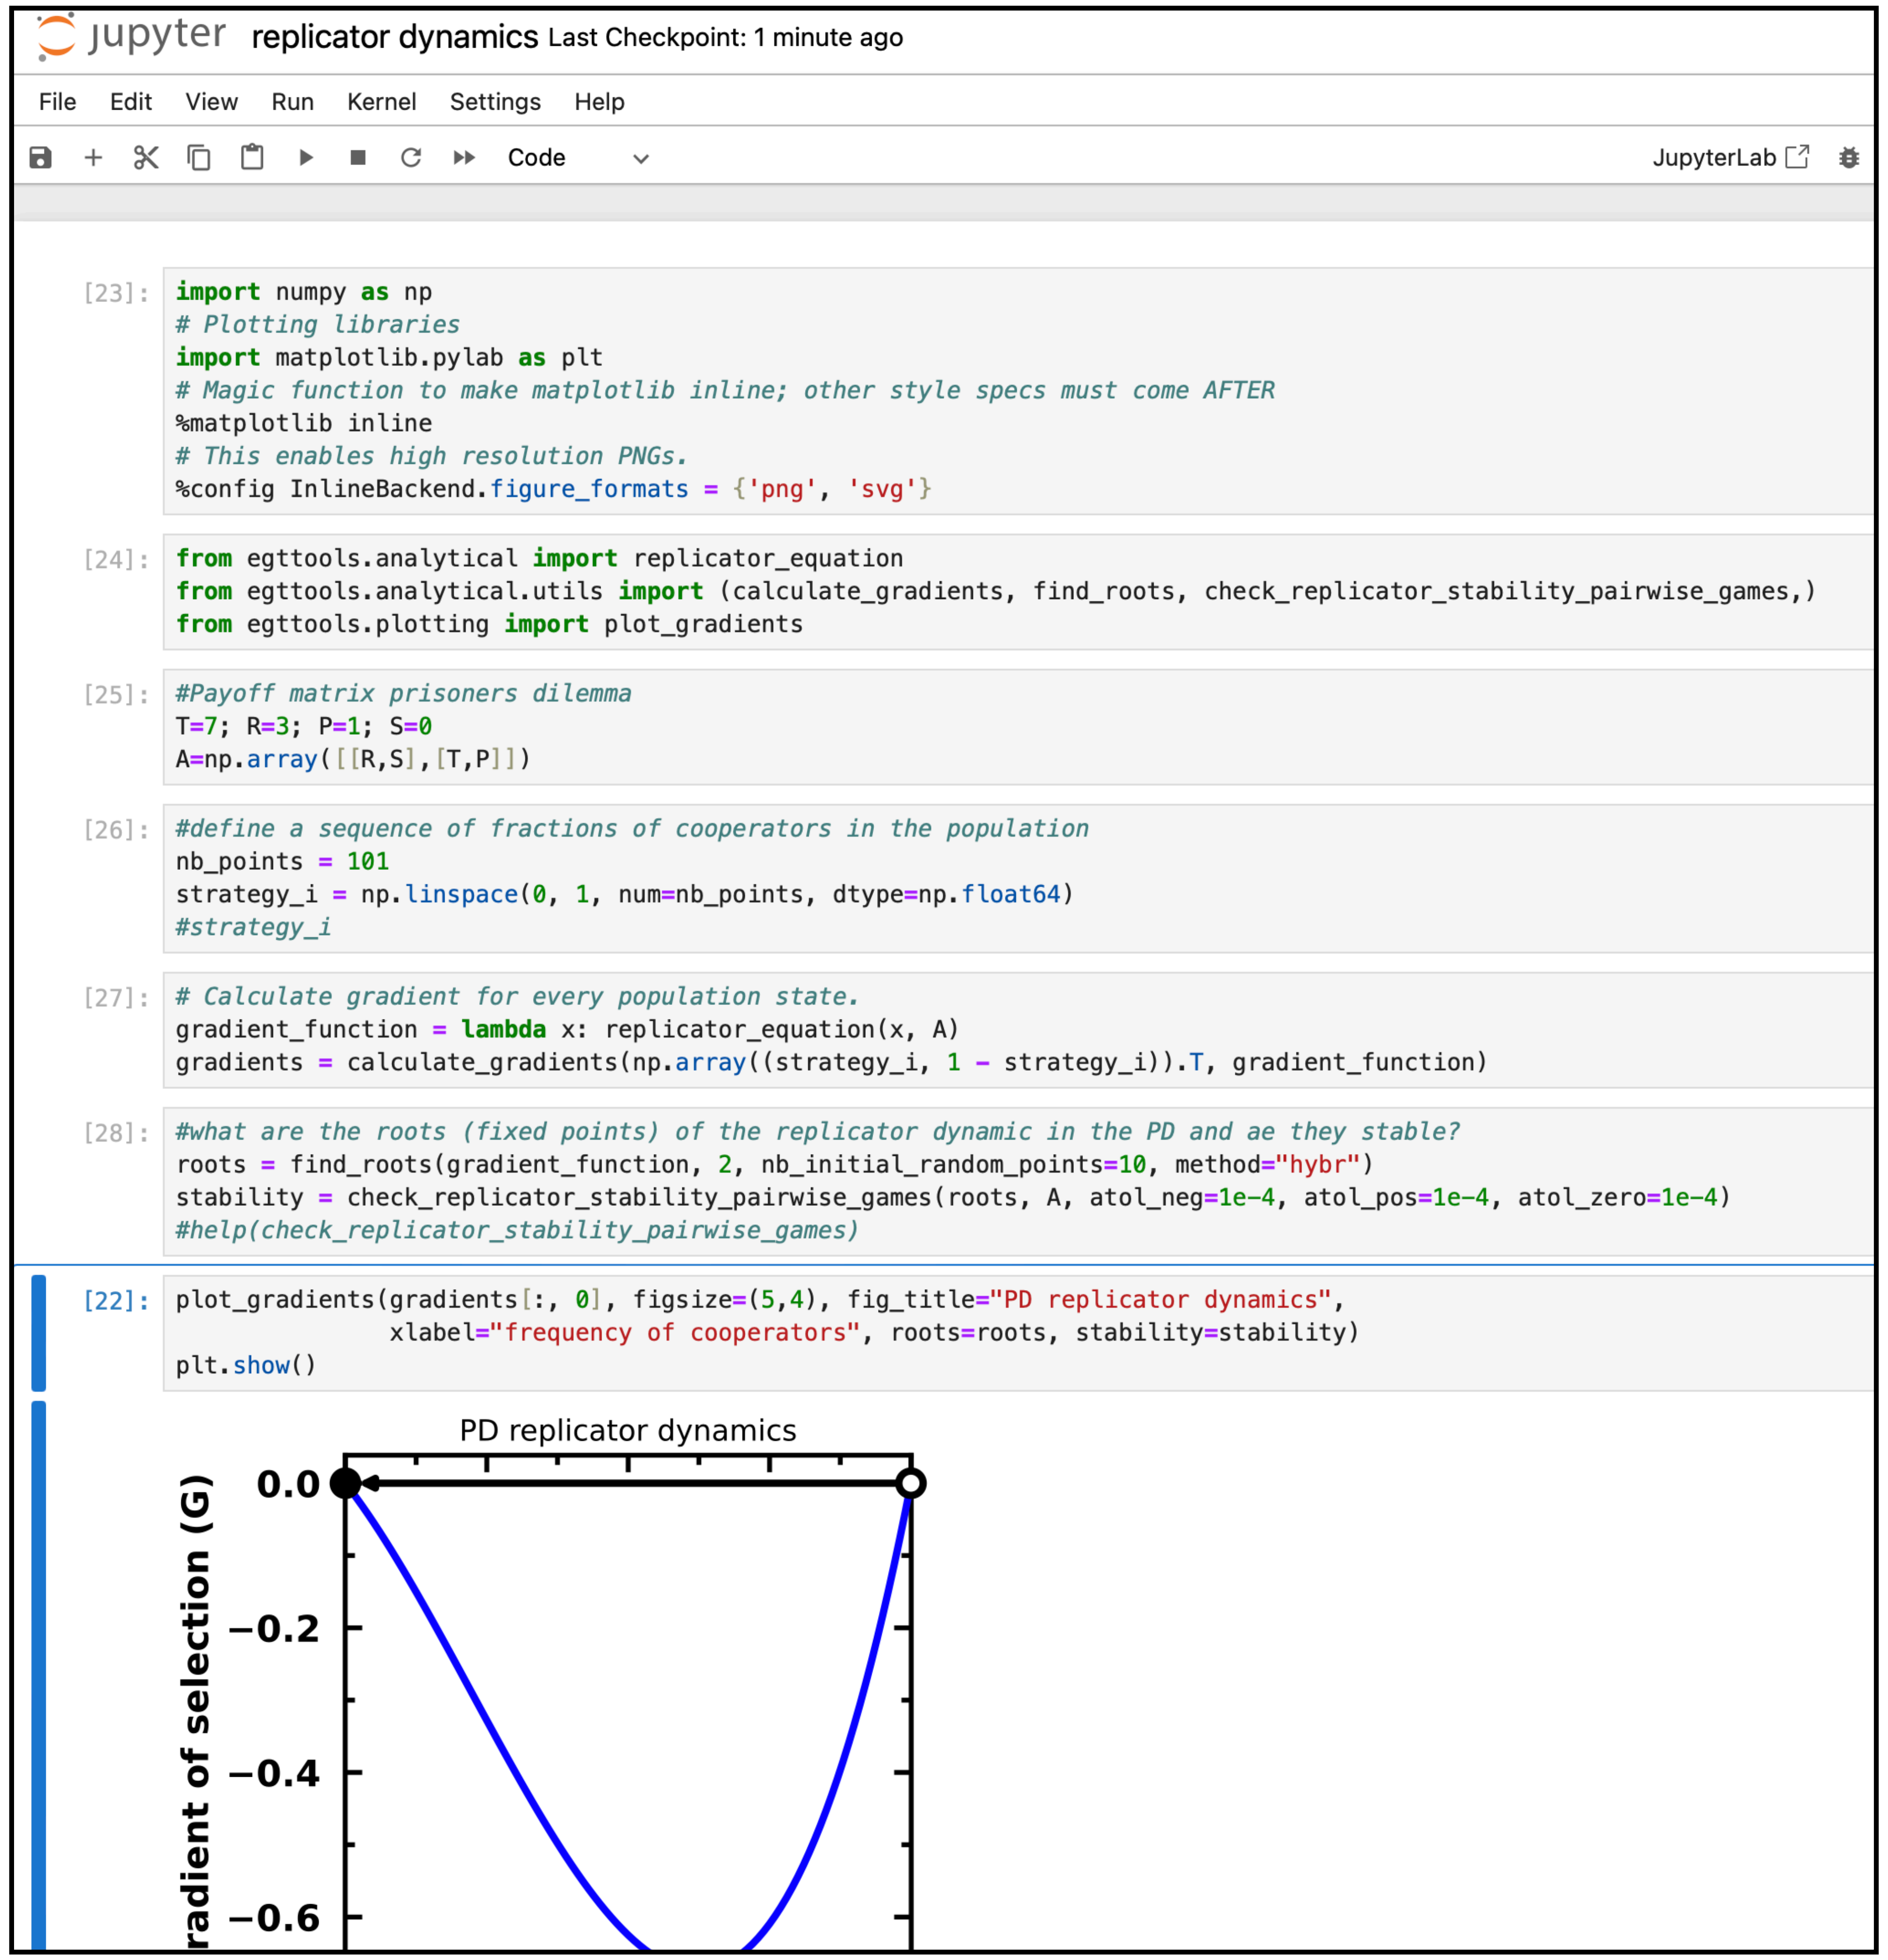

Details on the calculation

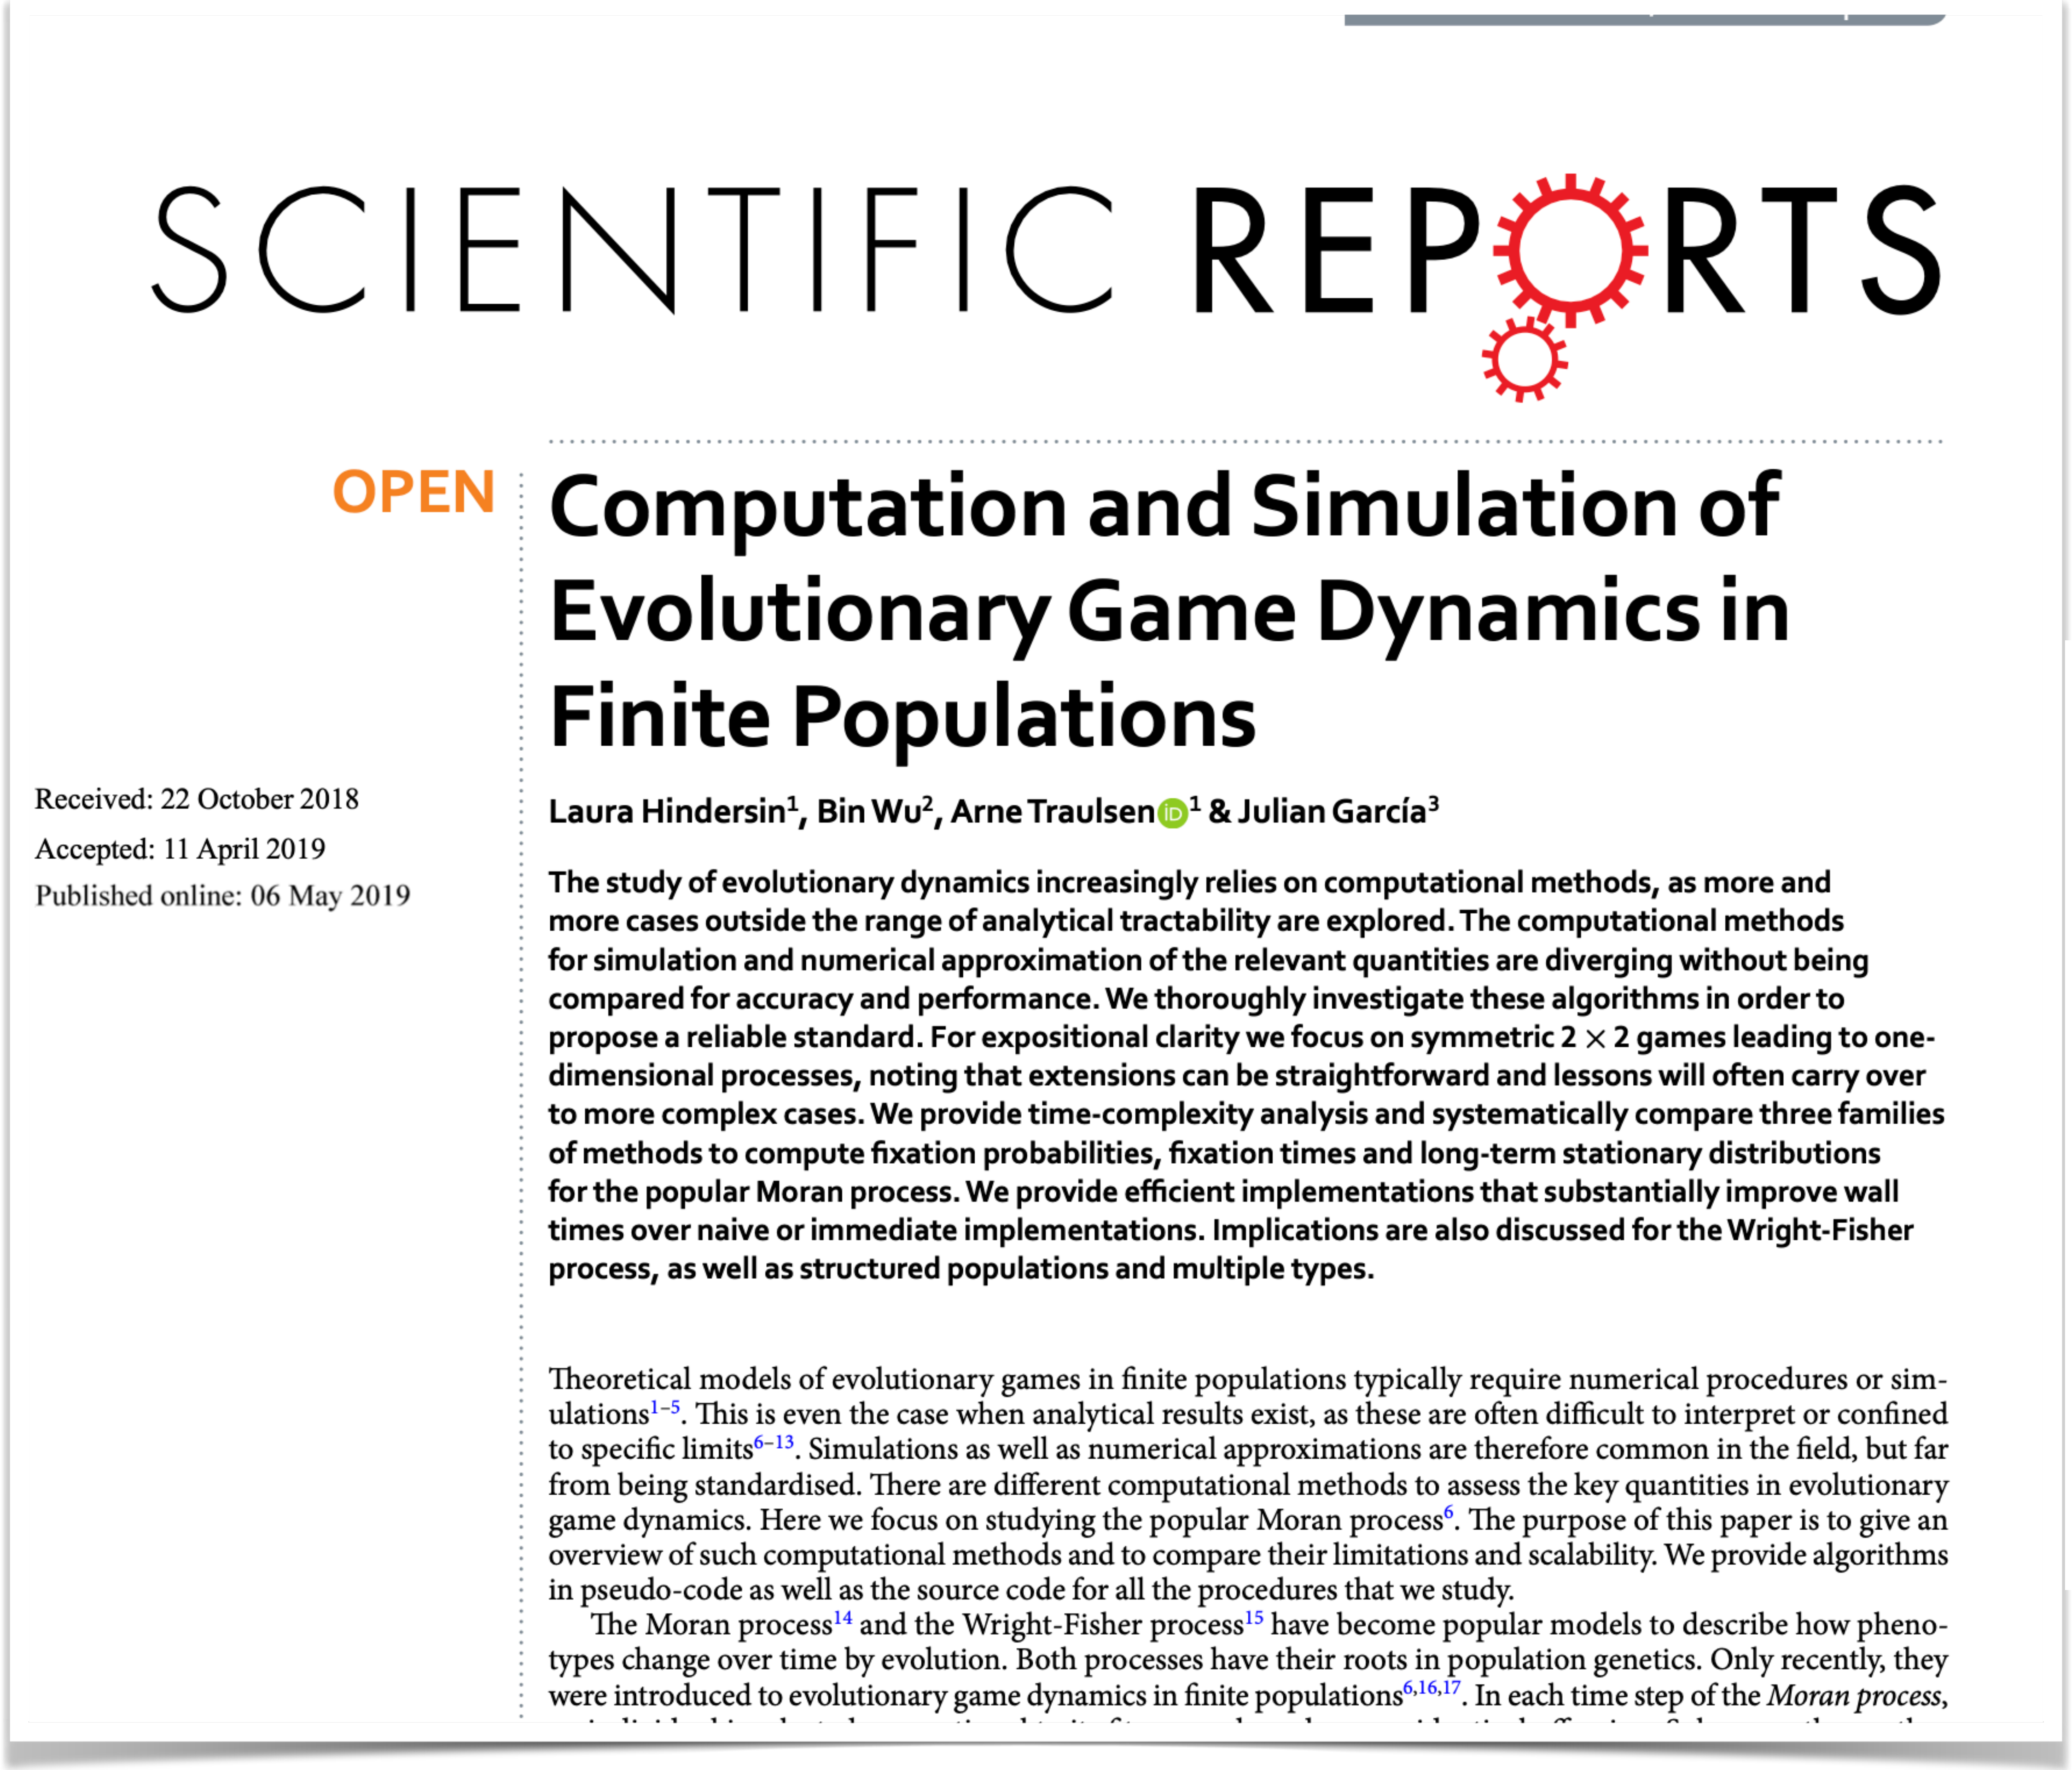

SCIENTIFIC REPORTS

OPEN

Computation and Simulation of Evolutionary Game Dynamics in Finite Populations

Received: 22 October 2018
Accepted: 11 April 2019
Published online: 06 May 2019

Laura Hindersin[1], Bin Wu[2], Arne Traulsen[1] & Julian García[3]

The study of evolutionary dynamics increasingly relies on computational methods, as more and more cases outside the range of analytical tractability are explored. The computational methods for simulation and numerical approximation of the relevant quantities are diverging without being compared for accuracy and performance. We thoroughly investigate these algorithms in order to propose a reliable standard. For expositional clarity we focus on symmetric 2×2 games leading to one-dimensional processes, noting that extensions can be straightforward and lessons will often carry over to more complex cases. We provide time-complexity analysis and systematically compare three families of methods to compute fixation probabilities, fixation times and long-term stationary distributions for the popular Moran process. We provide efficient implementations that substantially improve wall times over naive or immediate implementations. Implications are also discussed for the Wright-Fisher process, as well as structured populations and multiple types.

Theoretical models of evolutionary games in finite populations typically require numerical procedures or simulations[1–5]. This is even the case when analytical results exist, as these are often difficult to interpret or confined to specific limits[6–13]. Simulations as well as numerical approximations are therefore common in the field, but far from being standardised. There are different computational methods to assess the key quantities in evolutionary game dynamics. Here we focus on studying the popular Moran process[6]. The purpose of this paper is to give an overview of such computational methods and to compare their limitations and scalability. We provide algorithms in pseudo-code as well as the source code for all the procedures that we study.

The Moran process[14] and the Wright-Fisher process[15] have become popular models to describe how phenotypes change over time by evolution. Both processes have their roots in population genetics. Only recently, they were introduced to evolutionary game dynamics in finite populations[6,16,17]. In each time step of the *Moran process*,

Hindersin, L., Wu, B., Traulsen, A., & García, J. (2019). Computation and simulation of evolutionary game dynamics in finite populations. Scientific reports, 9(1), 6946.

Method	Advantages	Disadvantages
Direct	Wall time is independent of the game	Limited to birth-death processes
	Extendable to other birth-death processes, such as pairwise comparison processes (same complexity)	Not extendable to general graphs
Matrix-based	Wall time is independent of the game	Strongly limited by population size due to size of transition matrix
	Extendable to processes with dense transition matrix, such as Wright-Fisher (increased complexity)	
Simulations	Extendable to Fermi and Wright-Fisher	Wall time depends on the game and the selection intensity
	Extendable to games on graphs and multi-player games	Large number of realisations might be necessary

Table 1. Overview of the three methods discussed here. This table lists their limitations and possible extensions.

Replicator equation results for all social dilemmas

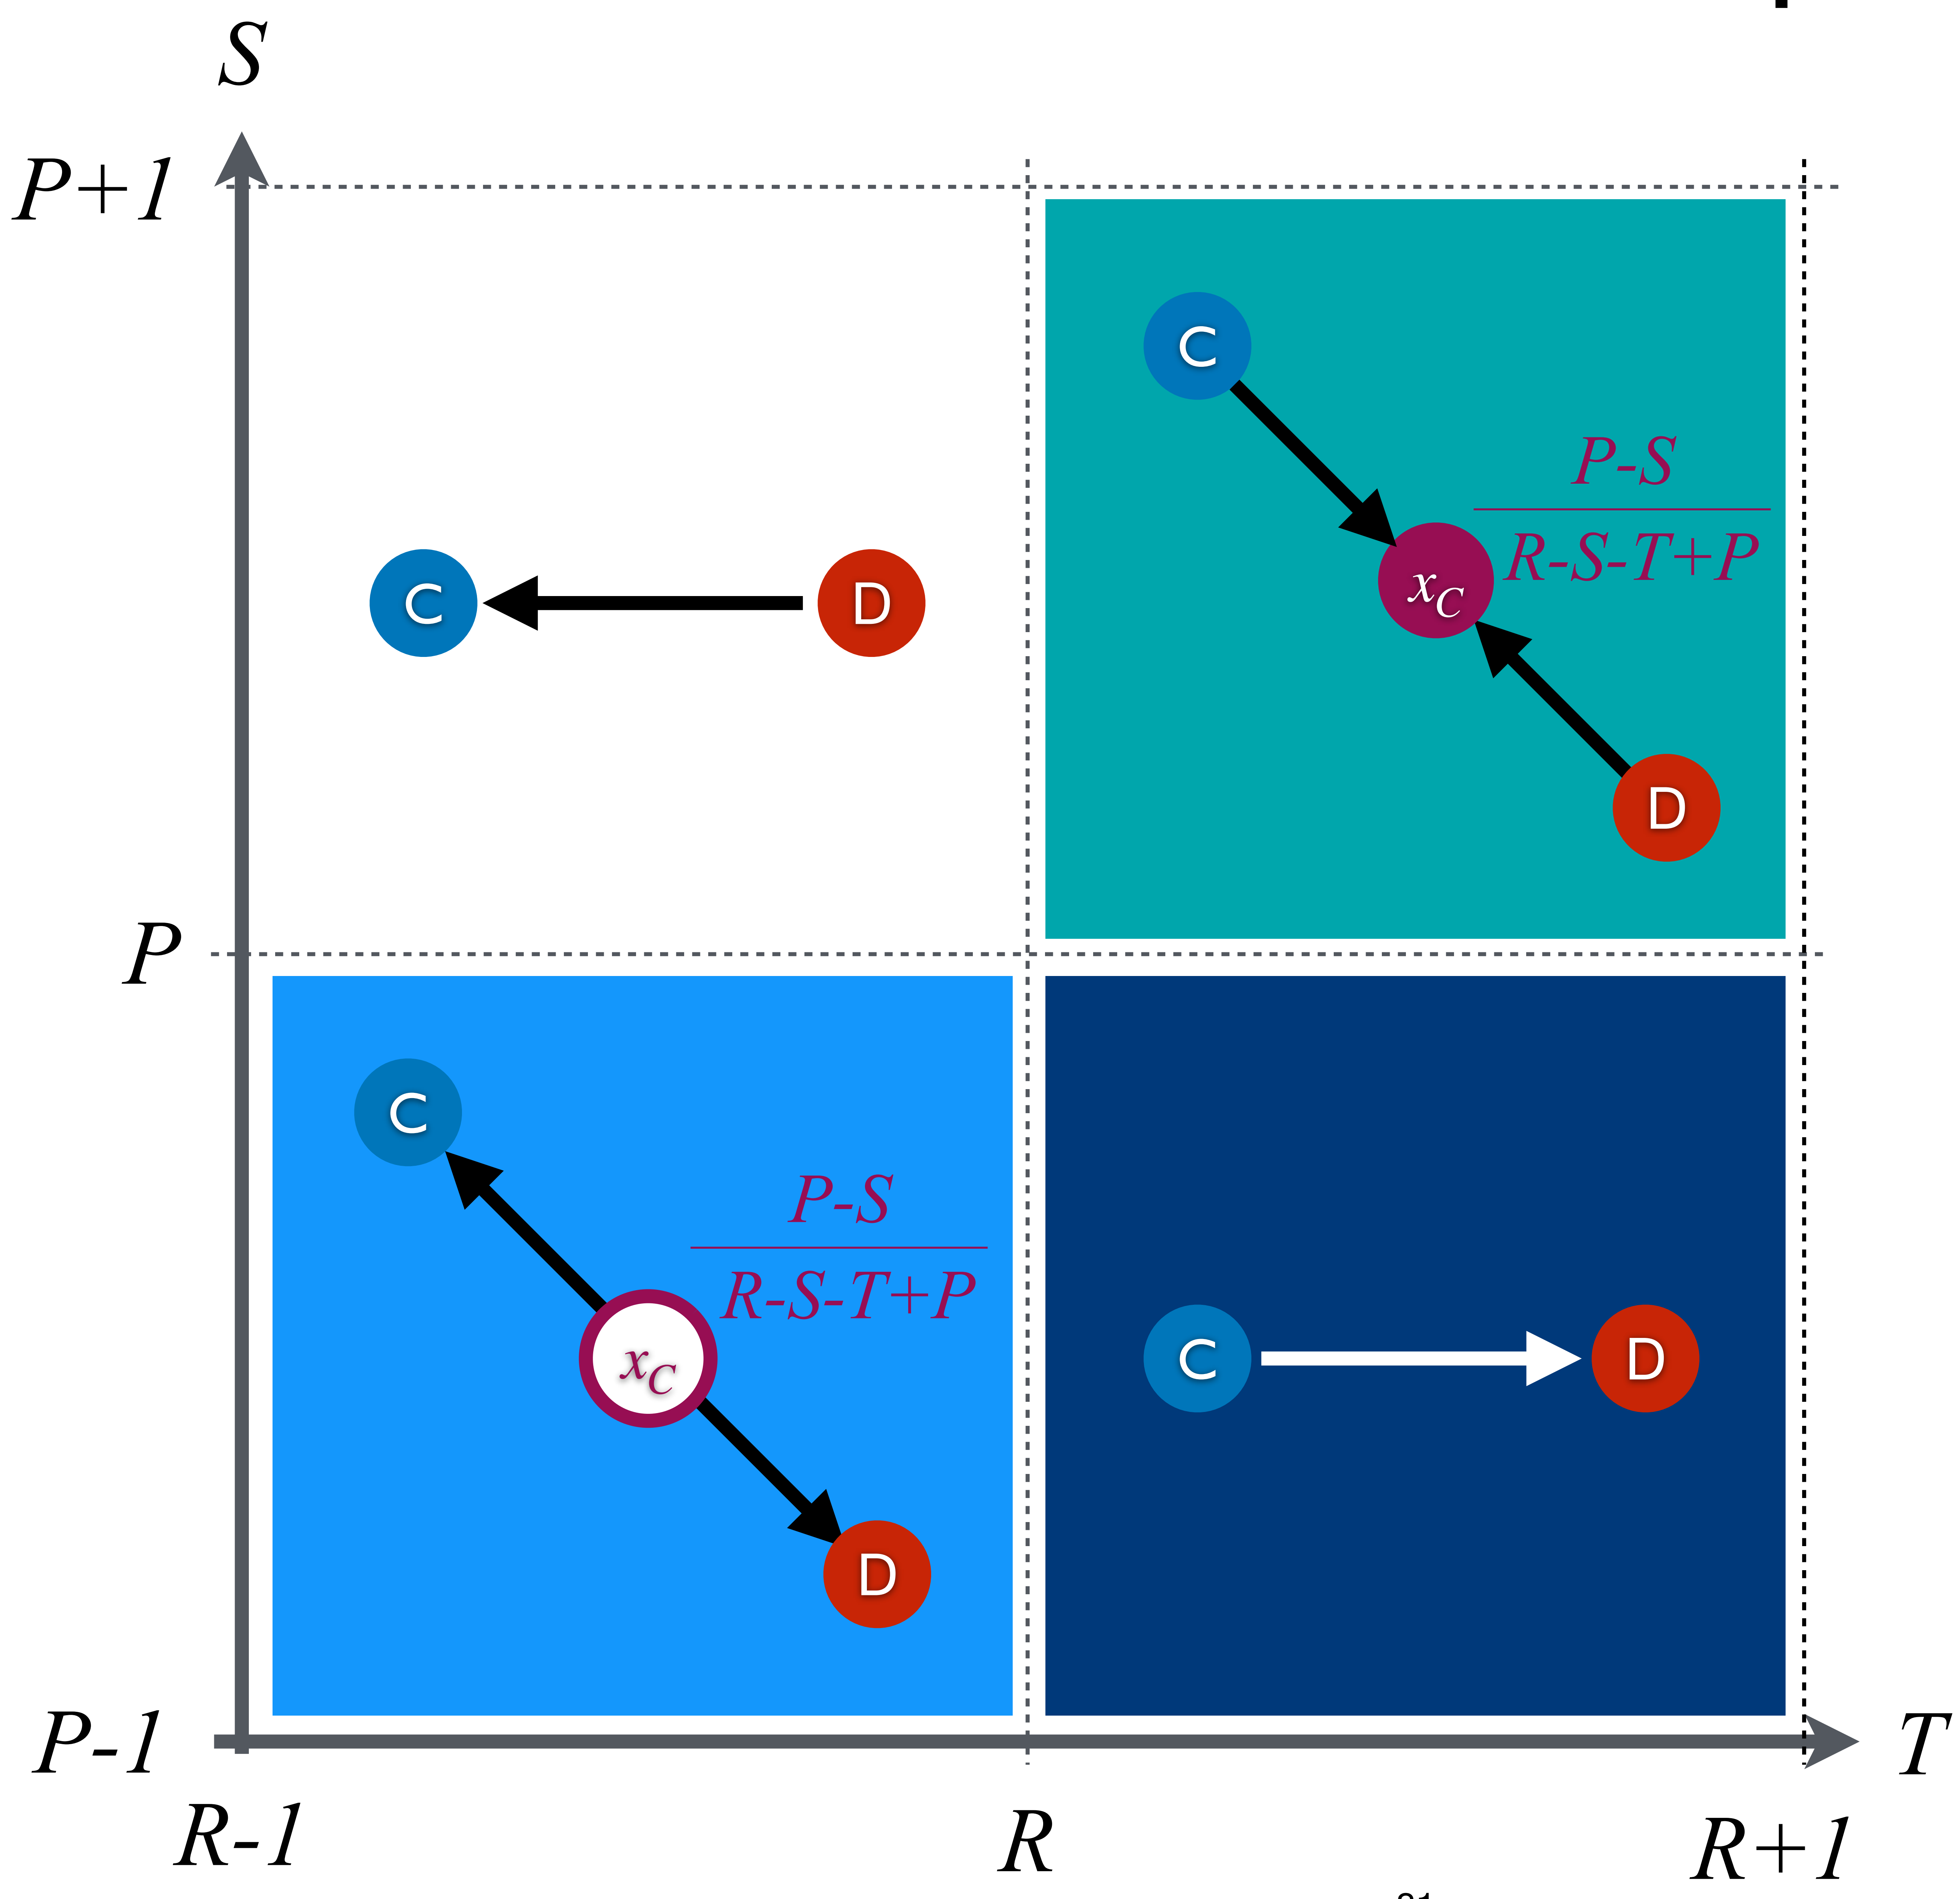

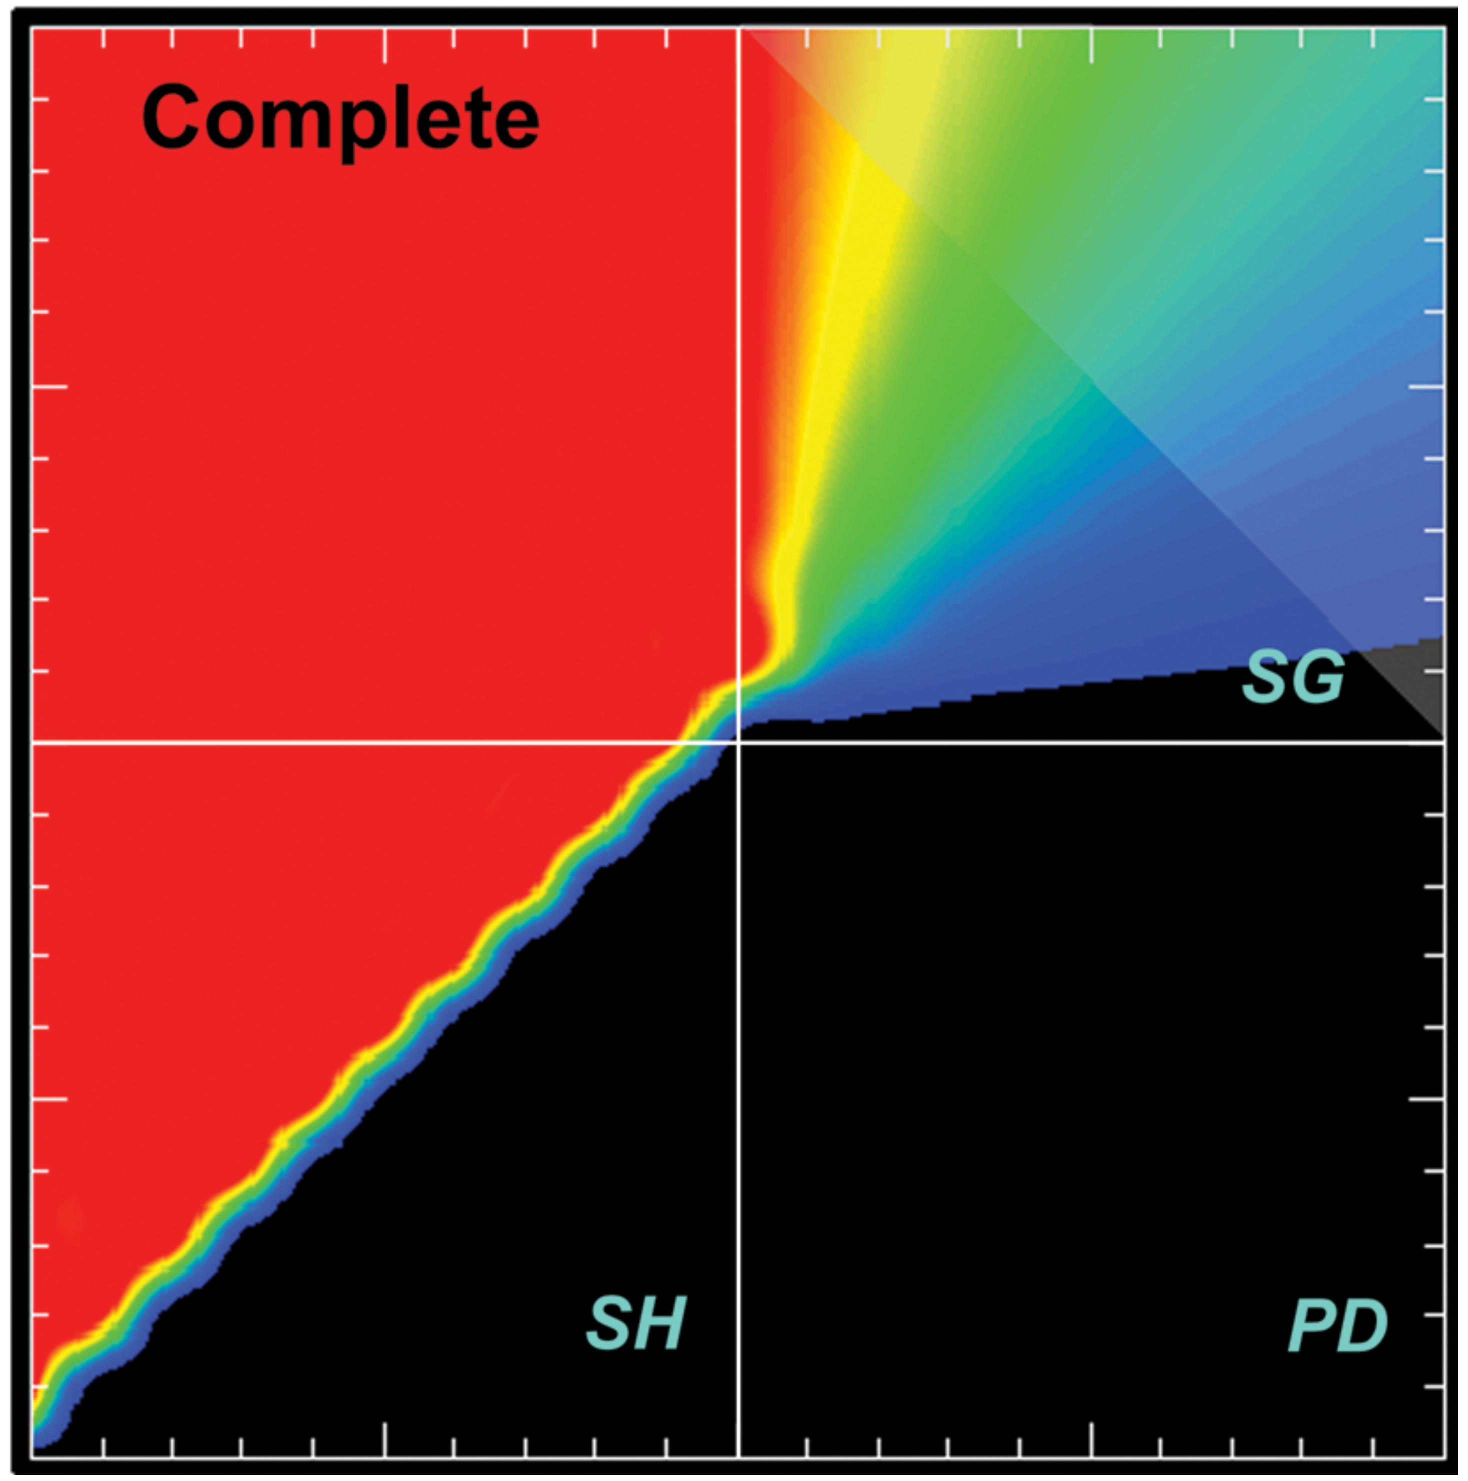

Mechanisms for the evolution of cooperation

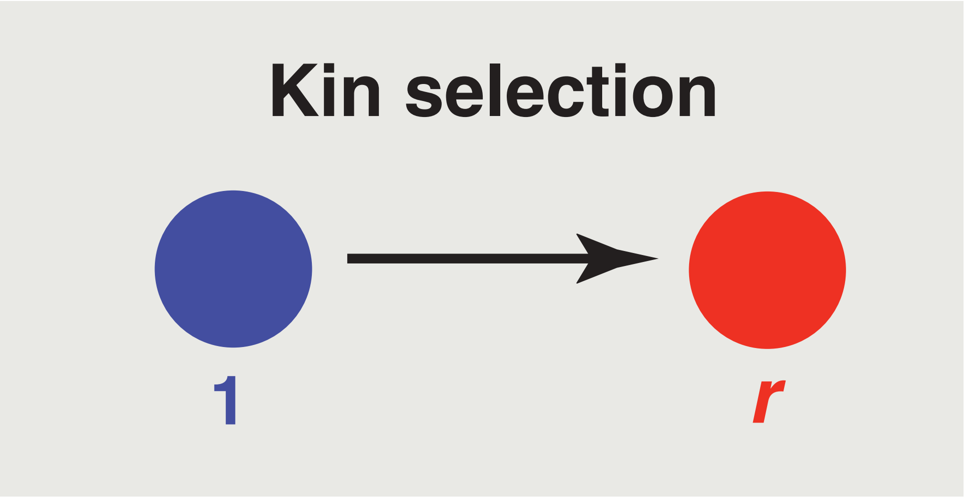

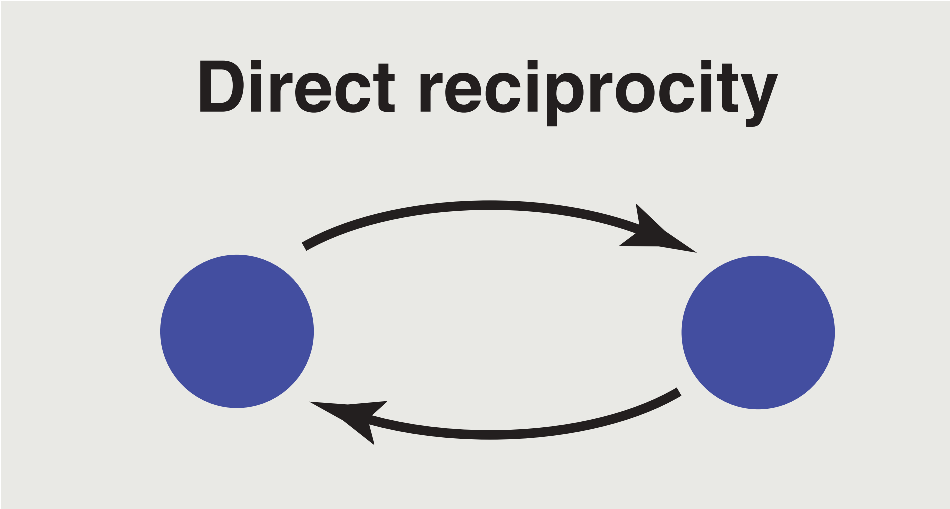

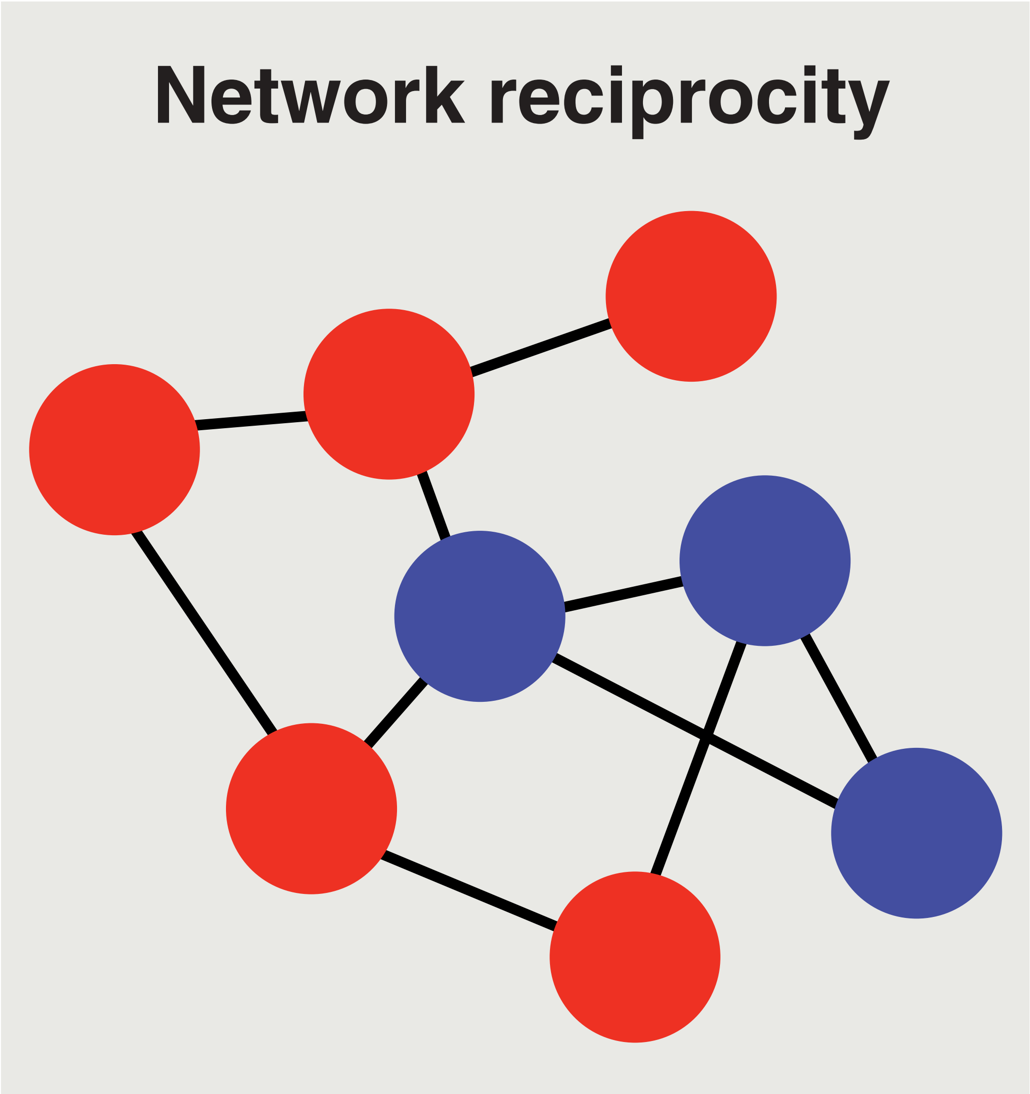

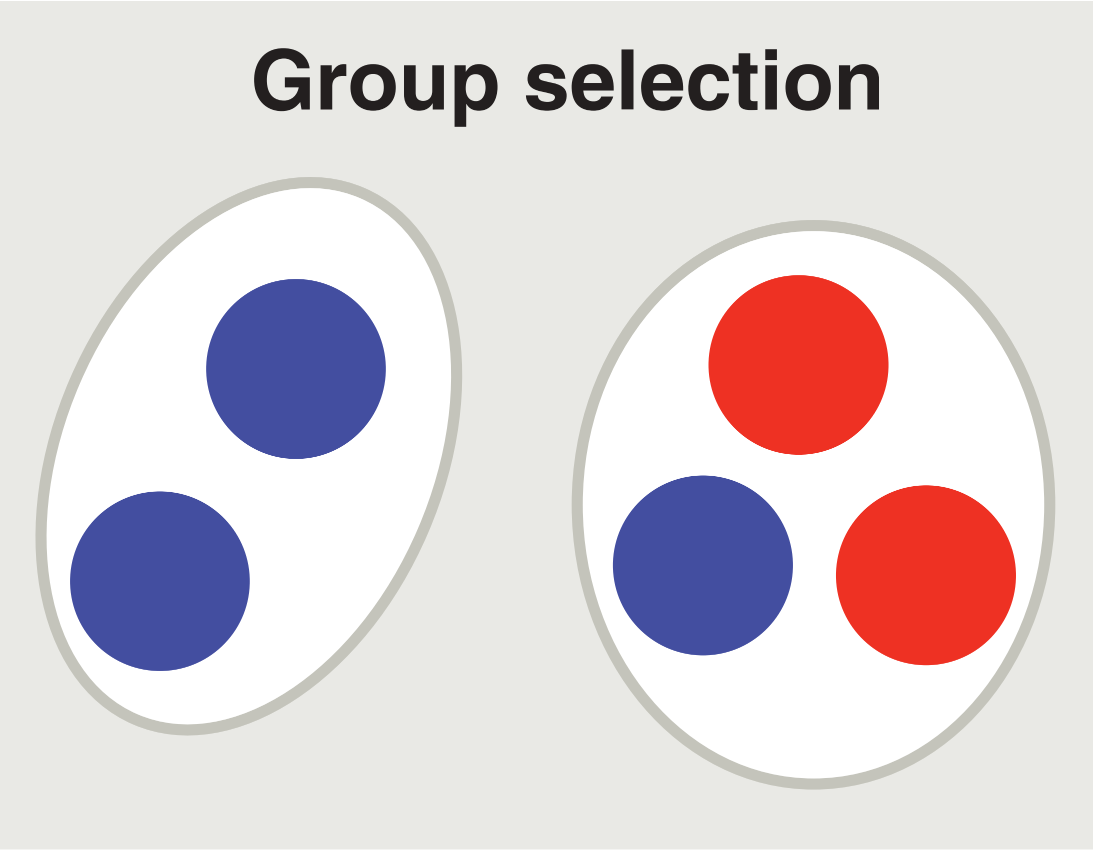

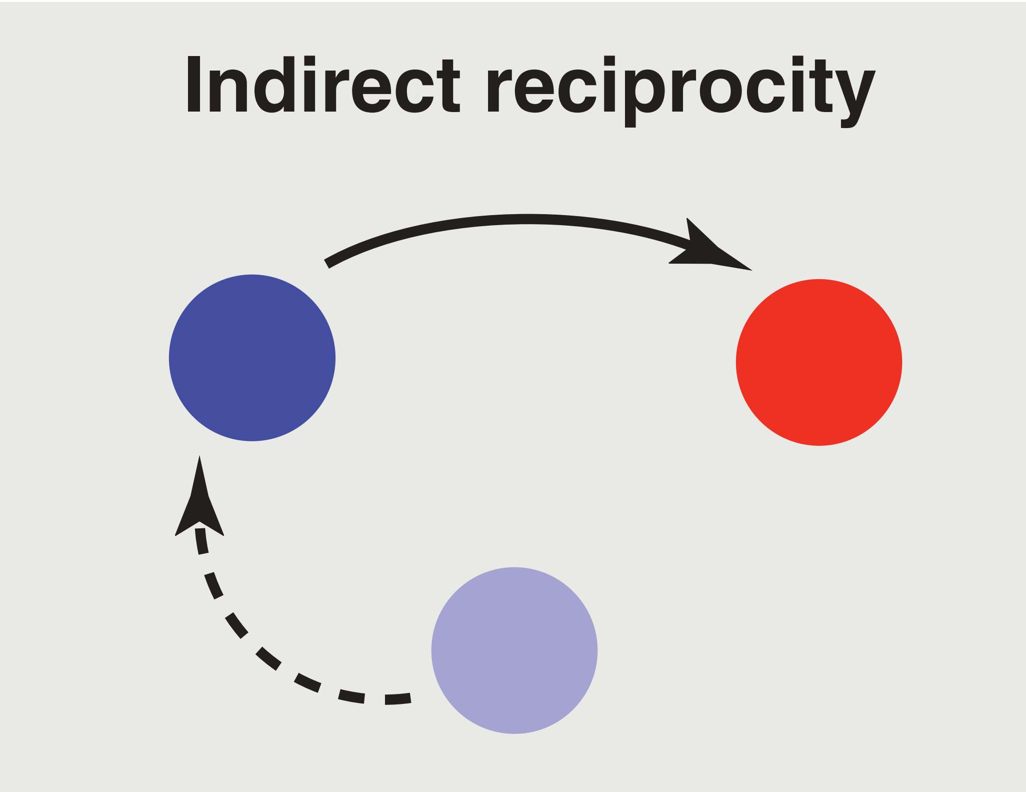

Part 4: Tutorial on how to reproduce an EGT paper

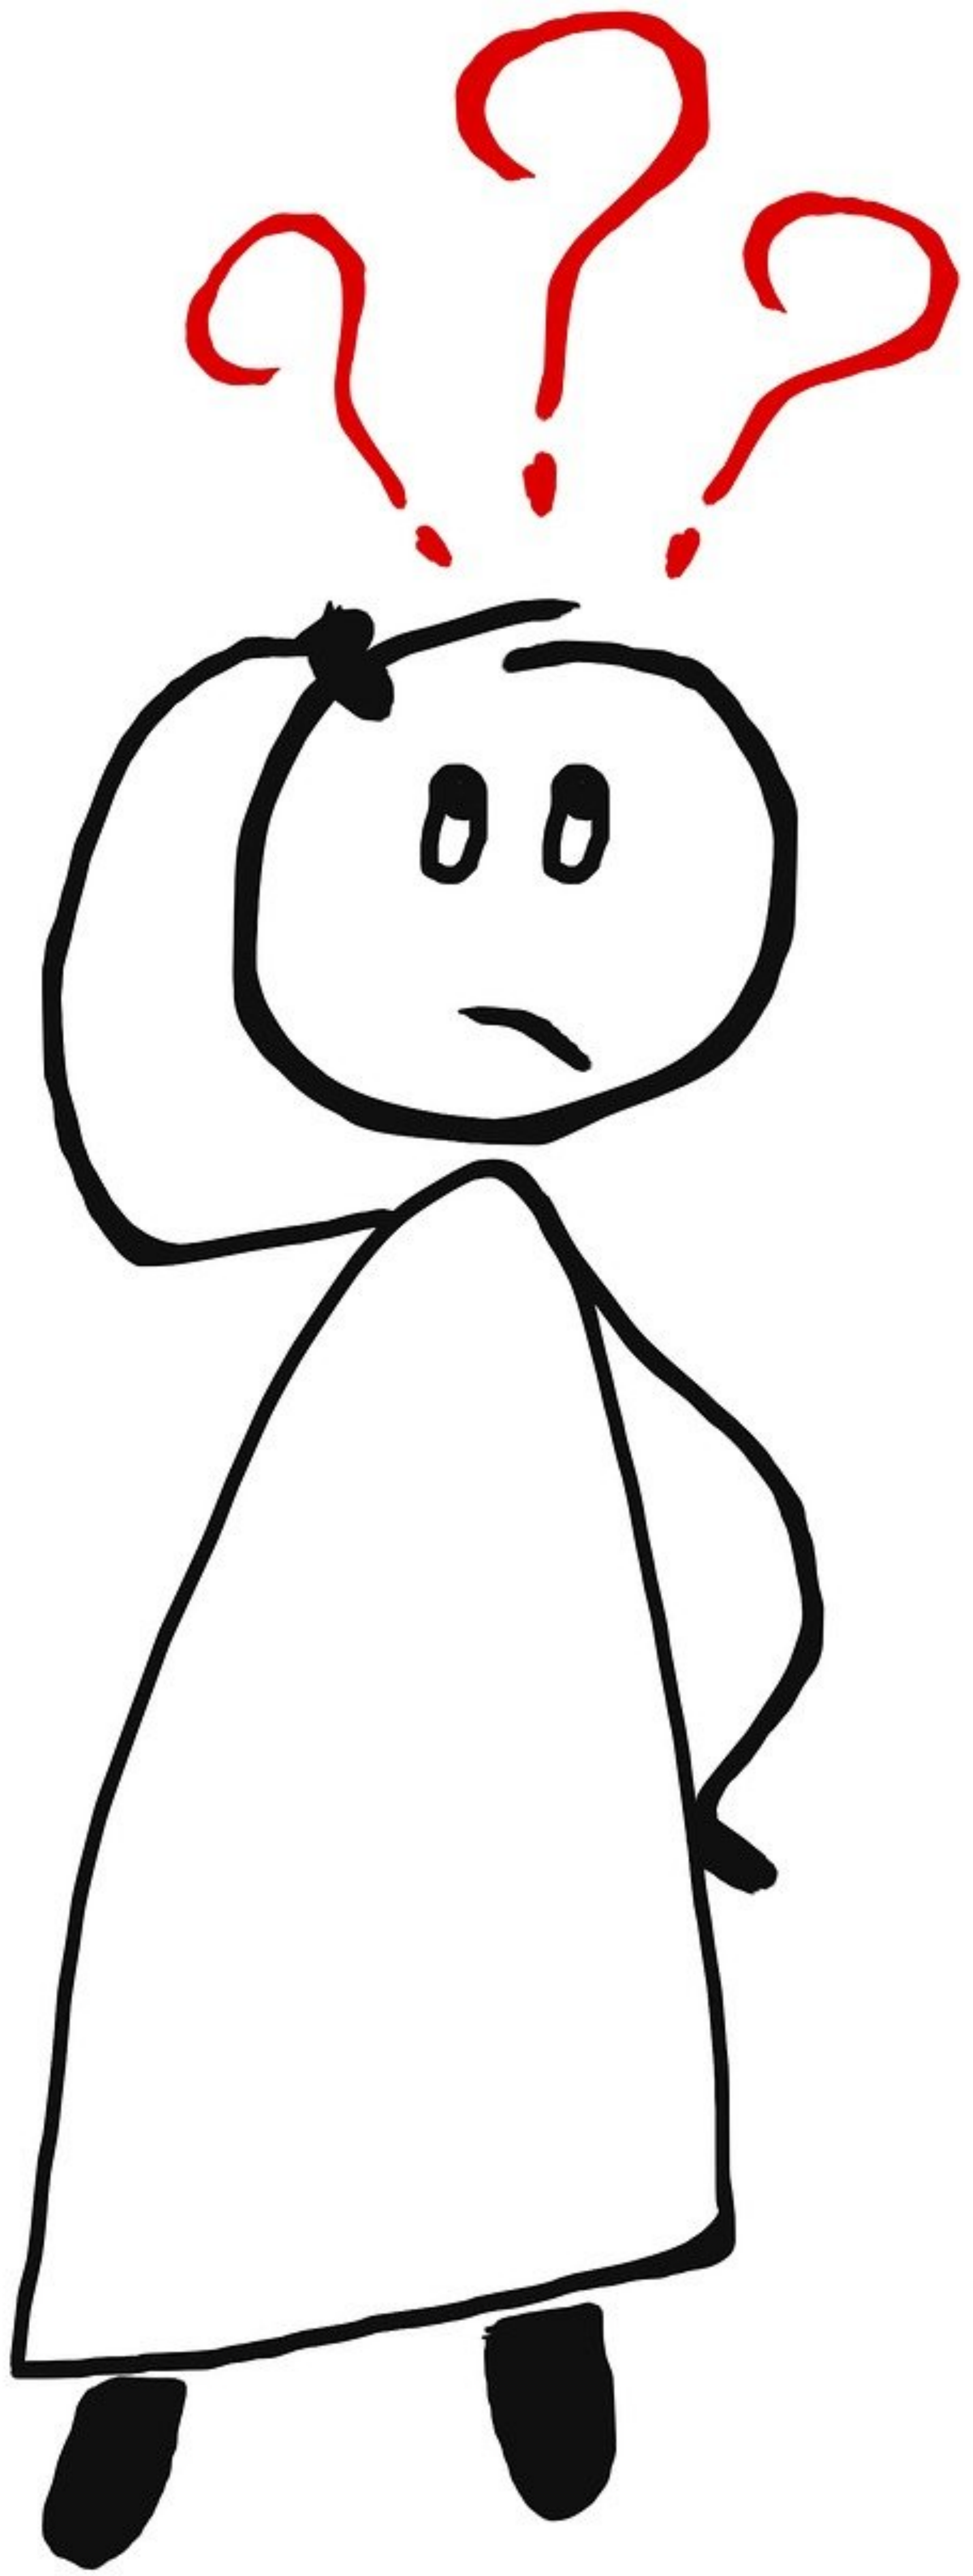

Questions ?

elias.Fernández.Domingos@ulb.be
@esocrats
https://github.com/Socrats