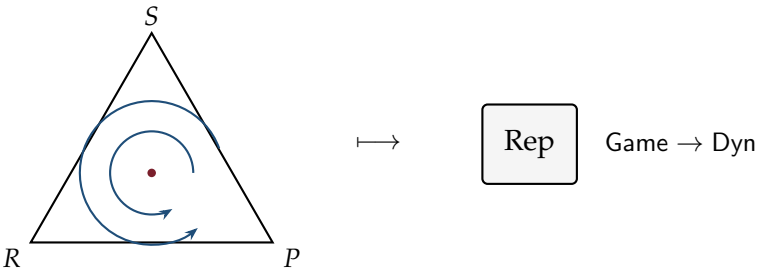
Higher Categorical Semantics

Evolutionary Game Theory in Higher Categories

Stability, replicators, recurrence,
finite populations and local interaction



Anonymous

2026

Version 1.0, 17 November 2026.

This volume is a treatment, in categorical language, of the material surveyed in J. McKenzie Alexander, *Evolutionary Game Theory* (Cambridge Elements in Decision Theory and Philosophy, Cambridge University Press, 2023). It reproduces none of that book's text, figures or tables; where a numerical example of the source is recomputed, this is said. An appendix gives a section-level concordance.

Every numbered statement carries a status tag: [classical] a classical result; [literature; source] a result of the cited literature; [reformulation; cf. source] a known result restated in the language of this volume; [observation] an observation of this volume, proved here; [programme] a proposal. Tags of the form [R4] name a check in the companion script `verify-evolutionary-games.py`; tags of the form [source, Thm. 6] point to the source.

This book was drafted with the assistance of an AI system (Anthropic's Claude), which also wrote and ran the companion script. The bibliography was assembled from memory and should be checked before it is relied on.

Preface

Evolutionary game theory replaces the rational player of classical game theory by a population. Strategies are inherited or imitated, payoffs are rates of reproduction or adoption, and the questions change from "what should one do?" to "what persists?". Alexander's Element [Alex23] is a compact map of the subject: the static theory of evolutionarily stable strategies and its set-valued refinements; the replicator equation and its relatives; finite populations; and models in which interaction is local. It is written for philosophers, and it is candid about where the models are fragile.

This volume retells that material with one question in mind: which of the constructions are *natural*, in the technical sense that they commute with the evident maps between games? The answers organize the subject more sharply than we expected.

- The replicator equation is a functor from games and payoff-preserving maps to dynamical systems (Theorem 3.2). Invariance of faces, the lumping of duplicate strategies, and the action of symmetries are three instances of one statement.
- Evolutionary stability of a *point* is not preserved by pulling back along a map that duplicates strategies; evolutionary stability of a *set* is (Theorem 2.8). The set-valued refinements that the source surveys are what the point-valued notion is forced to become.
- The equilibria of a regular game carry indices that sum to one, the Euler characteristic of the simplex (Theorem 4.2); behind the number is a homotopy type, the Conley index, and behind recurrent behaviour a decomposition whose obstruction is a first cohomology class (Chapter 4).
- A population is an open system, with payoffs as input and a state as output, and a game is a wiring of such systems (Chapter 5). The replicator is one box among several, and the source's comparison of dynamics is a comparison of boxes.
- Local interaction models are natural with respect to coverings of the interaction graph and equivariant under its symmetries (Theorem 7.1), which accounts for the rotating patterns the source exhibits.

Quantitative statements accompany the qualitative ones throughout: rates

of extinction, escape times, eigenvalues, fixation probabilities, thresholds. All of them are recomputed by the companion script.

How much of this is "higher" is assessed in Chapter 8. The short answer: most of it is ordinary category theory; homotopy enters through the index and the Conley index; cohomology enters in degree one; and the compositional structure of Chapter 5 is two-dimensional. Nothing here is a new theorem of evolutionary game theory. A companion paper applies the framework to machine learning.

The volume belongs to a series with [GIHC], on classical game theory, [ESCHC], on information theory, and [DTHC], on differential topology, and uses each of them once.

Contents

I Statics 1

1 Games, Populations, and the Simplex 2
- 1.1 Two readings of one object . . . 2
- 1.2 Maps of games . . . 3
- 1.3 Equilibria . . . 4

2 Evolutionary Stability 6
- 2.1 The invasion form . . . 6
- 2.2 Properties . . . 7
- 2.3 Sets, and why they are forced . . . 8
- 2.4 Relative entropy . . . 8

II Dynamics 10

3 The Replicator Equation 11
- 3.1 Derivation as a quotient . . . 11
- 3.2 Naturality . . . 11
- 3.3 Two strategies . . . 12
- 3.4 Geometry: gradients and entropy . . . 12
- 3.5 The folk theorem and elimination . . . 13
- 3.6 Cycles . . . 14

4 Index, Recurrence, and Decomposition 16
- 4.1 The index theorem . . . 16
- 4.2 Recurrence and the Conley index . . . 17
- 4.3 Decomposition . . . 18

5 Revision Protocols, Other Dynamics, and Open Systems 20
- 5.1 Revision protocols . . . 20
- 5.2 Imitation is natural and innovation is not . . . 21
- 5.3 Populations as open systems . . . 22
- 5.4 Two populations . . . 23

III Structure 25

6 Finite Populations 26
6.1 The Moran process . . . 26
6.2 Two limits that do not commute . . . 27
6.3 Weak selection . . . 28
6.4 Rare mutations . . . 29

7 Local Interaction 30
7.1 The model . . . 30
7.2 Coverings . . . 30
7.3 Why the ring rotates . . . 31
7.4 Cooperation at the boundary . . . 32
7.5 Graphs that amplify and suppress selection . . . 33

IV Assessment 35

8 Assessment 36
8.1 What the language did . . . 36
8.2 What it did not do . . . 36
8.3 How much of this is higher . . . 37
8.4 Programme . . . 38

A Toolkit 39

B Concordance with the Source 41

C Computations 43

D Solutions to the Exercises 46

References 50

Part I

Statics

Chapter 1

Games, Populations, and the Simplex

1.1 Two readings of one object

A point of the simplex $\Delta(S)$ over a finite set S of strategies can be read in two ways. It is a *mixed strategy*: one player randomizing over S. It is also a *population state*: a large population in which the fraction x_s of individuals plays the pure strategy s. The source moves between the readings freely and remarks that the equivalence of two definitions of evolutionary stability rests on their agreement [source, §2.1]. The agreement is a fact about maps, and it is the starting point of this volume.

Let FinStoch be the category of finite sets and stochastic matrices, and let Δ be the finite distribution monad on Set: $\Delta(S)$ is the set of probability vectors on S, a function $f\colon S \to T$ induces the pushforward $f_*\colon \Delta(S) \to \Delta(T)$, $(f_*x)_t = \sum_{f(s)=t} x_s$, and the algebras of Δ are the convex spaces. The real line is a convex space, so:

Proposition 1.1 (Mixed extension). [classical] *Every function $u\colon S \times S \to \mathbb{R}$ extends uniquely to a map $\Delta(S) \times \Delta(S) \to \mathbb{R}$ that is affine in each variable separately. In coordinates, with $A = (u(s,s'))$, the extension is $(x,y) \mapsto x^{\mathsf{T}}Ay$.*

Definition 1.2. A ***symmetric game*** is a pair (S, A) of a finite set S and a matrix $A \in \mathbb{R}^{S\times S}$; the entry $A_{ss'}$ is the payoff to a player using s against a player using s'. For $x, y \in \Delta(S)$ write

$$\pi(x \mid y) = x^{\mathsf{T}}Ay$$

for the payoff of x against y, and $f(y) = Ay$ for the vector of payoffs of the pure strategies against y.

Biaffinity is what identifies the two readings: the expected payoff against an opponent drawn at random from a population in state y is the payoff against the mixed strategy y, because both are the affine extension of the same function. Whenever the source appeals to "no difference from a mathematical point of view", it is appealing to Proposition 1.1.

1.2 Maps of games

Definition 1.3. A ***payoff-preserving map*** $f\colon (S,A) \to (T,B)$ is a function $f\colon S \to T$ with $B_{f(s)f(s')} = A_{ss'}$ for all $s, s' \in S$. Symmetric games and payoff-preserving maps form a category Game.

Three kinds of map generate everything.

Inclusions. An injective f exhibits (S,A) as the *subgame* of (T,B) on the strategies in its image.

Clone maps. A surjective f exhibits (S,A) as (T,B) with some strategies duplicated: s and s' with $f(s) = f(s')$ have equal rows and equal columns in A.

Symmetries. An automorphism of (S,A) is a relabelling of strategies that preserves every payoff.

Proposition 1.4. [observation] *For a payoff-preserving $f\colon (S,A) \to (T,B)$ and $x, y \in \Delta(S)$,*

$$\pi(f_*x \mid f_*y)_B = \pi(x \mid y)_A, \qquad (Bf_*y)_{f(s)} = (Ay)_s \ \text{for every } s \in S.$$

Proof. $(Bf_*y)_{f(s)} = \sum_t B_{f(s)t} \sum_{f(s')=t} y_{s'} = \sum_{s'} B_{f(s)f(s')} y_{s'} = \sum_{s'} A_{ss'} y_{s'}$. The first identity follows by pairing with x. □

So Δ lifts to a functor from Game to convex spaces with a biaffine form, and every construction of this volume can be asked whether it respects that functor.

Example 1.5 (Rock–Paper–Scissors and its relatives)**.** Let $S = \mathbb{Z}/n$ with n odd, and let $A_{ij} = 1$ if $i - j \in \{1, \dots, \frac{n-1}{2}\}$, -1 if $j - i \in \{1, \dots, \frac{n-1}{2}\}$ and 0 if $i = j$. For $n = 3$ this is Rock–Paper–Scissors and for $n = 5$ its five-strategy extension with which the source opens [source, §1.1]. Translation $i \mapsto i + 1$ is a symmetry, so $\mathbb{Z}/n$ acts on the game. The subgame of the five-strategy game on $\{0, 1, 3\}$ is a copy of Rock–Paper–Scissors, an inclusion in Game; the subgame on $\{0, 1, 2\}$ is not, since there 2 beats both others.

A second kind of map acts on payoffs and not on strategies.

Definition 1.6. A ***column shift*** replaces A by $A + \mathbf{1}c^{\mathsf{T}}$ for a vector $c \in \mathbb{R}^S$: a constant $c_{s'}$ is added to every payoff against s'. A ***rescaling*** replaces A by λA with $\lambda > 0$.

A column shift changes $\pi(x \mid y)$ by $c^{\mathsf{T}}y$, which does not depend on x; it changes no comparison between strategies. Every game is a column shift of a unique game with zero diagonal, and the source's normalization to positive payoffs [source, §4.1] is a column shift by a constant vector.

1.3 Equilibria

Definition 1.7. A state $x \in \Delta(S)$ is a ***(symmetric) Nash equilibrium*** of (S, A) if $\pi(x \mid x) \geq \pi(y \mid x)$ for all $y \in \Delta(S)$; equivalently, if $(Ax)_s \leq x^{\mathsf{T}} Ax$ for all s, with equality for $s \in \operatorname{supp} x$. It is ***strict*** if $\pi(x \mid x) > \pi(y \mid x)$ for all $y \neq x$, and ***quasi-strict*** if $(Ax)_s < x^{\mathsf{T}} Ax$ for all $s \notin \operatorname{supp} x$.

Theorem 1.8. [classical; Nash 1951] *Every symmetric game has a symmetric Nash equilibrium. If a finite group G acts on (S, A) by symmetries, the set of Nash equilibria is G-invariant and contains a G-fixed point.*

Corollary 1.9. *If a group of symmetries acts transitively on S, the uniform state is a Nash equilibrium. In particular the uniform state is an equilibrium of every game of Example 1.5.*

How equilibria behave under the maps of Definition 1.3 is recorded once and used repeatedly.

Proposition 1.10. [observation] *Let $f \colon (S, A) \to (T, B)$ be payoff-preserving.*

(a) If f is surjective, then x is a Nash equilibrium of (S, A) if and only if $f_* x$ is a Nash equilibrium of (T, B).

(b) If f is injective, then $f_* x$ Nash in (T, B) implies x Nash in (S, A); the converse holds if and only if no strategy outside the image of f earns more than $\pi(x \mid x)$ against $f_* x$.

Proof. By Proposition 1.4 the pure payoffs against $f_* x$ at strategies in the image of f are the pure payoffs against x, and the mean payoffs agree. The equilibrium condition on T is the equilibrium condition on S together with a condition at strategies outside the image, which is vacuous when f is surjective. □

The asymmetry between (a) and (b) recurs at every level of the theory: clone maps are invisible to selection, inclusions are visible exactly through the strategies they omit.

Notes

The distribution monad and convex spaces as its algebras go back to Stone and to Świrszcz; a modern account is in [Fri09]. Theorem 1.8 is from Nash's paper on non-cooperative games [Nas51]. The category Game is not standard. Categories of games with strategy-changing morphisms appear in the literature on potential games and on open games [GHWZ18]; those morphisms are richer and are not needed here. The classical theory with rational players is

the subject of [GIHC]. The invariance of evolutionary dynamics under column shifts is used systematically in [Zee80, HoSi98].

Exercises

Exercise 1.1. Let $y_\epsilon = (1-\epsilon)x + \epsilon\mu$. Show that $\pi(x \mid y_\epsilon) - \pi(y_\epsilon \mid y_\epsilon) = \epsilon\,\big(\pi(x \mid y_\epsilon) - \pi(\mu \mid y_\epsilon)\big)$.

Exercise 1.2. Determine the group of symmetries of Rock–Paper–Scissors and of the five-strategy game of Example 1.5.

Exercise 1.3. For the Hawk–Dove game with $A = \left(\begin{smallmatrix} (V-C)/2 & V \\ 0 & V/2 \end{smallmatrix}\right)$, find the zero-diagonal normal form under column shifts and read off the equilibria for $V < C$ and for $V > C$.

Exercise 1.4. Let (S, A) be the Hawk–Dove game with the strategy Dove duplicated. Describe all its Nash equilibria for $V < C$.

Chapter 2

Evolutionary Stability

2.1 The invasion form

A Nash equilibrium asks that no deviation does better against the incumbent. In a population that is too weak: a deviation that does equally well against the incumbent is not selected against, and may spread if it does better against itself. The source's first example [source, §2.1] is the game $\left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)$, in which the second strategy is an equilibrium that the first invades by drift. All of the static theory can be phrased with one function.

Definition 2.1. The ***invasion form*** of a state x is the function on $\Delta(S)$

$$Q_x(y) = \pi(x \mid y) - \pi(y \mid y),$$

the advantage of x over y in a population in state y.

Definition 2.2. [classical; Maynard Smith and Price 1973; [source, Defs. 1–4]] A state x is

(i) an ***evolutionarily stable strategy*** (ESS) if for every $\mu \neq x$ either $\pi(x \mid x) > \pi(\mu \mid x)$, or $\pi(x \mid x) = \pi(\mu \mid x)$ and $\pi(x \mid \mu) > \pi(\mu \mid \mu)$;

(ii) ***locally superior*** if $Q_x(y) > 0$ for all $y \neq x$ in a neighbourhood of x;

(iii) protected by a ***uniform invasion barrier*** if there is $\bar{\epsilon} > 0$ such that $\pi(x \mid y_\epsilon) > \pi(\mu \mid y_\epsilon)$ for all $\mu \neq x$ and $0 < \epsilon < \bar{\epsilon}$, where $y_\epsilon = (1-\epsilon)x + \epsilon\mu$.

Theorem 2.3. [classical; Hofbauer, Schuster and Sigmund 1979; [Wei95]] *The three conditions are equivalent.*

The source presents the equivalence as evidence that the definition is not arbitrary. In the present language two thirds of it is one line: by Exercise 1.1,

$$Q_x(y_\epsilon) = \epsilon\,\big(\pi(x \mid y_\epsilon) - \pi(\mu \mid y_\epsilon)\big), \tag{2.1}$$

so the barrier inequality along the ray from x towards μ *is* local superiority along that ray. What remains is uniformity in μ, which is compactness of the simplex, and the passage from (i) to a neighbourhood, which uses that $\epsilon \mapsto \pi(x \mid y_\epsilon) - \pi(\mu \mid y_\epsilon)$ is affine with value $\pi(x \mid x) - \pi(\mu \mid x) \geq 0$ at 0 and slope governed by the second clause of (i).

Example 2.4 (Hawk–Dove)**.** With A as in Exercise 1.3 and $V < C$, let $x = (p, 1-p)$ with $p = V/C$ and $y = (q, 1-q)$. The payoff advantage of Hawk over Dove against y is $(V - qC)/2$, so

$$Q_x(y) = (p - q)\,\frac{V - qC}{2} = \frac{C}{2}\,(p - q)^2.$$

The mixed state is globally, and not only locally, superior. [S1]

Example 2.5 (Drift)**.** For $A = \left(\begin{smallmatrix} 2 & 1 \\ 1 & 1 \end{smallmatrix}\right)$ and x the second pure strategy, $Q_x(y) = 1 - (1 + q^2) = -q^2 < 0$ when y plays the first strategy with probability $q > 0$. The equilibrium is not merely not superior; it is inferior to every nearby state. [S4]

2.2 Properties

The source collects five facts [source, Thms. 1–5]. Each is a statement about Q.

Theorem 2.6. [classical; Bishop and Cannings 1976; Haigh 1975]

(a) A game need not have an ESS: Rock–Paper–Scissors has none. [S2]

(b) A weakly dominated strategy is not an ESS. [S3]

(c) If x is an ESS and $y \neq x$ is a Nash equilibrium, then $\operatorname{supp} y \not\subseteq \operatorname{supp} x$.

(d) A game with n strategies has at most $\binom{n}{\lfloor n/2 \rfloor}$ ESS.

(e) An ESS in the interior of the simplex is the only ESS, and it is globally superior.

(f) A quasi-strict equilibrium x is an ESS if and only if $z^{\mathsf{T}} A z < 0$ for all $z \neq 0$ with $\operatorname{supp} z \subseteq \operatorname{supp} x$ *and* $\sum_s z_s = 0$.

Proof. (a) By Corollary 1.9 the uniform state u is an equilibrium, and it is the only one. For a pure strategy e, $\pi(u \mid e) = 0 = \pi(e \mid e)$ and $\pi(e \mid u) = 0 = \pi(u \mid u)$, so both clauses of (i) fail. (b) If μ weakly dominates x then $\pi(\mu \mid x) \geq \pi(x \mid x)$ and $\pi(\mu \mid \mu) \geq \pi(x \mid \mu)$, so $Q_x(\mu) \leq 0$ with μ an alternative best reply. (c) If $\operatorname{supp} y \subseteq \operatorname{supp} x$ then y is a best reply to x, since every pure strategy in the support of x is. The second clause of (i) gives $\pi(x \mid y) > \pi(y \mid y)$, contradicting that y is an equilibrium. (d) By (c) the supports of distinct ESS form an antichain of subsets of S; apply Sperner's theorem. (e) If x is interior, every y is a best reply to x, so $Q_x(y) > 0$ for all $y \neq x$; then no other state is an equilibrium by (c). (f) For y near x write $y - x = z + w$ with z supported in $\operatorname{supp} x$; quasi-strictness makes the first-order term in w strictly positive, and for $w = 0$, $Q_x(y) = -z^{\mathsf{T}} A z$ because $z^{\mathsf{T}} A x = 0$. □

Criterion (f) is the one used by the companion script to decide evolutionary stability exactly.

2.3 Sets, and why they are forced

A population may be stable as a whole while drifting inside a set of equivalent states. The source surveys Thomas's evolutionarily stable sets and three variants due to Balkenborg and Schlag [source, §2.3]. We use the following form.

Definition 2.7. [classical; Thomas 1985; [Wei95]] A non-empty closed set $L \subseteq \Delta(S)$ is an ***evolutionarily stable set*** if every $x \in L$ has a neighbourhood U with $Q_x(y) \geq 0$ for all $y \in U$, and $Q_x(y) > 0$ for $y \in U \setminus L$.

A one-point set $\{x\}$ is an evolutionarily stable set exactly when x is an ESS, and every point of an evolutionarily stable set is a Nash equilibrium.

Theorem 2.8 (Stability is a property of sets)**.** [observation] *Let $f\colon (S,A) \to (T,B)$ be a payoff-preserving surjection and let $L \subseteq \Delta(T)$ be an evolutionarily stable set. Then $f_*^{-1}(L)$ is an evolutionarily stable set of (S,A). If f is not injective and x is an ESS of (T,B) whose support meets the set of duplicated strategies, then no point of $f_*^{-1}(x)$ is an ESS.*

Proof. By Proposition 1.4, $Q_\sigma(\mu) = Q_{f_*\sigma}(f_*\mu)$ for all $\sigma, \mu \in \Delta(S)$. Let $\sigma \in f_*^{-1}(L)$ and let U be the neighbourhood of $f_*\sigma$ from Definition 2.7. Then $f_*^{-1}(U)$ is a neighbourhood of σ on which $Q_\sigma \geq 0$, with strict inequality off $f_*^{-1}(L)$; and $f_*^{-1}(L)$ is closed and non-empty. For the second statement, the fibre over x contains a segment through each of its points, on which Q_σ vanishes identically. □

Example 2.9. Duplicate Dove in the Hawk–Dove game. The fibre over the ESS is the segment $L = \{(p, s, 1-p-s)\}$, and $Q_\sigma(\mu) = \frac{C}{2}(p - \mu_{\mathrm{H}})^2$ for $\sigma \in L$: zero exactly on L, positive elsewhere. [S6]

Theorem 2.8 explains a development that the source records historically. Duplicating a strategy changes nothing about the biology or the sociology of a model; it is a change of description. A solution concept that is destroyed by it cannot be the right concept, and the least repair is to pass from points to sets. The three variants of Balkenborg and Schlag differ in the order of quantifiers over the neighbourhood (one threshold for each invader, for each incumbent, or for all) and so in how they behave when L is not compact in the relevant topology; for the closed sets of Definition 2.7 we do not need the distinction.

2.4 Relative entropy

The invasion form has a second life. For $x, y \in \Delta(S)$ with $\operatorname{supp} x \subseteq \operatorname{supp} y$ let

$$D(x \,\|\, y) = \sum_{s} x_s \log \frac{x_s}{y_s}$$

be the relative entropy. It is the quantity singled out in [ESCHC], following Baez and Fritz, as the unique functorial measure of information lost by a stochastic map. In Chapter 3 we shall see that along the replicator dynamics

$$\frac{d}{dt} D(x \,\|\, y(t)) = -Q_x(y(t)),$$

so that local superiority of x says precisely that relative entropy to x decreases strictly along nearby trajectories. Evolutionary stability is the statement that a population near x *gains information about* x.

Notes

The definition of an ESS is from [MSP73]; the equivalences of Theorem 2.3 are in [HSS79] and [Wei95, Ch. 2]. Parts (c)–(e) of Theorem 2.6 are due to Bishop and Cannings [BiCa76], part (f) to Haigh [Hai75]; the sharp maximal number of ESS for given n is a harder problem than the Sperner bound suggests. Evolutionarily stable sets were introduced by Thomas [Tho85] and refined in [BaSco1]. The observation that they are the closure of the point concept under duplication of strategies is, as far as we know, not stated in this form in the literature, though it is implicit in the treatment of "clones" in the literature on stability of sets. The characterization of relative entropy is in [BaFr14].

Exercises

Exercise 2.1. Show that a strict Nash equilibrium is an ESS, and that in a two-strategy game with $A_{11} \neq A_{21}$ and $A_{22} \neq A_{12}$ every ESS is either strict or interior.

Exercise 2.2. For the generalized Rock–Paper–Scissors game with payoff $a > 0$ for a win and $-b < 0$ for a loss, show that the uniform state is an ESS if and only if $a > b$.

Exercise 2.3. Exhibit a game with three strategies and three ESS, attaining the bound of Theorem 2.6(d).

Exercise 2.4. In the game $\left(\begin{smallmatrix} 1 & 100 \\ 0 & 100 \end{smallmatrix}\right)$ of the source's discussion of weak dominance, compute Q_x for each pure strategy x and decide which is an ESS.

Part II

Dynamics

Chapter 3

The Replicator Equation

3.1 Derivation as a quotient

Let $n_s(t) > 0$ be the number of individuals using strategy s, and suppose each type grows at a per-capita rate equal to its payoff against the current population: $\dot{n}_s = n_s f_s(x)$ with $x = n/|n|$, $|n| = \sum_s n_s$, $f(x) = Ax$. Differentiating $x_s = n_s/|n|$ gives the ***replicator equation*** [TaJo78]

$$\dot{x}_s = x_s\big((Ax)_s - x^{\mathsf{T}}Ax\big), \qquad s \in S. \tag{3.1}$$

Proposition 3.1. [reformulation; cf. [HoSi98]] *The growth system on the open cone $\mathbb{R}^S_{>0}$ is invariant under the scaling action of $\mathbb{R}_{>0}$ up to a change of speed along orbits, and (3.1) is the vector field it induces on the quotient $\mathbb{R}^S_{>0}/\mathbb{R}_{>0} \cong \operatorname{int}\Delta(S)$. The field extends smoothly to the closed simplex.*

The source derives (3.1) twice, from births and deaths and from imitation [source, §3.2.1–3.2.2]. The first is Proposition 3.1; the second is Chapter 5. We write $\mathrm{Rep}(S, A)$ for the dynamical system on $\Delta(S)$ defined by (3.1), and Dyn for the category whose objects are compact spaces with a continuous flow and whose morphisms are continuous maps commuting with the flows.

3.2 Naturality

Theorem 3.2 (The replicator equation is a functor)**.** [observation] *If $f\colon (S, A) \to (T, B)$ is payoff-preserving, then $f_*\colon \Delta(S) \to \Delta(T)$ maps solutions of $\mathrm{Rep}(S, A)$ to solutions of $\mathrm{Rep}(T, B)$. Hence $\mathrm{Rep}\colon \mathsf{Game} \to \mathsf{Dyn}$ is a functor.*

Proof. Let $x(t)$ solve (3.1) for A and put $y = f_*x$. By Proposition 1.4, $(Ax)_s = (By)_{f(s)}$ and $x^{\mathsf{T}}Ax = y^{\mathsf{T}}By$. Therefore

$$\dot{y}_t = \sum_{f(s)=t} x_s\big((Ax)_s - x^{\mathsf{T}}Ax\big) = \Big(\sum_{f(s)=t} x_s\Big)\big((By)_t - y^{\mathsf{T}}By\big) = y_t\big((By)_t - y^{\mathsf{T}}By\big)$$

for t in the image of f, and $y_t = 0$ identically otherwise, which also solves the equation. □

Three familiar facts are the three generating cases of Definition 1.3. [R1]

Corollary 3.3. *(a)* Faces are invariant. *For every $S' \subseteq S$ the face $\Delta(S')$ is invariant, and the flow on it is the replicator flow of the subgame. In particular strategies absent initially stay absent.*

(b) Clones lump. *If strategies are duplicated, the total frequency of each class of duplicates follows the replicator equation of the game without duplicates.*

(c) Symmetries act. *The group of symmetries of (S, A) acts on $\Delta(S)$ by automorphisms of the flow; the fixed-point set of every subgroup is invariant.*

On the side of payoffs, (3.1) depends only on differences $(Ax)_s - (Ax)_{s'}$, so:

Proposition 3.4. [classical] *A column shift does not change the replicator field, and a rescaling $A \mapsto \lambda A$ multiplies it by λ. The functor* Rep *factors through games modulo column shifts.*

3.3 Two strategies

By Proposition 3.4 a two-strategy game may be taken to be $\left(\begin{smallmatrix} 0 & a \\ b & 0 \end{smallmatrix}\right)$. With x the frequency of the first strategy, (3.1) reads

$$\dot{x} = x(1-x)\big(a - (a+b)x\big). \tag{3.2}$$

signs of (a, b)	in the source	rest points	behaviour
$a > 0 > b$ or $b > 0 > a$	Prisoner's Dilemma	$0, 1$	one strategy wins
$a, b > 0$	Hawk–Dove	$0, 1, \frac{a}{a+b}$	interior attracts
$a, b < 0$	coordination	$0, 1, \frac{a}{a+b}$	interior repels
$a = b = 0$	"game of life"	every point	nothing moves

The drift game of Chapter 2 is the boundary case $a = 0 > b$: there $\dot{x} = x^2(1-x)$ for the invading strategy, and the invader's frequency passes from x_0 to x_1 in time

$$\frac{1}{x_0} - \frac{1}{x_1} + \log \frac{x_1(1-x_0)}{x_0(1-x_1)}.$$

An invader that is merely not selected against escapes in time of order $1/x_0$; one that is selected for escapes in time of order $\log(1/x_0)$. From $x_0 = 0.05$ to $\frac{1}{2}$ the drift game takes time 20.94. [S4]

3.4 Geometry: gradients and entropy

The ***Shahshahani metric*** on the open simplex is $g_x(u, v) = \sum_s u_s v_s / x_s$ for tangent vectors u, v (vectors with $\sum_s u_s = 0$). It is the Fisher information metric of statistics.

Theorem 3.5. [classical; Shahshahani 1979; Kimura 1958] *If A is symmetric, (3.1) is the gradient flow of $\frac{1}{2}x^{\mathsf{T}}Ax$ for the Shahshahani metric, and mean payoff increases at the rate*

$$\frac{d}{dt}\,x^{\mathsf{T}}Ax = 2\,\mathrm{Var}_x(Ax) = 2\sum_s x_s\big((Ax)_s - x^{\mathsf{T}}Ax\big)^2.$$

This is Fisher's fundamental theorem in the form relevant to games in which the two players' interests coincide. [R2]

Theorem 3.6. [classical; [HSS79, Zee80]] *Along any solution $y(t)$ of (3.1) in the interior of a face containing* $\operatorname{supp} x$*,*

$$\frac{d}{dt}D(x\,\|\,y(t)) = -Q_x(y(t)).$$

Proof. $\frac{d}{dt}D(x\,\|\,y) = -\sum_s x_s\dot{y}_s/y_s = -\sum_s x_s\big((Ay)_s - y^{\mathsf{T}}Ay\big) = -(\pi(x\mid y) - \pi(y\mid y))$. □

Corollary 3.7. [source, Thm. 6.3] *An ESS is asymptotically stable; an interior ESS attracts the whole interior of the simplex. More precisely, if $Q_x > 0$ on $U\setminus\{x\}$, every sublevel set $\{y : D(x\,\|\,y) \le c\}$ contained in U lies in the basin of attraction of x. For Hawk–Dove, $\frac{d}{dt}D(x\,\|\,y) = -\frac{C}{2}(p-q)^2$.* [R3]

3.5 The folk theorem and elimination

Theorem 3.8. [classical; [HoSi98]; [source, Thm. 6]] *(1) Every Nash equilibrium is a rest point of (3.1). (2) Every Lyapunov stable rest point is a Nash equilibrium. (3) Every ESS is asymptotically stable. No converse holds.*

For (1) see [R7]. A pure state is always a rest point by Corollary 3.3(a) and need not be an equilibrium. For the failure of the converse of (3) the source gives the game $\left(\begin{smallmatrix}1&1&1\\0&0&3\\2&1&0\end{smallmatrix}\right)$: its only equilibrium is the uniform state, which is asymptotically stable and is not an ESS, since criterion (f) of Theorem 2.6 fails. [R8]

Theorem 3.9 (Elimination, with a rate)**.** [classical; Akin 1980; [source, Thm. 7]] *Suppose the pure strategy s is strictly dominated by $\mu\in\Delta(S)$ with margin $\delta>0$: $\pi(\mu\mid z) - (Az)_s \ge \delta$ for all z. Then along every interior solution*

$$x_s(t) \le \frac{x_s(0)}{\prod_r x_r(0)^{\mu_r}}\,e^{-\delta t}.$$

Proof. $\frac{d}{dt}\big(\log x_s - \sum_r \mu_r \log x_r\big) = (Ax)_s - \pi(\mu\mid x) \le -\delta$, and $\prod_r x_r^{\mu_r} \le 1$. □

When s is a twin of another strategy s' with every payoff lowered by δ, the proof gives an identity: $\log(x_s/x_{s'})$ decreases at the constant rate δ. [R9] The hypothesis that the solution be interior is the source's proviso that all strategies be initially present, and by Corollary 3.3(a) it cannot be dropped.

3.6 Cycles

Theorem 3.10. [classical; [HoSi98]] *Let A be antisymmetric with an interior equilibrium x^*. Then $D(x^* \,\|\, y)$ is constant along every interior solution of (3.1). No interior solution converges to x^* or to the boundary.*

Proof. $y^{\mathsf{T}}Ay = 0$ and $x^{*\mathsf{T}}Ay = -y^{\mathsf{T}}Ax^* = 0$, so $Q_{x^*} \equiv 0$; apply Theorem 3.6. □

For Rock–Paper–Scissors this gives the family of closed orbits around the centre. [R4]

Proposition 3.11 (Generalized Rock–Paper–Scissors). [classical; [Zee80]] *Let a win pay $a > 0$ and a loss $-b < 0$. The linearization of (3.1) at the uniform state has eigenvalues*

$$\lambda_{\pm} = \frac{b-a}{6} \pm i\,\frac{\sqrt{3}\,(a+b)}{6}.$$

The centre is asymptotically stable and an ESS for $a > b$, a centre for $a = b$, and a repeller for $a < b$, when interior orbits approach the boundary. [R5, R6]

The cases $a > b$ and $a < b$ are the source's "good" and "bad" versions. By Corollary 3.3(c) the uniform state is a rest point for every a, b, being the fixed point of the $\mathbb{Z}/3$ of symmetries. For three strategies the possible phase portraits were classified by Zeeman and by Bomze [Zee80, Bom83]; they are assembled from rest points and cycles. With four strategies the state space is three-dimensional and chaotic attractors occur [ASS80]; the source illustrates one [source, §3.2.5]. Two statements survive chaos.

Theorem 3.12. [classical; [HoSi98]] *If a solution of (3.1) stays at a positive distance from the boundary, then every limit point of its time averages $\frac{1}{T}\int_0^T x(t)\,dt$ is an interior equilibrium. If the interior equilibrium is unique, the time averages converge to it.*

The second surviving statement is the index theorem of the next chapter.

Notes

The replicator equation was introduced in [TaJo78] and named in [ScSi83]. The standard reference is [HoSi98]; [San10] treats it within a general theory of population dynamics. Theorem 3.5 is in [Sha79, Aki79]. The functorial statement

Theorem 3.2 is elementary and its three cases are classical; we have not found it stated as a single naturality property. Relative entropy as a Lyapunov function, and the reading of the replicator equation as an information-geometric flow, are developed in [BaPo16, Haro9]. The equivalence of the replicator equation on n strategies with a Lotka–Volterra system in $n-1$ variables is due to Hofbauer [Hof81].

Exercises

Exercise 3.1. Derive (3.2) and the escape-time formula for the drift game.

Exercise 3.2. Prove the formula for $\frac{d}{dt} x^{\mathsf{T}} A x$ in Theorem 3.5, and show that for general A the rate is $2\,\mathrm{Cov}_x(Sx, Ax)$, where $S = \frac{1}{2}(A + A^{\mathsf{T}})$.

Exercise 3.3. Verify Proposition 3.11 by computing the trace and determinant of the linearization on the tangent space.

Exercise 3.4. A population plays Rock–Paper–Scissors with Rock present in two indistinguishable variants, in initial state $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$. Describe the motion of the class frequencies.

Chapter 4

Index, Recurrence, and Decomposition

Three questions cannot be answered one equilibrium at a time. How many equilibria can there be, and of what kinds? What replaces an equilibrium when the dynamics cycles? Which part of a game is responsible for the cycling? Each has an answer in which topology does the work.

4.1 The index theorem

Write F_A for the replicator field. At a rest point x its derivative preserves the tangent space $T = \{v : \sum_s v_s = 0\}$ of the simplex.

Definition 4.1. A game is ***regular*** if every Nash equilibrium is quasi-strict and $D_x F_A|_T$ is invertible at each. The ***index*** of a Nash equilibrium x of a regular game is

$$\operatorname{ind}(x) = \operatorname{sign}\det\big(-D_x F_A|_T\big) \in \{+1, -1\}.$$

Regular games are open and dense among all games. An equilibrium at which every eigenvalue of $D_x F_A|_T$ has negative real part has index $+1$.

Theorem 4.2 (Index theorem). [classical; Hofbauer and Sigmund; Poincaré–Hopf] *For a regular game,*

$$\sum_{x \text{ Nash}} \operatorname{ind}(x) = 1.$$

Sketch. For small $\epsilon > 0$ the field $F_A(x) + \epsilon(u - x)$, with u the uniform state, points strictly inward on the boundary of the simplex. Its rest points converge, as $\epsilon \to 0$, to the rest points of F_A that are *saturated*, that is, to the Nash equilibria: a rest point with a strategy outside its support that earns more than average is pushed out of the simplex and disappears. Regularity makes the correspondence a bijection that preserves the sign of the determinant. The negative of an inward-pointing field points outward, and by the Poincaré–Hopf theorem the sum of the indices of its zeros is the Euler characteristic of the simplex, which is 1 because the simplex is contractible; see [DTHC, Ch. 7] for the three descriptions of the Euler characteristic used here. □

The theorem is checked on four hundred random games with three and four strategies. [l1]

Corollary 4.3. *In a regular game: (a) the number of Nash equilibria is odd; (b) every ESS has index $+1$* [I2]*; (c) if there are k equilibria of index $+1$ there are exactly $k-1$ of index -1; in particular two ESS force a third equilibrium, which is not stable.* [I3]

In the coordination game $\left(\begin{smallmatrix} 2 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ the two conventions have index $+1$ and the mixed equilibrium, which separates their basins, has index -1. The index theorem is the reason a model with two stable conventions always has a threshold between them. It holds whatever the dynamics does elsewhere, cycles and chaos included.

4.2 Recurrence and the Conley index

When trajectories do not converge, the object that plays the part of an equilibrium is a recurrent set.

Definition 4.4. For a flow φ on a compact metric space, an (ϵ, τ)***-chain*** from x to y is a sequence $x = x_0, \dots, x_k = y$ with times $t_i \geq \tau$ such that $d(\varphi_{t_i}(x_i), x_{i+1}) < \epsilon$. A point is ***chain recurrent*** if for all ϵ, τ there is a chain from it to itself; two chain recurrent points are equivalent if chains run both ways. The classes are the ***chain components***.

Theorem 4.5 (Conley's decomposition)**.** [classical; Conley 1978] *Every flow on a compact metric space has a continuous function that is strictly decreasing along orbits off the chain recurrent set and constant on each chain component. The flow is gradient-like between chain components.*

Chain components generalize equilibria in exactly the way the source's examples require.

Example 4.6. For the Prisoner's Dilemma the chain components are the two pure states, and All-Defect is the only one that attracts. For a coordination game there are three. For Rock–Paper–Scissors with $a = b$ every state is chain recurrent and the whole simplex is a single chain component: arbitrarily small errors carry the population from any closed orbit to any other, and to the boundary. For $a < b$ the chain components are the centre and the boundary cycle, and the boundary cycle attracts.

Papadimitriou and Piliouras propose the chain components of a learning dynamic, and the sink components among them, as what a game "means" when equilibria do not predict behaviour [PaPi19].

To each isolated invariant set Conley attaches more than a number.

Definition 4.7. An ***index pair*** for an isolated invariant set K is a pair of compact sets $L \subseteq N$ such that K is the largest invariant set in $N \setminus L$, L is positively

invariant in N, and orbits leave N only through L. The ***Conley index*** of K is the pointed homotopy type $h(K) = [N/L]$.

Theorem 4.8. [classical; Conley 1978; [Sal85]] *The Conley index does not depend on the index pair. It is invariant under continuation: if a family of flows has a common isolating neighbourhood, the index is constant along the family. A hyperbolic rest point with k unstable directions has index the pointed sphere S^k; an attracting closed orbit has the index of a circle with a disjoint base point, whose reduced homology has rank one in degrees 0 and 1. If $h(K)$ is not the homotopy type of a point, K is non-empty.*

Remark 4.9 (What is higher here). The index of Theorem 4.2 is the Euler characteristic of the Conley index: an equilibrium with k unstable directions contributes $(-1)^k$. The Conley index remembers k, and for recurrent sets it remembers a homotopy type that no count recovers. Continuation invariance is the precise sense in which a qualitative conclusion drawn from a model survives perturbation of its payoffs: as long as an attractor stays isolated, its index, and hence its existence, persists. An ESS has index S^0 and persists under small changes of the game; the family of cycles of Rock–Paper–Scissors at $a = b$ is not isolated and does not.

For computation one discretizes. The ***invasion graph*** of a game has the pure strategies as vertices and an edge $i \to j$ when $A_{ji} > A_{ii}$, that is, when j invades a population of i. Its sink strongly connected components are a finite shadow of the attracting chain components. [I7] The Markov–Conley chains of α-Rank [OPP19] are random walks on such graphs, weighted by fixation probabilities; we return to them in Chapter 6.

4.3 Decomposition

Every matrix splits as $A = S + K$ with $S = \frac{1}{2}(A + A^{\mathsf{T}})$ symmetric and $K = \frac{1}{2}(A - A^{\mathsf{T}})$ antisymmetric: a game of common interest plus a zero-sum game. By Theorems 3.5 and 3.10 the replicator field is the sum of a gradient field and a field that, when K has an interior equilibrium, conserves a relative entropy. Mean payoff obeys $\frac{d}{dt} x^{\mathsf{T}} A x = 2\,\mathrm{Cov}_x(Sx, Ax)$. [I6]

The antisymmetric part decomposes further, and here the statement is cohomological. Regard an antisymmetric matrix K as a 1-cochain on the complete graph with vertex set S: K_{ij} is the advantage of i over j.

Theorem 4.10 (Hodge decomposition of a zero-sum game). [literature; [JLYY11, CMOP11]] *Every antisymmetric K decomposes uniquely as*

$$K_{ij} = (r_i - r_j) + C_{ij}, \qquad r_i = \tfrac{1}{n}\textstyle\sum_j K_{ij}, \qquad \sum_j C_{ij} = 0 \ \text{for all } i.$$

The two parts are orthogonal for the Frobenius inner product. The gradient part satisfies the cocycle condition $G_{ij} + G_{jk} + G_{ki} = 0$ on every triangle, and conversely every antisymmetric matrix satisfying it is a gradient. The cyclic part C has the same curl $K_{ij} + K_{jk} + K_{ki}$ as K on every triangle.

The converse statement is the vanishing of the first cohomology of the full simplex on S: a closed cochain is exact. [14, 15] We call $\tau(K) = \|G\|^2 / \|K\|^2 \in [0,1]$ the ***transitivity index***. Rock–Paper–Scissors has $\tau = 0$ and curl ± 3; a game of pure rating differences has $\tau = 1$; a random antisymmetric game on five strategies had $\tau = 0.56$.

Remark 4.11. If only some pairs of strategies are ever compared, the complete graph is replaced by a comparison graph with its triangles, and its first cohomology need not vanish. A third, *harmonic*, component appears: cyclic advantage around a loop that bounds no triangle, locally consistent and globally not a rating. This is the only place in the volume where a cohomology group is not zero for a trivial reason.

Notes

The index theorem for the replicator equation is in [HoSi98, Ch. 13]; the perturbation argument is Hofbauer's. Conley's theory is in [Con78], with an accessible account of the index in [Sal85]. Chain recurrence as a solution concept for games is argued in [PaPi19]. The Hodge decomposition of pairwise comparison data is from [JLYY11]; its game-theoretic form for general finite games, into potential, harmonic and non-strategic parts, is from [CMOP11]; the use of the transitive and cyclic parts in the evaluation of agents is from [BTP18].

Exercises

Exercise 4.1. Compute the index of each equilibrium of the three generic classes of two-strategy games in the table of Chapter 3, and verify Theorem 4.2.

Exercise 4.2. Show that an antisymmetric K is a gradient if and only if its curl vanishes on all triangles through one fixed vertex.

Exercise 4.3. Compute the Hodge decomposition and the transitivity index of the antisymmetric game on three strategies with $K_{12} = 1$, $K_{23} = 1$, $K_{13} = t$, as a function of t.

Exercise 4.4. Determine the invasion graph and its sink components for the Hawk–Dove game with $V < C$. What do they miss?

Chapter 5

Revision Protocols, Other Dynamics, and Open Systems

The replicator equation is one dynamic among many. The source derives it a second time from imitation, introduces two rivals, and reports a theorem that separates them [source, §3.2.2, §3.3]. This chapter places all three in one frame and then takes the frame apart into components.

5.1 Revision protocols

A ***population game*** is a continuous map $F\colon \Delta(S) \to \mathbb{R}^S$; the linear case $F(x) = Ax$ is the symmetric game (S, A). Following Sandholm [San10], a ***revision protocol*** is a continuous map

$$\rho\colon \mathbb{R}^S \times \Delta(S) \longrightarrow \mathbb{R}^{S\times S}_{\geq 0},$$

where $\rho_{ij}(\pi, x)$ is the rate at which a player of strategy i switches to j when payoffs are π and the state is x; the ***mean dynamic*** it induces is

$$\dot{x}_i = \sum_j x_j\,\rho_{ji}(F(x), x) - x_i \sum_j \rho_{ij}(F(x), x). \tag{5.1}$$

Remark 5.1. [reformulation; cf. [San10]] For a fixed population state the matrix ρ is the rate matrix of a continuous-time Markov chain on S, and (5.1) is its forward equation $\dot{x} = xL$. A revision protocol is thus a map from states to generators, $x \mapsto L(x)$, and the mean dynamic is $\dot{x} = x\,L(x)$: a Markov chain whose transition rates depend on its own law. The chain describes one agent; the equation describes the population. This is the sense in which the continuous models of the source are limits of the finite ones of Chapter 6.

Proposition 5.2 (Imitation gives the replicator equation). [classical; [Sch98, San10]] *Under* pairwise proportional imitation, $\rho_{ij} = x_j[\pi_j - \pi_i]_+$, *the mean dynamic is the replicator equation.* [O1]

Proof. $\dot{x}_i = x_i \sum_j x_j\big([\pi_i - \pi_j]_+ - [\pi_j - \pi_i]_+\big) = x_i \sum_j x_j(\pi_i - \pi_j) = x_i(\pi_i - \bar{\pi})$. □

The factor x_j is the probability of meeting a j-player; the agent needs to know nothing about the game except what it and the player it meets have earned. Two rivals, the Brown–von Neumann–Nash (BNN) dynamic and the Smith dynamic, dispense with the meeting.

dynamic	protocol ρ_{ij}	mean dynamic $\dot{x}_i$
replicator	$x_j[\pi_j - \pi_i]_+$	$x_i(\pi_i - \bar{\pi})$
BNN	$[\pi_j - \bar{\pi}]_+ =: k_j$	$k_i - x_i \sum_j k_j$
Smith	$[\pi_j - \pi_i]_+$	$\sum_j \big(x_j[\pi_i - \pi_j]_+ - x_i[\pi_j - \pi_i]_+\big)$

5.2 Imitation is natural and innovation is not

Hofbauer and Sandholm isolate four properties of a dynamic $\dot{x} = V_F(x)$ [HoSa11]; the source states them as the hypotheses of its Theorem 8. *Continuity*: V_F depends continuously on F. *Positive correlation*: $V_F(x) \cdot F(x) > 0$ whenever $V_F(x) \neq 0$. *Nash stationarity*: $V_F(x) = 0$ exactly at Nash equilibria. *Innovation*: if x is not a Nash equilibrium and an unused strategy is a best reply to x, that strategy grows.

	continuity	pos. correlation	Nash stationarity	innovation
replicator	yes	yes	no	no
BNN	yes	yes	yes	yes
Smith	yes	yes	yes	yes

At All-Cooperate in the Prisoner's Dilemma the replicator field vanishes and the other two do not [O2]; at every Nash equilibrium of sixty random games both of the others vanish. [O3]

Theorem 5.3 (Hofbauer–Sandholm). [literature; [HoSa11]; source, Theorem 8] *For every dynamic satisfying the four properties there is a game with a strictly dominated strategy that survives along solutions from most initial conditions.*

The example has the same shape for every such dynamic: a game with a cycle, and a twin of one strategy made slightly worse. In bad Rock–Paper–Scissors (win 1, loss 2) with a twin of Scissors that earns 0.1 less, the twin's time-averaged share under the Smith dynamic is 0.116; under the replicator equation it is zero, and Theorem 3.9 gives the rate. [O4]

The functor of Chapter 3 explains the difference.

Proposition 5.4 (Names and individuals). [observation]

(a) *Every imitative protocol $\rho_{ij} = x_j\, r(\pi_i, \pi_j)$ has a mean dynamic that is natural for all payoff-preserving maps: faces are invariant and clones lump, exactly as in Corollary 3.3.*

(b) *A dynamic that is natural for subgame inclusions leaves every face invariant, so every pure state is a rest point. In any game with a pure state that is not a Nash equilibrium it violates Nash stationarity and Innovation.*

(c) *The Brown–von Neumann–Nash and Smith dynamics are natural for symmetries and for neither of the other two generating cases of Definition 1.3: an absent strategy can appear, and duplicating a strategy changes the motion of the class totals.* [O7]

Proof. (a) The mean dynamic is $\dot{x}_i = x_i \sum_j x_j \big(r(\pi_j, \pi_i) - r(\pi_i, \pi_j)\big)$. The factor x_i gives invariance of faces, and the sum depends on x only through the distribution of payoffs, which is unchanged when clones are merged. (b) is immediate from the definitions. (c) In both protocols the inflow into strategy i does not carry the factor x_i. A strategy with two names receives the inflow twice: for Rock–Paper–Scissors with Rock duplicated the discrepancy between the pushed-forward field and the field of the quotient game reached 0.56 for Smith and 0.38 for Brown–von Neumann–Nash over fifty random states. □

Remark 5.5. Imitation counts individuals; a direct protocol counts names. Under imitation a strategy is as visible as it is common, so a twin of a strategy that exists is invisible as a separate thing, and a dominated twin dies at the rate of its penalty. Under a direct protocol every name is a candidate at every moment whatever its frequency, so a twin is a second target for switching, and it is replenished each time the cycle makes it nearly optimal. Theorem 5.3 is the price of Nash stationarity: one cannot have both the invariance of Theorem 3.2 and rest points that are exactly the Nash equilibria.

5.3 Populations as open systems

In (3.1) the game and the learning rule are fused. They can be separated.

Definition 5.6. A ***population box*** on S is an open dynamical system with input a payoff vector $\pi \in \mathbb{R}^S$, state $x \in \Delta(S)$, output the state, and dynamics $\dot{x} = V(x, \pi)$. The ***replicator box*** has $V_i = x_i(\pi_i - x \cdot \pi)$.

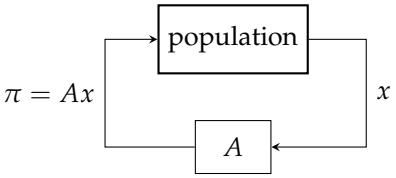
Proposition 5.7 (The replicator box integrates and normalizes)**.** [classical] *For any integrable input $\pi(\cdot)$ and interior initial state, the replicator box has the solution*

$$x(t) = \operatorname{softmax}\Big(\log x(0) + \int_0^t \pi(\tau)\, d\tau\Big).$$

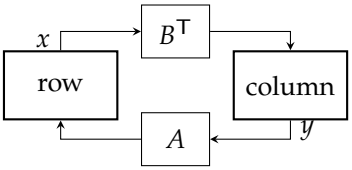
In the coordinates $z = \log x$ modulo constants the box is the integrator $\dot{z} = \pi$. [O6]

All the nonlinearity of the replicator equation is in the game and in the readout, the map softmax from $\mathbb{R}^S/\mathbb{R}\mathbf{1}$ onto the open simplex; the box itself is linear. The Lyapunov function of Theorem 3.6 is the Bregman divergence of the convex function whose gradient is this readout.

A game is then a way of closing boxes.



one population



two populations

Remark 5.8. [reformulation; cf. [SSV20, Mye21]] Open dynamical systems are the algebras of an operad of wiring diagrams: a diagram with boxes of given interfaces determines a closed system once each box is filled. A symmetric game is the one-box diagram with feedback through A; a bimatrix game is the two-box diagram; a local-interaction model (Chapter 7) has one box per vertex of a graph. Replacing the replicator box by a Smith box changes the dynamics and not the diagram. The statement that the source's models differ "only in the dynamics" or "only in the interaction structure" is the statement that a model is a pair of a diagram and a filling.

5.4 Two populations

For a bimatrix game (A, B) the two-box system with replicator boxes is

$$\dot{x}_i = x_i\big((Ay)_i - x^\mathsf{T}Ay\big), \qquad \dot{y}_j = y_j\big((B^\mathsf{T}x)_j - y^\mathsf{T}B^\mathsf{T}x\big)$$

on $\Delta(S) \times \Delta(T)$. The single-population equation for A is its restriction to the diagonal in the case $B = A^\mathsf{T}$: the diagonal is the fixed set of the symmetry exchanging the two boxes.

Theorem 5.9. [classical; [Sel80, RiWe95, HoSi98]] *In the two-population replicator dynamics a rest point is asymptotically stable if and only if it is a strict Nash equilibrium. In particular no mixed equilibrium is asymptotically stable.*

After a change of velocity the flow in the interior preserves volume, which excludes interior attractors. For the source's Battle of the Sexes, with payoffs $A = \left(\begin{smallmatrix} 2 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ to Row and $B = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 2 \end{smallmatrix}\right)$ to Column, the mixed equilibrium $(x, y) = (\frac{2}{3}, \frac{1}{3})$ is a saddle with eigenvalues $\pm\frac{2}{3}$, and orbits started near it end at one

of the two pure equilibria. [05] The population model selects conventions and never the mixed equilibrium, which is the content of the source's discussion [source, §3.4].

Notes

Revision protocols and mean dynamics are developed systematically in [San10]; proportional imitation is due to Schlag [Sch98] and, in a different form, to Björnerstedt and Weibull [BjWe96]. The Brown–von Neumann–Nash dynamic is from [BvN50] and the Smith dynamic from [Smi84]. Elimination of dominated strategies under imitative dynamics is in [Nac90, SaZh92]. For open dynamical systems and wiring diagrams see [SSV20, Mye21]; for games assembled compositionally from open pieces, [GHWZ18] and [GIHC].

Exercises

Exercise 5.1. Show that *imitation driven by dissatisfaction*, $\rho_{ij} = x_j(K - \pi_i)$ with K larger than every payoff, also has the replicator equation as its mean dynamic.

Exercise 5.2. Show directly that the rest points of the Brown–von Neumann–Nash dynamic are exactly the Nash equilibria.

Exercise 5.3. Prove Proposition 5.7.

Exercise 5.4. Write the two-population replicator equations for the Battle of the Sexes above in the coordinates (x, y), and verify the eigenvalues $\pm\frac{2}{3}$ at the mixed equilibrium.

Part III

Structure

Chapter 6

Finite Populations

A real population is finite, and in a finite population chance decides what the differential equation leaves open. The source's model is the frequency-dependent birth–death process [source, §4.1]. Its analysis is classical, and we shall say plainly at the end how little the categorical language adds to it. What it does add is a precise statement of the one point at which the source draws a conclusion that needs qualifying.

6.1 The Moran process

Two strategies A and B have payoffs $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, all positive. In a population of N with i players of type A, the expected payoffs against a randomly chosen other player are

$$F_i = \frac{a(i-1) + b(N-i)}{N-1}, \qquad G_i = \frac{c\,i + d(N-i-1)}{N-1},$$

and the fitnesses are $f_i = 1 - w + wF_i$, $g_i = 1 - w + wG_i$ for an intensity of selection $w \in (0,1]$. At each step one individual is chosen to reproduce with probability proportional to fitness, and its offspring replaces a uniformly chosen individual. The number of A-players is a Markov chain on $\{0,\dots,N\}$, that is, an endomorphism in FinStoch, with

$$p_i^+ = \frac{i f_i}{i f_i + (N-i) g_i} \cdot \frac{N-i}{N}, \qquad p_i^- = \frac{(N-i) g_i}{i f_i + (N-i) g_i} \cdot \frac{i}{N},$$

and 0 and N absorbing.

Theorem 6.1 (Fixation). [classical; [Mor58, Now06]] *The probability that a single A-player takes over the population is*

$$\rho_A = \Big(1 + \sum_{k=1}^{N-1} \prod_{i=1}^{k} \frac{g_i}{f_i}\Big)^{-1}.$$

Proof. The probability h_i of absorption at N from i is the harmonic function with $h_0 = 0$, $h_N = 1$. Harmonicity gives $h_{i+1} - h_i = (p_i^- / p_i^+)(h_i - h_{i-1}) = (g_i/f_i)(h_i - h_{i-1})$, and the differences telescope. □

For constant fitness r of the mutant relative to the resident this is $\rho = (1 - r^{-1})/(1 - r^{-N})$, and $\rho = 1/N$ in the neutral case. [F3]

The source tabulates the fixation probability ρ_B of a single defector in the Prisoner's Dilemma $a = 3$, $b = 1$, $c = 4$, $d = 2$ with $w = 1$. The formula above, with self-interaction excluded as in the definition of F_i and G_i, reproduces the table; if a player may meet itself the values differ in the second decimal place at $N = 10$ and have the same limit. [F1]

N	10	10^2	10^3	10^4	10^5
source	0.34263	0.25879	0.25088	0.25009	0.25001
computed, no self-play	0.34263	0.25879	0.25088	0.25009	not run
computed, self-play	0.30163	0.25500	0.25050	0.25005	not run

Proposition 6.2 (The limit). [classical] *If B strictly dominates A ($c > a$ and $d > b$) and $w = 1$, then $\rho_B \to 1 - a/c$ as $N \to \infty$.*

Proof. For ρ_B the ratios in Theorem 6.1 are those of resident to mutant fitness. With j mutants and $N \to \infty$ the j-th ratio tends to a/c, and by dominance all ratios are bounded by some $\theta < 1$ for large N, uniformly in the composition. Dominated convergence gives $\rho_B^{-1} \to \sum_{k \geq 0} (a/c)^k = (1 - a/c)^{-1}$. □

With $a = 3$ and $c = 4$ the limit is $\frac{1}{4}$. [F2] The number $1 - a/c$ is the survival probability of a branching process in which each mutant has relative fitness c/a: while mutants are rare they meet only residents, and their lineage either dies out early or escapes.

6.2 Two limits that do not commute

The source concludes from $\rho_B \to \frac{1}{4}$ that "the replicator dynamics cannot be the limit of the birth-death process", since under the replicator dynamics any positive fraction of defectors takes over with certainty [source, §4.1]. The observation is correct. The conclusion holds for one order of limits and fails for the other.

Theorem 6.3 (Mean-field limit). [classical; [Kur70, TCH05]] *Fix an initial* fraction $x_0 \in (0, 1)$ *and a time horizon T, measured in generations of N steps. As $N \to \infty$ the frequency of A in the Moran process converges in probability, uniformly on $[0, T]$, to the solution of the* adjusted replicator equation

$$\dot{x} = \frac{x(1 - x)\big(F(x) - G(x)\big)}{xF(x) + (1 - x)G(x)},$$

with fluctuations of order $N^{-1/2}$.

The adjusted equation has the same orbits as (3.2) and a different clock. For a Hawk–Dove game with background fitness the root-mean-square deviation from the solution at $T = 4$ was 0.158, 0.084 and 0.040 at $N = 50$, 200 and 800: it halves when N is quadrupled. [F5]

Remark 6.4 (Which limit). [observation] Write $\Phi_N(x_0, t)$ for the law of the frequency at generation t from initial fraction x_0. The replicator dynamics is the limit $N \to \infty$ with x_0 and t fixed. Fixation of a single mutant is the limit $t \to \infty$ first, at $x_0 = 1/N$, and then $N \to \infty$; here the initial fraction tends to zero, the limiting differential equation starts at the rest point $x = 0$ and never leaves, and the outcome is decided in a boundary layer of width $O(1/N)$ where the population is a branching process and not a flow. The two descriptions compose: a defector lineage escapes the layer with probability $1 - a/c$ and, having escaped, is carried to fixation by the flow with probability tending to one. The square

$$\begin{array}{ccc} \Phi_N(x_0, t) & \xrightarrow{N\to\infty} & \varphi_t(x_0) \\ {\scriptstyle t\to\infty}\downarrow & & \downarrow{\scriptstyle t\to\infty} \\ \text{fixation law} & \xrightarrow[N\to\infty]{} & ? \end{array}$$

commutes for each fixed $x_0 > 0$ in the Prisoner's Dilemma and fails along $x_0 = 1/N$. The replicator dynamics is a limit of the birth–death process, and it is not the limit that computes invasion from one mutant.

6.3 Weak selection

Theorem 6.5 (The one-third law). [literature; [NSTF04]] *For $w \to 0$, selection favours the fixation of A, in the sense $\rho_A > 1/N$, if and only if*

$$a(N-2) + b(2N-1) > c(N+1) + d(2N-4).$$

For a coordination game ($a > c$, $d > b$) and large N this holds if and only if the unstable equilibrium $x^* = (d-b)/(a-c+d-b)$ of the replicator equation, the frequency of A at the barrier between the two conventions, is below $\frac{1}{3}$.

The condition was confirmed against exact fixation probabilities in 299 random coordination games at $N = 60$. [F4] The constant has a simple origin: to first order in w the advantage $h(x) = F - G$ is weighted by the number of visits that a neutral lineage started from one mutant pays to frequency x, which is proportional to $1 - x$, and $\int_0^1 (\alpha x - \beta)\, 2(1-x)\, dx = \alpha/3 - \beta$. A convention can thus be invaded by a single mutant, with probability above the neutral one, exactly when its basin under the deterministic dynamics is smaller than a third; the replicator equation alone says only that it cannot be invaded at all.

6.4 Rare mutations

Add mutation at a small rate μ. The chain becomes ergodic, and as $\mu \to 0$ its stationary distribution concentrates on the monomorphic states, where it is the stationary distribution of the *embedded chain* that jumps from all-i to all-j with probability proportional to the fixation probability $\rho_{i\to j}$ of a single j-mutant [FuImo6]. For large populations and strong selection this chain moves only along the edges of the invasion graph of Chapter 4 and its stationary distribution is carried by the sink components: all of Rock–Paper–Scissors, either convention of a coordination game, defection in the Prisoner's Dilemma. [17] This is the finite-population counterpart of the chain components of Theorem 4.5, and it is the construction used to rank learning agents by Markov–Conley chains [OPP19]. Stochastic stability in the sense of [KMR93, You93] is the same idea for best-reply dynamics, where the selected convention is the risk-dominant one.

Remark 6.6 (What the language adds here). Not much. A finite-population model is a morphism in FinStoch; fixation is a harmonic function; the embedded chain is a functor from games to Markov chains that is natural for symmetries and, like the direct protocols of Proposition 5.4, not for cloning (an example is computed in the companion paper [EGML]). The content of the chapter is analytic, and the one structural point is Remark 6.4.

Notes

The frequency-dependent Moran process and its fixation probabilities are in [NSTF04, Now06]; the original process is Moran's [Mor58]. Density-dependent Markov chains and their differential-equation limits are due to Kurtz [Kur70]; the identification of the limit of the Moran process with the adjusted replicator equation is in [TCH05]. The one-third law is from [NSTF04]. The small-mutation limit is treated in [FuImo6].

Exercises

Exercise 6.1. Show without the formula that neutral fixation has probability $1/N$.

Exercise 6.2. Derive $\rho = (1 - r^{-1})/(1 - r^{-N})$ from Theorem 6.1.

Exercise 6.3. Deduce the large-population form of the one-third law from the inequality in Theorem 6.5.

Exercise 6.4. For the Prisoner's Dilemma of the source's table, show that the fixation probability of a single cooperator is smaller than $(3/4)^{N-1}$.

Chapter 7

Local Interaction

In the last part of the source the population lives on a graph: each player meets only its neighbours and learns only from them [source, §4.2]. The models are simulated, and the source is candid that their behaviour is sensitive to details. This is the chapter in which structure helps most, because a graph has maps, and the maps explain patterns that the simulations display.

7.1 The model

Let Γ be a finite graph with vertex set V, and (S, A) a game. A ***configuration*** is a map $\sigma\colon V \to S$. The score of v is $u_v(\sigma) = \sum_{w \sim v} A_{\sigma(v)\sigma(w)}$. Under ***imitate-the-best*** every player simultaneously adopts the strategy of the highest scorer among itself and its neighbours. Ties need a rule, and we take one that uses no labels: a player keeps its strategy if that strategy is among the best, or if the best scorers disagree. This defines a map

$$U_\Gamma\colon S^V \longrightarrow S^V,$$

a discrete dynamical system on a finite set. Every orbit is eventually periodic, and the questions are which periods occur and why.

7.2 Coverings

A ***covering*** $p\colon \tilde{\Gamma} \to \Gamma$ is a graph homomorphism that maps the neighbours of each vertex bijectively onto the neighbours of its image. The ring C_{kn} covers C_n for $n \geq 3$; the infinite line covers every ring; a torus covers a smaller torus; the square lattice covers them all.

Theorem 7.1 (Naturality under coverings). [observation] *For a covering $p\colon \tilde{\Gamma} \to \Gamma$ the pullback $p^*\colon S^V \to S^{\tilde{V}}$, $\sigma \mapsto \sigma \circ p$, commutes with the update:*

$$U_{\tilde{\Gamma}} \circ p^* = p^* \circ U_\Gamma.$$

Hence $\Gamma \mapsto (S^V, U_\Gamma)$ is a contravariant functor from graphs and coverings to dynamical systems. The same holds for every update rule in which a player's new strategy is a function of the labelled tree of walks of some fixed length starting at that player; imitate-the-best uses walks of length two.

Proof. Since p is bijective on neighbourhoods, $u_{\tilde{v}}(\sigma \circ p) = u_{p(\tilde{v})}(\sigma)$ for every $\tilde{v}$. Applying this at $\tilde{v}$ and at each of its neighbours, the pairs (strategy, score) seen by $\tilde{v}$ are those seen by $p(\tilde{v})$, with multiplicity. The update depends on nothing else. □

Corollary 7.2. *(a) Every orbit on Γ lifts to an orbit on every cover. The configurations on the ring C_{kn} that repeat with spatial period n are exactly the pullbacks of configurations on C_n, and they evolve as they do there.* [L2]

(b) The automorphisms of Γ and the symmetries of the game act on S^V and commute with U_Γ; the fixed set of every group of such symmetries is invariant. [L3]

(c) What happens on a small ring or torus happens on every larger one that covers it, and on the infinite lattice. A phenomenon found in simulation on a 30×30 torus that survives restriction to a 3×3 or 6×6 torus can be analysed there by hand.

Remark 7.3. [literature; [DeLe15, GoSt06]] For continuous-time systems on networks this is the theorem of DeVille and Lerman that fibrations of graphs induce maps of dynamical systems, and the invariant subspaces of (b) are the synchrony subspaces of Golubitsky and Stewart. On a Cayley graph the statement that U commutes with translations is half of the Curtis–Hedlund–Lyndon theorem: the continuous translation-equivariant self-maps of S^G are exactly the cellular automata [Hed69]. Local-interaction models on lattices are cellular automata whose local rule happens to come from a game.

7.3 Why the ring rotates

The source places Rock–Paper–Scissors on a ring of nine players in three clusters, *RRR PPP SSS*, and observes that each cluster moves one place per generation, for ever [source, §4.2.1]. [L1]

Proposition 7.4. [observation] *Let τ be rotation of C_9 by three places and c the cyclic relabelling $R \mapsto P \mapsto S \mapsto R$, which is a symmetry of the game. The configuration $\sigma_0 = RRR\,PPP\,SSS$ is fixed by the* twisted *symmetry $g\colon \sigma \mapsto c \circ \sigma \circ \tau^{-1}$. The set of configurations fixed by g has 27 elements and is invariant under U. Nine of them are the rotations of σ_0, on which U acts as rotation by one place, with period nine; the other eighteen are mapped into this orbit by one application of U.* [L8]

Proof. Invariance is Corollary 7.2(b). A fixed configuration is determined by its values on three consecutive vertices, whence 27. For the motion, score a win 1, a loss -1 and a tie 0. A player inside a cluster scores 0; at the boundary between the clusters of R and P the R-player scores -1 and the P-player $+1$.

The losing boundary player sees a neighbour with the top score and adopts its strategy; the winning one keeps its own; the middle player of a cluster sees only its own strategy. Each cluster therefore advances one place into the cluster it beats and retreats one place from the cluster that beats it. The statement about the other eighteen configurations is by enumeration. □

The pattern is a discrete rotating wave: a relative equilibrium, fixed by the dynamics up to a symmetry. The twisted symmetry accounts for the shape of the pattern, three clusters in cyclic order, and the local computation for its speed. By Corollary 7.2(a) the same wave runs on every ring of $9k$ players, and patterns with clusters of other sizes are waves for other twisted symmetries. Within the symmetric set the wave is the only possible long-run behaviour, which is more than symmetry alone could give. The source remarks that on a ring the only outcomes are cycling and uniformity. With the tie rule adopted here there is a third: the configuration *RPS RPS RPS*, in which every player beats one neighbour and loses to the other, is a fixed point, since all scores are equal. We found no tie rule stated in the source, and the remark may be correct for the rule used there; the example shows that it depends on one.

7.4 Cooperation at the boundary

Now let $S = \{C, D\}$ with the Prisoner's Dilemma payoffs $T > R > P > S$ (temptation, reward, punishment, sucker). On the simple ring a cluster of cooperators cannot grow. On the $1,2$-***circulant***, where each player has the two nearest players on each side as neighbours, it can [source, §4.2.2].

Proposition 7.5 (Thresholds on the $1,2$-circulant)**.** [observation] *Consider a run of at least five cooperators next to a run of at least five defectors. Under imitate-the-best the boundary between them*

moves into the defectors, one place per round,	*if* $3R + S > 2T + 2P$*;*
is stationary	*if* $3R + S < 2T + 2P < 4R$*;*
moves into the cooperators, two places per round,	*if* $2T + 2P > 4R$*.*

Proof. Number the cooperators $1, 2, 3, \dots$ from the boundary and the defectors $0, -1, -2, \dots$. The scores are

$$u_1 = 2R + 2S, \qquad u_2 = 3R + S, \qquad u_3 = 4R,$$
$$u_0 = 2T + 2P, \qquad u_{-1} = T + 3P, \qquad u_{-2} = 4P.$$

Defector 0 sees players $-2, \dots, 2$: the best defector it sees is itself and the best cooperator is player 2, so it switches exactly when $3R + S > 2T + 2P$. Defector -1 sees player 0, who outscores the only cooperator in view, and stays; the

remaining defectors see no cooperator. Cooperators 1 and 2 both see defector 0 and cooperator 3, and stay exactly when $4R > 2T + 2P$; cooperator 3 sees only cooperators and player 0 is out of its sight. □

With $R = 1$, $P = S = 0$ and $T = b$ the thresholds are $b = \frac{3}{2}$ and $b = 2$. Starting from six cooperators among thirty players, the cluster grew by two per round for $b = 1.1$, stayed at six for $b = 1.7$, and vanished in two rounds for $b = 2.1$. [L6] For the source's matrix $(T, R, P, S) = (1.1, 1, 0, -0.1)$ the first inequality reads $2.9 > 2.2$, which are the two scores the source quotes for the cooperator and the boundary defector. The simulation then reproduces the rest of the source's account: the cooperators spread until one defector remains; the lone defector scores $4T = 4.4$ and converts its four neighbours; the block of five defectors shrinks to three and to one; and the number of defectors cycles $1, 5, 3$ for ever. On the simple ring the same cluster is frozen for this matrix and eroded for the source's second matrix $(1.5, 1, 0.6, 0)$, as the source states. [L7]

Remark 7.6. The mechanism is visible in the proof: cooperation spreads when the cooperator one step behind the front, who has three cooperating neighbours, outscores the defector at the front, who exploits two. On the simple ring there is no such player. The source's caution that results of this kind depend on the neighbourhood, the payoffs and the update rule is exactly right, and the proposition says how: through three linear inequalities that change with each of them.

7.5 Graphs that amplify and suppress selection

For constant fitness the Moran process can be run on a weighted directed graph: an individual at i, chosen in proportion to fitness, places its offspring at j with probability w_{ij}, where W is a stochastic matrix.

Theorem 7.7 (Isothermal theorem). [literature; [LHN05]] *The fixation probability of a single mutant of fitness r, placed uniformly at random, equals the Moran value $(1 - r^{-1})/(1 - r^{-N})$ for every r if and only if W is doubly stochastic.*

A doubly stochastic matrix is a stochastic map that preserves the uniform distribution, so the isothermal graphs are the endomorphisms of the probability space $(V, \text{uniform})$ in FinStoch. For them the number of mutants is itself a birth–death chain with ratio r (Exercise 7.4), and structure is invisible to selection: on the complete graph and on the cycle the computed fixation probabilities agree with the Moran value to ten digits. [L4] Other graphs change the balance between selection and drift.

$N = 6$	well-mixed	star	directed line
advantageous mutant, $r = 1.5$	0.3654	0.4432	0.1667
disadvantageous mutant, $r = 0.7$	0.0571	0.0333	0.1667

The star amplifies selection in both directions, and on the directed line only a mutant at the root can take over, so fixation is $1/N$ whatever the fitness. [L5] The values are exact, from the linear system for the absorption probabilities on all 2^6 subsets.

Notes

Local-interaction models of this kind go back to Nowak and May [NoMa92]; the source's examples are in the tradition of Skyrms's and Alexander's work on the evolution of social structure [Sky04, Alex07]. Evolutionary graph theory, the isothermal theorem, amplifiers and suppressors are from [LHN05]. The source ends with a simulation of the centipede game on a lattice, in which the socially optimal outcome occurs 13.3 per cent of the time, close to the figure of about 15 per cent in the experiments of McKelvey and Palfrey [source, §4.2.3]; we have not reproduced it, since it depends on the distribution of initial conditions, which the functorial statements above do not see.

Exercises

Exercise 7.1. Show that reduction modulo n is a covering $C_{kn} \to C_n$ for $n \geq 3$ and not for $n = 2$.

Exercise 7.2. Verify that the configurations on C_9 fixed by the twisted symmetry of Proposition 7.4 number 27, that each uses every strategy exactly three times, and that $RPS\,RPS\,RPS$ is a fixed point of U that is not among them.

Exercise 7.3. Repeat the analysis of Proposition 7.5 for the simple ring, and recover the source's two examples.

Exercise 7.4. Show that for doubly stochastic W the ratio of the probabilities that the number of mutants increases and decreases is r in every state. Why is the star not isothermal?

Part IV

Assessment

Chapter 8

Assessment

8.1 What the language did

Five things, in decreasing order of how much they changed our understanding of the source.

It isolated the invariance that separates the dynamics. The replicator equation is natural for all payoff-preserving maps (Theorem 3.2); the Brown–von Neumann–Nash and Smith dynamics are natural only for symmetries (Proposition 5.4). The survival of dominated strategies, which the source presents as a surprising theorem about a class of dynamics, becomes the visible cost of giving up naturality for inclusions and clones: a dynamic that sees names instead of individuals keeps reintroducing a bad name.

It showed why stability must be a property of sets. The source moves from evolutionarily stable strategies to evolutionarily stable sets as a refinement motivated by examples. Theorem 2.8 gives a reason that is independent of examples: stability of a point is destroyed by duplicating a strategy, stability of a set is preserved, and no notion that depends on how strategies are named can be the right one.

It explained simulated patterns. The rotating clusters on the ring are forced by a twisted symmetry and a covering (Proposition 7.4, Theorem 7.1); the growth of cooperation on the circulant is three linear inequalities (Proposition 7.5). Both were found by asking what the maps of the model are.

It separated the game from the learner. Treating a population as an open system (Proposition 5.7, Remark 5.8) makes "same game, different dynamics" and "same dynamics, different structure" into statements about a diagram and its fillings, and it is the form in which the subject is used in machine learning, as the companion paper [EGML] develops.

It located a non-commuting square. The source's claim that the replicator dynamics is not the limit of the birth–death process is a statement about two orders of limits (Remark 6.4).

8.2 What it did not do

It proved no new theorem of evolutionary game theory. The observations tagged as such are short, and some of them are surely known to people

who work with these models, whether or not they are written down in this form. The fixation formulas, the mean-field limit and the one-third law gain nothing from being called morphisms in FinStoch. Chaos in the replicator equation, which the source illustrates with a four-strategy example [source, §3.2.5], is untouched: we have no categorical handle on it beyond the remark that a strange attractor is a chain component with a non-trivial Conley index. Bomze's classification of the phase portraits for three strategies is likewise only cited. The centipede simulation is not reproduced. And the source's philosophical conclusion, that these models give how-possibly explanations whose robustness must be argued case by case [source, §5], is one that the present treatment supports in a specific way and does not replace: naturality tells you which features of a model cannot matter, and so narrows the list of those that might.

8.3 How much of this is higher

Source	Structure used here	Level
§2 stable strategies and sets	Game; pullback along clone maps	1-categorical
§3.2 replicator dynamics	the functor $\mathsf{Rep}\colon \mathsf{Game} \to \mathsf{Dyn}$	1-categorical
§3.2.4 stability, cycles	index theorem; Conley index	homotopy types
§3.2.4 Rock–Paper–Scissors	Hodge decomposition	cohomology in degree 1
§3.3 other dynamics	naturality and its failure	1-categorical
§3.3–3.4 dynamics, populations	open systems and wirings	operadic; a double category
§4.1 birth–death process	FinStoch; a non-commuting square	1-categorical
§4.2 local interaction	coverings; equivariance	1-categorical, with groupoids of symmetries

In summary, the subject at the level of the source is mostly ordinary category theory. Three things point beyond it. Games modulo column shifts and positive rescalings form a groupoid over each strategy set, and Rep is invariant along it up to reparametrization of time; treating these equivalences as 2-cells is the natural setting for Proposition 3.4. The Conley index is a homotopy type, functorial for the inclusions of a Morse decomposition, and its continuation invariance is a statement about families of games. And the first cohomology that vanishes for the complete comparison graph does not vanish for partial comparisons, which is where the decomposition of Theorem 4.10 acquires a

harmonic part.

8.4 Programme

[programme] Four questions seem worth the effort. (1) Is there a homotopy-invariant of a game, computable from the payoff matrix, that determines the Conley indices of the chain components of its replicator flow, as the index of Theorem 4.2 determines their Euler characteristics? (2) Which dynamics, among those arising from revision protocols, are functors on Game? Proposition 5.4 shows that all imitative ones are; a characterization would say what naturality costs in terms of the properties of Theorem 5.3. (3) Can the covering functor of Theorem 7.1 be extended from coverings to fibrations and to stochastic update rules, so that it applies to the Moran process on graphs and explains amplifiers as it explains waves? (4) Is there a setting in which the two limits of Remark 6.4 are the two legs of one universal construction?

Notes

For the source's own assessment of what evolutionary models explain, see its concluding section. A treatment of classical game theory in the style of this volume is [GIHC]; relative entropy as a functor, used in Chapter 3, is developed in [ESCHC]; the index theory used in Chapter 4 is in [DTHC].

Appendix A

Toolkit

This appendix collects the general notions used in the text, in the order in which they are needed.

Categories and functors. A category has objects, morphisms and an associative composition with identities; a functor preserves them. A statement of the form "the construction $X \mapsto F(X)$ is natural" means that F is a functor: every map $X \to Y$ of the relevant kind induces a map $F(X) \to F(Y)$, compatibly with composition. A group acting on X then acts on $F(X)$, a subobject of X gives an invariant subobject of $F(X)$, and a quotient of X gives a quotient of $F(X)$.

Convex spaces and the distribution monad. For a finite set S, $\Delta(S)$ is the simplex of probability distributions on S. A map $f\colon S \to T$ induces the pushforward $f_*\colon \Delta(S) \to \Delta(T)$, and Δ is a monad on finite sets whose algebras are convex spaces [Fri09]. A map $\Delta(S) \to \mathbb{R}$ is affine exactly when it is the expectation of a function on S; a payoff is a function on $S \times S$, and its biaffine extension is the mixed extension.

Stochastic maps. FinStoch is the category of finite sets and stochastic matrices. A Markov chain on a finite set is an endomorphism in it. A state is absorbing if its row is a point mass; a function is harmonic if it is fixed by the matrix acting on functions.

Dynamical systems. A flow on a compact space X is a continuous action of $\mathbb{R}$; a map of flows is a continuous map commuting with the actions. A rest point is *Lyapunov stable* if orbits starting near it stay near it and *asymptotically stable* if in addition they converge to it. A *Lyapunov function* is a function that decreases along orbits; a strict one certifies asymptotic stability.

Index and Euler characteristic. For a vector field with non-degenerate zeros on a compact manifold with boundary on which it points outward, the sum over zeros of the sign of the determinant of the derivative is the Euler charac-

teristic (Poincaré–Hopf). A simplex is contractible, so its Euler characteristic is 1.

Cochains on a graph. A 0-cochain is a function on vertices and a 1-cochain an antisymmetric function on ordered pairs of adjacent vertices. The differential sends r to $(i, j) \mapsto r_i - r_j$. A 1-cochain is closed if its sum around every triangle vanishes. On the complete graph every closed 1-cochain is a differential.

Coverings and symmetries. A covering of graphs is a homomorphism bijective on the neighbourhood of each vertex. An automorphism of a graph is a covering of the graph by itself. A map between dynamical systems on configuration spaces that commutes with the updates sends orbits to orbits, periodic orbits to periodic orbits of period dividing the original, and invariant sets to invariant sets.

Appendix B

Concordance with the Source

The table lists the sections of J. McKenzie Alexander, *Evolutionary Game Theory* [Alex23], with the place in this volume where their content is treated and the form it takes there. The source was read from a text extraction of its PDF; its definitions and theorems were located by search and read in context, and the remaining text was read selectively, section by section. Figures were not consulted except through their captions and the surrounding discussion.

Source	Here	Form
1.1 Rock, Paper, Scissors (Lizard, Spock)	Example 1.5	circulant games; their automorphism groups
1.2–1.3 Game theory; fundamentals	Chapter 1	the distribution monad; the category Game; Nash's theorem for symmetric games
2.1 Basic concepts (Definitions 1–4)	Theorem 2.3	one quadratic form Q_x; the ray identity (2.1)
2.1 Drift game; Hawk–Dove	Example 2.4, Chapter 3	$Q = C(p-q)^2/2$; escape time from a neutrally stable state
2.2 Properties (Theorems 1–5)	Theorem 2.6	stated with proofs or references; the Bishop–Cannings bound
2.3 Evolutionarily stable sets	Theorem 2.8	pullback along clone maps preserves stable sets and destroys stable points
3.1 Introduction	Chapter 3	the category Dyn
3.2 Replicator dynamics (Definition 11)	Proposition 3.1, Theorem 3.2	quotient of the growth equation; functoriality
3.2.1 Biological derivation	Proposition 3.1	as stated
3.2.2 Cultural derivation	Proposition 5.2	mean dynamic of an imitative revision protocol
3.2.3 Simple games	Chapter 3, two strategies	one normal form with three generic cases
3.2.4 Stability (Theorems 6, 7)	Theorem 3.6, Theorem 3.8, Theorem 3.9	relative entropy as Lyapunov function; elimination with its rate
3.2.4 Rock–Paper–Scissors variants	Theorem 3.10, Proposition 3.11, Theorem 4.10	conserved quantity; eigenvalues; cyclic part of a game

Source	Here	Form
3.2.4 Classification for three strategies	Chapter 3	cited only
3.2.5 Chaos	Chapter 8	not treated
3.3 Other dynamics (Theorem 8)	Theorem 5.3, Proposition 5.4	naturality separates imitation from innovation
3.4 Multipopulation games	Theorem 5.9	two boxes and a wiring
4.1 Birth–death process	Theorem 6.1, Proposition 6.2, Remark 6.4	the source's table recomputed; two limits
4.2 Local interaction	Theorem 7.1	functoriality for coverings
4.2.1 Rock–Paper–Scissors on a ring	Proposition 7.4	twisted symmetry; rotating wave
4.2.2 Prisoner's Dilemma on graphs	Proposition 7.5	thresholds; the source's run reproduced
4.2.3 Centipede game	Chapter 7, Notes	not reproduced
5 Conclusion	Chapter 8	what naturality contributes to robustness
—	Theorem 4.2, Theorem 4.5, Theorem 4.8	added: index theorem, chain recurrence, Conley index
—	Theorem 6.3, Theorem 6.5	added: mean-field limit, one-third law
—	Theorem 7.7	added: isothermal theorem, amplifiers and suppressors

Appendix C

Computations

The script `verify-evolutionary-games.py` that accompanies this volume uses only the Python standard library, fixes its random seed, and runs in about ten seconds. It performs 42 checks, each printing `PASS` or `FAIL` with the computed values, and exits with status zero if and only if all pass. The checks recompute the numerical statements of the text and test finite or numerical instances of the structural ones. They are evidence that the examples are stated correctly; they do not prove the theorems. Differential equations are integrated with a fourth-order Runge–Kutta scheme; fixation probabilities are computed exactly, from the product formula or by solving the linear system for absorption.

Check	What is computed
S1	Hawk–Dove with $V = 2$, $C = 4$: $Q_{x^*}(y) = C(p-q)^2/2$ on a grid of states (Example 2.4)
S2	Rock–Paper–Scissors: the uniform state is the only Nash equilibrium and is not an ESS
S3	the source's game with a weakly dominated equilibrium strategy: not an ESS (Theorem 2.6)
S4	drift game: the neutral equilibrium is left by the replicator flow; time from 0.05 to $\frac{1}{2}$ is 20.944
S5	Hawk–Dove with $V < C$ has exactly one ESS, the mixed one
S6	Theorem 2.8: over the mixed ESS of Hawk–Dove with Dove duplicated, the fibre is a stable set and no point of it is an ESS
R1	Theorem 3.2 on random games: pushforward along a clone surjection and a subgame inclusion carries solutions to solutions
R2	Theorem 3.5: $\frac{d}{dt}$(mean payoff) $= 2\,\mathrm{Var}$ for symmetric matrices
R3	Theorem 3.6: $\frac{d}{dt} D(x^* \,\|\, x) = -Q_{x^*}(x)$ along numerical solutions
R4	Theorem 3.10: the relative entropy to the centre is conserved, drift 1.3×10^{-15} over time 40
R5	Proposition 3.11: eigenvalues $(b-a)/6 \pm i\sqrt{3}(a+b)/6$ against a numerical Jacobian
R6	good Rock–Paper–Scissors converges to the centre; bad Rock–Paper–Scissors leaves every compact subset of the interior
R7	every Nash equilibrium of random games is a rest point
R8	the source's three-strategy game: the unique equilibrium is uniform, asymptotically stable, and not an ESS

Check	What is computed
R9	Theorem 3.9: the log-ratio of a dominated twin to its original falls at exactly the penalty
I1	Theorem 4.2 on 400 random games with three or four strategies: indices sum to 1, the number of equilibria is odd
I2	all 590 ESS found (by Haigh's criterion) have index $+1$
I3	coordination game: indices $+1, +1, -1$
I4	Theorem 4.10: Rock–Paper–Scissors is purely cyclic, a rating difference purely transitive, the parts orthogonal; $\tau = 0.558$ for a random game
I5	the gradient part is a cocycle on every triangle; the cyclic part carries all the curl
I6	$\frac{d}{dt} x^{\mathsf{T}} A x = 2\,\mathrm{Cov}_x(Sx, Ax)$ on random games
I7	sink components of the invasion graph for three standard games; the small-mutation chain concentrates on them
O1	Proposition 5.2 on random games and states
O2	at All-Cooperate the replicator field vanishes and the two direct dynamics do not
O3	both direct dynamics vanish at every Nash equilibrium of sixty random games
O4	Theorem 5.3: time-averaged share 0.116 of a dominated twin under the Smith dynamic; extinction under the replicator
O5	Battle of the Sexes with two populations: the mixed equilibrium is a rest point, and forty nearby orbits end at the two pure equilibria
O6	Proposition 5.7 against a numerical solution
O7	Proposition 5.4(c): defect of naturality under cloning, 4×10^{-17} for the replicator, 0.555 for Smith, 0.380 for Brown–von Neumann–Nash
F1	the source's table of fixation probabilities, with and without self-interaction
F2	Proposition 6.2: the limit $1 - a/c = \frac{1}{4}$
F3	neutral and constant-fitness fixation
F4	Theorem 6.5 against exact fixation in 299 random coordination games, $N = 60$
F5	Theorem 6.3: root-mean-square deviation 0.158, 0.084, 0.040 at $N = 50, 200, 800$
L1	ring of nine: rotation by one place per round, period nine
L2	Theorem 7.1 for $C_9 \to C_3$, all 27 configurations
L3	equivariance under rotations, reflections and relabelling on random configurations; the twisted symmetry fixes the wave
L4	Theorem 7.7: complete graph and cycle
L5	star and directed line at $N = 6$ (exact)
L6	Proposition 7.5 at $b = 1.1, 1.7, 2.1$
L7	the source's matrix on the $1, 2$-circulant and on the ring; the cycle $1, 5, 3$
L8	Proposition 7.4: all 27 symmetric configurations; a frozen configuration outside them

Two discrepancies with the source were looked for and one convention was found. The source's fixation table is reproduced to all printed digits only if a player does not interact with itself (F1). No tie rule for imitate-the-best was found in the source; with the rule of Chapter 7 its described runs are reproduced (L1, L7), and a frozen configuration exists that its text does not mention (L8).

Appendix D

Solutions to the Exercises

Solutions to all exercises follow, by chapter. Where a solution relies on a computation, the computation is also carried out by `verify-evolutionary-games.py`.

Chapter 1

1.1. By affineness in the first argument, $\pi(y_\epsilon \mid y_\epsilon) = (1-\epsilon)\pi(x \mid y_\epsilon) + \epsilon\,\pi(\mu \mid y_\epsilon)$. Subtract from $\pi(x \mid y_\epsilon)$.

1.2. A symmetry is a permutation that preserves who beats whom. For Rock–Paper–Scissors the three rotations do, and a transposition reverses the cycle, so the group is $\mathbb{Z}/3$. For the five-strategy game the translations give $\mathbb{Z}/5$. A symmetry is an automorphism of a tournament, and such groups have odd order, since an involution would exchange two strategies and so reverse the result between them. A subgroup of the symmetric group on five letters of odd order containing a 5-cycle has order 5 or 15, and there is no subgroup of order 15, since it would be cyclic. The group is $\mathbb{Z}/5$.

1.3. Subtract $(V-C)/2$ from the first column and $V/2$ from the second: $\left(\begin{smallmatrix} 0 & V/2 \\ (C-V)/2 & 0 \end{smallmatrix}\right)$, so $a = V/2$ and $b = (C-V)/2$. For $V < C$ both are positive: neither pure state is an equilibrium and the mixed equilibrium has Hawk at frequency $a/(a+b) = V/C$. For $V > C$, $a > 0 > b$: Hawk strictly dominates and All-Hawk is the only equilibrium.

1.4. By Proposition 1.10(a) applied to the clone map onto Hawk–Dove, the equilibria are exactly the states whose image is the mixed equilibrium: $(V/C,\, s,\, 1 - V/C - s)$ with $0 \le s \le 1 - V/C$. They form a segment.

Chapter 2

2.1. If x is strict, then $\pi(x \mid x) > \pi(y \mid x)$ for $y \neq x$, which is the first clause of the definition. In a two-strategy game let x be a pure ESS, say the first strategy. It is an equilibrium, so $A_{11} \ge A_{21}$, and by hypothesis the inequality is strict; then x is strict. An ESS that is not pure is interior.

2.2. The uniform state is an interior equilibrium, so by Theorem 2.6(f) it is an ESS if and only if $z^{\mathsf{T}}Az < 0$ for $z \neq 0$ with $\sum_s z_s = 0$. Here $z^{\mathsf{T}}Az = (a-b)(z_1z_2 + z_2z_3 + z_3z_1)$ and $z_1z_2 + z_2z_3 + z_3z_1 = -\tfrac{1}{2}|z|^2$ on that plane, so $z^{\mathsf{T}}Az = -\tfrac{a-b}{2}|z|^2$, negative exactly when $a > b$.

2.3. The identity matrix: each pure strategy is a strict equilibrium, hence an ESS, and $\binom{3}{1} = 3$.

2.4. Write $y = (1-q, q)$, with q the weight on the second strategy. For x the first pure strategy, $\pi(x \mid y) = 1 + 99q$ and $\pi(y \mid y) = (1-q)(1+99q) + 100q^2$, so $Q_x(y) = q(1+99q) - 100q^2 = q(1-q) > 0$ for $0 < q < 1$: an ESS. For x the second, with $p = 1-q$ the weight on the first, $Q_x(y) = -p^2 < 0$: not an ESS, although it is an equilibrium.

Chapter 3

3.1. With payoffs $\left(\begin{smallmatrix} 0 & a \\ b & 0 \end{smallmatrix}\right)$ and state $(x, 1-x)$, the two strategies earn $a(1-x)$ and bx, and $\dot{x} = x(1-x)\big(a(1-x) - bx\big)$. For $\dot{x} = x^2(1-x)$, partial fractions give $\int \frac{dx}{x^2(1-x)} = -\frac{1}{x} + \log\frac{x}{1-x}$; evaluate between x_0 and x_1.

3.2. $\frac{d}{dt}x^{\mathsf{T}}Ax = \dot{x}^{\mathsf{T}}(A + A^{\mathsf{T}})x = 2\sum_s x_s\big((Ax)_s - x^{\mathsf{T}}Ax\big)(Sx)_s = 2\,\mathrm{Cov}_x(Sx, Ax)$, because $\sum_s x_s\big((Ax)_s - x^{\mathsf{T}}Ax\big) = 0$. For symmetric A this is twice the variance of Ax.

3.3. At the uniform state the linearization on the tangent plane is $\frac{1}{3}A$ composed with the projection along $\mathbf{1}$. In the basis of the two complex eigenvectors of the cyclic shift, A acts by $a\omega + (-b)\bar{\omega}$ and its conjugate, with $\omega = e^{2\pi i/3}$, so the eigenvalues are $\frac{1}{3}(a\omega - b\bar{\omega})$ and its conjugate: real part $\frac{1}{3}\cdot\frac{-a+b}{2}$, imaginary part $\pm\frac{1}{3}\cdot\frac{\sqrt{3}}{2}(a+b)$. The trace is $(b-a)/3$ and the determinant $\big((b-a)^2 + 3(a+b)^2\big)/36$.

3.4. By Corollary 3.3(b) the class frequencies follow the replicator equation of Rock–Paper–Scissors from the initial state $(\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$. By Theorem 3.10 they move on the closed curve on which the relative entropy to the centre is constant, equivalently on which the product of the three frequencies equals $\frac{1}{32}$. The two variants of Rock keep equal shares.

Chapter 4

4.1. Let $F(x) = x(1-x)\big(a - (a+b)x\big)$; the index at an equilibrium is the sign of $-F'$. For $a > 0 > b$ the only equilibrium is $x = 1$, with $F'(1) = b < 0$: index $+1$. For $a, b > 0$ the only equilibrium is $x^* = a/(a+b)$, with $F'(x^*) = -(a+b)x^*(1-x^*) < 0$: index $+1$. For $a, b < 0$ the equilibria are 0, 1 and x^*,

with $F'(0) = a < 0$, $F'(1) = b < 0$ and $F'(x^*) > 0$: indices $+1, +1, -1$. The sum is 1 in each case.

4.2. A gradient has zero curl on every triangle. Conversely, if the curl vanishes on triangles through the vertex 0, put $r_i = K_{i0}$. Then $K_{ij} + K_{j0} + K_{0i} = 0$ gives $K_{ij} = K_{i0} - K_{j0} = r_i - r_j$.

4.3. The row means are $r = \left(\frac{1+t}{3}, 0, -\frac{1+t}{3}\right)$, so $G_{12} = G_{23} = \frac{1+t}{3}$, $G_{13} = \frac{2(1+t)}{3}$, and $C_{12} = C_{23} = \frac{2-t}{3}$, $C_{13} = -\frac{2-t}{3}$. The curl is $2 - t$, and

$$\tau = \frac{2(1+t)^2}{3(2+t^2)}.$$

It equals 1 at $t = 2$, where the game is a rating difference, and 0 at $t = -1$, which is Rock–Paper–Scissors.

4.4. Dove invades Hawk, since $0 > (V - C)/2$, and Hawk invades Dove, since $V > V/2$. The graph is a two-cycle and its only sink component is the whole of it. This records correctly that neither pure state persists, and it misses the outcome: the flow converges to the mixed state, which is not a vertex.

Chapter 5

5.1. $\dot{x}_i = \sum_j x_j x_i (K - \pi_j) - x_i (K - \pi_i) \sum_j x_j = x_i (K - \bar{\pi}) - x_i (K - \pi_i) = x_i (\pi_i - \bar{\pi})$.

5.2. At a Nash equilibrium no strategy earns more than the mean, so $k = 0$ and the field vanishes. Conversely, let $\dot{x} = 0$ and suppose $\kappa = \sum_j k_j > 0$. Then $x_i = k_i/\kappa$ for all i, so the support of x consists of strategies earning strictly more than the mean, which contradicts $\sum_i x_i (\pi_i - \bar{\pi}) = 0$. Hence $k = 0$, which is the equilibrium condition.

5.3. Let $z(t) = \log x(0) + \int_0^t \pi$ and $x = \mathrm{softmax}(z)$. Then $\frac{d}{dt} \log x_i = \pi_i - \frac{d}{dt} \log \sum_j e^{z_j} = \pi_i - \sum_j x_j \pi_j$.

5.4. Row's two strategies earn $2y$ and $1 - y$; Column's earn x and $2(1 - x)$. Hence $\dot{x} = x(1 - x)(3y - 1)$ and $\dot{y} = y(1 - y)(3x - 2)$. At $(\frac{2}{3}, \frac{1}{3})$ the Jacobian has zero diagonal and both off-diagonal entries equal to $3 \cdot \frac{2}{9} = \frac{2}{3}$; its eigenvalues are $\pm\frac{2}{3}$.

Chapter 6

6.1. With equal fitnesses the number of A-players is a bounded martingale, so its expectation at absorption equals its initial value: $N\rho = 1$. Alternatively,

each of the N initial individuals is equally likely to be the ancestor of the final population.

6.2. Here $g_i/f_i = r^{-1}$ for every i, and $1 + \sum_{k=1}^{N-1} r^{-k} = (1 - r^{-N})/(1 - r^{-1})$.

6.3. Dividing by N and letting $N \to \infty$ gives $a + 2b > c + 2d$, that is, $a - c > 2(d - b)$. With $a > c$ and $d > b$ this is $x^* = (d - b)/\big((a - c) + (d - b)\big) < \frac{1}{3}$.

6.4. With i cooperators, $\frac{g_i}{f_i} = \frac{4i+2(N-i-1)}{3(i-1)+(N-i)} = \frac{2i+2N-2}{2i+N-3}$, which decreases in i and at $i = N - 1$ equals $\frac{4N-4}{3N-5} > \frac{4}{3}$. The last term of the sum in Theorem 6.1 alone exceeds $(4/3)^{N-1}$.

Chapter 7

7.1. The neighbours $i \pm 1$ of i go to $(i \pm 1)$ mod n, which are distinct and are the neighbours of i mod n when $n \geq 3$. For $n = 2$ they coincide, and C_2 is in any case a single edge.

7.2. The condition is $\sigma(i + 3) = c(\sigma(i))$ for all i. The values at $0, 1, 2$ are free and determine the rest consistently, because c^3 is the identity: 27 configurations. Each orbit $\{i, i + 3, i + 6\}$ of the rotation carries a, $c(a)$, $c^2(a)$, so each strategy occurs once per orbit and three times in all. In *RPS RPS RPS* every player beats one neighbour and loses to the other, all scores are 0, and under the tie rule nobody changes. It satisfies $\sigma(i + 3) = \sigma(i) \neq c(\sigma(i))$ and so is not among the 27.

7.3. Number the cooperators $1, 2, \ldots$ and the defectors $0, -1, \ldots$ from the boundary. The scores are $u_1 = R + S$, $u_2 = 2R$, $u_0 = T + P$, $u_{-1} = 2P$. Defector 0 outscores both its neighbours and stays. Cooperator 1 sees defector 0 and cooperator 2, and stays exactly when $2R > T + P$. So a cluster never grows; it is frozen if $2R > T + P$ and loses one player at each end per round if $2R < T + P$. For $(1.1, 1, 0, -0.1)$, $2 > 1.1$: frozen. For $(1.5, 1, 0.6, 0)$, $2 < 2.1$: eroded.

7.4. Let M be the set of mutants and $m = \sum_{i,j \in M} w_{ij}$. The number of mutants increases with probability proportional to $r \sum_{i \in M,\, j \notin M} w_{ij} = r(|M| - m)$, using row sums, and decreases with probability proportional to $\sum_{i \notin M,\, j \in M} w_{ij} = |M| - m$, using column sums, with the same constant. The ratio is r. On the star each leaf sends its offspring to the centre, so the column of the centre sums to $N - 1$ and that of a leaf to $1/(N - 1)$.

References

Entries were assembled from memory, except [OPP19], whose authors, title and venue were checked against the publisher's page on 17 November 2026. Volume and page numbers in particular should be verified before they are relied on.

[Aki79] E. Akin, *The Geometry of Population Genetics*, Lecture Notes in Biomathematics 31, Springer, 1979.

[Alex07] J. M. Alexander, *The Structural Evolution of Morality*, Cambridge University Press, 2007.

[Alex23] J. M. Alexander, *Evolutionary Game Theory*, Elements in Decision Theory and Philosophy, Cambridge University Press, 2023. The source.

[ASS80] A. Arneodo, P. Coullet, C. Tresser, Occurrence of strange attractors in three-dimensional Volterra equations, *Phys. Lett. A* 79 (1980), 259–263.

[BaFr14] J. C. Baez, T. Fritz, A Bayesian characterization of relative entropy, *Theory Appl. Categ.* 29 (2014), 422–456.

[BaPo16] J. C. Baez, B. S. Pollard, Relative entropy in biological systems, *Entropy* 18 (2016), no. 2, 46.

[BaSc01] D. Balkenborg, K. H. Schlag, Evolutionarily stable sets, *Internat. J. Game Theory* 29 (2001), 571–595.

[BiCa76] D. T. Bishop, C. Cannings, Models of animal conflict, *Adv. Appl. Probab.* 8 (1976), 616–621.

[BjWe96] J. Björnerstedt, J. W. Weibull, Nash equilibrium and evolution by imitation, in *The Rational Foundations of Economic Behaviour*, Macmillan, 1996, 155–171.

[Bom83] I. M. Bomze, Lotka–Volterra equation and replicator dynamics: a two-dimensional classification, *Biol. Cybernet.* 48 (1983), 201–211.

[BTP18] D. Balduzzi, K. Tuyls, J. Pérolat, T. Graepel, Re-evaluating evaluation, in *Advances in Neural Information Processing Systems* 31, 2018.

[BvN50] G. W. Brown, J. von Neumann, Solutions of games by differential equations, in *Contributions to the Theory of Games I*, Ann. of Math. Stud. 24, Princeton University Press, 1950, 73–79.

[CMOP11] O. Candogan, I. Menache, A. Ozdaglar, P. A. Parrilo, Flows and decompositions of games: harmonic and potential games, *Math. Oper. Res.* 36 (2011), 474–503.

[Con78] C. Conley, *Isolated Invariant Sets and the Morse Index*, CBMS Regional Conference Series in Mathematics 38, American Mathematical Society, 1978.

[DeLe15] L. DeVille, E. Lerman, Modular dynamical systems on networks, *J. Eur. Math. Soc.* 17 (2015), 2977–3013.

[DTHC] Anonymous, *Differential Topology in Higher Categories*, Higher Categorical Semantics, 2026. Companion volume.

[EGML] Anonymous, Evolutionary game theory in machine learning and AI: naturality, regret, and recurrence, Higher Categorical Semantics, 2026. Companion paper.

[ESCHC] Anonymous, *Entropy, Sources, and Channels in Higher Categories*, Higher Categorical Semantics, 2026. Companion volume.

[Fri09] T. Fritz, Convex spaces I: definition and examples, preprint, arXiv:0903.5522.

[FuIm06] D. Fudenberg, L. A. Imhof, Imitation processes with small mutations, *J. Econom. Theory* 131 (2006), 251–262.

[GHWZ18] N. Ghani, J. Hedges, V. Winschel, P. Zahn, Compositional game theory, in *LICS 2018*, ACM, 2018, 472–481.

[GIHC] Anonymous, *Games and Information in Higher Categories*, Higher Categorical Semantics, 2026. Companion volume.

[GoSt06] M. Golubitsky, I. Stewart, Nonlinear dynamics of networks: the groupoid formalism, *Bull. Amer. Math. Soc.* 43 (2006), 305–364.

[Hai75] J. Haigh, Game theory and evolution, *Adv. Appl. Probab.* 7 (1975), 8–11.

[Har09] M. Harper, Information geometry and evolutionary game theory, preprint, arXiv:0911.1383.

[Hed69] G. A. Hedlund, Endomorphisms and automorphisms of the shift dynamical system, *Math. Systems Theory* 3 (1969), 320–375.

[Hof81] J. Hofbauer, On the occurrence of limit cycles in the Volterra–Lotka equation, *Nonlinear Anal.* 5 (1981), 1003–1007.

[HoSa11] J. Hofbauer, W. H. Sandholm, Survival of dominated strategies under evolutionary dynamics, *Theor. Econ.* 6 (2011), 341–377.

[HoSi98] J. Hofbauer, K. Sigmund, *Evolutionary Games and Population Dynamics*, Cambridge University Press, 1998.

[HSS79] J. Hofbauer, P. Schuster, K. Sigmund, A note on evolutionary stable strategies and game dynamics, *J. Theoret. Biol.* 81 (1979), 609–612.

[JLYY11] X. Jiang, L.-H. Lim, Y. Yao, Y. Ye, Statistical ranking and combinatorial Hodge theory, *Math. Program.* 127 (2011), 203–244.

[KMR93] M. Kandori, G. J. Mailath, R. Rob, Learning, mutation, and long run equilibria in games, *Econometrica* 61 (1993), 29–56.

[Kur70] T. G. Kurtz, Solutions of ordinary differential equations as limits of pure jump Markov processes, *J. Appl. Probab.* 7 (1970), 49–58.

[LHN05] E. Lieberman, C. Hauert, M. A. Nowak, Evolutionary dynamics on graphs, *Nature* 433 (2005), 312–316.

[Mor58] P. A. P. Moran, Random processes in genetics, *Proc. Cambridge Philos. Soc.* 54 (1958), 60–71.

[MSP73] J. Maynard Smith, G. R. Price, The logic of animal conflict, *Nature* 246 (1973), 15–18.

[Mye21] D. J. Myers, Double categories of open dynamical systems (extended abstract), in *Applied Category Theory 2020*, EPTCS 333, 2021, 154–167.

[Nac90] J. H. Nachbar, "Evolutionary" selection dynamics in games: convergence and limit properties, *Internat. J. Game Theory* 19 (1990), 59–89.

[Nas51] J. Nash, Non-cooperative games, *Ann. of Math.* 54 (1951), 286–295.

[NoMa92] M. A. Nowak, R. M. May, Evolutionary games and spatial chaos, *Nature* 359 (1992), 826–829.

[Now06] M. A. Nowak, *Evolutionary Dynamics: Exploring the Equations of Life*, Harvard University Press, 2006.

[NSTF04] M. A. Nowak, A. Sasaki, C. Taylor, D. Fudenberg, Emergence of cooperation and evolutionary stability in finite populations, *Nature* 428 (2004), 646–650.

[OPP19] S. Omidshafiei, C. Papadimitriou, G. Piliouras, K. Tuyls, M. Rowland, J.-B. Lespiau, W. M. Czarnecki, M. Lanctot, J. Perolat, R. Munos, α-Rank: multi-agent evaluation by evolution, *Sci. Rep.* 9 (2019); arXiv:1903.01373.

[PaPi19] C. Papadimitriou, G. Piliouras, Game dynamics as the meaning of a game, *ACM SIGecom Exchanges* 16 (2019), no. 2, 53–63.

[RiWe95] K. Ritzberger, J. W. Weibull, Evolutionary selection in normal-form games, *Econometrica* 63 (1995), 1371–1399.

[Sal85] D. Salamon, Connected simple systems and the Conley index of isolated invariant sets, *Trans. Amer. Math. Soc.* 291 (1985), 1–41.

[San10] W. H. Sandholm, *Population Games and Evolutionary Dynamics*, MIT Press, 2010.

[SaZh92] L. Samuelson, J. Zhang, Evolutionary stability in asymmetric games, *J. Econom. Theory* 57 (1992), 363–391.

[Sch98] K. H. Schlag, Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits, *J. Econom. Theory* 78 (1998), 130–156.

[ScSi83] P. Schuster, K. Sigmund, Replicator dynamics, *J. Theoret. Biol.* 100 (1983), 533–538.

[Sel80] R. Selten, A note on evolutionarily stable strategies in asymmetric animal conflicts, *J. Theoret. Biol.* 84 (1980), 93–101.

[Sha79] S. Shahshahani, A new mathematical framework for the study of linkage and selection, *Mem. Amer. Math. Soc.* 17 (1979), no. 211.

[Sky04] B. Skyrms, *The Stag Hunt and the Evolution of Social Structure*, Cambridge University Press, 2004.

[Smi84] M. J. Smith, The stability of a dynamic model of traffic assignment: an application of a method of Lyapunov, *Transportation Sci.* 18 (1984), 245–252.

[SSV20] P. Schultz, D. I. Spivak, C. Vasilakopoulou, Dynamical systems and sheaves, *Appl. Categ. Structures* 28 (2020), 1–57.

[TaJo78] P. D. Taylor, L. B. Jonker, Evolutionarily stable strategies and game dynamics, *Math. Biosci.* 40 (1978), 145–156.

[TCH05] A. Traulsen, J. C. Claussen, C. Hauert, Coevolutionary dynamics: from finite to infinite populations, *Phys. Rev. Lett.* 95 (2005), 238701.

[Tho85] B. Thomas, On evolutionarily stable sets, *J. Math. Biol.* 22 (1985), 105–115.

[Wei95] J. W. Weibull, *Evolutionary Game Theory*, MIT Press, 1995.

[You93] H. P. Young, The evolution of conventions, *Econometrica* 61 (1993), 57–84.

[Zee80] E. C. Zeeman, Population dynamics from game theory, in *Global Theory of Dynamical Systems*, Lecture Notes in Math. 819, Springer, 1980, 471–497.