Evolutionary Game Theory in Higher Categories

Stability, stochastic selection,
and compositional models

Anonymous

November 2026

Contents

I Fitness, invasion, and stability

1 Population states and payoff maps **1**

1.1 Distributions and relabeling 1

1.2 Random matching and expected fitness 2

1.3 Exact maps between games 3

1.4 Parameterized populations 3

2 Evolutionary stability and invasion **5**

2.1 Best replies and the resident population 5

2.2 The invasion test and the two inequalities 5

2.3 Neutral stability and sets of solutions 6

2.4 Transporting the comparison 7

2.5 A population formulation of best replies 7

3 Local superiority and mixed strategies **9**

3.1 The critical cone 9

3.2 Uniform superiority 10

3.3 Interior residents and the quadratic test 11

3.4 Persistence in payoff families 11

3.5 A local margin and continuation on a face 12

4 Payoff equivalence and spaces of games **13**

4.1 The quotient by common fitness increments 13

4.2 Relabeling as a groupoid action 14

4.3 Aggregation and local labels 15

4.4 Games with selected residents 16

II Dynamics on population states

5 Replicator flows and mutation **17**

5.1 From abundance to proportion 17

5.2 The simplex and its faces 17

5.3 Two mutation conventions 19
5.4 Time changes and the strength of mutation 20

6 Invasion and Lyapunov geometry 21
6.1 Relative entropy near a resident state 21
6.2 Selection as a metric gradient 22
6.3 Dissipation and a conservative obstruction 23

7 Revision protocols and response dynamics 25
7.1 Response as a differential inclusion 25
7.2 Pairwise comparison 26
7.3 Logit revision and regularization 27

8 Interacting populations and spatial structure 29
8.1 Product simplices and assortative encounters 29
8.2 Exact reduction along equitable partitions 30
8.3 A potential from detailed balance 31
8.4 Migration and exponential synchronization 31

III Stochastic populations and aggregation

9 Finite populations and Markov dynamics 33
9.1 Count states and fitness 33
9.2 Birth, mutation, and replacement 34
9.3 Fixation probabilities retain the update law 36

10 Rare mutation and stochastic stability 37
10.1 A theorem for stationary weights 37
10.2 Leading powers and stable states 38
10.3 Eliminating transient excursions 39
10.4 Rare steps and fixation 40

11 Mean-field limits and interaction structure 41
11.1 A generator on the frequency simplex 41
11.2 Recovering deterministic selection 43
11.3 When counts close over a graph 43

12 Stochastic categories and population aggregation **45**
12.1 Kernels as morphisms 45
12.2 Maps of dynamics 46
12.3 Strong lumpability 46
12.4 Generators and approximate quotients 47

IV Games and dynamics with interfaces

13 Open games and the continuation interface **49**
13.1 Lenses as the deterministic interface 49
13.2 Strategies and contextual responses 50
13.3 Sequential composition 50
13.4 A compositional equilibrium test 52
13.5 Deterministic adapters 52

14 Parallel games and higher comparisons **53**
14.1 Parallel composition 53
14.2 Strategy maps and two-dimensional structure 54
14.3 Probability and continuation transport 55
14.4 An equilibrium interface and a revision rule 55

15 Open dynamics and network wiring **57**
15.1 Input and output ports 57
15.2 Replicator modules 58
15.3 Composition by state identification 58
15.4 Boundary behavior of steady states 59

16 Semiconjugacy and deterministic aggregation **61**
16.1 Flow maps and projectable fields 61
16.2 Partition quotients of selection 62
16.3 Comparison with linear lumpability 63
16.4 Observed equilibria and stability 63
16.5 Finite-time error under a nonzero defect 63

V Descent, compression, and learning

17 Fitness systems over varying contexts **65**
17.1 An internal strategy object 65
17.2 Local generators and their gluing 66

17.3 Parametrization and effective descent 67
17.4 Monodromy and restricted parameter choices 68

18 Selection under stochastic compression **69**
18.1 Columns as conditional distributions 69
18.2 Composition of two compressions 70
18.3 Trajectory error on a bounded time interval 71
18.4 The tangent and cotangent interpretation 71

19 Action groupoids and symmetric selection **73**
19.1 The fixed simplex and equilibrium 73
19.2 What the quotient records 74
19.3 Coherent actions in an ∞-topos 75

20 Learning rules and populations of agents **76**
20.1 An update from relative entropy 76
20.2 The pathwise regret bound 77
20.3 Averaged joint play 77
20.4 Outer selection and the evaluation interface 78

A Notation **79**

B Machine-checked selection identities **80**
B.1 Drift and the normalized update 80
B.2 Deterministic aggregation 80
B.3 Reproducing the check 81

C Subject index **82**

Bibliography **83**

Part I Fitness, invasion, and stability

CHAPTER 1

Population states and payoff maps

Evolutionary game theory uses a payoff comparison to describe how the composition of a population changes the prospects of its members. The state of a population is a distribution of strategic types. The same set of types can support many different models, since changing the matching process or the rule for converting payoffs to growth can change the resulting dynamics. This chapter isolates the distribution and payoff operations used in the pairwise, randomly matched model [1, Chs. 1–2].

1.1 Distributions and relabeling

For a nonempty finite set I of types, the population state space is the simplex

$$\Delta(I) = \left\{ x \in \mathbb{R}^I : x_i \geq 0, \quad \sum_{i \in I} x_i = 1 \right\}.$$

Write δ_i for the state concentrated on i. The support of x is the set of its positive coordinates. For a nonempty subset $S \subseteq I$, the inclusion $\Delta(S) \to \Delta(I)$ inserts zero in the omitted coordinates. Its image is a face, and its relative interior consists precisely of the states of support S. The affine hull of the full simplex is the hyperplane $\sum_i x_i = 1$, whose translation space is $T_I = \{v : \sum_i v_i = 0\}$. Feasible directions at a boundary state are those $v \in T_I$ with $v_i \geq 0$ whenever $x_i = 0$. Indeed, this is necessary for $x + tv$ to have nonnegative coordinates at positive t. It is sufficient for small t because only finitely many of the positive coordinates can decrease.

A map of type sets $f : I \to J$ pushes a distribution forward by

$$(f_* x)_j = \sum_{f(i)=j} x_i.$$

Its linear extension uses the same formula on all vectors. If $g : J \to K$, collecting the terms in the sums indexed by each fiber of gf gives $(gf)_* = g_* f_*$. The identity map gives the identity on distributions. Thus simplex formation is a functor from finite sets to convex sets and affine maps. It extends on all sets to $\mathsf{D}I$, the set of finitely supported probability distributions. A map on Dirac

states extends uniquely to an affine map out of $\mathsf{D}I$, since every element is a finite combination $\sum_i x_i\delta_i$ and its coefficient function is uniquely determined by x. Accordingly a stochastic assignment of an output distribution to each type is determined by its values on types and then acts affinely on populations.

There is also a canonical way to combine mixtures of populations. For $w \in \mathsf{DD}I$, its barycenter has coordinates $\mu_I(w)_i = \sum_x w_x x_i$. Taking the barycenter of a Dirac distribution on a population returns that population. Mixing Dirac distributions on individual types returns their original coefficient distribution. The two orders of flattening three layers of mixtures agree because in either order the coefficient of i is obtained by adding the same products of three weights. This gives the unit and multiplication of the finite-distribution monad. The construction uses only finite sums even when I is infinite. It differs from fixing a population size, since the weights may be any nonnegative reals adding to one.

1.2 Random matching and expected fitness

A symmetric two-player game on I assigns the row player's payoff A_{ij} when the row type is i and its opponent is j. The opponent receives A_{ji} in that encounter. Symmetry is the exchange of player roles in this description. No symmetry of the entries of A is required. If the players draw types independently with distributions p, q, the expected payoff is

$$u_A(p,q) = \sum_{i,j} p_i A_{ij} q_j = p^{\mathsf{T}} A q.$$

Its two separate affine laws follow by collecting the coefficients of the product distribution. In particular, a randomly sampled opponent from a population x gives the fitness vector Ax, while the population average is $\overline{u}_A(x) = x^{\mathsf{T}}Ax$. A correlated matching prescription instead supplies a conditional distribution of opponents for each row type. If these distributions are $q(i,x)$, the corresponding fitness is $F_i(x) = \sum_j A_{ij} q(i,x)_j$. General matching rules therefore need not produce a linear fitness map, even when the encounter payoffs are fixed.

The pairing in u_A is functorial with respect to probability aggregation. It uses the transpose of the aggregation map to transport a score. For a score $b \in \mathbb{R}^J$, define $(f^*b)_i = b_{f(i)}$. Then

$$\sum_i x_i (f^*b)_i = \sum_j (f_*x)_j b_j.$$

This follows by collecting the left-hand terms over each fiber of f. The formula distinguishes the direction in which population weights move from the direction in which a fitness score is read back.

1.3 Exact maps between games

Consider the category whose objects are pairs (I, A) and whose maps $(I, A) \to (J, B)$ are type maps $f : I \to J$ satisfying $A_{ij} = B_{f(i),f(j)}$ for every i, j. Identity maps satisfy the same equation. If f and g are payoff-preserving, the comparison for gf follows by applying the two equalities successively. This category describes exact games with strategy maps. There are useful coarser notions of comparison, including maps that preserve only payoff differences, developed in Chapter 4.

Proposition 1.1 (Fitness under exact aggregation)**.** *If $f : (I, A) \to (J, B)$ is payoff-preserving, then*

$$Ax = f^*(Bf_*x), \qquad u_A(p, q) = u_B(f_*p, f_*q).$$

Conversely, either the fitness identity at all population states or the payoff identity at all pairs of distributions implies payoff preservation.

Proof. Regroup the terms of $(Ax)_i = \sum_k B_{f(i),f(k)}x_k$ over the fibers of f. This gives the claimed fitness identity. The identity for payoffs then follows from the pairing between aggregation and score pullback. Conversely, evaluating the first identity at δ_j gives equality in each component i. Evaluating the second identity at (δ_i, δ_j) also gives equality of entries. □

$$\begin{array}{ccc} \Delta(I) \times \Delta(I) & \xrightarrow{f_* \times f_*} & \Delta(J) \times \Delta(J) \\ \Big\downarrow{\scriptstyle u_A} & & \Big\downarrow{\scriptstyle u_B} \\ \mathbb{R} & = & \mathbb{R} \end{array}$$

Figure 1.1: An exact type map commutes with expected payoff.

A type bijection satisfying the comparison is a change of labels. A noninjective type map can identify distinct population states, while a nonsurjective map may leave types of the target unrepresented. If all target types have preimages, every target state can be represented by lifting its mass to selected types above its support. Such a lift need not be unique. Expected payoff agrees for all lifts by Proposition 1.1. A selection rule using the labels of the lifts can still see the difference. Passing to aggregated population states must therefore be justified for the particular rules one wants to transport.

1.4 Parameterized populations

Suppose payoffs vary continuously with an environmental parameter s in a space S, giving an array $A(s)$. Evaluation $(s, p, q) \mapsto u_{A(s)}(p, q)$ is continuous by its

finite-sum formula. Substituting a map $r : S' \to S$ pulls back the family to $A(r(s'))$, and evaluations agree after the same substitution. If local descriptions use different names for the types, a transition is a bijection of the corresponding type sets. An exact transition then gives the comparison in Figure 1.1, and these comparisons compose because pushforward does. The simplex may consequently vary as a bundle with transition maps permuting its coordinates. Expected payoff is then a single continuous function on pairs of states over the same parameter, as long as the local payoff matrices obey the transition comparisons. The groupoid of labels keeps track of both the state and its change of presentation.

This setup will support two further comparisons. Static stability asks which distributions do well against populations with a small mutant component. A dynamics asks where trajectories go after an actual change in the population. The payoff operation is common to the two questions, and additional assumptions on the trajectory rule will enter later.

The exact payoff condition is closed under pullback of the type labels. Given B on J and any $f : I \to J$, there is a unique matrix on I for which f is exact. Its entries are $B_{f(i),f(j)}$. Pulling back twice gives the same matrix as pulling back along the composite. Thus the payoff spaces form a contravariant system over type sets, whereas the population spaces form a covariant system. Their comparison in the payoff square explains why treating both as covariant under a quotient of labels can change the model.

CHAPTER 2

Evolutionary stability and invasion

Fix a population simplex $\Delta(I)$ and an expected payoff u_A as in Chapter 1. We use the symmetric Nash condition

$$u_A(p,p) \geq u_A(q,p) \qquad (q \in \Delta(I)).$$

It says that no alternative earns more against p than a resident earns against p. An evolutionary stability test also asks what happens when the alternative changes the population against which these payoffs are evaluated. Maynard Smith and Price introduced the stability idea in the study of animal conflict [3]. Alexander develops its different formulations in [1, Ch. 2].

2.1 Best replies and the resident population

Let $M_p = \max_i (Ap)_i$. The value $u_A(p,p)$ is the p-weighted average of the numbers $(Ap)_i$. A pure maximizer earns M_p against p, while the average cannot exceed M_p. Hence p satisfies the symmetric Nash condition if and only if $u_A(p,p) = M_p$. Equivalently, every type in the support of p attains M_p. If any type with positive weight earned less, the weighted average would be strictly smaller than M_p.

This calculation also identifies the complete best-reply correspondence at p. It is the face spanned by those pure types at which Ap reaches its maximum. Strict symmetric Nash means all other population states $q \neq p$ earn less against p. In the mixed strategy simplex this forces the best-reply face to consist of a single point. In particular, a resident supported on several types can satisfy the Nash condition without satisfying its strict version.

2.2 The invasion test and the two inequalities

For a resident p, a competitor $q \neq p$, and a mutant share $0 < \varepsilon < 1$, put

$$x_\varepsilon = (1-\varepsilon)p + \varepsilon q.$$

The pair p, q may consist of mixed strategies. The parameter ε specifies the share of the competitor in the population. It does not describe the probabilities within either of those two strategies. Define the payoff advantage in the resulting population by

$$D_{p,q}(\varepsilon) = u_A(p, x_\varepsilon) - u_A(q, x_\varepsilon).$$

The resident p is evolutionarily stable if for each $q \neq p$ there exists a threshold $b(q) > 0$ such that $D_{p,q}(\varepsilon) > 0$ whenever $0 < \varepsilon < \min\{1, b(q)\}$. The threshold in this definition can depend on q. A threshold valid for all competitors is a different quantifier placement, addressed in Chapter 3.

Theorem 2.1 (Invasion criterion for pairwise games)**.** *A resident p is evolutionarily stable if and only if it is symmetric Nash and, for each $q \neq p$ with $u_A(p,p) = u_A(q,p)$, one has*

$$u_A(p,q) > u_A(q,q).$$

Proof. For each fixed competitor write $d_0 = u_A(p,p) - u_A(q,p)$ and $d_1 = u_A(p,q) - u_A(q,q)$. The affine law in the opponent argument gives

$$D_{p,q}(\varepsilon) = (1-\varepsilon)d_0 + \varepsilon d_1.$$

If this is positive on a sufficiently small interval of positive ε, continuity gives $d_0 \geq 0$. When $d_0 = 0$, the same interval condition requires $d_1 > 0$ by the displayed identity. Conversely, if $d_0 > 0$, the advantage remains positive for all sufficiently small ε by continuity. If $d_0 = 0$ and $d_1 > 0$, the advantage is $\varepsilon d_1 > 0$ throughout the range of positive mutant shares. Applying the argument to every $q \neq p$ gives the equivalence. □

The two terms are evaluated in different environments. The first evaluates both strategies against a purely resident population. The second evaluates both against the competing strategy. In the event of a first-order tie, the strictly positive second comparison explains why the resident gains a strict advantage when a positive mutant share is present. Every strict symmetric Nash equilibrium is evolutionarily stable, since none of its competitors ties in the first comparison.

For a general continuous fitness map the first comparison still gives a necessary condition for positive advantage near zero mutant share. The exact tie test in Theorem 2.1 uses affine dependence on the population. With nonlinear dependence, the advantage along the segment can have a nontrivial shape. An endpoint comparison then requires additional hypotheses before it determines the sign near the resident.

2.3 Neutral stability and sets of solutions

Neutral stability uses the weak inequality $D_{p,q}(\varepsilon) \geq 0$ in the invasion test. This standard variant allows a competing strategy to match the resident's score when its share is small [2, Ch. 2]. The same proof, with a weak inequality on the tie condition, shows that p is neutrally stable exactly when it is symmetric Nash and $d_1 \geq 0$ for each distinct q with $d_0 = 0$. Thus the solution predicates satisfy

$$\mathrm{ESS}(A) \subseteq \mathrm{NS}(A) \subseteq \mathrm{Nash}(A).$$

These are subsets of the state space. An equality between local payoffs in the invasion test need not select one of the tied states. Persistence or attraction of the resident depends on the dynamics chosen to resolve or retain such ties.

The symmetric Nash condition defines a closed relation between payoff matrices and population states. Indeed, for each fixed pure type i it is the nonnegativity of the continuous difference $u_A(p,p) - (Ap)_i$, and the pure inequalities imply the conditions for all mixtures by averaging. An evolutionary stability relation adds a conditional strict inequality. For instance, along the parameter ray $A(t) = tA$ for $t > 0$, all comparisons in Theorem 2.1 retain their sign. At $t = 0$ all comparisons vanish. If the simplex has more than one point, every resident is then neutral and no resident is strictly evolutionarily stable. Whenever the positive ray has an ESS this supplies a limiting failure of strict stability. Thus a family of games with a continuous payoff map may have different solution fibers at limiting parameter values.

2.4 Transporting the comparison

A label change transports all these predicates. To see this, apply a payoff-preserving bijection $f : (I, A) \to (J, B)$. The map f_* respects the mixture x_ε by affinity and preserves its payoff by Proposition 1.1. It also gives a bijection between the competitors of p and of f_*p. The inequalities and their quantifiers therefore agree. For a family with label transitions the stable locus can be described on each chart and glued as a subset of the simplex family by this argument. It retains the topology inherited from that family when the payoffs vary in a topological base.

For a payoff-preserving surjection f, the Nash predicate obeys the same comparison, since every competitor of f_*p can be lifted. Strict stability has an additional obstruction. If $q \neq p$ has $f_*q = f_*p$, the payoff identity makes $D_{p,q}(\varepsilon) = 0$ for every share. The lifted matrix can thus preserve all coarse payoff information and have additional tied alternatives upstairs. This is one reason to record the class of game maps before using it to compare solution sets.

2.5 A population formulation of best replies

The Nash relation has a useful description by the subset of types with maximal score. Let $B \subseteq I$ be nonempty, and consider the part of the simplex where $B_p = B$. On this part, belonging to the best-reply face is equivalent to the coordinate condition $p_i = 0$ for all $i \notin B$. The region $B_p = B$ is specified by equality of the scores on B and strict inequalities between the common score and each score outside B. For a fixed matrix these are linear equality and inequality constraints on p. Thus the support condition can be checked on these regions

without a search over all mixed competitors. Evolutionary stability requires a further comparison on the face of tied competitors. That is a quadratic condition and need not reduce to checking only its vertices.

If p is strictly Nash, its payoff gap on the region of competitors at distance at least $r > 0$ from p has a positive minimum whenever that region is nonempty. This follows from continuity on a closed subset of the compact simplex. There is generally no positive minimum over every distinct competitor, since a sequence of competitors can approach p. The existence of a uniform invasion barrier is instead a uniform statement about the relative contributions of the first and second payoff differences. The next chapter establishes it by examining the directions in which the first difference vanishes.

CHAPTER 3

Local superiority and mixed strategies

The invasion test fixes a competitor and then varies its share. Local superiority compares the resident with every nearby population against that population itself. In finite-dimensional matrix games, critical directions at the resident relate these two descriptions. They also isolate the part of the payoff matrix that controls stability of a mixed resident. The equivalence with a uniform invasion barrier is discussed in [1, Ch. 2] and [2, Ch. 2].

3.1 The critical cone

Write $T_I = \{v : \sum_i v_i = 0\}$. For $p \in \Delta(I)$ let $C_p = \{v \in T_I : v_i \geq 0 \text{ if } p_i = 0\}$. Every displacement $q-p$ belongs to C_p, and every $v \in C_p$ is an actual displacement up to multiplication by a small positive scalar. Introduce the linear and quadratic forms

$$L_p(v) = -v^{\mathsf{T}} A p, \qquad Q_A(v) = v^{\mathsf{T}} A v$$

and the critical cone $K_p = C_p \cap \ker L_p$. The Nash condition is exactly $L_p \geq 0$ on C_p. In one direction one scales each feasible vector to an actual displacement and applies Nash. In the other direction one applies the inequality to $q - p$ for arbitrary q.

If $q = p + v$, the two terms in the invasion criterion are $d_0 = L_p(v)$ and $d_1 = L_p(v) - Q_A(v)$. It follows from Theorem 2.1 and scaling of the quadratic form that

$$p \text{ is an ESS} \quad \Longleftrightarrow \quad L_p|_{C_p} \geq 0, \qquad Q_A(v) < 0 \quad (0 \neq v \in K_p).$$

The critical cone consists of directions with no initial loss against the resident. The quadratic condition applies only on that cone. At a pure resident with a unique best response, for example, there is no nonzero critical direction. At a boundary resident with additional tied best responses, the cone can include directions that add previously absent types. Restricting the calculation to the resident's support can omit those directions.

3.2 Uniform superiority

Call p locally superior if $u_A(p,x) > u_A(x,x)$ for every $x \neq p$ sufficiently near p in the simplex. A uniform invasion barrier is one positive threshold that works for all distinct competitors in the definition of evolutionary stability.

Theorem 3.1. *For a matrix game on a finite type set, evolutionary stability, local superiority, and the existence of a uniform invasion barrier are equivalent.*

Proof. Suppose p is an ESS. For $v = x - p$ the superiority difference is

$$u_A(p,x) - u_A(x,x) = L_p(v) - Q_A(v).$$

If there is no nonzero feasible direction, the local condition is vacuous. Otherwise put $S = C_p \cap \{v : \|v\| = 1\}$, a compact set. On its critical subset $S \cap K_p$, the quadratic form is negative. The relatively open subset of S on which $Q_A < 0$ contains this critical set. On its complement, if nonempty, L_p has a strictly positive minimum m by compactness and the fact that its zeros are critical. Also $|Q_A|$ has a finite maximum M on S. For $w \in S$ and $r > 0$ we have

$$L_p(rw) - Q_A(rw) = r\big(L_p(w) - rQ_A(w)\big).$$

It is positive wherever $Q_A(w) < 0$ because $L_p \geq 0$. On the complementary subset it is positive for $0 < r < m/(1+M)$. This proves local superiority, including the case in which the complementary subset is empty.

Now assume local superiority on a neighborhood of radius $\delta > 0$. The simplex has finite diameter D. For each mutant $q \neq p$, the invader mixture $x_\varepsilon = (1-\varepsilon)p + \varepsilon q$ satisfies $0 < \|x_\varepsilon - p\| \leq \varepsilon D$. The identity

$$u_A(p,x_\varepsilon) - u_A(x_\varepsilon,x_\varepsilon) = \varepsilon\big(u_A(p,x_\varepsilon) - u_A(q,x_\varepsilon)\big)$$

follows from affinity in the first argument. Hence every share below $\min\{1, \delta/(1+D)\}$ gives positive invasion advantage independently of q. This yields a uniform invasion barrier. A uniform threshold gives the pointwise thresholds of evolutionary stability immediately. □

Two features of the proof have distinct roles. Compactness of the unit directions controls approach to the critical cone. Boundedness of the simplex controls the passage from nearby populations to a common mutant share for every q. General strategy spaces can lose either property. The finite-dimensional proof cannot then be transferred without a replacement estimate.

3.3 Interior residents and the quadratic test

In the relative interior of $\Delta(I)$, the cone C_p is the whole tangent space T_I. A linear form nonnegative on a vector space is zero. The Nash condition therefore says that Ap is constant in the type index. For an interior Nash resident, $K_p = T_I$, and evolutionary stability is equivalent to the restriction of the symmetric part $(A + A^{\mathsf{T}})/2$ being negative definite on T_I. Indeed $v^{\mathsf{T}}Av$ only sees that symmetric part. When $A = A^{\mathsf{T}}$, this is the Hessian test for the quadratic function $x \mapsto x^{\mathsf{T}}Ax/2$ restricted to the affine hull of the simplex. Its differential at a Nash state vanishes along T_I, and its exact expansion has quadratic term $Q_A/2$. When A is not symmetric, the argument establishes the sign of the critical quadratic form. The derivative of this same quadratic function contains the symmetric part and in general differs from the fitness vector Ax.

Negative definiteness on T_I also implies uniqueness of a symmetric Nash state whenever one exists. If p, q both satisfy Nash, adding $u_A(p,p) \geq u_A(q,p)$ and $u_A(q,q) \geq u_A(p,q)$ gives $(p-q)^{\mathsf{T}}A(p-q) \geq 0$. Since $p - q \in T_I$, this forces equality of the states under negative definiteness.

3.4 Persistence in payoff families

The interior test has a direct parameter consequence. Let Π be the Euclidean projection onto T_I. If $Q_A < 0$ on $T_I \setminus \{0\}$, the operator $\Pi A|_{T_I}$ is invertible. Indeed, a vector v in its kernel satisfies $v^{\mathsf{T}}Av = 0$. Injectivity then gives invertibility on the finite-dimensional space T_I. For any fixed b in the affine hull of the simplex, the unique solution of $\Pi Ap = 0$ in that affine hull is

$$p_A = b - (\Pi A|_{T_I})^{-1}\Pi Ab.$$

The expression varies continuously with A near a matrix for which the operator is invertible. Strict positivity of each coordinate is open. Negative definiteness is also open, by taking the maximum of Q_A over the compact unit sphere in T_I when $T_I \neq 0$. Thus an interior ESS satisfying this test persists through sufficiently small payoff perturbations, and in that neighborhood the displayed formula gives the unique Nash state. Relabeling commutes with the zero-mean projection, and uniqueness makes these locally computed states agree wherever the labels change. This is an explicit instance of transport of a solution in a parameter family.

3.5 A local margin and continuation on a face

The compact-direction argument yields a useful quantitative version. If p is an ESS and the feasible unit set is nonempty, there are constants $c, \delta > 0$ such that

$$u_A(p, x) - u_A(x, x) \;\geq\; c\|x - p\|^2 \qquad (\|x - p\| < \delta).$$

Indeed, if the critical unit set is nonempty, choose $a > 0$ with $Q_A \leq -a$ there. On the directions where $Q_A < -a/2$, the difference at $v = rw$ is at least $(a/2)r^2$. On the closed complement the minimum of L_p is some $m > 0$. For $r < \min\{1, m/(2(1+M))\}$, the difference is at least $(m/2)r \geq (m/2)r^2$. If either region is empty only the other bound is needed. When the critical unit set is empty the second estimate applies on all directions. The lower bound is an estimate on payoffs. Its use as an estimate on motion requires a dynamics whose drift is related to those payoffs.

At a Nash point, let $B_p = \{i : (Ap)_i = M_p\}$, where $M_p = \max_i (Ap)_i$. Since $\sum_i v_i = 0$, one has

$$L_p(v) = \sum_{i \notin B_p} (M_p - (Ap)_i) v_i.$$

All terms on the right are nonnegative for feasible directions because every omitted index has zero weight in p. Consequently $v \in K_p$ precisely when $v \in C_p$ and $v_i = 0$ outside B_p. If every absent type has a strictly lower payoff, then $B_p = \operatorname{supp}(p) = S$. The critical cone is simply the tangent space to $\Delta(S)$. An ESS of this kind with a strict quadratic test on that tangent space persists on the same face under small matrix perturbations: solve the interior equation for the restricted matrix, extend by zero, and use continuity to preserve the payoff gaps to all absent types. Any Nash state sufficiently near p for a sufficiently small perturbation still has positive coordinates on S, and the same payoff gaps prevent extra positive coordinates outside S. Its uniqueness therefore follows from the restricted affine equation. Here the strict gaps and the curvature on the face control different aspects of the same continuation.

CHAPTER 4

Payoff equivalence and spaces of games

The payoff array contains information about the absolute score of each encounter. The static comparisons in Chapter 2 and the usual replicator field use only relative scores against the same population. This difference suggests an equivalence relation on games with fixed type labels. Once labels are also allowed to vary, the resulting space of games has a groupoid presentation. Its arrows identify numerical presentations of the same relative-fitness model.

4.1 The quotient by common fitness increments

Fix a nonempty finite type set I, put $n = |I|$, and write $\mathbf{1}$ for its constant vector. Let

$$\mathcal{A}_I = \mathrm{End}(\mathbb{R}^I), \qquad W_I = \{\mathbf{1}b^{\mathsf{T}} : b \in \mathbb{R}^I\}.$$

If $B = A + \mathbf{1}b^{\mathsf{T}}$, then

$$Bx = Ax + (b^{\mathsf{T}}x)\mathbf{1}.$$

The increment is common to every row type, though its size can depend on the opponent distribution x. For p, q, x in the simplex we have $(p-q)^{\mathsf{T}}(B-A)x = 0$, since p, q have the same total mass. Thus comparing p and q against x gives identical advantage in the two games. This includes both comparisons in the ESS test.

Define the zero-mean projection $\Pi_I v = v - \mathbf{1}(\sum_i v_i)/n$. Its kernel is the span of $\mathbf{1}$ and its image is T_I. The map $A \mapsto \Pi_I A$ has kernel W_I: each column is killed precisely when that column is constant across the rows. It maps onto the linear maps $\mathbb{R}^I \to T_I$, since the projection acts as the identity on their image. It follows that

$$\mathcal{Q}_I = \mathcal{A}_I / W_I \simeq \mathrm{Hom}(\mathbb{R}^I, T_I).$$

Each class has a canonical representative with zero column sums. Alternatively, knowing the relative scores against every population state determines this class. Indeed the relative scores give all row differences in $(A-B)x$. If these vanish for all x, evaluation at each pure state shows $A - B \in W_I$.

Proposition 4.1 (The quotient determined by replicator motion)**.** *Write X_A for the polynomial vector field on the simplex given by*

$$(X_A(x))_i = x_i((Ax)_i - x^{\mathsf{T}}Ax).$$

Then $X_A = X_B$ on the whole simplex if and only if $A - B \in W_I$. If $B = \lambda A + \mathbf{1}b^{\mathsf{T}}$ for constant $\lambda > 0$, then $X_B = \lambda X_A$ and the static stability predicates agree.

Proof. The common fitness increment cancels in the expression for X_A because the weights of x add to one. Scaling every fitness by λ scales both terms in the difference and every payoff comparison by the same positive number. This proves the forward formulas.

Conversely, take x in the relative interior of the simplex. Each $x_i > 0$, so equality of the fields implies $((A - B)x)_i = x^{\mathsf{T}}(A - B)x$ for all i. Thus $\Pi_I(A - B)x = 0$ on the relative interior. Continuity gives zero on its closure, and in particular on every pure state. The column argument then places $A - B$ in W_I. □

If a solution curve for X_A is parameterized by t, changing its parameter to λt gives a solution for X_B on the corresponding time interval. Positive scaling therefore gives a comparison of time units. Keeping time fixed gives the sharper quotient $\mathcal{Q}_I$. Neither equivalence identifies the unadjusted absolute payoffs. Those retain a role if another process, such as a baseline contribution to fitness, is sensitive to their absolute levels.

4.2 Relabeling as a groupoid action

Let $\mathfrak{S}_I$ be the permutation group of the type set. A permutation σ acts by $A \mapsto P_\sigma A P_\sigma^{-1}$, where $P_\sigma \delta_i = \delta_{\sigma(i)}$. It preserves W_I, since it fixes $\mathbf{1}$. It also commutes with Π_I, since it preserves the sum of the coordinates. The action therefore passes to $\mathcal{Q}_I$. Its action groupoid has objects $q \in \mathcal{Q}_I$ and arrows $q \to q'$ given by permutations σ for which $q' = \sigma q$. Two matrices in the same orbit determine isomorphic objects. The permutations fixing a class are the automorphisms of that object. Passing to the set of orbits forgets this group. Accordingly, the orbit set classifies a single game up to relabeling, while the groupoid can also record changes of presentation along a parameter space.

One can equivalently keep the raw matrices and allow translation by W_I in the arrows. The group $W_I \rtimes \mathfrak{S}_I$ acts on $\mathcal{A}_I$ by $(w, \sigma)A = \sigma A + w$. Its multiplication uses the evident permutation action on W_I. Consider the functor of action groupoids

$$\mathcal{A}_I // (W_I \rtimes \mathfrak{S}_I) \longrightarrow \mathcal{Q}_I // \mathfrak{S}_I$$

sending A to its class and an arrow (w, σ) to σ. It is fully faithful: if the classes of A, B satisfy $[B] = \sigma[A]$, then $w = B - \sigma A$ belongs to W_I and is the unique translation giving an arrow above σ. It is essentially surjective by the canonical normalized representative $\Pi_I A$ in each class. A quasi-inverse uses this representative, sending each permutation to an arrow with translation zero. The comparison back to a raw matrix is translation by its difference from its normalized form. The formulas are continuous and even smooth as functions of the matrices. Consequently they also give equivalent presentations for families with continuous or smooth transition data. Here the removal of redundant payoff coordinates gives an equivalence of groupoids because the arrows over each permutation are recovered uniquely.

$$\begin{array}{ccc} \mathcal{A}_I & \xrightarrow{\sigma} & \mathcal{A}_I \\ {\scriptstyle \Pi_I(-)}\downarrow & & \downarrow{\scriptstyle \Pi_I(-)} \\ \mathcal{Q}_I & \xrightarrow{\sigma} & \mathcal{Q}_I \end{array}$$

Figure 4.1: Canonical normalization respects every permutation of labels.

4.3 Aggregation and local labels

For a parameter base S, a family with local labels can be given on an open cover by maps to $\mathcal{Q}_I$ along with permutations on overlaps satisfying the cocycle law. The comparison on overlaps is the one in Figure 4.1, and population states transform by the same permutations. Thus the payoff family and population bundle can carry nontrivial label transport. For example, a loop can act by a permutation fixing the payoff class and moving some points of the simplex. A solution family that is required to be a section must respect that transport. A numerical solution set at a single parameter point does not specify how to continue the individual solutions around such a loop.

A payoff-preserving surjection of type sets is a different operation. It can remove duplicate labels by summing their population weights. Its fibers over interior mass assignments can have positive dimension. It therefore fails to identify the underlying state spaces even when the expected payoff descends. In particular, if a state puts positive mass above a target type that has at least two preimages, a nonzero transfer of mass within that fiber leaves its image unchanged. Chapter 2 shows that such a transfer prevents the specified lift from being an ESS. The groupoid equivalence for normalization uses actual comparisons with inverses. A coarse aggregation needs a separate argument for each invariant transported along it.

4.4 Games with selected residents

The total solution space is more informative for families than a list of stable residents for each matrix. For the ESS predicate, form the subspace

$$\mathcal{S}_I = \{(q, p) \in \mathcal{Q}_I \times \Delta(I) : p \text{ is evolutionarily stable at } q\}.$$

The definition is independent of the representative of q. Permuting the game and the resident together preserves this subspace, by the payoff computation in Chapter 2. Projection therefore gives a functor $\mathcal{S}_I//\mathfrak{S}_I \to \mathcal{Q}_I//\mathfrak{S}_I$. In a parameter family, a selected resident is a lift along this projection. In local coordinates it consists of a choice on each chart and compatibility of the choices on every overlap.

Over a chosen game class with a specified identification of its labels, the stable residents form the ordinary set $\mathrm{ESS}(q)$. If one instead restricts to the full subgroupoid of games isomorphic to q, then the category of marked residents is equivalent to the action groupoid of $\mathrm{Stab}(q)$ on $\mathrm{ESS}(q)$. Indeed, bringing a game in that orbit back to q takes its resident to one at q, and different choices of a permutation differ by a stabilizer element. The groupoid records how residents can be identified by changes of labels. This construction only uses the stated stability predicate and its invariance under permutations. It applies equally to the Nash and neutral predicates.

CHAPTER 5

Replicator flows and mutation

5.1 From abundance to proportion

Let S be a nonempty finite strategy set and write $\Delta(S)$ for its probability simplex. A fitness assignment is a map $F : \Delta(S) \to \mathbb{R}^S$. Its i-th entry is the rate associated with strategy i, evaluated in the population x. For a matrix A, the assignment $F(x) = Ax$ evaluates the expected payoff of each strategy against the population. The derivation below applies to nonaffine assignments as well. Taylor and Jonker give an early dynamical treatment, and Alexander explains the abundance and revision interpretations of the equation [4, 1].

Suppose n_i is the abundance of strategy i and that $\dot{n}_i = n_i F_i(x)$, where $x_i = n_i / \sum_j n_j$. Writing $M = \sum_j n_j$ and $\phi(x) = \sum_j x_j F_j(x)$, one obtains $\dot{M} = M\phi(x)$. The quotient rule gives

$$\dot{x}_i = x_i\big(F_i(x) - \phi(x)\big). \tag{5.1}$$

The subtraction of mean fitness is thus the correction due to normalization. Conversely, if x solves (5.1), then $M(t) = M(0)\exp(\int_0^t \phi(x(s))\,ds)$ recovers a positive total abundance and $n_i = Mx_i$ solves the abundance equation. The abundance formulation uses independence of the fitness assignment from the total. If the rates depend on total abundance, the proportions generally retain that dependence and need an additional state variable.

There are two simple symmetries of (5.1). Adding $c(x)$ to every $F_i(x)$ adds $c(x)$ to $\phi(x)$ and leaves the field fixed. Relabeling S transports x and F together. For an affine assignment, adding the same row vector to each row of the payoff matrix is one instance of the first symmetry. Equality of vector fields records a specific observational equivalence of fitness data.

5.2 The simplex and its faces

Assume the fitness assignment extends to a locally Lipschitz map on a neighborhood of $\Delta(S)$. On that neighborhood form $\phi(x) = \sum_j x_j F_j(x)$ without requiring $\sum_j x_j = 1$. For a local solution x set $m(t) = \sum_j x_j(t)$. Then $\dot{m} = (1 - m)\phi(x)$.

If $m(0) = 1$, the integrating factor for this scalar equation implies $m(t) = 1$ throughout its interval. Independently, each coordinate satisfies

$$x_i(t) = x_i(0) \exp\left(\int_0^t \left(F_i(x(s)) - \phi(x(s)) \right) ds \right). \tag{5.2}$$

This identity includes coordinates starting at zero, because the scalar linear equation along a given solution has the zero solution for such an initial value.

Proposition 5.1 (Existence and restriction to faces)**.** *For the locally Lipschitz assignment just specified, (5.1) has a unique solution for every initial state of $\Delta(S)$, defined for all real times. Throughout any finite time interval its support equals the initial support. On the face indexed by a nonempty subset $T \subseteq S$, the solution is the replicator solution for the restricted assignment on $\Delta(T)$.*

Proof. Local existence and uniqueness follow from the locally Lipschitz field. Conservation of mass and (5.2) keep a solution in the simplex throughout its interval of existence, in either direction of time. A solution confined to this compact subset of the domain can be continued at each finite endpoint. Support preservation follows from the positivity of the exponential. On the face the terms $x_j F_j(x)$ with $j \notin T$ vanish from the mean, and each missing coordinate has zero velocity. Uniqueness identifies the restricted trajectory with the solution calculated directly in the face. □

In particular, if $j_T : \Delta(T) \to \Delta(S)$ extends a state by zero, and F_T restricts $F \circ j_T$ to the coordinates in T, the solution operators obey the commutative square

$$\begin{array}{ccc} \Delta(T) & \xrightarrow{\Phi_t^{F_T}} & \Delta(T) \\ \downarrow & & \downarrow \\ \Delta(S) & \xrightarrow{\Phi_t^{F}} & \Delta(S). \end{array}$$

Here both vertical maps are j_T. The group law $\Phi_{s+t} = \Phi_s \circ \Phi_t$ follows by uniqueness applied after a change of time origin. A relabeling yields the same square with the vertical maps replaced by permutations. These comparisons fix the sense in which the construction respects the data of a game. A map between strategy sets that identifies distinct types needs separate hypotheses to carry the differential equation through the induced map on states.

A rest state has equal fitness $\phi(x)$ on its support. This condition is sufficient because the velocity of every absent type is already zero. A Nash state also bounds the payoffs of absent types above by $\phi(x)$. The distinction is caused by support preservation. In the interior the two conditions coincide. For any pair of present types,

$$\frac{d}{dt} \log \frac{x_i}{x_j} = F_i(x) - F_j(x).$$

Thus a constant lower bound $\gamma > 0$ for the payoff difference along an interval $[0, t]$ gives $x_i(t)/x_j(t) \geq e^{\gamma t} x_i(0)/x_j(0)$. The inference concerns the relative growth of two present types. It yields no positive mass for a type initially absent.

5.3 Two mutation conventions

A revision generator Q has nonnegative off-diagonal entries and zero row sums, with Q_{ji} denoting the rate from j to i for $j \neq i$. Independent mutation during the evolution of the proportions gives

$$\dot{x}_i = x_i(F_i(x) - \phi(x)) + \sum_j x_j Q_{ji}. \tag{5.3}$$

The generator contributes zero to the mass derivative. On the simplex the contribution to a vanishing coordinate is the sum of the incoming rates weighted by the masses. This property yields forward invariance even when the initial state lies on the boundary. To prove it, in the ambient equation replace x_j by $\max(x_j, 0)$ in the incoming terms with $j \neq i$. The modified problem is locally Lipschitz. In the equation for coordinate i the terms involving x_i are linear in x_i along a given solution, and the remaining forcing is nonnegative. Variation of constants proves that all coordinates stay nonnegative for $t \geq 0$. The modification consequently agrees with the original equation on this solution. The mass identity then retains $m = 1$, and compactness permits global forward continuation.

For a constant Q, a positive coordinate remains positive by the same variation formula. If $Q_{ji} > 0$ and coordinate j is positive for $t > 0$, its forcing makes i positive for $t > 0$ as well. Induction along directed edges shows that irreducibility, meaning directed reachability between all strategies, produces full support at every strictly positive time. Mutations change the invariant face structure present in (5.1). They also destroy its backward invariance on the closed simplex in general.

Another convention incorporates mutation into reproduction. Let R be a row stochastic matrix of offspring probabilities and let $b_j(x) \geq 0$ be an absolute birth rate. With fitness used in these units the normalized equation is

$$\dot{x}_i = \sum_j x_j b_j(x) R_{ji} - x_i \sum_j x_j b_j(x). \tag{5.4}$$

This follows by normalizing the abundance equation with incoming births $\dot{n}_i = \sum_j n_j b_j(x) R_{ji}$. Setting R equal to the identity recovers selection, and the sum of the velocities on the simplex is zero. Under the same local Lipschitz hypothesis, forward invariance follows by the incoming-rate argument. For that argument the rates can be extended off the simplex by their positive parts so that incoming forcing remains nonnegative. For the diagonal term the coefficient is $b_i R_{ii} - \sum_j x_j b_j$.

A common baseline added to b_j changes (5.4) by $c(R^{\mathsf{T}}x - x)$. This describes mutation occurring during additional births. In (5.3), the corresponding baseline change to F has no effect. The two conventions therefore require different information about which transformations of fitness preserve the model. The formulation of an evolution operator should specify where mutation enters before passing to equivalence classes of payoff data [6].

5.4 Time changes and the strength of mutation

Multiplying a fitness assignment by a positive constant a changes the replicator solution from $x(t)$ to $x(at)$. The geometric trajectories are the same and their parametrizations differ. The same rule holds for a positive continuous scalar $a(x)$ on the simplex. Along a trajectory of the rescaled field, the clock for the original field has rate $a(x(t))$. On the compact simplex, the positive minimum of a ensures that this clock advances without bound as physical time grows. Thus this kind of rescaling preserves forward limit sets as well as rest points.

In the additive mutation model, the same clock rescales both selection and the generator. Holding Q fixed while rescaling F can change the stationary states. More precisely, if x is stationary for both fitness F and fitness aF with the same Q and a constant $a \neq 1$, subtraction of the two stationary equations gives $R_F(x) = 0$. The original equation then gives $Q^{\mathsf{T}}x = 0$. Conversely, a common zero of these two fields is stationary for every relative scaling. When Q is irreducible this test picks out at most its unique stationary distribution [6]. Consequently, fitness units and mutation rates carry joint information when the model includes both mechanisms.

CHAPTER 6

Invasion and Lyapunov geometry

6.1 Relative entropy near a resident state

A state $p \in \Delta(S)$ is locally superior for F if

$$(p - x) \cdot F(x) > 0$$

for every distinct x sufficiently close to p in the simplex. This formulation compares the resident with a nearby population in the environment of that population. For symmetric matrix games it is equivalent to the usual evolutionarily stable strategy condition discussed in Chapter 2. The implication to dynamical stability is the classical result studied by Hofbauer, Schuster, and Sigmund [5, 1].

Set $U_p = \{x \in \Delta(S) : x_i > 0 \text{ whenever } p_i > 0\}$. The quantity

$$D(p\|x) = \sum_{i:p_i>0} p_i \log \frac{p_i}{x_i}$$

is finite on U_p, including when p lies in a proper face. The bound $-\log z \geq 1 - z$ at $z = x_i/p_i$ yields

$$D(p\|x) \geq 1 - \sum_{i:p_i>0} x_i \geq 0.$$

Equality demands $x_i = p_i$ on the support of p, which also forces the remaining entries of x to vanish. The derivative along (5.1) is

$$\frac{d}{dt} D(p\|x(t)) = x(t) \cdot F(x(t)) - p \cdot F(x(t)). \tag{6.1}$$

The cancellation uses $\sum_i p_i = 1$. Support preservation makes the logarithms legitimate for all finite times whenever they are defined initially.

Theorem 6.1 (Local stability from superiority)**.** *For the replicator equation of Chapter 5, a locally superior state p is a locally asymptotically stable rest point relative to $\Delta(S)$.*

Proof. For any $q \in \Delta(S)$, approach p along the segment $x = p + \epsilon(q - p)$. Dividing the superiority inequality by $\epsilon > 0$ and taking its limit shows $(q - p) \cdot F(p) \leq 0$. Thus p is Nash and is a rest point of the replicator field. Take a closed relative ball around p contained in U_p and in the neighborhood where superiority holds.

Taking the ball small enough is possible even when p lies on the boundary because only its positive coordinates need to stay positive. If the simplex is a singleton, stability is immediate. Otherwise, for a sufficiently small radius the relative sphere is nonempty, and the minimum there of $D(p\|-)$ is positive by compactness. Any trajectory starting inside with entropy below this minimum stays inside. By taking the radius arbitrarily small we obtain Lyapunov stability.

For such a trajectory, the entropy is nonincreasing and bounded below. If a point $z \neq p$ were a limit point, its entropy derivative would be strictly negative. The trajectory has bounded speed on the closed ball. It would therefore spend intervals of a fixed positive duration where the derivative is bounded above by a negative constant, on arbitrarily late visits near z. Taking disjoint visits contradicts convergence of the entropy values. The limit points are consequently all p. The trajectory lies in a compact set and hence converges to p. □

The same argument supplies a global conclusion on U_p if p is a Nash state and

$$(x - p) \cdot \big(F(x) - F(p)\big) < 0 \qquad (x \neq p)$$

on the whole simplex. The Nash inequality implies $(x - p) \cdot F(x) < 0$. A finite sublevel set of $D(p\|-)$ stays away from zeros of the required coordinates, so the compact-set argument applies there. For completeness, each term with $p_i > 0$ is bounded below by $p_i \log p_i$, while the term for x_i tends to infinity at $x_i = 0$. This gives positive lower bounds for the required coordinates on any one sublevel set. This strict inequality is a special sufficient condition. Hofbauer and Sandholm establish broader geometric criteria for convergence in stable population games [8].

6.2 Selection as a metric gradient

At an interior state x, tangent vectors are vectors u with $\sum_i u_i = 0$. The Shahshahani metric is

$$g_x(u, v) = \sum_i \frac{u_i v_i}{x_i}.$$

The same definition applies to the relative interior of each face by restricting the coordinates to its support [7, 6]. Define the fitness one-form on a face by $\alpha_x(u) = \sum_i F_i(x) u_i$. If R_F denotes the replicator vector field, its metric pairing with u is

$$g_x(R_F(x), u) = \sum_i (F_i(x) - \phi(x)) u_i = \alpha_x(u).$$

The field is the metric dual of the payoff one-form. Adding a state-dependent baseline to all payoffs leaves the one-form unchanged, because all tangent vectors have zero sum.

Pointwise, every tangent covector has a representative in the ambient strategy coordinates, and two such representatives differ by a multiple of $\mathbf{1}$. Writing X for the relative interior of the full simplex, the fitness, one-form, and field correspond by

$$\begin{array}{ccc} \mathbb{R}^S/\mathbb{R}\mathbf{1} & \xrightarrow{[f]\mapsto f\cdot(-)} & T_x^*X \\ & \searrow^{R_x} & \downarrow g_x^{-1} \\ & & T_xX. \end{array}$$

Here R_x denotes the linear map that projects a fitness vector by its resident mean and multiplies by the current frequencies. At a proper face the corresponding metric duality uses only supported coordinates. All payoff coordinates outside the support then disappear from the tangent covector, even though they are needed to evaluate invasion by absent types.

If $\alpha = dV$ for a differentiable potential V, this duality says that R_F is the Riemannian gradient of V. The rate of increase of the potential is

$$\frac{d}{dt}V(x(t)) = g_x(R_F, R_F) = \sum_i x_i(F_i(x) - \phi(x))^2. \tag{6.2}$$

Only supported coordinates enter this identity. For a potential continuously differentiable on the closed simplex in its affine hull, the right side extends continuously to every face. Compactness bounds the potential and the speed of the trajectory. As in the entropy argument, a nonrest limit point would give repeated fixed-duration intervals of strictly positive increase, which is impossible. This shows that each limit point is a rest point. It leaves convergence to a particular rest point open when the rest set has several possible limit points.

For matrix payoffs, write $A = S + K$ with S symmetric and K skew-symmetric. When $K = 0$, one can take $V(x) = x^{\mathsf{T}}Sx/2$. More generally A defines a potential on the interior precisely when the bilinear form $u^{\mathsf{T}}Av$ is symmetric on tangent vectors. Indeed, the exterior derivative of the payoff one-form on constant tangent vectors is $v^{\mathsf{T}}Au - u^{\mathsf{T}}Av$. This must vanish for a potential, and it suffices because the open simplex is convex. Thus a skew part matters for integrability only through its restriction to tangent directions.

6.3 Dissipation and a conservative obstruction

At an interior rest point p of a matrix game, $Ap = c\mathbf{1}$. Consequently

$$(x - p) \cdot Ax = (x - p)^{\mathsf{T}}S(x - p).$$

The right side depends only on the symmetric part because $u^{\mathsf{T}}Ku = 0$. Strict negative definiteness of S on the tangent hyperplane gives strict entropy dissipation. The skew part can nevertheless change the direction of the velocity.

For a purely skew field $A = K$ with $Kp = 0$, one instead has $\phi(x) = 0$ and $p^{\mathsf{T}}Kx = 0$ everywhere. Equation (6.1) makes $D(p\|x)$ a first integral on the interior. Every trajectory from a distinct interior state has positive constant entropy, so it cannot converge to p. This obstructs asymptotic stability regardless of whether individual level sets carry periodic motion.

Suppose that p is interior and that there is a strict uniform bound of the form $(x - p) \cdot F(x) \leq -c\|x - p\|^2$ near p, with $c > 0$. The entropy Hessian in the ambient positive orthant is $\mathrm{diag}(p_i/x_i^2)$. This Hessian is bounded above and below by positive definite diagonal matrices on a small closed neighborhood of p. The first derivative of entropy at p vanishes on the tangent space. Taylor's theorem therefore gives positive constants a, b such that

$$a\|x - p\|^2 \leq D(p\|x) \leq b\|x - p\|^2.$$

An initial state close enough stays in this neighborhood by Lyapunov stability. Then (6.1) implies exponential decay at rate at least c/b for the entropy. This argument uses the full support of the resident. At a boundary resident, redistributing mass to an absent strategy can give a linear term in the expansion. For a fixed q with mass m outside the support of p, direct differentiation yields

$$D(p\|p + \epsilon(q - p)) = \epsilon m + O(\epsilon^2) \qquad (\epsilon \downarrow 0).$$

The general stability implication accommodates this boundary term through the entropy sign and compactness.

These statements concern the vector field associated with a given payoff rule. Equation (6.2) uses the exact derivative of a potential and (6.1) uses the comparison with a resident. They give separate ways to turn a strategic inequality into a dynamical conclusion, with distinct hypotheses at the boundary and on the rest set.

CHAPTER 7

Revision protocols and response dynamics

7.1 Response as a differential inclusion

A payoff assignment gives current incentives. A revision protocol determines which of those incentives are used when a strategy changes. This extra datum affects boundary motion, the identification of rest points, and convergence. Alexander and Sandholm develop the passage from revision protocols to population dynamics [1, 10].

Write

$$B(x) = \arg\max_{y \in \Delta(S)} y \cdot F(x).$$

Because the objective is linear in y, this is the nonempty convex face spanned by the maximizers of $F_i(x)$. Best-response revision uses the differential inclusion

$$\dot{x}(t) \in B(x(t)) - x(t).$$

The solution convention is an absolutely continuous path satisfying the inclusion almost everywhere. If F is continuous, the graph of B is closed. In fact, taking a limit in the maximizing inequalities for $y_n \in B(x_n)$ shows that any limit y maximizes against the limit x. Compactness of the simplex then gives upper semicontinuity. One may extend the correspondence to the ambient Euclidean space by precomposing B with the nearest-point projection onto the simplex. The standard existence theorem for locally bounded upper semicontinuous differential inclusions with nonempty compact convex values supplies a local solution [11].

For a solution, $b(t) = \dot{x}(t) + x(t)$ belongs to the simplex almost everywhere. The integrating factor gives

$$x(t) = e^{-t}x(0) + \int_0^t e^{-(t-s)}b(s)\,ds.$$

The coefficients form a probability weighting. This keeps the path in the simplex and hence within the region where the original correspondence is defined. The velocity is bounded there, so the path extends to every forward time by continuation. Existence requires no choice of best response at a tie. There may be several paths from the same initial state. A point admits a constant path exactly when

$x \in B(x)$, the Nash condition. Convergence results for this inclusion require substantial further structure. For instance, Hofbauer and Sorin study continuous concave–convex zero-sum games and identify convergence to the saddle-point set [12].

7.2 Pairwise comparison

Suppose a reviser currently using j switches towards i at rate $\rho(F_i(x) - F_j(x))$, absorbing any common uniform sampling frequency into time. Let ρ be nonnegative and locally Lipschitz, strictly positive for positive arguments and zero for nonpositive arguments. The drift of x_i is the difference between incoming and outgoing flux:

$$W_i(x) = \sum_j x_j \rho(F_i(x) - F_j(x)) - x_i \sum_j \rho(F_j(x) - F_i(x)). \tag{7.1}$$

Each individual transfer contributes with opposite signs to the two relevant coordinates. Summing therefore gives zero, and the rate at an absent coordinate is nonnegative. If F is locally Lipschitz, the incoming-term variation argument of Chapter 5 shows that the resulting equation has a unique global forward trajectory from each simplex state.

Proposition 7.1 (Rest states under comparison)**.** *A state is a rest point of* (7.1) *if and only if it is Nash.*

Proof. Reindexing the sum of the outgoing terms gives

$$F(x) \cdot W(x) = \sum_{i,j} x_j \rho(F_i(x) - F_j(x))(F_i(x) - F_j(x)) \geq 0.$$

Every summand is nonnegative, and it is positive exactly when $x_j > 0$ and $F_i(x) > F_j(x)$. At a rest point no such pair exists. Thus every used strategy maximizes payoff, so $x \in B(x)$. Conversely, if $x \in B(x)$, every transition away from a used strategy has rate zero. Both flux terms vanish. □

This is the positive-correlation argument for pairwise comparison dynamics [9]. It also yields a strict increase of the potential along nonstationary paths when the payoff one-form is exact. In a related protocol, sample the candidate from the current population and use $[F_i - F_j]_+ = \max(F_i - F_j, 0)$ as the switch rate conditional on the candidate. The flux from j to i is then $x_j x_i [F_i - F_j]_+$. Combining it with its reverse flux and using $[a]_+ - [-a]_+ = a$ recovers (5.1). This latter sampling convention cannot create an absent type. The state-independent availability of candidates is the hypothesis that makes the preceding characterization include the boundary Nash inequalities.

7.3 Logit revision and regularization

Fix a noise parameter $\eta > 0$. The logit response is

$$B_i^{\eta}(x) = \frac{\exp(F_i(x)/\eta)}{\sum_j \exp(F_j(x)/\eta)}.$$

The smooth response map from fitness vectors to probabilities preserves common payoff baselines and is equivariant under strategy relabeling. For a bijection $S \to S'$ with permutation map P the latter statement is

$$\begin{array}{ccc} \mathbb{R}^S & \xrightarrow{\operatorname{softmax}_\eta} & \Delta(S) \\ \downarrow & & \downarrow \\ \mathbb{R}^{S'} & \xrightarrow{\operatorname{softmax}_\eta} & \Delta(S'). \end{array}$$

The two vertical maps apply P. The result follows because the normalization is a sum over the strategy set. A many-to-one map need not commute with the same response, since several copied entries contribute several terms to that sum.

In detail, let $q\colon S' \twoheadrightarrow S$ be a surjection and pull a payoff vector f back to $f \circ q$. The mass assigned to the fiber of i is proportional to $|q^{-1}(i)| \exp(f_i/\eta)$. The effective coarse payoff includes a shift of $\eta \log |q^{-1}(i)|$, which cancels from the response when the fiber sizes agree. More generally, using positive reference weights w_j on S' makes the mass of a fiber proportional to $\exp(f_i/\eta) \sum_{j:q(j)=i} w_j$. The weighted response therefore commutes with lumping for pulled-back fitness vectors when one pushes the reference weights forward. To extend this calculation to a semiconjugacy of the revision equations, the fine fitness must be constant within each fiber and must be computed from the coarse state by the same rule at every instant.

One can derive the logit rule by regularizing the response optimization. Let $H(y) = -\sum_i y_i \log y_i$, using the continuous extension $0 \log 0 = 0$. For fixed x, write $Z = \sum_j \exp(F_j(x)/\eta)$. Substituting the definition of B^{η} gives

$$y \cdot F(x) + \eta H(y) = \eta \log Z - \eta D(y \| B^{\eta}(x)).$$

The identity is valid also for y on a proper face because the second argument of entropy has full support. By the nonnegativity established in Chapter 6, the unique regularized maximizer is $B^{\eta}(x)$. Compare its objective value with that of a pure maximizing strategy, which has entropy zero. Since $H(y) \le \log |S|$,

$$0 \le \max_i F_i(x) - B^{\eta}(x) \cdot F(x) \le \eta \log |S|.$$

The entropy bound follows by computing relative entropy against the uniform distribution. In particular, each logit rest point is an $\eta \log |S|$-approximate

Nash state, in the sense that this is a bound on the gain from any unilateral population-level deviation.

The dynamics $\dot{x} = B^\eta(x) - x$ have a unique forward solution for locally Lipschitz F. The same integrating factor used for best response shows forward invariance. For a fixed η, continuity on the compact simplex gives a positive lower bound for every response coordinate. Thus the solution is interior for all $t > 0$, with each coordinate bounded uniformly away from zero after any positive time. Brouwer's fixed point theorem ensures a rest point for continuous F. Every rest point is interior and satisfies $x = B^\eta(x)$.

If $\eta_k \to 0$ through positive parameters and $x^{(k)} = B^{\eta_k}(x^{(k)})$ converges, its limit is Nash. For a strategy strictly suboptimal at the limit, continuity gives a payoff gap of at least some $\delta > 0$ compared with a maximizing strategy there, for all large k. Its probability in $x^{(k)}$ is at most $\exp(-\delta/\eta_k)$, hence is zero in the limit. Only maximizers can retain mass. This argument does not choose which Nash states arise along branches of fixed points.

The potential case gives a convergence result at a fixed noise level. Let V be a potential in the sense of Chapter 6, continuously differentiable on the simplex in its affine hull. On the interior the logit dynamics satisfy

$$\begin{aligned}\frac{d}{dt}(V + \eta H)(x) &= \eta \sum_i (B_i^\eta(x) - x_i)(\log B_i^\eta(x) - \log x_i) \\ &= \eta\big[D(B^\eta(x)\|x) + D(x\|B^\eta(x))\big].\end{aligned}$$

The first equality uses $F_i(x) = \eta \log B_i^\eta(x) + \eta \log \sum_j e^{F_j(x)/\eta}$ and the fact that the velocity sums to zero. Thus $V + \eta H$ strictly increases except at rest points. The interior bounds after positive time make the compactness argument of Chapter 6 applicable. If V is concave, then $V + \eta H$ is strictly concave on the simplex. Each interior logit rest point satisfies its constrained first-order condition, so there is only one. All trajectories converge to it. The existence, limit, and convergence statements use different hypotheses on fitness and on the response rule. More general perturbation and revision theories are developed by Sandholm [10].

CHAPTER 8

Interacting populations and spatial structure

The comparison neighborhood in Chapter 7 consisted of individuals drawn from the same population. An individual can instead interact frequently with nearby populations while drawing its behavioral replacement from its own population. An individual who moves changes the composition of both the source and destination. These mechanisms give different equations and lead to different forms of reduction. We study exact synchronized subspaces for a deterministic interaction model and then derive a quantitative synchronization criterion for migration.

8.1 Product simplices and assortative encounters

Let I be a finite set of populations and S_r a finite set of types for each $r \in I$. A state in $\mathcal{D}_I = \prod_{r \in I} \Delta(S_r)$ records conditional frequencies within the populations. The payoff to type $i \in S_r$ is $F_i^r(x)$, potentially a function of the full profile. Local normalization gives

$$\phi_r(x) = \sum_{i \in S_r} x_i^r F_i^r(x), \qquad \dot{x}_i^r = x_i^r \big(F_i^r(x) - \phi_r(x)\big).$$

Each population has its own conserved mass. Assuming a locally Lipschitz extension near the product, the proof of Chapter 5 applies simultaneously. The equation for each population's mass receives its own continuous coefficient $\phi_r(x)$, and each coordinate still has an integrating factor. Compactness of the product yields the global flow. A choice of one face per population gives an invariant product face. If every F^r depends solely on x^r, the flow is the product of the separate population flows. If payoffs couple populations, a product face can remain invariant while the flow within it is coupled.

Take a common type set S, a matrix A of row-player payoffs on S, and a row stochastic matrix $W = (W_{rs})$ on I. An encounter by a resident of r draws its partner's population from row r and then draws the type from that population. The vector of expected payoffs is

$$F_W^r(x) = A \sum_{s \in I} W_{rs} x^s.$$

Diagonal weights quantify within-population encounters. Directed weights allow asymmetric access to partners. The state equations apply the selection step within each destination of the payoff map, retaining separate local averages even when all populations see the same distribution of partners. At a stationary profile each resident type in each population has the same payoff. Comparison with absent types is required to promote this equality to a population Nash condition.

Here W is part of a deterministic model posed directly at the population level. For an individual-based graph process, replacing, for example, the conditional distribution of neighbors of type i by the unconditional population mixture is an additional modeling step. Ohtsuki and Nowak derive update-dependent replicator approximations for regular graphs using weak selection and pair approximation [13]. That derivation treats local correlations explicitly. The present equation follows directly from the specified partner sampling.

8.2 Exact reduction along equitable partitions

Let $\pi\colon I \twoheadrightarrow J$ partition I into blocks, and write $L\colon \mathcal{D}_J \longrightarrow \mathcal{D}_I$ for replication of block values, $(Lz)^r = z^{\pi(r)}$. The partition is equitable for W if the total weight from r to block b depends only on the block $a = \pi(r)$. Define $\bar{W}_{ab}$ to be that total. This gives a row stochastic matrix $\bar{W}$ and, if L also denotes the incidence matrix $L_{ra} = \mathbf{1}_{\pi(r)=a}$, the identity $WL = L\bar{W}$.

Proposition 8.1. *Write R_W and $R_{\bar{W}}$ for the replicator fields built with the same payoff matrix A. For an equitable partition, $R_W(Lz) = LR_{\bar{W}}(z)$. In particular the image of L is invariant, and the restriction of the flow to that image is conjugate to the reduced flow.*

Proof. For r in block a, the expanded profile has payoff $F^r_W(Lz) = A\sum_b \bar{W}_{ab} z^b = F^a_{\bar{W}}(z)$. Its resident frequencies are z^a, so its local mean payoff is the same as the block mean payoff. Substituting those two equalities makes all derivatives agree with the expanded block derivatives. A reduced solution lifts to a full solution, which is unique from its initial condition. □

For the flows Φ, this reduction is an intertwining relation.

$$\begin{array}{ccc} \mathcal{D}_J & \xrightarrow{\Phi^{\bar{W}}_t} & \mathcal{D}_J \\ {\scriptstyle L}\downarrow & & \downarrow{\scriptstyle L} \\ \mathcal{D}_I & \xrightarrow{\Phi^{W}_t} & \mathcal{D}_I \end{array}$$

If $\bar{W}L' = L'\widetilde{W}$ for a second equitable partition, then $W(LL') = (LL')\widetilde{W}$. Thus the embeddings and their intertwining equations compose. The one-block parti-

tion is equitable for any stochastic W. Its synchronous dynamics is the ordinary replicator equation for A.

For a general profile, a weighted block average retains less information than is needed to compute its derivative. For probability weights μ on a specified block, put $\bar{x}_i = \sum_r \mu_r x_i^r$ over that block and $h_i^r = F_i^r(x) - \phi_r(x)$. Then

$$\dot{\bar{x}}_i = \bar{x}_i \sum_r \mu_r h_i^r + \sum_r \mu_r (x_i^r - \bar{x}_i) \left(h_i^r - \sum_s \mu_s h_i^s \right).$$

The second term is a covariance. Its appearance precedes any further attempt to express the averages of relative payoffs through the coarse state. Equitability controls the aggregate encounter weights, while the embedding supplies the synchronized state on which that control implies an exact closed system.

8.3 A potential from detailed balance

Suppose W admits a positive probability vector μ satisfying $\mu_r W_{rs} = \mu_s W_{sr}$. If $A = A^{\mathsf{T}}$, define

$$P(x) = \frac{1}{2} \sum_{r,s} \mu_r W_{rs} (x^r)^{\mathsf{T}} A x^s.$$

Differentiation of the first factor gives $\frac{1}{2} \sum_r \mu_r \langle u^r, F_W^r(x) \rangle$. In the term from the second factor exchange r and s, then apply detailed balance and symmetry of A. This gives the other half of $dP_x(u) = \sum_r \mu_r \langle u^r, F_W^r(x) \rangle$.

On the relative interior of each product face, give the tangent space the weighted product of the Shahshahani metrics from Chapter 6. Its inner product is $\sum_r \mu_r \sum_{i: x_i^r > 0} u_i^r v_i^r / x_i^r$. For $v = R_W(x)$ the local means cancel against $\sum_i u_i^r = 0$. Therefore R_W is the gradient of P in this metric. The rate of increase of P equals the weighted sum of the within-population payoff variances, and vanishes at stationary profiles. This argument uses reciprocity of the encounter-weighted payoff couplings. A directed row stochastic matrix by itself supplies the reduction theorem but does not supply the symmetry used in the derivative of P.

8.4 Migration and exponential synchronization

Let $V = R_F$ be a common local replicator field with a Lipschitz constant C on $\Delta(S)$ in Euclidean norm. Consider the additional coupling

$$\dot{x}^r = V(x^r) + \kappa \left(\sum_s W_{rs} x^s - x^r \right), \qquad \kappa \geq 0.$$

This term diffuses relative frequencies between equal-normalization populations. It mixes a particular type across populations even when its local frequency is

initially zero. If the replicator has the ambient form above, the total mass vector m satisfies $\dot{m}_r = (1 - m_r)\phi_r + \kappa((Wm)_r - m_r)$ along any solution. The constant unit vector solves this linear inhomogeneous equation, with the coefficients taken along that solution. Uniqueness preserves each mass. The incoming linear terms have nonnegative off-diagonal coefficients for every type. Variation of constants, or local Picard iteration of the resulting positive integral equation for each type, preserves nonnegativity. Thus the model has a global flow on $\mathcal{D}_I$.

Proposition 8.2. *Let $|I| \geq 2$ and suppose W is irreducible and reversible with a positive stationary distribution μ. Let $\lambda > 0$ be the smallest positive eigenvalue of $I - W$ in the μ-weighted inner product. If $\kappa\lambda > C$, the squared distance to the synchronous subspace, in the product norm with weights μ, decays at least at rate $2(\kappa\lambda - C)$.*

Proof. For the mean $\bar{x} = \sum_r \mu_r x^r$ set $\delta^r = x^r - \bar{x}$ and $E = \frac{1}{2}\sum_r \mu_r \|\delta^r\|^2$. Since the deviations have weighted mean zero, the term from $\dot{\bar{x}}$ in $\dot{E}$ vanishes. The local contribution is $\sum_r \mu_r \langle \delta^r, V(x^r) - V(\bar{x})\rangle \leq 2CE$. The migration contribution equals $-\kappa$ times the Dirichlet form

$$\langle \delta, (I - W)\delta \rangle_\mu = \frac{1}{2}\sum_{r,s} \mu_r W_{rs} \|\delta^r - \delta^s\|^2.$$

Reversibility makes $I - W$ self-adjoint. Irreducibility forces every zero of the displayed form to be constant across populations. Applying the spectral gap on the mean-zero subspace bounds the form below by $2\lambda E$. The claim follows by integrating $\dot{E} \leq -2(\kappa\lambda - C)E$. □

Migration cancels from the mean equation by stationarity of μ. The mean therefore satisfies $\|\dot{\bar{x}} - V(\bar{x})\| \leq C\sqrt{2E}$. The estimate controls the deviation from the local law along the mean path, and the stability of any particular synchronous orbit is a further dynamical question. The partition theorem gives an exact invariant subspace. The migration estimate gives an attraction criterion for one such subspace. Both properties specify what a spatial reduction preserves and under which hypotheses it does so.

CHAPTER 9

Finite populations and Markov dynamics

Alexander starts the finite-population analysis from individual replacement and then asks what changes when interaction is local [1, Sec. 4]. To make this distinction precise, separate the population state, the law of an update, and the statistic that describes the next state. A stochastic update assigns a probability to each possible successor. Its support determines what can happen, and its probability weights determine the relative likelihoods. Both are needed to describe extinction, fixation, and stationary behavior.

9.1 Count states and fitness

Let S be a nonempty finite strategy set. Fix a population size $N \geq 2$. A labeled configuration is a function $z : \{1, \dots, N\} \to S$. Its count vector is $n(z)$, with $n_i(z)$ the number of individuals using strategy i. The set of count states and the associated frequencies are

$$C_N(S) = \left\{ n \in \mathbb{N}^S : \sum_{i \in S} n_i = N \right\}, \qquad x_i = \frac{n_i}{N}.$$

Count states are orbits of the action of label permutations on configurations. The induced action on an update law can be used to test whether the orbit statistic itself evolves by a Markov chain. Under exchangeable sampling, the law of the updated counts depends only on the current counts. On a graph, two configurations in the same orbit can place the same strategy at vertices with different neighborhoods. The Markov property for counts then requires a separate argument.

Suppose A_{ij} is the payoff to strategy i against j, and that an individual's opponent is sampled uniformly from the other $N - 1$ individuals. The expected payoff of a present type i is

$$u_i^N(n) = \sum_{j \in S} \frac{n_j - \mathbf{1}_{\{i=j\}}}{N - 1} A_{ij}.$$

If $n_i = 0$, there is no individual of type i whose payoff is needed. Give each currently present type a strictly positive reproductive weight $f_i(n)$. For instance a positive response $f_i(n) = \phi(u_i^N(n))$ implements a prescribed mapping from payoffs to reproductive probabilities. The choice is part of the model. Adding the same constant to all payoffs preserves many static comparisons. Adding the same constant to positive reproductive weights changes ratios f_i/f_j in general and hence can change the stochastic dynamics. A common positive multiplication of all weights at a state leaves their normalized ratios unchanged. The distinction parallels the several meanings of payoff equivalence in Chapter 4.

9.2 Birth, mutation, and replacement

The birth–death process updates one position at a time. It belongs to the overlapping-generation tradition of Moran's model [14]. Here the parent is chosen with probability proportional to its reproductive weight, the offspring type is sampled from a mutation kernel, and a uniformly chosen member of the current population is replaced. The parent is eligible for replacement in this convention. Let M_{ik} be the probability that a type i parent produces type k, with each row of M summing to one. Define

$$b_i(n) = \frac{n_i f_i(n)}{\sum_j n_j f_j(n)}, \qquad c_k(n) = \sum_i b_i(n) M_{ik}, \qquad d_j(n) = \frac{n_j}{N}.$$

For an absent type we set $n_i f_i(n) = 0$ by convention. The offspring distribution is $c(n)$. Let e_i denote the count vector with a unit at i and zero elsewhere. If $k \neq j$ and $n_j > 0$, the transition from n to $n + e_k - e_j$ has probability $c_k(n) d_j(n)$. The probability of remaining at n is $\sum_j c_j(n) d_j(n)$. These probabilities sum to one because all ordered pairs of offspring and removed types have been accounted for. A death rule depending on the parent or on its location replaces the product $c_k d_j$ by a joint kernel. In that case the joint law is part of the dynamics.

The update can also be written as a composition of kernels, retaining only the information required by the next stage. With $D_n = \{(k, j) \in S^2 : n_j > 0\}$, each arrow below is well defined on its displayed states:

$$\{n\} \xrightarrow{c \otimes d} D_n \xrightarrow{\text{replacement}} C_N(S).$$

The final map is deterministic. This factorization makes conservation of total population size explicit. It also gives a route to composition with other mechanisms, such as scheduling several local updates, by composing their kernels in time. The conditional mean of the frequency increment is

$$\mathbb{E}[x_{\text{next}} - x \mid n] = \frac{1}{N}(c(n) - x). \tag{9.1}$$

For this identity, simply average $e_k - e_j$ against the joint law. Independence supplies the stated c and d, while the mean-increment formula needs only their marginals. The second moments can depend on the joint law even when the drift agrees.

In the independent case, writing $\mathrm{diag}(c)$ for the diagonal matrix with entries c_k, the conditional covariance of the frequency increment is

$$\frac{1}{N^2}\big(\mathrm{diag}(c) - cc^{\mathsf{T}} + \mathrm{diag}(d) - dd^{\mathsf{T}}\big),$$

with c, d taken as columns in this formula. Each removed or added type is represented by a random coordinate vector. Independence makes their cross covariance zero, so the two covariance matrices add. The matrix annihilates the all-ones vector, expressing that a frequency increment is tangent to the affine simplex. This covariance and the conditional drift answer different questions about the same update kernel.

Proposition 9.1 (Fixation and recurrence)**.** *For the update above with strictly positive reproductive weights, suppose first that M is the identity. The chain reaches a monomorphic population almost surely, and each monomorphic population is absorbing. If M instead has strictly positive entries, the chain is irreducible and aperiodic. It has a unique stationary law and converges to that law from every initial distribution.*

Proof. When M is the identity, a currently absent type cannot be produced. In a monomorphic state all replacements preserve the type. Starting from any state containing i, choose a parent of type i and remove an individual of another type at each successive step until reaching the monomorphic state of type i. Each step has strictly positive conditional probability and the path has length at most $N - 1$. The state set is finite, so there is a uniform positive lower bound on the probability of absorption within $N - 1$ steps, considering one such path for each state. By the Markov property, the probability of avoiding absorption over successive blocks of $N - 1$ steps decays geometrically.

If M is strictly positive, $c_k(n) > 0$ for every k and every n. A move removing a present type and adding any specified type is therefore possible. Starting from n, repeatedly move a unit from a coordinate in surplus relative to a target n' to one in deficit. This reaches n' in at most N steps. Every state also has a self-loop because at least one d_j is positive and every c_j is positive. Thus the N-step transition matrix is strictly positive. A stationary law exists as a limit point of Cesàro averages. The N-step matrix has a common strictly positive minorant in every row. Removing it leaves the same residual mass in every row and gives a total-variation contraction by that residual mass, which is strictly less than one. This proves uniqueness and convergence for the full chain. □

9.3 Fixation probabilities retain the update law

When mutation is absent, let $h_i(n)$ be the probability of fixation at the monomorphic state of type i. Conditioning on the first update makes h_i harmonic away from the absorbing states. Its boundary values are one at fixation of i and zero at fixation of any other type. These conditions characterize it. Indeed, a harmonic function with these boundary values has constant conditional mean until absorption, and the absorption argument above gives almost sure convergence to the boundary values. Boundedness on the finite state space permits taking expectations in the limit.

Under neutral reproduction all present reproductive weights are equal, so $c(n) = x$ in the mutation-free process. Equation (9.1) gives harmonicity of n_i/N, and the boundary values give $h_i(n) = n_i/N$. For nonneutral reproduction these fixation probabilities depend on the transition law. The drift alone can be insufficient when different updates induce different hitting distributions, and a deterministic trajectory is a different observable from the absorption probability. The next chapter studies selection by rare mutation, in which fixation between departures from recurrent states supplies the effective transition information. Chapter 11 compares the frequency process with its deterministic limit on a fixed time interval.

CHAPTER 10

Rare mutation and stochastic stability

A small mutation probability changes the meaning of a long-run prediction. The mutation-free chain can have many recurrent classes. With positive mutation the chain can leave those classes, and its stationary law can weight them differently even when mutation tends to zero. The stationary limit first takes a long-time average at positive mutation and then decreases the mutation parameter. This order is intrinsic to the prediction.

10.1 A theorem for stationary weights

Let E be a nonempty finite state set and P an irreducible transition matrix on E, with rows indexed by the starting state. An in-tree rooted at i assigns one outgoing edge to each vertex other than i, such that repeated traversal of the edges brings any vertex to i. The root has no outgoing edge. Define

$$W_i(P) = \sum_{\tau \text{ rooted at } i} \prod_{(x,y)\in\tau} P(x,y).$$

The one-state case uses the weight one for the empty tree. Each weight is positive because the support of P has a path from every vertex to i. Choosing at each vertex an edge decreasing the shortest such path length gives a tree of positive weight.

Theorem 10.1 (Markov chain tree formula)**.** *The unique stationary distribution of P is $\pi(i) = W_i(P)/\sum_j W_j(P)$.*

Proof. It suffices to show that the weighted inflow to i from distinct states equals the weighted outflow from i:

$$\sum_{j\neq i} W_j(P)P(j,i) = W_i(P)\sum_{k\neq i} P(i,k).$$

A term on the left adds the edge $j \to i$ to a tree rooted at j. This produces a directed graph with one outgoing edge at every vertex and a unique directed cycle, of length at least two, containing i. A term on the right adds $i \to k$ to a tree rooted at i. It gives a graph of the same kind. From such a graph, deleting

the edge entering i on the cycle recovers the construction on the left. Deleting the outgoing edge at i recovers the construction on the right. The products are the same because they use all edges of the corresponding graph. The diagonal terms of stationarity cancel, so this bijection proves the displayed identity. Normalize and use uniqueness of the stationary law for a finite irreducible chain. □

This is the rooted-tree form of the Markov chain tree theorem [16]. It accounts for indirect paths between states. The stationary ranking can depend on several paths taken together and therefore need not coincide with a ranking of the individual edges leaving each state.

10.2 Leading powers and stable states

Suppose P^ϵ is irreducible for every sufficiently small $\epsilon > 0$. Assume for each off-diagonal edge that its weight is either identically zero for small ϵ, or

$$P^\epsilon(x, y) = c_{xy}\epsilon^{r_{xy}} + o(\epsilon^{r_{xy}}), \qquad c_{xy} > 0, \quad r_{xy} \geq 0.$$

Call r_{xy} the resistance. Omit identically zero edges and set

$$\gamma_i = \min_{\tau \text{ rooted at } i} \sum_{(x,y)\in\tau} r_{xy}.$$

For each root define $a_i > 0$ by summing $\prod_{(x,y)\in\tau} c_{xy}$ over its minimizers. Each tree product has exponent given by the sum of its edge resistances. The tree sum is finite and its coefficients at the smallest exponent are positive. Hence

$$W_i(P^\epsilon) = \epsilon^{\gamma_i}(a_i + o(1)).$$

Writing $\gamma_* = \min_i \gamma_i$, the stationary mass tends to zero for $\gamma_i > \gamma_*$ and tends to $a_i / \sum_{\gamma_j=\gamma_*} a_j$ for $\gamma_i = \gamma_*$. A state with positive limiting mass is called stochastically stable. Thus the minimum resistance determines the stable set, and the coefficients determine the limiting weights within it. The leading powers depend on the revision and mutation mechanism. An ordering of payoffs alone does not provide them.

This argument is a direct finite-state counterpart to resistance and stochastic potential in Young's theory of regular perturbed Markov processes [15]. Mutation mechanisms with weights decaying faster than all powers may require a different asymptotic scale. Aperiodicity is needed for general convergence from arbitrary initial laws along discrete time. The stationary weights in the tree theorem are defined for every irreducible kernel.

10.3 Eliminating transient excursions

The state set of a population chain can be large. An exact reduction separates states recurrent under no mutation from transient excursion states. Suppose $P^\epsilon \to P^0$ and write $E = K \amalg T$, where K is the set of recurrent states of P^0 and T is the set of transient states. Define successive visits to K by counting each ordinary time step spent in K, including consecutive time steps there. With matrix blocks indexed by the starting and ending sets, this trace kernel is

$$\widehat{P}^\epsilon = P^\epsilon_{KK} + P^\epsilon_{KT}(I - P^\epsilon_{TT})^{-1}P^\epsilon_{TK}. \tag{10.1}$$

If T is empty, only the first term is needed. The geometric series for the inverse sums incomplete excursions according to the number of steps that stay in T. For each $m \geq 0$ the term with m such steps follows the block sequence

$$K \xrightarrow{P^\epsilon_{KT}} T \xrightarrow{(P^\epsilon_{TT})^m} T \xrightarrow{P^\epsilon_{TK}} K.$$

Transience makes the series converge at $\epsilon = 0$, since the probability of staying in T decays geometrically in a finite chain. The inverse therefore exists and is continuous for sufficiently small ϵ.

The stationary equation on T is $\pi^\epsilon_T = \pi^\epsilon_K P^\epsilon_{KT} + \pi^\epsilon_T P^\epsilon_{TT}$. Solving it and substituting into the equation on K gives $\pi^\epsilon_K \widehat{P}^\epsilon = \pi^\epsilon_K$. Thus the trace stationary distribution is the normalized restriction of π^ϵ. The trace is irreducible for positive ϵ, since an original positive-probability path between visited states decomposes into consecutive trace steps. Moreover, $P^0_{KT} = 0$, because each recurrent class is closed at zero mutation. The same stationary identity then gives $\pi^\epsilon_T \to 0$. This permits computing stochastic stability using only the trace, when its transitions satisfy the requisite asymptotics. A trace records return events in the original chain. Its stationary probabilities describe fractions of visits, and the normalization accounts for the time omitted in excursions.

The error in dropping transient mass can be estimated before computing the trace stationary law. For a nonnegative matrix B, let $\|B\|_\infty$ be its maximum row sum. The stationary equation gives

$$\pi^\epsilon(T) \leq \|P^\epsilon_{KT}\|_\infty \left\|(I - P^\epsilon_{TT})^{-1}\right\|_\infty.$$

The second factor bounds the expected duration of an excursion from any state in T, and remains bounded as $\epsilon \to 0$ at fixed population size. For example, when the probability of leaving K in one update is uniformly $O(\epsilon)$, the total stationary time outside K is $O(\epsilon)$.

10.4 Rare steps and fixation

If the recurrent classes of P^0 are singleton absorbing states and $P^\epsilon = P^0 + \epsilon R + o(\epsilon)$ entrywise, formula (10.1) gives

$$\widehat{P}^\epsilon = I + \epsilon\big(R_{KK} + R_{KT}(I - P^0_{TT})^{-1}P^0_{TK}\big) + o(\epsilon).$$

The off-diagonal first-order terms are direct changes between absorbing states and changes that enter T once at order ϵ, then reach an absorbing state according to P^0. The latter factor contains the unperturbed fixation probabilities. If several rare steps are needed, first order alone cannot recover all resistances. More generally the trace need not lump several recurrent states into a single class at positive mutation. One can either keep all states of K, check an exact lumpability criterion as in Chapter 12, or invoke a separate asymptotic class reduction with its hypotheses.

Call the first-order coefficient G. Nonnegativity of off-diagonal trace probabilities and normalization of their rows imply that G is a continuous-time generator on K. If it is irreducible, the normalized trace stationary laws converge to its unique stationary law. Indeed, divide their stationary equations by ϵ after subtracting the identity term. Every accumulation point annihilates G. Irreducibility makes that accumulation point unique, and compactness gives convergence. This argument identifies the scope of the first-order approximation.

These conclusions address a fixed population size. Chapter 11 treats a different limiting operation, in which population size increases over a bounded time interval. Neither proof interchanges large population, long time, and vanishing mutation. Such an interchange requires additional bounds, because absorption and mixing time can themselves depend strongly on population size and on mutation.

CHAPTER 11

Mean-field limits and interaction structure

A deterministic evolution can describe the leading motion of a stochastic population over a bounded time interval. The connection is a convergence statement about processes, with a specified scale for the frequency increments and the update rate. It requires bounds on the microscopic drift and on fluctuations. Kurtz gave a general framework for the approximation of differential equations by jump processes [17]. Benaïm and Weibull studied the approximation and time scales for evolutionary games [18]. The result below isolates assumptions that can be checked by evaluating the local update law.

11.1 A generator on the frequency simplex

Let $D \subseteq \mathbb{R}^d$ be a compact convex set and let $D_N \subseteq D$ be a finite state set. The simplex lattice $C_N(S)/N$ is the main application. Fix a finite set $\mathcal{V}$ of possible jump directions. Write the generator as

$$\mathcal{L}_N \varphi(x) = N \sum_{v \in \mathcal{V}} \beta_v^N(x)[\varphi(x + v/N) - \varphi(x)].$$

Each intensity β_v^N is nonnegative and is zero whenever the indicated transition is infeasible. Terms of zero intensity are omitted. Suppose

$$B = \sup_{N, x \in D_N} \sum_v \beta_v^N(x) \|v\|^2 < \infty, \qquad F_N(x) = \sum_v \beta_v^N(x) v.$$

The vector $F_N(x)$ is the instantaneous conditional drift. The rate of quadratic variation from jumps is of order N^{-1} since each jump squares the factor N^{-1} and the generator contributes N. If F_N converges uniformly on its state sets to a field F, Taylor's formula gives $\mathcal{L}_N \varphi \to D\varphi[F]$ for a smooth φ with bounded second derivative on a neighborhood of D. The Taylor remainder is at most $B \sup \|D^2 \varphi\|/(2N)$. The convergence of generators on such functions identifies the candidate equation.

Theorem 11.1 (Bounded-time approximation)**.** *Assume F is Lipschitz on D with constant L, and set $\delta_N = \sup_{x \in D_N} \|F_N(x) - F(x)\|$. Suppose $\delta_N \to 0$ and*

$X_N(0) \to x_0$ in probability. If the solution of $\dot{x} = F(x)$ from x_0 remains in D on $[0,T]$, then

$$\sup_{0\le t\le T} \|X_N(t) - x(t)\| \longrightarrow 0 \quad \textit{in probability.}$$

Proof. The coordinate martingales of the generator give the vector decomposition

$$X_N(t) = X_N(0) + \int_0^t F_N(X_N(s))\,ds + M_N(t).$$

For the squared Euclidean norm of the martingale, the compensator of the sum of squared jumps contributes at most BT/N in expectation up to time T. In detail, the jump v/N contributes $\|v\|^2/N^2$ at rate $N\beta_v^N(X_N(s))$. Their integrated sum bounds $\mathbb{E}\|M_N(T)\|^2$. Thus Doob's quadratic inequality, applied to the submartingale $\|M_N(t)\|$, gives

$$\mathbb{E} \sup_{0\le t\le T} \|M_N(t)\|^2 \le 4BT/N.$$

Subtracting the deterministic integral equation and using the Lipschitz bound yields, by Grönwall's inequality,

$$\sup_{0\le t\le T} \|X_N(t) - x(t)\| \le e^{LT}\left(\|X_N(0) - x_0\| + T\delta_N + \sup_{0\le t\le T}\|M_N(t)\|\right).$$

Every term in parentheses tends to zero in probability. □

The estimate is quantitative once drift and initial-state errors are bounded. The factor e^{LT} also locates the difficulty of extending the comparison to long time. The theorem controls complete paths on each fixed interval. It does not compare stationary distributions, which involve behavior on a time scale that can increase with N.

For deterministic initial states with error $a_N = \|X_N(0) - x_0\|$, choose any $\eta > e^{LT}(a_N + T\delta_N)$. The proof gives

$$\mathbb{P}\left(\sup_{t\le T}\|X_N(t) - x(t)\| > \eta\right) \le \frac{4BT}{N\left(e^{-LT}\eta - a_N - T\delta_N\right)^2}.$$

The expression separates the initial discrepancy and the systematic drift error from the probability of a fluctuation. Different microscopic kernels can share the same limit F while differing in B or in the conditional covariance from Chapter 9.

For an update chain with transition matrix P_N, an independent Poisson clock of rate N gives the generator $N(P_N - I)$. The chain is evaluated at the number of clock arrivals by time t, so a unit of macroscopic time corresponds on average to N updates. A self-loop consumes a clock arrival and produces a zero displacement. This convention makes the drift of the continuous-time chain N times the mean increment in one discrete update. It also preserves stationary laws because $\pi N(P_N - I) = 0$ is equivalent to $\pi P_N = \pi$.

11.2 Recovering deterministic selection

Use the birth–death update of Chapter 9 at rate N. Its mean field is $c(n) - x$ by Equation (9.1). Suppose the weights of present types converge uniformly to bounded strictly positive functions $f_i(x)$, with a common positive lower bound. The limiting offspring fractions and field are

$$c_k(x) = \sum_i \frac{x_i f_i(x)}{\overline{f}(x)} M_{ik}, \qquad F_k(x) = c_k(x) - x_k, \qquad \overline{f}(x) = \sum_i x_i f_i(x).$$

If the weights are Lipschitz, the field is Lipschitz. Its components sum to zero, and $c_k(x) \geq 0$ at a face with $x_k = 0$. A direct form of the integral equation is

$$x_k(t) = e^{-t} x_k(0) + \int_0^t e^{-(t-s)} c_k(x(s))\, ds.$$

This exhibits the role of the offspring distribution. Invariance of the simplex also follows by constructing the solution locally as the fixed point of this integral map on paths in the simplex, where it is a contraction on short intervals and maps such paths to such paths. Short intervals can be concatenated because the simplex is compact and the constants can be chosen uniformly. With no mutation, $F_k(x) = x_k(f_k(x) - \overline{f}(x))/\overline{f}(x)$. Because $\overline{f}(x) > 0$, this has the replicator trajectories after a positive change of time. The rates with respect to calendar time differ by this common factor. With mutation, normalization and transmission are combined in c and need not give a particular additive mutation convention on the replicator equation.

The distinction between finite and limiting payoffs can be bounded explicitly. For $n_i > 0$ the payoff in Chapter 9 satisfies

$$u_i^N(n) - (Ax)_i = \frac{(Ax)_i - A_{ii}}{N - 1}.$$

If $f_i(n) = \phi(u_i^N(n))$ with a Lipschitz positive map ϕ on the interval covered by the payoffs, the difference from the weight $\phi((Ax)_i)$ is $O(N^{-1})$ uniformly over present types. The absent types contribute zero to the birth distribution. Positivity on the compact payoff interval gives a uniform lower bound for the normalization, so the corresponding drift approximation has order N^{-1}. The martingale estimate controls the stochastic error separately at order $N^{-1/2}$ in root mean square over fixed times.

11.3 When counts close over a graph

Suppose the population is an N-vertex weighted graph. For an oriented pair of vertices u, v, let $\lambda_{uv}^N \geq 0$ be a sampling weight with $\sum_{u,v} \lambda_{uv}^N = 1$. A type i at u

replaces a type j at v at rate $N\lambda^N_{uv}\varphi(i,j)$, for fixed nonnegative bounded biases φ. The relevant statistic is the oriented edge frequency

$$\theta^N_{ij}(z) = \sum_{u,v} \lambda^N_{uv} \mathbf{1}_{\{z(u)=i,\ z(v)=j\}}.$$

The density drift on a labeled state z is $\sum_{i,j} \theta^N_{ij}(z)\varphi(i,j)(e_i - e_j)$, where e_i is the coordinate unit at i. For a general graph the edge frequencies can vary when the vertex frequencies are fixed. A closed equation for vertex frequencies requires additional information about their evolution.

An explicit sufficient bound exists for regular weighted sampling. Suppose $\lambda^N_{uv} = W_{uv}/N$, with W a nonnegative matrix with row sums one, and write J_N for the matrix with all entries $1/N$. If $\mathbf{1}_i$ is the vertex indicator of type i, then

$$\left|\theta^N_{ij}(z) - x_i x_j\right| = \frac{1}{N}\left|\mathbf{1}_i^{\mathsf{T}}(W - J_N)\mathbf{1}_j\right| \le \|W - J_N\|_{2\to 2}\sqrt{x_i x_j}.$$

Thus $\|W - J_N\|_{2\to 2} \to 0$ is sufficient for uniform closure of the drift. The proof of Theorem 11.1 still applies to the projected coordinate processes, using the generator on labeled configurations to obtain their compensators. Uniform drift approximation substitutes for the exact Markov property of the projection.

$$\text{labeled state} \xrightarrow{\text{observe}} (x, \theta^N) \xrightarrow{\text{retain counts}} x$$

With uniform sampling of distinct ordered vertices, $W = (NJ_N - I)/(N-1)$. Then $W - J_N = (J_N - I)/(N-1)$ and its norm is $1/(N-1)$ for $N \ge 2$. For any graph family satisfying the norm bound, the limiting field is

$$F(x) = \sum_{i,j} x_i x_j \varphi(i,j)(e_i - e_j).$$

It is Lipschitz on the simplex. Its coordinates sum to zero and each $F_k(x)$ has the factor x_k, so the simplex is invariant. A density path on labeled configurations therefore converges to this solution under the initial-state hypothesis of Theorem 11.1. The omitted edge data may persist on sparse or strongly clustered graphs. One can retain additional observables or prove a separate estimate of their effect. Chapter 12 gives the criterion for exact Markov closure at each population size.

CHAPTER 12

Stochastic categories and population aggregation

12.1 Kernels as morphisms

A state description is useful only when its dynamics can be computed at the level of that description. For the exchangeable population rule in Chapter 9, counts describe a Markov chain exactly. In the network rule of Chapter 11, the count drift generally depends on edge frequencies. The category of stochastic kernels gives a common language for these two situations.

Let **FinStoch** have finite sets as objects and stochastic kernels as morphisms. A kernel $K : X \rightsquigarrow Y$ assigns to each $x \in X$ a probability $K(x, y)$ on Y. The composite of K and $H : Y \rightsquigarrow Z$ has entries

$$(KH)(x, z) = \sum_{y \in Y} K(x, y)H(y, z).$$

Matrix products follow the order in which updates are applied, since laws are row vectors. Identity kernels have entries $\mathbf{1}_{\{x=y\}}$. Matrix associativity is the rule for summing over consecutive unobserved states.

The monoidal product $X \times Y$ models joint state spaces, with product kernels given by

$$(K \otimes H)\big((x, x'), (y, y')\big) = K(x, y)H(x', y').$$

This product applies independent random draws to the two coordinates given their current states. The singleton is monoidal unit and terminal object. Every state space has a deterministic copy operation $\Delta_X(x) = (x, x)$. A kernel K preserves copy when $K\Delta_Y = \Delta_X(K \otimes K)$. On the left, a sample from $K(x, -)$ is copied. On the right, independent samples are drawn from the same row. Comparing off-diagonal probabilities shows that for distinct y, z one must have $K(x, y)K(x, z) = 0$. Thus each row is a point mass, and K comes from a function. Every function preserves copy. This is the deterministic subcategory of **FinStoch**, an instance of the distinction between general and deterministic morphisms in a Markov category [20].

12.2 Maps of dynamics

A discrete-time system is a pair (X, P) with $P : X \rightsquigarrow X$. A stochastic map of dynamics from (X, P) to (Y, Q) is a kernel $K : X \rightsquigarrow Y$ with $PK = KQ$. The equation says that updating the source and then observing it gives the same law as observing it and applying the target update.

$$\begin{array}{ccc} X & \xrightarrow{P} & X \\ {\scriptstyle K}\downarrow & & \downarrow{\scriptstyle K} \\ Y & \xrightarrow{Q} & Y \end{array}$$

Iteration gives $P^n K = KQ^n$ for every n. If $QH = HR$ as well, then $P(KH) = (KH)R$, so these maps compose. Equivalently, a system is a functor $B\mathbb{N} \to \mathbf{FinStoch}$, where $B\mathbb{N}$ has one object and the monoid of nonnegative time steps as its endomorphisms. A stochastic map of dynamics is a natural transformation. The categorical equation for all times is generated by its component at time one.

If $\mu P = \mu$, then $(\mu K)Q = \mu PK = \mu K$, so the map transports stationary laws. Independent dynamics can also be composed in parallel. For two maps satisfying $P_i K_i = K_i Q_i$, the kernel $K_1 \otimes K_2$ intertwines $P_1 \otimes P_2$ and $Q_1 \otimes Q_2$. Here the tensor products of the transitions represent synchronous updates with conditional independence.

12.3 Strong lumpability

Let $q : X \to Y$ be a surjection recording a deterministic statistic, and write $J(x, y) = \mathbf{1}_{\{q(x)=y\}}$. Its values identify blocks $B_y = q^{-1}(y)$. Strong lumpability requires that the observed process $q(X_n)$, for a P-chain X_n, have a fixed Markov update on Y for every initial law on X. This notion has a block characterization [19, §6.3].

Theorem 12.1 (Strong lumpability)**.** *For a kernel P on X, the following are equivalent.*

1. *There is a single kernel Q governing the process $q(X_n)$ for every initial law.*
2. *There is a kernel Q on Y satisfying $PJ = JQ$.*
3. *For every y, y' the sum $\sum_{x' \in B_{y'}} P(x, x')$ is constant for $x \in B_y$.*

When these conditions hold, Q is unique and $Q(y, y')$ is the common value in the third condition.

Proof. For $x \in B_y$, the (x, y')-entries of PJ and JQ are the block sum and $Q(y, y')$, respectively. This proves equivalence of the last two conditions because

the sums give nonnegative rows of total mass one. Given the original process history, the conditional probability that $q(X_{n+1}) = y'$ is the block sum for X_n. Under the third condition it is $Q(q(X_n), y')$. Conditioning on the observed history leaves that value unchanged, proving the first condition. Conversely, initialize at a point x and use the proposed update Q at time one. Its entry from $q(x)$ to y' must equal the block sum from x, hence the block sum agrees for all points in the block. □

A Markov description obtained for one specified initial law can depend on the microscopic conditional distributions within the blocks. This weaker statement does not imply an intertwining for deterministic J. For strong lumpability the intertwining suffices to give both the distributional identity $\mu P^n J = (\mu J)Q^n$ and the assertion about the full observed process in the theorem.

The quotient also has a factorization property for stochastic maps of dynamics. Suppose $H : (X, P) \to (Z, R)$ intertwines the updates and its rows are constant on each q-fiber. Then the unique kernel $\overline{H}$ defined by those common rows satisfies $H = J\overline{H}$. From $PH = HR$ and $PJ = JQ$ one gets $JQ\overline{H} = J\overline{H}R$. Surjectivity of q lets one cancel J since left multiplication by J simply repeats rows. Thus $Q\overline{H} = \overline{H}R$. Every observation of X determined by the quotient statistic, and compatible with its own time update, factors through the quotient as a map of dynamics.

Proposition 12.2 (Equivariance gives a quotient)**.** *If a group G acts on X and $P(gx, gx') = P(x, x')$ for every g, the orbit map $X \to X/G$ is strongly lumpable.*

Proof. Every target orbit B is preserved bijectively by g, so

$$\sum_{x' \in B} P(gx, x') = \sum_{u \in B} P(gx, gu) = \sum_{u \in B} P(x, u).$$

The sums therefore depend only on the source orbit. □

For site-independent replacement and permutation-invariant payoffs, site permutations act equivariantly on the labeled chain. Their orbits are precisely the configurations with a given strategy-count vector, recovering the population quotient in Chapter 9. In spatial dynamics the relevant permutation group is typically the automorphism group of the weighted interaction graph. Its orbits may record information beyond counts, such as whether a pair of strategy classes frequently meet across graph edges.

12.4 Generators and approximate quotients

A continuous-time finite-state chain has generator L_X with nonnegative off-diagonal entries and zero row sums. A deterministic quotient intertwines the

semigroups at every time if and only if $L_X J = J L_Y$. Differentiation at zero gives necessity. For sufficiency the generator equation gives $L_X^k J = J L_Y^k$ by induction, and the result follows from the exponential series. Thus a quotient generator exists when each generator block sum is constant on a source block. For off-diagonal blocks the contributing entries are nonnegative. Row sums of the quotient generator are zero by those of L_X. This condition can be checked on jump intensities directly, including when the total holding rate varies among states of a block.

When block sums almost agree, an error estimate controls the effect on a fixed observation horizon. For a proposed kernel Q on Y, assume

$$\sup_{x \in X} \left\| P(x, -)J - Q(q(x), -) \right\|_{\mathrm{TV}} \leq \varepsilon,$$

where total variation is one half the ℓ^1 norm. Writing $E = PJ - JQ$, an exact telescoping identity gives

$$\mu P^n J - \mu J Q^n = \sum_{j=0}^{n-1} \mu P^{n-1-j} E Q^j.$$

The row vector before E is a probability, so the norm of its product with E is at most ε by convexity. Propagation under Q^j is a contraction in total variation, giving a bound of $n\varepsilon$ for every starting law μ. This estimate concerns the one-time observed laws and is uniform on a specified range of time steps. A stationary error estimate requires more information, such as a quantitative contraction for the quotient dynamics. In the limit theorem of Chapter 11, the target is a deterministic flow and the uniform drift error is what enters the finite-time comparison.

CHAPTER 13

Open games and the continuation interface

Classical evolutionary game theory starts from payoffs on an entire population interaction. A component description needs an interface for an interaction whose surroundings are still unspecified. Open games supply that interface in the form of an input and a continuation. The continuation tells the component how each possible output would be evaluated by its surroundings. This makes off-equilibrium choices available for composition [21].

13.1 Lenses as the deterministic interface

Use (X, S) for an interface with forward type X and backward type S. A lens from (X, S) to (Y, R) consists of two maps

$$v : X \to Y, \qquad u : X \times R \to S.$$

Its forward action sends a state x to $v(x)$. Its backward action translates a return value r using the original state x. In game semantics a return value may contain utility information. The lens axioms here are the composition and identity laws for these pairs.

For a second lens $(v', u') : (Y, R) \to (Z, Q)$, the composite has forward map $v'v$ and backward map

$$(x, q) \longmapsto u\big(x, u'(v(x), q)\big).$$

For three lenses, both parenthesizations substitute the forward values in the same order and then the backward values in reverse order. The lens $(\mathrm{id}_X, (x, s) \mapsto s)$ is a two-sided identity. This gives an ordinary category. The product of two lenses uses cartesian products on both types and applies the two maps separately.

A continuation on the target is an arbitrary function $k : Y \to R$. An input at the source is an element $x \in X$. A lens sends an input forward by v and transports a continuation backward by

$$k \longmapsto \big(x \longmapsto u(x, k(v(x)))\big).$$

The backward transport reverses composition by direct substitution. Keeping k as a function matters. A single value $k(v(x))$ gives a return only at the chosen output. A component deciding whether to deviate may need the value at a different output.

13.2 Strategies and contextual responses

An open game $\mathcal{G} : (X, S) \to (Y, R)$ consists of a strategy set $\Sigma_{\mathcal{G}}$, a play map $P_{\mathcal{G}}$, a coplay map $C_{\mathcal{G}}$, and best-response relations:

$$\begin{aligned} P_{\mathcal{G}} &: \Sigma_{\mathcal{G}} \times X \to Y, \\ C_{\mathcal{G}} &: \Sigma_{\mathcal{G}} \times X \times R \to S, \\ B_{\mathcal{G}}(x, k) &\subseteq \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{G}}. \end{aligned}$$

Here $k : Y \to R$. The strategies are chosen for the whole component. Fixing a strategy σ gives a lens with maps $P_{\mathcal{G}}(\sigma, -)$ and $C_{\mathcal{G}}(\sigma, -, -)$. A pair $(\sigma, \sigma') \in B_{\mathcal{G}}(x, k)$ means that σ' is an admissible response to the incumbent σ in this context. The equilibrium set consists of the strategies σ with $(\sigma, \sigma) \in B_{\mathcal{G}}(x, k)$. An open game in this generality specifies responses as data. Its response relation can be constructed from an optimization problem, a restricted revision protocol, or another stated selection rule [21].

For a decision component with observed type X and choice type Y, take $\Sigma = Y^X$ and $P(\sigma, x) = \sigma(x)$. With real return type and backward source type 1, the standard maximizing relation allows (σ, σ') if

$$k(\sigma'(x)) \;\geq\; k(y) \qquad \text{for every } y \in Y.$$

The map C is the unique map to 1. Existence of a response requires attainment of a maximum. At a given observed x this relation constrains the strategy only at x. A stronger requirement at all possible observations needs another equilibrium test, as in subgame analysis. The distinction is visible in the typing of $B(x, k)$.

13.3 Sequential composition

Let $\mathcal{H} : (Y, R) \to (Z, Q)$ follow $\mathcal{G}$, with outer continuation $k : Z \to Q$. An incumbent (σ, τ) determines the intermediate forward value $y = P_{\mathcal{G}}(\sigma, x)$. The downstream strategy determines the continuation seen upstream:

$$k_\tau(y') = C_{\mathcal{H}}(\tau, y', k(P_{\mathcal{H}}(\tau, y'))), \qquad y' \in Y.$$

The upper path carries states forward. The lower arrow carries the continuation back across the downstream component.

$$X \xrightarrow{P_{\mathcal{G}}(\sigma,-)} Y \xrightarrow{P_{\mathcal{H}}(\tau,-)} Z$$

$$R^Y \xleftarrow{k \mapsto k_\tau} Q^Z$$

Here Q^Z and R^Y denote sets of continuations. A composite strategy is a pair. Its play and coplay are the composition of the fixed-strategy lenses, and its response relation is

$$\begin{aligned}&\big((\sigma,\tau),(\sigma',\tau')\big) \in B_{\mathcal{H}\circ\mathcal{G}}(x,k)\\ &\quad\iff (\sigma,\sigma') \in B_{\mathcal{G}}(x,k_\tau) \quad \text{and} \quad (\tau,\tau') \in B_{\mathcal{H}}(y,k).\end{aligned}$$

The two tests use the incumbent context. The second component sees the output generated by σ. The first component evaluates alternatives using τ as a fixed strategy, including its play and coplay at changed inputs. This is a prescription for testing responses with the other components' strategies held fixed.

Proposition 13.1 (Associativity of the response construction)**.** *Sequential composition is associative under the canonical reassociation of strategy products. Singleton-strategy identity lenses with the sole identity response give identities.*

Proof. Write $\mathcal{J} : (Z, Q) \to (W, T)$ for the third component and υ for its incumbent. For an outer continuation $\ell : W \to T$, put

$$k_\upsilon(z') = C_{\mathcal{J}}\big(\upsilon, z', \ell(P_{\mathcal{J}}(\upsilon, z'))\big).$$

With $y = P_{\mathcal{G}}(\sigma, x)$ and $z = P_{\mathcal{H}}(\tau, y)$, either parenthesization tests the three responses in the contexts

$$\begin{aligned}\mathcal{G} &: \quad \big(x,\ y' \mapsto C_{\mathcal{H}}(\tau, y', k_\upsilon(P_{\mathcal{H}}(\tau, y')))\big),\\ \mathcal{H} &: \quad (y, k_\upsilon),\\ \mathcal{J} &: \quad (z, \ell).\end{aligned}$$

For the first parenthesization these contexts result from first transporting through $\mathcal{J}$ and then through $\mathcal{H}$. For the second, the coplay of the downstream lens composite expands to the same substitution. The play and coplay agree by lens associativity. Inserting an identity leaves the transported state and continuation unchanged and adds the sole relation on a singleton. □

Taking strategy sets up to reindexing by these coherent bijections gives an ordinary category. Keeping actual strategy sets makes the reassociation visible as a comparison between games. The next chapter also keeps reindexings that need not be invertible.

13.4 A compositional equilibrium test

The diagonal of the response formula gives the equilibrium criterion

$$\sigma \in E_{\mathcal{G}}(x, k_\tau), \qquad \tau \in E_{\mathcal{H}}(P_{\mathcal{G}}(\sigma, x), k),$$

where $E_{\mathcal{G}}$ denotes the contextual equilibrium set. This is a system of mutually dependent tests. Retaining only $E_{\mathcal{G}}$ is enough for this particular diagonal calculation when all contexts and the fixed-strategy lenses are retained. Off-diagonal responses are needed to study transformations that preserve response possibilities or to introduce an update rule using those possibilities. For evolutionary models the selected continuation can include frequency-dependent payoffs. One must still specify how the current population determines that continuation, which strategy information is available, and how admissible responses change frequencies. The interface supplies a way to keep these choices explicit.

For instance, an adapter can expose the population state as an input and compute a payoff vector as its output. A downstream selection component then evaluates each strategy against that vector. Compositional optimization tests whether a strategy passes the resulting response relation. A dynamical model adds a state transition or vector field on populations, whose state change is an additional map. These two observations about the system can then be checked for agreement under the hypotheses of Chapter 6.

13.5 Deterministic adapters

The response formula gives a useful rule for changing the language in which an interface is described. Suppose $L : (X', S') \to (X, S)$ and $M : (Y, R) \to (Y', R')$ are lenses. Regard each as a singleton-strategy game with its sole response. For a context (x', k') of $M \circ \mathcal{G} \circ L$, the response test of the middle game has context

$$\left(v_L(x'),\ y \longmapsto u_M(y, k'(v_M(y)))\right).$$

This follows by first applying the lens substitution to the return from M, then carrying the input forward through L. The singleton comparisons impose no further response condition, and the same formula therefore computes equilibria of the adapted game. If a downstream adapter translates every return by the same amount for the relevant decision context, a real maximizing response test is unchanged. More generally a strictly increasing transformation of every return in that test preserves the ordering of alternatives. The key hypothesis here is a common transformation of the returns being compared. A translation depending on the output may alter their order and must be included in the continuation itself.

CHAPTER 14

Parallel games and higher comparisons

Sequence describes a change of context through an intermediate interface. Parallel composition retains two interfaces and asks how the environment evaluates their joint output. A joint continuation can couple the decisions even when their play maps factor.

14.1 Parallel composition

For $\mathcal{G} : (X, S) \to (Y, R)$ and $\mathcal{H} : (U, T) \to (V, Q)$, their product has object types $(X \times U, S \times T)$ and $(Y \times V, R \times Q)$, and strategy set $\Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$. Its play and coplay are the products of the fixed-strategy lenses. Given a context with forward state (x, u) and continuation $k : Y \times V \to R \times Q$, an incumbent (σ, τ) produces $y_* = P_{\mathcal{G}}(\sigma, x)$ and $v_* = P_{\mathcal{H}}(\tau, u)$. Define the two continuations seen by the components as

$$k_1(y) = \pi_R k(y, v_*), \qquad k_2(v) = \pi_Q k(y_*, v).$$

The parallel best-response relation requires $(\sigma, \sigma') \in B_{\mathcal{G}}(x, k_1)$ and $(\tau, \tau') \in B_{\mathcal{H}}(u, k_2)$. A trial change in one component thus sees the same incumbent output from the other. In particular, even a product of games whose input states are independent can have coupled equilibrium conditions when the continuation is joint. If $k(y, v) = (r(y), q(v))$, the continuations are r and q independently of σ, τ, so the equilibrium sets multiply.

Proposition 14.1 (Interchange)**.** *Parallel composition commutes with sequential composition up to the canonical exchange of strategy factors. Its associativity, unit and symmetry are given by the corresponding product bijections.*

Proof. Consider two chains, $\mathcal{G}$ followed by $\mathcal{H}$ and $\widetilde{\mathcal{G}}$ followed by $\widetilde{\mathcal{H}}$, with incumbent strategies σ, τ and $\widetilde{\sigma}, \widetilde{\tau}$, respectively. Let $y_*, \widetilde{y}_*$ be the intermediate outputs, $z_*, \widetilde{z}_*$ the final outputs, and k the joint outer continuation. In either order of assembly the test at $\mathcal{H}$ uses y_* and $z \mapsto \pi_1 k(z, \widetilde{z}_*)$. The test at $\mathcal{G}$ uses the continuation

$$y \longmapsto C_{\mathcal{H}}(\tau, y, \pi_1 k(P_{\mathcal{H}}(\tau, y), \widetilde{z}_*)).$$

Indeed, composing within the first chain propagates the first outer continuation through $\mathcal{H}$. Composing across the parallel pair first propagates k through the product lens, and then freezes the second intermediate value at $\widetilde{y}_*$. Its downstream output is precisely $\widetilde{z}_*$. The formulas for the other chain follow by using π_2 and holding z_* fixed. Thus all four response tests agree. The play and coplay maps agree because composition of product lenses is componentwise.

For any number of parallel factors the continuation for one factor fixes all other incumbent outputs and projects to the return for that factor. This gives the same formula for either parenthesization and after permutation of factors. The singleton unit introduces neither another output choice nor another return choice. Coherence follows from the corresponding equalities of tuples. □

14.2 Strategy maps and two-dimensional structure

One can preserve more information by keeping morphisms between games. For games $\mathcal{G}, \mathcal{G}'$ on the same interface, take a map $\phi : \Sigma_{\mathcal{G}} \to \Sigma_{\mathcal{G}'}$ that preserves play and coplay at every state and return. Require also that

$$(\sigma, \eta) \in B_{\mathcal{G}}(x, k) \implies (\phi(\sigma), \phi(\eta)) \in B_{\mathcal{G}'}(x, k).$$

This gives a category for each pair of interfaces. Composition of strategy maps preserves the relation by transitivity of implication, and identity maps preserve all the specified data. Products of these maps preserve the sequential and parallel response formulas. For the sequential case, play preservation gives the same intermediate incumbent state and coplay preservation gives the same transported continuation. One can then apply response preservation separately in the two games. The parallel argument uses preservation of the two incumbent outputs. This yields a bicategorical description after retaining the canonical product reassociations as invertible cells. For a more general morphism allowing lenses between the boundary interfaces, see the double category of open games in [22].

For sequential composition, the diagram below holds at every upstream state x. Set $E_x(\sigma, \tau) = P_{\mathcal{H}}(\tau, P_{\mathcal{G}}(\sigma, x))$, where Z is the final forward output space. Primes denote the corresponding maps for the target games.

$$\begin{array}{ccc} \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}} & \xrightarrow{E_x} & Z \\ {\scriptstyle \phi \times \psi}\downarrow & & \| \\ \Sigma_{\mathcal{G}'} \times \Sigma_{\mathcal{H}'} & \xrightarrow{E'_x} & Z \end{array}$$

The backward returns commute by coplay preservation. These two identities compare the boundary behavior independently of whether a pair of strategies satisfies the response tests.

Response-preserving maps carry equilibria to equilibria: one applies the implication to (σ, σ). A noninvertible map can identify equilibria or fail to reach some equilibria of the target. An isomorphism of games instead preserves and reflects the response tests in every context. The geometric information carried by a morphism is therefore different from a claim about the size or structure of its image on equilibrium sets. This distinction matters when strategy representations are quotiented or multiple agents share an interface.

14.3 Probability and continuation transport

For a fixed strategy, one may model forward play by a probability kernel. For finitely supported probabilities $P(y \mid x)$ and $Q(z \mid y)$, their forward composite is

$$(QP)(z \mid x) = \sum_{y} P(y \mid x)\, Q(z \mid y).$$

This has unit given by a point mass at the input. Associativity is the reassociation of the joint sum over intermediate values. If a continuation takes values in $\mathbb{R}$, this kernel transports it by expectation. Indeed

$$\sum_{y} P(y \mid x) \sum_{z} Q(z \mid y) k(z) = \sum_{z} (QP)(z \mid x) k(z).$$

Thus a component that compares expected returns can use the transported function of its possible inputs. Existence of a maximizing strategy is again an additional hypothesis.

Stochastic parallel composition requires a decision about joint laws. Tensoring two kernels using independent randomness gives the product law conditional on their inputs. That tensor does not have the universal property of a cartesian product in stochastic maps. A joint distribution is generally not determined by its two marginals. For example, a nondegenerate finite law paired with itself by sharing its sample and the product of two independent copies of the law have the same marginals. The continuation may detect their dependence. Consequently continual substitution at the level of marginal expected payoffs alone can lose information relevant for later decisions. Bayesian open games supply additional structure for stochastic play, hidden information and best responses [23]. The calculation above describes only the forward-kernel and expected-return part of that framework.

14.4 An equilibrium interface and a revision rule

Response relations are evaluated at a context rather than at a time. To extract a trajectory one supplies a protocol for when revisions occur, which states agents

observe, and how allowed revisions become a state update. For a fixed context, let $B(\sigma)$ denote the set of responses to σ. If it is nonempty at every incumbent, a deterministic revision rule can choose a single $U(\sigma) \in B(\sigma)$ at each incumbent. A fixed point of U is an equilibrium because its selected response is itself. The converse holds if the rule retains the incumbent whenever $\sigma \in B(\sigma)$. The relation alone leaves this decision open, as well as the choice among other responses and the possible regularity of the selector. A protocol can instead attach a probability law to responses, which adds another layer of dynamical data.

In continuous time a smooth population vector field F and its positive rescaling aF by a smooth function a have the same rest points. On any interval over which a solution of one field exists and the scalar remains positive, its trajectory can be reparametrized in time to a solution of the other. This follows from the chain rule by solving for the derivative of the new time parameter to equal a along the trajectory. Quantities measured at specified times can nevertheless change. A contextual equilibrium criterion does not specify the clock, the distribution of samples, or the size of a discrete update. These data enter the dynamical interface in the next chapter.

CHAPTER 15

Open dynamics and network wiring

An open dynamical system specifies how its current state changes under an input from its surroundings. Its output can supply another system with an input. A wiring diagram determines which values are passed and, when species or states are identified, where their rates of change are combined. Different choices of interface give different composition rules.

15.1 Input and output ports

Consider a smooth state space M, an input space A, and an output space B. Write a module as

$$v : M \times A \longrightarrow TM, \qquad h : M \to B, \qquad v(x,a) \in T_x M.$$

The state equation is $\dot{x} = v(x,a)$ and its output is $h(x)$. A single time-independent wiring of modules indexed by i can be expressed by maps

$$w : A_{\mathrm{ext}} \times \prod_i B_i \to \prod_i A_i, \qquad r : \prod_i B_i \to B_{\mathrm{ext}}.$$

At state $\boldsymbol{x} = (x_i)_i$, put $\boldsymbol{b} = (h_i(x_i))_i$ and $\boldsymbol{a} = w(a_{\mathrm{ext}}, \boldsymbol{b})$. The new state velocity has components $v_i(x_i, a_i)$ and the output is $r(\boldsymbol{b})$. A feedback wire can carry $h_i(x_i)$ to the same or another module. As the output depends only on the current state, closing such a wire requires no simultaneous algebraic solution for the inputs.

For nested networks the rule is associative: to compute one component v_i, evaluate the elementary outputs, substitute them into the wiring maps of each enclosing network, and apply the resulting input to v_i. Flattening the network substitutes the same functions into one composite. Associativity of function composition and the canonical comparison of product state spaces give equality of the component velocities. The same calculation applies to outputs. This is the readout form of the wiring algebra developed in [24]. If a readout is also allowed to depend on its instantaneous input, a directed loop may require a separate condition guaranteeing that its algebraic equations have a unique solution.

15.2 Replicator modules

For a population on a strategy simplex, expose a state readout and accept an input a that parametrizes a payoff field. In coordinates the local velocity is

$$v_i(p,a) = p_i \left(f_i(p,a) - \sum_j p_j f_j(p,a) \right).$$

The velocity sums to zero, and $v_i = 0$ on the face $p_i = 0$, for every a. Consequently feedback through the payoff input preserves these algebraic identities. With smooth, locally Lipschitz coefficients on a neighborhood of the simplex, the wired vector field has unique local solutions. The face identities and uniqueness prevent a coordinate from crossing zero in finite time, and the mass identity preserves the normalization. The same argument works for separately normalized populations on a product of simplices, since the identity applies to each local collection of coordinates. This is one way to give a dynamical meaning to interactions between modules. A relation that recognizes mutually admissible responses is a separate structure, as in Chapter 13.

For a varying external input the solution concept requires regularity of the input in time and local integrable bounds for the coefficients. Composition of wiring gives a formula for the field under those assumptions. It gives no single guarantee of convergence to a selected equilibrium. Even on a product of compact simplices the connections can couple the fields and alter stability. With a fixed external input and smooth wiring, compactness does ensure continuation of solutions for all forward times.

15.3 Composition by state identification

For chemical reactions and transport, a different interface uses labels for shared species or sites. Represent an open network by a cospan $I \to S \leftarrow O$ of finite sets. Its internal field is $v_S : \mathbb{R}^S \to \mathbb{R}^S$ on concentrations, or the restriction to nonnegative concentrations. A map $j : S \to P$ induces pullback j^* which duplicates concentrations and pushforward j_* which sums their velocities. The transported field on P is $j_* v_S j^*$.

Suppose $I \xrightarrow{i} S \xleftarrow{o} O$ is followed by $O \xrightarrow{i'} T \xleftarrow{o'} J$. The new species set is the pushout. Denote the structure maps into it by j_S, j_T. The field of the composite is

$$v_P = (j_S)_* v_S j_S^* + (j_T)_* v_T j_T^*.$$

Every species identified by a shared port gets the velocity contributed by both components. For a third composable cospan, the two iterated pushouts are canonically isomorphic. In the common pushout, the contributions from each

original module are the transported fields along its composite structure map. This follows from $(kj)_* = k_* j_*$, $(kj)^* = j^* k^*$, and distributivity of linear pushforward over addition. The zero field on the identity cospan is a unit. Disjoint union puts two networks side by side without adding identifications between them.

For a mass-action reaction with input complex $\alpha \in \mathbb{N}^S$, output complex $\beta \in \mathbb{N}^S$, and rate constant k, its field is $v_S(x) = kx^\alpha(\beta - \alpha)$. The identity $(j^*x)^\alpha = x^{j_*\alpha}$ shows term by term that transport of the field is the field for the reaction with complexes $j_*\alpha, j_*\beta$. Therefore gluing networks and then taking the mass-action equations gives the same field as gluing the module fields. The general construction of [25] retains boundary interfaces and gives a functor from open reaction networks to open dynamical systems.

The same rule applies to linear master equations with a column-conservative generator Q on $\mathbb{R}^S$. The induced matrix $j_* Q j^*$ has nonnegative off-diagonal entries since those entries sum coefficients between disjoint fibers of j. Its columns sum to zero since $\mathbf{1}^\top j_* = \mathbf{1}^\top$ on the source. Thus gluing in this sense preserves the generator conditions. A quotient of an existing time evolution needs an intertwining condition in addition. Open Markov processes and their coarse-graining are studied in [26].

Linear conservation laws also pass through the pushout when their coefficients agree on the shared species. Suppose $\ell_S^\mathsf{T} v_S = 0$ and $\ell_T^\mathsf{T} v_T = 0$ pointwise, and $o^*\ell_S = (i')^*\ell_T$. There is a unique coefficient vector ℓ_P restricting to both local vectors. The identity $\ell_P^\mathsf{T}(j_S)_* = (j_S^*\ell_P)^\mathsf{T}$, and its counterpart for T, annihilate the two terms of the composite velocity separately. Hence $\ell_P^\mathsf{T} x$ is constant along the closed flow. When boundary fluxes are supplied, its derivative equals the weighted incoming flux minus the weighted outgoing flux. This compatibility criterion applies to conserved species weights even when the components use different internal coordinates.

15.4 Boundary behavior of steady states

A steady state with boundary fluxes $a \in \mathbb{R}^I$ and $d \in \mathbb{R}^O$ satisfies

$$v_S(x) + i_* a - o_* d = 0.$$

Its external data are the concentrations i^*x, o^*x and the fluxes a, d. They define a relation between boundary data. The relation of the composite is obtained by composing the two relations across the shared concentrations and fluxes. Here is the exact reason.

The pushout diagram

$$\begin{array}{ccc} O & \xrightarrow{i'} & T \\ {\scriptstyle o}\downarrow & & \downarrow{\scriptstyle j_T} \\ S & \xrightarrow{j_S} & P \end{array}$$

means that a concentration vector on P is equivalent to two vectors on S, T agreeing after pullback to O. The shared flux d enters the two steady equations with opposite signs. Applying $(j_S)_*, (j_T)_*$ and adding cancels it and gives the composite steady equation. Conversely suppose the composite steady equation holds and restrict the concentrations to S, T. Form the two local residuals including their external fluxes but omitting the shared flux. These residuals lie in the kernel of the sum of the two pushforwards to P. The sequence

$$\mathbb{R}^O \xrightarrow{(o_*, -i'_*)} \mathbb{R}^S \oplus \mathbb{R}^T \xrightarrow{(j_S)_* + (j_T)_*} \mathbb{R}^P \longrightarrow 0$$

is exact since the free vector space functor preserves the pushout. Thus the residuals are obtained from some common flux with opposite signs, giving two local steady equations. This proves the relation formula for unrestricted real port fluxes. Additional inequalities on fluxes require their own feasibility check for such a factorization. The relation records possible external observations of steady states. Trajectories and the conditions for reaching those states remain part of the dynamical description.

CHAPTER 16

Semiconjugacy and deterministic aggregation

An aggregate description of selection groups several strategies under the same observed label. Its exactness is a differential question. For every microscopic population with a given aggregate, the instantaneous aggregate velocity must agree. A coarse field that fits a single trajectory need not describe nearby microscopic states with the same observed population.

16.1 Flow maps and projectable fields

Let F and G be smooth vector fields on M and N and let $q : M \to N$ be smooth. We call q a *semiconjugacy* when it carries solution curves of F to solution curves of G with the same time parameter. This is equivalent to

$$Dq_x F(x) = G(q(x)) \qquad (x \in M). \tag{16.1}$$

Indeed, for a source solution $x(t)$, the chain rule makes $q(x(t))$ a solution of the target initial-value problem. Uniqueness then gives commutation of the two flow maps on their common time domain. Conversely, differentiating this commutation at zero gives (16.1). On simplices we use smooth fields on relative neighborhoods whose forward flows preserve the simplices. In this case the derivative at zero may be taken from the right.

$$\begin{array}{ccc} M & \xrightarrow{\Phi_t} & M \\ {\scriptstyle q}\downarrow & & \downarrow{\scriptstyle q} \\ N & \xrightarrow{\Psi_t} & N \end{array}$$

A complete forward flow determines a functor $B(\mathbb{R}_{\geq 0}, +) \to \mathbf{Set}$. Here B makes the monoid of nonnegative times into a one-object category, and the functor sends addition of times to composition of evolution maps. The square says precisely that q is the component of a natural transformation. Smooth semiconjugacies form a category with ordinary identities and composition. The naturality condition uses a common clock. A state-dependent time change such as the positive rescaling in Chapter 14 requires a separate description.

Given F and a surjective submersion q, a suitable G exists exactly when $Dq_x F(x)$ is constant on every fiber $q^{-1}(y)$, as a vector in $T_y N$. Necessity is immediate from (16.1). For sufficiency, define the candidate at y to be this common value. A local smooth section of the submersion expresses it as a smooth field on a neighborhood of y, and the expressions from different sections agree. Surjectivity also gives uniqueness. For an affine quotient of simplices by a partition, a global affine section is available. Choose a probability vector supported in each block and lift coarse masses to these fixed proportions. Evaluating DqF on that section gives the smooth candidate whenever the fiber condition holds. The candidate preserves the coarse simplex whenever the original flow preserves the fine simplex, by applying the chain rule to a source solution through a lift of each coarse initial state.

16.2 Partition quotients of selection

Let $b : S \twoheadrightarrow T$ be a partition map and let $B_{\alpha i}$ be 1 if $b(i) = \alpha$, and 0 otherwise. Write $z = Bp$ for the image of the population p, so $z_\alpha = \sum_{b(i)=\alpha} p_i$. Let $F_i(p) = p_i(f_i(p) - \bar{f}(p))$ with $\bar{f}(p) = \sum_i p_i f_i(p)$. For every block with $z_\alpha > 0$ define

$$\widehat{f}_\alpha(p) = \frac{1}{z_\alpha} \sum_{b(i)=\alpha} p_i f_i(p).$$

Summing the microscopic equations and regrouping the mean gives

$$(BF(p))_\alpha = z_\alpha \left(\widehat{f}_\alpha(p) - \sum_{\beta : z_\beta > 0} z_\beta \widehat{f}_\beta(p) \right)$$

on the active blocks. The other block velocities vanish, since all of their microscopic masses are zero.

Proposition 16.1. *For a proposed coarse payoff field $g(z)$, the partition sends the replicator field for f to the replicator field for g if and only if $\widehat{f}_\alpha(p) - g_\alpha(Bp)$ is independent of α among the active blocks for every p.*

Proof. Subtract the two velocities on each active block and divide by its positive mass. The resulting equality requires the difference between the two payoffs to equal the weighted mean of those differences on the active blocks. Hence all of the differences are equal there. A common difference cancels in the centered payoffs and yields equality of the velocities. Inactive blocks contribute neither mass nor velocity to either calculation. □

A sufficient condition is $f_i(p) = g_{b(i)}(Bp) + c(p)$ with the same scalar $c(p)$ for all strategies. For matrix payoffs, the corresponding form is $A = B^{\mathsf{T}} \bar{A} B + \mathbf{1} c^{\mathsf{T}}$, where

c is a constant vector and $\bar{A}$ is the coarse payoff matrix. The proposition is weaker than pointwise payoff agreement within each block. Microscopic deviations from the proposed block payoff can cancel in the population average. This cancellation is tested at every microscopic population in the domain.

16.3 Comparison with linear lumpability

For a Markov equation $\dot{p} = Qp$ with a column-conservative generator, the same partition is exact for another generator $\overline{Q}$ if $BQ = \overline{Q}B$. The column of BQ at a microscopic state i records the total signed rate into each coarse block. A matrix $\overline{Q}$ exists precisely when these columns agree for all i in each source block. To see this, compare the matrices on each basis vector. If the equalities of columns hold, define the coarse column using any microscopic representative of the block. The sums of its entries vanish because the columns of Q sum to zero. Its off-diagonal entries are nonnegative because the corresponding target blocks contain no source state. Thus the resulting matrix is a Markov generator. Equivalently, the intertwining holds for the matrix exponentials at every nonnegative time, by applying it to each power and summing the exponential series. The implication in the reverse direction follows by differentiating the exponentials at zero. For open Markov processes, preservation of interfaces and fluxes adds constraints to the condition on the generators [26]. In the nonlinear selection equation the fiber test applies to the centered payoff averages, whose weights are themselves microscopic state variables.

16.4 Observed equilibria and stability

Semiconjugacy carries a rest point of F to a rest point of G. Over a target rest point the condition says $Dq_x F(x) = 0$, allowing motion tangent to a fiber. If V is a Lyapunov function for G, then $D(V \circ q)_x F(x) = DV_{q(x)} G(q(x))$. Thus the same monotonicity is inherited by the pullback, with strictness only where the coarse Lie derivative is strict. These facts identify the limit of equilibrium and stability information that can be recovered from aggregated coordinates alone.

16.5 Finite-time error under a nonzero defect

Suppose the target is represented in a Euclidean domain, as for an affine simplex. The defect of an observation is

$$r_q(x) = Dq_x F(x) - G(q(x)).$$

Suppose G is L-Lipschitz on a region containing both the observed source solution $q(x(t))$ and a target solution $z(t)$ over $[0, T]$. If the defect along the source solution

has norm at most ϵ, then for $0 \le t \le T$,

$$\|q(x(t)) - z(t)\| \le e^{Lt}\|q(x(0)) - z(0)\| + \begin{cases} \epsilon(e^{Lt} - 1)/L, & L > 0, \\ \epsilon t, & L = 0. \end{cases}$$

The integral equations for the two curves, followed by the triangle inequality, bound their distance at t by its initial value plus the integral of L times the current distance and ϵ. The integral form of Gronwall's inequality yields the claim. For polynomial fields, formal approximate differential equivalences give more specific ways to find and certify reductions [27].

If $h : N \to P$ compares G to a third field J, the chain rule gives

$$r_{h \circ q}(x) = Dh_{q(x)} r_q(x) + r_h(q(x)).$$

In particular a bound on Dh controls the magnification of the first reduction's defect under composition. The error estimate uses a specified time interval and hypotheses about the region visited by the comparison trajectories. Long-time conclusions need additional information. For example, if the target satisfies the contraction inequality

$$\langle u - v, G(u) - G(v) \rangle \le -\lambda \|u - v\|^2 \qquad (\lambda > 0)$$

on the comparison region, the error in Euclidean norm satisfies

$$\|q(x(t)) - z(t)\| \le e^{-\lambda t}\|q(x(0)) - z(0)\| + \frac{\epsilon}{\lambda}(1 - e^{-\lambda t}).$$

For positive distance, differentiate its square, apply the contraction inequality and Cauchy–Schwarz, and divide by twice the distance. At distance zero the upper right derivative of the distance is at most ϵ, since the difference of velocities is then the defect itself. Integrating the resulting differential inequality gives the bound. If the two trajectories remain in the region, the bound holds on their whole common forward time domain.

CHAPTER 17

Fitness systems over varying contexts

A population in an ecological or strategic context uses a set of available types, a state of its interactions, and a rule for evaluating prospective behavior. When contexts vary, these data can have nontrivial transition maps. A model can then be a family of games without having a single global numbering of the types. This chapter uses sheaves to state exactly what must agree between presentations.

17.1 An internal strategy object

Let B be a smooth manifold and let $p : S \to B$ be a smooth locally constant finite covering, with all fibers nonempty. Its sheets over an open set give types in that context. The set of types is a sheaf of finite sets over B. Think of an open subset as a stage on which type labels may be observed. On a trivializing open set U, the probability state space over U is $U \times \Delta(I_U)$, where I_U indexes the local sheets and

$$\Delta(I_U) = \left\{ x \in \mathbb{R}^{I_U} : x_i \geq 0, \quad \sum_i x_i = 1 \right\}.$$

On an overlap the sheets are related by a locally constant permutation $\sigma_{UV}(b) : I_V \to I_U$. The transition $(b, x) \mapsto (b, \sigma_{UV}(b)x)$ carries faces to faces and preserves both the mass and the nonnegativity conditions. The transitions assemble a simplex bundle over B. Equivalently, the simplex is a sheaf of states, using smooth plots on each context. The closed fibers have corners. Smoothness of a fitness map on such a fiber means that its local components extend smoothly to a neighborhood in the affine span.

A fitness operator is a section of the relative function object from the simplex bundle to the bundle of real type-indexed vectors. Concretely it is a collection of functions $f_U(b, x) \in \mathbb{R}^{I_U}$ with

$$f_U(b, \sigma_{UV}(b)x) = \sigma_{UV}(b) f_V(b, x).$$

This is a descent condition on maps between bundles. It can be applied equally to restricted subfamilies of fitness operators. For instance a collection of pairwise

payoff matrices $A_U(b)$ gives the operators $A_U(b)x$ when the matrices satisfy $A_U(b)\sigma_{UV}(b) = \sigma_{UV}(b)A_V(b)$. A constant affine shift of all payoffs leaves the replicator field unchanged, as in Chapter 4. It can be convenient to take the corresponding quotient of fitness data before imposing compatibility.

17.2 Local generators and their gluing

For each context b and a local vector x, the replicator operator assigns

$$R_{f_U}(b,x)_i = x_i(f_U(b,x)_i - \langle x, f_U(b,x)\rangle).$$

Its values are tangent to the affine hyperplane of mass one. Every zero coordinate remains zero in that field. These identities are pointwise in b and are independent of the trivialization.

Proposition 17.1 (Descent of replicator fields)**.** *Under the fitness transition law above, the replicator fields glue to a vertical vector field on the simplex bundle. The flow on any fiber is compatible with relabeling and with restriction to context opens. If the fitness fields are smooth up to the boundary in the specified sense, the flows exist for all nonnegative times on each simplex fiber.*

Proof. A permutation σ preserves the Euclidean pairing, preserves coordinatewise multiplication, and fixes the all-ones vector. Therefore

$$R_{f_U}(b,\sigma x) = \sigma R_{f_V}(b,x).$$

The derivative of the transition in the vertical coordinate is σ, so these equations are exactly the gluing condition for a vertical field. The equality is preserved under further open restriction.

Because the fitness is smooth in an affine neighborhood, the local differential equation has unique solutions. A solution for the V-presentation, carried across the overlap, is a solution in the U-presentation with the same initial condition as the one in that chart. Uniqueness gives the flow comparison. The normalized mass and support identities of Chapter 5 keep each trajectory in its compact simplex. The continuation theorem for smooth ordinary differential equations then gives existence for all nonnegative time at every fixed b. The mass calculation only uses the finite set indexing the simplex. It does not require a distinguished order on that set. □

The commutative square records how a transition acts on the whole flow, rather than on its rest points alone:

$$\begin{array}{ccc} \Delta(I_V) & \xrightarrow{\Phi_t^V} & \Delta(I_V) \\ {\scriptstyle\sigma}\downarrow & & \downarrow{\scriptstyle\sigma} \\ \Delta(I_U) & \xrightarrow{\Phi_t^U} & \Delta(I_U). \end{array}$$

Here t is elapsed time in a fixed context, not a path in B. A model in which context itself changes with time needs a rule for that motion. A connection or transition process would be additional data.

17.3 Parametrization and effective descent

In an ∞-topos $\mathcal{E}$, let X, Y, P be objects over a base B. A family of evaluations from X to Y with parameter P is a morphism $P \times_B X \to Y$. By the adjunction for the exponential in $\mathcal{E}_{/B}$, such a family is classified by $P \to [Y^X]_B$, where $[Y^X]_B$ denotes the fiberwise exponential. An inverse-image or presentation change can preserve evaluation while changing parameter coordinates. The pullback formula follows by applying the functor composing structure maps to the exponential adjunction: for $q : C \to B$ and $T \to C$,

$$q_!(T \times_C q^* X) \simeq q_! T \times_B X.$$

Thus $q^*[Y^X]_B \simeq [(q^*Y)^{q^*X}]_C$ by Yoneda. If $U \to B$ is an effective epimorphism with Čech nerve $U_\bullet$, then

$$\mathcal{E}_{/B} \simeq \lim_{[n]\in\mathbf{\Delta}} \mathcal{E}_{/U_n}.$$

This is descent for slices in an ∞-topos [28, §6.1.3]. Objects and morphisms over B can thus be specified locally along with a homotopy-coherent chain of overlap comparisons. The comparison extends to evaluation families since their parameter maps land in the compatible exponential object.

This gives a semantics for local games. A family can be observed by changing the base, and the result of interpreting that family is independent of a selected cover. For the concrete smooth model the relevant gluing equations reduce to equalities; for spaces with groupoid-valued states a commuting triangle may instead carry an invertible comparison. One must still specify which points are probability states and which maps define fitness. The categorical descent theorem glues these structures once they are defined in the chosen semantics. It does not supply a probability law from the base topos.

17.4 Monodromy and restricted parameter choices

A loop in B transports its strategy fiber by a permutation. The image of the resulting representation of the fundamental group is a monodromy group M. In a connected component, a flat state is a locally constant probability vector in the trivializations. It is specified by a vector x at one base point that is fixed by every permutation in M. To prove it, transport a choice from that point along paths. This is independent of the path exactly when every based loop fixes the vector. Flat states therefore have a common value within each orbit of M on types. If O is such an orbit and w_O is its mass, each of its coordinates has value $w_O/|O|$. The flat parameter region is consequently a simplex on monodromy orbits. A general smooth section of the bundle can vary within a context and needs only to obey transition laws.

A flat payoff matrix, by the same argument, is fixed under conjugation by the monodromy permutations. Its entry depends on an orbit of the diagonal action on ordered pairs of types. These conditions can force parameter tying, even when local populations have several distinguishable types. Conversely, the pair orbits may retain interaction patterns that are lost if one replaces all types by their orbit labels before giving the fitness. Context determines which changes of presentation are legitimate. The next chapter measures a different failure, when a coarse description is a randomized channel rather than a permutation.

CHAPTER 18

Selection under stochastic compression

An observer can group strategies by a classification rule. A deterministic grouping sends each strategy to one class, whereas a randomized classifier sends it to a distribution of classes. These two changes of description have different effects on replicator dynamics. When fitness is pulled back along the classifier, the difference is given by a conditional covariance.

18.1 Columns as conditional distributions

Let I, J be nonempty finite sets, and let A be a column-stochastic matrix from I to J. Its column a_i is the conditional distribution of the label $j \in J$ given a strategy $i \in I$. Write $x \in \Delta(I)$ and $z = Ax$. The pushforward of a distribution is Ax. The corresponding pullback of a payoff vector $\gamma \in \mathbb{R}^J$ is $A^\top \gamma$, since

$$\langle x, A^\top \gamma \rangle = \langle Ax, \gamma \rangle.$$

This is expectation along a Markov kernel, expressed in dual coordinates [20]. If $g : \Delta(J) \to \mathbb{R}^J$ is the observed fitness, consider the fine fitness $f(x) = A^\top g(Ax) + c(x)\mathbf{1}_I$, where c is a common payoff offset. This assumption controls the dependence of a fine type's payoff on all its possible labels. It must be checked in a proposed aggregation.

For a probability vector w, set $M(w) = \operatorname{diag}(w) - ww^\top$. The replicator field is $R_f(w) = M(w)f(w)$. This form includes the constraint that shifts by a common payoff have zero field, since $M(w)\mathbf{1} = 0$. For a channel A define

$$C_A(x) = \operatorname{diag}(Ax) - A \operatorname{diag}(x) A^\top.$$

Theorem 18.1 (Covariance of a compressed replicator field)**.** *Under the payoff hypothesis above,*

$$R_g(Ax) - AR_f(x) = C_A(x)\, g(Ax).$$

The matrix $C_A(x)$ is positive semidefinite. At a given x the defect vanishes exactly when $g(Ax)$ is constant on the support of every column a_i with $x_i > 0$. In particular, exact pushforward for all x and all coarse fitness maps holds precisely when A represents a deterministic function $I \to J$.

Proof. The common offset disappears in $M(x)f(x)$. Using $Axx^\top A^\top = (Ax)(Ax)^\top$ gives

$$M(Ax) - AM(x)A^\top = C_A(x).$$

Multiplication by $g(Ax)$ gives the claimed difference. For an arbitrary test vector $u \in \mathbb{R}^J$,

$$C_A(x) = \sum_i x_i(\operatorname{diag}(a_i) - a_i a_i^\top),$$

$$u^\top C_A(x)u = \sum_i x_i \left(\sum_j (a_i)_j u_j^2 - \left(\sum_j (a_i)_j u_j \right)^2 \right).$$

The summands are conditional variances and are nonnegative. A variance is zero precisely when u is constant on the support of the corresponding column. For a positive semidefinite real symmetric matrix C, $u^\top Cu = 0$ is equivalent to $Cu = 0$ by diagonalization. Taking $u = g(Ax)$ proves the equality criterion. Every deterministic column has zero covariance. Conversely choose x concentrated on any specified i and allow a constant target fitness taking an arbitrary vector value. Universal vanishing then forces every coordinate indicator to have zero variance under a_i, so every coordinate probability is either zero or one. Their sum is one, proving determinism. □

Deterministic compression groups types without spreading any type across distinct observed labels. Its exactness is the coordinate expression of summing replicator derivatives over the fibers of a function, when the payoffs agree within fibers. In stochastic compression, equality can still hold for particular payoffs or particular support faces. It is then a condition on the conditional distribution and the payoff together.

18.2 Composition of two compressions

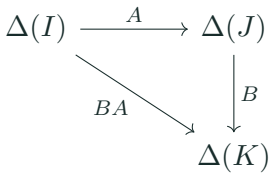
For a second channel $B : J \to K$ the defects obey

$$C_{BA}(x) = C_B(Ax) + BC_A(x)B^\top.$$

Indeed, first add and subtract $B \operatorname{diag}(Ax)B^\top$ in the expression for $C_{BA}(x)$. The two terms are the covariance from the second channel and the covariance of its conditional means induced by the first. This is the law of total covariance for the original sampling sequence. The formula treats loss of a selection signal as a two-stage correlation calculation.

The two summands are positive semidefinite. If the final fitness is $v \in \mathbb{R}^K$, exact pushforward through the composite holds at x precisely when $C_B(Ax)v = 0$ and $C_A(x)B^\top v = 0$. To see this, test the displayed expression in the quadratic

form at v. It is a sum of two nonnegative numbers, each with the zero criterion of Theorem 18.1. The two equality conditions correspond to asking that the final payoff be constant on the occupied conditional supports of the second stage and that its first-stage conditional mean be constant on the relevant supports of the first. This criterion can be applied before composing the channels.



If each step is deterministic, the maps and the exact fitness pullbacks compose without defect. This yields a class of morphisms of selection systems. The same category of stochastic maps can be used for compression with a nonzero defect. In that case the formula identifies the specific obstruction to the map being a morphism of replicator dynamics.

18.3 Trajectory error on a bounded time interval

Write $x(t)$ for the fine solution and $z(t)$ for the coarse solution starting at $Ax(0)$. Suppose both remain on their simplices and that R_g is Lipschitz there with constant L for a fixed norm. Define $c_t = C_A(x(t))g(Ax(t))$. The projected fine solution satisfies

$$\frac{d}{dt}Ax(t) = R_g(Ax(t)) - c_t.$$

Subtracting the coarse equation and applying the integral form of Grönwall's inequality yields

$$\|Ax(t) - z(t)\| \leq \int_0^t e^{L(t-s)}\|c_s\|\, ds.$$

One derivation is to integrate the difference of the two derivatives, bound the target field difference by $L\|Ax - z\|$, and iterate the resulting integral inequality. Its Volterra series sums to the displayed exponential. If the defect is identically zero, uniqueness gives exact semiconjugacy. If the defect is merely small, the estimate controls finite-time approximation. It does not give a uniform long-time guarantee when the bound grows with t. For other mechanisms that yield semiconjugacy and lumpability, see Chapters 12 and 16.

18.4 The tangent and cotangent interpretation

On the relative interior of a simplex, the metric

$$\langle v, w\rangle_x = \sum_i \frac{v_i w_i}{x_i}, \qquad \sum_i v_i = \sum_i w_i = 0$$

identifies the centered payoff covector with its replicator vector. In fact $\langle R_f(x), v\rangle_x = \sum_i f_i(x)v_i$. The metric is the familiar Shahshahani metric [7]. Its inverse on covectors modulo constant vectors has matrix $M(x)$. If Ax is on the boundary of the target simplex, the same construction applies on the relative interior of its support face, with its inverse metric represented in the full coordinates by zero extension. The difference between the target inverse metric and the pushforward of the source inverse metric is $C_A(x)$. The covariance is thus also the change in the quadratic form used to convert a payoff into a selection direction. This observation explains why contraction of information under sampling need not make evolution commute with sampling. The statement proved here uses ordinary finite-dimensional distributions and an exact payoff pullback, which are strong enough to identify the defect in closed form.

CHAPTER 19

Action groupoids and symmetric selection

A permutation group can act on a game by relabeling its types. The resulting symmetry has a fixed locus, an orbit space, and an action groupoid. Each answers a different question. This chapter puts these constructions alongside the probability simplex and uses them to separate symmetry of a state from the quotient of a family of states.

19.1 The fixed simplex and equilibrium

Let a group G act by permutations of a nonempty finite set I, and suppose $f : \Delta(I) \to \mathbb{R}^I$ is continuous and equivariant. The image of G in the permutation group is finite. The fixed simplex $\Delta(I)^G$ consists of the states with equal coordinates on each orbit, so it is naturally a simplex on the orbit set. If p is fixed, $f(p)$ is also fixed. The replicator field and its flow preserve this fixed simplex by the descent calculation in Chapter 17, whenever the field satisfies the existence and uniqueness hypotheses there.

There is an equilibrium in the fixed simplex under the given continuity hypothesis, even before choosing a dynamic solution.

Proposition 19.1 (An invariant Nash state)**.** *Under these hypotheses there is a G-fixed p such that $\langle y, f(p)\rangle \le \langle p, f(p)\rangle$ for every $y \in \Delta(I)$.*

Proof. For any x put $a_i(x) = \max\{f_i(x) - \langle x, f(x)\rangle, 0\}$ and $s(x) = \sum_i a_i(x)$. The continuous map

$$F_i(x) = \frac{x_i + a_i(x)}{1 + s(x)}$$

is equivariant because each operation in its definition commutes with permutation. It maps the fixed simplex to itself. By Brouwer's theorem it has a fixed point p there. The fixed point equation gives $a_i(p) = s(p)p_i$. If $s(p) > 0$, each occupied type has strictly higher fitness than the mean. Taking the mean over all occupied types would give a strict inequality with itself. Hence $s(p) = 0$, and the individual inequalities $f_i(p) \le \langle p, f(p)\rangle$ give the result for all y by linearity in y. □

This calculation uses the ambient equilibrium inequalities. One can also compress along the orbit function and work with orbit masses, provided the fitness assigned to each orbit is the correct one. A restriction of the dynamics to symmetric states only tests attraction in the fixed simplex. For smooth fitness, at an interior fixed point the derivative of the field is a linear map commuting with the permutation action on the tangent space. The averaging projector onto the fixed vectors commutes with that map. The tangent space consequently splits into the fixed part and the kernel of averaging, both invariant under the derivative. Stability on the fixed part leaves the transverse block to be tested separately. At the boundary, feasible perturbations are governed by the tangent cone of the ambient simplex and can enter larger faces.

19.2 What the quotient records

For an action on a space X, write $X//G$ for the realization of the action bar object $[n] \mapsto X \times G^n$. It maps to the delooping BG and the fiber over its distinguished point is X. If X is an ordinary discrete set, this is the groupoid of elements and symmetry arrows. After choosing orbit representatives x, it is equivalent to the coproduct of BG_x over those representatives, where G_x is the subgroup fixing x. Every orbit is a connected groupoid and inclusion of the full subgroupoid at one object is an equivalence. The set of connected components is the coarse orbit set. The fuller quotient retains the automorphism groups of pure states.

The whole probability simplex makes the role of the acted-on object clear. It has an equivariant contraction $\lambda x + (1 - \lambda)u$ to the uniform state u. Its homotopy quotient over BG is therefore equivalent to BG. The pure strategy set, a rest-point subspace, or a space of trajectories may have different quotient types. Equivariant contractibility of the whole simplex gives its quotient type. Each subspace requires a separate computation.

For a smooth equivariant field, the space of trajectories on a fixed time interval carries the action by transport of each trajectory. Initial and terminal evaluation maps are equivariant. In a well-posed deterministic system the initial evaluation is an equivariant homeomorphism to the initial state space, using continuous dependence on initial conditions. Terminal evaluation is then the flow map expressed through this identification. For a set-valued best response or a multivalued update (see Chapter 7), the corresponding initial evaluation can have multiple trajectories over a state. In that situation its fibers encode a choice of resolution of ties or admissible continuations.

19.3 Coherent actions in an ∞-topos

An internal group object in an ∞-topos $\mathcal{E}$ has a reduced bar groupoid $G_\bullet$ and a delooping BG. Effectivity identifies $G_\bullet$ with the Čech nerve of $1 \to BG$. An action on X can be specified by a cartesian transformation into $G_\bullet$ with degree-zero object X. Its realization is the map $X//G \to BG$. Conversely any object over BG pulls back along the Čech nerve to a cartesian transformation. Descent for slices shows that these assignments are inverse [28, 29, §6.1.3]. In particular,

$$\mathrm{Act}_G(\mathcal{E}) \simeq \mathcal{E}_{/BG}.$$

The corresponding square

$$\begin{array}{ccc} X & \longrightarrow & X//G \\ \downarrow & & \downarrow \\ 1 & \longrightarrow & BG \end{array}$$

is a homotopy pullback. This describes quotient families while retaining their lifts and all compatibility conditions for change of context.

An evaluator $P \times X \to Y$ between equivariant objects is represented in this slice, so its local evaluations must transform both their parameters and inputs. For a discrete group acting on spaces, the space of equivariant maps from X to Y is the homotopy fixed-point space of $\mathrm{Map}(X, Y)$ with conjugation transport. For example, the cosimplicial description has in degree n a family of maps for each string in G^n. Transporting the domain along the string identifies it with the usual cosimplicial diagram for the homotopy limit of the conjugation action. Totalization imposes naturality and its higher coherence. The assertion concerns the space of possible equivariant maps. The existence of a natural dynamical invariant, a distinguished equilibrium, or a numerical update requires its own mathematical conditions.

In these constructions the context can be an ordinary parameter space, a classifying groupoid, or an internal base for locally varying types. Equivariance and descent describe how locally constructed fitness and revision laws assemble. They do not determine a revision law without assumptions on its behavioral or statistical meaning. Weight updates provide a particularly direct case where those assumptions and the resulting equations can be specified together.

CHAPTER 20

Learning rules and populations of agents

Selection can act on candidate policies or programs as well as on inherited traits. In a learning system, three choices enter explicitly: what counts as a strategy, how its score is obtained, and how that score changes its frequency. The replicator field is an infinitesimal revision rule. A discrete-time learner can approximate it or use a different update with its own statistical guarantee.

20.1 An update from relative entropy

Let x be a probability vector with positive coordinates on a finite set I. Given a payoff vector a and a step $\eta > 0$, exponential weights defines

$$x_i^+ = \frac{x_i e^{\eta a_i}}{\sum_j x_j e^{\eta a_j}}.$$

A common additive shift of the payoff cancels. The update uniquely maximizes on the simplex

$$y \longmapsto \langle y, a\rangle - \eta^{-1} D(y\|x), \qquad D(y\|x) = \sum_i y_i \log(y_i/x_i),$$

where a zero summand is interpreted as zero. Indeed, substitution of x^+ writes the objective as a constant minus $\eta^{-1} D(y\|x^+)$. Nonnegativity of relative entropy and its equality condition follow from Jensen's inequality for log, with terms of zero mass handled by continuity. If the initial state has zeros, the same definition is used on its support face. This is a multiplicative weights update [30, 31].

Differentiation at $\eta = 0$ gives

$$\left.\frac{\partial x_i^+}{\partial \eta}\right|_{\eta=0} = x_i(a_i - \langle x, a\rangle).$$

Putting $a = f(x)$ yields the replicator field. If f is Lipschitz on the compact simplex, it is bounded there and the local truncation error relative to the Euler step is uniformly $O(\eta^2)$. Indeed the denominator is bounded away from zero for small η and the second derivative with respect to η is uniformly bounded. On a bounded time interval, compare with the Lipschitz ordinary differential

equation. The local flow also differs from its Euler step by $O(\eta^2)$. The discrete error then satisfies a recursion bounded by $(1 + L\eta)$ times the previous error plus $O(\eta^2)$. Summing the geometric series gives an $O(\eta)$ error at grid times for fixed elapsed time. This connects the update to continuous selection without drawing a conclusion about arbitrarily long training.

20.2 The pathwise regret bound

Suppose payoffs a^t are observed each round, with $|a_i^t| \le M$ for every coordinate and round. Their choice may depend on previous weights. Let x^t be the pre-update state and η a fixed step size.

Proposition 20.1 (External regret of exponential weights)**.** *For a comparator p supported on the initial support, after T updates one has*

$$\sum_{t=0}^{T-1} \langle p - x^t, a^t \rangle \le \frac{D(p\|x^0)}{\eta} + \frac{\eta M^2 T}{2}.$$

Proof. Write $Z_t = \sum_i x_i^t e^{\eta a_i^t}$. The update equation gives

$$D(p\|x^{t+1}) - D(p\|x^t) = -\eta \langle p, a^t \rangle + \log Z_t.$$

For a fixed round, the second derivative of $\log \sum_i x_i^t e^{s a_i^t}$ is the variance of a^t under the distribution tilted by s. It is at most M^2 since all of its values belong to $[-M, M]$. The first derivative at zero is $\langle x^t, a^t \rangle$. Integrating the second derivative bound gives $\log Z_t \le \eta \langle x^t, a^t \rangle + \eta^2 M^2/2$. The claimed inequality follows by summing and dropping the nonnegative final relative entropy. □

This is a comparison to a fixed behavior in hindsight. With an appropriate step as a function of the horizon, the bound divided by T tends to zero. It allows an arbitrary sequence of bounded revealed scores and does not require that a single game produce them. For continuous-time regularized reinforcement learning and its relation to replicator and other dynamics, see [32].

20.3 Averaged joint play

The regret theorem yields a statement about a game when the scores are the actual counterfactual payoffs of that game. For a finite action product, let u_i be player i's utility. Each round the players independently use their mixed strategies conditioned on the history, giving a product law ν_t. If player i receives the payoff vector of expected utilities of each of its possible actions against the other players' current choices of distributions, its payoff under its own distribution is $\mathbb{E}_{\nu_t} u_i(a)$. Its payoff from a fixed unilateral deviation b_i is $\mathbb{E}_{\nu_t} u_i(b_i, a_{-i})$. Set

$\overline{\nu}_T = T^{-1} \sum_t \nu_t$. Then the coarse correlated equilibrium inequalities take the form

$$\mathbb{E}_{\overline{\nu}_T}[u_i(b_i, a_{-i}) - u_i(a)] \leq \frac{\mathrm{Regret}_i(T)}{T}$$

for every i, b_i, with regret defined as the maximum over the fixed deviations (or a bound on that maximum). This is obtained by adding the per-round deviations. When these bounds tend to zero every accumulation point of the averaged joint laws is a coarse correlated equilibrium, by continuity of the finitely many linear inequalities. The conclusion concerns joint laws, which can carry correlation through their common time index. A Nash statement on the last product distribution would require other hypotheses. This distinction is important in agent populations whose behavior continues to cycle or react to changing opponents.

20.4 Outer selection and the evaluation interface

A program strategy can be represented by a fixed artifact, with an evaluation map from its identity and the current context to a score. The test distribution, budget, counterfactual comparison class, and timing determine whether the score estimates the intended fitness. An inner process can generate or train candidate artifacts. An outer process can then keep a distribution over verified candidates and update that distribution using scores. Exponential weights keeps the distribution in the simplex and is insensitive to score shifts common to all candidates in a round. These facts are properties of the update. They alone say nothing about the behavior or validity of any candidate.

There is also a compositional concern. If an evaluator reports the same payoff for every strategy in a coarse fiber, deterministic aggregation can commute with selection. If it reports only a randomized label, the calculation of Chapter 18 identifies a covariance defect even when the fine payoffs are conditional expectations of coarse scores. This gives a test of a proposed reporting interface before using it as a summary of population behavior. Delayed reports, context changes, or evaluation mismatch introduce separate errors. A learning guarantee is conditional on specifying those mechanisms and bounding their effect.

The constructions in the book put equilibrium, dynamics, stochastic modeling, and presentation changes into compatible types. A model of interacting agents adds an interpretation of strategies and a mechanism for producing scores at the chosen interface. The resulting learning statements remain relative to that mechanism.

APPENDIX A

Notation

I, $\Delta(I)$, x_i
: Type set, its probability simplex, and mass assigned to type i. See Chapter 1.

$f(x)$, $\langle x, f(x)\rangle$
: The current fitness vector and its mean under x.

$R_f(x)$, $M(x)$
: The replicator field and the matrix $\operatorname{diag}(x) - xx^\top$ that converts a payoff covector into a selection direction. See Chapters 5 and 18.

$D(p\|x)$
: Relative entropy of a comparator p from the state x, finite when the support of p is contained in that of x. See Chapters 6 and 20.

$\operatorname{supp}(x)$
: The types with positive mass in x. The associated support face contains trajectories from x for the replicator field without mutation.

P, Q
: In stochastic chapters, transition kernels or generator matrices with conventions specified at their use. In parametrized games, P can also denote a parameter object.

$A^\top$, $C_A(x)$
: The pullback of payoff vectors along a column-stochastic channel A, and its conditional covariance $\operatorname{diag}(Ax) - A\operatorname{diag}(x)A^\top$. See Chapter 18.

(X, S), (Y, R)
: Forward and backward types of an open-game interface. See Chapter 13.

$\mathcal{E}_{/B}$, $P \times_B X$
: The slice over a context object B and its fiber product. See Chapter 17.

BG, $X//G$
: The delooping of an internal group object and the homotopy quotient of a G-object. See Chapter 19.

APPENDIX B

Machine-checked selection identities

The accompanying Lean project formalizes the algebra that makes deterministic aggregation compatible with selection. It uses Lean 4.19.0, pinned mathlib, and 23 audited theorems. The index types are arbitrary finite types. The statements apply uniformly to every cardinality, including reindexing along a map that leaves some target coordinates empty. See [33] for the proof assistant.

B.1 Drift and the normalized update

Let $x, u : I \to R$, where I has a finite type structure. The formal definitions are

$$m(x) = \sum_i x_i, \qquad \mu_x(u) = \sum_i x_i u_i, \qquad F(x,u)_i = x_i\big(u_i - \mu_x(u)\big).$$

Over a commutative ring the kernel checks

$$\sum_i F(x,u)_i = (1 - m(x))\mu_x(u).$$

It follows that the drift is tangent to the mass-one hyperplane. The zero-coordinate property, invariance under adding a common payoff on that hyperplane, and linearity under uniform payoff scaling are also checked. These are identities of algebraic expressions used in Chapter 5.
Over a field the project defines the reproductive update

$$S(x,u)_i = \frac{x_i u_i}{\mu_x(u)}.$$

It verifies unit mass when the denominator is nonzero. Over a linearly ordered field with x in the simplex and strictly positive fitness u, Lean derives positivity of the denominator and proves that the output lies in the simplex. For affine reproductive weights $w_i = 1 + \eta u_i$, under $m(x) = 1$ and $1 + \eta\mu_x(u) \neq 0$, it also checks the exact discrete-to-drift identity

$$S(x,w)_i - x_i = \frac{\eta}{1 + \eta\mu_x(u)} F(x,u)_i.$$

No time limit is needed to obtain this equality.

B.2 Deterministic aggregation

For a map $q : I \to J$, the Lean definition of aggregation is $(q_* x)_j = \sum_{q(i)=j} x_i$. The project proves preservation of mass and of simplex membership, functoriality under composition, and the pullback-pushforward pairing

$$\mu_x(v \circ q) = \mu_{q_* x}(v).$$

For $v : J \to R$ it proves the two commutation identities

$$q_* F(x, v \circ q) = F(q_* x, v), \qquad q_* S(x, v \circ q) = S(q_* x, v).$$

A state-dependent form and a specialization to pullbacks of payoff matrices make the first equation applicable to the deterministic semiconjugacy of Chapter 16. Each result uses finite sums and function extensionality. The second identity is algebraic as an equality of field terms. Its population interpretation uses the nonzero normalizer as in the preceding section.

B.3 Reproducing the check

From the extracted Lean source, run `lake exe cache get`, `lake build`, and `lake env lean Audit.lean`. The manifest pins the mathlib revision. The last command prints the transitive axioms of the audited declarations. The supplied audit records standard Lean axioms and no admitted proofs. The appendix covers the algebraic vector field and reproductive update. The main text gives the analytical and higher-categorical arguments, including existence of flows, stochastic limits, and the channel covariance formula in Chapter 18.

APPENDIX C

Subject index

References give the chapters or appendices where each term is developed.

Aggregation 12, 16, 18.
Assortative interaction 8.
Best response 7.
Coarse correlated equilibrium 20.
Context 17, 19.
Continuation 13, 14.
Coplay 13.
Covariance defect 18.
Descent 17, 19.
Equilibria 2, 19.
Evolutionarily stable strategy 2, 3, 6.
Exponential weights 20.
Finite populations 9, 11.
Fitness 1, 4.
Formalization B.
Homotopy quotient 19.
Invasion 2, 3.
Lens 13.
Logit dynamics 7.
Lumpability 12.
Lyapunov functions 6.
Markov kernel 9, 12.
Mass action 15.
Mean-field limit 11.
Monodromy 17.
Multipopulation dynamics 8.
Mutation 5, 10.
Open dynamics 15.
Open game 13, 14.
Payoff shift 4, 20.
Pushout gluing 15.
Regret 20.
Revision protocol 7, 14.
Replicator equation 5, 6.
Semiconjugacy 16, 18.
Shahshahani metric 6, 18.
Stochastic stability 10.
Support face 1, 5.
Symmetry 17, 19.

Bibliography

[1] J. McKenzie Alexander. *Evolutionary Game Theory*. Elements in Decision Theory and Philosophy, Cambridge University Press, 2023. doi:10.1017/9781108582063.

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

[3] J. Maynard Smith and G. R. Price. The logic of animal conflict. *Nature* 246 (1973), 15–18. doi:10.1038/246015a0.

[4] P. D. Taylor and L. B. Jonker. Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences* 40 (1978), 145–156. doi:10.1016/0025-5564(78)90077-9.

[5] J. Hofbauer, P. Schuster, and K. Sigmund. A note on evolutionary stable strategies and game dynamics. *Journal of Theoretical Biology* 81 (1979), 609–612. doi:10.1016/0022-5193(79)90058-4.

[6] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998. doi:10.1017/CBO9781139173179.

[7] S. Shahshahani. *A New Mathematical Framework for the Study of Linkage and Selection*. Memoirs of the AMS 17, no. 211 (1979). doi:10.1090/memo/0211.

[8] J. Hofbauer and W. H. Sandholm. Stable games and their dynamics. *Journal of Economic Theory* 144 (2009), 1665–1693.e4. doi:10.1016/j.jet.2009.01.007.

[9] W. H. Sandholm. Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium. *Games* 1 (2010), 3–17. doi:10.3390/g1010003.

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

[11] J.-P. Aubin and A. Cellina. *Differential Inclusions: Set-Valued Maps and Viability Theory*. Springer, 1984. doi:10.1007/978-3-642-69512-4.

[12] J. Hofbauer and S. Sorin. Best response dynamics for continuous zero-sum games. *Discrete and Continuous Dynamical Systems B* 6 (2006), 215–224. doi:10.3934/dcdsb.2006.6.215.

[13] H. Ohtsuki and M. A. Nowak. The replicator equation on graphs. *Journal of Theoretical Biology* 243 (2006), 86–97. doi:10.1016/j.jtbi.2006.06.004.

[14] P. A. P. Moran. Random processes in genetics. *Mathematical Proceedings of the Cambridge Philosophical Society* 54 (1958), 60–71. doi:10.1017/S0305004100033193.

[15] H. P. Young. The evolution of conventions. *Econometrica* 61 (1993), 57–84. doi:10.2307/2951783.

[16] F. T. Leighton and R. L. Rivest. Estimating a probability using finite memory. *IEEE Transactions on Information Theory* 32 (1986), 733–742. doi:10.1109/TIT.1986.1057250.

[17] T. G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability* 7 (1970), 49–58. doi:10.2307/3212147.

[18] M. Benaïm and J. W. Weibull. Deterministic approximation of stochastic evolution in games. *Econometrica* 71 (2003), 873–903. doi:10.1111/1468-0262.00429.

[19] J. G. Kemeny and J. L. Snell. *Finite Markov Chains*. Springer, 1976. Publisher.

[20] T. Fritz. A synthetic approach to Markov kernels, conditional independence, and theorems on sufficient statistics. *Advances in Mathematics* 370 (2020), 107239. doi:10.1016/j.aim.2020.107239.

[21] N. Ghani, J. Hedges, V. Winschel, and P. Zahn. Compositional game theory. *Proceedings of LICS* (2018), 472–481. doi:10.1145/3209108.3209165.

[22] J. Hedges. Morphisms of open games. *Electronic Notes in Theoretical Computer Science* 341 (2018), 151–177. doi:10.1016/j.entcs.2018.11.008.

[23] J. Bolt, J. Hedges, and P. Zahn. Bayesian open games. *Compositionality* 5 (2023), no. 9. doi:10.32408/compositionality-5-9.

[24] D. Vagner, D. I. Spivak, and E. Lerman. Algebras of open dynamical systems on the operad of wiring diagrams. *Theory and Applications of Categories* 30 (2015), 1793–1822. Journal.

[25] J. C. Baez and B. S. Pollard. A compositional framework for reaction networks. *Reviews in Mathematical Physics* 29 (2017), 1750028. doi:10.1142/S0129055X17500283.

[26] J. C. Baez and K. Courser. Coarse-graining open Markov processes. *Theory and Applications of Categories* 33 (2018), 1223–1268. Journal.

[27] L. Cardelli, G. Squillace, M. Tribastone, M. Tschaikowski, and A. Vandin. Formal lumping of polynomial differential equations through approximate equivalences. *Journal of Logical and Algebraic Methods in Programming* 134 (2023), 100876. doi:10.1016/j.jlamp.2023.100876.

[28] J. Lurie. *Higher Topos Theory*. Annals of Mathematics Studies 170, Princeton University Press, 2009. Author manuscript.

[29] T. Nikolaus, U. Schreiber, and D. Stevenson. Principal ∞-bundles: general theory. *Journal of Homotopy and Related Structures* 10 (2015), 749–801. doi:10.1007/s40062-014-0083-6.

[30] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences* 55 (1997), 119–139. doi:10.1006/jcss.1997.1504.

[31] S. Arora, E. Hazan, and S. Kale. The multiplicative weights update method: a meta-algorithm and applications. *Theory of Computing* 8 (2012), 121–164. doi:10.4086/toc.2012.v008a006.

[32] P. Mertikopoulos and W. H. Sandholm. Learning in games via reinforcement and regularization. *Mathematics of Operations Research* 41 (2016), 1297–1324. doi:10.1287/moor.2016.0778.

[33] L. de Moura and S. Ullrich. The Lean 4 theorem prover and programming language. *Automated Deduction – CADE 28* (2021), 625–635. doi:10.1007/978-3-030-79876-5_37.