



TESI DOCTORAL UPF / 2023



Towards a Statistical Physics of Major Evolutionary Transitions
Information, Thermodynamics and Ecology at the Mathematical Foundations of Evolution

Jordi Piñero Fernandez

TESI DOCTORAL UPF / 2023



Universitat
Pompeu Fabra
Barcelona

Towards a Statistical Physics of Major Evolutionary Transitions Information, Thermodynamics and Ecology at the Mathematical Foundations of Evolution

Jordi Piñero Fernandez

Towards a Statistical Physics of Major Evolutionary Transitions

Information, Thermodynamics and Ecology at
the Mathematical Foundations of Evolution

Jordi Piñero Fernandez

TESI DOCTORAL UPF / ANY 2023

Director de la tesi
Prof. Ricard Solé
Department of Medicine and Life Sciences (MELIS)



A la memòria de la meva mare,
Clementa Fernandez Pérez

Acknowledgements

A Complex Trek

This journey begun around the end of 2014. I learned then that in Barcelona, the city where I pursued most of my university studies, there is a legendary group of dauntless scientists for which curiosity has no limits, who dare to ask the kinds of questions that fascinated me ever since I borrowed Sagan's *Cosmos* from my secondary school maths teacher: What is life? Are there laws of biology? Is it truly possible to express all of the complexity in the universe in mathematical form?

These questions had brought me to study and love physics and later high energy physics and relativity. But a crave for something *more*, something *different* kept me searching, and thus I found myself in that strangely looking building in La Barceloneta, where I was scheduled to meet Ricard, who captains the exploration vessel known as the Complex Systems Lab. Little did I know then, what an incredible adventure laid ahead. Since that day, my mind has been exposed to all manner of incredible discoveries that had been made by the most daring complex systems scientists over the past decades. And I became hooked, captured by the mathematical, the philosophical and the transdisciplinary essence of their theories.

In these past years I have had the privilege to meet incredible minds; computer scientists, biochemists, virologists, mathematicians, biologists, ecologists, statistical physicists, engineers and linguists; I had the fortune to visit the Santa Fe Institute on multiple occasions, an enchanted land of intellectual pursuit and a true modern *Bibliotheca Alexandrina* for theory and knowledge. What comes next is my best attempt at transmitting the joy and beauty of what I have come to learn during this odyssey.

Agraïments

A tu, Nora, la música de la meua vida, i a tu, Dolors, la companya de tots els meus viatges, l'amor, avui, demà i ahir, gràcies. I a la mama, gràcies per la vida, per l'amor, i els tantíssims i preuats records. Et portem al cor. Al meu pare, que sempre m'ha recolzat, acompanyat i estimat. Al meus germans Martí i Francesc, per la infància, el paradís perdut, però que guia els nostres passos, us estimo. A tots, gràcies per compartir la felicitat amb mi.

A la meua tieta Matilde (la 'tieta'), que es el pal de paller, i al meu tiet Antonio (el 'tiu'), per la seva memòria d'elefant. A la meua abuela Rosario, que ens cuidava com qui més; i a la meua iaia Sierra, per la nostra complicitat extraordinària que avui encara compartim.

Vull agrair als meus amics, companys i camarades. Al Pepe Ródenas, per tantes i talment dublineses aventures, al Lluís Hernández, pel Club Rashevsky, al Josep Carreres, per que sempre hi puc comptar, a l'Adrià, allà on siguis, a l'Evelín, i a l'Arturo; a l'Albert Sánchez, a l'Ignasi, (el 'pàjaru') i en Genís, el Pebrot, en Quim i el Martí Sala, a tots ells, per les *trobades de trobadors trobaires*, la gresca i el vi fins la matinada. No puc oblidar-me dels companys de la facultat: la Laia, la Mariona i en Ricard, l'Ale, l'Artemi, la Claudia i en Valentí. I als genials JovAAS: en Dani, el Didac, el Sergi, l'Albert i l'Andreu.

Vull donar les gràcies a totes aquelles persones que em van ensenyar: a la Roser, al Gregori, al Pepe, a la Neus, l'Àngels i la Maria Dolors de l'institut; i als professors de la facultat: en Fiol, en Casademunt, en Llosa i el Serra, de qui tots en guardo un bon record per tot allò après i gaudit. I al gran Tong, per les seves *lectures* esplèndides a Cambridge. I no oblidar-me de'n Sid i en Sam, companys de batalla allà al bell nord.

M'agradaria agrair a tots els companys i amics amb qui he compartit aquesta aventura: als meus camarades, en Guim Aguadé i el Víctor Maull, per les nostres gestes i periples, i per una amistat que la complexitat ha unit ('we are fine!'). A l'Aina Ollé-Vila, per la intel·ligència i la bondat. Al Carlos Toscano, per totes les converses i tangents, i tantes hores i cafès debatint; i al incomparable Dani Amor, per aquella fantàstica visita

a Boston. Al Bernat Corominas-Murtra, per ser tant geni i per convidar-me a Vienna a debatre sobre física. A l'inigualable, l'increïble, genial i inconfusible: Luiño. A en Raúl i el Salva, per ser uns pous de saviesa. I a l'Annamarie, per aquells mesos de recerca al CSL. Gràcies a la Núria i en Jordi Pla. I als imprescindibles: la Regina, la Natàlia, l'Alejandro i la Neus.

Vull fer una menció especial per l'Oriol Vall. Per la teva insaciable curiositat, per cuidar-me tant i per ser un referent, per la teva visió, la teva sensibilitat i per la teva gran humanitat.

I also want to thank the invaluable support and wisdom of the great Sidney Redner. And the mighty Simon Levin, who welcomed me at his lab in Princeton during that fall of 2019. Thanks to Dieter, for taking us bird-watching outside of the lab for a change. To Stu Kauffman, for our Thanksgiving dinner at his beautiful home in Santa Fe, and for his genius and creativity. Thanks for all the hospitality by everyone at the Santa Fe Institute, and to the awesome friends I made during my visits there: Deepak Bhat and Jacopo Grilli, And to everyone from the CSSS'19. To my dearest friend, Harun Šiljak; thank you so much for so much. And to Artemy Kolchinsky: thank you for your brilliance and for firing up that lighthouse without which we would stumble in the dark.

Per acabar, vull agrair molt especialment al Professor Ricard Solé. Gràcies per acollir-me al Complex Systems Lab. Per descobrir-me la Complexitat. Per fer-me gaudir de tantes i meravelloses lectures –i la pila va creixent, fins a la criticalitat. Per tota la honestat, la velentia i la humanitat. Pel sentit de l'humor. Per la confiança. Per l'amistat. Per ensenyar-me el valor de les preguntes en ciència. Per abordar i atrevir-se a explorar les idees més intrèpides i audaces; i per compartir-ho amb tots i contagiar-nos d'aquesta gran inquietud que es la curiositat científica. Pel sagrat peregrinatge al Santa Fe Institute, i les 'night crews'. Gràcies per tots aquests anys. Uns anys que m'han canviat la vida i m'han portat coneixement, felicitat, saviesa i una gran il·lusió i per seguir mirant endavant. Ha estat un absolut privilegi, un autèntic plaer i una aventura trepidant. Gràcies.

Abstract

This thesis aims at establishing the basis for a statistical physics theory of major evolutionary transitions. Its structure is based on three major pillars: nonequilibrium thermodynamics and information theory, stochastic processes and cooperator ecosystems, and cognitive systems from a statistical physics perspective. We argue how each of these pillars bound and condition the emergence of major evolutionary transitions. Through a thermodynamic based analysis, we derive bounds on entropy production for primordial replicators and a universal bound on energy harvesting applicable to any molecular system. We study the effects of stochasticity in neutral cooperative ecosystems and arrive at system assembly structures that emerge independently from selective forces. We introduce a morphospace of possible cognitive architectures and propose a unifying conjecture for all such systems.

Resum

Aquesta tesi té com a objectiu establir les bases per una teoria de la física estadística de les grans transicions evolutives. La seva estructura consta de tres pilars bàsics: la termodinàmica de fora de l'equilibri i la teoria de la informació, els processos estocàstics en els ecosistemes de cooperadors, i els sistemes cognitius vistos des de la perspectiva de la física estadística. S'argumenta per què aquests pilars condicionen i limiten l'emergència de les grans transicions evolutives. A través d'una anàlisi termodinàmic, es deriven els límits de producció d'entropia dels replicadors primitius i s'estableix un límit universal a la capacitat d'extracció d'energia aplicable a qualsevol sistema molecular. S'estudien els efectes de l'estocasticitat en ecosistemes de cooperadors neutres i se'n demostra l'emergència d'estructures d'assemblatge de sistema independents de forces selectives. S'introdueix el morfoespai de possibles arquitectures cognitives i es proposa una conjectura unificadora per aquests sistemes.

PREFACE

Life is riddled by complexity. From microscopic organisms to large scale ecosystems, complex patterns of organization shape and connect our living world at all levels. How did this complexity arise? Is there order in the entangled bank?

There is a recurrent joke amongst statistical physicists: ‘if you pick any positive number, as large as you want, it will always be closer to zero than infinity; yet we insist on assuming that the thermodynamic limit exists’. This pun hints at the phenomenon of *more is different*, a motto coined by Phil Anderson in the homonymous 1972 essay. Here, Anderson warns us of the perils of a reductionist view and stresses that, while fundamental laws pertaining to any generic agent (be it an electron, a gene or an ant) uncover the elementary symmetries that govern it, if large numbers of such units are put together, then the microscopic symmetries begin to *break* and new symmetries are revealed. This phenomenon pervades all manner of systems: from magnets to traffic jams and the universe itself, and the idea of a *phase transition* is its conceptual cornerstone. In any case, it is the *increasing* of the number of degrees of freedom of our system, i.e. the *thermodynamic limit*, that begets symmetry breaking. In many ways, this dissertation concerns with the problem of such a limit in its different forms.

But what about life? Isn’t life composed of a hierarchical structure of systems upon systems? How does such a multilevel tapestry of forms and function unfolds? In evolution, is it also true that ‘more is different’? The history of Life of Earth is marked by a small number of major shifts in complexity: from the Origins of Life itself, to the formation of the first multicellular organisms or the emergence of language and societies; these grand innovation events brought to bear by evolution are known as the *major evolutionary transitions*, and they each have left a mark in all the eras of our living planet and across the eons. Our quest here is to understand them.

In this dissertation we bring together three pillars of statistical physics; nonequilibrium thermodynamics, stochastic processes and information theory, wrapped up by the unifying thread of phase transitions. Our goal is to put together a coherent framework in which to study the preconditions for major evolutionary transitions to take place and outline the space of what is possible and what not in the most remarkable shifts in complexity that took place over evolutionary history.

Contents

List of Figures	XVI
List of Tables	XVII
List of Boxes	XVIII
List of Abbreviations	XIX
1. INTRODUCTION	1
1.1. Information, complexity and evolution	3
1.2. Major evolutionary transitions	16
1.3. Major transitions as phase transitions?	21
1.4. Key Questions	25
2. THERMODYNAMICS, INFORMATION AND EVOLUTION	27
2.1. Thermodynamics and information	31
2.1.1. The demon's story	32
2.1.2. Stochastic thermodynamics	35
2.1.3. The thermodynamic nature of information	38
2.2. Replication, information and the 2nd law	39
2.3. Entropic bounds for elementary replicators	40
2.3.1. The three replicator types	42
2.3.2. A landscape of entropic bounds	45
2.4. Universal bounds for nonequilibrium energy harvesting systems	47

2.4.1.	Evolution and energy transduction	49
2.4.2.	Universal bounds on energy harvesting	51
2.4.3.	Illustrative examples	52
3.	STOCHASTIC HOLOBIONTS	56
3.1.	Diversity and complexity	62
3.1.1.	Neutral Theory of Ecology	64
3.2.	Back to Cooperation	67
3.2.1.	Two-Species Cooperator Systems	68
3.3.	Multi-Species Cooperator Systems	74
3.3.1.	Deterministic system	74
3.3.2.	Stochastic system	78
3.3.3.	Open ecosystems	82
4.	LIQUID BRAINS AND COGNITIVE SPACES	88
4.1.	Brains, neurons and solid cognition	92
4.1.1.	Solid brains: criticality	94
4.2.	Liquid brain attractors	97
4.3.	Liquid brain criticality	101
4.4.	The immune system as a liquid brain	103
4.5.	Mapping liquid versus solid cognition	108
5.	DISCUSSION	115
5.1.	Thermodynamic bounds on replication and energy harvesting	115
5.2.	Precursor patterns in cooperative systems	117
5.3.	Towards a unifying theory of cognitive networks	118
6.	CONCLUSIONS	122
7.	PAPERS	

List of Figures

1.1.	Koltsov, Schrödinger & Lotka	2
1.2.	Turing, von Neumann and Shannon.	6
1.3.	Information channel	10
1.4.	Gould, Alberch & Kauffman	13
1.5.	Major conceptual pillars	14
1.6.	Maynard-Smith, Szathmáry & Jablonka	15
1.7.	Sketch sequence of a major evolutionary transition . .	18
1.8.	Phase Transitions in Physics	22
2.1.	Prigogine, BZ reaction & the Krebs cycle	28
2.2.	Boltzmann, entropy and the arrow of time.	31
2.3.	Szilard's Engine.	33
2.4.	Thermodynamics of measurement.	38
2.5.	Coarse-graining the replicators phase space	41
2.6.	Three early replicator types	43
2.7.	Entropy production landscape for replicators	48
2.8.	A diagram of an energy harvesting system.	50
2.9.	Unicycle model	54
3.1.	Darwin and Cooperation	58
3.2.	Merezkhovsky, Margulis and Axelrod	60
3.3.	<i>In silico</i> and <i>in vivo</i> examples of cooperation	61
3.4.	MacArthur, Wilson and Hubbell.	65
3.5.	Diagrammatic descriptions of Island Biogeography and Hubbell's Neutral Theory	68

3.6.	Competitors vs. cooperators I	69
3.7.	Competitors vs. cooperators II	71
3.8.	Attractors in the simplex representation.	77
3.9.	Closed neutral cooperators system.	80
3.10.	Extinction dynamics.	81
3.11.	Survival species and RSA	85
3.12.	Species turn-over patterns.	87
4.1.	Examples of cognitive networks	90
4.2.	Real and formal neurons.	93
4.3.	Criticality in solid-liquid brains.	98
4.4.	Population-level attractors and task allocation	100
4.5.	Modelling ant colony interactions on a lattice	102
4.6.	Perelson, Weisbuch and Parisi	104
4.7.	Frustration and Fitness Landscapes	105
4.8.	Early models of the adaptive immune response	107
4.9.	Collective intelligence	109
4.10.	Morphospace for liquid and solid cognition	110
7.1.	Guesser model	222
7.2.	Optimal correctness.	226
7.3.	Guesser's maximum power increase	227
7.4.	Guesser phases.	230
7.5.	Guesser efficiency patterns.	231

List of Tables

1.1. Major Evolutionary Transitions	20
---	----

List of Boxes

1.1	Shannon's conditions for the definition of Entropy	11
2.1.1	Szilard's Engine The diffusion problem	33
3.1	Diversity indexes	64
3.1.1	Neutral theory of <i>competitors</i>	67
3.2.1	The WKB Method	74
3.2.2	Fluctuations of a simple oscillator	80
4.3	Spin glasses: what & why?	106

List of Abbreviations

IT - Information Theory
MET - Major Evolutionary Transition
DoL - Division of Labour
OoL - Origins of Life
M-S&S - Maynard-Smith & Szathmáry
ST - Stochastic Thermodynamics
LEB - Low Entropy Bound
NT - Neutral Theory
RSA - Relative Species Abundance
SAR - Species-Area Relation
LD - Lifetime Distribution
STD - Species Turnover Distribution
IS - Immune System
WKB - Wentzel-Kramers-Brillouin
FP - Fokker-Planck
BZ - Belousov-Zhabotinsky
LHS - Left Hand Side
RHS - Right Hand Side

Chapter 1

INTRODUCTION

*Nothing makes sense in biology
except under the light of evolution*

T. Dobzhansky

Evolution underpins life's astonishing diversity, its myriad of functions and its complex structures at all scales. Since the revolutionary insights of Darwin and Wallace over 150 years ago [Dar59, DW58], entire generations of scientists have expanded the reach of the original theory in many directions, from genomics to evolutionary game theory. This advances have benefited from groundbreaking discoveries, in particular the discovery of the structure of DNA by Watson, Crick and Franklin as the information-carrying molecule that pervades the molecular fabric of all life on Earth. Nowadays, evolutionary thinking is at the core of all the Life Sciences. Among other implications of these developments is the definition of a tree of life connecting every single living agent in the biosphere to a common, original ancestor.




Figure 1.1: **Koltsov, Schrödinger & Lotka** Three key figures that contributed to the development of theoretical principles of life include (from left to right) Nikolai Koltsov, Erwin Schrodinger and Alfred Lotka. The first two contributed in diverse directions to highlight the nature of living information (ahead of the revolution of molecular biology) and Schrodinger pointed to thermodynamics as a crucial component to understand living systems. Lotka, on the other hand, was a pioneer in suggesting that, also at the ecosystem level, energetics and evolution should be included in our causal description of biological complexity.

Although the path from Darwin and Wallace’s work to the modern synthesis is primarily grounded in experimental approximations (from the fossil record to comparative genomics), theoretical ideas and mathematical models have played a crucial role. Darwin himself¹ acknowledged the influence of Malthus work on the limits of growth in establishing the principle of Natural Selection. Not surprisingly, the development in the 20th century of the evolutionary framework that defined the so-called “New Synthesis” of evolution was based on mathematical modelling. Such models were originally population-based constructs, represented by the interplay between ecological interactions and the introduction of *fitness* [HS88]. But life is also deeply tied to information and the laws of physics. How can these be incorporated into the Darwinian framework?

¹Moreover, Darwin also understood the power of mathematical thinking, as he mentioned in his autobiography. He wrote that he wished he had learned the basic principles of mathematics, “for men thus endowed seem to have an extra sense.”

1.1. Information, complexity and evolution

One central issue within the theory (and a significant hole in Darwin's narrative) was the nature of information itself. Where is the substrate of heredity and how is it propagated from generation to generation? The exact way in which information is stored at the molecular level was elucidated in 1953, but some scholars already advanced the mechanisms using purely theoretical arguments. One of them was Nikolai Koltsov (Figure 1.1a), who, in 1927, hypothesized that inherited traits had to be stored in some class of 'special' double-stranded giant molecules. In a somewhat visionary way, he also concluded that such molecules would replicate in a semi-conservative fashion, using each strand as a template [Soy01].

Koltsov's brilliant prediction was pretty much forgotten, but in 1944 a small book entitled "What is Life?", written by one of the founding fathers of quantum physics, Erwin Schrödinger, brought back the problem of the nature of the hereditary material under a physics perspective [Sch44]. The book was the outcome of a series of lectures delivered in 1943 at the Dublin Institute of Advanced Studies. At that time, there was a debate on whether proteins or nucleic acids could be the right candidates. At some point, Schrödinger conjectured that, whatever the final solution was, it had to be some kind of 'aperiodic crystal'. It was a profound (and eventually correct) insight, and inspired and stimulated the research that ended up in the discovery of the DNA double helix. However, Schrödinger had a wider vision of the problem. He was interested in understanding how a physics of living systems could be defined, given the apparently different nature of biological agents². He pointed out that there were no paradoxes in the high-order structures found in biology, which apparently contradict standard thermodynamics. In a nutshell, if the second law of thermodynamics tells us that the entropy always increases (and so 'disorder' is expected to increase) this would appear as a

²Schrödinger himself had been influenced by the ideas of another famous physicist that contributed to the rise of molecular biology: Max Delbrück. At the time, many theoretical physicists turned to biology and evolution in search of fresh intellectual adventures, e.g. Leo Szilard or Walter Gilbert.

paradox when dealing with life. This is of course not the case: the second law applies to closed systems, while evolved living structures are out of equilibrium, open systems. Such a distinction marked a crucial point for the approach of physics to biological organization.

The importance of thermodynamics also received the attention of scholars within higher-level living systems, far beyond molecules and cells. This was the case of Alfred Lotka, a well-known mathematician whose work on the dynamics of pairwise ecological interactions had an enormous influence in ecology [Lot20, Vol26]. His research (which also inspires part of the developments of this thesis) is best known for the insights it provided about the nature of predator-prey and competitive exchanges [Lot56, SB12, May19, Cas99]. But Lotka was also interested in a wider perspective, in particular on the evolution of stable, diverse communities and the general principles that guided their organization. To get there, Lotka conjectured that universal laws should be constructed from thermodynamics. In one of his papers [Lot22], he starts:

It has been pointed out by Boltzmann' that the fundamental object of contention in the life-struggle, in the evolution of the organic world, is available energy. In accord with this observation is the principle that, in the struggle for existence, the advantage must go to those organisms whose energy-capturing devices are most efficient in directing available energy into channels favourable to the preservation of the species.

and in a follow-up paper entitled “Natural selection as a physical principle” where he advocates for a new approach to evolution of communities where physical universal principles are integrated in a unified theory. For such a theory to succeed, organisms and their interactions should be amenable for a thermodynamics view grounded in engines and flows. He explicitly writes:

For the battle array of organic evolution an assembly of armies of energy transformers-accumulators (plants) and engines (animals); armies composed of multitudes of similar units: the

individual organisms. The similarity of units invites to statistical treatment, the development of a statistical mechanics of which the units shall be, not simple material particles in of the type familiar in the kinetic theory [...]; the units in the new statistical mechanics will be energy transformers subject to irreversible collision which trigger action is a dominant feature.

In the light of modern physics of active matter, this corresponds to an ambitious program to bring together physics, agency and population dynamics. In this list we do not include information in an explicit way. It was too soon for Lotka to have had in mind a computational picture of biological agents (his “units”) as not only engines but also information machines.

The new wave of conceptual ideas that started to shape a complexity view of biology started in the early 1940s with the rise of Cybernetics [Wie19]. Under the visionary perspective of Norbert Wiener a powerful, interdisciplinary field rapidly grew over the next decade with a major ambition: understand biological systems (at any scale) in terms of regulatory feedback processes. Wiener, Arturo Rosenblueth and Julian Bigelow published an influential paper entitled “Behavior, Purpose and Teleology” where they outline the view of a theoretical attack to the problem of behaviour and agency [RWB43]. The same year, another groundbreaking paper was put forward by Warren McCulloch and Walter Pitts: “A Logical Calculus of the Ideas Immanent in Nervous Activity”, [MP43], published in the *Bulletin of Mathematical Biophysics*³. McCulloch-Pitts theory is nothing less than the first serious attempt to formalize neural networks under a computational perspective. Using threshold units as the basic elements for their systems approach to brain function, they showed how

³In 1939, Nicolas Rashevsky (1899-1972) founded *The Bulletin of Mathematical Biophysics*, the first scientific journal explicitly devoted to the theoretical study of Biology and Evolution from a mathematical and physical perspective. Rashevsky was deeply inspired by earlier pioneering work such as D’Arcy Thomson’s *On Growth and Form* (1917) [Tho17], who himself was a contemporary of the late Charles Darwin. Ever since, these seminal ideas have grown into a fully-fledged scientific field [NAS22].



Figure 1.2: **Turing, von Neumann and Shannon.** Alan Turing, John von Neumann and Claude Shannon contributed in fundamental ways to our understanding of the nature of computation and information. Interestingly, all of them explored felt interested in pushing forward the boundaries of the fields they created to ask themselves about living systems, from neural networks and brains to problem-solving, robotic agents (Shannon appears holding his artificial mouse, ‘Theseus’, that he used as part of a maze-solving problem).

Boolean logic functions could be efficiently implemented by neural-like elements. This model has prevailed since, and defines nowadays the basis of most models of neural computation and artificial intelligence.

In parallel with the rise of cybernetics, the British mathematician Alan Turing had published in 1936 his famous paper “On Computable Numbers, with an Application to the Entscheidungsproblem”, [Tur36], where the formal bases of a theory of computation were established. In his paper Turing proved that there is a Universal Computing Machine (now called Universal Turing Machine, UTM) capable of performing any conceivable mathematical computation [Hop84]. The implications of Turing work are still of fundamental relevance in many fields and had enormous implications for the problem of mathematical completeness⁴.

⁴Many scholars consider Turing’s paper one of the most influential works in the history of science. In particular, Turing showed that there was no solution to the decision problem by first showing that the halting problem for Turing machines is undecidable: it is not possible to decide algorithmically whether a Turing machine will ever halt.

Another fundamental actor in this historic period of time, who was well-acquainted with the rest, was the Hungarian-American mathematician John von Neumann, who also contributed to the problem of robust computation in brains versus computers⁵. We owe to von Neumann the definition and formalization of the computer architecture that became the standard until our days (the so called von Neumann architecture). But he was also fascinated by biology, from self-replicating systems to brains. Shortly before his death, he published a small book entitled “The Computer and the Brain”, [VNK12] where he made an attempt to compare both. Von Neumann was particularly interested in a problem that will also be relevant in the last part of this dissertation (see Chapter 4), namely: how can we build a reliable machine made of unreliable units? At the time of this work, computers were gigantic machines and vacuum tubes were prone to failure, which caused the failure of the whole computing process. Evolved brains had obviously solved that, but how?

The picture of this revolutionary period needs to be completed with the contributions made by Claude Shannon. He is considered the father of information theory (IT), which will be a central piece in the problems discussed here both within the context of thermodynamic bounds on evolution and in cognition. In 1948, and based on the pioneering work of George Boole [Boo54] he published⁶, “The Mathematical Theory of Communication” in The Bell System Technical Journal, [Sha48]. There he established the bases of information theory starting from the concept of a communication channel, as depicted in Figure 1.3. In any given communication system, there is a sender and a receiver of a given message, which is transmitted through a medium (such as the air or a wire). In the process, the message needs to be encoded at the beginning and decoded at the end of the linear chain. In general, each of the components

⁵Although not known for many decades, Turing also explored (around 1948) some of the problems associated to neural architectures, in contrast with his own, sequential machine concept. Turing speculated in his notes about the potential of what he called “Unorganized Machines”. For a review of these ideas, see [CP99].

⁶A more accessible description of the theory was also published in the same year as a book, [SW49].

of this process is affected by noise. In other words, a given message can fail to be properly sent due to one or several sources of noise. Shannon was motivated by the problem of defining the best encoding strategy to send a message between a given sender and a receiver, provided a set of constraints derived from the noisiness of the channel. One important part of the theory relies on the definition of a quantity, \mathcal{H} , that relates to the amount of “information” of a message. Shannon was thinking of a sender that would generate sequences of n symbols at random, each one having a probability $\{p_i\}_{i \in \{1, \dots, n\}}$ of being generated. How can we define the amount of information necessary to transmit a given symbol (or an array of such) to the receiver?

One way of looking at the problem is to think in what do we exactly mean by “information” when dealing with a certain number of choices. Let us forget for a moment about the noisiness of the communication channel. Consider, a system of n symbols, e.g. the words in the dictionary. We are here interested in transmitting a single word (regardless of context). Now, some words are *rarer* than others, for example ‘flabbergast’ is much rarer than ‘time’. Thus, we can assign a probability value to each word in the English dictionary, namely p_i for $i = 1, \dots, n$. If we put the focus on the rarest of words, we might as well consider a generic function that depends on the inverse probabilities:

$$h(p_i) = f\left(\frac{1}{p_i}\right). \quad (1.1)$$

In a way, we would like this function to *measure* the degree of ‘surprise’ given by the utterance of the i -th word. Overall, we could say that the average amount of ‘surprise’ in the dictionary is given by

$$\mathcal{H}(p) = \langle h(p) \rangle = \sum_{i=1}^n p_i h(p_i). \quad (1.2)$$

At this point, Shannon proposed three conditions for the function \mathcal{H} (and indirectly for h) to accomplish (see Box 1.1), which lead him to the natural logarithm as an adequate candidate, i.e. $h(p_i) = \ln\left(\frac{1}{p_i}\right)$. Accordingly,

the so called Shannon’s *entropy*, which will be a central concept in this dissertation, was defined as

$$\mathcal{H}(p_1, \dots, p_n) = \langle -\ln p \rangle = - \sum_{i=1}^n p_i \ln p_i \quad (1.3)$$

as the appropriate measure of overall uncertainty in a system of n symbols⁷. As Shannon knew well, entropy was not a new concept. Ludwig Boltzmann and Gibbs (see Chapter 2) already described this quantity as is used in statistical physics. However, we haven’t yet arrived at a definition of *information* but rather only at a measure of uncertainty assigned to a symbol or a dictionary as a whole. To understand information, let us play a simple game. Suppose a sender, Alice, picks up a word from the dictionary completely at random. Importantly, once Alice has picked up a word, say i^* , then she *knows* what exactly word it is, i.e. $p_i^A(t_0) = \delta_{ii^*}$. However, at that instant, Bob (the receiver) has no notion of what the word is, i.e. $p_i^B(t_0) = 1/n, \forall i$. Suppose they build some communication process such that Bob can end up knowing exactly what Alice is holding in her hand (e.g. a ‘yes-or-no’ algorithm spanning all the dictionary). At the end of such process, Bob’s uncertainty about Alice’s word is reduced to zero, this is $p_i^B(t_f) = \delta_{ii^*}$. Formally, we write:

$$I = \mathcal{H}(p^B(t_0)) - \mathcal{H}(p^B(t_m)) = \ln n, \quad (1.4)$$

where the sign convention is determined *a priori*. This is what Shannon called *information*, since it measures all the information acquired by Bob in the process of communication with Alice⁸. A simpler example is given by the so-called ‘binary-choice’. Suppose Alice instead just tosses a coin and it lands either heads or tails, again, she knows exactly what the state of coin is, and now Bob asks ‘Is it heads?’, to which Alice responds (unambiguously) either ‘yes’ or ‘no’. Then, in this case $I = \ln 2 := 1\text{bit}$, the

⁷Advanced and complete approximations to this quantity have been derived recently [HCM23].

⁸It is important to note that *information* only makes sense here when there are two or more types of interacting agents, in this case these Alice and Bob, and in others, one could consider an evolving agent and its environment, etc.

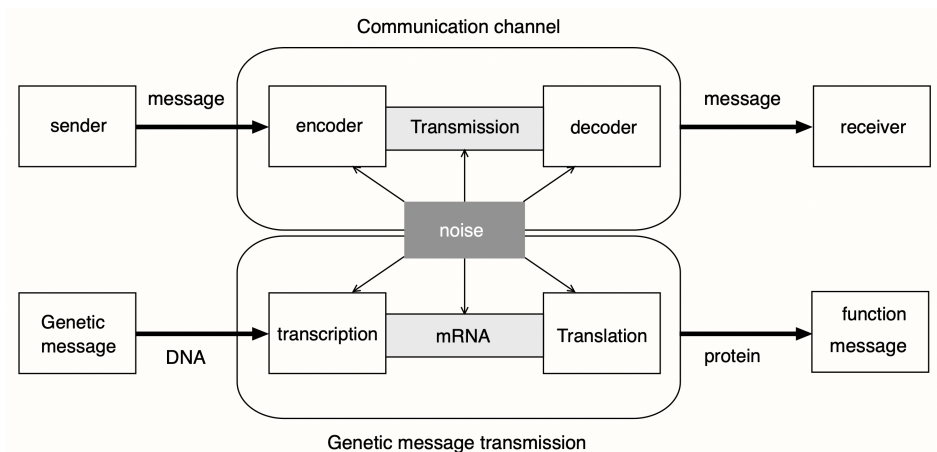


Figure 1.3: **Information channel.** Information transfer through a channel as defined in Shannon’s original work versus genetic information transmission. Diagram adapted from Eigen [Eig13].

fundamental unit of information, with which we are now so familiar after the computer revolution. Shannon did not stop here, and in his application of ‘entropy’ he proved two theorems concerning the fundamental limits of channel capacity, i.e. the maximum amount of information that can be sent through a noisy medium. For a formal introduction on Information Theory, see [CT91].

Remarkably, the theoretical architecture derived from Shannon’s IT can be applied to several research contexts: from linguistics to biology. Consider the case of the genetic code and the processes of transcription and translation depicted in the bottom part of Figure 1.3. Here Alice is picking up a gene, which frequency will depend on the organism’s genotype, while Bob is attempting to *build* the corresponding protein, but first, this information needs to be transmitted. Thus, under this framework transcription and translation are now the code or algorithm that is used to transmit the message, and measuring the distributions over the genotype by Alice and over proteins by Bob gives an account of the information transferring that is taking place [Smi99, Smi00].

As we will elaborate in Chapter 2, information has revealed itself not only as a mathematical means to derive optimality bounds in codes or a method to deal with swaths of data but it became an integral part of physics itself [PHS15]. In fact, information and thermodynamic quantities (such as heat or work) are shown to be intimately related to one another in ways that were, perhaps, unexpected by Shannon himself. This realisation implies that information can be put on the same footing as the energetics that Lotka used to characterize the power flows in living systems. This problem is one of the central topics of this dissertation, and we will come back to it in Sections 2.1.3-2.4.

BOX 1.1: Shannon's conditions for the definition of Information

Three key properties should be fulfilled by such a measure, namely:

1. \mathcal{H} must be a continuous function.
2. If all the probabilities p_i are identical, i. e. $p_i = 1/n$, then \mathcal{H} should be a monotonic increasing function of n .
3. If a choice is broken down into two successive ones, the original \mathcal{H} should be the weighted sum of the individual values of \mathcal{H} .

After their respective proposals, ideas emerging from both cybernetics and information theory had a major influence in molecular biology, ecology and evolution. As aforementioned, one obvious example in such influence is the terms adopted by the early molecular geneticists such as (genetic) code, transcription, translation, or proofreading [Smi99, Smi00]. Within cells, fluctuations can also introduce molecular noise that makes it difficult to propagate the information from DNA to the resulting protein that can perform (or help perform) a given function. But how do they fit into the evolutionary narrative? If life is so deeply informational, how can these ideas and concepts be helpful in describing evolutionary patterns? And what is the role of information thermodynamics?

Finally, a crucial factor here is the potential role played by contin-

gency and randomness versus determinism. This is an old problem within evolutionary biology and has often been cited within the context of Stephen J. Gould's idea of evolution as a highly historical process [Gou90, Gou02]. In Gould's view, the path-dependent nature of evolutionary change has a major impact on how evolutionary patterns and processes unfold over large (macroevolutionary) time scales. To illustrate his point, Gould suggested a mental experiment in his classic and controversial book entitled *Wonderful life*. Think of the biosphere 550 million years ago, before the Cambrian Explosion event. If it were possible to "re-run the tape of evolution", what kind of biosphere would be observable? Following the contingent picture, it is reasonable to conclude that the biosphere would be very different, almost alien. Is that so?

An alternative view is provided by two critical observations: (a) the convergence in evolution [Mor03] and (b) the existence of fundamental constraints that limit the space of the possible. This view is supported by the rigorous study of complex adaptive systems, which posits the existence of higher-level laws of organisation that pervade the structural and dynamical properties of broad classes of nonlinear systems, as shown by the statistical physics approach [Sol11]. In particular, the hierarchic structures evoked by complexity involve the presence of qualitative shifts (phase transitions) which sit at the core of its explanatory power. Here, we will take this stance.

Despite the necessarily historical nature of the process, many observable structures (such as the camera eye design) are re-invented within diverse clades in totally independent ways [Lan10]. Far from anecdotal, this is a widespread pattern in evolution across scales [SMDN15]. Within the context of developmental paths, which are largely responsible for the complexity of the genotype-phenotype mapping, the late Pere Alberch (Figure 1.4 middle) suggested that fundamental constraints exist in accessing the space of the possible forms [Alb89]. One central idea emerging from this work has to do with the possible and the actual. A major contribution first proposed by palaeontologist David Raup [Rau66] is the *morphospace*, i.e. a space of the possible that can help locate the observed systems under consideration (shells, species, artefacts, organs, networks



Figure 1.4: **Gould, Alberch & Kauffman** Is evolution contingent? Although Stephen J. Gould (left) suggested that path dependence would strongly influence future evolutionary outcomes, other scholars, such as Pere Alberch and Stuart Kauffman, have argued against such views. While Alberch showed how developmental constraints might strongly limit the space of the possible, Kauffman introduced the concept of “order for free”: evolution might be more than natural selection.

or cancer types, to name a few) within a two- or three-axis coordinate system. The morphospace helps to chart the diversity of solutions associated with one of these classes of systems and detect patterns of possibilities. Moreover, gaps and voids in this space provide unique windows into the potential limits to the generative rules involved. Sometimes, sharp boundaries separate different designs or an occupied domain next to an empty one. This sharp boundary can be connected with required innovations that allow crossing from one domain to another.

The second pillar comes from a physics perspective of complexity that is deeply connected to the concept of *universality*. As we discuss below, there is a long tradition in physics concerning the study of phases of matter and the transitions between them. In this field of phase transitions [Yeo92, Sol11] the theory (and extensive experimental work) allows defining simple models concerning the ways atoms or molecules interact under different (tunable) external parametric conditions. Each phase has well-defined *qualitative* properties that are markedly different from oth-

ers. The transition from one phase to another is located at a specific points or lines in parameter space, and different transitions may show the *exact* same statistical properties. Such findings had enormous consequences for our views of the nature of change and the presence of fundamental laws.




Figure 1.5: **Major conceptual pillars.** A Venn diagram that sketches the three major conceptual pillars that are considered in this dissertation. At the bottom left, the limitations of information thermodynamics on early living systems and the origins of life; at the top, the role of noise-induced assembly of cooperative interacting systems; and at the bottom right the landscape of cognitive networks about their physicality and computational function.

Along these lines, the pioneering work of theoretical biologist Stuart Kauffman (one of the founding fathers of complexity science) revealed a crucial concept that departs from the Darwinian paradigm: “order for free” [Kau91, Kau93, Kau00]. By working on a diverse range of problems, including gene regulatory networks, rugged landscapes or autocat-

alytic cycles [Kau69a, Kau71, Kau86, KL87] Kauffman discovered that some fundamental properties regarding stability, evolvability and robustness could be achieved “at the edge of chaos” [KK⁺95]. Such an idea can be encapsulated in the very nature of *critical points* from physics. Order for free must be understood as an unexpected source of predictable organization that cannot be explained in terms of natural selection. There are intrinsic, *emergent* properties that strongly affect the space of the possible and cannot be reduced to the microscopic details. This irreducible order is at the core of complex adaptive systems across scales [SG00].



Figure 1.6: **Maynard-Smith, Szathmáry & Jablonka** Evolutionary biologists John Maynard-Smith (left) and Eörs Szathmáry (center) published “The Major Transitions in Evolution” in 1995, where they introduced the concept. Eva Jablonka (right) and Marion J. Lamb focused the problem of METs on the emergence of modes of transmission of information; hereditary and otherwise [JL06]. Jablonka and Lamb later published “Evolution in four dimensions”, in which their novel and visionary ideas were put together [JL14].

The answers to our previous questions are found in the paths towards complexity that life has traced for billions of years. Major innovation events mark these evolutionary trajectories, and understanding such events from a theoretical perspective pervades the motivation of this thesis. In many cases, new forms of replication emerged, while forms of information evolved into others. Replicators and messages are usually connected since the emergence of more complex replicative entities often came with novel forms of gathering, storing and processing information. As dis-

cussed in the following chapters (particularly within the context of cognitive agents), information will be the central concept that allows approaching the problem of cognition and its landscape of possibilities. Along with the mathematical theory outlined above, information has also been crucial to understanding another pillar of biocomplexity: thermodynamics. When dealing with early replicators, thermodynamic constraints are expected to be very relevant. We will elaborate these in Chapter 2. Finally, the ecological context is essential to understand evolution. As Evelyn Hutchinson would say, an evolutionary play occurs in the ecological theatre [Hut65]. We will come back to this crucial point in Chapter 3. The three key domains used here are outlined in Figure 1.5 using a Venn diagram. Many fundamental processes that define evolutionary complexity occur at the intersections between them. But before we proceed, we need first to define the major innovations, their meaning and the theoretical challenges that are explored in this dissertation.

1.2. Major evolutionary transitions

The term *Major Evolutionary Transition* was first coined by Maynard-Smith and Szathmary (M-S&S) [SMS95, MSS95]. In this seminal work, the two evolutionary biologists defined METs as profound innovation events in evolution. Agreement on complete definition of METs is arguably an open problem, and other scholars have extended or adapted M-S&S’s view to enclose other evolutionary innovations [Mor04, DD05, Lan10, JL14, Lan15]. Nonetheless, and for the most part, our discussion here can be firmly grounded on M-S&S’s formulation. Let us now quote [MSS95]:

“[after a major transition takes place], entities that were capable of independent replication before the transition can replicate only as part of a larger whole after it[...].”

This tentative definition puts the focus on the process of *replication*, which is one of the aspects studied in Chapter 2. To help us clarify what does

such a definition entail, let us examine Table 1.1, which resumes the METs considered in [MSS95]. Here, we note that each transition involves the emergence of a new kind of replicator that is built up from previously existing replicative units, which no longer can operate at an individual level: e.g. an insect colony replicates as a colony, not ant by ant [Gor99]. The only exception to this is provided by the earliest MET, i.e. the origins of replication itself.

As mentioned, a remarkable feature after each MET is the emergent information-processing structure. For example, within the Origins of Genetic Information (see Table 1.1), from simple replicator molecules with no translational function, to a large replicator entity based on translation and coding. Here, we identify the emergence of genes themselves, which provide the character of *information* in evolution. I.e., a higher replicator structure that procures new ways of inheritance, regulation, variation and selection. But how do these phenomena arise? And, perhaps more importantly, can a purely adaptive view of evolution explain the emergence of such major transitions?

To advance in these questions, the bulk of the work by Maynard-Smith, Szathmary and later Jablonka and Lamb [JL06, JL14] focuses on the search for universals: a deep study of the METs that reveals the commonalities between all of the transitions. This is a necessary step towards a general theory, which has to result from an interdisciplinary synthesis involving ideas of evolutionary theory and systems biology but also the principles of self-organization, information theory, complexity and phase transitions.

Let us give a brief account of how we can sketch any MET in its simplest form. Suppose there are a number of darwinian agents, which internal complexity allows them to evolve, adapt to an environment and, of course, to replicate (Figure 1.7a). These agents interact amongst themselves. Interactions are the result of information processes that establish an interdependence; these may come from the formation of cooperative (autocatalytic) loops that permit collectives to surpass the *error catastrophe* [ES77, ES82] while avoiding parasitic actors [BH95] (Figure 1.7b), and are precursors to the higher-level entities to emerge.

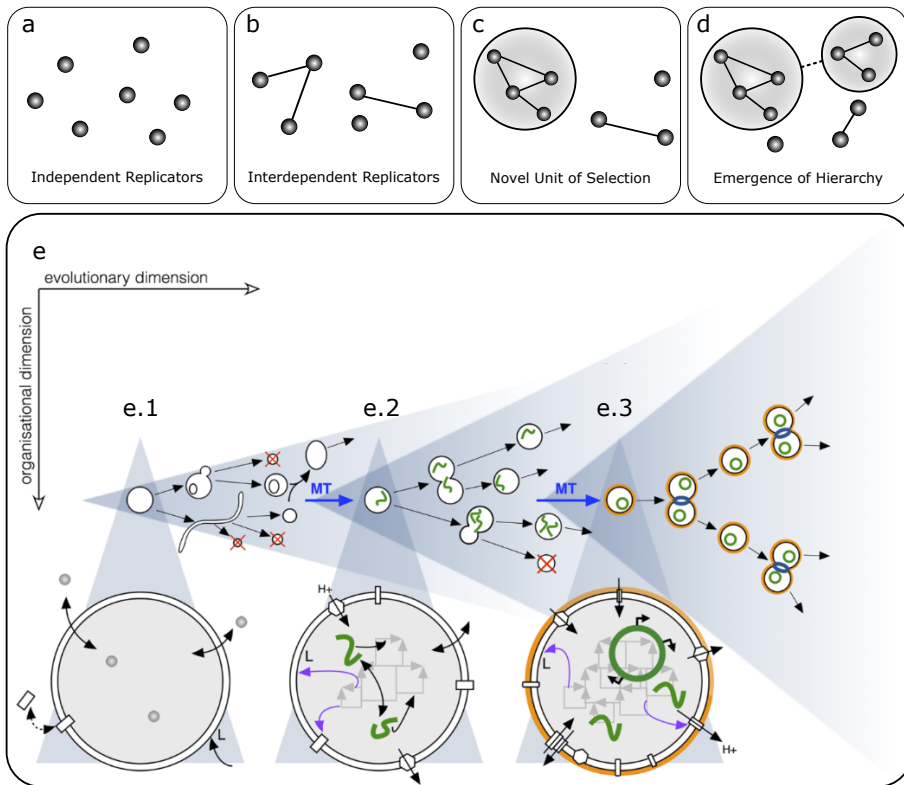


Figure 1.7: Sketch sequence of a major evolutionary transition. Consider a set of replicative agents (a), which establish relations of inter-dependency through some interactions (b). Precursor processes may push these interacting agents to assemble into a higher-level replicator unit, which components are now operating in an integrated manner (c). This gives rise to a hierarchic structure in which higher and lower level agents interact yet their evolutionary landscapes become detached. (e) corresponds to a major transition model from protocell precursors to LUCA (Last Universal Common Ancestor) that is proposed in [SEMSRM17]. In short: (e.1) consists of simple self-assembled vesicles via some cyclic nonequilibrium process involving various molecular species, (e.2) is a secondary stage where self-replication has emerged and (e.3) corresponds to an advanced stage that includes internal complex structures that regulate information and metabolic flows.

Through a process of specialization of the internal degrees of freedom within these collectives, their components undergo *division of labour* (DoL), which finally yields a novel entity that replicates as a whole (Figure 1.7c). This establishes a new level on the hierarchy, in which new interactions will continue to evolve at multiple scales (Figure 1.7d).

Throughout this process, information transmission itself undergoes transitions [JL06]. From replication at the lower-level to the high-level entity, this process requires an integration of information from all its internal components to be passed on into its progeny. A case that clearly exemplifies this phenomenon is the emergence of chromosomes, which are higher-level entities that comprise a large number of replicator units (genes). These new integrated information entities often permit novel ways of exploring complex adaptive landscapes; e.g. through recombination. Finally, it is important to distinguish between complexification and progress: the emergence of a hierarchic structures through the major innovations events in evolution do *not* imply an *a priori* superiority for an evolutionary stage over another. In any case, the hierarchic structure reveals a system of systems, with bottom-up and top-down causation.

Major Evolutionary Transitions [MSS95]	Lower Level	Higher Level	Special Features
Origins of Replication	Non-replicative molecules and/or biocompartments	Replicative molecules and/or biocompartments	Autocatalysis Nonequilibrium
Origins of Genetic Information	Non-translational molecules	Translational-molecules	RNA World Genetic Codes
Origin of Eukarya	Prokaryote	Eukaryote	Endosymbiosis Mutualism
Origins of Sex	Asexual clones	Sexual popul.	Recombination Germline and soma
Origins of Multicellularity	Unicellular populations	Multicellular organisms	Differentiation Specialization
Origins of Eusociality	Solitary individuals	Colonies	Division of Labour Cultural Evolution

Table 1.1: **Major Evolutionary Transitions.** Adapted from [MSS95]. M-S&S also include the transition from primitive societies to human civilizations, pointing at language as one of its main characteristics.

1.3. Major transitions as phase transitions?

We started this section by briefly reviewing some key theoretical results of particular relevance to this dissertation. These pertain to the physics of phase transitions, and all share an intrinsic simplicity in their formulation yet an enormous explanatory power. This is a common feature of the physics approach to complexity: look first at the fundamental rules and principles that pervade the dynamics of a given system, then, try to capture the basic rules that drive interactions among its parts and derive its expected emergent patterns [SMCL⁺96]. This approach has been tremendously successful within statistical physics, in particular in relation with phase transitions, a class of phenomena that has arguably deep ties with evolutionary innovations [Sza15].

In general, phase transitions are associated to *qualitative*, dynamical and structural shifts of organization of the states of matter when an external parameter is tuned. Such changes are regularly captured by a so-called *order parameter*, which essentially reduces the symmetries of the system into a single variable that undergoes a dramatic change whenever a phase transition occurs. The term phase transition has been generalized to many other domains [Sol11]. A typical example is the phase change from liquid to vapor in water, as shown in Figure 1.8e. Each phase is here defined by the density of the system, which acts as an order parameter and experiences a dramatic shift whenever the pressure, and temperature values cross some concrete thresholds (see Figure 1.8b), whereas within a given phase no major differences are observed after manipulation of the specific parameters (in this case, pressure or temperature). E.g., in water at atmospheric pressure, an increase from 2°C to 25°C gives rise to a density change of about 2%, while crossing the boiling temperature produces a change which is about 1600-fold. This is an example of a “first-order” transition, and occurs in a myriad of systems; from polymers to climate change (tipping points), and many others [Sol11]. Biological systems also display these type of phenomena, for example, ant colonies exploring their environment undergo processes that are very well captured by the formalism of phase transitions [BSR01, PS19], and so do flocks of

birds [VZ12] and virus assembly processes [SSE21].




Figure 1.8: Phase Transitions in Physics (a) magnetization of a ferromagnetic, as described by the two-dimensional Ising model. Here we simulate a lattice of spins using different temperatures and plot the average magnetization. Two well-defined phases are observed, separated by a phase transition temperature T_c . In the ferromagnetic phase, symmetry breaking occurs between two equivalent macroscopic states (a,b), while a disordered, unique state is observable in the paramagnetic phase (d). At T_c , (c), coexistence and complex patterns of fluctuations are observable. Images (e) shows the ‘first-order’ phase transition of water from liquid to vapor, and (f) shows the transition that takes place in bird flocking [VZ12].

A different class of transition, known as a second-order phase transition, involves yet another qualitative change but, this time, follows a smooth ‘transition’ order parameter as parameters vary. Consider the behaviour of ferromagnetic materials: when a magnet is heated, it initially maintains its magnetization (the order parameter), which decays until it vanishes at a critical temperature T_c . In these transitions, remarkably rich behaviour is observed, including both very high variance in both structural and dynamical traits. Divergent behaviour is also measured for some macroscopic properties (such as susceptibility and heat capacity) as we approach criticality. A microscopic description that allows formulating a simple model capable of capturing most relevant features of the real transition is the so called *Ising model* or, as physicist Nigel Goldenfeld calls it, “the *Drosophila* model of statistical mechanics”.

Imagine a magnet as a lattice of units (atoms), each one describable as a microscopic magnet (*spin*) itself with two possible orientations: up and down. Mathematically, this entails a collection of N units whose states are indicated as $S_i \in \{-1, +1\}$, for $i = 1 \dots, N$, and are located on a geometric grid. The lowest energy state correspond to that when nearest atoms have the same spin states. Denote the average number of up spins (+1) by N_+ and down spins (-1) by N_- , and define the order parameter of global magnetization as $M = N_+ - N_-$. This value gives a measure of the force with which the magnet can attract other ferrous materials. As the temperature grows, the coupling between spins becomes less and less strong. If temperature T is below a critical value, labeled T_c , interactions among nearest spins favour a dominant direction: either a majority of up spins or the alternative dominated by down spins. However, if $T > T_c$, disorder wins, and spins essentially behave randomly. This is thus a transition between two equally likely ordered states with $M \neq 0$ and a disordered phase with $M = 0$. While the order parameter does not experience a sudden shift, its derivative with respect to T is nonetheless divergent. In fact, other macroscopic measurables such as susceptibility (and heat capacity) still diverge at T_c (see Figure 1.8b).

All such phenomena is captured by the Ising model, which involves

an *energy function*, H , given by:

$$H(\{S_i\}) = -J \sum_{(j,k)} S_j S_k, \quad (1.5)$$

where J is a coupling constant that weights the strength of spin-spin interactions, and brackets (i, k) indicate ordered pairs restricted to nearest neighbors in the lattice [CM05]. Note that, for a pair of spins, energy is minimal when $S_j S_k = 1$, i.e. if both are equally oriented. The mathematical and numerical simulations of (1.5) reveal a number of profound insights on the physics around critical points [Yeo92]. On one hand, when a spin changes its configuration $S_i \rightarrow -S_i$ while the system is at the vicinity of T_c , then the effect it triggers will amplify to the full scale of the system. This is known as the correlation length divergence [Yeo92, Set21], and it implies that, at an effective level, the system is now capable of long range correlations thus allowing it to access a very large set of states (see complex structures forming at T_c in Figure 1.8); this is an essential precursor for processing larger amounts of information (much larger than the individual agent's information capacity, in this case, a single bit per S_i). Over the last decades, this phenomenon has been proposed as a characteristic feature shared by most complex adaptive systems [SMCL⁺96], while some scholars have put forward mechanisms by which systems poise themselves towards criticality in a self-organized manner [BTW88, BP95, Bak13, MB11].

On the other hand, taking (1.5) and extending it to non-trivial topologies beyond the lattice, or allowing interactions to be spin dependent, i.e. via an interaction matrix $(\mathbf{J})_{ij} = J_{ij}$, result in one of the cornerstones of modern statistics physics, namely the theory of spin glasses [SN13, PUZ20]. This generalized version of the Ising model has the capacity adequately model a vast array of complex systems from neural networks to the immune system. In Chapter 4 of this dissertation, we will return to this model in order to explore the space of all possible cognitive networks which involves both mobile and static neural agents. Few are the seemingly innocuous mathematical expressions such as the Ising model that carry within it riches beyond imagination.

1.4. Key Questions

In this Chapter, we have laid out some of the basic ideas that inspired the research paths taken in this dissertation. In the following chapters, we will unfold the problems, open questions and potential theoretical perspectives that are at the core of the research areas that have been addressed here. Instead of studying specific METs, we focus on three major domains (see Figure 1.5) of limits and universals pervading METs. These three theoretical pillars are captured by our key questions:

Question I *How can the thermodynamics of information be used to unveil the feasibility of early living systems?*

The origins of this question have to be found in our interest to understand the minimal requirements for survival and replication of a simple autonomous agent. As discussed above, information is a central problem in approaching the evolution of life, and any attempt to define general laws and bounds to such agents necessarily requires the use of thermodynamics.

Question II *Are there constraints derived from stochastic-induced phenomena in the assembly patterns of cooperative ecosystems?*

Cooperation is a key player in most major evolutionary transitions. This question follows Lotka's intuition on how to achieve minimal descriptions for a collective of interacting species. Moreover, work done at the Complex Systems Lab has revealed that, before the control of genetic regulation, cooperative interactions are responsible for many complex structures that permit METs (such as the transition to multicellularity). Following these motivations, we decided to go to the baseline and define the simplest cooperative ensemble and analyse its properties. Against our intuition, such as neutral, stochastic system reveals order for free.

Question III *Are there limitations to the possible cognitive architectures reachable by evolution? How may the physics of threshold-based neural interactions constrain the space of possible cognitions? What is the impact of the “state” (liquid or solid) on cognitive complexity?*

We will answer this question(s) with a broad modern exploration of the literature which leads us to the formulation of a new conceptual landscape which we term the ‘liquid brain hypothesis’. The goal of this project was to find universal properties that pervade complexity in all cognitive systems from the neural cortex to ant colonies or the immune system. During a workshop at the Santa Fe Institute on “Liquid Brains, Solid Brain”; the necessity of mapping the space of possible cognitions (a morphospace of possible ‘brains’) became clear. Thus, we decided on a systematic effort in this direction. As a result, we wrote a comprehensive paper on the statistical physics of liquid versus solid brains, which unveiled a multiverse of phase transitions at the fundamental theoretical core. Here, we return to this ideas with renewed incentive and push forward this new emerging field by presenting our cognitive morphospace together with a conjecture that, we hope, will bear its fruits in the near future.

Chapter 2

THERMODYNAMICS, INFORMATION AND EVOLUTION

Modern cosmology began with Darwin and Wallace

Leonard Susskind

The quote above is from a legendary physicist, Leonard Susskind, writing in his book *The Cosmic Landscape* about the historical development of cosmology. Cosmology can be seen, in many ways, as the physics counterpart of biological evolution. This well-established field has been growing to become a solid theoretical framework that successfully explains a plethora of phenomena. These include a vast array of scales in space and time, from the first minutes of the universe to the formation of galaxies and galactic clusters to potential scenarios for expansion [Bau22]. It is worth noting that we know now that our own changing universe underwent several major transitions from inflation to the dark ages (after recombination) and onto the production of galaxies, stars and our speck of dust we call Earth. But is this cosmological framework detached from Darwinian evolution? Physicists like Lee Smolin do not believe so and have speculated about a natural selection process shaping the




Figure 2.1: **Prigogine, BZ reaction & the Krebs cycle.** Ilya Prigogine (a) posed that is made of dissipative structures. Simpler, but far from trivial examples from chemistry also reveal the potential for complex dynamics out of chemical reactions. This is the case (b) of the Belousov-Zhabotinsky reaction, which generates large-scale periodic webs in space and time. Biological systems display complex dissipative structures that organize and sustain the flows of matter, energy and information. One example is the Krebs cycle (c), an image by David Goodsell.

space of possible universes [Smo92, Smo04]. Scholars such as Harold Morowitz provide their view of major transitions by merging cosmological and biological change [Mor04]. Within this philosophical framework, major evolutionary transitions as described in the previous chapter, are the natural follow-up of cosmic evolution. What are a theory's most relevant components that encapsulate biology, physics and information?

In this dissertation, one of the main pillars of our approximation to evolution is based on thermodynamics. Some visionary scholars carried out some early attempts at such an interdisciplinary task [Ras60, Vol83, Vol13a, Vol13b], who pioneered the surge of mathematical biophysics and begun to connect the fields of biology, physics and chemistry through the tapestry of thermodynamics. Major advances were put forward by Ilya Prigogine (Figure 2.1a) and his collaborators [NP89] introduced the concept of self-organization which brought new and provocative ideas to tackle questions such as the emergence of order "out of chaos", while

authors like Morowitz [Mor68, Mor13] emphasized the role of entropy production and information in the molecular processes carried by living systems.

One central idea within the early nonequilibrium thermodynamics developed within the Brussels school (and coined by Prigogine) was the concept of *dissipative structure*. The theoretical framework involves an equation for entropy production, δ_i associated with irreversible processes within the system under consideration. Specifically, such entropy production is written as

$$\delta_i S = \int_V \left(\sum_{\mu=1}^n X_\mu J_\mu \right) dV$$

where the integral domain is the volume V of the system, and the term within parenthesis is the specific entropy production. The sum involves the products between thermodynamic forces X_i (such as gradients) and the fluxes J_i associated to each force.

In a nutshell, a dissipative structure is an open system with a well-defined dynamical organization resulting (in most cases) from symmetry breaking. Prigogine's developments provided the basis for a linear theory of nonequilibrium thermodynamics, and in collaboration with Grégoire Nicolis. They made a unique synthesis of ideas that connected this field with nonlinear dynamics and bifurcation theory. Paradigmatic examples of such structures (which can unfold in space and time) can be found across physical, chemical and biological scales. Two examples are displayed in Figures 2.1b,c. The first is the Belousov-Zhabotinsky reaction. Here, chemical reactions taking place at the microscopic scale and spatial diffusion provide the basis for amplification processes that can be observed as giant spiral waves. This is an excellent example of emergent structures and the underlying irreducible order at the core of complex systems: microscopic interactions taking place in a stochastic context can generate large-scale phenomena that are orders of magnitude away in space and time from the "molecular chaos". One crucial ingredient responsible for this is a process that allows long-range order to be propagated: diffusion. Another set of examples is given by biological rhythms

and a whole family of metabolic cycles driven by a closet set of enzymatic reactions, as it is the case of Kreb's cycle (Figure 2.1c).

What is the relative importance of information, ecology and thermodynamics when dealing with evolutionary innovations? One concept that we consider fundamental here concerns a central problem in our understanding of how complexity originates. Since information-processing entities will necessarily require flows of matter of energy, each step towards a more complex organism will have a cost. Thinking in natural selection and replicators, one could easily conclude that the expected biosphere should be made of simple entities. But the answer to this problem might remain precisely in the capacity of natural systems to adapt and even anticipate their environments. As pointed out by physicist Jorge Wagensberg [Wag00]:

Any real object divides the world into two parts: itself and the rest of the universe. Both portions influence each other mutually through a real or imagined common surface: the boundary. Changes in one induce changes in the other. Some objects of this world display an extremely rare property: they tend towards becoming independent of the uncertainty of their environment. This is equivalent to another trend: the perpetuation of some identity. Let us call these objects living individuals: A living individual is a part of the world with some identity that tends to become independent of the uncertainty of the rest of the world.

In other words, attempts to build an explanatory theory for the emergence of complexity must consider both agent properties and their interactions with their worlds. All of them will involve nonequilibrium states, and a full understanding of their feasibility and robustness will require connecting information and thermodynamics. In this chapter, we review some of the key ideas that are nowadays driving the development of a new theory grounded in stochastic thermodynamics.




Figure 2.2: **Boltzmann, entropy and the arrow of time.** In this drawing, we sketch the concept of the second principle of thermodynamics and the probabilistic description of entropy. From left to right, time increases and macroscopic phenomena occur in a well-defined manner (and not in reverse). This includes biological systems undergoing replication. Drawing courtesy of R. Solé.

2.1. Thermodynamics and information

James C. Maxwell was arguably the most revolutionary physicist of the nineteenth century. After procuring his famous unification of the theory of Electromagnetism, Maxwell turn his efforts into establish the building blocks of Statistical Physics and Kinetic Theory [Max72]. It was around this time that Maxwell came up with a most peculiar idea, a paradox, which later became the source heated debate and inspiration towards a theory of nonequilibrium physics and the nature of information. This now-famous *paradox* is known as the “Maxwell’s demon”, and it eventually led thermodynamics into the realm of the mesoscale: quantum computers, nano-devices and cellular machinery. In this Section, we will give a brief account of Maxwell’s paradox, how it is resolved and what implications stem from considering information as a thermodynamic object.

2.1.1. The demon's story

In 1854, Rudolf Clausius introduced the first statement of the second law of thermodynamics as follows:

Heat can never pass from a colder to a warmer body without some other change, connected therewith, occurring simultaneously.

Later, Lord Kelvin and Max Planck independently postulated equivalent statements of the second law in terms of heat reservoirs and cyclic engines. And just a decade afterwards, James C. Maxwell proposed a *gedanken* experiment that seemed to violate this fundamental law by introducing a “finite being” (later dubbed a *daemon*/demon) capable of letting heat flow from a cold to a warm body while the demon is left unscathed. In short, the system consists of two compartments: on the right, high-temperature particles (rapid jitter) are bouncing on the compartment walls whilst on the left, cold particles are doing so at much lower (average) speeds. The demon can open and close a trap on the wall between the two compartments. Moreover, it can *measure* the speed at which a particle from either side is travelling towards the trap door and open or close it at will. This way, the demon can facilitate the accumulation of hot (high-speed) particles on the hotter compartment and vice-versa, thus violating Clausius's statement of the second law.

This apparent paradox by Maxwell ultimately led to the unraveling of the thermodynamic nature of information. In the year 1929, Leo Szilard [Szi29, LR90] proposed an engine that was capable of extracting $k_B T \ln 2$ units of work after *consuming* a single bit of information on the demon's part 2.1.1. Two decades later, Szilard's work led to Brillouin exploring the formal relation between thermodynamic and information entropies [LR02], the former introduced by Ludwig Boltzmann around 1870 and later refined by Gibbs [Jay65], and the latter proposed by Claude Shannon [Sha48] within the context of communication systems.

BOX 2.1.1 - Szilard's Engine

Szilard's engine is a thought experiment proposed by Leo Szilard [Szi29] and based on a single-particle version of the Maxwell's demon. Consider a box containing a single particle at temperature T . The box can be separated in two compartments using a mobile (massless) wall, which is placed and removed inside the box by the demon such that, when put in place, separates the box into two equal volume spaces.




Figure 2.3: **Szilard's Engine** in four stages (in order): measurement, expansion and work extraction, return to initial state through thermalization at temperature $T = 1/k_B\beta$.

The demon *measures* (i.e. acquires information about) the position of the particle, either on the left (L) or the right (R). If it is at L, then the demon places the wall attached to a piston pointing to extract work out of the expansion of the L compartment from $V/2 \rightarrow V$ (see Fig. 2.3), then remove the wall, thermalize, and restart the cycle. The total work extracted is $W = k_B T \ln V_f/V_i = k_B T \ln 2$, which (apparently) came at no cost. In reality, the demon had first to receive an amount of information equal to $I = \ln 2$, since it is a binary choice (L or R). As later shown by Landauer, the erasure of this bit of information implies the dissipation of $Q = k_B T I = k_B T \ln 2$, which yields $\Delta U = W + Q = 0$, and the thermodynamic cycle is effectively closed.

In a nutshell, Boltzmann considered the number of possible states attainable by a system¹, namely Ω , and postulated that Clausius state func-

¹This is often denoted as the *cardinality* of the phase space.

tion, entropy, is given by²

$$S = k_B \ln \Omega , \quad (2.1)$$

where $k_B = 1.384 \times 10^{-23} \text{ JK}^{-1}$ is the so-called Boltzmann constant. This expression cannot generally be proven as a theorem, but it can be shown to follow all of the properties desired in Clausius' theory, which is mainly concerned with entropy differences rather than exact entropy values [Pel11, Ton12]. In his seminal paper of 1948, [Sha48], Shannon used an analogue of the Boltzmann-Gibbs's entropy to measure the *information* that a message carries (see Chapter 1).

Solving Szilards and Brillouin's questions took over several decades and involved a merge of novel and provocative ideas [PHS15]: the reason why the demon is not able to violate the 2nd Law is because all the information required by it in order to decide whether to open or close the trap will eventually dissipate into heat, more accurately, the demon will dissipate a total of $Q = k_B T I$ (in joules). Therefore, the demon needs to import energy in the form of work so that it can function cyclically, and so that $\Delta U = W + Q$ eventually balances out. This indeed resolves the paradox, but begs the question: why information? How does an abstract (and somewhat subjective) concept such as information acquire a physical (energetic) character? And, if so, what are the implications for information processing as a thermodynamically constrained practice?

Rudolf Landauer advanced these questions in showing that the erasure of one bit of information from any possible memory device involves a heat dump of $k_B T \ln 2$ of the system (i.e. the demon) into the heat bath [Lan61]. This became known as the Landauer bound, and has been since amply demonstrated both theoretically and experimentally [TSU⁺10, BAP⁺12]. On the other hand, Bennett combined the newly found thermodynamics of information with logic, reversibility and computation, while Chaitin and Zurek took these ideas beyond by integrating thermodynamics and Turing's theory of computation, a combination that ultimately

²Boltzmann killed himself in 1906 in Duino, Italy. Buried in Vienna, on his tombstone this formula is engraved.

gave rise to the field of algorithmic information theory [Ben73, Zur89]³.

2.1.2. Stochastic thermodynamics

Gearing back into physics, right before the turn of the century, Chrisoff Jarzynsky and Gavin Crooks realized an extension of statistical physics capable of reproducing a large array of mesoscopic phenomena. Famously, they derived an extended version of the 2nd Law applicable to stochastic (non-deterministic) trajectory ensembles on the phase space [Jar97, Cro98], thus moving beyond the paradigm of Liouville and Poincaré. These ideas lead to the discipline of *stochastic thermodynamics* (ST). In a nutshell, this formalism is capable of giving a notion of energetic balances for systems that evolve according to stochastic equations of motion (e.g. Wiener process). As such this new formalism requires generalized definitions of the core thermodynamical variables such as *heat*, *work* and other state-functions like free energy and entropy (see [PP21] pp. 36-66).

Stochastic work (w) and heat (q) are operators on stochastic trajectories, $x_\lambda(t)$, where λ denotes an external parameter (e.g. position of a piston or external magnetic field), which may itself vary over time. Here x represents the set of variables that constitute the full phase space. Consider a discrete system (coupled to a thermal bath) with energy states $\{\varepsilon_x\}(\lambda)$. This non-equilibrium energetics are typically coupled to externally driving forces that break detailed balanced conditions. During a jump from state $x \rightarrow x'$ (at fixed parameter values) the heat exchanged with the reservoir is given by $q(x \rightarrow x') = (\varepsilon_{x'}(\lambda) - \varepsilon_x(\lambda)) + \phi_{xx'}$, where the term $\phi_{xx'}$ is the driving term that breaks detailed balance condition and depends intrinsically on the direction of the jump (either $x \rightarrow x'$ or $x' \rightarrow x$). This represents, for example, the energetic contribution of ATP hydrolysis in biochemical reactions, or the external action of optical

³For a comprehensive introduction to the history, development and principles behind the theory of logical reversibility and the Landauer bounds, see [Sag14]. For a multidisciplinary compilation of contributions from authors in different fields using the precepts of information theory, Landauer bounds and algorithmic information theory see [Zur90].

tweezers on the system.

By extension, work is defined over a trajectory by

$$w(x_\lambda(t)) = \int_{\lambda_0}^{\lambda_f} d\lambda \frac{\partial \varepsilon_x}{\partial \lambda} + \sum_{(x,x') \in CP[x_\lambda(t)]} \phi_{x'x}, \quad (2.2)$$

with CP indicating the time-ordered consecutive pairs of states followed by the trajectory. Over a full trajectory, the combination of the both definitions for stochastic heat and work exchange yield

$$w(x_\lambda(t)) + q(x_\lambda(t)) = \varepsilon_{x_f} - \varepsilon_{x_0}, \quad (2.3)$$

which can be regarded as a generalized 1st Law, which is now applied at the level of stochastic trajectories. On the other hand, an outright candidate for a definition of *stochastic entropy*(see [PP21] pp. 47-53)⁴:

$$s_x(t) = -k_B \ln p_x(t), \quad (2.4)$$

which is also an operator on a stochastic trajectory $x(t)$. Finally, the connection with the physical entropy of a system from the statistical ensemble perspective is evaluated as the average over the stochastic entropy, i.e.

$$S(t) = \langle s(t) \rangle_t = -k_B \sum_x p_x(t) \ln p_x(t), \quad (2.5)$$

where the subscript t indicates that this is evaluated at a fixed time. At a glance, ST does not seem to constitute an huge jump from the usual precepts of statistical physics, however, the consideration of *trajectory* ensembles instead of the state averaging (which is permitted under the ergodic principle assumed in most equilibrium set-ups) provides a theoretical handle on an enormous array of nonequilibrium systems before intractable [Sei08, Sei12, PP21].

⁴Here onward, wherever is not explicitly required, we shall drop the subscript λ .

Extended 2nd Law and Fluctuation-Dissipation

In order to gain some intuition into the utility of the formalism of ST, let us introduce a few results without a proof. An interesting example, which can be directly derived from the previous considerations, is the relation between probability measures on the ensemble of trajectories and stochastic entropy differences [PP21]:

$$s^{\text{total}}(x_\lambda(t)) = k_B \ln \frac{\mathcal{P}(x_\lambda(t))}{\mathcal{P}(x_{\hat{\lambda}}(\hat{t}))}, \quad (2.6)$$

where $\hat{\cdot}$ indicates time-reversed. This entails an interpretation in terms of reversibility in trajectory space, in other words, the total entropy in ST as an operator on $x_\lambda(t)$ recounts its irreversibility. Considering averages over *forward* ($\langle \cdot \rangle_F$) and *backward* ($\langle \cdot \rangle_B$) trajectories, it is easy to show that

$$\langle \exp[-s^{\text{total}}(x_\lambda)/k_B] \rangle_F = 1 \implies \langle s^{\text{total}}(x_\lambda) \rangle_F \geq 0. \quad (2.7)$$

This result offers a direct relation between the extended formalism of ST and classic thermodynamics, and can be understood as a derivation of the classic 2nd Law from ST axioms.

Jarzynski and Crooks proved that for stochastic trajectories with equilibrium end-states [Cro99, Cro00]⁵

$$\text{Jarzynski equality: } \langle e^{-w} \rangle_F = e^{-\Delta F}, \quad (2.8)$$

$$\text{Crooks equality: } \frac{p(w; \lambda)}{p(-w; \hat{\lambda})} = e^{w - \Delta F}, \quad (2.9)$$

where w is the work done on the system following a given (forward) trajectory of the conditioned ensemble, while ΔF is the equilibrium free-energy difference between its end-states, and $p(w; \lambda)$, $p(-w; \hat{\lambda})$ are the probabilities extracting work w following a forward(λ) and backward($\hat{\lambda}$) driving protocols, respectively. Crook's result is a form of a nonequilibrium fluctuation-dissipation theorem.

⁵Initially, Jarzynski and Crooks proved this relations using equilibrium end-states. Hatano and Sasa later showed that equivalent expressions apply with nonequilibrium end-states, [Hat99, HS01].

2.1.3. The thermodynamic nature of information

Armed with the tools of ST, physicists were capable of bringing closure to the Maxwell's demon problem. In particular, it is easy to prove that a system consisting of an object (the compartment in Szilard's example) and a device or observer (the demon) while in contact with a heat bath can become correlated in a way that a measurement is done and free energy in the form of work can then be extracted while preserving the generalized 2nd Law. Suppose that the object is described by $x \in X$ with




Figure 2.4: **Thermodynamics of measurement.** Generalization of a measurement process from an ST approach: X represents the states of the measuring device, Y the environment or the object which is going to be measured and $\beta = 1/k_B T$ corresponds to the heat reservoir.

energy ε_x associated and, at initial time, it is at equilibrium with a temperature reservoir at $T = 1/k_B \beta$. Any measurement requires that the object and the device (which is described by $y \in Y$) interact in some way. For simplicity, consider that all states of the device (e.g. $y = L, R$ in the Szilard's engine) are energetically equivalent, say $\varepsilon_y = 0$, for all $y \in Y$. By definition, the mutual information between x and y at any given moment in time is simply given by

$$I(X : Y) = \sum_{x,y} p_{xy} \ln \frac{p_{xy}}{p_x p_y}.$$

Now, if object and device are uncorrelated before the measurement ($t = t_0$), then $p_{x,y}(t_0) = p_x^0 p_y^0$. Assume that the measurement does *not* affect the state of the object. The process of measuring will generate some correlation between X and Y . Denote t_m the time immediately after measurement takes place. Then, using (2.5) we derive:

$$\begin{aligned} \Delta S^{\text{sys.}} &= -k_B \sum_{x,y} [p_{xy}^m \ln p_{xy}^m - p_{xy}^0 \ln p_{xy}^0] \\ &= -k_B \sum_{x,y} p_{xy}^m \ln \frac{p_{xy}^m}{p_x^m p_y^m} + \sum_y [p_y^0 \ln p_y^0 - p_y^m \ln p_y^m] \\ &= -k_B I(\text{obj.} : \text{dev.})_m + \Delta S^{\text{dev.}}, \end{aligned}$$

where we used the assumption $p_x^0 = p_x^m$. Note that $I(X : Y)_m$ is the mutual information once the device has measured the object. This result indicates that, if the process is performed adiabatically ($\Delta S^{\text{sys.}} = 0$) then $\Delta S^{\text{dev.}} = k_B I(\text{obj.} : \text{dev.})_m$, i.e. that if information has been acquired from the object, then the device entropy increases. Finally, if the device needs to operate cyclically, then its entropy must be reduced again, which implies a heat dump of value $Q = k_B I(\text{obj.} : \text{dev.})_m$. This neatly brings closure to the Maxwell's demon paradox, and it does so by embracing the physical principles of ST.

2.2. Replication, information and the 2nd law

The conceptualization of information as a natural thermodynamic quantity has lead to renewed interest in the search for mathematical principles in evolutionary theory and early living systems [SM16, VWKK22]. A pressing question pertaining to any evolutionary process under the light of the second law is the entropy production patterns that replication entails. This problem has occupied physicists and complex systems scientists for decades [Ras60, Mor68], and recent progress has come from the theoretical territory of Stochastic Thermodynamics. In England (2013) [Eng13], the author combines the Crooks identity (2.9) and coarse-graining to de-

rive a generalized second law that reads as

$$\Delta S_{\mathcal{A} \rightarrow \mathcal{B}}^{\text{tot.}} = \beta \langle q \rangle_{\mathcal{A} \rightarrow \mathcal{B}} + \Delta S_{\mathcal{A} \rightarrow \mathcal{B}}^{\text{sys.}} \geq \ln \frac{\Pi(\mathcal{A} \rightarrow \mathcal{B})}{\Pi(\mathcal{B} \rightarrow \mathcal{A})}, \quad (2.10)$$

where \mathcal{A} and \mathcal{B} are macroscopic (coarse-grained) states (see Figure 2.5) the averaging $\langle \cdot \rangle_{\mathcal{A} \rightarrow \mathcal{B}}$ corresponds to an average over all possible trajectories that map initial microstates that belong to \mathcal{A} to final microstates at \mathcal{B} , and $\Pi(\mathcal{A} \rightarrow \mathcal{B})$ to the effective transition rate of all such trajectories. This approach allows us to connect the microscopic biochemical level to the larger macroscale (be it molecular, cellular or otherwise).

This refining of the second law permits a treatment of self-replication under the umbrella of stochastic thermodynamics, which was studied in [Eng13] for the simple self-replicator. Here $\mathcal{A} = \bullet$ equals the presence of a single replicator, and $\mathcal{B} = \bullet + \bullet$ indicates the presence of two such replicators, hence, $\mathcal{A} \rightarrow \mathcal{B}$ captures the forward replicative step: $\bullet \rightarrow \bullet + \bullet$.

England's (2013) results connect the reproductive fitness of a simple replicator to its metabolic efficiency (through $\langle q \rangle_{\bullet \rightarrow \bullet\bullet}$) and internal state distributions ($\Delta S_{\bullet \rightarrow \bullet\bullet}^{\text{sys.}}$). In general, the entropic bound in (??) provides a thermodynamic metric that can be used to compare different replicator types that played a crucial role around the first METs, such as the Origins of Life. In the following section, we study their thermodynamic limitations depending on their coarse-grained (effective) coupling parameter values and summarize our findings in [PS18].

2.3. Entropic bounds for elementary replicators

Formally, living systems can be viewed as nonequilibrium dissipative structures that utilize matter, energy and information [Hop94] to self-maintain, replicate and adapt. Most introductory literature on population dynamics and evolution focuses its theoretical efforts on formalizing elementary replicator systems from a kinetic approach, either determinis-




Figure 2.5: **Coarse-graining the replicators phase space.** A generic phase space Ω is parsed in two subsets (not necessarily joint), \mathcal{A} and \mathcal{B} . In the examples proposed in [PS18], such subsets are chosen so that they represent a replicator before and after producing a copy of itself. Transition rates then map to experimental rates at the macroscopic regime (see [PS18] pp. 2-3).

tic [Sch11] or stochastic [KRB10]. However, the overreduction of the “replicator” concept may lead to a loss of a contextual (materialist) view in which components (external to the replicative entity) are too required to complete its reproductive cycle⁶. Thus, and in the spirit of pushing forward towards an all-encompassing theoretical baserock for evolution, an in-between stage in which finer-grained degrees of freedom are not fully swept away is required.

Recent results in the field on ST have proved valuable to advance such questions in a more generalized (thermodynamically framed) context for generic replicator systems [Eng13, PME16] as well as molecular information copying systems [OTW17].

Our paper [PS18] is a theoretical exploration that aims to bring together the aforementioned mathematical tools in order to revisit the problem of the emergence of early replicators and derive simple thermodynamic constraints to compare three general replicator types on the same

⁶See [SM16] pp. 20-21 and also Chapter 8 for a comprehensive critique to the concept of “replicator”.

footing. This is achieved without a large departure from the usual theoretization of the concept of replicators by exploiting novel results on the statistical physics of self-replication and adaptation [Eng13, PME16].

Hence, in the paper [PS18], we set out to explore the thermodynamic feasibility of the three elementary replicator motifs, i.e. *simple*, *hyperbolic* and *parabolic* replicators [SS97, Sza06, ES79, SS01, vK86, ZO87, PJ04]. In particular, we contextualize this decades-old problem by employing the following aspects belonging to the theory of ST:

- The possibility to invoke statistical measures over single trajectories rather than ensembles of trajectories [PP21].
- The so-called extended second law, which can be derived from ST [PME16, BCLP16].
- The coarse-graining scheme derived in [Eng13].

The results obtained allow us to go beyond the regular dynamic systems approach while maintaining a simple mapping from a thermodynamic set-up to a stochastic process view of the three replicator systems.

In this section we briefly account for the basic definitions, methods and results discussed in [PS18] and refer the reader to Sec.2.1-2.2 in for detailed derivations.

2.3.1. The three replicator types

In their paper on the Evolution from replicators to reproducers, [SS97] (and later in [Sza06]) Szathmáry and Maynard-Smith propose a pathway for the origins of heredity, which traces its essential steps to the three basic replicator types (parabolic, simple and hyperbolic) and do so following a particular evolutionary argument which connects *fidelity*, *growth* and *cooperation*. Let us elaborate by considering, at the same time, the mathematical definitions of the replicator types and their implications for establishing a reasonable route to the emergence of the first reproducers.




Figure 2.6: **Three early replicator types.** From left to right: the parabolic, simple (exponential) and hyperbolic replicators. Classified by Szathmary and Maynard-Smith (1997) to form an evolutionary pathway from replicators to reproducers.

Parabolic replication

The term *parabolic* comes from the inherent kinetic growth patterns that a template-based replication displays when a rapid transition between bounded and unbounded templates exists (see left in Figure 2.6). After separating replicative time-scales and the complementary binding reaction, the effective growth equation is given by (see Appendix A in [PS18] for a formal proof of this statement),

$$\frac{dx}{dt} = k\sqrt{x} \Rightarrow x(t) = \left(\frac{kt}{2} + x_0^{1/2} \right)^2. \quad (2.11)$$

The relevant features to be highlighted owing to this replicator type are:

- Growth in its common exponential form is inhibited due to the presence of the complementary binding reaction.
- Diversity is preserved due to the long stability of the bound template compounds and the slow growth patterns, which precludes a strong competitive exclusion between molecular species.

This scenario is proposed by S-MS as a precursor case in the Origins of Life arguing that: stability is central in a context of high en-

ergetic flows (which have the tendency to destabilize), diversity is enhanced (which helps explore the primordial evolutionary landscape); however, heredity is limited, specifically because variations cannot spread rapidly due to the slow growth of the parabolic replicator.

Simple replication

Simple replication is the most common notion of replicator that we employ in biology and in physics. It pertains to the basic ideas of exponential growth, examples of which range from nuclear chain reactions to epidemic spread (see centre Figure 2.6). Its kinetic equation is written as

$$\frac{dx}{dt} = k'x \Rightarrow x(t) = x_0 e^{k't}. \quad (2.12)$$

In the context of prebiotic replicators, it is considered by S-MS as a step forward in their evolutionary pathway from replicators to reproducers, posterior to the regime of parabolic dominance. The argument is simply that their exponential nature boosts the process of competition between variants such that *selection* takes over in an overwhelming manner. However, this phase is not fully stable and it hinders the complexification of the replicative molecules over evolutionary time. This is fundamentally because of Eigen's error threshold or *error catastrophe* [SS82] which imposes a limit on the heritable information carried by the replicator.

Hyperbolic replicator

The later stage in the evolution to reproducers is conducted by *cooperation*. Extensive work on this type of replicators was put forward by Eigen and Schuster (1977) [ES79]. In particular, for a simple autocatalytic (self-cooperative) replicator entity (see right Figure 2.6, its kinetic growth patterns follow:

$$\frac{dx}{dt} = k'x^2 \Rightarrow x(t) = \left(\frac{1}{x_0} - k''t \right)^{-1}, \quad (2.13)$$

which, under no other restrictions, will grow singularly as one approaches a finite time asymptote $t \rightarrow t_s = (k''x_0)^{-1}$. This, of course, is not realistic, but speaks of the essential kinetic difference between this and the former two replicator types.

From an evolutionary perspective, the hyperbolic replicator is connected to two key aspects that pervade all METs:

- Division of Labour (DoL): in cooperating entities, the aggregate begins to distribute functions which interactions maintain the replicative entity.
- Resistance to parasitism: the facilitation of an increase in complexity - which contrasts with the simple replicator type - permits the avoidance of parasites (which tend to lag behind in this race to higher and higher complexity).

As detailed in [Sza06] (and references therein), one of the important drawbacks of hyperbolic replicators comes when one looks further from the simplified single variable models such as Eq. (2.13), and considers autocatalytic cycles of n steps. Here, it is possible to show that the larger the n the stronger the metabolic kinetic directionality of the cycle is required for the system to sustain itself. We will return to this particular problem when we discuss our results on a second paper, in Section 2.4.3 of this Thesis

In the following sections, we aim to embed this evolutionary view of early replicators (and reproducers) into the paradigm of modern nonequilibrium thermodynamics, which provides a different mode of looking at these evolutionary transitions through the arrow of irreversibility.

2.3.2. A landscape of entropic bounds

Considering the arguments made in Sections 2.2 and 2.3.1, we are now in a position to outline the main results of our paper [PS18]. Firstly, we notice that the RHS on (2.10) corresponds to the average total entropy production of the system at the macroscopic (coarse-grained) level undergoing the (macro-)transition $\mathcal{A} \rightarrow \mathcal{B}$ in units of $k_B = 1$ (see (2.6)

and [PP21] pp. 67-70). Hence, we denote by *lower entropic bound* the expression:

$$LEB(\mathcal{A}, \mathcal{B}; \Omega) := \ln \frac{\Pi_{\tau}(\mathcal{A} \rightarrow \mathcal{B})}{\Pi_{\tau}(\mathcal{B} \rightarrow \mathcal{A})}. \quad (2.14)$$

Note that this is a lower bound on the average work performed by the system to undergo the coarse-grained transition from $\mathcal{A} \rightarrow \mathcal{B}$; which is an inescapable cost for any type of evolved or adaptive system. Our goal for the rest of this section is to study how this cost translates into each of the three replicator types when undergoing a single ‘replication’ event.

We consider the following assumptions⁷:

- Replicators are isolated and independently modelled as a population of individuals (n) in a closed urn (fixed maximum population size, N) with free-floating monomers that act as substrates ($N - n$).
- Renewal of ‘seed’ replicators is present but negligible. To all effects and purposes, this allows us to disregard absorbing points and centre the analysis on the thermodynamic restrictions on each replicative strategy rather than its first-passage properties.
- All replicator types are subject to the same effective (macroscopic) decay rate, δ , which is a fixed parameter.
- Particular only to the parabolic replicator: two time scales dominate the system, namely the association-dissociation timescale (τ_0) and the replication (template-copying) timescale (τ_1). By definition, $\tau_0 \ll \tau_1$.

Next, we focus our study at the difference of LEB between each replicator type, i.e.

$$\Delta LEB(\mathbf{r}'|\mathbf{r}) := LEB(\mathbf{r}') - LEB(\mathbf{r}).$$

⁷Details of the derivations of the LEB for each replicator are found in [PS18] Sec. 2.2, 2.3 and Appendix B.

Here \mathbf{r} stands for a replicator type. We use \mathbf{s} , \mathbf{h} and \mathbf{p} for the *simple*, *hyperbolic* and *parabolic* replicators, respectively. We introduce coupling constants for each of the replicators and reduce all constants against the simple replicator replication rate ([PS18] pp. 6-8). We can now define thermodynamic *dominance* (symbol \succ) for a pair of replicator types such that

$$\mathbf{r}' \succ \mathbf{r} \Leftrightarrow \Delta LEB(\mathbf{r}'|\mathbf{r}) > 0,$$

which indicates when a replicator is thermodynamically more favourable than another, respectively. Note that this analysis only focuses on its replicative properties, i.e. self-maintenance, robustness or stability are not considered in this study. This gives rise to a landscape of *LEB* differences, Figure 2.7, that allows for a characterization of the most thermodynamically feasible replicator type depending on the couplings of the system and the initial population, which we termed $x = n(0)/N$.

2.4. Universal bounds for nonequilibrium energy harvesting systems

*Those are some of the things that molecules do
given four billion years of Evolution.*

Carl Sagan

Evolution drives innovation at all scales; however, nothing escapes the laws of thermodynamics. How can molecular systems adapt and increase their information-processing capacities while adequately balancing its thermodynamic costs? Early living systems required harvesting and channeling energy to maintain themselves out-of-equilibrium via error-correction or self-assembly processes. What type of trade-off must any such system satisfy between the entropy-increasing forces pulling it to equilibrium forces and the energy harvesting configuration that keeps it in a nonequilibrium state?




Figure 2.7: Entropy production landscape for replicators. Diagram for the dominance spaces of each replicator types (adapted from Fig.3 in [PS18]). The depictions for *simple*, *hyperbolic* and *parabolic* replicators (shown in the diagram’s legend) appear in the diagram whenever they acquire dominance. The diagram’s horizontal and vertical axis are the hyperbolic and parabolic growth rates divided by the simple replicator growth rate, respectively. Dominance shifts in each of the regions depend on the initial condition x , for example, in the upper right corner of the diagram, dominance shifts from parabolic to hyperbolic as we move the initial condition, x , from 0 to 1. All separatrices are obtained exactly and given in [PS18] pp. 7-8.

To answer this kinds of questions, we turn back to the generalized version of the 2nd Law that the method fo Stochastic Thermodynamics offer. In [PSK23] we provide a rigorous derivation of a universal trade-off between entropic-increasing forces and energetical rewards after systems that poise themselves out-of-equilibrium and couple to another nonequilibrium environment from which energy is extracted. Here, we spare the reader of the technical details and attempt a heuristic review of the results found in our paper as well a general interpretation of their meaning and importance in the context of the emergence of early living systems.

2.4.1. Evolution and energy transduction

Energy transduction is the process via which a system can extract energy from an environment at a given rate while maintaining itself in a nonequilibrium steady state, e.g. a cycle that uses some energy to maintain itself running while acquiring some energy from an external source in the process. Such types of processes pervade all manner of systems: from elementary (molecular) biophysical cycles to artificial stochastic engines⁸ [Hi13, SET98, AW13, ABMM16].

From an evolutionary perspective, the *invention* of energy transduction (or simply, energy harvesting) is intimately linked to key research questions of the field of Origins of Life and METs (see [SM16] pp. 24-25 & 28-31, [Mor68, Lan22]). On one hand, the emergence of energy transduction and storage itself is regarded as a phase transition enabling a coarse-grained description of *system* vs. *environment* which yields a precursor stage of what Morowitz termed the paradigm of *chemical self-organization* (see [SM16] Sec. 6.1.3 pp. 349-355). On the other hand, we consider the emergence of *control* mechanisms, which we have vaguely defined as those mechanisms internal to the system whose function is to push its state to the optimal energy-harvesting configuration. These can be thought of *catalyzers*, which utilize the present protometabolic architecture to manipulate the system. From the theory of early living systems, control mechanisms correspond to the earliest precursor of genes, since their adaptability and mutability is a necessary conditions for their optimizing function. However, such evolutionary processes, from elementary chemical systems to the earliest protocells and protometabolisms, must take place by traversing a thermodynamically favourable pathway (within permissible stochastic deviations).

From our analysis in [PSK23], we may consider a stage of early living systems in which both a background protometabolism (which we term *baseline* mechanisms) and control processes (a space of possible catalyzers) are present. Thus, our analysis concerns with the universal

⁸See also [SM16] Sec. 5.4.1 and 5.4.2 pp. 322-325 for a short summary of the role of energy transduction and storage and the Origins of Early Living Systems.

bounds applicable to such class systems. In actuality, such a process classification can represent many other systems that operate on a molecular level and which internal dynamics are separable into two categories, namely baseline and control. In general, the former is concerned with the entropy-increasing (dissipative) mechanisms while the latter imbues correlations with external (environmental) information, which facilitates transduction and thus maximizes power yield. As an example, consider a glucose transporter on a cell membrane. Its production and degradation rates can be included into its baseline processes, while the sensing of external conditions (i.e. external glucose concentration) and molecular signalling pathways to control the production of said transporter enzymes are considered to control mechanisms⁹




Figure 2.8: **A diagram of an energy harvesting system.** The system's states are depicted as \bullet . The system harvests energy and stores it in a work reservoir, hereby represented as a pulley holding a weight. Energy is extracted from a nonequilibrium environment, which may or may not be dynamic itself. The states of the environment are depicted by \circ . The conjunction of the system and nonequilibrium energy source (environment) is in contact with a thermal bath at $T = 1/k_B\beta$. Baseline processes (e.g. diffusion or degradation) are depicted in solid lines, while control mechanisms (pertaining to the system's self-maintenance processes) are dashed lines.

⁹In the article [PSK23], we explicitly explore a similar biological example in Section 5 of the SM.

2.4.2. Universal bounds on energy harvesting

The central result in [PSK23] comes in the form of a convex optimization problem that reads as:

$$\Delta\dot{W}^* := \max_p -\partial_t^R \mathcal{S}(p) + \langle p - \pi, d - q \rangle. \quad (2.15)$$

For a complete detailed derivation and rigorous analysis of (2.15), see the main text in [PSK23]. Here, we will focus on the interpretation and applications of (2.15). Firstly, the expression (2.15) can be regarded as a theorem whose assumptions are (1) the separability of the full system into baseline and control mechanisms, (2) the 1st and 2nd laws of Stochastic Thermodynamics; and (3) the irreducibility of the baseline processes (i.e. all states of the system are reachable, even if highly unlikely). Now, let us unfold (2.15) term by term:

- $\Delta\dot{W}^*$: it corresponds to the the maximum increase of energy flow into the reservoir that can be achieved via means of *whichever* control mechanisms evolution can possibly concoct.
- $-\partial_t^R \mathcal{S}(p)$: this is the entropic cost due to pushing the system away from equilibrium into a distribution (p) that facilitates energy transduction. Note that if p is very far from equilibrium, then entropy will increase dramatically, hence the minus derivative in front of the entropy functional.
- $\langle p - \pi, d - q \rangle$: this is an average over the distribution reached by the control system (p) from the distribution without any control (π), weighed with the power gain through external driving from the energy source (d) minus the dissipative cost released by the system in the form of heat (q).
- Let us note that the RHS of (2.15) contains **no** terms that depend on the specifics of the control mechanisms employed by the system. Thus, it is *universal* in the sense that the maximum attainable power increase, $\Delta\dot{W}^*$, is a true unbreakable maximum given the baseline description of the system.

Solving (2.15) in general is not always possible, but can be resolved using standard numeric recipes [DB16]. Closed-form solutions can be obtained in three physically meaningful regimes, which we term *linear response* (LR), corresponding to the situation in which the deviation of configurations between controlled and uncontrolled systems, $p - \pi$, is very small; *macroscopic* (M), which occurs when the pay-off is very large, and the system can ‘afford’ a high entropic cost; and *far-from-equilibrium* (FE), that results from a perturbative extension of the latter, i.e. when entropic costs are small but non-negligible.

A distinction to be made between our analysis in [PSK23] and the work done in the context of the so-called *information engines* (e.g. [Szi29, MJ12]), is that in the latter the system is set up to consume information as a *fuel*, just like in the Szilard engine, where the knowledge of a bit can yield up to $k_B T \ln 2$ joules, whereas our formalism focuses on information as a *catalyzer*, i.e. the system utilizes information (here interpreted as the deviation of the baseline steady-state distribution π) to maximize its coupling to an external power source which can yield arbitrary amounts of power yield. This is the reason why we believe this principled approach is more coherent in the context of early living systems and the origins of life.

2.4.3. Illustrative examples

In order to motivate our theoretical results, in [PSK23] we provide the reader with some examples. Here we will recover one significant case inspired in the logic of biochemical cycles, dubbed the *unicyclic model*. In Appendix 2 of this dissertation, we provide another example dubbed the *minimal thermodynamic guesser* offering a full derivation of the system within the three regimes outlined above.

Unicyclic model

Consider a ring of n states that is coupled to a nonequilibrium environment out of which the energy can be drawn when transitions between two

specific states occur (see Figure 2.9a). This example can be interpreted as a biochemical cycle that captures an energetic molecule (with some value Θ in units of energy). Now, control mechanisms could be provided by any possible catalytic molecules that facilitate the directionality of the cycle to maximize the energy output rate. In absence of such control mechanisms, and evolving only under baseline (diffusive or degradation) processes, the cycle tends to equilibrate such that no net flux is generated and the power gain zeroes out on average.

However, the addition of catalytic forces also comes at a cost: e.g. production of catalyzers (paid off by a heat release) or internal entropy costs (derived from the internal microscopic configuration of the catalyzers). Remarkably, the second law allows us to capture all of these costs in a term that does *not* depend on whichever specific catalyzers are at play. Hence, the following analysis will hold true whatever the catalytic system evolution can come up with over time.

In [PSK23], we solve for the LR, M and FE regimes respectively, compare our results with exact (numerical) solutions obtained with `cvxpy` [DB16] (see pp. 4-5 in [PSK23] and Section 3 of the SM) and extend the analysis to the case of a fluctuating environment, for which the reaction that brings in the energetic flow moves along the cycle in a random manner.

Here, we recover some of these results and ponder its meaning in the context of precursor conditions for the early metabolic systems. As argued above, a key point in the analysis of this model is that it gives rise to a strong constraint on the evolutionary capabilities of a cyclic system that transduces energy from a nonequilibrium environment in a simple manner. For example, if the system is poised at the LR regime, then the problem can be exactly solved and the maximum amount of power increase attainable by any set of control mechanisms (catalyzers) is given by the expression:

$$\Delta\dot{W}^* = \beta\Theta^2 \frac{(n-1)}{4n^2}, \quad (2.16)$$

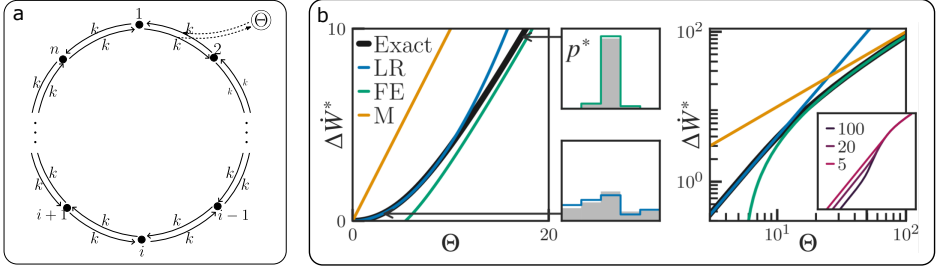


Figure 2.9: Unicycle model. Fig. (a) depicts a symmetric unicyclic model coupled to a static environment that allows energy to flow when transition $1 \rightarrow 2$ takes place. Plots in (b) show how the three regimes (LR, M and FE) approximate the exactly computed solution for $\Delta\dot{W}^*$ as well as the corresponding optimal distributions. The right figure shows the same in log-log scales and the inset shows how the value for $\Delta\dot{W}^*$ changes in the LR regime for different n values.

and the deviation from baseline steady-state follows

$$\Delta p_i^* = [p_i^* - \pi]_i = \frac{\beta\Theta}{4n^2} \begin{cases} 2(i-1) - (n+1) & \text{for } i = 2, \dots, n \\ (n-1) & \text{for } i = n \end{cases}. \quad (2.17)$$

Moreover, in this case, it is possible to assess the validity for the LR approximation by setting $\|\Delta p^*\| \ll 1$, which means that the system will remain in the LR regime as long as

$$\Theta \ll \Theta_{\text{LR}} := 4\sqrt{3} \sqrt{\frac{n^3}{n^2-1}} k_B T. \quad (2.18)$$

Together with (2.16), these are general results that carry a remarkable interpretation. Suppose that the number of elements in the cycle is quite large, i.e. $n \gg 1$, which would correspond to large biochemical cycles in chemical configuration space. Thus, if energy harvesting is to emerge out of the exploitation of such cycles, these processes will more likely be bound to the LR regime, since $\Theta_{\text{LR}} \sim n^{1/2}$. Therefore, this restricts the total energy output due to its characteristically high entropic costs.

Moreover, the longer the cycle, the higher the cost and the smaller the *returns* on the maximum power increase possible, which follows

$$\Delta\dot{W}^* \sim \frac{1}{n}.$$

This can be regarded as a direct thermodynamic limitation on the evolutionary prospects of any molecular system which engages -by evolutionary or selective means- in a process of maximizing energy harvesting, a necessary element for the emergence of larger more complex entities.

Chapter 3

STOCHASTIC HOLOBIONTS

*The species in which peace and mutual support are the rule,
prosper, while the unsociable species decay*

Pyotr Kropotkin

*Life is a symbiotic and cooperative union that allows those
who associate to succeed*

Lynn Margulis

Many major evolutionary transitions are grounded in the emergence of new forms of cooperation capable of holding together higher-order entities from simpler ones. Cooperation pervades the rise of molecular systems capable of overcoming mutation thresholds, multicellular assemblies incorporating division of labour, or insect societies' appearance. Each of these structures incorporates new properties that cannot be observed at the level of its parts. Despite the burden involved in sustaining the new, larger entity, the advantage of staying together can overcome, under some circumstances, the cost of the association.

Cooperation can be achieved, in particular through closed catalytic loops. For example, mutualistic interactions pervade ecological communities at many scales, from bacterial communities to microbiomes and

large-scale ecosystems [Bro15]. Charles Darwin already outlined the presence of these reciprocal relations in one of his memorable studies on the ecology of earthworms, his last book [Dar92, Wil07]. The diagram of Figure 3.1(a) summarises the underlying idea. Darwin realised the presence of such cooperative connection, and it is worth mentioning that worms (not considered fascinating animals at that time) are included in Darwin's famous quote about the "endless forms":

It is interesting to contemplate a tangled bank, clothed with many plants of many kinds, with birds singing on the bushes, with various insects flitting about, and with worms crawling through the damp earth, and to reflect that these elaborately constructed forms, so different from each other, and dependent upon each other in so complex a manner, have all been produced by laws acting around us. These laws, taken in the largest sense, being Growth with reproduction; Inheritance which is almost implied by reproduction; Variability from the indirect and direct action of the conditions of life, and from use and disuse; a Ratio of Increase so high as to lead to a Struggle for Life, and as a consequence to Natural Selection, entailing Divergence of Character and the Extinction of less improved forms. Thus, from the war of nature, from famine and death, the most exalted object which we are capable of conceiving, namely, the production of the higher animals, directly follows. There is grandeur in this view of life, with its several powers, having been originally breathed by the Creator into a few forms or into one; and that, whilst this planet has gone circling on according to the fixed law of gravity, from so simple a beginning endless forms most beautiful and most wonderful have been, and are being evolved.

In his work, which included -as usual- many experiments and field observations, Darwin realized that earthworms improve soil porosity and organic content that helps plants grow, resulting in more organic matter




Figure 3.1: Darwin and Cooperation. Darwin mentioned the existence of cooperative loops in his work on natural selection within the context of the impact of earthworms (which was often mocked by contemporary cartoonists) on soils and grass (a). Such closed cooperative loops are widespread in ecosystems and involve several species (b) and have been formalized in terms of "Hypercycles", as indicated in (c). These catalytic systems are prone to collapse under parasites (d), which might have driven the evolution of complexity.

and mechanisms of soil preservation (favouring the earthworm population). This positive feedback can be sketched in terms of a simple, closed cooperative loop as shown in 3.1b. But the loops can be more complex, as we outline in 3.1c. Here the vegetation is grazed by animals, whose activity enhances the survival of invertebrates, improving soil quality, thus favouring plant growth. In general, ecosystems are characterised by multiple feedback loops, and therefore interactions might be more complex [Bro15].

With Darwin's grand vision came a picture of evolution where species would be connected through a divergent tree to be rooted in a common ancestor. Some new branches appear as speciation occurs, while others are

terminated as species become extinct. But some significant events do not fit so well this tree-like pattern. As mentioned in Chapter 1, one is the emergence of complex cells, which required the cooperative (symbiotic) mixing of two previously independent organisms [Sag67, Mar81]. Scholars such as Merezhkovsky (Figure 3.2a) had previously hinted at this idea (i.e. that the so-called tree of life was rather a meshed network instead), but it was Lynn Margulis (Figure 3.2b) who definitively provided robust and solid evidence for symbiosis as a cooperative explanation for the origin of eukaryotic cells. The problem of cooperation is of course, much wider than this. Cooperation needs to be achieved again when multicellularity first emerges, and it is also a fundamental component of the success of humans in evolution. Not surprisingly, it has also been a very active area within the social sciences. Here the problem is how cooperation among selfish agents might emerge stably. Such problems can be formulated in mathematical terms under the formalism of evolutionary game theory, for which Maynard Smith, Robert Axelrod (Figure (Figure 3.2c) and Peter Schuster (among other scholars) contributed to its building.

An elegant description of this class of cooperative loops is the *hypercycle*, first suggested within the context of prebiotic evolution [ES77, ES78a, ES78b, Kau71, Kau86, SS97, Sza06, HL15, Sch16]. Here a simple catalytic system is defined forming a closed graph where the replication of each component is catalysed by a previous one in the loop. At the same time, it also catalyses the replication of the next. The simplest case is shown in Figure 3.1c for a two-member system [ES12, SS06]. If we indicate by Φ_1 and Φ_2 their population sizes, a pair of coupled equations allows us to represent the hypercycle model as follows:

$$\begin{aligned} \frac{d\Phi_1}{dt} &= \alpha_{12}\Phi_1\Phi_2 \left(1 - \frac{\Phi_1 + \Phi_2}{K}\right) - \delta_1\Phi_1 \\ \frac{d\Phi_2}{dt} &= \alpha_{21}\Phi_1\Phi_2 \left(1 - \frac{\Phi_1 + \Phi_2}{K}\right) - \delta_2\Phi_2 \end{aligned} \tag{3.1}$$

where K stands for the system's carrying capacity, δ is the degradation/death rate of both species, and the replication rates of the cross-catalytic loop are indicated by $\alpha_{ij} > 0$. Thus, due to the second-order



Figure 3.2: **Merezhkovsky, Margulis and Axelrod.** Three major players in the development of a theory of cooperation and evolution: Konstantin Merezhkovsky, Lynn Margulis and Robert Axelrod. The first two contributed to develop the modern theory of symbiosis while Axelrod provided the basis for a game-theoretic approach to the transition to cooperation.

kinetics, no proliferation of any of the two partners will occur without the other.

The hypercycle is closely related to the hyperbolic replicator discussed in the previous chapter. It can out-compete other non-cooperative species [ES12, Sza06]. Still, a major drawback is that it can also be easily threatened by a parasite (Figure 3.1d) capable of destabilising the whole system [SS79]. Interestingly, mathematical and computer models indicate that this problem can be limited by diffusion in a spatial domain¹ [BBN96, BH95, AS06, SS07].

Cooperation is formally describable as one kind of nonlinear interaction that reveals several universal properties. Although much has been developed regarding the statistical physics of cooperation, some impor-

¹Hypercycles displaying spatial structures (Figure 3.1f) are obtained from $n > 4$ loops capable of exhibiting oscillations. In a nutshell, the spatial structure imposes a limitation to the spread of the parasite, which can even go extinct if the inaccessibility of its target species, combined with its death rate, makes it non-viable [NGM91].




Figure 3.3: *In silico* and *in vivo* examples of cooperation. Three examples of systems displaying pre-multicellular or proto-organismal properties are based on the presence of competitive and cooperative interactions. The first two (a-b) are emergent structures resembling organismal entities that emerge from simple, indirect cooperation [DNBM⁺16]. The third (c) is a close up to a gut microbiome, formed by an ecological assembly of diverse species that share and process together a collective metabolic network.

tant problems remain. One concerns the pre-conditions associated with the rise of cooperative assemblies. In this context, it has been shown that some essential traits related to the emergence of multicellularity might have predated this transition in its pre-genetic era. This is the case, for example, of the self-organization of *proto-organisms* [DNS15] or the multicellular-like organization of some microbiome assemblies (figure 3.3).

Evolved cooperation between two classes of cells emerges under a suitable set of environmental conditions associated with stress [OVS19]. In these cases, we do not have actual multicellular organisms. Still, the emergence of some traits tied to *organismality* is already there, from the division of labour between different cell types to a spatial patterning that allows separating an internal medium from the external world. Perhaps the crucial idea emerging from our growing understanding is the understanding of organisms as *holobionts*, i.e. organisms made of a predictable ensemble of interacting or simply coexisting species. As pointed in [GST12]:

animals can no longer be considered individuals in any sense of classical biology: anatomical, developmental, physiological, immunological, genetic, or evolutionary. Our bodies must be understood as holobionts whose anatomical, physiological, immunological, and developmental functions evolved in shared relationships of different species. Thus, the holobiont, with its integrated community of species, becomes a unit of natural selection whose evolutionary mechanisms suggest complexity hitherto largely unexplored.

A diverse range of models suggests that a complex landscape of organization pervades true multicellular complexity. What is the minimal class of model we can define that captures such pre-conditions? In other words, what is the minimal model for a (stochastic) holobiont? The path followed in this dissertation is to present and comprehensively analyse the simplest, and yet significant case scenario: a stochastic, neutral model of cooperators. In this case study, a population of cooperative species is considered with no predefined phenotypic heterogeneity. Instead, all species are equal and compete for a limited space. Surprisingly, a distinct, coherent (and stable over long times) core of cooperators is shown to spontaneously emerge. This finding reveals that even under this minimal set of assumptions and considering stochastic neutral rules of interaction, a systems-level order beyond the species level is found.

3.1. Diversity and complexity

The approach taken here requires several well-established ideas and measures from population ecology. We will need them to present our theoretical framework. Despite their broad range of types, ecosystems across the biosphere share key quantitative similarities², we will distinguish non-spatially and spatially explicit cases. Examples of non-spatial patterns are: (1) The *diversity*, sometimes described as the number of species S

²For a theoretical introduction to the diversity patterns employed in the field of Ecology, see [Mag88] and Table I in [ASG⁺16] for a shortlist.

present at a given time, but also using weighted quantities such as Simpson's index (D) or a variation of Shannon's entropy (H) (see Box 3.1). (2) The *relative species-abundance* (RSA) patterns, i.e. the probability that a given species is represented by a fraction x out of the total number of individuals of all species inhabiting a given region; (3) The *species turnover distribution* (STD), that weights the variation over time of the abundance of a given species within the larger interacting community and (4) The *lifetime distribution* (LD), which characterizes the time between first appearance and extinction of a given species.

BOX 3.1 - Diversity indexes

Introduced by Edward H. Simpson in 1949 [Sim49], the idea behind the so-called Simpson's index (D) is to go a step beyond a simple 'counting' of the number of species within an ecosystem and introduce some sense of *abundance* equality (or lack thereof). Shortly after, and with the advent of Shannon's Theory of Information, the Shannon-Wiener index to measure ecological diversity was proposed (H). Suppose that a given ecosystem contains S , whose relative abundances are given by $\{x_i\}_{i \in \{1, \dots, S\}}$. Then define

$$D = \sum_{i=1}^S x_i^2, \quad H = - \sum_{i=1}^S x_i \ln x_i \quad (3.2)$$

We notice that $D \in [1/S, 1]$, where the lower bound is achieved when $x_i = 1/S, \forall i$ (equidistribution) and the upper bound when $\exists i^*$ such that $x_{i^*} = 1$ and $x_{i \neq i^*} = 0$ (exclusion of $S - 1$ species or maximal inequality). Some texts use the convention $D_s = 1 - D$. On the other hand, $H \in [0, \ln S]$, with $H = 0$ when maximal inequality is reached and $H = \ln S$ for equidistribution. A thorough introduction to these topics and related is provided by Magurran (2003) [Mag03].

Spatial patterns often involve more sophisticated definitions since they require two characteristic scales to be considered: spatial and temporal. A usual example is the *species-area relationship* (SAR), which studies the

number of species inhabiting an area $A \subset \mathcal{R}$ where here \mathcal{R} represents the whole of the area occupied by the ecosystem. SAR typically displays a power-law of the form $S(A) = cA^z$, where S here is the observed number of species inside of A and z is a characteristic exponent that takes values around $z \sim 0.2 - 0.3$. In this Thesis, we will only focus on studying non-spatial patterns. Finally, in this text, we will deal with ecosystems of *fixed* size, i.e. systems in which the total number of individuals populating it is fixed *a priori*. Note that this does not imply the system is necessarily closed (e.g. migration processes can occur via a replacement rule).

The quantitative approach led by ecologists has pushed the field into developing an impressive theoretic corpus whose core philosophy is to unveil the underlying mechanisms that give rise to such universal patterns. From the physics perspective, two observations are to be made at this point:

- (I) The dynamics underlying these systems differ from usual Hamiltonian dynamics studied in statistical physics insofar as evolutionary processes are present (selective and non-selective forces, extinction, absorbing points, etc.).
- (II) Whichever microscopic mechanisms one may concoct, statistical physics teaches us that not all of these will be relevant for the emergence of large-scale patterns. Thus, it is our task to distil the minimal set of interactions that conduce to the macroscopic behaviour of the system that is relevant at our observational scale.

3.1.1. Neutral Theory of Ecology

In 1967, MacArthur and Wilson (Figure 3.4) published “The theory of island biogeography”, which constituted a pioneering attempt to explain the macroecological patterns observed by ecologists from a stochastic process viewpoint [MW01]. Particularly, they assessed the balance of immigration and extinction within extant metacommunities on an island that is close to a ‘mainland’ (M), which acts as a species reservoir. They showed that taken as purely stochastic processes, the aforementioned rates




Figure 3.4: **MacArthur, Wilson and Hubbell.** Left to right: Robert H. MacArthur (1930-1972), Edward O. Wilson (1929-2021), Stephen Hubbell (1942-). MacArthur and Wilson introduced the concept of Island Biogeography in 1967 [MW01] as a pioneering approach to a stochastic-based theory of ecology. In 2001, Hubbell developed the Neutral theory of Ecology [Hub11], drawing a parallel to Kimura’s Neutral theory of Evolution [Kim83].

construed the shape of the ecosystem diversity in dependence on parameters such as distance from mainland and neighbouring islands, island size and human intervention (see scheme in Figure 3.5a).

The work by MacArthur and Wilson influenced ecological thinking for generations and eventually produced a schism in the ecology community between the so-called *Niche theory* and what later became known as the *Neutral theory* (NT) of ecology. In a nutshell, the former school argues that the principle of competitive exclusion and resource-consumer dynamics shape the structure and patterns of ecosystems, while the latter assumes a top-down approximation in which the details pertaining to the microscopic (individual-to-individual) interactions are coarse-grained and the system is analysed through the methods of stochastic processes. In short, the goal of the NT of ecology is to derive testable expressions for the observed macroecological patterns outlined in Section 3.1.

BOX 3.1.1 - Neutral theory of *competitors*

It is possible to derive the RSA of a fixed-size neutral ecosystem (urn) with S species and N individuals that is coupled to an external species reservoir (mainland, \mathcal{M}). Consider an auxiliary parameter $\lambda \in (0, 1]$ and introduce the following iterative rules (Figure 3.5b):

1. With probability $1 - \lambda$: a first individual is chosen at random, then a second one: if their colours do not match, then the second pick is changed by the colour of the first. If the colours match, nothing changes. Then the two particles are put back into the urn.
2. With probability λ : an individual is selected at random from \mathcal{M} and replaces another (randomly selected) individual from the urn.

If the mainland species are equidistributed, then the system is fully neutral (competitive equivalence). It is possible to solve this problem exactly in the discrete [MAS00]. An approximate solution in the continuum for the stationary regime is easy to derive using a Fokker-Planck approximation. Due to neutrality, any given species can act as the representative of the whole ensemble. This reduces the problem to a single variable $x = n/N \in [0, 1]$ which, for $N \gg 1$, evolves under transition rates

$$\begin{aligned}\omega(x + \Delta x|x) &= (1 - \lambda)x(1 - x) + \lambda(1 - x)/S, \\ \omega(x - \Delta x|x) &= (1 - \lambda)x(1 - x) + \lambda(1 - 1/S)x.\end{aligned}$$

The steady-state distribution, i.e. the RSA, becomes:

$$P_{\text{st.}}(x) \sim x^{-\beta} e^{-2N\lambda x}, \quad \beta = 1 - 2\lambda(N/S).$$

In 2001, Stephen Hubbell, who was deeply inspired by MacArthur and Wilson's theory, published "The Unified Neutral Theory of Biodiversity and Biogeography" [Hub11]. Here, Hubbell focuses on monotrophic communities, e.g. coral reefs, birds or plants in a forest; and makes the

following assumptions to proceed:

- a) Competitive equivalence b/w interacting species (*neutral hypothesis*): at the dynamical scales considered, all species display the same fitness value, i.e. compete with the same effectiveness.
- b) Stochasticity at the individual level: birth-death processes, replication, replacement, or random drift dominate the dynamics of the ecosystem and its assembly.
- c) Predictive macroscopic patterns from theory: the theoretical output is based on *global* measures such as diversity, relative species-abundance or SAR.

The vision of neutral theory was met with staunch criticism, considered unorthodox and often dubbed unrealistic [Ric06] (see [ASG⁺16] pp. 3-4 and references therein). Others embraced these ideas, developed the theory [MAS00, SAM02, MAS04] and probed it via serious empirical tests [VBHM03, AEM06]. The theoretical precepts of Hubbell and others later came to be known as the Neutral theory of ecology. In this Thesis, we base our work on these assumptions.

Finally, from a the Physics viewpoint, it is necessary to highlight how the NT is much closer to the philosophical basis provided in (i) and (ii) above. On one hand, the microscopic dynamics is now reduced to simpler processes laid down in (a,b), thus losing the deterministic handle on the system but –in some cases– gaining solvability. On the other hand, a coarse-graining is necessarily imposed on the system in order to obtain (c); this puts the focus on the large-scale statistical patterns, rather than in the microscale details. The latter point brings the NT closer to nonequilibrium thermodynamics of ecology.

3.2. Back to Cooperation

Our main research question for this Chapter is the development of a Neutral theory of Cooperative Ecosystems. As we will argue, such a theory is not merely an extension of the classic NT, but it actually involves




Figure 3.5: Diagrammatic descriptions of Island Biogeography and Hubbell's Neutral Theory. (a,b) Diagram and abstraction of the Theory of Island Biogeography [MW01]. Arrows denote migration from mainland (M) to an island or urn (I, II or III) or between islands (e.g. $I \rightarrow II$). Each color correspond to a species. (b,c) Abstraction of the replacement rule in Hubbell's NT of ecology.

a whole new set of methods and ideas that imbue this problem with an idiosyncrasy of its own. The aim is thus to arrive at a top-down (macroscopic) analysis of cooperation in large sets of interactive agents (cooperators); study their feasibility, stability and global features, while emphasizing the role of randomness, coexistence, turnover and other large-scale patterns.

We will restrict our analysis to fixed-size ecosystems, i.e., ecosystems that contain a fixed number of individuals (N). We will consider a minimal model for which only two species of cooperators coexists, deduce its principal stochastic properties and carry over several ideas when moving forward to the study of the general multi-species case. In this Thesis, we will only focus on the non-spatial properties of such systems, with emphasis on patterns such as relative abundance distributions (RSA), turnover (STD) and extinction patterns (LD) (see Section 3.1). Extension to spatially constraint or fluctuating-size cooperative ecosystems are left for future work.

3.2.1. Two-Species Cooperator Systems

In an experimental set-up due Müller *et al.* (2014) [MNNM14], the authors showed how a synthetically engineered strain of bacteria, in which

mutualism is forced via metabolic interaction between two mutants (bio-marked in blue and yellow, respectively), reveals astonishing macroscopic growth pattern differences (see Figure 3.6). In particular, in the top panel of the Figure 3.6a, the system is strongly affected by competitive exclusion, hence the radial color stripes which are separated at some characteristic angular scale. On the other hand, the mutualist-pair type evolution, shown in the lower panel of Figure 3.6a, breaks down the radial pattern displayed by the above and facilitates coexistence: proximity between species, and equidistribution is favoured.




Figure 3.6: Competitors vs. cooperators I. In (a) we show two experiments of competitive and cooperative bacterial populations growing on a Petri dish. Competitive exclusion gives rise to single-species radial stripes, whereas engineered mutualism (b) gives a mixed pattern, with local coexistence (from [MNNM14]). (c) Our model of this system considers an urn with two types of cooperator species. Replication occurs only when two randomly selected individuals are one from each of the strains.

To advance in our proposal for a NT of a holobiont dominated by cooperators, it is necessary to consider this class of systems in which two species of cooperators coexists in a fixed-size system while, at the same time, are subject to stochastic noise and random migration fluxes. This has been the object of our paper Piñero, Redner & Solé (2022) [PRS22], which main findings we will followingly outline.

Closed system

Let us begin by considering the closed system. Suppose that N individuals are placed in an urn (see Figure 3.6c) each belonging to species A or species B . The iterative rules hereby imposed are:

1. Pick a random pair of individuals.
2. If the pair is of the type AB , one member of the pair reproduces. Otherwise, nothing changes.
3. If 2. was successful, then the respective offspring replaces one random individual from the urn.

Sections 2, 3 and 4 in [PRS22] provide an extensive set of results concerning the problem of closed systems. Let us next recover and discuss a few of them. Since the system size is fixed, denoting by n the number of individuals of type A readily gives the same for B as $N - n$. We also note that, as per the system defined by rules 1-3, the distribution is equivalent for either A or B , hence, this two-species system is both neutral and cooperative. At the same time, due to the replacement rule 3, there is an intrinsic competition for “space” in the urn. Eventually, the system will fall under the absorbing states $n = N$ or $n = 0$. Statistical properties of these two absorbing points (e.g. exit probability or average first passage time) are identical due to symmetry.

From the deterministic view, and ignoring absorbing points for the moment, the system contains a single fixed point at $n = N/2$. Moreover, fluctuations around this fixed point are of the order of \sqrt{N} (see Eq. (6) in [PRS22]). Solely the fact that there exists a fixed point around which some fluctuations take place represents a strong departure from the competitive counterpart. It is easy to show that for the latter no fixed point exists, and only the usual absorbing points are present. In particular, the size of the fluctuations for the competitive counterpart are population-dependent³, namely $\text{Var}(n) \simeq N\langle n \rangle (1 - \langle n \rangle)$. Thus, the competitor

³This result can be derived from the example shown in Box 3.1.1 –and imposing $\lambda = 0$ – or simply considering equation (5b) in page 6 of [PRS22] where one substitutes $a_n = b_n = (n/N)(1 - n/N)$.

two-species system will display a *modulated* (a.k.a. demographic) noise, whilst the cooperators fluctuations are strongly constrained around its fixed point. This lesson will have to be carried over in the following steps when considering systems of $S > 2$ cooperators.

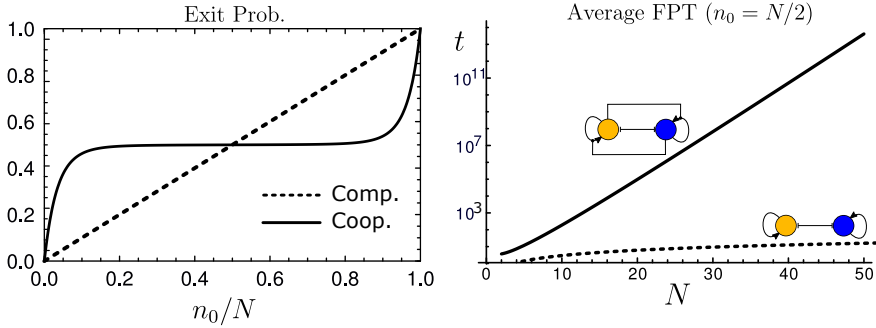


Figure 3.7: **Competitors vs. cooperators II**: a comparison of analytical expressions for either type of two-species ecosystem for (left) the exit probability, i.e. the probability that a given species takes over the system given an initial condition, computed for $N = 2^4$; and the average first-passage time (in logarithmic scale), i.e. the average time to reach fixation at either $n = 0$ or $n = N$ with initial condition $n_0 = N/2$. Dashed/straight lines correspond to competitors/cooperators cases.

The previous observations regarding the fluctuation behavior of either system are also reflected in the patterns of exit probability (i.e. the probability that a species occupies the whole system given its initial condition) and the mean fixation time (i.e. the average time for a species to reach either absorbing point $n = 0, N$). This implies that the fixation statistics in either cases (competitors vs. cooperators) will differ substantially. In particular, taking units of time to be such that every N iterations of rules (1-3) then time is increased by 1, i.e. time increases $\delta t = 1/N$ for every iteration, then it is possible to compute the average fixation time in each case (neutral competitors/neutral cooperators) assuming that the

initial condition is $n_0 = N/2$,

$$T_{\text{comp.}} \left(n_0 = \frac{N}{2} \right) \simeq \frac{N}{4}, \quad T_{\text{coop.}} \left(n_0 = \frac{N}{2} \right) \simeq \frac{\sqrt{8\pi}}{Ne} 2^N. \quad (3.3)$$

The second result in (3.3) is obtained in Section 4 in [PRS22] using the WKB method (see Box 3.2.1). We immediately note that $T_{\text{comp.}}$ and $T_{\text{coop.}}$ differ enormously in their respective scalings: the former is linear with the system size while the latter scales exponentially with N .

The reason behind the necessity of WKB is precisely another remarkable feature of the presence of a fixed point in cooperative ecosystems. As discussed above, the fluctuations around this point are of order \sqrt{N} but, if $n_0 \sim N/2$, then the deviation that needs to take place in order to reach fixation is of size $N \gg \sqrt{N}$ (at the large N limit). This implies that, at initial times, i.e. before fixation, the probability distribution is peaked around $n_0 = N/2$. If we denote this pseudo-steady distribution by $\tilde{P}(x)$ (where $x = n/N$), and since it resembles a delta-function, it will be well represented by an exponential function of the type

$$\tilde{P}(x) \sim e^{Ns(x)}$$

As discussed in Box 3.2.1, this requires giving up the usual Fokker-Planck approximation (which involves a perturbation of P around x), for an approach which studies the perturbations of the function $s(x)$ instead; i.e. the WKB method. This is another important lesson to carry over into the extension to many-species cooperator ecosystems, where we expect a similar phenomena to occur.

Open system

Finally, including an immigration component into the two-species cooperator system is explored in Section 5 of [PRS22], where we showed that migration is a strong system stabilizer⁴ In particular, when $\lambda > 1/(N + 1)$

⁴We refer the reader to this section for an exhaustive analysis.

BOX 3.2.1 - The WKB Method

Named after Wentzel, Kramers and Brillouin; the WKB method was many times independently discovered over history. The earliest known case is due to the italian astronomer Francesco Carlini (1817) [Sac20]. The method reached ample fame as an approach to solve wave-function problems in Quantum Mechanics.

In Stochastic Processes, we commonly use the Fokker-Planck (FP) approximation to derive a PDE for our distribution $P(x, t)$. Within the FP equation lays a perturbative assumption, namely that $\Delta P \simeq P' \Delta x + \dot{P} \Delta x / D + \frac{1}{2} P'' (\Delta x)^2 + \dots$, where $D = (\Delta x)^2 / \Delta t$ is a finite constant. However, if dealing with systems of fixed number of particles, $N \gg 1$, it is common to have

$$P \sim e^{Ns(x,t)}, \quad (3.4)$$

where $x = n/N$ is the fraction of the number of individuals (n) of a certain species, and $\Delta x = 1/N$. For such a class of systems, the aforementioned expansion breaks down completely. Heuristically, this is because a small deviation in x now yields an exponential shift on P , due to the nature of a large exponential argument. This is a common phenomenon in the theory of large deviations [AM17]. In this context, WKB arises as a successful approximation by letting

$$s(x, t) = s_0(x, t) + s_1(x, t) \Delta x + O(\Delta x^2). \quad (3.5)$$

For a one-dimensional markovian process with transition rates $\omega(x - \Delta x | x) =: \omega_-(x)$ and $\omega(x + \Delta x | x) =: \omega_+(x)$, the stationary solution for $P_{\text{st.}}(x)$ through $s^{\text{st.}}(x)$ is (see Appendix C of [PRS22]):

$$s_0^{\text{st.}}(x) = \int^x dx' \frac{\omega_+(x')}{\omega_-(x')}, \quad s_1^{\text{st.}}(x) = -\frac{1}{2} \ln [\omega_+(x) \omega_-(x)]. \quad (3.6)$$

the system becomes unimodal around its fixed point, i.e. equidistribution

is guaranteed. This stabilization role is expected to follow qualitatively when considering the S -species system.

3.3. Multi-Species Cooperator Systems

In this section, we will discuss ongoing work which, at the moment of submission of this Thesis, it has not yet been published. For the purposes of completion, we will outline a series of preliminary results and discuss both our current work and possible future directions.

In order to approach a theoretical model for systems of cooperators, we first consider the general deterministic expression (belonging to a Lotka-Volterra class of equations)

$$\frac{dn_k}{dt} = \alpha_k n_k + \sum_{l=1}^S \frac{\gamma_{kl}}{N} n_k n_l - n_k \Phi(\mathbf{n}), \quad (3.7)$$

with $k \in \{1, \dots, S\}$ and $\sum_{k=1}^S n_k = N$, which is imposed via

$$\Phi(\mathbf{n}) = \sum_k n_k \left(\alpha_k + \sum_{l=1}^S \frac{\gamma_{kl}}{N} n_k n_l / N \right)$$

Assuming neutrality corresponds to setting

$$\alpha_k = \alpha, \quad \gamma_{kl} = \gamma, \quad \forall k, l \text{ with } k \neq l \text{ and } \gamma_{kk} = 0 \quad (3.8)$$

Classic NT is recovered by turning off all the second order interactions, i.e. $\gamma = 0$. However, motivated by the questions posed above, here we will focus our efforts on studying the properties of a fully cooperative neutral ecosystem, by taking (3.8) and assigning $\alpha = 0$ and $\gamma > 0$.

3.3.1. Deterministic system

Following an equivalent procedure as before, it is instructive to begin by studying the deterministic set of equations resulting after applying the

assumptions above to obtain

$$\frac{dn_k}{dt} = \frac{1}{N} \sum_{\substack{l=1 \\ l \neq k}}^S n_k n_l \left(1 - \frac{1}{N} \sum_{\substack{m=1 \\ m \neq l}}^S n_m \right) = n_k \left[\mathcal{D}(\mathbf{n}) - \left(\frac{n_k}{N} \right) \right], \quad (3.9)$$

where we have shifted time as $t \rightarrow \gamma t$ and introduced the Simpson's functional as

$$\mathcal{D}(\mathbf{n}) = \sum_{k=1}^S \left(\frac{n_k}{N} \right)^2 = \frac{\mathbf{n} \cdot \mathbf{n}}{N^2}.$$

Note that this functional comes out of the treatment of equations (3.7) with only the assumptions concerning neutrality and cooperative interactions. As discussed in Section 3.1, the Simpson's index is bound within the interval $[1/S, 1]$ and measures the diversity of the system, with $\mathcal{D} = 1/S$ corresponding to equidistribution. Next, it is possible to show that equation (3.9) contains $2^S - 1$ fixed points and S true absorbing points. To make sense of this statement, it is practical to turn to a geometrical visualization.

The manifold that represents the closed ecosystem containing S species is given by the S -simplex, which we denote by \mathcal{M}_S which is an $(S-1)$ -dimensional space defined by

$$\mathcal{M}_S = \left\{ \mathbf{x} \in [0, 1]^S \mid \sum_k x_k = 1 \right\},$$

where we defined $x_k = n_k/N$ for all k . Now, the simplest of all the 2^S fixed points is given by equidistribution among all S species, i.e. $\mathbf{x}^* = (1/S)\mathbf{1}^T$. This results immediately from (3.9), since the Simpson's functional takes the value $\mathcal{D} = 1/S$, and thus the right parenthesis in (3.9) vanishes for all k . Assume a system that is set to begin evolving according to some stochastic process from a state close (or equal to) the fixed point $\mathbf{n}^* \leftrightarrow \mathbf{x}^* \in \mathcal{M}_S$. If the evolution is stochastic (see below), then some deviation will push the system into the boundary of the S -simplex, $\partial\mathcal{M}_S$. Whenever this happens, then the system *jumps* to a

one lesser dimensional simplex, \mathcal{M}_{S-1} , with one of the dimensions now turned off, i.e. $x_{k'} = 0$ for some k' . The surviving (non-zero) species will continue to evolve over time on a reduced system that will contain another fixed point, namely $n_{k \neq k'}^* = N/(S-1)$ and $n_{k'}^* = 0$. See Figure 3.8. This process, which we term as the *extinction* of a single species, will proceed onwards for the stochastic closed neutral system of cooperators until fixation occurs for a given species, i.e. one species takes all. Observe that, while we clearly have that

$$\mathcal{M}_1 \subset \cdots \subset \mathcal{M}_{S-1} \subset \mathcal{M}_S,$$

there are $\binom{S}{i}$ possible \mathcal{M}_{S-i} manifolds with $i = 0, \dots, S-1$, each of which has a fixed point at its respective center. These add up to a total of $2^S - 1$ fixed points. Importantly, each of these fixed points is only truly attractive when restricting the system into its corresponding submanifold. In other words, the attractiveness of all these fixed points remains ‘hidden’ until the system’s trajectory hits its corresponding submanifold $\mathcal{M}_{1 \leq r \leq S}$. On the other hand, the system can only be finally absorbed (fixed) into S different \mathcal{M}_1 simplexes, one for each species.

Before moving on to a full stochastic version of our neutral cooperator ecosystem dynamics, let us study a generalized potential function as perceived by any representative species on the ecosystem. Recall that, owing to neutrality, any species equally represents the dynamics of any other. If working on the continuum, i.e. for $x_k = n_k/N$ for all k , then from (3.9), we read:

$$\frac{dx_k}{dt} = x_k [\mathcal{D}(\mathbf{x}) - x_k], \quad (3.10)$$

with now $\mathcal{D}(\mathbf{x}) = \mathbf{x} \cdot \mathbf{x}$. Equation (3.10) bares remarkable similarities with the replicators equation [Sch11], where $\mathcal{D}(\mathbf{x})$ would correspond to a global (mean-field) positive feedback, and the remaining term between brackets acts as a regulator that pushes all species to maintain diversity. In dynamical systems jargon, $x_k = 0$ (extinction) and $x_k = \mathcal{D}(\mathbf{x})$ are the two fixed point equations. Note that, if the remaining species are such that $x_k = 1/R$ for some $k = 1, \dots, R \leq S$ and $x_{k' \neq k} = 0$ for the rest




Figure 3.8: **Attractors in the simplex representation.** (a) Shows the \mathcal{M}_4 simplex where A, B, C and D denote each of the $S = 4$ species, respectively. Each corner represents fixation (full occupancy) of every respective species, pairs such as AB represent coexistence between only two species, thrices such as ABC , coexistence between three species, and $ABCD$ is the central fixed point corresponding to equidistribution. (b) Depicts two stochastic trajectories, one occurring at the bulk of \mathcal{M}_4 and another taking place after the *extinction* or collapse of species A , which constrains the trajectory at the \mathcal{M}_3 with BCD as a center attractor.

of the species, then we achieve the latter fixed point; this represents all central points of every simplex \mathcal{M}_R . In particular, $R = 1$ correspond to the S possible fixation points (vertices of \mathcal{M}_S). Now, consider, the quasi-potential function:

$$\mathcal{V}(\mathbf{x}) = -\frac{1}{4} [\mathcal{D}(\mathbf{x})]^2 + \frac{1}{3} \sum_{k'} x_{k'}^3, \quad (3.11)$$

then it is easy to show that

$$\frac{dx_k}{dt} = -\frac{\partial \mathcal{V}(\mathbf{x})}{\partial x_k}.$$

Which fits the conventional definition of a potential function. Representing this function is rather impractical, since its dimension is very large.

However, if we situate the system in some \mathcal{M}_R , with $1 \leq R \leq S$, then, for a given representative species, namely $\tilde{x} \neq 0$, the rest of the non-zero densities will approximately take values around $x \sim 1/R$, which gives $\mathcal{D}(\mathbf{x}) \approx 1/R$. Under this approximation, by taking the equation for \tilde{x} in (3.10), with $\tilde{x} > 0$, we can rederive a single-variable generalized potential, namely

$$\mathcal{V}(\tilde{x}) \approx -\frac{\tilde{x}^2}{2R} + \frac{\tilde{x}^3}{3} = -\frac{1}{6R^3} + \frac{y^2}{2R} + \frac{y^3}{3}, \quad (3.12)$$

where $y := \tilde{x} - (1/R)$, i.e. a shift to the fixed point. Expression (3.12) is a very rough approximation, since it acts as if blocking all degrees of freedom but \tilde{x} , which is highly unrealistic. However, it gives us a perspective on how the stochastic counterpart to our deterministic problem will behave. For example, from (3.12) we read, from the quadratic term, a *Hooke* coefficient $k = 1/R$. This means that we should expect fluctuations around $\langle y \rangle = 0$ (which correspond to $\langle \tilde{x} \rangle = 1/R$) to be of order

$$\sigma_y = \sqrt{\langle y^2 \rangle} \sim \sqrt{\frac{R}{N}}. \quad (3.13)$$

Although a very rough estimate, expression (3.16) implies that every time an extinction occurs, i.e. $R \rightarrow R - 1$, then the fluctuations decrease by $\delta\sigma_y(\mathcal{M}_R \rightarrow \mathcal{M}_{R-1}) \approx -1/2\sqrt{RN}$. This result implies a slow-down effect as extinctions occur over time, i.e. the smaller the number of remaining species (R) the smaller the fluctuation size around the simplex's center. All of the above results can be stated without involving a full stochastic model of the dynamics of the neutral cooperators ecosystem. However, in order to assess properties such as extinction times in each simplex, \mathcal{M}_R , or a more wide description of the species-abundance distribution (RSA), then a stochastic extension is required.

3.3.2. Stochastic system

Having learned some of our system's elementary properties from the deterministic approach, we now turn to its stochastic counterpart. In this

This we will introduce an iterative model and explore it via computer simulation. Then we will outline the difficulties entailing a respective theoretical treatment of the probability distributions and enumerate which approximations are and are not successful in providing analytical insight on the problem.

BOX 3.3.2 - Fluctuations of a simple oscillator

A lot in Physics is based in exploiting the niceities of the oscillator dynamics. Consider a simple harmonic oscillator with Hooke constant k under thermal noise, $\beta = 1/k_B T$. Its hamiltonian is simply given by

$$H(p, x) = \frac{p^2}{2m} + \frac{1}{2}kx^2.$$

Using the Gibbs measure, $P(p, x) \sim e^{-\beta H(p, x)}$, we study the marginal distribution for $P(x) = \int dp P(p, x)$, which gives

$$P(x) = \frac{1}{Z} e^{-\frac{\beta}{2}kx^2}, \quad Z = \int dq e^{-\frac{\beta}{2}kq^2} = \sqrt{\frac{\beta k}{2\pi}}. \quad (3.14)$$

Thus, $\sigma_x = 1/\sqrt{\beta k}$. However, for systems under demographic noise the common scaling for the distribution is coupled not with an inverse temperature but with the system size, i.e. $\beta \rightarrow N$. Here, $P(x) \sim e^{-Ns(x)}$, where $s(x)$ is a generalized (stochastic) potential function. If $s(x)$ contains a second order term with coupling $k/2$, then the fluctuations will scale as

$$\sigma_x \sim \frac{1}{\sqrt{kN}}.$$

Consider a finite-sized urn with a fixed population of N individuals from S species (colors) which are all cooperators (as followingly defined). Let us consider the following iterative rules (see Figure 3.9):

1. Pick a random pair of individuals.
2. If this pair belong to distinct species, then randomly replicate one of

them.

3. Return the initial pair to the urn, and, if in 2 there was a newly created individual, then return it to the urn by replacing a randomly chosen element from it.




Figure 3.9: **Closed neutral cooperators system.** Simple depiction of the algorithm for closed neutral cooperative ecosystems. Here the urn represents the finite-sized system with each color representing a different species. Only species of distinct colors can interact in pairs to produce offspring, which, if successful, will replace a random element from the urn. Every iteration, $\delta t = 1/N$ units of time are consumed.

We choose our units of time such that, for each iteration of the algorithm above, then t is incremented by $\delta t = 1/N$. Heuristically, since all species but one random species will eventually become extinct, it is clear that the closed system will follow a trajectory that has ‘memory’ on which species have previously died out and finished up ‘choosing’ a single species. However, as we have established in Section 3.2.1, this process will take in exponential time in N . In fact, simulating this algorithm gives rise to a non-trivial set of trajectories towards extinction that shift the central point of attraction (as predicted by the analysis above). Denote by T_R the average time between an extinctions $\mathcal{M}_{R+1} \rightarrow \mathcal{M}_R$ and $\mathcal{M}_R \rightarrow \mathcal{M}_{R-1}$, and study look at the trajectories of $R(T)$, with every T_R (for each trajectory) marking an extinction event, which corresponds to a unit jump downwards of the graph given in Figure 3.10.

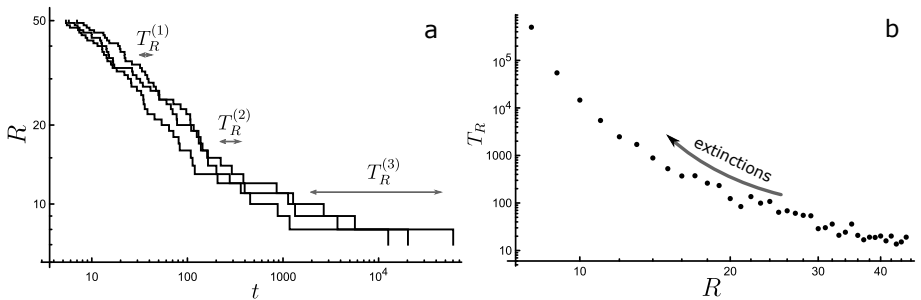


Figure 3.10: Extinction dynamics. (a) In log-log scales, three trajectories of the coarse-grained state of $R(t)$ for $N = 10^3$ and $S = 50$. Note that $T_R^{(1)} \ll T_R^{(2)} \ll T_R^{(3)}$, which means that extinction times grow faster than exponentially with R . (b) Also in log-log scales, shows the average extinction times from $R \rightarrow R - 1$, i.e. T_R , as a function of R , for $N = 10^3$ and $S = 50$. We note the superexponential increase in average extinction time at the left tailing end of the graph, while the noisy fluctuations at the right end are only due to rapid extinction rates happening at larger R .

Thus, we observe how the relaxation times become longer at each sub-simplex, while the quasi-steady state distribution is stabilized for longer periods of time, that is, until an extinction point is reached and the system chooses a next (smaller dimensional) simplex⁵. This is not too surprising given that, heuristically, we have shown that in each subsimplex the fluctuations decay (3.16). These observations lead us to conclude that the closed ecosystem behaves as a *glass* [SN13]. The study of glasses (or spin glasses) constitute today an enormous field of research within the statistical physics community with multiple ramifications into Complex Systems, Ecology, Evolutionary Theory and more (see [SN13] for an introduction to spin glass theory and [PUZ20] for a technical primer). In particular, here we are interested in the glassy features of: long relaxation times and contingent trajectories.

An immediate shortcoming from the fact that the system is glassy is that its high level of interactions renders it hard to study analytically.

⁵Note that, from our simple algorithm, transitions corresponding to double extinctions cannot occur in a single time-step.

In the first place, the application of the Fokker-Planck approximation to obtain some estimation on the quasi-steady state distributions emerging in between extinction events becomes too rough. The reason is that the strong memory dependence of the trajectories (either in full population density vector space or in the coarse-grained remaining species variable space) makes the usual *markovian* assumptions behind equations of the Fokker-Planck type irredeemably hopeless. Even though this was sorted out in the two-species case by using the WKB approximation, neither will this method be capable of bringing any good estimation of the average first-passage times observed in Figure 3.10.

In contrast, we must turn to the relatively recent methods of approximation for the treatment of glassy systems with long relaxation times and attempt to arrive at a solution while redefining the problem from the general theoretical angle to our ecosystem's view.

Here, we stop our argument and outline future steps towards the resolution of this problem. We have observed that the closed system displays non-trivial relaxation dynamics, even though its strict asymptotic limit ($t \rightarrow \infty$) is the trivial solution of equal probability of fixation for any of the equidistributed initial number of species (S). However, due to the exponential scaling to reach such true fixed states, the relevant features of the system reside only in its intermediate extinction/relaxation dynamics. These relaxation times cannot be easily captured by a similar analysis as in the case of $S = 2$ due to the irreducibility of the global interactions encapsulated in the term $\mathcal{D}(\mathbf{x})$ in (3.11).

Followingly, we will consider the implications of opening up the ecosystem to migration. We will do so for small migration rates, which will result in an interesting intermediate regime between the closed system and the migration dominated system (see [MW01]).

3.3.3. Open ecosystems

After arriving at an analytical hold-on on the intricacies of the neutral theory of cooperators for the closed ecosystem case, we conclude this chapter by introducing migration into the ecosystem. In particular, we

will show the main role of migration is that it acts as a system’s “stabilizer”: how it breaks some of the glassy properties discussed above and allows for better analytical treatment of the main ecological features that we are interested in.

In the first place, implementing migration into the deterministic view of the theory of neutral cooperators for fixed-size ecosystems is similar to what is done in the classic NT of competitors [SAM02]. More specifically, one can easily extend (3.10) to

$$\frac{dx_k}{dt} = x_k (\mathcal{D}(\mathbf{x}) - x_k) + \mu \left(\frac{1}{S} - x_k \right), \quad (3.15)$$

where $\mu \in \mathbf{R}^+$ is some coupling reflecting the migration rate with respect to the cooperation interaction rates. We note that now only one fixed point exists in the system, i.e. the homogeneous state $x_k = 1/S$ for all $k = 1, \dots, S$. In other words, the carved-out multi-simplex space with holes in each simplex centre is now smoothen such that only an attractor point persists at the centre of \mathcal{M}_S , and no absorbing points are left in the system⁶.

Perhaps more importantly, the open ecosystem scenario offers the possibility to study a series of ecological observables that is much closer to what experimental data can support, namely, species-abundance distributions (RSA), species survival timescales (LD) and turnover times (STD) (see Section 3.1).

To approach the problem from a heuristic (analytical) perspective. Let us consider that the migration rate coupling is very small when compared with the cooperation interaction forcing, which in (3.15) we set it equal to 1, i.e. $\mu \ll 1$. Following a similar hand-wavy argument as we did in (3.16), in this case, we argue that, on average, the system still displays a fast (transient) phase of extinctions leading to an equilibrium value of extant species⁷. Thus, we denote R as the average number of extant species and proceed similarly as before by applying our approach to $\mathcal{D}(\mathbf{x}) \approx 1/R$

⁶This fact is not too surprising. After all, migration is also required in the classic NT in order to arrive at sensible observables [MAS00, SAM02, MAS04].

⁷As we will discuss later, the actual species that survive this first transient do not per-

and studying the fluctuations over the generalized potential:

$$\mathcal{V}(\tilde{x}) \approx -\frac{\tilde{x}^2}{2R} + \frac{\tilde{x}^3}{3} - \frac{\mu}{S}\tilde{x} + \frac{\mu}{2S}\tilde{x}^2. \quad (3.16)$$

This time, our fixed point is shifted as

$$\tilde{x}^* \approx \frac{1}{R} - \mu \left(1 - \frac{R}{S}\right),$$

where we have neglected higher orders of μ . This yields to a characteristic fluctuation size of

$$\sigma \sim \sqrt{\frac{S}{N}} \left(\frac{1}{r} - \nu + 2r\nu\right), \quad (3.17)$$

where $r := R/S$ is the fraction of remaining species and $\nu := \mu S$, which is no longer an infinitesimal value. Thus, if migration is strong, then we expect to have $r \rightarrow 1$ and $\sigma \sim \sqrt{S/N}(1 + \nu)$; whereas if the system is such that $\mu \rightarrow 0$, then we expect $R \rightarrow O(1)$ and $r \rightarrow O(1/S)$, which implies that, under these circumstances, $\sigma \sim O(S/\sqrt{N})$.

However, it is important to note that, while this analysis helps in gaining intuition for how the system fluctuates depending on μ , it does *not* provide a concrete analytical answer to the question of what *actual* diversity (i.e. $r = R/S$) we expect to observe for a given migration rate. An answer to this question requires a better understanding of the relaxation dynamics of the system, which displays a glassy behaviour for low values of μ . Here, we will limit ourselves at providing a procedural argument

Notice that, for small μ values, the system presumably displays a behaviour proximal to that of the closed system theory. On the other hand, a rough estimation of the number of species inhabiting the ecosystem can be obtained by comparing extinction (Γ_μ) and species-reentry ($\tilde{\Gamma}_\mu$) rates, hence equilibrium is reached when

$$\Gamma_{N,S}(R; \mu) \simeq \tilde{\Gamma}_{N,S}(R; \mu). \quad (3.18)$$

sist indefinitely over time; instead, slow turnover takes place while keeping the average number of species present in the system relatively constant.

The key point is to assume that extinction rate $\Gamma_{N,S}(R; \mu) \approx \Gamma_{N,S}(R; \mu = 0)$ for sufficiently small values of μ , which corresponds to the inverse glassy relaxation time discussed in the previous section. On the other hand, migration reentry rates are easily computed from

$$\tilde{\Gamma}_{N,S}(R; \mu)\delta t = \mu \left(1 - \frac{R}{S}\right). \quad (3.19)$$

Thus, combining (3.19) with (3.18) should yield an approximated equation for $R_{N,S}^*(\mu)$. This should correspond to the numerical curves shown in Figure 3.11a, and would constitute a prediction for the expected *diversity* of the ecosystem.

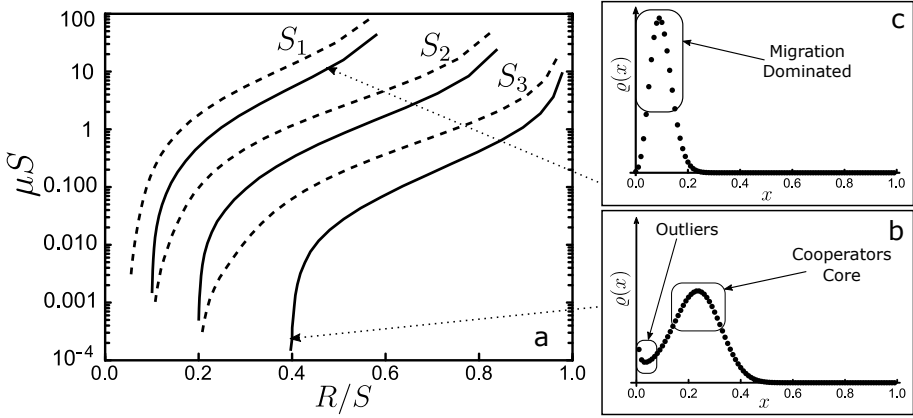


Figure 3.11: **Survival species and RSA.** Parameter space $(r, \nu) = (R/S, \mu S)$ is shown in (a), where the curves of $R^*(\mu)$ obtained from numerical simulations corresponding to $N = 20, 40$ (dashed, straight) and $S_1 = N$, $S_2 = N/2$ and $S_3 = N/4$, respectively. As the mutation rates increase, the system undergoes a nonlinear transformation from a bi-modal (cooperator core plus outliers, (b)) phase into a uni-modal one (migration dominated, (c)).

A core of cooperators

The right plots in Figure 3.11 reveal two types of underlying relative-abundance distributions, one that reflects the high-migration regime in

which the system fluctuates around $R^* \sim S$; and a second one, occurring at low migration rates, which shows a two-level system of species, namely *outliers* and *core cooperators* species. Let us elaborate on our latter definitions by focusing on the second type of RSA:

- *Core Cooperators*: this set of roughly $R < S$ species appears in large abundances. They display feedback dynamics in which, via cooperation, they maintain their core dominance over large periods of time.
- *Outliers*: this abundance peak corresponds to the new entry of species driven by migration processes. There is a flow of species between the outliers and the cooperators that we study below.

This is one of the main results of our analysis. On one hand, it provides a strong constraint on the kinds of ecosystem assemblies adopted by cooperative agents; it predicts the presence of a cluster of species that ‘take over and maintain themselves in the ecosystem despite the presence of a second (highly unstable) a cluster of species with very low abundances that arise from migratory processes.

Species turnover patterns

One way to measure and compare the robustness of the core of cooperators and the cluster of outliers is to study the species turnover times conditional to belonging to each of these subsets. Let us define τ_{ext} as the time any given species takes to go extinct. At first, and owing to neutrality, we do not indicate whether a species will belong to either cluster (cooperators or outliers). Figure 3.12a,b shows how we can think about this problem using the approximated potential, (3.12), which, if plotted, shows that the potential barrier for a core cooperator species to go extinct is much larger than an outlier. Hence, we expect that the outliers will quickly go to extinction and undergo rapid turnover; while the cluster of cooperators will go extinct at much longer times with a low turnover frequency. Figure 3.12c shows that this is precisely how the dynamics occur.

Note the log-log scales of the plot and observe that the count for extinctions at small timescales (the first peak from left to right) is exponentially larger than that pertaining to large timescales (second peak). This indicates that the system dynamics implemented (stochastically) *classifies* the extant species into the two classes described above, despite the lack of any species-specificity (neutrality). Finally, as expected, when the migration parameter is turned on ($\mu \gg 1$), the system recovers the statistics of a ‘migration-only’ model with a single peak of turnover time statistics that depends only on the diversity of the species pool (Figure 3.12d).

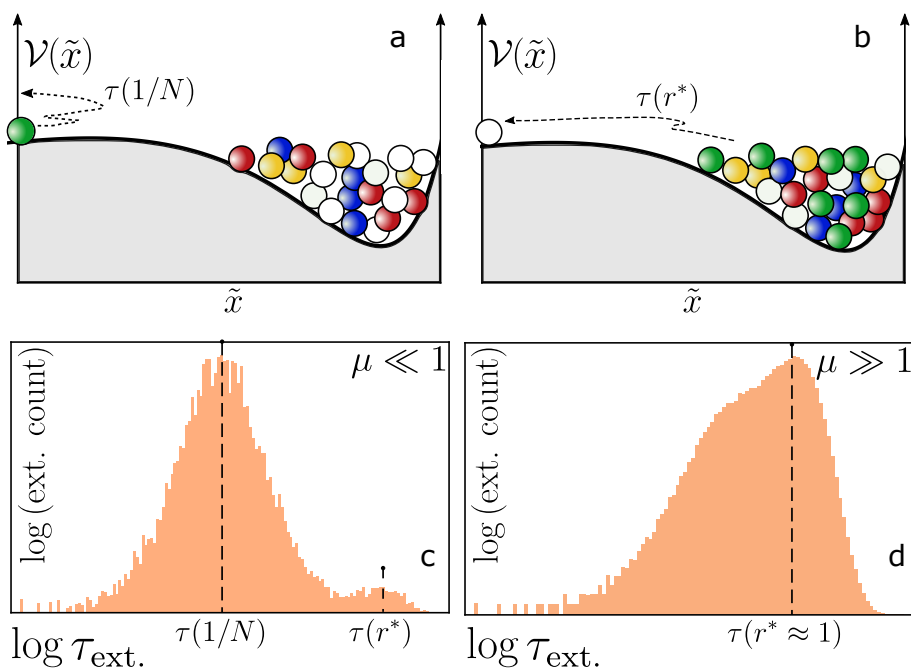


Figure 3.12: Species turn-over patterns. The top figures depict the time spent by any given species at a non-zero abundance state from its re-entry. At low μ : (a) shows rapid and slow lifetimes for outlier and core-cooperator species. The bottom figures are histogram plots in for a simulation with $N = 1000$ and $S = 500$ at small (c) and large (d) μ . Particularly, (c) shows this two-level clustering of rapid-frequent vs slow-rare turnover.

Chapter 4

LIQUID BRAINS AND COGNITIVE SPACES

Can a machine be made super-critical?

Alan Turing

The dominant role of information processing in living systems makes biology different from physics. As physicist John Hopfield pointed out, it is computation that makes biology different from physics [Hop94]. We have previously discussed the importance of information à la Shannon and its connection to the formalism of stochastic thermodynamics. We have shown how such approaches allow for an understanding of certain thermodynamic bounds on adaptive agents and replicators. While such bounds provide a strong handle on some of the evolutionary constraints for a very general class of systems, other barriers to achieving higher complexity are present in the history of evolution. In particular, *cognition* as information processing is essential for building up the statistical physics theory of METs. How many information-based transitions are there? What is the diversity of solutions associated with the resulting complexity?

The conquering of the planet by biological cognitive agents has been

achieved across scales. It involves several *major cognitive transitions*¹. A first important research undertaking consists of tracing the origins of the evolutionary steps that allowed information to be a crucial part of early replicators and the prerequisites for such early cognitive systems to emerge and become selectively relevant. One may conjecture that one of the reasons why information processing became so widespread and tied to life is that, while cognitive complexity is costly, it has the potential to yield arbitrarily large advantages since it allows a reduction of environmental uncertainty [Wag00, SS18]. When dealing with the large-scale networks of interconnected agents (which here we loosely define as *brains*), one of their main adaptive traits is their capability to operate as prediction engines. But what needs to be predicted? How complex a prediction *machine* has to be? And how does an information-processing system constrain its internal agent's complexity?

A central problem that only begun to be seriously considered with the rise of evolutionary theory is: why brains? Darwin's groundbreaking theory of evolution by natural selection transformed how scientists looked at the natural world [Dar59]. Biological structures are no longer mere parts of a classification scheme organized into categories. Darwin himself contributed to locating the human mind as part of the evolutionary path that connects us with the rest of the cognitive systems in the biosphere. To Darwin, cognition was widespread, perhaps present even in plants, and a question of degree, not quality. Here we will look at how it pervades complexity and whether or not qualitative properties separate the diverse forms in which "brains" (or cognitive matter) can be defined.

For multicellular animals, the presence of a brain resulted from the advantages of perception and an active exploration of their worlds. The *moving hypothesis* posits that active exploration of an organism's spatial environment was a key step in the evolutionary trajectory that produced brains [Lli02]. Under this viewpoint, prediction is both a cause and a consequence of animal movement, and its implementation requires learning

¹The definition of cognition is currently debated by scholars from multiple disciplines. However, this debate is beyond the scope of this dissertation. For further reading, see [LD20].




Figure 4.1: **Examples of cognitive networks** classified with labels solid/liquid and neural/aneural: (a) standard neurons, (b) flatworm nervous system, (c) computer chips, (d) placozoans, (e) roots, (f) leaf stomata, (g) ant colonies, (h) termite societies, (i) robot swarms, (j) slime mould (*Physarum*), (k) immune system and (l) microbial communities. Image extracted from [SMF19].

networks to emerge. But the ways evolution has made possible cognitions diverge considerably when looking beyond standard brain architectures, particularly vertebrate brains, remain an open problem. A very different class of cognitive network is exemplified by ant colonies, which are often referred to as superorganisms. The label is a reminder that they exhibit collective behavior and operate as a single entity, much like an individual organism. These highly organized social structures have enabled ants to thrive in virtually every corner of the planet. As it occurs with multicellular animals as individuals, which display development and cellular differentiation, one of the key features of ant colonies is their division of labor. That is, different ants have different roles within the colony, sometimes as fixed castes, and sometimes in a flexible way; with some specializing in foraging, some in nest maintenance, and others in defense. This specialization enables the colony to efficiently utilize its resources and respond to changing conditions. Moreover, ant colonies manipulate information by a combination of both chemical and tactile (ant-to-ant) signals to communicate and respond to colony needs and environmental challenges. The architecture is thus fluid, the system is a network of interconnected brains and they use (in the vast majority of species) a nest that also carries information. How relevant is individual versus collective complexity? Is it appropriate to talk about *one* brain here?

The answer to the previous questions can be achieved by (1) finding universal principles of information processing and computation in network-based agents and (2) mapping the space of cognitive systems systematically. Networks are an appropriate framework since the fabric of the best-known cognitive systems -brains- is made of large numbers of interconnected neurons that gather, process and respond to their worlds. To this end, in this chapter, we will highlight a few case studies presented in a synthesis work on the biology and statistical physics of liquid brains [SMF19, PS19]. We start with a summary of how solid brains can be modelled, compared with two case studies: collective behaviour in social insects and a model of the immune system grounded in the physics of spin glasses. Finally, we define a tentative space of cognition (Solé and Piñero, unpublished).

4.1. Brains, neurons and solid cognition

We use the term “brains” in a generic way: it will refer to ensembles of interconnected neurons (or neural-like agents). Since its inception in the mid-twentieth century, the field of neural networks has experienced massive growth and has moved on in multiple directions. Nonetheless, two key turning points were the classical paper by Hopfield [Hop82] and the development of multi-layer, feed-forward networks capable of complex recognition tasks. These became the basis for a statistical physics description of neural networks that emerged and primarily marked the development of this class of systems, and all were based on the simplest neural model first due to Warren McCulloch & Walter Pitts (1943) [MP43].

The pictorial conceptualization behind the McCulloch-Pitts neuron is sketched in Figure 4.2. This *formal neuron* is a simple Boolean system that reacts to a set of N inputs from a set of input neurons and responds to the sum of these inputs in a threshold manner. The state of each formal neuron is indicated as S_k and can take two possible values: each input and the output can be a 0 or 1 (we can also use -1 and $+1$). These two states are commonly associated with the resting and firing states: either the neuron remains inactive or is active and sends signals to other neurons. But how each neuron influences another is weighted by the nature and strength of the connection between them. These weights are indicated by ω_{ki} and have continuous values: positive or negative for activator or inhibitory connections, respectively.

In the McCulloch-Pitts approach, a neuron (S_i) that receives the input from others integrates all of them, each weighted by the value and sign captured by ω_{ki} . The result is the following discrete equation:

$$S_i(t+1) = \phi \left(\sum_{k=1}^N \omega_{ki} S_j(t) - \theta_i \right), \quad (4.1)$$

where an additional number, θ_i is introduced. This is a threshold for response: the neuron will fire if the sum in the RHS is larger than θ . The




Figure 4.2: Real and formal neurons. In (a), a pyramidal neuron is shown to be compared with the toy model of a neuron, as suggested by McCulloch and Pitts (b), where the minimal components are preserved at the logic level. A set of “input neurons”, $\{S_1, \dots, S_n\}$, send their signals to neuron S_i in which the sum of all inputs, weighted by their specific links ω_{ki} , is performed and compared to an internal threshold number θ_i . The decision of firing is then made using a transference function $\varphi(\Sigma)$.

function $\phi(x)$ introduces this all-or-none response:

$$\phi(x) = \begin{cases} 1 & ; \text{if } x > 0 \\ 0 & ; \text{otherwise} \end{cases}. \quad (4.2)$$

The shape of this function is displayed in Figure 4.2. It is a step-like response, but more general response functions can be defined. As an example, we could use a saturating form with intermediate continuous values, namely $\phi(x) = 1/(1 + \exp(-\beta x))$. Here β approaches the step-function limit as β grows to high values. A major result shown in [MP43] is that every Boolean gate can be mapped into a neural circuit described as a set of connected formal neurons, which implies that all possible computer operations are theoretically executable by brains [Arb12].

4.1.1. Solid brains: criticality

Dynamical activity in cortical networks constitutes a significant line of active research and has been central to accounting for the enormous potential of neural nets for optimal information processing. In particular, the *critical brain* hypothesis [Chi04, Chi10, SP13, Mun18] states that the optimal computational state of activity patterns for the cortex coincides with a critical point in neural activity. Criticality refers to a state where the brain is poised at *the edge of chaos*, with neural activity exhibiting complex patterns that are neither completely random nor completely regular. Here scaling laws are a characteristic pattern found both in the spatial and temporal organization of brain activity.

One of the key reasons for the relevance of criticality is that it has been associated with a range of cognitive tasks, such as perception, attention, memory, and decision-making. It has been conjectured that criticality may play a fundamental role in how the brain processes information. Moreover, it is a relevant feature that might explain how coordinated functional traits emerge and are optimized. Notably, the departure from the critical state is likely connected with neurological disorders. Studies have shown that pathological conditions such as epilepsy, schizophrenia,

and Alzheimer's disease are associated with changes in neural activity that disrupt the critical state of the brain. This suggests that criticality in neural activity may be essential for maintaining healthy brain function and preventing neurological disorders.

Modelling the complexities of the whole brain dynamics is far from simple, and a complete approximation should consider the interplay between spatial dynamics, synchronization and noise that is intrinsic to cortical networks. An example of the kind of model that is commonly used is a dynamical system of local membrane depolarization potentials, $V_i(t)$, of a set of spiking neurons (see [DJR⁺08] and references therein) that reads:

$$\tau \frac{dV_i}{dt} = -(V_i(t) - V_l) + \tau \sum_{j=1}^N J_{ij} \sum_k \delta(t - t_j^{(k)}) \quad (4.3)$$

where V_l is the resting potential, the matrix J_{ij} introduces the couplings (the synaptic efficiency) between neurons and the last term on the right-hand side introduces the synaptic current from presynaptic neurons. This class of *integrate and fire* models² involve a highly nonlinear dynamics, and analytic results are seldom easy to be derived. However, a simple model helps understand the presence of well-defined phases and the presence of criticality. As it occurs with other problems considered here, discretization allows for simplifying the system representation, while keeping the most fundamental properties [HG14].

Consider a system of N agents (neurons) with two possible states, say active and inactive. Suppose we indicate by $S_i(t)$ the state of a given neuron, with $S_i \in \{0, 1\}$. In that case, we can consider the previous

²In the integrate and fire model, the neuron is represented as a simple electrical circuit, with a capacitor that stores charge and a resistor that allows charge to leak out. When a synaptic input arrives, it causes a small change in the voltage across the capacitor. If the voltage reaches a certain threshold, the neuron fires an action potential, and the capacitor is discharged. After firing, the neuron is temporarily refractory, meaning it cannot fire again until a certain amount of time has passed. Despite their simplicity, these models have been widely used to represent both single-cell and network-level features of neural assemblies.

McCulloch-Pitts threshold model as the dynamical framework to describe the dynamical changes. Suppose we restrict ourselves to a system of excitatory, homogeneous synaptic couplings, i.e. $J_{ij} = J > 0$. In that case, it is possible to consider a random network where the j -th neuron is connected with k_j neighbours. Since N is fixed, we have $N = N_a + N_i$ where N_a and N_i are the numbers of active and inactive neurons, respectively. An equation for the dynamics of N_a can be easily derived from the following:

$$\frac{dN_a(t)}{dt} = \sum_{j=1}^N J \left[\sum_{i=1}^{k_i} \mathbb{P}(S_j(t) = 0 \cap S_i(t) = 1) - \frac{1}{\tau} S_j(t) \right] \quad (4.4)$$

We introduce a characteristic decay time, τ , for the spontaneous inactivation of a given unit. Here $\mathbb{P}(S_j = 0 \cap S_i = 1)$ indicates the joint probability of having the units S_j and a given neighbour S_i . Assuming the network is homogeneous, we use $k_i = \langle k \rangle$. Additionally, let us ignore correlations (as in a standard mean field approximation) and assume stochastic independence, i.e. $\mathbb{P}(S_j = 0 \cap S_i = 1) = \mathbb{P}(S_j = 0)\mathbb{P}(S_i = 1)$ to obtain:

$$\frac{dN_a}{dt} = J \left(\sum_{j=1}^N \mathbb{P}(S_j = 0) \right) \left(\sum_{i=1}^{k_i} \mathbb{P}(S_i = 1) \right) - \frac{1}{\tau} N_a \quad (4.5)$$

Under these assumptions, using $A = N_a/N$, and dividing both sides of the previous equation by N , we obtain a dynamical equation for the activity³:

$$\frac{dA}{dt} = -\frac{1}{\tau} A + J \langle k \rangle A(1 - A). \quad (4.6)$$

The fixed points for this dynamical equation are obtained from $dA/dt = 0$ and are $A^* = 0$ (no activity) and

$$A^* = 1 - \frac{1}{\tau J \langle k \rangle} \quad (4.7)$$

³An extra approximation is to consider the network is such that $\langle k \rangle$ is small.

It is easy to show (see Section 2b in [PS19]) that the activity fluctuations will relax following a scaling of the form

$$\delta A(t) \sim \exp \left[-\frac{t}{\tau} (1 - J\langle k \rangle) \right]. \quad (4.8)$$

This implies that, if $J\langle k \rangle \rightarrow 1$, then the effective relaxation time, T , diverges following a scaling law

$$T \sim |1 - J\langle k \rangle|^{-1}$$

which is a hallmark of criticality.

It has been shown that the plasticity and rapid dynamical response present in the critical regime of activity constitutes an optimal regime for fast response to external stimuli and optimal information processing [Beg08]. From a clinical perspective, moving away towards lower connectivity corresponds to an unconscious brain state, while the opposite produces high activation modes that are akin to epilepsy [Zim20]. As shown below, this critical point's presence and adaptive relevance might be a universal, shared property of cognitive networks. In practical terms, the presence of criticality is observable by the rapid increase in the time series variance associated with the dynamical state.

4.2. Liquid brain attractors

As mentioned at the beginning of the chapter, a whole class of living systems, along with some artificial ones, can be described as liquid brains. These include different types of social insects (ants, termites, bees and social wasps), immune systems (formed by many cells moving and interacting within a fluid medium), and flocks of birds and fish schools. The fluid dynamics that dominate each example within these classes is a characteristic trait: the spatial location of each agent changes over time. As a consequence, their interactions (direct or indirect) also change. This contrasts sharply with the spatial correlations intrinsic to solid brains, where agents (neurons) remain in the same positions, and cell-cell connections

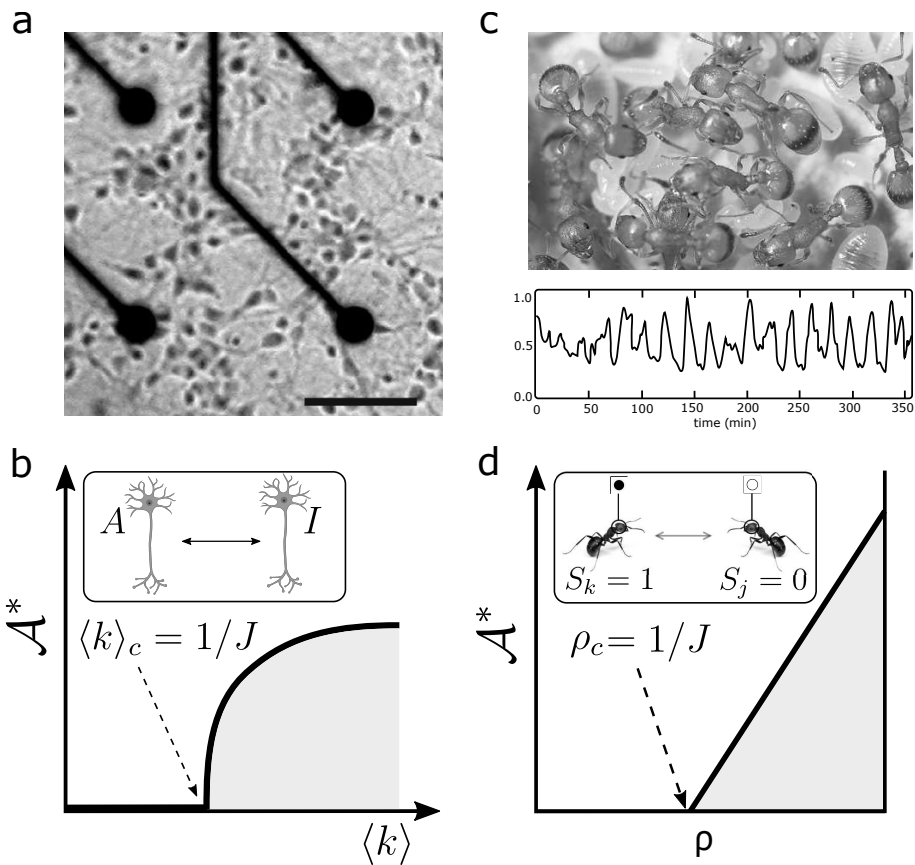


Figure 4.3: Criticality in solid-liquid brains. (a) shows a section of neural cortex with micro-electrodes attached to record activity (adapted from [EFG⁺07]), such cortical patterns of activity are shown to be poised near criticality (b), which is modelled after a phase transition displayed in the propagating activity model (4.6) and controlled by the connectivity parameter $\langle k \rangle$. (c) shows an ant species from the genus *Leptothorax* that displays an oscillatory pattern of activity which is too captured by letting the system approach criticality in a phase transition model (4.13), which is controlled by the density of ants ρ . Adapted from [PS19].

evolve in time due to learning. Moreover, there are different kinds of agent complexities involved. Cells in the immune system are to be compared (in terms of complexity) with neurons in brains, while ants, termites or birds are each equipped with a whole brain. In that respect, the latter are brains made of brains.

There is a vast literature on self-organization in ant and termite colonies that we cannot summarize here [CDF⁺20]. Studying these systems involves various modelling approximations, from simple non-linear differential equations to pattern-formation models incorporating the physical context and boundary conditions as part of the problem.

In this section, we address, using a specific model, one first important question: How do liquid brain attractors differ from those associated with solid brains? To answer such a question, let us consider an ant colony where a finite number of tasks must be executed. Our example is inspired by the colonies of harvester ants (belonging to the genus *Pogonomyrmex*), so-called because they collect, store and eat seeds. A colony living in a dryland environment performs various tasks outside the nest, several of which are easily identified. Foragers, for example, leave the nest on shared foraging trails, collect food, and return it to the nest. Nest maintenance workers maintain interior tunnels, carry out sand, and clear vegetation from the nest mound. On the other hand, patrollers recruit foragers towards new food sources. Finally, midden workers sort and maintain the colony refuse pile. The main point is that, through ant-to-ant interactions, ants can switch from one task to another when the colony's needs and/or as a response to environmental states change.

One expected result is that the number of attractors will be somehow reduced by the inevitable degeneracy associated with the nature of interactions. The reason has to be found in how we define the "states" in our system. It is clear that each ant has a brain, and thus, at the individual scale, neuron states are relevant and would define individual behavioural conditions. But there is a functional class, too: the one that describes the task being performed by each ant. In this context, observations can be used to code the transition probabilities between tasks in a dynamical system defined as a threshold neural network [GGT92].

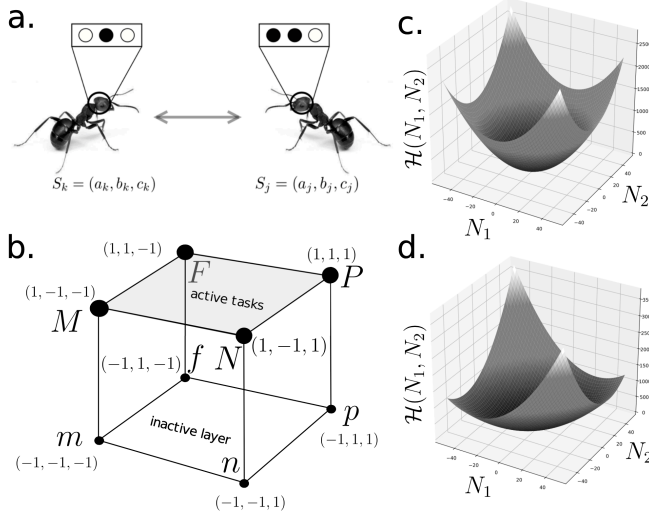


Figure 4.4: Population-level attractors and task allocation. Here virtual ants (a) each carrying a 3-spin internal description. The total state space is a three-dimensional Boolean cube (b). The model exhibits attractor dynamics with an associated energy function (see text), as illustrated in (c-d), for a two-task system with $\alpha = 1$ and $\beta = 0.1$ in (c) and $\beta = 0.5$ in (d). Figure from [PS19].

Let us consider our colony as formed by a set of N ants interacting through a matrix J_{jk}^μ . Considering ants as mobile agents, their states are updated following a threshold function:

$$S_j^\mu(t+1) = \Theta(h_j^\mu(t)) = \Theta\left(\sum_k J_{jk}^\mu S_k^\mu(t)\right). \quad (4.9)$$

The (internal) state of the j -th ant belonging to the task μ , S_j^μ will remain the same provided that $S_j^\mu h_j^\mu > 0$.

As is done in models of solid neural networks, an energy function is defined accordingly as follows:

$$\mathcal{H}(\{S_k^\mu\}, \{J_{ij}^\mu\}) = -\frac{1}{2} \sum_\mu \sum_{i,j} J_{ij}^\mu S_i^\mu S_j^\mu. \quad (4.10)$$

But this is a highly degenerate state, and The observable state is the number of ants performing each task, which leads to an energy is defined in terms of the set $\{n_k\}$:

$$\mathcal{H}(\{n_k, \Gamma_{ij}\}) = -\frac{1}{2} \sum_{i,j} \Gamma_{ij} n_i n_j , \quad (4.11)$$

with a new set of parameters $\{\Gamma_{ij}\}$ that depend on the microscopic couplings (see [GGT92]). The system's equilibrium states (attractors) are the minima of this function and, for $\Gamma_{11} = \Gamma_{22} = \alpha$ and $\Gamma_{12} = \Gamma_{21} = \beta$:

$$\mathcal{H}(N_1, N_2) = -\frac{1}{2} (\alpha N_1^2 + \alpha N_2^2 - 2\beta N_1 N_2) , \quad (4.12)$$

as it is illustrated in Figure 4.4. The crucial point to be made here is that, in our liquid context, the high-dimensional nature of neural networks derived from the connectivity matrix (which distinguishes between specific pairs of neurons) is no longer valid. Changing connectivity matrices and the lack of agent individuality (beyond its task string) imply low-dimensional attractors.

4.3. Liquid brain criticality

Although the nature of the attractors in liquid brains clearly departs from those associated with solid ones, a common property seems shared by both: the exploitation of critical points. It seems to be present in fish schools and flocks of birds and has been described in some ant colonies. Here we consider this example, which has been reported for ant colonies of the genus *Leptothorax* [SMG93, SM95] Such critical dynamics are revealed by recording the number of individuals that are active (and thus perform tasks and move around). Such recordings indicate critical-like behaviour, with low activity levels punctuated by synchronized bursts of activity. Such bursts are triggered by the local activation of an inactive individual by an active one. Additionally, inactive ants can spontaneously activate with some probability. The statistics are consistent, particularly

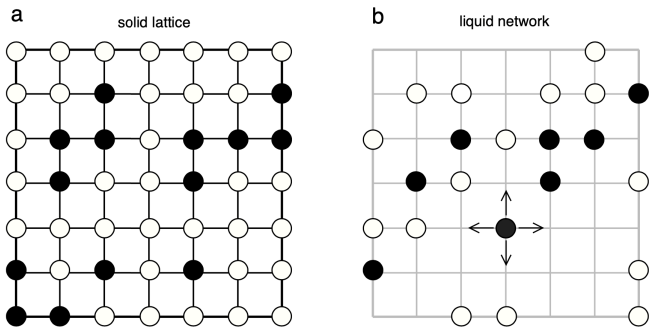


Figure 4.5: **Modelling ant colony interactions on a lattice.** Instead of a fixed distribution of elements placed at regular locations (a) the proper modelling of an ant colony should consider instead a liquid system (b) where active elements (filled circles) can move to the nearest locations, while inactive ones (open circles) remain in the same place.

at the time scale level, with scale-free fluctuations and scaling behaviour. The observed patterns indicate that there are no regular dynamics at the individual level, and thus collective activations are not just a set of coupled oscillators that get in sync.

The previous observations of the collective and single-ant behaviour suggest that the problem is somewhat similar to the one explored above on the propagation of activity in cortical systems. An important observation to be made here is that ants can adjust their nest density by changing their boundaries and that a critical density of ants can sustain the critical-like fluctuations found in nature. This suggests that they might get self-organized in a critical state. A simple model helps us capture this phenomenon and allows us to compare the two cases, solid and liquid brains, under the light of self-organization and criticality.

Consider a two-dimensional lattice. At each point, an ant can be present or not; if present, it can be either active or inactive. If left alone, an ant maintains activity during some characteristic time τ . In general, the update of individual states will occur using our familiar threshold function update, similar to the one discussed in previous sections. If we reduce

our scope to the only-activation scenario, when nearby ants are active an inactive ant will activate with some coupling J . Following the same steps, we can write [PS19] a simple mean-field model that now reads:

$$\frac{dA}{dt} = \frac{J}{\tau}A(\rho - A) - \frac{1}{\tau}A, \quad (4.13)$$

where A is the fraction of active ants in the lattice and ρ is the total density of ants. Notice that instead of average connectivity, the likelihood of local interactions is weighted by the ant density parameter. Equivalently as in the analysis done in Section 4.1.1, we obtain fluctuations of order

$$\delta A(t) \sim \exp\left[-\frac{t}{\tau}(1 - J\rho)\right], \quad (4.14)$$

around the zero activity point, thus criticality is reached $J\rho \rightarrow 1$, with a relaxation time diverging like $T \sim |1 - J\rho|^{-1}$. At the active phase, $A^* = \rho - 1/J$. See Figure 4.3c,d.

After this exercise, we observe that either system (solid or liquid) behaves similarly from an activity dynamics perspective. Despite their apparent physical differences, we say that these systems belong to the same *universality class*. The fact that both self-organize to criticality [MB11] is arguably a major feature of evolution involving interacting systems whose function is to perform a cognitive task.

4.4. The immune system as a liquid brain

Our last example deals with a major evolutionary innovation, namely the emergence of the immune system. This liquid network is made of several diverse cell types (allowing, once again, division of labour) that responds to pathogens to protect the individual from infection. This is a complex adaptive system capable of memory and learning. Major contributions to theoretical immunology came from the statistical physics analysis of its memory states and the immune repertoire, and a great account of them was written by Alan Perelson and Gerard Weisbuch (Figure 4.6)



Figure 4.6: Perelson, Weisbuch and Parisi. Three pioneers in the development of models of immune system’s complexity, some involving the connectionist approach [Far90]. From left to right: Alan Perelson, Gerard Weisbuch and Giorgio Parisi. Perelson and Weisbuch made major contributions to theoretical immunology, including work on immune networks. Some of these models are deeply connected with spin glasses, a field developed by Parisi, who won a Nobel Prize in Physics in 2022.

in their classic paper *Immunology for physicists*, [PW97]. Some of their work was based on previous work in statistical physics concerning disordered systems. In 1990, Parisi [Par90] presented a simple model of the IS that explored its memory capacity in ways similar to those used in the statistical physics of neural networks. In this particular work, Parisi was inspired by the study of idiotypic networks [Jer74] and used a spin glass analogy.

Parisi considered a system such as (4.15) in which every degree of freedom corresponds to an idio type, and $\{J_{ij}\}$ are suitably defined (see Section 3b(iii) in [PS19]). This analysis yielded a prediction for the number of possible memory states attainable by the system, which is of the order $2^{\lambda N}$, with $\lambda \sim 0.3$, where N is the total number of idiotypes (which are around $10^6 - 10^7$ in humans). These findings fuelled the application of statistical physics to the analysis of the immune system.

BOX 4.4: Spin glasses: what & why?

Spin glass theory surged in the context of condensed matter around the 1970's^a [SN13], but was independently discovered in the context of evolution by Stuart Kauffman [Kau69b, KL87, Kau93]. Spin glasses pervade a wide range of problems, from material sciences to artificial neural nets. Its theoretic core is difficult to synthesize, however, some intuition can be gained by studying the (seemingly innocuous) energy function:

$$\mathcal{H}(\{J_{ij}\}, \{S_i\}) = - \sum_{(i,j)} J_{ij} S_i S_j. \quad (4.15)$$

Here, $\{S_i\}$ are the states of the $i = 1, \dots, N$ degrees of freedom (dubbed *spins*), $(\mathbf{J})_{ij} = J_{ij}$ its pair-wise interaction matrix, and the brackets (i, j) indicate some type of connection topology between the spins. The system spontaneously tends to its lower energy values. When $J_{ij} = J$, we have the well-known Ising model [Set21], which captures the essential phase transition of a magnet or liquid-vapor. When J_{ij} becomes somewhat more complicated, a radically new outlook emerges. Here, a key concept is given by *frustration*. Consider a two-state case, $S_i = \{\uparrow, \downarrow\}$, and suppose that $J_{ij} = J < 0$. Then, if the system's topology is triangular (see Fig. 4.7), we recognize that top spin cannot relax into a stable state.




Figure 4.7: **Frustration and Fitness Landscapes.** Left: frustration of an $N = 3$ (antiferromagnetic) spin system. Right: Waddington landscape [Wad14], which displays ultrametricity.

^aSee [And88a, And88b, And88c, And88d, And89a, And89b, And90] for an introductory set of essays by the brilliant Phil Anderson.

This is a microscopic example of frustration for a very small setup. Scaling up this phenomenon into a fully-fledged spin glass, with, e.g. randomly generated J_{ij} values, results in a rich variety of phenomena. Frustration can represent the evolutionary aspects of *epistasis*, which involve the non-linear fitness responses after mutation for coupled genes, and it ultimately gives rise to an *ultrametric* characterization of all the accessible spin glass states, which, in evolutionary theory, constrains the possible paths the system may evolve through. In development, this phenomenon is captured by Waddington’s landscape (Fig. 4.7).

For example, Percus *et al.* [PPP92, PPP93] combined the statistical analysis of idiotypic interactions to propose an answer to the *self vs. non-self* distinguishability problem in immune systems, as well as better accounting for the completeness of the immune repertoire as a shared memory system between the individual agents (idiotypes) and the collective network, i.e. the networks spin-glass energy minima (see Section 3b in [PS19], and [PW97] for a comprehensive introduction).

Parisi’s model belongs to the discrete, symmetric picture inspired by Hopfield networks, but many different problems can be considered, given the complexity and diversity implicit in IS organization. Let us consider, for example, the dynamics of the IS by studying a simple model of the population x_i associated with the i -th B-cell clonal type [dB88, DBH89, DBKP90]. Suppose that it can be described as the sum of three different terms, which constitute a balance between a source (cell production from the bone marrow), a linear decay term associated with cell death and the proliferation resulting from the stimulation by other clones coupled to a cell proliferation rate p . Formally:

$$\frac{dx_i}{dt} = m - \delta x_i + p\phi(h_i)x_i. \quad (4.16)$$

The factor $\phi(h_i)$ in the last term involves information-processing mecha-

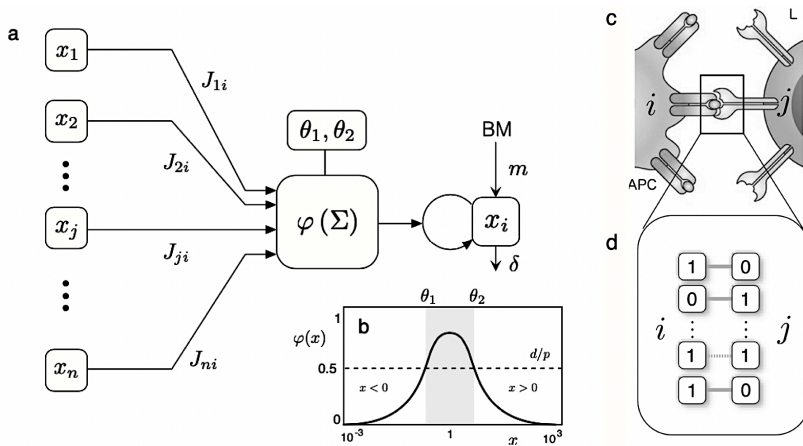


Figure 4.8: **Early models of the adaptive immune response** were based on a neural network picture (a), but involved transfer functions with two thresholds (b). On the molecular scale, the IS is a pattern-recognition system where molecular interactions between cells (c) are mediated by a matching process (d).

nisms akin to the McCulloch-Pitts neuron, while we set

$$h_i = \sum_{j=1}^n J_{ij}x_j + A_i x_i, \quad (4.17)$$

where n is the total immune clonal repertoire, and the weight matrix J_{ij} encodes for all the interactions between the clonal types; which can be either inhibitory or excitatory. Here, A_i is the i -associated antigen concentration in the environment, which acts as an external field. In this case, the nature of the function $\phi(h)$ is a continuous-valued population level, instead of a step-wise one (4.2). In particular, De Boer and Hogeweg (1989) used a log-bell-shaped response function following experimental observations (Figure 4.8b, see [PW97] pp. 1237-1239 and references therein) Other models explore such couplings in correspondence with biochemical affinities for epitope-paratope molecular interaction [BA10] (Figure 4.8c,d, see Section 3b(iv)-(vi) in [PS19]).

We note that, albeit the system described by (4.16) is much more complicated than the approximations given in (4.6) and (4.13), they are deeply

connected in a formal but also in a universal manner. Thus, in spite of the fact that the immune system's computational objectives (essentially the distinction of a foreign or malign antigen in the body) differ widely from those that arise in insect colonies or brain function; it is possible to conjecture that its operational optimality is too self-organized around criticality.

4.5. Mapping liquid versus solid cognition

The conceptual framework advanced in this chapter asks for a new synthesis. Solid and liquid brains are both the outcome of evolutionary forces, are shaped by some common constraints and display several dissimilarities that might be related to their adaptive response to environmental needs. A successful synthesis will need a complex systems perspective. The examples discussed above, along with all those considered in the review [PS19], deal with systems displaying a diverse range of information processing capacities, each describable by a different class of physical state. Moreover, some are tied to complex organs within organisms resulting from a developmental process, while others, such as ant or termite societies, cannot be mapped into the organ picture, despite they experience some kind of developmental process. A tentative space of possible cognitive systems should thus include as potential axes: (1) computational complexity, (2) physical state and (3) to what extent the system is the result of a developmental process. Let us consider some case studies that will help us in determining the bounds of our space and locate the different classes of cognition.

The solid, vertebrate brain that has been intensively studied over the last century might have resulted from an evolutionary pressure to deal with a challenging environment to be actively explored. Thus, movement would be a major driver for solid brains to merge. Such brains, as we have discussed above, are describable topologically by a static graph with spatially located nodes (which can degrade and disappear) connected by means of a weighted matrix. In terms of dynamics, various threshold-

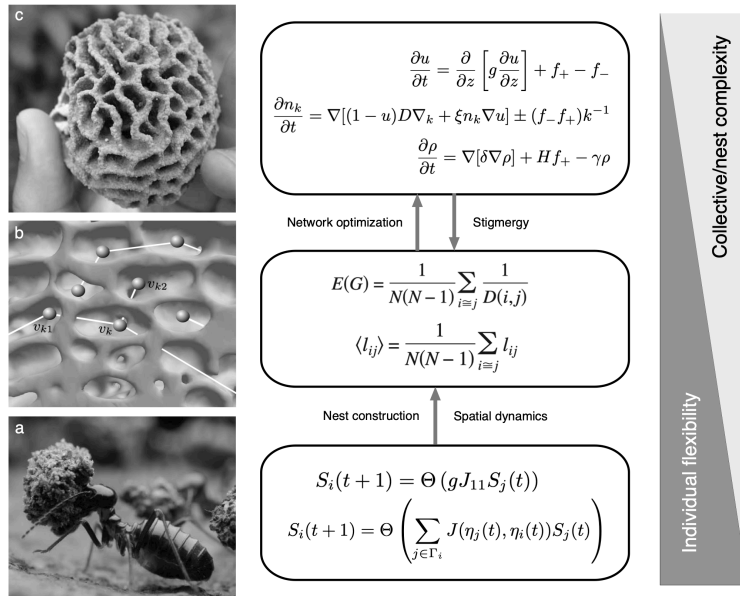


Figure 4.9: Collective intelligence is an emergent phenomenon where interactions among individuals (a) need to be understood under the context of the scaffold they create (the nest), which involves a network of interactions (b) that occurs and adapts over time to a changing set of physical boundary conditions (c). The arrows between models are a reminder that emergent phenomena occur when moving between levels, and that there is an irreducible feedback loop between colony organization and individual degrees of freedom.

based dynamical systems exist that can do the job of modelling different kinds of phenomena, from synchronization to pattern recognition. However, if a collective of agents is at work, and the agents are mobile, we have a rather different situation. Starting from early, pre-social insects, superorganisms emerged as a consequence of cooperative interactions that included: (a) remain close to eggs and (b) take care of them for a while. Such care was typically taking place in a spatial context defined as a precursor of a nest architecture. The advantages of this early cooperation are clear: by protecting young larvae or pre-adult offspring from predators or external uncertainties. Since the underlying organization is that of a

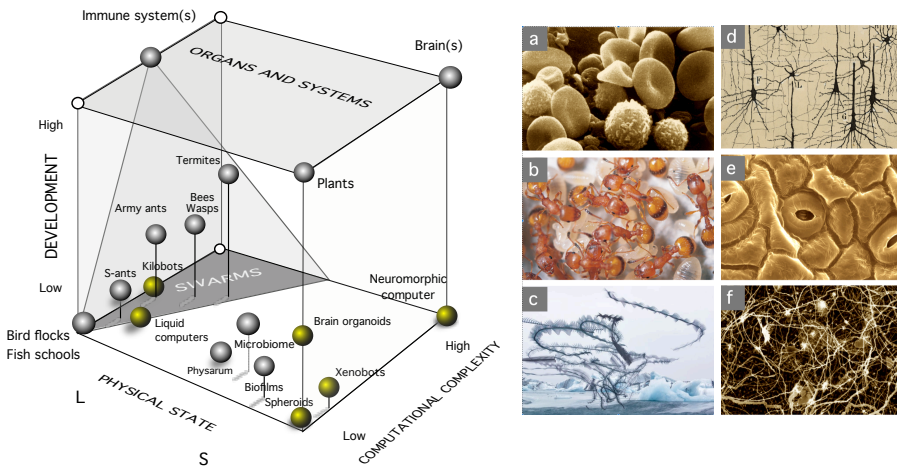


Figure 4.10: Morphospace for liquid and solid cognition. The left image is a tentative space of the possible where the axes are defined by: (i) the computational complexity of the system, (ii) the physical state (from liquid to solid) and (iii) the presence of a developmental process that generates the final network. Six examples are shown: (a) blood cells (both red and white), (b) *Leptothorax* sp. Ants, (c) bird flocks (image by Xavi Bou), (d) cortical neurons in a vertebrate brain (drawing by Ramon y Cajal), (e) stomata in a plant leaf and (f) a random nanowire neural network.

collective of individuals, decision making is no longer a centralized phenomenon, to be located in the network. The emergence of the superorganism (i.e. a cooperative group of agents equipped with brains) necessarily entails to deal with the constraints derived from multi-agency.

With the exception of army ants, one of the universal traits of collective intelligence is the presence of a nest. Nest complexity allows to define a set of complexity levels associated to different traits (and modelling approximations) from individuals to whole colonies, as illustrated in Figure 4.9 in the case study of termites. They are constructed by exploiting self-organization principles that include Turing-like instabilities [TBD98] Nests can be simple (such as the one of *Leptothorax* colonies) or extremely complex, as it is the case of termite societies. Termite nests

are engineered structures that can be found in many regions of the world [Tur11]. The construction of a termite nest is a complex process that involves the cooperation and coordination of thousands of individual termites whose individual size can be three orders of magnitude smaller than the nest. One of the key features of termite nests is their ability to regulate temperature and humidity levels throughout the day and night. All these features gave support to understand the colony as a whole (nest and individuals) as a class of homeostatic organism [Tur19]

The time development follows a predictable sequence of events, which include the initial formation of a pattern of pillars and arches. While the dynamics of individuals can be described, in lab conditions, as a threshold neural network, as soon as we consider termites in the nest context, a feedback loop between the emergent physical scaffold and termite behavior emerges [TBD98]. Such process can also be described in terms of networks of connected chambers and evidence suggests that these nets are close to a critical point. This time, criticality is associated to percolation: chambers are connected by tunnels in such a way that it spans the nest volume in an efficient way with a sparse connectivity [VCMP⁺09]. All in all, there is a subset in our morphospace, as sketched in Figure 4.10, where swarms, the immune system and some artificial collectives (such as *kilobots*) occupy a domain of close-to-liquid structure, low or intermediate development (with the exception of the immune system) and limited cognitive (computational) complexity.

In our cognition cube we easily locate solid brains in a corner of the space. This is of course an oversimplification, since we cannot collapse a whole range of brain architectures in a single point. What we want to point out is that they define an upper bound that combines full development (as an organ within an organism), high computational complexity. It shares this top roof of the space with organs (including the immune system) that also gather and respond to information but in much simpler ways and usually dominated by negative control feedbacks. Among the general traits observed in most of these solid brains is the presence of some kind of intermediate processing elements. Within the context of cortical systems, these interneurons have been known since Cajal's early studies of

brain architecture. They are a specially important element from the point of view of the evolution of cognitive complexity. They do not directly receive inputs from the environment nor directly respond to such inputs. Instead, they are intermediate steps in the non-trivial information processing allowed by a solid neural network. We emphasize here this class of element because it is fundamentally absent (or only indirectly present) in the liquid phase of our morphospace. Moreover, located on the same plane but in the low-cognition side we locate plants, whose "solid" architecture [DNB19], sessile nature⁴ and lack of intermediate processing units makes them a successful group that might not have required cognitive powers to get there. However, we also know that some plant subsystems (plants are highly modular, clonal organisms) are capable of responding to specific external signals (such as water and CO_2 levels) in a distributed manner with critical-like behaviour [PWMM04].

What is there in the domain of no development and solid cognition. In this case, several examples can be described that are the result of artificial designs. Here we only indicate a few, from neuromorphic computers (where one goal is to get inspiration from cortical neurons to design neural circuits) to xenobots. The latter are Xenobots can be defined as a new class of synthetic living machine made from frog embryo cells[KBLB20, EL21]. They are designed using an evolutionary algorithm that predicts optimal shapes able to perform given functionalities (such as cleaning up a given area). One of the remarkable outcomes of this research is the fact that "behaviour" can be disconnected from any neural equipment and instead be strongly tied to morphology (see [BL16] and references therein).

Along with the examples outlined in this chapter and reviewed in [PS19], some remarkable organisms appear to inhabit a rather special place in the cognitive universe. This is the case of the slime mould *Physarum*, which has been studied extensively for its remarkable ability

⁴Seeds might be the exception to this rule. They are the "moving" part of plants and are equipped in many cases with a gene regulatory network that allows decision making possible. This is particularly relevant in the context of germination, which requires integration of temporal information.

to exhibit complex, self-organized behaviour [Ada10, DLBS10, NYT00]. This organism is a single-celled (but multi-nucleated) eukaryote that can grow up to several centimetres in size and moves by extending its cytoplasmic mass in a rhythmic, pulsating manner. Despite its simple structure and lack of a nervous system, *Physarum* can solve many different problems⁵, from finding the shortest path between two points to constructing efficient transport networks. The latter is based on the ability of this organism to form networks that can optimize transport. We have here an adaptive network of connections that is connected to shape shifts and has been modelled by a threshold-like transport model where an original graph (with nodes occupied by nutrient) is formed by a uniform spread of the mould followed by a selection of shortest paths that can be essentially understood as a least action process [BMV12, TKN07]. The final computation of this shape-shifting, unique organism is nothing but morphology. The shape of the final set of living connected paths is the solution to the input problem. Interestingly, some fundamental similarities exist between this and the ways ant colonies use pheromone trails to explore their worlds.

A final point can be made here regarding the empty spaces of our morphospace. Consistently with our conjecture that liquid systems will have highly degenerate attractors and thus a much lower dimensionality associated to their attractor landscape, a complex cognition comparable to the one of solid brains is unlikely to be sustained. In this context, although it is possible to suggest a hybrid model that combines ants and fields to obtain some kind of Hopfield-like network [IMS91] the needs of a colony might be so reduced that an efficient and diverse task allocation mechanism might largely suffice to deal with environmental uncertainties. In this context, the ecological context that we have explored in this

⁵It can solve mazes and other complex spatial problems. When presented with a maze, *Physarum* will initially explore the available paths at random, extending its tendrils in all directions. However, over time, the organism will converge on the shortest path between the entrance and the exit of the maze, forming a network that optimizes the flow of nutrients. This process is driven by the chemical signals that *Physarum* releases as it moves through the maze, which reinforce the paths that are already being used and inhibit the exploration of less efficient routes.

dissertation, where cooperation plays a specially relevant role, might play a fundamental role in shaping the architecture of liquid interactions.

Chapter 5

DISCUSSION

5.1. Thermodynamic bounds on replication and energy harvesting

In Chapter 2 we have established a framework embedded into the language of nonequilibrium thermodynamics that allows for a comparison between different replicator classes (Section 2.3.1). Such analysis is grounded on the extended second law and previous results on coarse-graining for mesoscopic systems [Eng13] (Section 2.2).

Our results contrast with previous work, which focused only on dynamic or kinetic arguments [SS97], and connect the replicator processes to thermodynamic statements about the minimal entropy production required (here dubbed *LEB* for *low entropy bounds*) per replication event in each of the three basic replicator types. This established a thermodynamic metric by which it is possible to compare these replicators on energetic rather than kinetic grounds.

A deeper exploration for such dissipative bounds are object of future work. In particular, the following are open questions to be resolved:

- Application of the *LEB* framework in explicit competition set-ups. This problem consists on finding entropy production rates for systems that couple a coarse-grained description of a transition of pairs

of replicators, for example: $(\mathbf{s}, \mathbf{p}) \rightarrow (\mathbf{s}, \mathbf{s})$.

- Introduction of absorbing points. If a replicator can go extinct, how can this formalism be implemented given that detailed balance would break down? Are there ways to extend the current mainstream nonequilibrium physics so that an even more complete second law is achieved and, from it, derive appropriate bounds on entropy production for the different replicator architectures?

Also in Chapter 2, by considering a very general set-up for systems capable of harvesting energy from a nonequilibrium source, we have demonstrated the existence of universal bounds on the power yield in the space of all possible controlled energy transduction architectures. Our proof is grounded on a rigorous analysis of the interplay between energy extraction and information-entropic forces in molecular energy harvesting systems. We applied our formalism to basic architectures inspired by the logic of early biochemical cycles and arrived at following two major conclusions (1) early living systems biochemistry are constrained in achieving operational control of energy sources and (2) larger cycles, which could have been reachable in chemistry space, might have been selected against due to the higher entropic costs of maintenance as shown by (2.16) and (2.18) and their scaling behavior with n , the cycle size.

Further research arising from these second set of results should be directed at answering questions like:

1. The study Pareto-fronts on the trade-off landscape given by (2.15). Can these give rise to phase transitions [SS13]? And, is so:
 - a) Is there a particular type transitions to be expected?
 - b) Are there general restrictions that shape of the objective function?
2. If the system contains some nonlinearities, e.g., in the manner in which energy is extracted, then, does the regime analysis provided (LR, M and FE regimes) still hold?

3. Is it then possible to model major transitions on energy harvesting (e.g. endosymbiosis) using our framework?
4. In absence of a good criterion for the separation of degrees of freedom, are there endogenous processes that self-organize the system into *agent* and *environment*? If so, then how does this reflect on the type of optimization problem proposed here?

Finally, as usual in the any formal framework pertaining to statistical physics, it is important to recall that all of our underlying theoretical background is ultimately based a theory of probability [Jay65]. This observation is useful when aiming to translate advances of the kind we proposed here into other paradigms, such as the Statistical Physics of Evolution [Pel11]. Here, thermal baths are replaced by mutational backgrounds and the degrees of freedom are substituted by genes or bit-arrays, while the fitness landscapes take the role of a free-energy functional. In this context, an extension of our formalism might lead to a better understanding of universal bounds on evolutionary speed limits, a topic that has recently gained a lot of attention [SFS18, Ito18, ID20, IDC+21].

5.2. Precursor patterns in cooperative systems

In Chapter 3, we have turned to *cooperation* as a universal feature of METs. We have borrowed the tools and methods of stochastic processes to construct models of two and multiple species of cooperators. We have considered a minimal set of assumptions, which led us to the analysis of a Neutral theory of cooperative ecosystems. Out of this analysis, we concluded that at low migration rates:

- The system spontaneously *breaks* the symmetries imposed by neutrality and gives rise to a two-level description of the extant species: *outliers* and *core cooperators*. The emerging patterns reveal that:
 - A species belonging to the outliers set is short-lived and is kept at low abundance rates.

- A species belonging to the cooperators core remains relatively stable for long periods of time at large abundance rates, and fluctuations to extinction are rare.
- Equilibrium between migration and extinction is heavily dominated by the long relaxation times observed in the closed system of cooperators. Out of this equilibrium emerges a non-linear relation between the extant number of species (R), the system's size (N), the external species diversity (S) and the migration rate (μ). Such non-trivial relation is the object of further study and it is central to the understanding of how cooperative ecosystems assemble and self-stabilize.

In summary, our findings highlight the strength of the stochastic effects on a system of cooperative agents. In particular, it shows that there are precursor patterning phenomena that shape how the system assembles *before* genetic control or other selective forces take any action. This is a characteristic feature of what complex systems scholars like Stuart Kauffman dubbed 'order for free' [Kau93]. In short: the system of cooperators, once established, construes a scaffold upon which adaptation will operate. This is an example of how the physics of stochastic processes may facilitate the appearance of a MET.

5.3. Towards a unifying theory of cognitive networks

In Chapter 4, we reviewed the historic approach to cognitive networks from simple neural models to the various biological cognitive systems of all types (from ant colonies to slime mould). We highlight the statistical physics framework that model such systems in a unifying way, and provide a new characterization based on a morphospace with axes spanning the *physical state* (liquid to solid), *degree of development* and the *computational complexity* of any given cognitive network. These are represented

qualitatively and offer a holistic view of the space of all possible cognitive networks, which allows us to consider cases that are not explored by evolution, and, if so, ask why.

Throughout Chapter 4 we emphasize the role of criticality, as introduced in Chapter 1, and argue that it plays an essential role in matching the collective system's dynamics and the emerging computational complexity at the system's level, which corresponds to two of the theoretical pillars of METs considered in Chapter 1 (see Figure 1.5).

On the other hand, we study another limiting case of a highly liquid system, i.e. the immune system. We notice the following (see also Table 1 and Section 3b(vii) in [PS19]):

- (a) Connectivity is physically *liquid*, i.e., there's a high degree of interaction between the agents. Following the spin-glass metaphor 4.4, the topology of (i, j) is closer to a clique than a sparse network. However, as a counterpart, the set of $\{J_{ij}\}$, which correspond to the epitope-paratope couplings, becomes highly specific.
- (b) Memory and learning processes are mediated by a combination of population-based set of attractors plus a synaptic (epitope-paratope) based system of allocation. This parallels with the degeneracy of memory states observed in ant colonies, particularly in their division of labour (see Figure 4.4).
- (c) Following the analysis of [BA10] (see Section 3b in [PS19]), a dual classification of idiotypes emerges: strongly and weakly interacting. This heterogeneity is exploited by the system in solving the self vs. non-self identification problem. This heterogeneity in interaction weights constitutes a strong analogy to the heterogeneity found in the cortex neuron firing frequency distribution [BM14, LKR11], which reveals two types of neurons at play: generalists (high activity) and specialists (low activity). In contrast with the previous argument in (b), this reveals a connection the immune system with cortical networks.

These observations lead to an astonishing conclusion, namely that, despite the underlying physical support of either system (the brain cortex and the immune system), evolution has shaped their computational strategies in strikingly similar ways. While, at the same time, the immune system uses population proliferation processes to determine its functional state, specially when activating its antigenic response. This reinforces the hypothesis that not every kind of computational architecture is possible or, at least, not easily reachable by evolutionary means.

We finish this discussion with a conjecture. Consider any system made out of a large number of agents that interact pair-wisely and do so under three constraints:

- **Interactions:** connections between two agents, (i, j) , are established and broken at some rate σ . This relates with the *liquidity* of the system.
- **Weights:** interaction forcing couplings, J_{ij} , have some degree of heterogeneity¹, γ . This relates to the synaptic adaptability and specificity of the system.
- **Activity:** a characteristic activity decay timescale given, τ , which determines the microscopic computational timescale of the agents²

Then, our conjecture can be heuristically formulated as follows:

For any brain-like system to be poised at an optimal regime for cognition, then

$$\mathbf{c} := \sigma\gamma\tau = \mathcal{O}(1). \tag{5.1}$$

We term this conjecture the *liquid brain hypothesis*, and it states that a system in a optimal computational regime will display a trade-off governed

¹Here, we do not define this parameter in any explicit form, but only heuristically. A proper definition of heterogeneity in cognitive networks is left for future work.

²Recall this τ was used in Sections 4.1.1 and 4.3 when we discussed about criticality in cognitive networks.

by a constant c , which is adimensional, of order one, and its value will depend on each system considered and the methods used to determine $\{\sigma, \gamma, \tau\}$. We posit that, for a consistent definition of these measures, most (if not all) evolved cognitive networks will operate at the manifold defined by (5.1).

To make matters clear, consider a few heuristic examples:

- Suppose a system, like the cortex, for which we have some degree of heterogeneity in the neural weights, γ . However, the rate of establishing new interactions, σ is much slower than the neural activity decay rate, τ^{-1} . Thus, we expect combination of the three terms gives rise to some constant c of order 1.
- Similarly, ants of the same colony may display homogeneous interaction couplings, i.e. a simple ant-to-ant pheromone signal, which means that $\gamma \ll 1$. To attain a sufficient level of computational capacity, ant colonies show a fast interaction rate when compared to its activity decay rate, i.e. $\sigma\tau \gg 1$, which compensates the RHS of (5.1). As is the case explores in Section 4.3, here σ will too depend on the colony density, ρ .

To proof our liquid brain hypothesis requires a more rigorous framework. Such a mathematical theory will necessarily combine the aspects discussed in Chapter 4 regarding self-organization towards criticality, spin glass statistics and reaction-diffusion patterns in different spatial layouts. Other lessons from Chapter 3 regarding cooperation and population distributions may also play an important role. This constitutes a program of research that will doubtlessly be part future research efforts.

Chapter 6

CONCLUSIONS

The goal of this PhD project was to advance towards a unifying framework for understanding the complexity shifts pertaining to the major evolutionary transitions. This framework is based on three pillars: nonequilibrium thermodynamics, stochastic processes and ecological assembly and the theory of information and computation.

In Chapter 1, we set out to tackle three concrete objectives proposed as *key questions* (Section 1.4), the first of which relates to the bounds on early living systems imposed by thermodynamic laws. Here, we found that:

- Macroscopic replicator kinetics (hyperbolic, parabolic or simple) and the internal minimal entropy production values are intrinsically related via the extended 2nd law and the process of coarse-graining.
- Entropic bounds on each replicator type reveal a thermodynamic metric with which one can argue about the feasibility of replicators over the course of evolution through irreversibility arguments, rather than kinetic ones.
- Energy harvesting can be constrained by the information or entropy reducing costs of controlling internal cycles in a system to correlate with the external nonequilibrium source of power, the environment.

- Information can be used as a *catalyzer* in an energy harvesting system, rather than *fuel*. The latter use is limited to produce energetic yields of around $\sim k_B T \ln 2$, while the former is unbounded and depends on the system at hand.
- Early biochemical cycles were likely constrained by their information-associated costs. This results in a non-adaptive force towards smaller cycles as the primordial metabolic networks at the Origins of Life.

The second key question concerns with stochastic-induced constraints on the assembly of cooperative ecosystems, and the potential role played by these as precursor stages for selection to operate on top. Here, we found that it is possible to build a neutral theory of cooperator ecosystems, which display the following properties:

- System states decay in exponential time with an argument that scales as the system size, N .
- The long relaxation times, which correspond to long extinction processes, are nonlinear with S and belong to a certain class of glassy systems.
- Understanding diversity and assembly requires in cooperative ecosystems requires a strong handle on its extinction dynamics.
- Simulations show that a two-level system emerges ‘for free’ out of the stochastic noise, namely *core cooperators* (with high degree of stability and perseverance) and *outliners* (low species lifetime and large turnover rates). We argue that such phenomenon constitutes a strong precursor for cooperative systems to turn into a larger scale unit, a major universal feature of METs.

The third key question pertains to the evolvability and feasibility of possible cognitive architectures and their constrains both physical and computation. We based our discussion on the conceptualization brought by in the special issue on ‘Liquid brains, solid brains’ [SMF19], that took place at the Santa Fe Institute in 2018. In our analysis, we found that:

- Criticality and self-organization around criticality appear as a universal feature across known cognitive architectures, which leads us to extend the *critical brain hypothesis* into the larger space of cognitive networks which includes liquid brains.
- Stark physical differences exist between the memory allocation mechanisms present in solid vs. liquid neural nets. In particular, solid brains display a high degree of heterogeneity and memory is encoded in the individual neurons; while liquid brains like an ant colony use a distributed memory system based on population levels, which implies a high degeneracy of states with respect to the individual neurons (ants).
- It is possible to arrive at a conjecture about a trade-off that any type of cognitive system undergoes in optimizing its computational or functional targets, i.e. that there is a balance between interaction-formation rates, heterogeneity of weights and natural activity decay time of a neuron. We posit that a rigorous definition of this balancing is possible and captures a universal pattern in all cognitive neural systems reachable by evolutionary processes.

Chapter 7

PAPERS

This chapter includes the four papers that produced out of the Thesis project that scaffold this dissertation:

- Piñero, J., & Solé, R. (2018). Nonequilibrium entropic bounds for Darwinian replicators. *Entropy*, **20**(2), 98. DOI: 10.3390/e20020098
- Piñero, J., Solé, R., & Kolchinsky, A. (2023). Universal bounds and thermodynamic tradeoffs in nonequilibrium energy harvesting *arXiv*, 2303.04975. DOI: 10.48550/arXiv.2303.04975
- Piñero, J., Redner, S., & Solé, R. (2022). Fixation and fluctuations in two-species cooperation. *Journal of Physics: Complexity*, **3**(1), 015011. DOI: 10.1088/2632-072X/ac52e7
- Piñero, J., & Solé, R. (2019). Statistical physics of liquid brains. *Philosophical Transactions of the Royal Society B*, **374**(1774), 20180376. DOI: 10.1098/rstb.2018.0376

Article

Nonequilibrium Entropic Bounds for Darwinian Replicators

Jordi Piñero ^{1,2,*,†}  and Ricard Solé ^{1,2,3,*,†}

¹ ICREA-Complex Systems Lab, Universitat Pompeu Fabra, Doctor Aiguader 88, 08003 Barcelona, Spain

² Institut de Biologia Evolutiva, Psg. Barceloneta 37-49, 08003 Barcelona, Spain

³ Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

* Correspondence: jordi.pinero@upf.edu (J.P.); ricard.sole@upf.edu (R.S.); Tel.: +34-93-316-0000 (R.S.)

† These authors contributed equally to this work.

Received: 25 November 2017; Accepted: 26 January 2018; Published: 31 January 2018

Abstract: Life evolved on our planet by means of a combination of Darwinian selection and innovations leading to higher levels of complexity. The emergence and selection of replicating entities is a central problem in prebiotic evolution. Theoretical models have shown how populations of different types of replicating entities exclude or coexist with other classes of replicators. Models are typically kinetic, based on standard replicator equations. On the other hand, the presence of thermodynamical constraints for these systems remain an open question. This is largely due to the lack of a general theory of statistical methods for systems far from equilibrium. Nonetheless, a first approach to this problem has been put forward in a series of novel developments falling under the rubric of the extended second law of thermodynamics. The work presented here is twofold: firstly, we review this theoretical framework and provide a brief description of the three fundamental replicator types in prebiotic evolution: parabolic, malthusian and hyperbolic. Secondly, we employ these previously mentioned techniques to explore how replicators are constrained by thermodynamics. Finally, we comment and discuss where further research should be focused on.

Keywords: evolution; replicators; life; entropy; thermodynamics

1. Introduction

Biology follows the laws of physics, and yet it remains distinctive from many standard physical systems in a number of ways. In the first place, life's self-replicating mechanisms stand as a major difficulty when approaching it from a simple physical setup. On the other hand, life too differs from physics in its computational nature: all living forms conduct some sort of computation as a crucial component of their adaptive potential [1]. The success of life over chemistry is largely associated to the emergence of prebiotic molecular mechanisms that, in turn, allowed for a template-based landscape to become dominant over the whole biosphere. How this took place is one of the most fundamental questions in science [2–4].

Life forms are out-of-equilibrium structures capable to employ available matter, energy and information to propagate some type of identity. Most theoretical approaches to the evolution of replicators have been grounded on a kinetic description. Under such framework, interactions between (typically molecular) agents are represented by nonlinear differential equations, known as replicator equations [5]. They provide a deterministic view of Darwinian dynamics. However, as pointed out by Smith and Morowitz, “the abstraction of the replicator, which reduces Darwinian dynamics to its essentials, also de-emphasizes the chemical nature of life” [6]. The same can be concluded in relation with the lack of a thermodynamical context. Despite early efforts towards the development of a physics of evolutionary dynamics [6–9] a more satisfactory formalism has yet to emerge. In particular,

life propagation processes require an entropy production and balance equations can be defined [9–11]. However, a more general non-equilibrium statistical physics approach suitable for the problem of self-replication has been missing until recently [12–19]. How can this novel approach apply to the fundamental problem of replicator dynamics in the early stages of Life on Earth? Beyond the self-replicating potential of cells and molecules, several replication strategies are at work in living systems, also involving multiple scales [20–22]. The basic growth dynamics followed by each class has remarkably different consequences for selection. The simplest class is the Malthusian (exponential) growth dynamics exhibited by cellular systems growing under unlimited resources. Two other types of replicators are observed in Nature. One is associated to the emergence of cooperation dynamics, with different classes of replicators helping each other and forming a mutualistic assembly [23]. The second is related to a template-based replication mechanism that we can identify in living systems as the standard mechanism of nucleic acid replication. This mechanism has been shown to lead to the “survival of everyone”: it provides a mechanism capable of sustaining very diverse populations of replicators [24–26].

From the physics perspective, these systems involve large number of internal degrees of freedom interacting in an out-of-equilibrium context. In turn, this interplay in the microscopic level leads to a macroscopic emergent (coarse-grained) dynamics. A thermodynamical connection between these two levels can be made following the statistical physics methods cited above. The work presented here is an attempt to delineate these fundamental thermodynamical constraints for the three elemental types of prebiotic replicators.

2. Entropic Bounds for Replicators

Let us begin by reviewing the theoretical framework upon which the analysis of the problem will unfold [15,17–19]. Here, we outline a simplified version of theoretical basis behind this non-equilibrium approach. We also comment on the generalizations of the so-called extended second law [19]. Then, we summarize the elemental classes of replicators and their essential aspects [21], together with a series of implications regarding selection and adaptation. Finally, we lay out an approach to the question of how non-equilibrium thermodynamical bounds arise in these types of systems and how such constraints might have affected early evolutionary scenarios.

2.1. The Extended Second Law

Consider a classical time-evolving system described by its microscopical trajectory in the phase space $x(t) \in \Omega$ plus a set of controlled parameters $\lambda(t)$ evolving in a time interval $t \in [0, \tau]$ that act like external drivers for any given trajectory. Assume that the system remains in contact with a heat bath at temperature $T = 1/\beta$ throughout the entire trajectory. Denote the transition probability from a microscopical state x to y in the time interval ϵ by $\pi_\epsilon[x \rightarrow y]$. Now, if we slice time as $t_{i+1} - t_i = \epsilon$, with $t_n = \tau = n\epsilon$ and $t_0 = 0$, then, for sufficiently small ϵ , the microscopical reversibility condition implies [13,14]:

$$\frac{\pi_\epsilon[x^*(\tau - t)]}{\pi_\epsilon[x(t + t_{n-1})]} \dots \frac{\pi_\epsilon[x^*(t_1 - t)]}{\pi_\epsilon[x(t)]} = \exp \left\{ -\beta \sum_{i=0}^{n-1} Q_{i \rightarrow i+1}^b \right\}, \tag{1}$$

where the superscript * denotes momentum-reversed microstates, and $Q_{i \rightarrow i+1}^b$ denotes the heat exchange in going from from states $x(t_i)$ to $x(t_{i+1})$ as measured from the heat bath. Heuristically, (1) is interpreted as the composed detailed balance condition on each time-slice of the trajectory $x(t)$ (see Figure 1a). This can be represented by the functional relation:

$$\frac{\pi_\tau[x^*(\tau - t)]}{\pi_\tau[x(t)]} = \exp \{ -\beta Q_b[x(t)] \}. \tag{2}$$




Figure 1. Scheme of the formal approach to expressions (1)–(4). (a) A time-discretization is implemented in order to characterize the microscopical reversibility condition; (b) A qualitative scheme of possible trajectories between macrostates on the global phase space. The macroscopic coarse-grained states, \mathcal{A} (dark shaded region) and \mathcal{B} (light shaded region) are defined as disjoint ($\mathcal{A} \cap \mathcal{B} = \emptyset$) sections on the phase state Ω . The set of forward paths of duration τ constrained to start in \mathcal{A} and finish in \mathcal{B} is denoted by \mathbf{x}_τ .

Next, let us introduce two macrostates which can be interpreted as two disjoint sections of the phase space, $\mathcal{A}, \mathcal{B} \subset \Omega$ (see Figure 1b). Let us introduce notation for macrostate bounded trajectories in Ω by defining the set of forward trajectories $\mathbf{x}_\tau = \{x(t), t \in [0, \tau] \mid x(0) \in \mathcal{A} \wedge x(\tau) \in \mathcal{B}\}$, i.e., the set of possible trajectories subject to condition that the initial microstate is in \mathcal{A} and the final must be in \mathcal{B} . Then, construct the formal coarse-grained transition rate from \mathcal{A} to \mathcal{B} as

$$\Pi_\tau(\mathcal{A} \rightarrow \mathcal{B}) = \int_{\mathbf{x}_\tau} \mathcal{D}[x(t)] \pi_\tau[x(t)], \tag{3}$$

while, equivalently, denote $\mathbf{x}_\tau^* = \{x^*(\tau - t), t \in [0, \tau] \mid x^*(\tau) \in \mathcal{B} \wedge x^*(0) \in \mathcal{A}\}$ as the set of reversed macrostate bounded trajectories, driven by the reverse protocol $\bar{\lambda}(\tau - t)$ (details on the derivation can be found in [19]), and compute the inverse coarse-grained transition rate from \mathcal{B} to \mathcal{A} as

$$\Pi_\tau(\mathcal{B} \rightarrow \mathcal{A}) = \int_{\mathbf{x}_\tau^*} \mathcal{D}[x^*(\tau - t)] \pi_\tau[x^*(\tau - t)]. \tag{4}$$

Here onwards, let use bracket notation $\langle \cdot \rangle$ to denote averages over forward paths \mathbf{x}_τ . Under this theoretical framework, it can be shown [17,19] that the following relation must hold:

$$\left\langle \exp \left\{ -\Delta H[x(t)] - \beta Q_b[x(t)] + \log \left[\frac{\Pi_\tau(\mathcal{A} \rightarrow \mathcal{B})}{\Pi_\tau(\mathcal{B} \rightarrow \mathcal{A})} \right] \right\} \right\rangle = 1, \tag{5}$$

where we have defined the path-dependant observable:

$$\Delta H[x(t)] = -\log \left[\frac{p_\tau(x(\tau))}{p_0(x(0))} \right], \tag{6}$$

with $p_\tau(x(\tau))$ and $p_0(x(0))$ standing for the probability of landing at a certain $x(\tau) \in \mathcal{B}$ at time $t = \tau$ and departing from $x(0) \in \mathcal{A}$ at time $t = 0$. Notice that (6) is a functional that depends on the boundary conditions of the trajectory $x(t)$. Let us define,

$$\beta \mathcal{W}[x(t)] \equiv \Delta H[x(t)] + \beta Q_b[x(t)], \tag{7}$$

as a functional observable over the sample of forward paths \mathbf{x}_τ . On the one hand, a first order expansion on (5) imposes the following boundaries to the fraction of the coarse-grained transition rates:

$$\log \left[\frac{\Pi_\tau(\mathcal{A} \rightarrow \mathcal{B})}{\Pi_\tau(\mathcal{B} \rightarrow \mathcal{A})} \right] \leq \beta \langle \mathcal{W}[x(t)] \rangle = \langle \Delta H[x(t)] \rangle + \beta \langle Q_b[x(t)] \rangle . \tag{8}$$

This results implicitly allude to the Landauer bounds on heat production for bit erasure [27–29]. Inequality (8) constraints the irreversibility of the macroscopic process $\mathcal{A} \rightarrow \mathcal{B}$ with respect to the average generalized entropy produced internally, ΔH , and externally (into the bath), βQ_b , and it is dubbed the Extended/Bayesian Second Law (ESL) [17,19]. One interpretation is that macroscopic irreversibility increases the minimum dissipated energy during the process $\mathcal{A} \rightarrow \mathcal{B}$. Interestingly, expression (8) formalizes a bound on entropy production in relation to the coarse-grained properties of the process, such as the macroscopic transition rates. This result is of particular interest since, under many experimental circumstances, these are the only measurable quantities for a given system. We will come back to this point in the following sections.

On the other hand, a general perturbative analysis using the cumulant expansion [30] onto (5) leads to

$$\log \left[\frac{\Pi_\tau(\mathcal{A} \rightarrow \mathcal{B})}{\Pi_\tau(\mathcal{B} \rightarrow \mathcal{A})} \right] = \sum_{l \geq 1} (-\beta)^{l-1} \frac{\omega_l}{l!} , \tag{9}$$

where ω_l stands for the l -th cumulant of the distribution of $\beta \mathcal{W}[x(t)]$. In fact, (9) allows for a more sophisticated view of

$$\log \left[\frac{\Pi_\tau(\mathcal{A} \rightarrow \mathcal{B})}{\Pi_\tau(\mathcal{B} \rightarrow \mathcal{A})} \right] = \beta \langle \mathcal{W} \rangle - \Phi_\tau(\beta) , \tag{10}$$

where, formally

$$\Phi_\tau(\beta) = \frac{\beta^2}{2} \langle \mathcal{W} \rangle_c^2 - \frac{\beta^3}{6} \langle \mathcal{W} \rangle_c^3 + \dots \tag{11}$$

with the subscript c indicating cumulant expressions. Combining Equations (5) and (9), it can be shown that $\Phi_\tau \geq 0$. Indeed, Φ_τ is a measure the fluctuations of the distribution associated to observable $\mathcal{W}[x(t)]$. Thus, Equation (10) represents an extended fluctuation-dissipation theorem, where the LHS reflects the macroscopic (coarse-grained) irreversibility property and the RHS a balance between dissipated work and fluctuations over the \mathbf{x}_τ sample.

This result is of particular interest when a system is arranged such that a choice between two macroscopical end-states is forced. In such cases, fluctuation discrepancies might break symmetry thus favoring certain macroscopical transitions or suppressing others [18].

Moreover, these theoretical results can be generalized to less constrained versions of the ESL where no equilibrium trajectory end-points are required plus the system needs not to be at a fixed temperature, eventhough there is still contact with a heat bath (cf. [19]). Under this generalized lens, relations (5), (8)–(10) are formally equivalent, only now the space of possibilities over which averages are taken is constrained by the implemented coarse-grain. On the other hand, this implies that the operators in (7) are too redefined owing to the coarse-graining imposed in the system.


In the following sections we will revisit the paradigm of prebiotic replicators, and focus on how to minimally embed this problem into the formalism discussed above. Subsequently, we will argue how these entropic constraints may have coupled to prebiotic selection and added pressure to in an evolutionary context.

2.2. Replicators & Reproducers

Several fundamental replication strategies are at play in living systems. These strategies are present in multiple scales, from molecular replicators to cells and beyond. Each class of replicating


agent is characterized by a kinetic pattern, which dynamics entail distinct selective implications. Here, we will focus on three characteristic replicator classes [20,21].

Simple replicators: commonly known as Malthusian agents, correspond to systems whereby a single component A is capable of making a copy of itself by using the available resources, namely E , generating a certain waste product, W . Schematically,



Assuming a large repository of resources, the kinetics of this process can be reduced to a linear dynamical equation (see Table 1). Systems following this mechanism obey exponential growth laws.

Hyperbolic replicators: one of the most relevant novelties in evolution [31,32] is the concept of autocatalysis. This mechanism is a precursor of self-replicating entities that largely define the nature of living structures. It has been put forward by several authors [33–36] as a central process in the chemistry of prebiotic systems involving the emergence of cooperative agents (see Figure 2a).



Again, under well-mixed and unlimited resource conditions, the hyperbolic replicator kinetics is reduced to a second order equation (see Table 1). Autocatalytic growth is characterized by displaying a finite-time singularity at $t_c = 1/hx_0$ [21].





Figure 2. Hyperbolic and parabolic replicators. In (a) we display a simplified scheme of an experimental implementation of a catalytic set of ribozymes forming a cooperative loop. Here each component of the system helps the next to replicate. Dashed lines indicate weaker catalytic links (modified from [37]). The parabolic system outlined in (b) is based on complementary (template) peptide chains involving a ligation mechanism (adapted from [38]).

Parabolic replicators: this type of replicator arises from a combination of molecular reactions. In particular, oligonucleotides are known to exhibit such behaviour [26,39–41]. The minimal scheme where this particular dynamics is observed consists of the set of processes (see Figure 2b).



which, under conditions $a \gg b \gg c$ is reduced to a parabolic law $\dot{x} = \rho\sqrt{x}$, where x denotes the total concentration of the molecular component A regardless of the configuration, it being either associated (AA) or dissociated (A) (see Appendix A). Parameter $\rho = c\sqrt{2b/a}$.

Table 1. Summary of the minimal expressions for the kinetics of the three replicator classes discussed above. We have denoted as x the gross concentration of replicating molecules A , independently of the configuration.

Replicator Class	Reaction Scheme	Effective Dynamics
Simple	$A + E \rightarrow A + A$	$\dot{x} = gx$
Hyperbolic	$A + A + E \rightarrow A + A + A$	$\dot{x} = hx^2$
Parabolic	$A + E \rightarrow AA \leftrightarrow A + A$	$\dot{x} = \rho x^{1/2}$

2.3. Coarse-Grained Dynamics of Replicators

The dynamics of the three types of replicators discussed above are taking place on the macroscopic level. Molecular replicators encapsulate a whole system rich in complexity and structure, thus the measurable transition rates, such as g, h or ρ above, are emergent features of the interplay of the many internal degrees of freedom of the system. However, the statistical properties of these phenomena are non-ergodic, since replicating is constrained by an initial and a final coarse-grained states. As discussed in Section 2.1, averages reflecting the macroscopic transition rates are taken over a section of the space of possibilities, specifically over the subset of possible microscopical trajectories with an initial number of replicators $n - 1$ and a final number n (given a time scale τ), as detailed below.

To begin with, suppose that a system is composed of a fixed number of molecular templates or chains, N , which can either be internally ordered such that they behave as a replicators (A), namely active chains, or simply act as substrate (E), namely inactive chains. The goal here is to define an unambiguous coarse-graining measure capable of distinguishing two meaningful macroscopic states of the system. To do so, we will consider three such systems which replicators' act accordingly with the three replicator classes summarized in Table 1. We will also suppose that all replicators undergo equivalent decay processes. This assumption is taken so that we are able to probe the thermodynamical bounds purely for the processes involving replication. For simplicity, we use open systems (source flowing in) but finite (fixed total number of particles).

Following a markovian approach [42,43], each set of reaction rules allows defining transition probabilities and a master equation that in general will read:

$$\frac{dP(n,t)}{dt} = \sum_{m \neq n} \omega(n|m) P(m,t) - \sum_{m \neq n} \omega(m|n) P(n,t), \tag{15}$$

which gives the probability $P(n,t)$ of observing n active chains at time t . Here the $\omega(i|j)$ terms introduce the transition probabilities associated to each rule, duely determined by the corresponding Malthusian, hyperbolic and parabolic cases. The three urn-like systems analysed here are chemostat models since, when an element (replicator) decays, it is replaced by newly available source particles E (see Appendix B for details). In summary,

$$\frac{dP(n,t)}{dt} = g \left(\frac{n}{N}\right) \left(1 - \frac{n}{N}\right) [P(n-1,t) - P(n,t)] - \delta \left(\frac{n}{N}\right) [P(n,t) - P(n+1,t)], \tag{16}$$

$$\frac{dP(n,t)}{dt} = h \left(\frac{n}{N}\right)^2 \left(1 - \frac{n}{N}\right) [P(n-1,t) - P(n,t)] - \delta \left(\frac{n}{N}\right) [P(n,t) - P(n+1,t)], \tag{17}$$

$$\frac{dP(n,t)}{dt} = \frac{bc}{2a} \left(\sqrt{1 + \frac{4an}{bN}} - 1\right) \left(1 - \frac{n}{N}\right) [P(n-1,t) - P(n,t)] - \delta \left(\frac{n}{N}\right) [P(n,t) - P(n+1,t)]. \tag{18}$$

Notice that (16)–(18) are non-equilibrium macroscopic representations of the replicating dynamics. Here, the internal interactions that produce the effective behaviour described by the previous set of equations are all integrated out into its corresponding coupling constants. Thus, within this macroscopical framework we shall define the phase space subsets:

- \mathcal{A} —state in which the system contains a total number of $n - 1$ active chains.
- \mathcal{B} —state in which the system contains a total amount of n active chains.

Let us focus on the explicit bounds given by the LHS in expression (5). We first introduce notation for these lower entropic bounds,

$$LEB_r(x) := \log \left[\frac{\Pi_\tau(\mathcal{A} \rightarrow \mathcal{B})}{\Pi_\tau(\mathcal{B} \rightarrow \mathcal{A})} \right], \tag{19}$$

where the subscript $r \in \{\mathbf{s}, \mathbf{h}, \mathbf{p}\}$ indicates the replicator type (simple, hyperbolic and parabolic respectively), while $x := n/N$ in each case. Therefore, considering that the transition rates $\Pi_\tau(\mathcal{A} \rightarrow \mathcal{B})$ and $\Pi_\tau(\mathcal{B} \rightarrow \mathcal{A})$ for the defined coarse-grained states \mathcal{A} and \mathcal{B} correspond to the prefactors in each master equation above,

$$LEB_s(x) = \log \left[\frac{g}{\delta}(1-x) \right], \quad LEB_h(x) = \log \left[\frac{h}{\delta}x(1-x) \right], \tag{20}$$

$$LEB_p(x) = \log \left[\frac{c}{\delta} \frac{\alpha}{x} \left(\sqrt{1 + \frac{2x}{\alpha}} - 1 \right) (1-x) \right], \tag{21}$$

where we have defined $\alpha := b/2a$. Finally, introduce notation $\Delta LEB(r|r') := LEB_r(x) - LEB_{r'}(x)$ in order to compare each replicator type. Hence, for \mathbf{h} and \mathbf{p} against \mathbf{s} we derive

$$\Delta LEB(\mathbf{h}|\mathbf{s}) = \log \left(\frac{h}{g}x \right), \tag{22}$$

$$\Delta LEB(\mathbf{p}|\mathbf{s}) = \log \left[\frac{c}{g} \frac{\alpha}{x} \left(\sqrt{1 + \frac{2x}{\alpha}} - 1 \right) \right], \tag{23}$$

while, $\Delta LEB(\mathbf{h}|\mathbf{p}) = \Delta LEB(\mathbf{h}|\mathbf{s}) - \Delta LEB(\mathbf{p}|\mathbf{s})$. Notice that, since all replicators decay mechanism has been chosen to be equivalent (see Appendix B), then relative bounds $\Delta LEB(r|r')$ are δ -independent. Figure 3a-f show various curves (22) and (23) against the density value x .

Focusing on the limiting cases where the lower bounds between distinct replicators coincide, $\Delta LEB(r|r') = 0$, it is possible to derive the density values for which the LEB for replicator r exceeds that of replicator r' and viceversa. This is an interesting exercise since minimal entropy production can provide a guideline for thermodynamically advantageous processes. Bare in mind that exploring LEBs does not include the full picture, as fluctuations can shift the average dissipated energy and unbalance irreversibility as discussed above (cf. [18]).

Thus, let us define the LEB crossover density $x_{rr'}$ from r -LEB dominance to r' -LEB dominance, or, simply, $\Delta LEB(r|r')|_{x_{rr'}} = 0$. Working with reduced variables $\tilde{h} := h/g$ and $\tilde{c} := c/g$ we derive $x_{rr'} = x_{rr'}(\tilde{h}, \tilde{c})$ following (22) and (23):

$$x_{sh} = \tilde{h}^{-1}, \quad x_{ps} = 2\alpha\tilde{c}(\tilde{c} - 1), \quad x_{ph}^3 + \frac{2\alpha\tilde{c}}{\tilde{h}} \left(x_{ph} - \frac{\tilde{c}}{\tilde{h}} \right) = 0, \tag{24}$$

where the equation for x_{ph} , the density value where LEB dominance shifts from parabolic hyperbolic is given in an implicit form (Algebraic analysis shows that the equation for x_{ph} contains a single real root.). On the other hand, $0 < x_{rr'} < 1$ must be held, as it stands for a density variable.

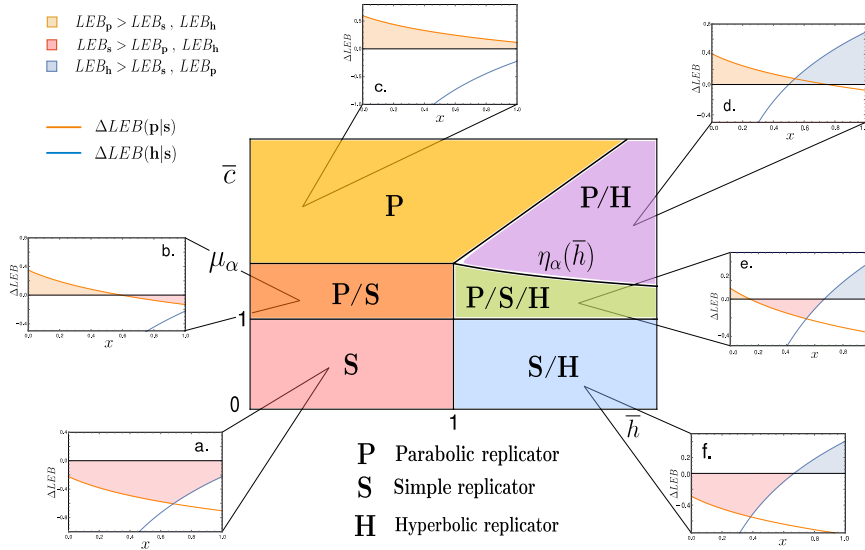


Figure 3. The central diagram corresponds to the space spanning the reduced variables (\bar{h}, \bar{c}) . We distinguish six phases depending on the dominance of the LEB of each replicator type, $\{S, H, P\}$. (a) **S**-dominant (the simple replicator LEB exceeds that of both parabolic and hyperbolic); (b) **P/S** i.e., at low densities, it is **S**-dominant, while, for $x > x_{ps}$ we observe **S** dominance; (c) **P**-dominant at all density values; (d) **P/H P** at low densities and **H**-dominant for $x > x_{ph}$; (e) **P/S/H** where the three replicators share dominance at some point, jumping orderedly at density values $x_{ps} < x_{sh}$; (f) here simple replicators have a higher LEB at low densities than parabolic ones, but hyperbolic ones take over at high densities, $x > x_{sh}$. Numerical values of (\bar{h}, \bar{c}) for each plot are: (a) (0.8, 0.8); (b) (1.42, 0.8); (c) (1.82, 0.8); (d) (1.5, 2); (e) (1.125, 1.5); (f) (0.75, 1.5), while $\alpha = 0.5$ for all graphs.

These considerations allow for a construction of a diagram (\bar{h}, \bar{c}) where space is divided into sections characterised by the replicator-types that display a dominant LEB. For instance, for $\bar{h}, \bar{c} < 1$ the simple replicator’s lower entropic production bound is always larger than the other two types, we denote this sector of the phase space by **S** (red shaded region in Figure 3). Most regions, however, will display dominance of entropy production by one type of replicators for a range of densities, and shift dominance over another type for another range of x values (see Figure 3b,d–f).

The lines separating sections of LEB dominance are given by the following set of inequalities, all derived from the results above:

$$P \Leftrightarrow \{ \bar{c} > 1 \ \& \ 0 < \bar{h} \} , \tag{25}$$

$$S \Leftrightarrow \{ \bar{c} < \mu_\alpha \ \& \ \bar{h} < 1 \} \cup \{ \bar{c} < \eta_\alpha(\bar{h}) \ \& \ \bar{h} > 1 \} , \tag{26}$$

$$H \Leftrightarrow \{ \bar{c} < \mu_\alpha \bar{h} \ \& \ \bar{h} > 1 \} , \tag{27}$$

with the associated functions

$$\mu_\alpha := \frac{1}{2} \left(1 + \sqrt{1 + \frac{2}{\alpha}} \right) , \quad \eta_\alpha(\bar{h}) := \frac{1}{2} \left(1 + \sqrt{1 + \frac{2}{\alpha \bar{h}}} \right) . \tag{28}$$

Notice that, in several patches of the space of parameters depicted in Figure 3, LEB dominance is dependent on specific density values. Also, $\Delta LEB(r|r')$ functions behave such that LEB dominance

always appears ordered as **P**, **S** and **H**, respectively. This ordered sequence can be understood as an indication of an underlying thermodynamical constraint for these pre-biotic replicating systems. Finally, notice that this analysis has been performed with fixed value of α . Nonetheless, shifting the values of this internal parameter does not substantially modify the structure of the phase space given in Figure 3, in fact, its topological arrangement will remain invariant.

Hence, from macroscopical considerations involving both coarse-grained values for the coupling constants $\{g, h, c\}$ and internal parameter α , we are able to derive a phase space compartmentalisation that allows a classification based on the lower (generalized) entropy production bounds for each replicator type. A qualitative tendency emerges from this picture: the parabolic replicator generates more entropy at low densities while so does the hyperbolic at high x values, leaving the simple replicator in between.

3. Discussion

A significant gap in our understanding of evolution, particularly in relation with early events and simple living systems, stems from the lack of a physical theory incorporating a thermodynamic description of replication dynamics. Self-replication stands as the one characteristic feature of living matter and its singular character was early appreciated by theoreticians when comparing cells and machines [44,45]. This work was an important step towards an understanding of the logic and computational nature of self-replicating agents. But a physical equivalent addressing the fundamental physics bounds to replication has been missing.

Recent work has addressed this problem revealing a powerful connection between entropy production and the transition probabilities underlying a stochastic, microscopic description [17,46]. Such connection can be efficiently exploited to analyse, under the coarse-graining described above, the general tendency of a Darwinian replicator to replicate itself. In this way, it is possible in particular to compare the efficiency of different classes of replicators by looking at their relative lower entropy bounds.

Instead of a direct comparison of the systems' measurable replication rates, this framework focuses on how, via a coarse-graining procedure, these parameters are resulting from the interplay of the many internal degrees of freedom. This technique ultimately leads to the estimation of the lower entropic bounds for each replicator. We interpret these non-equilibrium thermodynamic bounds as a consistent way of comparing and evaluating the likelihood of observing different classes of replicators. This is summarised in the phase diagram shown in Figure 3 where the relative dominance of each class is indicated. Notice that the analysis above does not involve competition between the replicator classes. All computations for the entropic bounds are done by considering the replicators to be evolving separately (see Appendix B for details).

Congruent approaches have been recently put forward following an equivalent theoretical formalism studying the non-equilibrium costs of production and destruction of polymers [47]. On the other hand, the present approach ought to be regarded as a minimal theoretical setup, and a number of issues can be raised. For instance, the fact that prebiotic systems might have exploited physical environments where sharp gradients are present, as it occurs with water-air interfaces [48]. Further developments in non-equilibrium statistical physics are needed in order to tackle these types of heterogeneities.

Even at this level of description, we can see how the coarse-graining predicts what to expect for the constraints operating on the classes of replicators in early evolutionary stages. The diagrams reveal the threshold conditions that would allow particular types of replicators to thrive or coexist in a competing scenario. In some domains only Malthusian dynamics are thermodynamically dominant, while, in others, parabolic replicators seem to be more efficient at generating entropy. Also, in some regions, a combination of parabolic and hyperbolic (cooperator) agents would share dominance. Overall, there is a robust characterisation of dominance related to the density of the system, revealing a preferential order as we move from low to high densities.

Future work should be aimed at the construction of theoretical microscopic models such that coarse-graining operations can be unambiguously defined and subsequent operations may be computed in order to obtain the emergent transition rates. This would yield a deeper understanding of both the coarse-graining process and how some biological systems seem to be able to operate at the edge of what is possible. Such an approach can lead to novel insights into the problem of how major evolutionary transitions (which are often tied to the emergence of novel classes of replicators) occur.

Acknowledgments: The authors thank Thomas Ouldridge for personal communication, and the Complex Systems Lab members for fruitful discussions. Special thanks to Gerda Taro for her inspiring ideas. This work was supported by the Botín Foundation by Banco Santander through its Santander Universities Global Division, the Spanish Ministry of Economy and Competitiveness, grant FIS2015-67616-P and the Santa Fe Institute.

Author Contributions: Jordi Piñero & Ricard Solé conceived, designed, performed the mathematical analysis, and wrote the paper.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

MDPI	Multidisciplinary Digital Publishing Institute
DOAJ	Directory of open access journals
ESL	Extended Second Law
LEB	Lower Entropic Bound
LHS	Left Hand Side
RHS	Right Hand Side

Appendix A

The argument for the effective kinetic law for the parabolic replicator goes as following: let $y = [AA]$ (concentration of associated molecules), and $z = [A]$ (concentration of dissociated molecules). Thus, define $x = 2y + z$ as the gross stoichiometric concentration of molecules of type A , regardless of configuration. Assuming that the time-scale of the replication reaction (here modulated by ratio c) is much larger than the association/dissociation processes, then, by focusing on the dynamics of replication, we can assume balanced equilibrium

$$by = az^2 \Leftrightarrow z = \sqrt{\frac{b}{a}} y^{1/2}. \quad (\text{A1})$$

Then, analysing the dynamics of the replication reaction, which is modulated by the parameter c ,

$$\frac{dy}{dt} = cz = c \sqrt{\frac{b}{a}} y^{1/2}, \quad (\text{A2})$$

while the kinetics for the gross concentration x is obtained by using (A1) as

$$\frac{dx}{dt} = 2\frac{dy}{dt} + \frac{dz}{dt} = \frac{cb}{2a} + 2c \sqrt{\frac{b}{a}} y^{1/2}, \quad (\text{A3})$$

but, as $a \gg b$, then the equilibrium of the association/dissociation reaction is very much unbalanced in favour of the associated molecular configuration AA , which implies that $x \approx 2y$. Thus, we conclude that the kinetics for x is given by

$$\frac{dx}{dt} = \frac{cb}{2a} + c \sqrt{\frac{2b}{a}} x^{1/2}. \quad (\text{A4})$$

Truncating at leading terms in (b/a) , we derive

$$\frac{dx}{dt} \approx \rho x^{1/2}, \tag{A5}$$

with $\rho = c \sqrt{\frac{2b}{a}}$.

Appendix B

Consider a well-mixed urn filled with N elements that can be characterised as dead or alive. Notice that such a characterisation embodies some kind of coarse-grained measure, since we are deliberately ignoring (integrating) all internal degrees of freedom for each element. Denote by $n < N$ the number of active (alive) elements in the urn at a given time t . In the following sections we will derive the coarse-grained (mesoscopic) time-dependent dynamics. Hence, for each replicator type, let us construct a master equation of the form

$$\frac{dP(n,t)}{dt} = \sum_{m \neq n} \omega(n|m) P(m,t) - \sum_{m \neq n} \omega(m|n) P(n,t), \tag{A6}$$

while restricting the dynamics to a first-step process and introducing a natural (single-particle) decay process modulated by parameter δ that will be equivalent to all replicating motifs. It is important to remark that these systems are implicitly open. This is because every time an active element turns into an inactive one what really is happening it is flowing out of the system and letting new source (E) flow in. In this sense, the proposed models are analogous to chemostats.

Appendix B.1

Beginning with the simple replicator, introduce the following rules (see Figure A1):

1. Pick an element of the urn at random.
2. If active, with probability g , pick a second element at random and (if not active) activate.
3. Pick an element at random again.
4. If active, with probability δ , deactivate.




Figure A1. A summary of the rules of replication in an urn model. Active chains are drawn as filled balls and inactive chains are white balls. (a) represents the action of selecting an active chain replicating following the simple replicator mechanism; (b) shows the replicating process of a hypercyclic replicator; (c) corresponds to the decay which, for the purpose of this work, is supposed to act equivalently in each replicator-type.

For the simple replicator, using the notation on Table 1, plus adding a single-particle decay process,

$$\omega(n|n-1) = \left(\frac{n-1}{N}\right) \left(\frac{N-n+2}{N}\right) g ; \quad \omega(n|n+1) = \left(\frac{n+1}{N}\right) \delta \tag{A7}$$

Then, for sufficiently large N , it is possible to approach the dynamics by

$$\frac{dP(n,t)}{dt} = g \left(\frac{n}{N} \right) \left(1 - \frac{n}{N} \right) [P(n-1,t) - P(n,t)] - \delta \left(\frac{n}{N} \right) [P(n,t) - P(n+1,t)] , \quad (A8)$$

Appendix B.2

Consider now the dynamics of hyperbolic replicators. Following the rules summarized in Table 1, introduce the algorithm (see Figure A1):

1. Pick an element of the urn at random.
2. If active, pick a second element at random.
3. If active, with probability h , pick a third element at random and (if not active) activate.
4. Pick an element at random again.
5. If active, with probability δ , deactivate.

Hence, restricting the dynamics to a first-step process, we may deduce the following transition probabilities (cf. [49,50])

$$\omega(n|n-1) = \left(\frac{n-1}{N} \right) \left(\frac{n-2}{N} \right) \left(\frac{N-n+3}{N} \right) h ; \quad \omega(n|n+1) = \left(\frac{n+1}{N} \right) \delta \quad (A9)$$

which, for $N \gg 1$, lead to the master equation

$$\frac{dP(n,t)}{dt} = h \left(\frac{n}{N} \right)^2 \left(1 - \frac{n}{N} \right) [P(n-1,t) - P(n,t)] - \delta \left(\frac{n}{N} \right) [P(n,t) - P(n+1,t)] , \quad (A10)$$

Appendix B.3

Finally, let us derive the macroscopical dynamics for a parabolic replicator by implementing the following set of rules on an urn of N elements (see Figure A2):

1. Pick an element of the urn at random. If active, then: (i) if in associated state (AA) then, with probability a , dissociate and iterate; (ii) if dissociated, pick a second element and, if active, with probability b , associate. Iterate this process until equilibrium is reached for association/dissociation reaction.
2. Pick an element of the urn at random. If active, pick a second element at random, if empty, with probability c , replicate.
3. Pick an element of the urn at random. If active, with probability δ , deactivate.




Figure A2. This diagram shows how the urn model of parabolic replicators is implemented. (a,b) correspond to the rapid association/dissociation reactions, which are supposed to equilibrate in much shorter time-scales than the replicating process, which is shown in (c), i.e., $\tau_0 \ll \tau_1$. The process of equilibration (left box) is iterated a large number of times before the loop goes into the replication process (right box).

The situation for the parabolic replicator is a peculiar one, for one thing, it involves two characteristic time-scales, a rapid one, concerning the association/dissociation process (see Appendix A), and the replication process. In order to approximate the transition rates let us define k as the number of associated pairs, AA , and m as the number of dissociated active elements in the urn, A . Let N be the total number of elements in the urn, including associated, dissociated and deactivated elements. Let us denote by n the total number of active elements, regardless of configuration, then, $n = 2k + m$. Now, assuming rapid equilibration of the association/dissociation reaction in (14),

$$\left(\frac{2k}{N}\right) b = \left(\frac{m}{N}\right)^2 a \Leftrightarrow m^2 = \frac{2b}{a} kN, \quad (\text{A11})$$

which can be related to the number n by

$$m^2 = \frac{b}{a} (n - m)N \Leftrightarrow m(n) = \frac{bN}{2a} \left(\sqrt{1 + \frac{4an}{bN}} - 1 \right), \quad (\text{A12})$$

where we neglect the negative root, as it is non-physical. Hence, it is now possible to construct a master equation for the first-step process of replication as in (14), with

$$\omega(n|n-1) = \left(\frac{m+1}{N}\right) \left(\frac{N-n+1}{N}\right) c; \quad \omega(n|n+1) = \left(\frac{n+1}{N}\right) \delta \quad (\text{A13})$$

which, for $N \gg 1$, and using (A12) lead to

$$\begin{aligned} \frac{dP(n,t)}{dt} &= \frac{bc}{2a} \left(\sqrt{1 + \frac{4an}{bN}} - 1 \right) \left(1 - \frac{n}{N} \right) [P(n-1,t) - P(n,t)] \\ &- \delta \left(\frac{n}{N} \right) [P(n,t) - P(n+1,t)], \end{aligned} \quad (\text{A14})$$

References

- Hopfield, J.J. Physics, computation, and why biology looks so different. *J. Theor. Biol.* **1994**, *171*, 53–60.
- Maynard-Smith, J.; Szathmáry, E. *The Origins of Life*; Oxford University Press: Oxford, UK, 1999.
- Dyson, F. *Origins of Life*; Cambridge University Press: Cambridge, UK, 1999.
- Kauffman, S.A. *The Origins of Order: Self-Organization and Selection in Evolution*; Oxford University Press: New York, NY, USA, 1993.
- Hofbauer, J.; Sigmund, K. *Evolutionary Games and Population Dynamics*; Cambridge University Press: Cambridge, UK, 1998.
- Morowitz, H.; Smith, E. *The Origin and Nature of Life on Earth: The Emergence of the Fourth Geosphere*; Cambridge University Press: Cambridge, UK, 2016.
- Kauffman, S.A. *Investigations*; Oxford University Press: New York, NY, USA, 2000.
- Babloyantz, A. *Molecules, Dynamics and Life: An Introduction to Self-Organization of Matter*; John Wiley & Sons: New York, NY, USA, 1986.
- Nicolis, G.; Prigogine, I. *Exploring Complexity: An Introduction*; W.H. Freeman & Co, Ltd.: New York, NY, USA, 1989.
- Glandsdorff, P.; Prigogine, I. *Thermodynamic Theory of Structure, Stability and Fluctuations*; John Wiley & Sons, Ltd.: London, UK, 1971.
- Wagensberg, J. Complexity versus uncertainty: The question of staying alive. *Biol. Philos.* **2000**, *15*, 493–508.
- Jarzynski, C. Nonequilibrium equality for free energy differences. *Phys. Rev. Lett.* **1997**, *78*, 2690–2693.
- Crooks, G.E. Nonequilibrium measurements of free energy differences for microscopically reversible Markovian systems. *J. Stat. Phys.* **1998**, *90*, 1481–1487.
- Crooks, G.E. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Phys. Rev. E* **1999**, *60*, 2721–2726.

15. Gomez-Marin, A.; Parrondo, J.M.R.; van den Broeck, C. Lower bounds on dissipation upon coarse graining. *Phys. Rev. E* **2008**, *78*, 011107.
16. Still, S.; Sivak, D.A.; Bell, A.J.; Crooks, G.E. Thermodynamics of prediction. *Phys. Rev. Lett.* **2012**, *109*, 120604.
17. England, J. Statistical physics of self-replication. *J. Chem. Phys.* **2013**, *139*, 121923.
18. Perunov, N.; Marsland, R.A.; England, J. Statistical physics of adaptation. *Phys. Rev. X* **2016**, *6*, 021036.
19. Bartolotta, A.; Carroll, S.M.; Leichenauer, S.; Pollack, J. Bayesian second law of thermodynamics. *Phys. Rev. E* **2016**, *94*, 022102.
20. Szathmáry, E. The origin of replicators and reproducers. *Philos. Trans. R. Soc. B* **2006**, *361*, 1761–1776.
21. Szathmáry, E. From replicators to reproducers: The first major transitions leading to life. *J. Theor. Biol.* **1997**, *187*, 555–571.
22. Solé, R. Synthetic transitions: Towards a new synthesis. *Philos. Trans. R. Soc. B* **2016**, *371*, 20150438.
23. Eigen, M.; Schuster, P. *The Hypercycle: A Principle of Natural Self-Organization*; Springer: Berlin, Germany, 1979.
24. Szathmáry, E.; Gladkih, I. Sub-exponential growth and coexistence of non-enzymatically replicating templates. *J. Theor. Biol.* **1989**, *138*, 55–58.
25. Scheuring, I.; Szathmáry, E. Survival of replicators with parabolic growth tendency and exponential decay. *J. Theor. Biol.* **2001**, *212*, 99–105.
26. Ellington, A.D. Origins for everyone. *Evol. Educ. Outreach* **2012**, *5*, 361–366.
27. Landauer, R. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.* **1961**, *5*, 183–191.
28. Bennett, C.H. The thermodynamics of computation—A review. *Int. J. Theor. Phys.* **1982**, *21*, 905–940.
29. Parrondo, J.M.R.; Horowitz, J.M.; Sagawa, T. Thermodynamics of information. *Nat. Phys.* **2015**, *11*, 131–139.
30. Zwanzig, R.W. High-temperature equation of state by a perturbation method. I. Nonpolar gases. *J. Chem. Phys.* **1954**, *22*, 1420–1426.
31. Szathmáry, E. Simple growth laws and selection consequences. *Trends Ecol. Evol.* **1991**, *6*, 366–370.
32. Maynard-Smith, J.; Szathmáry, E. *The Major Transitions in Evolution*; Oxford University Press: Oxford, UK, 1995.
33. Eigen, M. Self-organization of matter and the evolution of biological macromolecules. *Naturwissenschaften* **1971**, *58*, 465–523.
34. Eigen, M.; Schuster, P. The hypercycle: A principle of natural self-organization. Part A: Emergence of the hypercycle. *Naturwissenschaften* **1977**, *64*, 541–565.
35. Eigen, M.; Schuster, P. The hypercycle: A principle of natural self-organization. Part B: The abstract hypercycle. *Naturwissenschaften* **1977**, *65*, 7–41.
36. Eigen, M.; Schuster, P. The hypercycle: A principle of natural self-organization. Part C: The realistic hypercycle. *Naturwissenschaften* **1977**, *65*, 341–369.
37. Vaidya, N.; Manapat, M.L.; Chen, I.A.; Xulvi-Brunet, R.; Hayden, E.J.; Lehman, N. Spontaneous network formation among cooperative RNA replicators. *Nature* **2012**, *491*, 72–77.
38. Lee, D.H.; Granja, J.R.; Martinez, J.A.; Severin, K.; Ghadiri, M.R. A self-replicating peptide. *Nature* **1996**, *382*, 525–528.
39. von Kiedrowski, G. A self-replicating, hexadeoxy nucleotide. *Angew. Chem. Int. Ed. Engl.* **1986**, *25*, 932–935.
40. Zielinski, W.S.; Orgel, L.E. Autocatalytic synthesis of a tetranucleotide analogue. *Nature* **1977**, *327*, 346–347.
41. Paul, N.; Joyce, G.F. Minimal self-replicating systems. *Curr. Opin. Chem. Biol.* **2004**, *8*, 634–639.
42. Méndez, V.; Campos, D.; Bartumeus, F. *Stochastic Foundations in Movement Ecology: Anomalous Diffusion, Front Propagation and Random Searches*; Springer: Heidelberg, Germany, 2014.
43. Redner, S. *A Guide to First-Passage Processes*; Cambridge University Press: Cambridge, UK, 2001.
44. von Neumann, J.; Burks, A.W. *Theory of Self-Reproducing Automata*; University of Illinois Press: Urbana, IL, USA, 1966.
45. Sipper, M.; Reggia, J.A. Go forth and replicate. *Sci. Am.* **2008**, *18*, 48–57.
46. Andrieux, D.; Gaspard, P. Nonequilibrium generation of information in copolymerization processes. *Proc. Natl. Acad. Sci. USA* **2008**, *105*, 9516–9521.
47. Ouldrige, T.E.; ten Wolde, P.R. Fundamental costs in the production and destruction of persistent polymer copies. *Phys. Rev. Lett.* **2017**, *118*, 158103.
48. Griffith, E.C.; Tuck, A.F.; Vaida, V. Ocean-atmosphere interactions in the emergence of complexity in simple chemical systems. *Acc. Chem. Res.* **2012**, *45*, 2106–2113.

49. García-Tejedor, A.; Castaño, A.R.; Morán, F.; Montero, F. Studies on evolution and selective properties of hypercycles using a Monte Carlo method. *J. Mol. Evol.* **1987**, *26*, 294–300.
50. García-Tejedor, A.; Sanz-Nuño, J.C.; Olarrea, J.; de la Rubia, F.J.; Montero, F. Influence of the hypercyclic organization on the error threshold: A stochastic approach. *J. Theor. Biol.* **1988**, *134*, 431–443.



© 2018 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).

Universal bounds and thermodynamic tradeoffs in nonequilibrium energy harvesting

Jordi Piñero,^{1,2} Ricard Solé,^{1,2,3} and Artemy Kolchinsky^{4,*}

¹ICREA-Complex Systems Lab, Universitat Pompeu Fabra, 08003 Barcelona, Spain

²Institut de Biologia Evolutiva (CSIC-UPF), 08003 Barcelona, Spain

³Santa Fe Institute, Santa Fe, New Mexico 87501, United States

⁴Universal Biology Institute, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan

Many molecular systems operate by harvesting and storing energy from their environments. However, the maintenance of a nonequilibrium state necessary to support energy harvesting itself carries thermodynamic costs. We consider the optimal tradeoff between costs and benefits of energy harvesting in a nonequilibrium steady state, for a system that may be in contact with a fluctuating environment. We find a universal bound on this tradeoff, which leads to closed-form expressions for optimal power output and optimal steady-state distributions for three physically meaningful regimes. Our results are illustrated using a model of a unicyclic network, which is inspired by the logic of biomolecular cycles.

Introduction.— Various kinds of molecular systems, ranging from biological organisms to artificial devices, operate by continuously harvesting and storing energy from their environment [1–3]. In order to sustain energy harvesting, however, such systems maintain an organized nonequilibrium state, which itself requires costly self-maintenance processes such as error correction and self-assembly. In addition, such systems may be regulated by information flows about changing internal and environmental conditions, which also have unavoidable thermodynamic costs [4, 5]. Thus, there is a fundamental tradeoff between the benefits of energy harvesting versus the costs of maintaining and regulating the required nonequilibrium state.

In this Letter, we study this tradeoff in a general setting. We consider a system in a nonequilibrium steady state which produces power by harvesting energy from its environment and storing it in an internal work reservoir. We suppose that the ability to harvest and store energy from the environment depends on the system’s fluctuating state (Fig. 1a). For example, in a bacterium, the ability to transport sugar from the environment depends on the fluctuating availability of certain proteins inside the cell. We also consider an extension of this setup, where energy harvesting may depend on the state of a fluctuating environment (Fig. 1b). This may represent, for example, the fluctuations of the sugar concentration in the bacterium’s environment.

In our setup, the system and environment undergo two types of processes: *baseline* and *control* [6]. Baseline processes are entropy-increasing (diffusion, degradation, etc.) and lead to a loss of information about optimal states for energy harvesting, while control processes carry out self-maintenance and regulation (synthesis, sensing of environmental states, etc.) required for maintaining the nonequilibrium state. We compare the thermodynamic benefit of the control processes, in terms of how much they increase power output, versus their fundamental costs, as determined by the Second Law.

Our main result is an expression of the maximum power output achievable by any set of control processes, given a fixed baseline. This bound is stated in terms of an optimization problem, which also determines the optimal steady-state distribution for maximizing power output. We find closed-form

expressions for optimal power output and optimal steady state in three physically meaningful regimes: the weakly driven linear response regime, the irreversible macroscopic regime, and an intermediate far-from-equilibrium regime. Our analysis illuminates the tradeoffs involved in energy harvesting in various kinds of nonequilibrium systems.

Setup.— We consider a physical system with n states coupled to a heat bath at inverse temperature β , an internal work reservoir, and a nonequilibrium environment which acts as a source of driving (work input). The probability distribution over states is indicated as p , where p_i is the probability of state $i \in \{1, \dots, n\}$. Each state i may be a microstate or a coarse-grained mesostate with an internal entropy s_i [7]. In case of coarse-graining, we assume a separation of scales so that the internal entropy does not depend on the mesostate distribution p (in the absence of coarse-graining, $s_i = 0$). Given distribution p , we write the overall entropy as $\mathcal{S}(p) := S(p) + \langle p, s \rangle$, where $S(p) := \langle p, -\ln p \rangle$ is the Shannon entropy and $\langle p, s \rangle$ is the internal contribution. (Here $\langle x, y \rangle := \sum_i x_i y_i$ and $\ln p$ indicates element-wise logarithm.)

The system undergoes two different types of processes: *baseline* and *control* [6]. The probability distribution evolves according to a master equation $\partial_t p = \bar{R}p$, where $\bar{R} = R + R'$ is the combination of the baseline R and control R' rate matrices. We typically use unadorned symbols for baseline quanti-




Figure 1. (a) We consider a fluctuating system X that can harvest and store energy from its environment (red arrows) in a state-dependent way. We study the tradeoff between the benefits of energy harvesting versus the cost of maintaining the required nonequilibrium state. (b) System X may be coupled to a fluctuating environment Y , such that the rate of energy harvesting also depends on the fluctuating state of the environment.

ties, primes for control quantities, and overline for combined quantities. The overall rate of change of system entropy is $\partial_t \mathcal{S}(p) = \partial_t^R \mathcal{S}(p) + \partial_t^{R'} \mathcal{S}(p)$, where

$$\partial_t^R \mathcal{S}(p) := \langle R p, -\ln p + s \rangle, \quad \partial_t^{R'} \mathcal{S}(p) := \langle R' p, -\ln p + s \rangle$$

are the baseline and control contributions.

We consider the thermodynamics of baseline and control processes via three thermodynamic observables: heat flow to the bath, power output to the internal work reservoir, and power input from the environment. For baseline processes, we write the expected heat flow, power output, and power input under distribution p as $\langle p, q \rangle$, $\langle p, w \rangle$, and $\langle p, d \rangle$, respectively. Here q_i indicates the heat output/sec when the system is in state i , and similarly for w_i (work output/sec) and d_i (work input/sec). These energy exchanges may be due to transitions between states or due to internal fluxes within mesostates. For instance, if the baseline dissipates Q_{ji} joules of heat for each transition $i \rightarrow j$ and \bar{Q}_i joules/sec due to internal fluxes in mesostate i , then $q_i = \bar{Q}_i + \sum_{j(\neq i)} R_{ji} Q_{ji}$. In a similar way, for control processes, we write the expected heat, power output, and power input as $\langle p, q' \rangle$, $\langle p, w' \rangle$, and $\langle p, d' \rangle$.

We assume that the baseline rate matrix R is irreducible, which implies that it has a unique steady state, π . In this case, the combined rate matrix \bar{R} must also be irreducible with some unique steady state $\bar{\pi}$. We assume that the control processes obey the Second Law,

$$\partial_t^{R'} \mathcal{S}(p) + \beta \langle p, q' \rangle \geq 0 \quad \text{for all } p. \quad (1)$$

We also assume that the baseline and control processes obey the First Law (conservation of energy), and in particular that heat and work flows balance in both steady states π and $\bar{\pi}$,

$$\langle \pi, w + q - d \rangle = 0, \quad \langle \bar{\pi}, w + w' + q + q' - d - d' \rangle = 0. \quad (2)$$

Cost and benefit of control.— We now study the thermodynamic benefits and costs of the control processes. We define the *benefit* as the increase of steady-state power output when control is present versus not present,

$$\Delta \dot{W}(\bar{\pi}) := \langle \bar{\pi}, w + w' \rangle - \langle \pi, w \rangle. \quad (3)$$

We define the *cost* of the control processes as the minimal thermodynamic cost needed to maintain the nonequilibrium steady state $\bar{\pi}$, given the fixed baseline. This minimal cost is given by the rate of entropy decrease due to control processes. Combining with (1) and the fact that entropy is constant in steady state, $\partial_t^{R'} \mathcal{S}(\bar{\pi}) + \partial_t^R \mathcal{S}(\bar{\pi}) = 0$, gives

$$\partial_t^R \mathcal{S}(\bar{\pi})/\beta \leq \langle \bar{\pi}, q' \rangle. \quad (4)$$

In words, $\partial_t^R \mathcal{S}(\bar{\pi})$ quantifies the amount of self-maintenance (error correction, repair, etc.) that must be performed by control in order to counteract the dissipative baseline in steady state, and this self-maintenance leads to a minimal thermodynamic cost in terms of heat flow. We note that (4) was

previously shown to be the minimal cost of maintaining a nonequilibrium steady state in Refs. [6, 8].

We emphasize that the particular way to divide the system's dynamics into baseline and control will depend on the scientific question being investigated. In many cases, the baseline will represent unavoidable dissipative processes while the control will represent beneficial power-increasing processes, in which case $\partial_t^R \mathcal{S}(\bar{\pi}) > 0$ and $\Delta \dot{W}(\bar{\pi}) > 0$, so there is a tradeoff between the two. However, the division into baseline and control may both be drawn in other ways, and in principle both $\Delta \dot{W}$ and $\partial_t^R \mathcal{S}$ may be either positive or negative.

Maximum power increase.— We now derive our first main result, a bound on the power increase that reflects the optimal tradeoff between cost and benefit of control. Our derivation makes use of one physically motivated assumption: that the control processes are only coupled to the heat bath and the system's internal work reservoir and do not directly exchange work with the environment, $d' = 0$.

Before proceeding, we first derive the bound

$$\Delta \dot{W}(\bar{\pi}) \leq -\partial_t^R \mathcal{S}(\bar{\pi})/\beta + \langle \bar{\pi} - \pi, d - q \rangle. \quad (5)$$

To do so, we subtract one term from another in (2) and then use $d' = 0$ to give $\Delta \dot{W}(\bar{\pi}) + \langle \bar{\pi}, q' \rangle + \langle \bar{\pi} - \pi, q - d \rangle = 0$. Plugging in (4) and rearranging leads to (5). This inequality becomes tight when the control processes are thermodynamically reversible, so that (4) is tight.

The RHS of (5) depends on the baseline, as well as the steady-state distribution $\bar{\pi}$ of the combined baseline and control rate matrix \bar{R} . We remove the dependence on the control by optimizing over possible steady states,

$$\Delta \dot{W}^* = \max_p -\partial_t^R \mathcal{S}(p)/\beta + \langle p - \pi, d - q \rangle, \quad (6)$$

where the maximization is over all probability distributions. It is clear that $\Delta \dot{W}^* \geq 0$, since the objective vanishes for $p = \pi$.

Eq. (6) is our first main result. It expresses the maximum benefit of control in a way that depends only on properties of the baseline. It may be used to compute the efficiency of any possible control processes as $\Delta \dot{W}(\bar{\pi})/\Delta \dot{W}^* \leq 1$. Maximum efficiency is achieved if the control is thermodynamically reversible and implements the optimal steady-state distribution $\bar{\pi} = p^*$, where p^* achieves the maximum in (6).

In addition, the optimal p^* shows how to construct control processes that achieve maximal efficiency. Assume for simplicity that there is no coarse-graining ($s = 0$). Suppose that the control processes obey local detailed balance (DB), $q'_i = \sum_j R'_{ji} \ln(R'_{ji}/R'_{ij})$, and that they are parameterized in an Arrhenius form, $R'_{ij} = \kappa e^{-\ln p^*_j}$. In the limit of fast control $\kappa \rightarrow \infty$, it may be verified that $\bar{\pi} \rightarrow p^*$, that the bound (4) becomes tight, and that $\Delta \dot{W}(\bar{\pi}) \rightarrow \Delta \dot{W}^*$ [6].

The bound (6) reflects a tradeoff between an information-theoretic cost $-\partial_t^R \mathcal{S}(\bar{\pi})/\beta$ and an energetic benefit $\langle p - \pi, d - q \rangle$. At first glance, this tradeoff resembles existing information/energy tradeoffs in Maxwellian demons, Szilard engines, and other kinds of "information engines" [9–16]. However, it

is important to distinguish our result from information engines. In a typical information engine, there is no external source of driving and information (low entropy) serves as *fuel*, that is a source of free energy that is consumed in the process of work extraction. In an information engine, the “work value of a bit” is microscopic, on the order of $\beta^{-1} \ln 2$ per bit. We instead consider the relationship between energy and information in the presence of external driving, which is a scenario much more similar to the one faced by biological organisms and other autonomous systems [17, 18]. Here information plays a fundamentally different role: it enables the system to tap into external sources of energy that would otherwise be inaccessible, and it generally acts as a *catalyst* rather than a fuel. In our setup, information does not necessarily get consumed during work extraction, and the work value of a bit (the ratio between the two terms in (6)) can be much larger than $\beta^{-1} \ln 2$.

In fact, our result is more closely related to recent research on dissipative heat engines [19–25], as well as power extraction from nonequilibrium heat baths, such as active baths with non-thermal noise [26–28] or with inter-bath correlations [29]. However, instead of considering specific models, we provide a general bound which is applicable to various kinds of nonequilibrium systems in steady state.

Limiting regimes.—The bound (6) involves the optimization of a convex function, which can be solved numerically using standard techniques. However, the solution does not have a closed-form expression in general, even for a simple two-state system. Nonetheless, as we show in our second set of results, closed-form expressions may be identified in three physically meaningful regimes.

In the first regime, termed linear response (LR), the optimal distribution p^* is close to the baseline steady state π . Suppose for the moment that baseline obeys detailed balance, $\pi_i R_{ji} = \pi_j R_{ij}$ for all i, j . Typically, this implies that π is an equilibrium state and that LR is the near-equilibrium regime of weak driving. We can then rewrite (6) as

$$\Delta \dot{W}^* = \max_p \frac{1}{\beta} \langle R(p - \pi), \ln p - \ln \pi \rangle + \langle p - \pi, \psi^{\text{LR}} \rangle \quad (7)$$

where, for convenience, we defined the LR “payoff vector”,

$$\psi^{\text{LR}} := d - q + R^T (\ln \pi - s) / \beta. \quad (8)$$

The first term in (7) vanishes to first order in $p - \pi$, thus ψ^{LR} determines power output near equilibrium. Since $p \approx \pi$ by assumption, we approximate $\ln p_i - \ln \pi_i \approx (p_i - \pi_i) / \pi_i$ in (7) to give

$$\Delta \dot{W}_{\text{LR}}^* = \max_p \frac{1}{\beta} \langle R(p - \pi), D^{-2}(p - \pi) \rangle + \langle p - \pi, \psi^{\text{LR}} \rangle$$

where D is a diagonal matrix with entries $D_{ii} = \sqrt{\pi_i}$. As shown in SM 1.1 [30], this expression is a convex quadratic optimization, which can be solved in closed form. To introduce this solution, let $u_{(a)}$ indicate the a th right eigenvector of R normalized as $\|D^{-1}u_{(a)}\| = 1$, and let λ_α indicate the

corresponding real-valued eigenvalue ($\lambda_1 = 0$). Then,

$$\Delta \dot{W}_{\text{LR}}^* = \beta \sum_{\alpha > 1} \frac{|\Omega_\alpha|^2}{-\lambda_\alpha}, \quad p^* = \pi + \sum_{\alpha > 1} \frac{\beta \Omega_\alpha}{-\lambda_\alpha} u_{(a)}, \quad (9)$$

where $\Omega_\alpha := \frac{1}{2} \langle u_{(a)}, \psi^{\text{LR}} \rangle$ is the loading of the payoff vector on mode a .

Eq. (9) has a simple and physically meaningful interpretation. The optimal distribution p^* shifts the baseline steady state π toward relaxation mode a in proportion to the payoff loading (Ω_α) times the relaxation timescale ($-1/\lambda_\alpha$). Similarly, $\Delta \dot{W}_{\text{LR}}^*$ contains a separate contribution from each relaxation mode, weighted by the payoff magnitude and timescale. In this way, the optimal strategy minimizes the struggle against dissipation while maximizing expected payoff.

Our analysis can be generalized to the case where R does not obey DB. Then (9) still holds, as long as we consider the eigen-decomposition of the rate matrix $A_{ij} = (R_{ij} + R_{ji}\pi_i/\pi_j)/2$ instead of R . The rate matrix A obeys DB while having the same steady state π and the dynamical activity on all edges as R [31]. $A = R$ iff R obeys DB, and it is sometimes called the “additive reversibilization” of R [32].

In the second regime, which we call the macroscopic (M), the objective (6) is dominated by terms linear in p . In this case, we can ignore the nonlinear term in the objective (see SM 1.2 [30]), turning it into a simple linear optimization,

$$\Delta \dot{W}_{\text{M}}^* = \max_p \langle p - \pi, \psi^{\text{M}} \rangle, \quad \psi^{\text{M}} := d - q - R^T s / \beta, \quad (10)$$

where ψ^{M} is the macroscopic regime “payoff vector”. The optimal distribution is a delta function centered at the optimal state $i^* = \arg \max_i \psi_i^{\text{M}}$, giving

$$\Delta \dot{W}_{\text{M}}^* = \psi_{i^*}^{\text{M}} - \langle \pi, \psi^{\text{M}} \rangle, \quad p_i^* = \delta_{i^*i}. \quad (11)$$

The M payoff vector ψ^{M} differs from the LR payoff vector ψ^{LR} by $R^T \ln \pi$, which reflects the fact that the baseline steady state π has no effect on the M regime optimum. In fact, unlike the LR solution, the M solution is well defined even when the baseline steady state π does not have full support. The two payoff vectors become equal when π is uniform.

We finish by considering a third regime intermediate between LR and M, which we call far-from-equilibrium (FE). FE is a perturbation of the macroscopic regime that accounts for the principal contributions of the nonlinear term in (6). We use $p_i^* \approx \delta_{i^*i}$ to approximate $\langle R p, \ln p \rangle \approx \sum_i R_{ii^*} \ln p_i$. Using $\ln p_{i^*} \approx -\sum_{i \neq i^*} p_i$, this can be further approximated as $\langle R p, \ln p \rangle \approx \sum_{i \neq i^*} (R_{ii^*} \ln p_i - R_{i^*i} p_i)$. Plugging these approximations into (6) decouples the values of p_i in the objective function, allowing us to solve the optimization in closed form (see SM 1.3 [30]):

$$\begin{aligned} \Delta \dot{W}_{\text{FE}}^* &= \Delta \dot{W}_{\text{M}}^* + \frac{1}{\beta} \sum_{i \neq i^*} R_{ii^*} (\ln p_i^* - 1) \\ p_{i(\neq i^*)}^* &= \frac{R_{ii^*}}{\beta(\psi_{i^*}^{\text{M}} - \psi_i^{\text{M}}) + R_{i^*i}}, \quad p_{i^*}^* = 1 - \sum_{i: i \neq i^*} p_i^* \end{aligned} \quad (12)$$

The FE solution (12) also has a physically meaningful interpretation. It is a perturbation of the M solution (11), which shifts probability towards states with high transition rates (large R_{ii^*}) and small payoff decreases ($\psi_{i^*}^M - \psi_i^M$). This optimal strategy balances loss of payoff against the cost of pumping probability back against R_{ii^*} . The logic of the tradeoff is similar to the LR solution (9), although it plays out somewhat differently in the near-equilibrium and far-from-equilibrium settings.

In SM 1 [30], we also discuss formal conditions when the LR, M, and FE approximations are valid. In SM 5 [30], we illustrate these conditions using a simple model of a sugar transporter, which is shown to be in the macroscopic regime.

Fluctuating environments.— In our last set of theoretical results, we consider a generalization of our previous approach in which the state space includes an explicitly modeled fluctuating environment. Here each state i is a tuple $i = (x, y) \in X \times Y$, where X is the system of interest and Y is the environment.

In principle, the analysis above extends immediately to this setup. However, the bound in (6) does not impose any constraints on how the control processes can interact with the environment Y , which may be unrealistic in scenarios where control cannot directly change the state of the environment. For example, in an organism, control processes like genetic regulation only affect the intracellular state; their impact on environmental conditions (like nutrient concentrations) is only indirect.

Our analysis can be extended to account for this kind of constraint, which we term “local control”. Let us suppose that the control rate matrix obeys $R'_{(x'y')(xy)} = 0$ whenever $y \neq y'$, meaning that control cannot directly change the state of the environment. In matrix notation, $BR' = 0$ where $B \in \mathbb{Z}^{|Y| \times |X| |Y|}$ has elements $B_{y'(xy)} = 1 - \delta_{y'y}$. Next, recall that the combined steady-state distribution $\bar{\pi}$ obeys $(R + R')\bar{\pi} = 0$. Multiplying both sides by B and simplifying gives $BR\bar{\pi} = 0$, which must be satisfied by any achievable $\bar{\pi}$ under local control. This gives the following bound on work gain under local control,

$$\Delta \dot{W}^* = \max_{p: BR\bar{\pi}=0} -\partial_t^R \mathcal{S}(p) / \beta + \langle p - \pi, d - q \rangle. \quad (13)$$

The value will be smaller than (6) since the optimization has more constraints, and the difference between the two quantifies the cost of local control. Note also that the entropic term in (13) may be decomposed into the entropy change of subsystem X versus the information flow between X and Y [5]. This may be used to study the costs/benefits of self-maintenance of X versus the costs/benefits of information flow between X and Y . We leave this investigation for future work.

In SM 2 [30], we show that (13) can be solved in closed form in the LR regime where $p^* \approx \pi$. The solution has the same form as (9), plus a correction due to the local control assumption. We also derive solutions for the FE and M regimes, under the additional assumption that the transitions of the environment Y do not depend on the state of the system X .




Figure 2. Left: Maximum power gain $\Delta \dot{W}^*$ for the unicyclic network across different values of energy input Θ ($n = 5, \beta = 1$). We show the exact results, along with the LR, FE, and M approximations. Middle: exact and approximate optimal solution p^* , using the FE (top) and LR (bottom) approximations. The optimal state 1 is in the middle of the histogram. Right: Log-log plot of $\Delta \dot{W}^*$ vs. Θ across more scales. LR is accurate at small Θ while the FE is accurate at larger Θ . Inset: $\Delta \dot{W}^*$ as a function of Θ for networks of different sizes, ranging from $n = 5$ to $n = 100$. At larger Θ , the dependence on network size disappears.

Example: unicyclic network.— We illustrate our approach on a minimal model where energy harvesting requires the maintenance of a steady-state current in a unicyclic network. It may be interpreted as an idealized model of a biomolecular cycle (e.g., an enzymatic cycle [33–35], transporter [36, 37], or oscillator [38, 39]) in which energy is acquired at a particular step of the cycle.

We consider a unicyclic system X with n states (without coarse-graining) where the baseline dynamics involve diffusion across a 1-dimensional ring (see Fig. 1):

$$R_{x'x} = \delta_{x'+1,x} + \delta_{x+1,x'} - 2\delta_{x,x'}, \quad (14)$$

where $n+1 \equiv n$. Each time the system completes a particular transition $1 \rightarrow 2$, Θ units of energy are transferred from the environment to the internal work reservoir (and vice versa for the reverse transition $2 \rightarrow 1$). Energy is harvested without dissipation ($q = 0$). There is no current under the baseline dynamics in steady state, so baseline power output vanishes.

We study the increase of power output $\Delta \dot{W}^*$ (6) in the LR, FE, and M regimes. We summarize our main results, leaving details for SM 3.2 [30]. In the near-equilibrium LR regime (9), which holds when Θ is small,

$$\Delta \dot{W}_{LR}^* = \beta \Theta^2 \frac{(n-1)}{4n^2}. \quad (15)$$

Thus, power output scales quadratically with Θ , and inversely with the number of states in the cycle. For very large Θ , the optimal distribution concentrates on the optimal state $i^* = 1$, so the M regime becomes relevant where $\Delta \dot{W}_M^* = \Theta$. At intermediate Θ , the FE regime (12) is relevant,

$$\Delta \dot{W}_{FE}^* = \Theta - \frac{1}{\beta} [2 + \ln 2 + \ln(\beta\Theta - 1)(\beta\Theta - 2)]. \quad (16)$$

Unlike the LR regime, the M and FE solutions do not depend on system size n . In the SM, we also find the optimal distributions

p^* for the three regimes. The LR optimal distribution p^* has a remarkably simple form, building up in equally-sized steps toward a “cliff” located at optimal state $i^* = 1$.

We illustrate these results in Fig. 2. We show the exact value of $\Delta\dot{W}^*$, found by numerical optimization of Eq. (6), as well as the LR, FE, and M approximations. We also show the exact optimal distribution p^* and its LR and FE approximations. On the inset on the right, we show that $\Delta\dot{W}^*$ does not depend on n for large Θ .

In SM 3.3, we also consider an extension of the unicyclic network model, where the system X is now coupled to a fluctuating environment Y . We assume that X and Y both have n states, and that their baseline dynamics are bipartite and independent. X has the same dynamics (14) as before, while Y has nonequilibrium unicyclic dynamics $R_{y'y}^Y = q((1 + \epsilon)\delta_{y'+1,y} + \delta_{y+1,y'} - (2 + \epsilon)\delta_{y,y'})$ ($n + 1 \equiv n$), where q is the transition rate and ϵ determines the baseline cyclic current in the environment. We assume that the environment controls energy harvesting in the following way: Θ units of energy are transferred to the internal work reservoir only when subsystem X jumps from state $x = y \rightarrow x' = y + 1$, where y is the state of the environment Y (see Fig. 1b).

We study the maximum power output under the “local control” assumption that the control cannot change the state of Y . In this case, $\Delta\dot{W}^*$ is given by the constrained optimization (13). The optimal strategy requires X to continuously measure the fluctuating state of the environment Y . In the SM, we analyze power output in LR, M, and FE regimes. The most illuminating result is the LR regime,

$$\Delta\dot{W}_{\text{LR}}^* = \beta\Theta^2 \frac{(n-1)}{4n^2} \frac{1}{1+q(1+\epsilon/2)}. \quad (17)$$

Thus, optimal power output is equal to the power output without a fluctuating environment (15), divided $1 + q(1 + \epsilon/2)$. As expected, power output decreases with q , the rate of environmental transitions, and (15) is recovered as $q \rightarrow 0$. Power output also decreases as the cyclic current of the environment increases, as determined by the parameter ϵ , since this makes the environment more difficult to track.

Acknowledgments.— We thank members of the Complex Systems Lab as well as Bernat Corominas-Murtra, Luis Seoane, and David Wolpert for useful discussions. AK thanks Sosuke Ito for support and encouragement. JP is supported by “María de Maeztú” fellowship MDM-2014-0370-17-2. This project was partly made possible through the support of Grant 62417 from the John Templeton Foundation. The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of the John Templeton Foundation.

* artemyk@gmail.com

[1] T. L. Hill, *Free energy transduction in biology: the steady-state*

kinetic and thermodynamic formalism. New York: Academic Press, 1977.

- [2] F. Schweitzer, W. Ebeling, and B. Tilch, “Complex motion of Brownian particles with energy depots,” *Physical Review Letters*, vol. 80, no. 23, p. 5044, 1998.
- [3] S. De and R. Klajn, “Dissipative self-assembly driven by the consumption of chemical fuels,” *Advanced Materials*, vol. 30, no. 41, p. 1706750, 2018.
- [4] P. Mehta, A. H. Lang, and D. J. Schwab, “Landauer in the age of synthetic biology: energy consumption and information processing in biochemical networks,” *Journal of Statistical Physics*, vol. 162, no. 5, pp. 1153–1166, 2016.
- [5] J. M. Horowitz and M. Esposito, “Thermodynamics with continuous information flow,” *Physical Review X*, vol. 4, no. 3, p. 031015, 2014.
- [6] J. M. Horowitz, K. Zhou, and J. L. England, “Minimum energetic cost to maintain a target nonequilibrium state,” *Physical Review E*, Mar. 2017. arXiv: 1703.04722.
- [7] M. Esposito, “Stochastic thermodynamics under coarse graining,” *Physical Review E*, vol. 85, no. 4, p. 041125, 2012.
- [8] J. M. Horowitz and J. L. England, “Information-theoretic bound on the entropy production to maintain a classical nonequilibrium distribution using ancillary control,” *Entropy*, vol. 19, no. 7, p. 333, 2017.
- [9] J. M. Parrondo, J. M. Horowitz, and T. Sagawa, “Thermodynamics of information,” *Nature Physics*, vol. 11, no. 2, pp. 131–139, 2015.
- [10] T. Sagawa and M. Ueda, “Minimal energy cost for thermodynamic information processing: measurement and information erasure,” *Physical review letters*, vol. 102, no. 25, p. 250602, 2009.
- [11] D. Abreu and U. Seifert, “Thermodynamics of genuine nonequilibrium states under feedback control,” *Physical review letters*, vol. 108, no. 3, p. 030601, 2012.
- [12] F. J. Cao and M. Feito, “Thermodynamics of feedback controlled systems,” *Physical Review E*, vol. 79, no. 4, p. 041118, 2009.
- [13] S. Ito and T. Sagawa, “Information thermodynamics on causal networks,” *Physical review letters*, vol. 111, no. 18, p. 180603, 2013.
- [14] D. Hartich, A. C. Barato, and U. Seifert, “Stochastic thermodynamics of bipartite systems: transfer entropy inequalities and a Maxwell’s demon interpretation,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2014, no. 2, p. P02016, 2014.
- [15] A. C. Barato, D. Hartich, and U. Seifert, “Efficiency of cellular information processing,” *New Journal of Physics*, vol. 16, no. 10, p. 103024, 2014.
- [16] P. Sartori, L. Granger, C. F. Lee, and J. M. Horowitz, “Thermodynamic costs of information processing in sensory adaptation,” *PLoS Comput Biol*, vol. 10, no. 12, p. e1003974, 2014.
- [17] J. Barham, “A dynamical model of the meaning of information,” *Biosystems*, vol. 38, no. 2-3, pp. 235–241, 1996.
- [18] A. Kolchinsky and D. H. Wolpert, “Semantic information, autonomous agency and non-equilibrium statistical physics,” *Interface Focus*, vol. 8, p. 20180041, Dec. 2018.
- [19] V. Blickle and C. Bechinger, “Realization of a micrometre-sized stochastic heat engine,” *Nature Physics*, vol. 8, no. 2, pp. 143–146, 2012.
- [20] I. A. Martínez, É. Roldán, L. Dinis, D. Petrov, J. M. Parrondo, and R. A. Rica, “Brownian Carnot engine,” *Nature Physics*, vol. 12, no. 1, pp. 67–70, 2016.
- [21] V. Holubec, S. Steffenoni, G. Falasco, and K. Kroy, “Active Brownian heat engines,” *Physical Review Research*, vol. 2, no. 4, p. 043262, 2020.

- [22] P. Pietzonka and U. Seifert, "Universal trade-off between power, efficiency, and constancy in steady-state heat engines," *Physical review letters*, vol. 120, no. 19, p. 190602, 2018.
- [23] E. Penocchio, R. Rao, and M. Esposito, "Thermodynamic efficiency in dissipative chemistry," *Nature communications*, vol. 10, no. 1, pp. 1–5, 2019.
- [24] A. E. Allahverdyan and Q. Wang, "Adaptive machine and its thermodynamic costs," *Physical Review E*, vol. 87, no. 3, p. 032139, 2013.
- [25] A. E. Allahverdyan, S. G. Babajanyan, N. Martirosyan, and A. Melkikh, "Adaptive heat engine," *Physical Review Letters*, vol. 117, no. 3, p. 030601, 2016.
- [26] P. Margaretti and H. Stark, "Szilard engines and information-based work extraction for active systems," Mar. 2022.
- [27] G. Paneru, S. Dutta, and H. K. Pak, "Colossal power extraction from active cyclic Brownian information engines," *arXiv preprint arXiv:2203.13538*, 2022.
- [28] T. K. Saha, J. Ehrich, M. Gavrillov, S. Still, D. A. Sivak, and J. Bechhoefer, "Information engine in a nonequilibrium bath," July 2022.
- [29] G. L. Zanin, M. Antesberger, M. J. Jacquet, P. H. S. Ribeiro, L. A. Rozema, and P. Walther, "Enhanced photonic maxwell's demon with correlated baths," *Quantum*, vol. 6, p. 810, 2022.
- [30] See Supplemental Material at [URL will be inserted by publisher], which includes references [5, 7, 40–48].
- [31] The dynamical activity refers to the back-and-forth fluxes $R_{ij}\pi_j + R_{ji}\pi_i$.
- [32] J. A. Fill, "Eigenvalue Bounds on Convergence to Stationarity for Nonreversible Markov Chains, with an Application to the Exclusion Process," *The Annals of Applied Probability*, vol. 1, no. 1, pp. 62–87, 1991.
- [33] A. C. Barato and U. Seifert, "Thermodynamic uncertainty relation for biomolecular processes," *Physical review letters*, vol. 114, no. 15, p. 158101, 2015.
- [34] J. R. Moffitt and C. Bustamante, "Extracting signal from noise: kinetic mechanisms from a michaelis-menten-like expression for enzymatic fluctuations," *The FEBS journal*, vol. 281, no. 2, pp. 498–517, 2014.
- [35] H. Wierenga, P. R. Ten Wolde, and N. B. Becker, "Quantifying fluctuations in reversible enzymatic cycles and clocks," *Physical Review E*, vol. 97, no. 4, p. 042404, 2018.
- [36] J. Schnakenberg, "Network theory of microscopic and macroscopic behavior of master equation systems," *Reviews of Modern physics*, vol. 48, no. 4, p. 571, 1976.
- [37] S. Yoshida, Y. Okada, E. Muneyuki, and S. Ito, "Thermodynamic role of main reaction pathway and multi-body information flow in membrane transport," *Physical Review Research*, vol. 4, no. 2, p. 023229, 2022.
- [38] A. C. Barato and U. Seifert, "Coherence of biochemical oscillations is bounded by driving force and network topology," *Physical Review E*, vol. 95, no. 6, p. 062409, 2017.
- [39] J. Pajmians, D. K. Lubensky, and P. R. Ten Wolde, "A thermodynamically consistent model of the post-translational kai circadian clock," *PLOS Computational Biology*, vol. 13, no. 3, p. e1005415, 2017.
- [40] S. Puntanen, G. P. Styan, and J. Isotalo, *Matrix tricks for linear statistical models: our personal top twenty*. Springer, 2011.
- [41] A. A. Maciejewski and C. A. Klein, "Obstacle avoidance for kinematically redundant manipulators in dynamically varying environments," *The international journal of robotics research*, vol. 4, no. 3, pp. 109–117, 1985.
- [42] R. M. Gray *et al.*, "Toeplitz and circulant matrices: A review," *Foundations and Trends® in Communications and Information Theory*, vol. 2, no. 3, pp. 155–239, 2006.
- [43] C. Naslund, "Mathematics stackexchange. question 104967," 2012. <https://math.stackexchange.com/questions/104967/>.
- [44] Svyatoslav and C. Leibovici, "Mathematics stackexchange. question 4567421," 2022. <https://math.stackexchange.com/questions/4567421/>.
- [45] Y. J. Chung, C. Krueger, D. Metzgar, and M. H. Saier Jr, "Size comparisons among integral membrane transport protein homologues in bacteria, archaea, and eucarya," *Journal of bacteriology*, vol. 183, no. 3, pp. 1012–1021, 2001.
- [46] M. Razavi, S. M. Saberi Fathi, and J. A. Tuszynski, "The effect of the protein synthesis entropy reduction on the cell size regulation and division size of unicellular organisms," *Entropy*, vol. 24, no. 1, p. 94, 2022.
- [47] J. L. England, "Statistical physics of self-replication," *The Journal of chemical physics*, vol. 139, no. 12, p. 121923, 2013.
- [48] R. Milo and R. Phillips, *Cell biology by the numbers*. Garland Science, 2015.

Supplemental Material: Universal bounds and thermodynamic tradeoffs in nonequilibrium energy harvesting

Jordi Piñero, Ricard Solé, and Artemy Kolchinsky

Contents

1. Maximum power increase, Eq. (6) (static environments)	8
1. Linear Response (LR) regime, Eq. (9)	8
2. Macroscopic (M) regime, Eq. (11)	10
3. Far-from-Equilibrium (FE) regime, Eq. (12)	12
2. Maximum power increase, Eq. (13) (fluctuating environments)	14
1. Additional assumptions	14
a. Local control	14
b. Independent environment	15
c. Independent agent-and-environment	16
2. Linear Response (LR) regime	16
a. General case	16
b. Simplifying assumption: Independent environment	17
c. Simplifying assumption: Independent-agent-and-environment	18
3. Macroscopic (M) regime	20
4. Far-from-Equilibrium (FE) regime	20
3. Unicyclic model	23
1. An Algebraic Aperitif: eigendecomposition of unicyclic rate matrix	23
2. Unicyclic model: Static environment	23
a. Static environment: LR regime	24
b. Static environment: FE Regime	26
3. Unicyclic model: Fluctuating environment	27
4. LR regime	27
5. Fluctuating environment: FE regime	30
4. Alternative unicyclic model	32
1. Static environment: LR Regime	33
a. FE Regime	34
5. Glucose Transporter	35
6. Quadratic optimization theorem	38

1. Maximum power increase, Eq. (6) (static environments)

We consider a system of n states indexed by $i \in \{1, \dots, n\}$ that evolves according to a Markovian process with rate matrix $\bar{R} = R + R'$, where R and R' correspond to baseline and control processes respectively. Here we study the optimization (6) in three limiting regimes.

1.1 Linear Response (LR) regime, Eq. (9)

We first consider (6) in the linear response regime, that is under the assumption that the optimal distribution is close to the baseline steady state $p^* \approx \pi$, thereby arriving at (9). We assume that R is irreducible, so that π has full support.

First, we rewrite the objective as

$$\Delta \dot{W}^* = \max_p \frac{1}{\beta} \langle Rp, \ln p - \ln \pi \rangle + \langle p - \pi, \psi^{\text{LR}} \rangle \quad (\text{S1})$$

$$= \max_p \frac{1}{\beta} \langle R(p - \pi), \ln p - \ln \pi \rangle + \langle p - \pi, \psi^{\text{LR}} \rangle, \quad (\text{S2})$$

where we defined the LR payoff vector as in the main text,

$$\psi^{\text{LR}} = d - q + R^T(\ln \pi - s)/\beta. \quad (\text{S3})$$

Eq. (S1) appears as (7) in the main text, while the second line follows from $R\pi = 0$.

Next, for convenience define the matrices $A, D, M \in \mathbb{R}^{n \times n}$,

$$A_{ij} := \left(R_{ij} + R_{ji} \frac{\pi_i}{\pi_j} \right) / 2 \quad D_{ij} := \delta_{ij} \sqrt{\pi_i} \quad M := D^{-1}AD \quad (\text{S4})$$

Recall that A is an irreducible rate matrix with stationary distribution π , which obeys detailed balance if and only if $A = R$. In addition, $M = D^{-1}AD$ implies that M and A are related via a similarity transformation, which means that for any right eigenpair (λ, u) of A , $(\lambda, D^{-1}u)$ is an eigenpair of the symmetric matrix M . In particular, since A has a unique right eigenvector π with eigenvalue 0, then M also has a unique eigenvector $D^{-1}\pi$ with eigenvalue 0. Moreover, since A is a rate matrix, it is negative semidefinite ($\lambda_a \leq 0$), and so M is also negative semidefinite. We indicate the right eigenvectors of A as $u_{(a)}$ and the eigenvectors of the symmetric matrix M as $m_{(a)} = D^{-1}u_{(a)}$. We assume that $m_{(a)}$ are normalized as $\|m_{(a)}\| = 1$, therefore $\|D^{-1}u_{(a)}\| = 1$.

Next, define the vectors $z, v \in \mathbb{R}^n$,

$$z := D^{-1}(p - \pi) \quad v := \frac{\beta}{2} D \psi^{\text{LR}} \quad (\text{S5})$$

Using $\ln x \simeq x - 1$ for $x \approx 1$, we approximate the first term in (S2) as

$$\sum_{i,j} R_{ij} (p_j - \pi_j) \ln \left(\frac{p_i}{\pi_i} \right) \simeq \sum_{i,j} R_{ij} (p_j - \pi_j) \left(\frac{p_i - \pi_i}{\pi_i} \right) \quad (\text{S6})$$

$$= \sum_{i,j} \frac{\sqrt{\pi_j}}{\sqrt{\pi_i}} R_{ij} z_j z_i = \sum_{i,j} \frac{\sqrt{\pi_j}}{\sqrt{\pi_i}} A_{ij} z_j z_i = \sum_{i,j} (D^{-1} A D)_{ij} z_j z_i. \quad (\text{S7})$$

Combining and using the definitions above, we can rewrite (S2) in matrix notation as

$$\Delta \dot{W}(p) \simeq \frac{1}{\beta} (z^T M z + 2z^T v). \quad (\text{S8})$$

Let us consider the maximization of (S8) with respect to z , under the constraints $\sum_i p_i = 1$ and $p_i \geq 0, \forall i \in \{1, \dots, n\}$. Since by assumption π is strictly positive and $p \approx \pi$, the second condition follows automatically. The first constraint can be expressed as $z^T D^{-1} \pi = 0$, which gives $\sum_i p_i = \sum_i \pi_i = 1$. Thus, our optimization problem is now given by

$$\Delta \dot{W}_{\text{LR}}^* = \frac{1}{\beta} \max_{z \in \mathbb{R}^n: (z, D^{-1} \pi) = 0} z^T M z + 2z^T v. \quad (\text{S9})$$

Eq. (S9) is a quadratic optimization problem, which can be solved using standard techniques from linear algebra.

For convenience we summarize these techniques in Theorem 1 in Sec. 6 of this SM. That theorem implies that

$$\Delta \dot{W}_{\text{LR}}^* = -\frac{1}{\beta} v^T M^+ v \quad p^* = \pi - D M^+ v \quad (\text{S10})$$

where $M^+ = \sum_{a>1} \lambda_a^{-1} m_{(a)} m_{(a)}^T$ is the Penrose-Moore pseudo-inverse of M . In applying Theorem 1, we used result (S118) and the relation $p^* = \pi + D z^*$.

We can rewrite (S10) using the eigensystem of M ,

$$\Delta \dot{W}_{\text{LR}}^* = \frac{1}{\beta} \sum_{a>1} \frac{|\langle v, m_{(a)} \rangle|^2}{-\lambda_a} \quad p^* = \pi + \sum_{a>1} \frac{\langle v, m_{(a)} \rangle}{-\lambda_a} D m_{(a)} \quad (\text{S11})$$

Finally, as in the main text we define

$$\Omega_a = \frac{1}{\beta} \langle v, m_{(a)} \rangle = \frac{1}{2} \langle D \psi^{\text{LR}}, D^{-1} u_{(a)} \rangle = \frac{1}{2} \langle \psi^{\text{LR}}, u_{(a)} \rangle. \quad (\text{S12})$$

Plugging into (S10) gives Eq. (9) in the main text,

$$\Delta \dot{W}_{\text{LR}}^* = \beta \sum_{a>1} \frac{|\Omega_a|^2}{-\lambda_a} \quad p^* = \pi + \sum_{a>1} \frac{\beta \Omega_a}{-\lambda_a} u_{(a)} \quad (\text{S13})$$

Note that the assumption $p^* \approx \pi$ holds when $\|p^* - \pi\| \ll 1$. We can write

$$\begin{aligned} \|p^* - \pi\| &= \beta \left\| \sum_{a>1} \frac{\Omega_a}{-\lambda_a} u_{(a)} \right\| \leq \sum_{a>1} \frac{\beta |\Omega_a|}{-\lambda_a} \|u_{(a)}\| \\ &\leq \sum_{a>1} \frac{\beta |\Omega_a|}{-\lambda_a} \|D\| \|m_{(a)}\| \leq \beta (n-1) \max_{a>1} |\Omega_a / \lambda_a| \max_i \sqrt{\pi_i}, \end{aligned} \quad (\text{S14})$$

where we used the triangle inequality and properties of the matrix norm. Therefore, the LR approximation is guaranteed to be valid when

$$\max_{a>1} |\Omega_a/\lambda_a| \ll \frac{1}{\beta(n-1) \max_i \sqrt{\pi_i}}. \quad (\text{S15})$$

In other words, the payoff loading on each eigenmode must be slower than relaxation modes, up to a factor that depends only on β , n , and the steady state π .

1.2 Macroscopic (M) regime, Eq. (11)

In the macroscopic (M) regime, we assume that the nonlinear terms in (6) are small compared to the linear terms. This allows us to approximate the optimal solution as the linear optimization

$$\Delta\dot{W}^* = \max_p \frac{1}{\beta} \langle Rp, \ln p \rangle + \langle p - \pi, \psi^M \rangle \quad (\text{S16})$$

$$\approx \max_p \langle p - \pi, \psi^M \rangle \quad (\text{S17})$$

$$= \psi_{i^*}^M - \langle \pi, \psi^M \rangle =: \Delta\dot{W}_M^*. \quad (\text{S18})$$

Here $\psi^M := d - q - R^T s / \beta$ is the macroscopic payoff vector. The maximum in the last line is achieved by $p_i^* = \delta_{ii^*}$ with $i^* = \arg \max_i \psi_i^M$, which appears as (11) in the main text.

We now discuss when this approximation is valid. First, define the parameter α ,

$$\alpha := \frac{K}{\beta \Delta\dot{W}_M^*}, \quad (\text{S19})$$

where $K := \max_i | -R_{ii} |$ is the largest escape rate. This parameter reflects the balance between diffusion out of the optimal state versus macroscopic power gain. Assuming $\alpha < 1$, we show that

$$\left| \frac{\Delta\dot{W}^*}{\Delta\dot{W}_M^*} - 1 \right| \leq \alpha (-\ln \alpha - \ln \gamma + n) \quad (\text{S20})$$

where we define the following constant which depends only on R :

$$\gamma := \min_{i: \sum_j R_{ij} > 0} \pi_i \quad (\text{S21})$$

In fact, γ is the minimal steady-state probability of any state that has incoming transitions.

The RHS of (S20) vanishes for $\alpha \rightarrow 0$, so the LHS tightens as $\Delta\dot{W}^*/\Delta\dot{W}_M^* \rightarrow 1$. Thus, the parameters α and γ determine the regime of validity of the M approximation. At the same time, we emphasize that the M approximation only becomes accurate in relative terms, not in additive terms (that is, it is not necessarily true that $\Delta\dot{W}^* - \Delta\dot{W}_M^* \rightarrow 0$).

To derive (S20), we first consider an upper bound on $\Delta\dot{W}^*$. Using $Rp = \sum_j R_{ij}p_j \geq R_{ii}p_i$ and $S(p) = -\sum_i p_i \ln p_i \leq \ln n$, we have

$$\langle Rp, \ln p \rangle = \sum_{i,j} R_{ij}p_j \ln p_i \leq \sum_i |R_{ii}|(-p_i \ln p_i) \leq KS(p) \leq K \ln n \quad (\text{S22})$$

At the same time, since $\Delta\dot{W}_M^*$ maximizes $\langle p - \pi, \psi^M \rangle$, we can plug into (S16) to write

$$\Delta\dot{W}^* \leq K(\ln n)/\beta + \Delta\dot{W}_M^* = \Delta\dot{W}_M^*(\alpha \ln n + 1). \quad (\text{S23})$$

To derive a lower bound on $\Delta\dot{W}^*$, define the distribution $p_i := (1 - \alpha)\delta_{ii^*} + \alpha\pi_i$, with α defined in (S19).

Plugging this distribution into the objective in (S16) yields

$$\begin{aligned} \Delta\dot{W}^* &\geq \frac{1}{\beta} \langle Rp, \ln p \rangle + \langle p - \pi, \psi^M \rangle \\ &= \frac{1}{\beta} (1 - \alpha) \sum_i (R_{ii^*} \ln p_i) + \langle p - \pi, \psi^M \rangle, \end{aligned} \quad (\text{S24})$$

where we used that $\sum_j R_{ij}((1 - \alpha)\delta_{ji^*} + \alpha\pi_j) = R_{ii^*}(1 - \alpha)$ since $R\pi = 0$. Using $R_{i^*i^*} \ln p_{i^*} \geq 0$, we bound the first term in the RHS of (S24) as

$$\begin{aligned} \sum_i R_{ii^*} \ln p_i &\geq \sum_{i:i^* \neq i} R_{ii^*} \ln p_i = \sum_{i:i^* \neq i} R_{ii^*} \ln(\alpha\pi_i) \\ &\geq \sum_{i \neq i^*} R_{ii^*} (\ln \alpha + \ln \gamma) \\ &= -R_{i^*i^*} (\ln \alpha + \ln \gamma) \geq K (\ln \alpha + \ln \gamma), \end{aligned} \quad (\text{S25})$$

where we used that $\ln \alpha + \ln \gamma < 0$ in the last inequality. For the second term in (S24),

$$\sum_i (p_i - \pi_i) \psi_i^M = (1 - \alpha) \sum_i (\delta_{ii^*} - \pi_i) \psi_i^M = (1 - \alpha) \Delta\dot{W}_M^*. \quad (\text{S26})$$

We can plug (S25) and (S26) into (S24) to give a lower bound on $\Delta\dot{W}^*$,

$$\begin{aligned} \Delta\dot{W}^* &\geq (1 - \alpha)(-R_{i^*i^*})(\ln \alpha + \ln \gamma)/\beta + (1 - \alpha)\Delta\dot{W}_M^* \\ &= (1 - \alpha)\Delta\dot{W}_M^*\alpha(\ln \alpha + \ln \gamma) + (1 - \alpha)\Delta\dot{W}_M^* \\ &\geq \Delta\dot{W}_M^*\alpha(\ln \alpha + \ln \gamma) + (1 - \alpha)\Delta\dot{W}_M^* \\ &= \Delta\dot{W}_M^*\alpha(\ln \alpha + \ln \gamma - 1) + \Delta\dot{W}_M^* \\ &\geq \Delta\dot{W}_M^*\alpha(\ln \alpha + \ln \gamma - n) + \Delta\dot{W}_M^*, \end{aligned} \quad (\text{S27})$$

where in the second line we used the definition (S19), and in the third line we used $\ln \alpha + \ln \gamma < 0$. Combining yields

$$\Delta\dot{W}_M^*\alpha(\ln \alpha + \ln \gamma - n) \leq \Delta\dot{W}^* - \Delta\dot{W}_M^* \leq \alpha \ln n, \quad (\text{S28})$$

which can be rearranged to give (S20).

1.3 Far-from-Equilibrium (FE) regime, Eq. (12)

In this regime, we consider a perturbation of the macroscopic regime: we assume that the optimal p^* is close to delta function distribution centered at $i^* = \operatorname{argmax}_i \psi_i^M$. If $p_i \approx \delta_{ii^*}$,

$$\dot{p}_i = \sum_j R_{ij} p_j \approx R_{ii^*}. \quad (\text{S29})$$

This allows us to approximate the entropic term in (S16) as

$$\begin{aligned} \langle R p, \ln p \rangle &= \sum_{i,j} R_{ij} p_j \ln p_i \approx \sum_i R_{ii^*} \ln p_i \\ &= R_{i^*i^*} \ln \left(1 - \sum_{i:i \neq i^*} p_i \right) + \sum_{i:i \neq i^*} R_{ii^*} \ln p_i \approx \sum_{i:i \neq i^*} (-R_{i^*i^*} p_i + R_{ii^*} \ln p_i) \end{aligned} \quad (\text{S30})$$

where we used $\ln(1-x) \approx -x$ when $x \approx 0$. Plugging into (S16) gives

$$\begin{aligned} \Delta \dot{W}(p) &\simeq \frac{1}{\beta} \sum_{i:i \neq i^*} (-R_{i^*i^*} p_i + R_{ii^*} \ln p_i) + \left(1 - \sum_{i:i \neq i^*} p_i \right) \psi_{i^*}^M + \sum_{i:i \neq i^*} p_i \psi_i^M - \langle \pi, \psi^M \rangle \\ &= \Delta \dot{W}_M^* + \sum_{i:i \neq i^*} [-p_i (\psi_{i^*}^M - \psi_i^M + \beta^{-1} R_{i^*i^*}) + \beta^{-1} R_{ii^*} \ln p_i], \end{aligned} \quad (\text{S31})$$

where $\Delta \dot{W}_M^* := \psi_{i^*}^M - \langle \pi, \psi^M \rangle$, which does not depend on p . After maximizing with respect to $\{p_i\}_{i:i \neq i^*}$ by taking derivatives and setting them to zero, one obtains:

$$p_i^* = \frac{R_{ii^*}}{\beta(\psi_{i^*}^M - \psi_i^M) + R_{i^*i^*}} \quad (i \neq i^*), \quad \text{and} \quad p_{i^*}^* = 1 - \sum_{i:i \neq i^*} p_i^*. \quad (\text{S32})$$

Plugging (S32) into (S31) gives

$$\Delta \dot{W}^* \simeq \Delta \dot{W}_{\text{FE}}^* = \Delta \dot{W}_M^* + \frac{1}{\beta} \sum_{i:i \neq i^*} R_{ii^*} (\ln p_i^* - 1). \quad (\text{S33})$$

Let us briefly discuss the validity of the expressions above. As stated above, (S32) is valid when $p_i \approx \delta_{ii^*}$. For convenience, let $i^- = \operatorname{argmax}_{i:i \neq i^*} \psi_i^M$ indicate the second most valuable state of the system after i^* . We define the *payoff gap* as

$$\Delta \psi_{\min}^M := \psi_{i^*}^M - \psi_{i^-}^M = \min_{i:i \neq i^*} \psi_{i^*}^M - \psi_i^M.$$

Now, considering the optimal probability for $i \neq i^*$ found in (S32), the FE approximation will be valid as long as

$$p_{i \neq i^*} \ll 1 \quad \Leftrightarrow \quad \psi_{i^*}^M - \psi_i^M \gg \frac{1}{\beta} (R_{ii^*} - R_{i^*i^*}),$$

which holds as long as

$$\Delta \psi_{\min}^M \gg \frac{1}{\beta} (R_{i^-i^*} - R_{i^*i^*}). \quad (\text{S34})$$

Expression (S34) implies that FE approximation is valid when the payoff gap is much larger than the rates at which either i^* flows into i^- or escapes.

2. Maximum power increase, Eq. (13) (fluctuating environments)

We now extend our setup to consider a fluctuating environment. Now we consider a combined system that contains two subsystems X and Y , representing the “agent” and “environment” respectively, taking values in $\{1, \dots, |X|\}$ and $\{1, \dots, |Y|\}$. The baseline and control dynamics of the joint system are specified by the rate matrix $R, R' \in \mathbb{R}^{|X||Y| \times |X||Y|}$.

Maximum power output in this setup is given by the constrained optimization problem (13). Here we consider this optimization problem in three limiting regimes.

2.1 Additional assumptions

Before proceeding, we introduce three additional assumptions that are employed in some of our results below.

2.1 Local control

As described in the main text, we use the assumption of *Local control* in all our results on fluctuating environments. In fact, it is due to this additional assumption that the fluctuating environment setup is not simply a special case of the basic setup presented above.

The assumption of local control states that under the control rate matrix, the environment Y cannot change state:

$$\mathbf{Local\ control:} \quad R'_{(x'y')(xy)} = 0 \text{ if } y' \neq y. \quad (\text{S35})$$

This constraint can be formalized using the matrix

$$B \in \mathbb{R}^{|Y| \times |X||Y|} \text{ with } B_{y'(xy)} = 1 - \delta_{y'y}. \quad (\text{S36})$$

It can be verified that $BR' = 0$ is equivalent to (S35). Recall that $\bar{\pi}$ was defined as the steady-state solution of the combination of baseline and control dynamics, $(R + R')\bar{\pi} = 0$. Multiplying both sides of this equality by B and then using $BR' = 0$ implies $BR\bar{\pi} = 0$. This equality must hold for any achievable steady state in (5). This gives the following expression for maximum power gain in the presence of local control:

$$\Delta \dot{W}^* = \max_{p: BRp=0} -\partial_t^R \mathcal{S}(p) / \beta + \langle p - \pi, d - q \rangle, \quad (\text{S37})$$

where the optimization is over all probability distributions over $|X||Y|$ outcomes. This expression appears as (13) in the main text.

Below, it will be convenient to rewrite (S37) in the form of (S1),

$$\Delta \dot{W}^* = \max_{p: BRp=0} \frac{1}{\beta} \langle Rp, \ln p - \ln \pi \rangle + \langle p - \pi, \psi^{\text{LR}} \rangle, \quad (\text{S38})$$

where $\psi^{\text{LR}} := d - q + R^T(\ln \pi - s)/\beta$ is the LR payoff vector.

2.1 Independent environment

Where explicitly noted, we sometimes introduce an additional *Independent environment* assumption. This assumption states that under the baseline rate matrix R , the dynamics of the agent and environment are bipartite [5], meaning either the agent or environment can change state in any one jump, and moreover that the jumps of the environment do not depend on the state of the agent. Formally, this means that the combined rate matrix has the general form

$$\mathbf{Independent\ environment:} \quad R_{(x'y')(xy)} = \delta_{yy'} R_{(x'y)(xy)}^{X|Y} + \delta_{xx'} R_{y'y}^Y. \quad (\text{S39})$$

R^Y is a rate matrix over environment Y which satisfies $\sum_{y'} R_{y'y}^Y = 0$ for all y , which we assume has a unique steady state π^Y . $R^{X|Y}$ is a kind of conditional rate matrix over X , where transitions of X can in principle depend on the state of Y . Using $\sum_{x',y'} R_{(x'y')(xy)} = \sum_{y'} R_{y'y}^Y = 0$, we have

$$\begin{aligned} \sum_{x'} R_{(x'y)(xy)}^{X|Y} &= \sum_{x'} R_{(x'y)(xy)}^{X|Y} + \sum_{y'} R_{y'y}^Y \\ &= \sum_{x',y'} \delta_{yy'} R_{(x'y)(xy)}^{X|Y} + \sum_{x',y'} \delta_{xx'} R_{y'y}^Y = \sum_{x',y'} R_{(x'y')(xy)} = 0. \end{aligned} \quad (\text{S40})$$

Recall that local control (S37) implies that $BRp = 0$. Combining with (S39) allows us to write the product BR as

$$\begin{aligned} [BR]_{y'',xy} &= \sum_{x',y'} (1 - \delta_{y'y''}) R_{(x'y')(xy)} = - \sum_{x',y'} \delta_{y'y''} R_{(x'y')(xy)} \\ &= - \sum_{x',y'} \delta_{y'y''} (\delta_{y,y'} R_{(x'y)(xy)}^{X|Y} + \delta_{x,x'} R_{y'y}^Y) \\ &= - \sum_{y'} \delta_{y'y''} (\delta_{y,y'} \left(\sum_{x'} R_{(x'y)(xy)}^{X|Y} \right) + R_{y'y}^Y) = -R_{y''y}^Y, \end{aligned} \quad (\text{S41})$$

In the last line, we used (S40). Given (S41), $BRp = 0$ implies that

$$[BRp]_{y''} = - \sum_{x,y} R_{y''y} p_{xy} = - \sum_y R_{y''y} p_y^Y = 0. \quad (\text{S42})$$

Therefore, the marginal distribution p^Y must be in steady state for the rate matrix R^Y ,

$$p^Y = \pi^Y. \quad (\text{S43})$$

Therefore, local control (S35) and independent environment (S39) jointly imply $p^Y = \pi^Y$.

2.1 Independent agent-and-environment

Where explicitly noted, we sometimes also introduce the *Independent agent-and-environment* assumption, which is stronger than the *Independent environment* assumption discussed above. This assumption states that under the baseline rate matrix R , the dynamics of the agent and environment are bipartite and both evolve independently of each other. Formally, this means that the combined rate matrix has the form

$$\textbf{Independent agent-and-environment: } R_{(x'y')(xy)} = \delta_{yy'} R_{x'x}^X + \delta_{xx'} R_{y'y}^Y. \quad (\text{S44})$$

Here R^X and R^Y are rate matrices over the agent X and environment Y which satisfy $\sum_{x'} R_{x'x}^X = 0$ for all x and $\sum_{y'} R_{y'y}^Y = 0$ for all y , and which have unique steady states π^X and π^Y .

2.2 Linear Response (LR) regime

We now derive the LR regime solution of (S38), that is the the maximum power gain $\Delta \dot{W}_{\text{LR}}^*$ for a system coupled to a fluctuating environment. We then show how imposing additional assumptions on the system's dynamics can simplify this solution dramatically. This simplification will be used in the example explored in Sec. 4 of this SM.

2.2 General case

We first derive LR solution of (S37) in the general case, assuming only *Local control* (S35). To begin, we expand the state space considered in Sec. 1.1 to include both X and Y and study the constrained optimization problem:

$$\Delta \dot{W}_{\text{LR}}^* = \frac{1}{\beta} \max_{z \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{Y}|}; (zD^{-1}\pi)=0} z^T M z + 2z^T v \quad \text{such that } BRDz = 0. \quad (\text{S45})$$

where we define z , D , A , and $M = D^{-1}AD$ as in (S4) and (S5) above, and B as in (S36). The derivation of this expression from (S37) is the same as the derivation of (S9), except that we included the constraint $BRp = 0$, which is equivalent to $BRDz = 0$ given $z = D^{-1}(p - \pi)$.

The optimization problem (S45) is an equality-constrained quadratic optimization problem, which can be solved using standard techniques. Using Theorem 1 from Sec. 6 of this SM (S116), the solution is given by

$$\begin{aligned} \Delta \dot{W}_{\text{LR}}^* &= -v^T M^+ v + v^T \Lambda M \Lambda v \\ z^* &= (\Lambda - M^+) v \quad \implies \quad p^* = \pi + D(\Lambda - M^+) v, \end{aligned} \quad (\text{S46})$$

where we defined the matrices $\Lambda, \Pi \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{Y}| \times |\mathcal{X}| \times |\mathcal{Y}|}$:

$$\Lambda := \left[I - \left((-M)^{1/2} \Pi \right)^+ (-M)^{1/2} \right] M^+ \quad \Pi := I - \begin{bmatrix} BRD \\ \pi^T D^{-1} \end{bmatrix}^+ \begin{bmatrix} BRD \\ \pi^T D^{-1} \end{bmatrix} \quad (\text{S47})$$

The matrix Λ encodes the effect of the local control constraints on the solution of the optimization problem. The solution (S46) may be compared with the unconstrained LR solution (S10); note that the solution agree when $\Lambda = 0$.

It is important to note that even though we focused specifically on the problem of power gain under local control, the solution (S46) is more general. In particular, it may be used to study maximum power gain under various types of linear constraints of the form $G(p - \pi) = 0$ for some arbitrary matrix G . For instance, this may be used to study maximum power gain under conservation laws, or other types of constraints on the allowed steady state distributions $\bar{\pi}$.

In the following subsections, we demonstrate that the solution (S46) can be simplified by introducing the *Independent environment* (S39) and *Independent agent-and-environment* (S44) assumptions.

2.2 Simplifying assumption: Independent environment

We show that under the assumption of *Independent environment* (S39), Π can be expressed in a simpler form than (S47), and in particular without invoking the pseudo-inverse. First, define the following matrix $C \in \mathbb{R}^{|X| \times |Y| \times |X| \times |Y|}$:

$$C := \sum_y \boldsymbol{\tau}^{(y)} \boldsymbol{\tau}^{(y)T}, \quad \boldsymbol{\tau}_{x'y'}^{(y)} = \sqrt{\pi_{x'|y'}} \delta_{yy'} \quad \text{for } y \in Y. \quad (\text{S48})$$

We will show that *Independent environment* (S39) implies that

$$\Pi = I - C. \quad (\text{S49})$$

The derivation of (S49) proceeds in the following way. First, observe that the vectors $\boldsymbol{\tau}^{(y)}$ are orthonormal, $\langle \boldsymbol{\tau}^{(y)}, \boldsymbol{\tau}^{(y')} \rangle = \delta_{yy'}$, and therefore that $CC = C$. In addition, C is symmetric, which implies that it is an orthogonal projector. Next, for convenience define the matrix $H = \begin{bmatrix} BRD \\ \pi^T D^{-1} \end{bmatrix}$, so that $\Pi = I - H^+H$. Below, we first show that the rowspace of H is a subspace of the column space of C . After that, we show that the dimension of the column space of C is equal to the dimension of the rowspace of H . Therefore, the two spaces must be equal, so C is the orthogonal projector onto the rowspace of H . But it is known that H^+H is the orthogonal projector onto the rowspace of H . Thus, it must be that $C = H^+H$, implying (S49).

Let us show that the rowspace of H is a subspace of the column space of C . The entries of C and H can be written explicitly as

$$C_{(xy)(x'y')} = \sqrt{\pi_{x|y}} \sqrt{\pi_{x'|y'}} \delta_{y'y'}, \quad H_{y''(xy)} = \begin{cases} R_{y''y}^Y \sqrt{\pi_{xy}} & y'' \in \{1, \dots, |Y|\} \\ \sqrt{\pi_{xy}} & y'' = |Y| + 1 \end{cases}. \quad (\text{S50})$$

Here we have used the expression of BR (S41) that holds under the *Independent environment assumption*. Then, each row of H is a linear combination of the columns of C ,

$$\begin{aligned} R_{y''y}\sqrt{\pi_{xy}} &= \sum_{x',y'} c_{x'y'}^{(1,y'')} \sqrt{\pi_{x|y}} \sqrt{\pi_{x'|y'}} \delta_{y'y} & c_{x'y'}^{(1,y'')} &= R_{y''y'}^Y \sqrt{\pi_{x'y'}} \\ \sqrt{\pi_{xy}} &= \sum_{x',y'} c_{x'y'}^{(2)} \sqrt{\pi_{x|y}} \sqrt{\pi_{x'|y'}} \delta_{y'y} & c_{x'y'}^{(2)} &= \sqrt{\pi_{x'y'}}. \end{aligned}$$

This means that the rowspace of H is a subspace of the column space of C .

Next, we show that the two spaces have the same dimension. It is clear that the column space of C has dimension $|Y|$ by construction (S48). To determine the rank of H , observe that

$$H = \begin{bmatrix} BRD \\ \pi^T D^{-1} \end{bmatrix} = \begin{bmatrix} BR \\ \mathbf{1}^T \end{bmatrix} D = (K \otimes \mathbf{1}\mathbf{1}^T) D,$$

where $K = \begin{bmatrix} R^Y \\ \mathbf{1}^T \end{bmatrix}$ and \otimes is the Kronecker product, where we again used (S41). By basic properties of rank,

$$\text{rank } H = \text{rank}(K \otimes \mathbf{1}\mathbf{1}^T) = (\text{rank } K)(\text{rank } \mathbf{1}\mathbf{1}^T) = \text{rank } K.$$

Finally, define $P_t = [tI_Y \ \mathbf{1}] K$ for $t > 0$ where I_Y is the $|Y| \times |Y|$ identity matrix, and note that $\text{rank } P_t \leq \text{rank } K = \text{rank } H$. For sufficiently small t , P_t is an $|Y| \times |Y|$ transition matrix which is strictly diagonally dominant, and hence has full rank. Therefore, $|Y| = \text{rank } P_t$, implying that $|Y| \leq \text{rank } H$. At the same time, since the rowspace of H is contained in column space of C , we must have $\text{rank } H \leq \text{rank } C = |Y|$. Combining implies $\text{rank } H = |Y| = \text{rank } C$.

2.2 Simplifying assumption: Independent-agent-and-environment

We now simplify the solution (S46) under the assumption of *Independent agent-and-environment* (S44). To begin, we use (S44) to write the matrix $M = D^{-1}AD$ in explicit form as

$$\begin{aligned} M_{(x'y')(xy)} &= \frac{1}{2} \left(\sqrt{\frac{\pi_{xy}}{\pi_{x'y'}}} R_{(x'y')(xy)} + \sqrt{\frac{\pi_{x'y'}}{\pi_{xy}}} R_{(xy)(x'y')} \right) \\ &= \frac{1}{2} \left(\sqrt{\frac{\pi_x^X}{\pi_{x'}^X}} R_{x'x}^X + \sqrt{\frac{\pi_{x'}^X}{\pi_x^X}} R_{xx'}^X \right) \delta_{y'y} + \frac{1}{2} \left(\sqrt{\frac{\pi_y^Y}{\pi_{y'}^Y}} R_{y'y}^Y + \sqrt{\frac{\pi_{y'}^Y}{\pi_y^Y}} R_{y'y'}^Y \right) \delta_{x'x} \\ &\equiv M_{x'x}^X \delta_{y'y} + M_{y'y}^Y \delta_{x'x}, \end{aligned} \tag{S51}$$

which for convenience we defined the matrix $M^X \in \mathbb{R}^{|X| \times |X|}$ and $M^Y \in \mathbb{R}^{|Y| \times |Y|}$ as above. We also used that (S44) implies that the steady state distribution factorizes as $\pi_{xy} = \pi_x^X \pi_y^Y$. Observe that (S51) can be written in matrix notation as

$$M = M^X \otimes I_Y + I_X \otimes M^Y, \tag{S52}$$

where I_X and I_Y are identity matrices of size $|X| \times |X|$ and $|Y| \times |Y|$.

Let us denote the orthonormal eigensystems of M^X and M^Y as $(\lambda_a^X, m_{(a)}^X)$ and $(\lambda_b^Y, m_{(b)}^Y)$ with $a \in \{1, \dots, |X|\}$ and $b \in \{1, \dots, |Y|\}$. Since M^X and M^Y are related to the rate matrices R^X and R^Y by a similarity transformation, the top eigenvectors for M^X and M^Y (with $\lambda_1^X = \lambda_1^Y = 0$) correspond to the marginal steady states $m_{(1)x}^X = \sqrt{\pi_x^X}$ and $m_{(1)y}^Y = \sqrt{\pi_y^Y}$, respectively. Moreover, it can be shown that for all for $a \in \{1, \dots, |X|\}$ and $b \in \{1, \dots, |Y|\}$, $m_{(ab)} := m_{(a)}^X \otimes m_{(b)}^Y$ is an eigenvector of M with eigenvalue $\lambda_{ab} = \lambda_a^X + \lambda_b^Y$:

$$\begin{aligned} Mm_{(ab)} &= (M^X \otimes I_Y + I_X \otimes M^Y) \left(m_{(a)}^X \otimes m_{(b)}^Y \right) = \lambda_a^X m_{(a)}^X \otimes m_{(b)}^Y + \lambda_b^Y m_{(a)}^X \otimes m_{(b)}^Y \\ &= (\lambda_a^X + \lambda_b^Y) m_{(a)}^X \otimes m_{(b)}^Y = \lambda_{ab} m_{(ab)}. \end{aligned} \quad (\text{S53})$$

To summarize, we can write the eigendecomposition of M as

$$M = \sum_{a,b} (\lambda_a^X + \lambda_b^Y) m_{(ab)} m_{(ab)}^T = \sum_{a,b} (\lambda_a^X + \lambda_b^Y) m_{(ab)} m_{(a)}^X m_{(a)}^{X T} \otimes m_{(b)}^Y m_{(b)}^{Y T}. \quad (\text{S54})$$

Next, because the steady state factorizes, we can rewrite C (S50) as $C_{(xy)(x'y')} = \sqrt{\pi_x^X} \sqrt{\pi_{x'}^X} \delta_{y'y'}$, or in matrix notation

$$C = m_{(1)}^X m_{(1)}^{X T} \otimes I_X = \sum_b m_{(1)}^X m_{(1)}^{X T} \otimes m_{(b)}^Y m_{(b)}^{Y T}. \quad (\text{S55})$$

In addition, the assumption *Independent agent-and-environment* implies *Independent environment* (S39), which in turn implies $\Pi = I - C$ (S49). Combining gives

$$\begin{aligned} \Pi = I - C &= \sum_{a,b} m_{(a)}^X m_{(a)}^{X T} \otimes m_{(b)}^Y m_{(b)}^{Y T} - \sum_b m_{(1)}^X m_{(1)}^{X T} \otimes m_{(b)}^Y m_{(b)}^{Y T} \\ &= \sum_{a>1,b} m_{(a)}^X m_{(a)}^{X T} \otimes m_{(b)}^Y m_{(b)}^{Y T} = \sum_{a>1,b} m_{(ab)} m_{(ab)}^T. \end{aligned} \quad (\text{S56})$$

In this way, we expressed Π in terms of the eigenvectors of M .

It is clear that M and Π commute. Therefore, we can apply (S117) from Theorem 1 in Sec. 6. It states that the solution of (S45) is given by

$$\begin{aligned} \Delta \dot{W}_{\text{LR}}^* &= -v^T \Pi M^+ \Pi v \\ z^* &= -(M^+) \Pi v \quad \implies \quad p^* = \pi - D M^+ \Pi v. \end{aligned} \quad (\text{S57})$$

Using the expression of the pseudo-inverse $M^+ = \sum_{a,b:(a,b) \neq (1,1)} (\lambda_a^X + \lambda_b^Y)^{-1} m_{(ab)} m_{(ab)}^T$ and our expression for Π (S56), we have

$$M^+ \Pi = \Pi M^+ \Pi = \sum_{a,b:a>1} (\lambda_a^X + \lambda_b^Y)^{-1} m_{(ab)} m_{(ab)}^T. \quad (\text{S58})$$

Finally, combining with (S57) allows us to write the solution as

$$\Delta \dot{W}_{\text{LR}}^* = \beta \sum_{a,b:a>1} \frac{|\Omega_{ab}|^2}{-\lambda_a - \lambda_b} \quad p^* = \pi + \sum_{a,b:a>1} \frac{\beta \Omega_{ab}}{-\lambda_a - \lambda_b} u_{(ab)} \quad (\text{S59})$$

where we defined

$$\beta \Omega_{ab} = \langle v, m_{(ab)} \rangle = \frac{\beta}{2} \langle D \psi^{\text{LR}}, D^{-1} u_{(ab)} \rangle = \frac{\beta}{2} \langle \psi^{\text{LR}}, u_{(ab)} \rangle, \quad (\text{S60})$$

exactly as in (S12), and $u_{(ab)}$ is an eigenvalue of the matrix A (see (S4) and paragraph after). This result is in complete analogy with the static environment LR analysis (S13) and Eq. (9) in the main text.

2.3 Macroscopic (M) regime

We now derive the Macroscopic (M) regime solution for a system coupled to a fluctuating environment. To do so, we will impose the assumptions of *Local control* (S35) and the *Independent environment* (S39).

Recall our definition of the macroscopic payoff vector $\psi^{\text{M}} := d - q - \beta^{-1} R^T s$ from Sec. 1.2. Also recall (S16),

$$\Delta \dot{W}^* = \max_{p:BRp=0} \frac{1}{\beta} \langle Rp, \ln p \rangle + \langle p - \pi, \psi^{\text{M}} \rangle, \quad (\text{S61})$$

Given *Local control* and *Independent environment*, $BRp = 0$ implies that $p^Y = \pi^Y$ (S43). In addition, as in Sec. 1.2, we assume that in the Macroscopic regime we may neglect the nonlinear first term in (S61), which gives the approximation

$$\Delta \dot{W}^* \approx \max_{p_{x|y}} \sum_y \pi_y p_{x|y} \psi_{xy}^{\text{M}} - \langle \pi, \psi^{\text{M}} \rangle. \quad (\text{S62})$$

This is a simple linear optimization, where the optimal conditional distribution is given by $p_{x|y}^* = \delta_{x_y^* x}$ with $x_y^* = \arg \max_x \psi_{xy}^{\text{M}}$, i.e., the x_y^* is the best response state of X given that the environment is in state $Y = y$. The approximate solution under this Macroscopic Regime is

$$\Delta \dot{W}_{\text{M}}^* = \sum_y \pi_y \psi_{x_y^* y}^{\text{M}} - \langle \pi, \psi^{\text{M}} \rangle \quad p_{xy}^* \approx \delta_{x_y^* x} \pi_y, \quad (\text{S63})$$

2.4 Far-from-Equilibrium (FE) regime

We now derive the Far-from-equilibrium (FE) regime solution for a system coupled to a fluctuating environment. To do so, we will impose the assumptions of *Local control* (S35) and the *Independent environment* (S39).

We derive the Far-from-Equilibrium (FE) approximation as an expansion around the Macroscopic Regime solution (S63). Consider the nonlinear term in (S61),

$$\langle Rp, \ln p \rangle = \sum_{x,y} [pR]_{xy} \ln p_{xy} = \sum_{x,y} [pR]_{xy} \ln p_{x|y},$$

where we use the that $\sum_y (\sum_x [pR]_{xy}) \ln p_y = 0$ since $p^Y = \pi^Y$. We approximate the remaining term by using the expression for p^* in (S63),

$$\begin{aligned} \sum_{x,y} \sum_{x',y'} p_{x'y'} R_{(xy)(x'y')} \ln p_{x|y} &\approx \sum_{x,y} \sum_{y'} \pi_{y'} R_{(x,y)(x'_y, y')} \ln p_{x|y} \\ &= \sum_{y,y'} \pi_{y'} \left(\sum_{x \neq x_y^*} R_{(x,y)(x'_y, y')} \ln p_{x|y} + R_{(x_y^*, y)(x'_y, y')} \ln p_{x_y^*|y} \right). \end{aligned} \quad (\text{S64})$$

Let us now expand the last term as follows

$$\ln p_{x_y^*|y} = \ln \left(1 - \sum_{x \neq x_y^*} p_{x|y} \right) \approx - \sum_{x \neq x_y^*} p_{x|y}.$$

Next, we define

$$\mathcal{R}_{xy} := \sum_{y'} \pi_{y'} R_{(x,y)(x'_y, y')}, \quad (\text{S65})$$

which indicates the rate of probability flow from all ‘‘optimal states’’ (x'_y, y') to state (x, y) . Thus, we obtain that

$$\langle Rp, \ln p \rangle \approx \sum_{y, x \neq x_y^*} (\mathcal{R}_{xy} \ln p_{x|y} - \mathcal{R}_{x_y^*y} p_{x|y}). \quad (\text{S66})$$

Next, let us consider the linear term in (S61) and use the fact that $p^Y = \pi^Y$ to write

$$\begin{aligned} \langle p, \psi^M \rangle &= \sum_{x,y} p_{x|y} \pi_y \psi_{xy}^M = \sum_y \pi_y \left(p_{x_y^*|y} \psi_{x_y^*y}^M + \sum_{x \neq x_y^*} p_{x|y} \psi_{xy}^M \right) \\ &= \sum_y \pi_y \left[\left(1 - \sum_{x \neq x_y^*} p_{x|y} \right) \psi_{x_y^*y}^M + \sum_{x \neq x_y^*} p_{x|y} \psi_{xy}^M \right] \\ &= \sum_y \pi_y \left[\psi_{x_y^*y}^M + \sum_{x \neq x_y^*} p_{x|y} (\psi_{xy}^M - \psi_{x_y^*y}^M) \right]. \end{aligned} \quad (\text{S67})$$

Then, combining (S63), (S66) and (S67),

$$\Delta \dot{W}^{*} \approx \max_{p: BRp=0} \Delta \dot{W}_M^* + \frac{1}{\beta} \sum_{y, x \neq x_y^*} \left[\mathcal{R}_{xy} \ln p_{x|y} - (\beta (\psi_{x_y^*y}^M - \psi_{xy}^M) \pi_y + \mathcal{R}_{x_y^*y}) p_{x|y} \right]. \quad (\text{S68})$$

This optimization problem can be solved by taking derivatives with respect to $p_{x|y}$, which, after imposing $p^Y = \pi^Y$, gives

$$p_{xy}^* = \pi_y \left[\delta_{x_y^*y} + \sum_{x' \neq x_y^*} \frac{(\delta_{xx'} - \delta_{xx_y^*}) \mathcal{R}_{x'y}}{\beta (\psi_{x_y^*y}^M - \psi_{x'y}^M) \pi_y + \mathcal{R}_{x_y^*y}} \right]. \quad (\text{S69})$$

Finally, inserting back into (S68) yields the FE approximation

$$\Delta \dot{W}_{\text{FE}}^* = \Delta \dot{W}_{\text{M}}^* + \frac{1}{\beta} \sum_{y, x \neq x_y^*} \mathcal{R}_{xy} (\ln p_{x|y}^* - 1), \quad (\text{S70})$$

which is a generalization of (S33).

3. Unicyclic model

Here we analyze the unicyclic model considered in the main text. We study both the case with a static environment in Sec. 3.2 and the case with a fluctuating environment in Sec. 3.3. Before proceeding, we briefly introduce some useful facts about the eigendecomposition of a unicyclic rate matrix.

3.1 An Algebraic Aperitif: eigendecomposition of unicyclic rate matrix

Consider a unicyclic rate matrix R ,

$$R = \begin{pmatrix} -2k & k & 0 & \cdots & k \\ k & -2k & k & \cdots & 0 \\ 0 & k & -2k & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ k & 0 & 0 & \cdots & -2k \end{pmatrix}, \quad (\text{S71})$$

which is a symmetric circulant matrix. Due to symmetry, its steady state is uniform: $\pi_i = \frac{1}{n}$, $\forall i$. The eigensystem for R is obtained from the theory of circulant matrices [42], which, for odd n , yields:

$$\lambda_a = -2k [1 - \Re(\omega_a)] = -2k \left[1 - \cos\left(\frac{2\pi(a-1)}{n}\right) \right], \quad (\text{S72})$$

where $\omega_a := e^{i2\pi(a-1)/n}$. These eigenvalues are all degenerate twice (except $\lambda_{(1)} = 0$, whose eigenvector is simply $(1, 1, \dots, 1)$). An orthonormal choice of eigenbasis is given by the set

$$\{m_{(a)}\} = \left\{ \frac{1}{\sqrt{n}} (1, \omega_a, \omega_a^2, \dots, \omega_a^{n-1}) \right\}. \quad (\text{S73})$$

3.2 Unicyclic model: Static environment

We now analyze the unicyclic model in a static environment, using techniques developed in Sec. 1. We consider a systems with n state arranged in a ring, where baseline transitions are symmetric with uniform kinetics: $i \xrightarrow{k} k \Leftrightarrow i+1$ mod n , $\forall i = 1, \dots, n$ (see Fig. S3 here and Fig. 1(a) in the main text). The baseline dynamics are equivalent to a random walk on a one-dimensional ring, and the baseline rate matrix is given by (S71).

We assume that a single transition is coupled to the extraction of Θ work, without heat generation or internal entropy change ($q = 0$ and $s = 0$). The reverse of this transition removes the same amount of energy from the system's internal work reservoir. Without loss of generality, we assume that such specific reaction is $1 \xrightarrow{k} k \Leftrightarrow 2$. Note




Figure S3. A unicyclic system with transitions across the n states with symmetric back-and-forth rates k , as discussed in the main text. Work Θ is gained every time the system undergoes the transition $1 \rightarrow 2$, while the work Θ is lost in the reverse transition $2 \rightarrow 1$.

that in the main text we have chosen to set the units of $k = 1$ without loss of generality. However, we have kept k explicit in the rest of the SM. For this system, the baseline steady state is uniform, $\pi_i = 1/n$, and the payoff vector is given by

$$\psi^{\text{LR}} = \psi^{\text{M}} = (k\Theta, -k\Theta, 0, \dots, 0). \quad (\text{S74})$$

In the rest of this section, we analyze this model in the LR and FE regime. We will also discuss the validity of these approximation in terms of the model parameters.

3.2 Static environment: LR regime

We first consider the LR solution (S13). Recall that work Θ is extracted from the environment during transitions from states $1 \rightarrow 2$, and its reverse reaction requires the same amount of work to be invested. Combining (S74) with (S5), $\pi_i = 1/n$, and $D = \frac{1}{\sqrt{n}}I_n$ gives $v = \frac{k\beta\Theta}{2\sqrt{n}}(1, -1, 0, \dots, 0)$. In addition, since R obeys detailed balance, $R = A$. Finally, it is easy to verify that $M = D^{-1}AD = A = R$. Thus, M has the same eigendecomposition as R . The normalized eigenvectors of M are given by (S73).

We can compute $\Omega_a = \langle v, m_{(a)} \rangle / \beta$ for $a = 2, \dots, n$, using the eigenvector set given in (S73), which simply yields

$$\Omega_a = \frac{k\Theta}{2n} (1 - \omega_a). \quad (\text{S75})$$

Next, we compute the maximum power difference attainable using (S13):

$$\begin{aligned} \Delta \dot{W}_{\text{LR}}^* &= \sum_{a=2}^n \frac{\beta |\Omega_a|^2}{-\lambda_a} = \beta k \left(\frac{\Theta}{2n} \right)^2 \sum_{a=2}^n \frac{|1 - \omega_a|^2}{2 \left[1 - \cos \left(\frac{2\pi(a-1)}{n} \right) \right]} \\ &= \beta k \left(\frac{\Theta}{2n} \right)^2 \sum_{a=2}^n 1 = \beta k \left(\frac{\Theta}{2n} \right)^2 (n-1). \end{aligned} \quad (\text{S76})$$

On the other hand, continuing with odd n , it is also possible to compute the deviation of the optimal distribution from stationary distribution using (S13),

$$\Delta p^* = p^* - \pi = \beta \sum_{a=2}^n \frac{\Omega_a^\dagger D m_{(a)}}{-\lambda_a}, \quad (\text{S77})$$

which is depicted in Fig. S4(a). (Note that, due to our choice of a complex-valued eigenbasis, we must be careful in adding the complex conjugate on the Ω_a , following the ordering of the operator M^+ in (S13).) Now, component by component, we can rewrite (S77) as

$$\begin{aligned} \Delta p_i^* &= \frac{\beta\Theta}{4n^2} \sum_{a=2}^n \frac{(1 - \omega_a)^\dagger \omega_a^{i-1}}{1 - \cos\left(\frac{2\pi(a-1)}{n}\right)} \\ &= \frac{\beta\Theta}{4n^2} \sum_{a=2}^n \frac{\cos\left(\frac{2\pi(a-1)(i-1)}{n}\right) - \cos\left(\frac{2\pi(a-1)(i-2)}{n}\right)}{1 - \cos\left(\frac{2\pi(a-1)}{n}\right)} \\ &= \frac{\beta\Theta}{4n^2} \left[-\sum_{a=2}^n \cos\left(\frac{2\pi(a-1)(i-2)}{n}\right) - \sum_{a=2}^n \frac{\sin\left(\frac{2\pi(a-1)(i-2)}{n}\right) \sin\left(\frac{2\pi(a-1)}{n}\right)}{1 - \cos\left(\frac{2\pi(a-1)}{n}\right)} \right]. \end{aligned} \quad (\text{S78})$$

In the first line, we used the expressions for Ω_a from (S75), λ_a from (S72), $m_{(a)}$ from (S73), and $D = I/\sqrt{n}$ and then simplified. In the second line, we expanded $(1 - \omega_a)^\dagger \omega_a^{i-1} = \omega_a^{i-1} - \omega_a^{i-2}$ into real and imaginary components and then used that the imaginary components cancel over the sum. In the last line, we used the identity $\cos(x+y) = \cos(x)\cos(y) - \sin(x)\sin(y)$ for $x = \frac{2\pi(a-1)(i-2)}{n}$ and $y = \frac{2\pi(a-1)}{n}$ and then simplified. It can be verified that the first series in (S78) sums to -1 . The second series is trickier but can be simplified using trigonometric methods as shown in Ref. [43]. Plugging in that solution and simplifying gives the very simple expression:

$$\Delta p_i^* = \frac{\beta\Theta}{4n^2} \begin{cases} 2(i-1) - (n+1) & \text{for } i = 2, \dots, n \\ (n-1) & \text{for } i = 1 \end{cases}. \quad (\text{S79})$$

Thus, in the LR regime, the optimal distribution builds up in equal increments starting at $i = 2$ until the optimal state $i = 1$, after which it falls off a cliff. This is shown in Fig. S4.

Note that it is possible to compute the probability current across the transition $1 \rightarrow 2$, which we also leave as an exercise to the reader:

$$k(p_1^* - p_2^*) = \frac{k\beta\Theta}{2n} \left(\frac{n-1}{n}\right). \quad (\text{S80})$$

The expressions (S14) and (S15) suggest which values of Θ will cause the optimal solution to be in the LR regime. In this case, however, it is possible to work out the norm $\|p^* - \pi\|$ exactly:

$$\|\Delta p^*\| = \frac{\beta\Theta}{4n^2} \left[\sum_{x=2}^n (2x - n - 3)^2 + (n-1)^2 \right]^{1/2} = \frac{\beta\Theta}{4\sqrt{3}} \sqrt{\frac{n^2 - 1}{n^3}}$$




Figure S4. The optimal Δp^* (deviation of optimal distribution from the uniform baseline steady state) given (S79). The optimal distribution exhibits a clockwise cyclic current given by (S80). The magnitude of Δp^* decays as $\sim n^{-1}$, so larger rings will result in optimal distributions that are closer to the baseline steady state.

which implies that the LR regime is valid when

$$\|p^* - \pi\| \ll 1 \Leftrightarrow \Theta \ll \frac{4\sqrt{3}}{\beta} \sqrt{\frac{n^3}{n^2 - 1}} \quad (\text{S81})$$

In this case, the larger the number of states n in the ring, the wider the interval of Θ values that will make the optimal solution fall into the LR regime, i.e. for large n we will obtain

$$\Theta \ll \frac{4\sqrt{3}}{\beta} \sqrt{n}.$$

3.2 Static environment: FE Regime

Consider the FE expressions (S32) with (S74). Notice that $i^* = 1$, $\Delta\psi_2^M = \psi_{i^*}^M - \psi_2^M = 2k\Theta$ and $\Delta\psi_n^M = \psi_{i^*}^M - \psi_n^M = k\Theta$, while $R_{21} = R_{n1} = k$, $R_{11} = -2k$ and $R_{j1} = 0$ for all $j \neq 2, n$. Hence, for $a > 1$, we obtain

$$p_a^* = \frac{R_{a1}}{\beta\Delta\psi_a^M + R_{11}} = \begin{cases} \frac{1}{2(\beta\Theta - 1)} & a = 2 \\ \frac{1}{\beta\Theta - 2} & a = n \\ 0 & \text{otherwise} \end{cases} \quad (\text{S82})$$

$$p_1^* = 1 - \sum_{a=2}^n p_a^* = 1 - \frac{3\beta\Theta - 4}{2(\beta\Theta - 1)(\beta\Theta - 2)}.$$

The maximum power gain is given by Eq. (12),

$$\Delta\dot{W}_{\text{FE}}^* = \psi_1^M + \beta^{-1} \sum_{a=2}^n R_{a1} (\ln p_a^* - 1) \quad (\text{S83})$$

$$= k(\Theta - 2/\beta) - k\beta^{-1} \ln [2(\beta\Theta - 1)(\beta\Theta - 2)]. \quad (\text{S84})$$

Here we used the fact that, for this model, $\langle \pi, \psi^M \rangle = 0$.

3.3 Unicyclic model: Fluctuating environment

We now extend our analysis by considering a more elaborate model, illustrated in Fig. 1(b) in the main text. Here, the rate of work extraction can depend on the state of an explicitly modeled environment, requiring the use of techniques developed in Sec. 2. We will derive bounds on maximum power gain under the *Local control* assumption (see Sec. 2 1 a).

Suppose that the system X and environment Y both have n states, and they individually undergo transitions over a ring of n states. The baseline dynamics for system X is given by the unbiased random walk as in discussed in Sec. 3.2. The baseline dynamics over environment Y are given by a biased random walk parameterized as in (S85) (see below). Crucially, the work gained by the system X depends on the state of Y . Specifically, Θ of work is delivered each time system X undergoes the transition $X = y$ to $X = y + 1$, where y is the state of environment Y .

We parameterize the transition rates over Y in the following manner:

$$R_{y'y}^Y = \begin{cases} q & ; y' = y + 1 \bmod n \\ (1 + \epsilon)q & ; y' = y - 1 \bmod n \\ -(2 + \epsilon)q & ; y' = y \\ 0 & ; \text{otherwise} \end{cases}, \quad (\text{S85})$$

for asymmetry parameter ϵ . When the states of Y are arranged as ring and numbered following the convention of a clock and $\epsilon > 0$, the state of the environment tends to cycle counterclockwise. It may be verified that the global dynamics are bipartite and independent, i.e. $R = R^X \otimes I_Y + I_X \otimes R^Y$, and $R \in \mathbb{R}^{n^2 \times n^2}$ is an irreducible rate matrix. Thus, the *Independent agent-and-environment* assumption (S44) is satisfied, since there is no explicit dependence of x in the transitions of the Y partition, R^Y , and vice-versa.

Owing to the symmetry of the setup, the baseline steady-state distribution is uniform, $\pi_{xy} = 1/n^2, \forall x, y$. Moreover, $\pi_{xy} = \pi_x \pi_y$ with $\pi_x = \pi_y = 1/n$, and therefore also $p_y = 1/n$.

3.4 LR regime

As discussed in Sec. 2 2 c, we consider the eigensystems of the matrices M^X and M^Y . In this example $D = (1/n)I_{n^2}$, and $M = A = (R + R^T)/2$. By symmetry $R^X = (R^X)^T$, while (for $\epsilon \neq 1$) $R^Y \neq (R^Y)^T$. Let us denote $Q^Y := (R^Y + (R^Y)^T)/2$ such that $M = R^X \otimes I_Y + I_X \otimes Q^Y$. In matrix form:

$$Q^Y = \begin{pmatrix} -(2+\epsilon)q & (1+\frac{\epsilon}{2})q & 0 & \cdots & (1+\frac{\epsilon}{2})q \\ (1+\frac{\epsilon}{2})q & -(2+\epsilon)q & (1+\frac{\epsilon}{2})q & \cdots & 0 \\ 0 & (1+\frac{\epsilon}{2})q & -(2+\epsilon)q & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ (1+\frac{\epsilon}{2})q & 0 & 0 & \cdots & -(2+\epsilon)q \end{pmatrix}. \quad (\text{S86})$$

Next, consider the orthonormal eigensystems $(\lambda_a^X, m_{(a)}^X)$ of R^X and $(\lambda_b^Y, m_{(b)}^Y)$ of Q^Y with indexes $a, b = 1, \dots, n$. As shown above, $\lambda_1^X = \lambda_1^Y = 0$, which correspond to the steady states of X and Y , i.e. $m_{(1)x}^X = \sqrt{\pi_x} = \sqrt{1/n} = \sqrt{\pi_y} = m_{(1)y}^Y$. The eigensystem of M is given by $\{\lambda_{ab} = \lambda_a^X + \lambda_b^Y, m_{(ab)} = m_{(a)}^X \otimes m_{(b)}^Y\}$.

From the properties of circulant matrices [42],

$$\lambda_{ab} = -2k \left[1 - \cos\left(\frac{2\pi(a-1)}{n}\right) \right] - q(2+\epsilon) \left[1 - \cos\left(\frac{2\pi(b-1)}{n}\right) \right] \quad (\text{S87})$$

Note that, for odd n (odd n^2), the eigenvalues are four times degenerate, since transformations like $a \rightarrow n+2-a$ or $b \rightarrow n+2-b$ yield the same eigenvalue of $1 < a, b < (n+1)/2$. For even n (even n^2) there is a set of eigenvalues that are only twice degenerate, namely those indexed as $\{a(n/2+1)\}$ with $1 < a < n/2+1$, and those indexed as $\{(n/2+1)b\}$ with $1 < b < n/2+1$, plus a single non-degenerate eigenvalue at $ab = (n/2+1)(n/2+1)$.

The orthonormal eigenvector system is given by $\{m_{(ab)} = m_{(a)} \otimes m_{(b)}\}$, where $m_{(c)} := n^{-1/2} (1, \omega_c, \omega_c^2, \dots, \omega_c^{n-1})$ and $\omega_c := e^{i2\pi(c-1)/n}$ for $c = a, b$. Hence,

$$m_{(ab)} = \frac{1}{n} (1, \omega_a, \omega_a^2, \dots, \omega_a^{n-1}) \otimes (1, \omega_b, \omega_b^2, \dots, \omega_b^{n-1}), \quad (\text{S88})$$

which can be written in xy coordinates as

$$[m_{(ab)}]_{xy} = \frac{1}{n} \omega_a^{x-1} \omega_b^{y-1} = \frac{1}{n} e^{\frac{2\pi i}{n}[(a-1)(x-1)+(b-1)(y-1)]}. \quad (\text{S89})$$

In this case, it is convenient to work in the complex form for the eigenvector system.

We are now in a position to compute $\Delta \dot{W}_{\text{LR}}^*$ and p^* . Recall the definition of v from (S5). By symmetry $R^T \ln \pi = 0$ and, by assumption, $q = s = 0$. Hence, $v = \frac{\beta}{2} Dd$. The driving vector $d \in \mathbb{R}^{n^2}$ is given by

$$d_{xy} = \begin{cases} k\Theta (\delta_{xy} - \delta_{(x+1)y}) & x = 1, \dots, n-1 \\ k\Theta (\delta_{xy} - \delta_{1y}) & x = n \end{cases}, \quad \forall y = 1, \dots, n. \quad (\text{S90})$$

Thus, the dot product

$$\langle v, m_{(ab)} \rangle = \frac{\beta k \Theta}{2n^2} \left[\sum_{x \neq n, y} (\delta_{xy} - \delta_{(x+1)y}) \omega_a^{x-1} \omega_b^{y-1} + \sum_y (\delta_{ny} - \delta_{1y}) \omega_a^{n-1} \omega_b^{y-1} \right]$$

$$\begin{aligned}
&= \frac{\beta k \Theta}{2n^2} \left[\sum_{x \neq n} (\omega_a^{x-1} \omega_b^{x-1} - \omega_a^{x-1} \omega_b^x) + \omega_a^{n-1} (\omega_b^{n-1} - 1) \right] \\
&= \frac{\beta k \Theta}{2n^2} \sum_x \omega_{a+b-1}^{x-1} (1 - \omega_b).
\end{aligned} \tag{S91}$$

Hence, the dot product squared can be written, with a bit of algebra, as

$$\begin{aligned}
|\langle v, m_{(ab)} \rangle|^2 &= \left(\frac{\beta k \Theta}{2n^2} \right)^2 |1 - \omega_b|^2 \left| \sum_x \omega_{a+b-1}^{x-1} \right|^2 \\
&= \left(\frac{\beta k \Theta}{2n} \right)^2 2 \left[1 - \cos \left(\frac{2\pi(b-1)}{n} \right) \right] \delta(a+b \bmod n, 2).
\end{aligned} \tag{S92}$$

Since we will sum $1 < a \leq n$ and $1 \leq b \leq n$, then $1 < a+b \leq 2n$ and hence $a+b \bmod n \equiv 2$ iff $b = n - a + 2$.

Therefore,

$$\Delta \dot{W}_{\text{LR}}^* = \beta \left(\frac{k\Theta}{2n} \right)^2 \sum_{a>1} \frac{1 - \cos \left(\frac{2\pi(n-a+1)}{n} \right)}{k \left[1 - \cos \left(\frac{2\pi(a-1)}{n} \right) \right] + q \left(1 + \frac{\xi}{2} \right) \left[1 - \cos \left(\frac{2\pi(n-a+1)}{n} \right) \right]}. \tag{S93}$$

Now observe that $\cos \left(\frac{2\pi(n-a+1)}{n} \right) = \cos \left(\frac{2\pi(a-1)}{n} \right)$, so the summation above reduces to

$$\Delta \dot{W}_{\text{LR}}^* = k\beta\Theta^2 \frac{(n-1)}{4n^2(1+\xi)}. \tag{S94}$$

where we used $\xi := q \left(1 + \frac{\xi}{2} \right) / k$ to simplify notation. This is our main result for the fluctuating environment model, which appears as Eq. (17) in the main text.

Some interesting limiting cases are

- $q \left(1 + \frac{\xi}{2} \right) \ll k$, where the agent's baseline dynamics is much faster than the environment's. Then,

$$\begin{aligned}
\Delta \dot{W}_{\text{LR}}^* &= k\beta\Theta^2 \frac{(n-1)}{4n^2} (1 - \xi + \mathcal{O}(\xi^2)) \\
&= \left[k - q \left(1 + \frac{\xi}{2} \right) \right] \beta\Theta^2 \frac{(n-1)}{2n^2} (1 + \mathcal{O}(\xi^2)),
\end{aligned} \tag{S95}$$

which coincides with (S76) if $q \rightarrow 0$.

- $q \left(1 + \frac{\xi}{2} \right) \gg k$ where the environment's baseline dynamics is much faster than the agent's. Then,

$$\Delta \dot{W}_{\text{LR}}^* \approx \beta\Theta^2 \frac{k}{\xi} \frac{(n-1)}{4n^2} = \beta\Theta^2 \frac{k^2}{\left(1 + \frac{\xi}{2} \right)} \frac{(n-1)}{4n^2}. \tag{S96}$$

These limiting cases are closely related to our results in Sec. 1.1. Essentially, if one of the two timescales (either the agent's or the environment's) is much faster than the other, then the system is equivalent to the static ring with n states. Within the LR regime, the approximated solution of Δp^* can also be found using (S59) as

$$\Delta p_{xy}^* = \frac{\beta\Theta}{4n^3} \sum_{a>1} \frac{(1 - \omega_a^{-1}) \omega_a^{x-y}}{(1 + \xi) \left[1 - \cos \left(\frac{2\pi(a-1)}{n} \right) \right]}, \tag{S97}$$

with $\omega_a := e^{i2\pi(a-1)/n}$. Similarly to Section 3.2, the solution can be exactly computed to give:

$$\Delta p_{xy}^* = \frac{\beta\Theta}{4n^3(1+\xi)} \begin{cases} 2(x-y) - (n+1) & \text{for } x-y = 1, \dots, n-1 \\ (n-1) & \text{for } x=y \end{cases}, \quad (\text{S98})$$

The proof of (S98), which is left as an exercise to the reader, follows the same logic as the derivation of (S79). Plotting $\Delta p_{xy_0}^*$ for some fixed value of y_0 yields a graph equivalent to Fig. S4 after transforming $\Theta \rightarrow \Theta/(1+\xi)$ and adding the respective n factors. Here, the sharp probability deviation drop occurs at $x = x_y^* = y_0 \rightarrow y_0 + 1$.

3.5 Fluctuating environment: FE regime

We now consider the FE regime in the fluctuating environment. We first compute the Macroscopic Regime approximation (S63),

$$\Delta \dot{W}_M^* = \sum_y \frac{\psi_{x_y^* y}^M}{n} - \langle \pi, \psi^M \rangle = k\Theta, \quad (\text{S99})$$

where we used $\pi_y = 1/n$, the expression (S90) for d , and the fact that $x_y^* = y$. We also used that $\langle \pi, \psi^M \rangle = 0$, which is easy to show given the symmetries of (S90). Then, using all of the above, plus the definition of the effective transition rates \mathcal{R}_{xy} in (S65),

$$\begin{aligned} \mathcal{R}_{xy} &= \frac{1}{n} \sum_{y'} R_{(x,y)(x_y^*=y',y')} = \frac{1}{n} \sum_{y'} (\delta_{yy'} R_{x_y^* y'}^X + \delta_{xy'} R_{y' y}^Y) = \frac{1}{n} (R_{xy}^X + R_{yx}^Y) \\ &= \begin{cases} (k+q)/n & y = x + 1 \bmod n \\ -[2k + (2+\epsilon)q]/n & x = y \\ [k + (1+\epsilon)q]/n & x = y + 1 \bmod n \\ 0 & \text{otherwise} \end{cases}. \end{aligned} \quad (\text{S100})$$

At this point, we apply (S69) to obtain:

$$p_{x \neq x_y^* | y}^* = \begin{cases} \frac{k+q}{k\beta\Theta - 2k - (2+\epsilon)q} & y = x + 1 \bmod n \\ \frac{k + (1+\epsilon)q}{2k\beta\Theta - 2k - (2+\epsilon)q} & x = y + 1 \bmod n \\ 0 & \text{otherwise} \end{cases}. \quad (\text{S101})$$

while, for $x = x_y^* = y$,

$$p_{x=y|y}^* = 1 - \sum_{x \neq y} p_{x|y}^* = 1 - \frac{k+q}{k\beta\Theta - 2k - (2+\epsilon)q} - \frac{k + (1+\epsilon)q}{2k\beta\Theta - 2k - (2+\epsilon)q}. \quad (\text{S102})$$

Joint distributions are simply obtained by $p_{xy}^* = p_y^* p_{x|y}^* = \pi_y p_{x|y}^* = p_{x|y}^*/n$. On the other hand, to solve for the maximum power gain in the FE regime, we use (S70) to obtain:

$$\begin{aligned} \Delta \dot{W}_{\text{FE}}^* &= \Delta \dot{W}_{\text{M}}^* + \sum_{\substack{y \\ x=y+1 \bmod n}} \mathcal{R}_{xy} \left(\ln p_{x|y}^* - 1 \right) + \sum_{\substack{x \\ y=x+1 \bmod n}} \mathcal{R}_{xy} \left(\ln p_{x|y}^* - 1 \right) \\ &= \Delta \dot{W}_{\text{M}}^* + n \mathcal{R}_{(y+1)y} \left(\ln p_{x=y+1|y}^* - 1 \right) + n \mathcal{R}_{(y-1)y} \left(\ln p_{x=y-1|y}^* - 1 \right). \end{aligned} \quad (\text{S103})$$




Figure S5. We consider a unicyclic system with transitions across the n states with symmetric back-and-forth rates k . A single state, labelled by $i = 1$, leads to a flux of work input with power θ .

4. Alternative unicyclic model

In this section, we consider an alternative unicyclic model (see Fig. S5), different from the one studied in Sec. 3. We restrict our analysis to the case of a static environment, using techniques developed in Sec. 1.

In this model, the baseline dynamics are once again unicyclic with symmetric rates, as in Sec. 3.2. However, the transmission of work to the system from the environment operates differently. In this model, power transfer is not coupled to transitions. Instead, a single state leads to an internal flux of work from the environment; without loss of generality, we choose this state to be $i = 1$. In other words, the symmetry of the ring is broken by the presence of a highly energetic state ($i = 1$), however there is no preferential direction for the probability current. The power output when the system is in state 1 is given by the parameter θ , which carries the same units as \dot{W} . The baseline steady state is uniform, $\pi_i = 1/n$. We assume that $q = 0$ and $s = 0$, so the payoff vector is given by

$$\psi^{\text{LR}} = \psi^{\text{M}} = (\theta, 0, \dots, 0). \quad (\text{S104})$$

This section provides another illustration of how our theoretical framework unfolds in different scenarios. Note that this model is not discussed in the main text, and only appears here in the SM. Below we consider both the LR and FE regimes.




Figure S6. The optimal Δp^* (the deviation of the optimal distribution from the uniform baseline steady state) from (S108). Note that there optimal distribution does not lead to a cyclic current.

4.1 Static environment: LR Regime

Using the eigensystem discussed in discussed in Sec. 3.1, the expression for LR maximum power gain is Let us begin by computing the expression for maximum power gain in the LR regime, recall,

$$\Delta \dot{W}_{\text{LR}}^* = \beta \sum_{a=2}^n \frac{\Omega_a^2}{-\lambda_a}, \quad (\text{S105})$$

with $\Omega_a = \langle v, m_{(a)} \rangle$, where the vectors $m_{(a)}$ correspond to the eigensystem discussed in discussed in Sec. 3.1, i.e. in (S73). On the other hand, here $v = \frac{\beta\theta}{2\sqrt{n}}(1, 0, \dots, 0)$. The latter is obtained by construction using (S4) and (S5). Thus,

$$\Omega_a = \frac{\theta}{2n}.$$

For large even n , it is possible to show that [44]

$$\sum_{a=2}^{n/2} \frac{1}{1 - \cos\left(\frac{2\pi(a-1)}{n}\right)} \simeq \frac{n^2}{12} + \mathcal{O}(1), \quad (\text{S106})$$

which allows us to write

$$\Delta \dot{W}_{\text{LR}}^* \simeq \frac{\beta}{12k} \left(\frac{\theta}{2}\right)^2. \quad (\text{S107})$$

If we wanted to approximate the result above for a large odd n , we would need to add a term proportional to $1/4n^2$, which is of second order, hence the leading order is still captured by expression (S107).

We can also study the optimal distribution using (S13),

$$\Delta p^* = \beta \sum_{a=2}^n \frac{\Omega_a^\dagger D m_{(a)}}{-\lambda_a} = \frac{\beta\theta}{4kn^2} \sum_{a=2}^n \frac{u_{(a)}}{1 - \cos\left(\frac{2\pi(a-1)}{n}\right)} \quad (\text{S108})$$

with $u_{(a)} := (1, w_a, w_a^2, \dots, w_a^{n-1})$. This is computed numerically and its behavior is shown in Fig. S4(b). We leave this computation as an exercise to the reader.

Expression (S15) suggests sufficiency conditions for when the optimum will be in the LR regime. Note that, in this case, $\max_{a>1} \Omega_a = \theta/2n$ while

$$\min_{a>1} -\lambda_a = 4k \min_{a>1} \sin^2 \left(\frac{\pi(a-1)}{n} \right) = 4k \sin^2 \left(\frac{\pi}{n} \right).$$

Then, using $\|D\| = 1/\sqrt{n}$, condition (S15) will read as:

$$\frac{\theta}{2n} \ll \frac{4k\sqrt{n}}{\beta(n-1)} \sin^2 \left(\frac{\pi}{n} \right) \Leftrightarrow \theta \ll \frac{8\sqrt{nk}}{\beta} \left(\frac{n}{n-1} \right) \sin^2 \left(\frac{\pi}{n} \right). \quad (\text{S109})$$

For sufficiently large n , we can approximate $\sin^2(\frac{\pi}{n}) \approx (\frac{\pi}{n})^2$, so the LR regime is valid when

$$\theta \ll \frac{8\pi^2 k}{n^{3/2} \beta}. \quad (\text{S110})$$

Expression (S15), however, is not necessarily a tight bound. That is, the result here obtained is a sufficient condition for the LR regime to be valid, but not always a necessary one.

4.1 FE Regime

We now analyze the FE regime of model shown in Fig. S5(b). Following a similar procedure as above, we have $i^* = 1$, $\Delta\psi_2^M = \Delta\psi_n^M = \theta$, $R_{21} = R_{n1} = k$, $R_{11} = -2k$ and $R_{j1} = 0$ for all $j \neq 2, n$. Thus, for $a > 1$, we obtain:

$$\begin{aligned} p_a &= \frac{R_{a1}}{\Delta\psi_a^M + R_{11}} = \frac{1}{\beta\theta/k - 2} (\delta_{2a} + \delta_{na}) \\ p_1^* &= 1 - \sum_{a=2}^n p_a^* = \frac{\beta\theta/2k - 2}{\beta\theta/2k - 1}. \end{aligned} \quad (\text{S111})$$

Note that, for this example, $\langle \pi, \psi^M \rangle = \theta/n$. Therefore,

$$\begin{aligned} \Delta\dot{W}_{\text{FE}}^* &= \psi_1^M - \langle \pi, \psi^M \rangle + \beta^{-1} \sum_{a=2}^n R_{a1} (\ln p_a^* - 1) \\ &= \theta \left(\frac{n-1}{n} \right) - 2k\beta^{-1} [1 + \ln(\beta\theta k^{-1} - 2)]. \end{aligned} \quad (\text{S112})$$

5. Glucose Transporter

We illustrate our approach on a biologically-inspired example of a membrane-bound transporter. Although this example is highly idealized and simplified, its goal is to show in general terms how our methods can be used to study biological questions.

We consider a system X and environment Y , where X is a membrane-bound transporter protein (e.g., in an *E. coli*) and Y indicates the absence/presence of extracellular glucose. Specifically, X is the coarse-grained state of n_{AA} amino acids that make up a transporter protein. We use an order-of-magnitude estimate of $n_{AA} = 500$ [45]. X can be in one of two macrostates: $x = 0$ when the amino acids are free-floating and in equilibrium in the intracellular medium, and $x = 1$ when the amino acids are polymerized into a protein. The unpolymerized state has a much larger equilibrium entropy, due to the sequence and spatial degrees of freedom of the free-floating amino acids. For each free-floating amino acid, the entropy increase due to sequence freedom is $\approx \ln 20$ while the entropy increase due to spatial freedom can be roughly approximated as ≈ 20 nats, as the log ratio of the volume of an *E. coli*, $\approx 10^{-18} \text{m}^3$, and a protein, $\approx 10^{-27} \text{m}^3$ [46, 47]. Combining gives internal entropy of $s_{xy} = \delta_{x,0} n_{AA} (\ln 20 + 20)$ [7]. The environment Y can also be in one of two states: $y = 0$ when there is no glucose in the environment, and $y = 1$ when there is glucose. When both glucose and transporter are present, $x = 1$ and $y = 1$, transport takes place at an uptake rate γ glucose/sec. This is represented by a power output/input observables $w_{xy} = d_{xy} = \delta_{x,1} \delta_{y,1} \gamma \eta$, where η is joules of energy acquired for each glucose molecule. We suppose that the cell is growing in anaerobic conditions and metabolizing 2 ATP per glucose via fermentation, so $\eta = 2 \times 20 k_B T$ ($k_B T \equiv 1/\beta$ is taken to be room temperature). We use a typical value of $n_{AA} = 500$ amino acids per transporter.

The baseline processes include two types of transitions. First, X spontaneously transitions from the polymerized to the unpolymerized macrostate due to degradation and dilution, $R_{(0,y)(1,y)} = k_d$ (we assume that this transition is effectively irreversible). We choose $k_d = 1/1200$ to represent dilution when the cell is growing at a doubling time of 20 minutes. Second, the presence of glucose in the environment randomly fluctuates, $R_{(x,1)(x,0)} = R_{(x,0)(x,1)} = k_f$. We take $k_f = k_d$, so that the extracellular concentration fluctuates once per cell lifetime. One can verify that $\pi_x = \delta_{x,0}$ in the baseline steady state due to the irreversible degradation transition, thus power output vanishes in the baseline steady state.

The control process involves the synthesis of transporters, which is regulated by the presence of extracellular glucose, $R_{(1,1)(0,1)} = k_s$ (we assume this reaction is effectively irreversible). This synthesis releases $\approx n_{AA} 80 kT$ of heat, due to the hydrolysis of 4 ATP molecules per amino acid and $20 kT$ of heat per ATP. We set $k_s = 1/10$ representing ≈ 10 s timescale necessary for transcribing and translating a protein [48]. For simplicity, we ignore

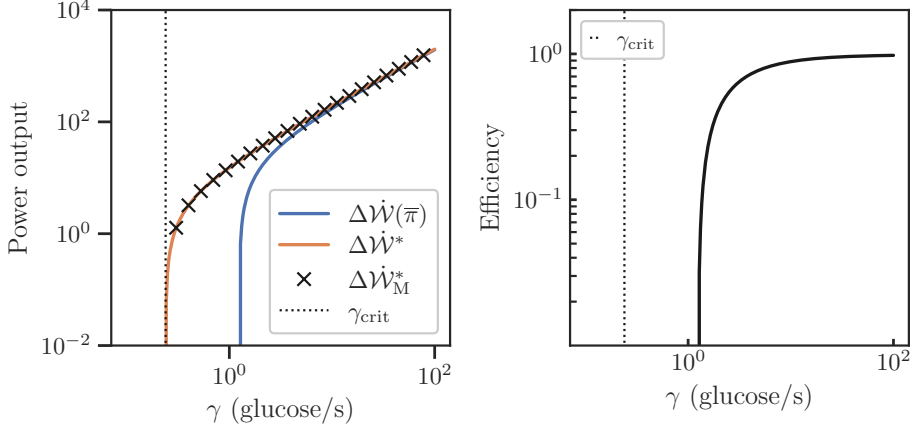


Figure S7. Model of transport of glucose. Left: In blue, the controlled output power difference for real parameter values in blue; in orange, the maximum power increase obtained from numerical optimization (S61) and, in black crosses, the Macroscopic Regime value (see Section 2.3). Inset shows the efficiency of the process the the same values of γ . Right: Shows the efficiency of the system measured against the maximum power increase, $\Delta\dot{W}(\bar{\pi})/\Delta\dot{W}^*$ as a function of γ .

other thermodynamic costs, such as those associated with active transport, regulation, transcription, etc.

In Fig. S7, we show results as a function of the uptake rate parameter γ (glucose/sec), whose changes can reflect varying concentration of glucose in macrostate $Y = 1$ as well as different levels of transporter efficiency. We explore a range from $\gamma = 0$ to the physiologically maximal $\gamma = 100$ [48]. The left plot shows the power increase achieved by the actual control processes, $\Delta\dot{W}(\bar{\pi})$. It also shows the optimal power increase $\Delta\dot{W}^*$ as determined by numerical optimization (S61), assuming the locality constraint that control processes cannot directly modify the state of Y (as in (S61)). We also plot the “best response” Macroscopic Regime prediction (see Section 2.3). There is a critical level of uptake,

$$\gamma_{\text{crit}} = k_d n_{AA} (\ln 20 + 20) / \eta \beta \approx 0.24, \quad (\text{S113})$$

for which benefit of the transporter outweighs its minimal cost.

In this model, the free energy scales are such that the nonlinear term can be ignored in our optimization problem, and the optimal distribution is close to the “best response” Macroscopic Regime distribution described in Section 2.3, where $x = y$ if $\gamma > \gamma_{\text{crit}}$ and $x = 0$ otherwise. Thus, the Macroscopic Regime prediction, $\Delta\dot{W}_M^* = \frac{\eta}{2} \max(\gamma - \gamma_{\text{crit}}, 0)$, is highly accurate. The inset on the left shows the efficiency calculated as $\Delta\dot{W}(\bar{\pi})/\Delta\dot{W}^*$. Efficiency is low at low uptake rates, due to the inefficiency of translation, but quickly increases to ≈ 1 as γ grows.

We may also consider accuracy of the Macroscopic approximation by evaluating the parameter α , as defined in (S19). For simplicity, consider the limit where the environment changes very fast, $k_f \rightarrow \infty$. In this case, the model is reduced to a two-state system representing the cell only ($X \in \{0, 1\}$), which is faced with a random environment where glucose is present half the time. The numerator of (S19) (maximum escape rate) is $k_d = 1/1200$, while the denominator (macroscopic power increase in units of entropy) is $\beta\Delta\dot{W}_M^* = \beta\psi_{i^*} = \gamma 40/2 - 500(\ln 20 + 20)/1200$ (the factor of $1/2$ enters because the glucose is only present half the time). At a realistic uptake value of $\gamma \approx 10$ glucose/sec, $\Delta\dot{W}_M^* \approx \gamma 40/2 = 200$ and $\alpha = (1/1200)/200 \approx 4 \times 10^{-6}$, as mentioned in the main text.

6. Quadratic optimization theorem

In this section, we provide a useful theorem for solving the quadratic optimization problems that occurs in our analysis of the linear response (LR) regime. It uses standard techniques from linear algebra.

Theorem 1. *Consider the maximization problem*

$$V = \max_{z \in \mathbb{R}^n, u^T z = 0, Gz = \mathbf{0}} z^T M z + 2z^T v, \quad (\text{S114})$$

where $M \in \mathbb{R}^{n \times n}$ is a negative semidefinite matrix with a single null eigenvector $u \in \mathbb{R}^n$, the matrix $G \in \mathbb{R}^{n \times k}$ encodes some set of k constraints, and $v \in \mathbb{R}^n$ is any vector. For convenience, define

$$\Pi := I - \begin{bmatrix} G \\ u^T \end{bmatrix}^+ \begin{bmatrix} G \\ u^T \end{bmatrix} \quad \Lambda := \left[I - \left((-M)^{1/2} \Pi \right)^+ (-M)^{1/2} \right] M^+. \quad (\text{S115})$$

The solution is given by:

$$\text{in the general case: } V = -v^T M^+ v + v^T \Lambda M \Lambda v \quad z^* = (\Lambda - M^+) v \quad (\text{S116})$$

$$\text{if } \Pi \text{ and } M \text{ commute: } V = -v^T \Pi M^+ \Pi v \quad z^* = -M^+ \Pi v \quad (\text{S117})$$

$$\text{if } G = 0: V = -v^T M^+ v \quad z^* = -M^+ v \quad (\text{S118})$$

Proof. We first consider the most general case, (S116). Define the block matrix $H := \begin{bmatrix} G \\ u^T \end{bmatrix}$, so that $\Pi = I - H^+ H$ is the orthogonal projector onto the kernel of H . Note that for any vector z , the vector Πz satisfies the constraints in (S114), $G \Pi z = 0$ and $u^T \Pi z = 0$, so for any allowed z we can assume $z = \Pi z$. We then eliminate the constraints in our optimization problem by replacing z by Πz in the objective,

$$V = \max_z z^T \Pi M \Pi z + 2(\Pi z)^T v. \quad (\text{S119})$$

Next, recall that M is a symmetric matrix with a single null eigenvector u . The orthogonal projector onto the range of M is

$$M M^+ = I - u u^T. \quad (\text{S120})$$

The range of M includes the kernel of H , which means

$$M M^+ \Pi = \Pi - u u^T (I - H^+ H) = \Pi, \quad (\text{S121})$$

as follows from $u u^T H^+ H = u([0 \ 1] H) H^+ H = u[0 \ 1] H = u u^T$. Using $M M^+ \Pi = \Pi$ from (S121), rewrite the objective in (S119) as

$$z^T \Pi M \Pi z + 2(\Pi z)^T v = z^T \Pi M \Pi z + 2z^T \Pi M M^+ v$$

$$\begin{aligned}
&= (\Pi z + M^+ v)^T M (\Pi z + M^+ v) - v^T M^+ v \\
&= -(\Pi z - \bar{M}^+ v)^T \bar{M} (\Pi z - \bar{M}^+ v) - v^T M^+ v \\
&= -v^T M^+ v - \left\| \bar{M}^{1/2} \Pi z - \bar{M}^{1/2} \bar{M}^+ v \right\|^2.
\end{aligned} \tag{S122}$$

where for convenience we defined the positive semidefinite matrix $\bar{M} = -M$. Note that for any A, x, y ,

$$\begin{aligned}
\|Ax - y\|^2 &= \|AA^+y - y + Ax - AA^+y\|^2 \\
&= \|(AA^+ - I)y\|^2 + \|A(x - A^+y)\|^2 \geq \|(AA^+ - I)y\|^2,
\end{aligned} \tag{S123}$$

where in the first line we added and subtracted AA^+y , and in the second line used the orthogonality $(AA^+ - I)A = 0$.

This lower bound is achieved by any vector of the form,

$$x^* = A^+y + (I - A^+A)g \quad g \in \mathbb{R}^n. \tag{S124}$$

Plugging (S123) into (S122) gives

$$\begin{aligned}
\left\| \bar{M}^{1/2} \Pi z - \bar{M}^{1/2} \bar{M}^+ v \right\|^2 &\geq \left\| [\bar{M}^{1/2} \Pi (\bar{M}^{1/2} \Pi)^+ - I] \bar{M}^{1/2} \bar{M}^+ v \right\|^2 \\
&= \left\| \bar{M}^{1/2} [(\bar{M}^{1/2} \Pi)^+ \bar{M}^{1/2} - I] \bar{M}^+ v \right\|^2 \\
&= \left\| \bar{M}^{1/2} [I - (\bar{M}^{1/2} \Pi)^+ \bar{M}^{1/2}] \bar{M}^+ v \right\|^2
\end{aligned} \tag{S125}$$

$$= v^T \Lambda \bar{M} \Lambda v = -v^T \Lambda M \Lambda v, \tag{S126}$$

where Λ is defined as in (S115). In the second line, we simplified slightly by using that

$$\Pi(A\Pi)^+ = (A\Pi)^+ \tag{S127}$$

for any A whenever Π is an orthogonal projector (see Appendix A in [41]).

The value of V in (S116) follows by combining (S119), (S122), and (S126). To find z^* , using (S124) we know that the minimum is achieved by

$$z^* = (\bar{M}^{1/2} \Pi)^+ \bar{M}^{1/2} \bar{M}^+ v + (I - (\bar{M}^{1/2} \Pi)^+ \bar{M}^{1/2} \Pi)g. \tag{S128}$$

Recall, however, that any allowed z in the original problem (S114) satisfies $z = \Pi z$. Multiplying both sides of (S128) by Π gives

$$\begin{aligned}
z^* &= \Pi(\bar{M}^{1/2} \Pi)^+ \bar{M}^{1/2} \bar{M}^+ v + \Pi(I - (\bar{M}^{1/2} \Pi)^+ \bar{M}^{1/2} \Pi)g \\
&= (\bar{M}^{1/2} \Pi)^+ \bar{M}^{1/2} \bar{M}^+ v
\end{aligned}$$

$$\equiv -((-M)^{1/2}\Pi)^+(-M)^{1/2}M^+v.$$

In the second line, we used (S127) in the first term, and cancelled the second term by using that $\Pi = (\bar{M}^{1/2})^+\bar{M}^{1/2}\Pi$ and a bit of rearranging.

Next, we consider the special case (S117). Note that if Π commutes with M , then

$$(\bar{M}^{1/2}\Pi)^+\bar{M}^{1/2} = \Pi(\bar{M}^{1/2})^+\bar{M}^{1/2} = \Pi\bar{M}^+\bar{M} = \bar{M}\bar{M}^+\Pi = \Pi,$$

where we used (S121). Plugging into (S115) gives $\Lambda = (I - \Pi)M^+$, and plugging into (S116) gives $z^* = -\Pi M^+v$, as in (S117). We also have

$$\begin{aligned} \Lambda M \Lambda &= (I - \Pi)M^+ M M^+(I - \Pi) \\ &= (I - \Pi)M^+(I - \Pi) \\ &= M^+ - 2M^+\Pi + \Pi M^+\Pi \\ &= M^+ - \Pi M^+\Pi, \end{aligned}$$

where we used that $\Pi M^+\Pi = M^+\Pi\Pi = M^+\Pi$. Plugging into (S116) and rearranging gives $V = -v^T\Pi M^+\Pi v$, as in (S117).

Finally, we consider the special case (S118) where $G = 0$. Then it can be verified from (S115) and (S120) that $\Pi = M M^+$. Plugging into (S117) and simplifying gives (S118). \square

OPEN ACCESS

PAPER



Fixation and fluctuations in two-species cooperation

RECEIVED

19 September 2021

REVISED

1 February 2022

ACCEPTED FOR PUBLICATION

8 February 2022

PUBLISHED

28 February 2022

Jordi Piñero¹ , S Redner^{2,*} and Ricard Solé^{3,*} ¹ ICREA-Complex Systems Lab, Universitat Pompeu Fabra, 08003 Barcelona, Spain² Institut de Biologia Evolutiva (CSIC-UPF), Psg Maritim Barceloneta, 37, 08003 Barcelona, Spain³ Santa Fe Institute, 1399 Hyde Park Road, Santa Fe NM 87501, United States of America

* Authors to whom any correspondence should be addressed.

E-mail: jordi.pinero@upf.edu, redner@santafe.edu and ricard.sole@upf.edu

Keywords: fixation, extinction, first passage, stochastic processes, cooperation

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



Abstract

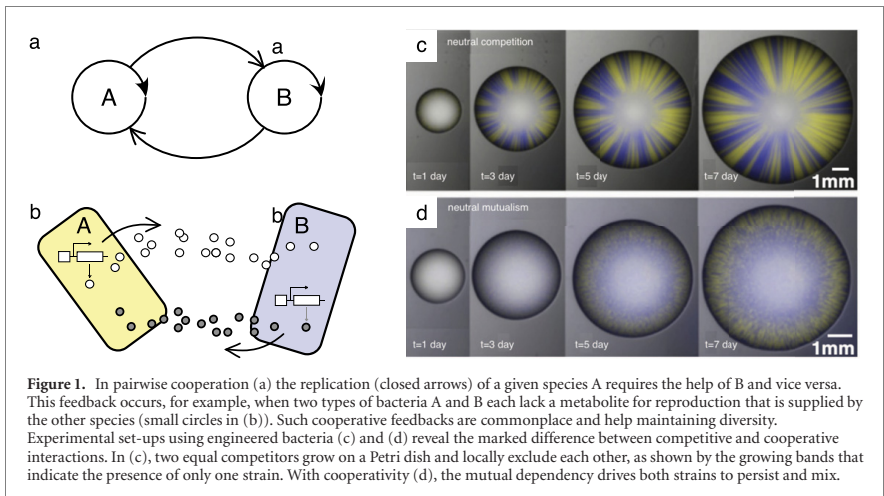
Cooperative interactions pervade in a broad range of many-body populations, such as ecological communities, social organizations, and economic webs. We investigate the dynamics of a population of two equivalent species A and B that are driven by cooperative and symmetric interactions between these species. For an isolated population, we determine the probability to reach fixation, where only one species remains, as a function of the initial concentrations of the two species, as well as the time to reach fixation. The latter scales exponentially with the population size. When members of each species migrate into the population at rate λ and replace a randomly selected individual, surprisingly rich dynamics ensues. Ostensibly, the population reaches a steady state, but the steady-state population distribution undergoes a unimodal to trimodal transition as the migration rate decreases below a critical value λ_c . In the low-migration regime, $\lambda < \lambda_c$, the steady state is not truly steady, but instead strongly fluctuates between near-fixation states, where the population consists of mostly A's or of mostly B's. The characteristic time scale of these fluctuations diverges as λ^{-1} . Thus in spite of the cooperative interaction, a typical snapshot of the population will contain almost all A's or almost all B's.

1. Introduction

Competitive interactions have played a prominent role in the literature of ecological and evolutionary dynamics, as well as in economics and sociology [1–3]. Resource limitations and their impact in defining the outcome of competition among species has shaped a large part of evolutionary thinking. A counterpoint to competition is cooperativity in which there are positive interactions and feedback loops between species. These mechanisms have received increasing attention recently [4, 5]. In fact, it is the presence of cooperative interactions, where positive reciprocal exchanges are at work, that appear to drive innovations in evolution and also maintain biodiversity in Nature [5].

Cooperation, or mutualism, has been part of mathematical models of populations since the formulation of Lotka–Volterra equations [2] and a variety of statistical physics models of human cooperation [6, 7]. In its most abstract form, two species (for example, A and B in figure 1(a)) ‘help’ each other by means of a mutual positive interaction; in some cases, both partners completely rely on one another for survival. This feature underlies the two-species system in figure 1(b), where a given species requires the other to replicate because each species needs a molecule that is produced by the partner species.

Recent experimental studies have shown that such cooperative populations can in fact be engineered. By following the scheme in figure 1(b), it is possible to create a completely symmetric dependence and make these mixed populations grow on a Petri dish [8–10]. Figure 1(c) shows the outcome of symmetric competition when each strain is marked with a different fluorescent protein: each strain locally out competes the other, thereby generating stripes of segregated domains. The cooperatively interacting population, on the other hand, constrains both species to remain in proximity, leading to a well-mixed population (figure 1(d)). These simple engineered, or synthetic, bacterial populations, which can be tuned so that they become virtually



symmetric, allow one to explore the fundamental dynamical features of interacting living consortia and also study the impact of stochasticity [11–13].

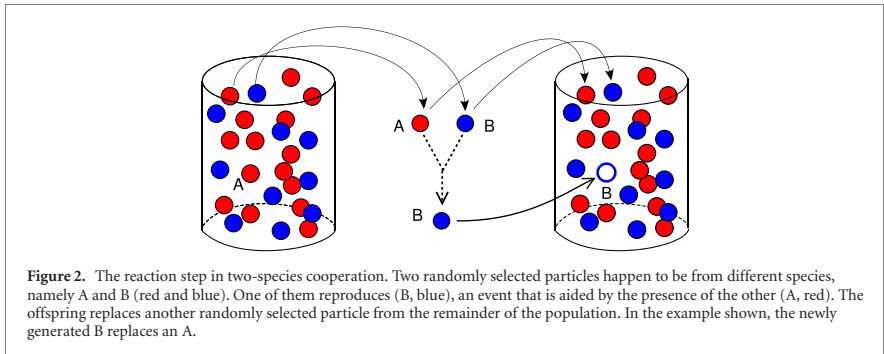
In this paper, we present an analytic approach to understand the role of stochasticity in a simple two-species stochastic model of cooperation. This model represents a special case of evolutionary game dynamics [14–17], with a specific and particularly simple form of the payoff matrix. We emphasize that we are not treating a general ecological model, but rather a simplified system in which only cooperative interactions occur. This scenario appears to be more relevant for the microbiome [18–20]. Current advances in microbial ecology involve experimental setups with a small number of interacting species [21, 22]. In this model, we treat both closed and open populations, in which there is either no migration or a finite rate of migration into the population, respectively. We define microscopic rules that incorporate both cooperativity, in which each species helps the other, as well as neutrality, in which neither species is preferred over the other. We first determine the steady state of the population in the absence of stochastic fluctuations. For a finite population, we then incorporate stochasticity and determine the time until fixation is reached for the situation where no migration can occur. When migration is allowed (with compensatory removal), the population now reaches a steady state; however, the character of this steady state dramatically changes as function of the migration rate. For large migration rate, both species are present in roughly equal abundances. However, for a sufficiently small migration rate, the population strongly fluctuates between consisting of nearly all A or all B. Thus a *typical* realization of the population has a completely different character than the *average* population. This change in behavior is mirrored by a bimodal to trimodal transition in the shape of the steady-state probability distribution of species abundances.

In section 2, we outline basic features of our two-species cooperation model in the absence of migration. We solve the model in the mean-field approximation and then include the role of stochasticity due to the finiteness of the population. For a finite population, we determine the fixation probability as a function of the initial population composition and the time until fixation, where only a single species remains. In section 5, we incorporate migratory inflow, with compensatory removal, so that the population size remains fixed and reaches a steady state. We discuss basic features of this steady state, including the intriguing feature that the species abundances can exhibit huge fluctuations, even though time-averaged properties are constant. We give some concluding remarks in section 6.

2. Two-species cooperation

We investigate a finite population of N particles, with n of species A and $N - n$ of species B. The population undergoes repeated reaction events in which each event consists of the following steps (figure 2):

- Pick a random pair of particles.
- If the pair is AB, one member of this pair reproduces; if the pair is AA or BB, nothing happens.
- The newly reproduced offspring replaces one randomly selected particle in the remainder of the population.



Thus interactions between members of different species are cooperative in nature, while members of the same species are non interacting. The replacement step (iii) ensures that the total population remains fixed. The lack of interactions between AA and BB pairs follows from the assumption of strict mutualism, i.e., replication occurs if and only if both species are present (as sketched in figure 2). After each update, time is incremented by $\frac{1}{N}$. This time increment corresponds to each particle undergoing, on average, steps (i)–(iii) in a single time unit. While this dynamics manifestly conserves the total number of particles, the composition of the population can change. When the population consists entirely of a single species—either all A’s or all B’s—there is no further dynamics and the population fixates.

If an A reproduces in a single interaction, then with probability $1 - \frac{n}{N} \equiv 1 - x$, the A offspring replaces a B and $n \rightarrow n + 1$. Conversely, with probability $x = \frac{n}{N}$, the A offspring replaces an existing A so that n does not change. The probability a_n at which $n \rightarrow n + 1$ therefore is

$$a_n = 2x(1 - x) \times \frac{1}{2} \times (1 - x) = x(1 - x)^2. \quad (1a)$$

Throughout, we use the variables n and $x = \frac{n}{N}$ interchangeably. In (1a), the factor $2x(1 - x)$ gives the probability a randomly selected pair is AB, the factor $\frac{1}{2}$ gives the probability that the A in this pair reproduces, and the factor $1 - x$ gives the probability that the A offspring replaces a B. By the same reasoning, when a B reproduces in an interaction, the probability at which $n \rightarrow n - 1$ is

$$b_n = 2x(1 - x) \times \frac{1}{2} \times x = x^2(1 - x). \quad (1b)$$

Here the trailing factor x in (1b) accounts for the probability that the B offspring replaces an A. Finally, the probability that the number of A’s and B’s does not change is given by

$$x^2 + (1 - x)^2 + 2x(1 - x) \left[\frac{1}{2}x + \frac{1}{2}(1 - x) \right] = 1 - x(1 - x) = 1 - a_n - b_n. \quad (1c)$$

The terms $x^2 + (1 - x)^2$ give the probability to pick either an AA or BB pair, for which no change in n occurs. For the last term, $2x(1 - x)$ is again the probability of picking an AB pair, while the factor in the square brackets is the probability that the offspring (either A or B with probability $\frac{1}{2}$) replaces its own kind so that n does not change.

3. Mean-field approaches

3.1. Rate equation

Using the probabilities in equation (1), the rate equation for the average number of A’s is

$$\dot{n} = N(a_n - b_n) = Nx(1 - x)(1 - 2x), \quad (2a)$$

or equivalently,

$$\dot{x} = x(1 - x)(1 - 2x). \tag{2b}$$

To keep the notation simple, n and x refer to average values in this section; that is, we do not write angle brackets. This rate equation has a stable fixed point at $x = \frac{1}{2}$ and unstable fixed points at $x = 0$ and $x = 1$. The stability of the fixed point at $x = \frac{1}{2}$ arises because the transition probabilities (1) tend to reduce population imbalances. Thus the steady state in this continuum description is a static population that consists of equal densities of A's and B's. That is, cooperativity promotes diversity in the mean-field description.

The solution to the rate equation (2b) may be straightforwardly obtained by first performing a partial fraction decomposition:

$$dt = \frac{dx}{x(1-x)(1-2x)} = dx \left(\frac{1}{x} - \frac{1}{1-x} + \frac{4}{1-2x} \right),$$

from which

$$t = \int_{x_0}^x dy \left(\frac{1}{y} - \frac{1}{1-y} + \frac{4}{1-2y} \right) = -4 \ln \frac{x(1-x)(1-2x)}{x_0(1-x_0)(1-2x_0)}.$$

We then obtain $x(t)$ by solving the resulting cubic equation. For $t \rightarrow \infty$, the limiting behavior is

$$x(t) \simeq \frac{1}{2} - 2x_0(1-x_0)(1-2x_0)e^{-t/4}, \tag{3}$$

so that the stable fixed point $x^* = \frac{1}{2}$ is approached exponentially quickly in time.

3.2. Master equation and its moments

In the stochastic dynamics where n is a discrete variable, the true fixed points are at $x = 0$ and $x = 1$. Even though the fixed point at $x^* = \frac{1}{2}$ is stable in the continuum limit, stochastic fluctuations allow the population to explore the full state space and eventually get trapped at either $x = 0$ or $x = 1$. This behavior is analogous to the extinction phenomena that arise, for example, in the logistic birth-death process, $A \rightarrow 2A$ and $2A \rightarrow 0$, as well as other reactions of this genre [23–26]. In these reactions, the rate equation predicts a steady population, N_s , which is determined by the balance between the birth and death rates. However, in the true stochastic dynamics, the population fluctuates around N_s , which actually is the *quasi* steady-state value. Ultimately, a sufficiently large fluctuation occurs that leads to extinction, from which there can be no escape, with an extinction time that scales exponentially in N_s [23–27].

To understand the stochastic dynamics for two-species cooperation, we study $P_n(t)$, the probability that the population consists of n A's and $(N - n)$ B's at time t . The time dependence of this probability distribution is given by

$$\dot{P}_n = N [a_{n-1}P_{n-1} + b_{n+1}P_{n+1} - (a_n + b_n)P_n]. \tag{4}$$

When the number of particles N is small, the set of equation (4) can be straightforwardly solved. For the initial condition of equal numbers of A's and B's, both $P_0(t)$ and $P_N(t)$ approach $\frac{1}{2}$ as $t \rightarrow \infty$, while all the other $P_n(t)$ vanish exponentially quickly in time. This direct approach quickly becomes tedious as N increases, however, and to gain insight into the long-time dynamics for general N , it is useful to study low-order moments of P_n . From equation (4) and using a_n and b_n from equation (1), the first moment obeys

$$\begin{aligned} \langle \dot{x} \rangle &= \frac{1}{N} \sum_n n \dot{P}_n = \sum_{1 \leq n \leq N} \{n a_{n-1} P_{n-1} + n b_{n+1} P_{n+1} - n(a_n + b_n) P_n\} \\ &= \sum_{1 \leq n \leq N} \{(n+1)a_n P_n + (n-1)b_n P_n - n(a_n + b_n) P_n\} \\ &= \sum_{1 \leq n \leq N} (a_n - b_n) P_n = \langle x(1-x)(1-2x) \rangle. \end{aligned} \tag{5a}$$

Here we now explicitly write angle brackets to denote average values. Under the assumption of no correlations, that is, $\langle x^k \rangle = \langle x \rangle^k$, (5a) reproduces the rate equation (2b).

Similarly, the equation of motion for the second moment is

$$\begin{aligned}
 \langle \dot{x}^2 \rangle &= \frac{1}{N^2} \sum_n n^2 \dot{P}_n = \frac{1}{N} \sum_{1 \leq n \leq N} \{n^2 a_{n-1} P_{n-1} + n^2 b_{n+1} P_{n+1} - n^2 (a_n + b_n) P_n\} \\
 &= \frac{1}{N} \sum_{1 \leq n \leq N} \{(n+1)^2 a_n P_n + (n-1)^2 b_n P_n - n^2 (a_n + b_n) P_n\} \\
 &= \frac{1}{N} \langle x(1-x) \rangle + 2 \langle x^2(1-x)(1-2x) \rangle.
 \end{aligned}
 \tag{5b}$$

It is more convenient to express (5a) and (5b) in terms of $z \equiv 2x - 1$, which lies in the range $[-1, 1]$. Doing so, we obtain

$$\begin{aligned}
 \langle \dot{z} \rangle &= -\frac{1}{2} \langle z(1-z^2) \rangle \\
 \langle \dot{z}^2 \rangle &= \langle (1-z^2) \left(\frac{1}{N} - z^2 \right) \rangle,
 \end{aligned}
 \tag{6}$$

which are both symmetric about $z = 0$. If we make the assumption of no correlations, that is, $\langle z^k \rangle = \langle z \rangle^k$, then the first equation reproduces the result that $z = 0$ is a stable fixed point. The second equation then predicts that the width of the distribution initially grows and eventually ‘sticks’ at the value \sqrt{N} . To check this point, we numerically integrated equation (4) for small values of N . From the resulting solution, we find that the width of the probability distribution initially grows with time and later approaches a nearly fixed value that is proportional to \sqrt{N} . However, at very long times, there is slow leakage of the probability distribution to the true stochastic fixed points at $z = \pm 1$. Thus the probability distribution eventually approaches two delta-function peaks at these fixed points. This behavior cannot be captured by low-order moment equations, such as (6). Instead, we need to study the full stochastic dynamics; this is done in the following section.

4. Fixation probability and fixation time

We now turn to two quantities of primary interest in the stochastic dynamics, namely, (i) the fixation (or exit) probability E_n , and (ii) the fixation time T_n . The fixation probability E_n is defined as the probability that a population of size N that initially contains n particles of type A reaches the static fixation state of all A’s. We use the backward Kolmogorov equation [28, 29] to compute the fixation probability. In this approach, E_n satisfies the recursion

$$E_n = a_n E_{n+1} + b_n E_{n-1} + (1 - a_n - b_n) E_n. \tag{7}$$

Since the process renews itself after each event, we can express the fixation probability from the state that contains n A’s in terms of the appropriately weighted average of the fixation probabilities after a single step to the states $n - 1$, n , and $n + 1$. The weights are merely the hopping probabilities to these respective states. Equation (7) is subject to the boundary conditions $E_0 = 0$ and $E_N = 1$. The first condition corresponds to the impossibility of reaching a population of all A’s if the initial state contains no A’s, while the second condition corresponds to the initial state coinciding with the desired final state of all A’s.

The solution to (7) is (see appendix A for the calculational details)

$$E_n = \sum_{m=0}^{n-1} \left[\binom{N-1}{m} \right]^{-1} / \sum_{m=0}^{N-1} \left[\binom{N-1}{m} \right]^{-1}. \tag{8}$$

Neither of these sums has a closed form, but for $N \rightarrow \infty$ the denominator approaches 2 [30]. For $N \gg 1$, E_n and its continuum counterpart $E(x)$ (see also appendix A) are nearly independent of $x = \frac{n}{N}$ when x is not close to 0 or 1. Figure 3 shows this dependence of E_n on n . Also shown are the corresponding results from discrete simulations of the fixation process. Simulations are carried out by setting a finite size (N) array of particles with two possible states. Each iteration the particles interact following rules (ii) and (iii) with the interaction rates given by expressions (1a) and (1b). Time is updated by $\Delta t = [N(a_n + b_n)]^{-1}$ after each iteration.

The anti-sigmoidal shape of E_n arises from the underlying drift that tends to drive any initial population towards $x = \frac{1}{2}$. Eventually, a large and rare stochastic fluctuation causes the population to escape this effective potential well and reach fixation. This anti-sigmoidal shape also strongly contrasts with the Moran process [31], which is symmetric (neutral), but non-cooperative. Here an AB pair equiprobably converts to either AA or BB. As a result of this lack of cooperativity, the fixation probability in the strictly neutral Moran process is simply the linear function $E(x) = x$ [29, 31–33].

We now investigate the fixation time T_n , which is defined as the average time for the population of N particles to first reach either of the two fixation states, $n = 0$ or $n = N$, when the population initially contains

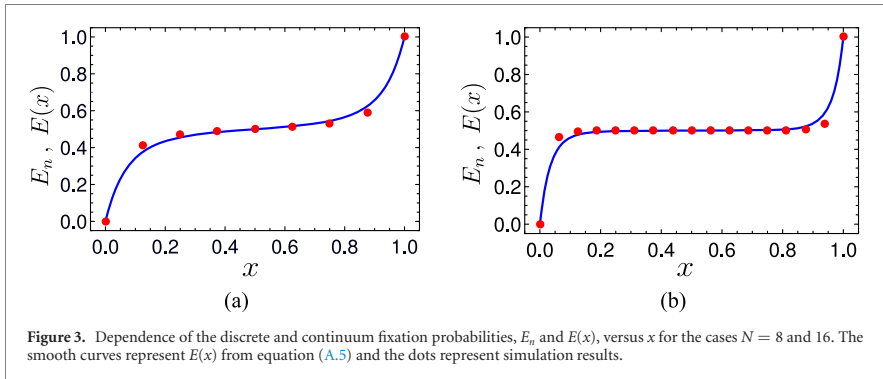


Figure 3. Dependence of the discrete and continuum fixation probabilities, E_n and $E(x)$, versus x for the cases $N = 8$ and 16 . The smooth curves represent $E(x)$ from equation (A.5) and the dots represent simulation results.

n A's. Within the same backward Kolmogorov framework as that used for the fixation probability, the fixation time satisfies [28, 29]

$$T_n = a_n T_{n+1} + b_n T_{n-1} + (1 - a_n - b_n) T_n + \delta t. \tag{9}$$

Again T_n may be expressed as the appropriately weighted average of the fixation time after a single hopping event to the states $n - 1$, n , and $n + 1$, plus the time $\delta t = \frac{1}{N}$ required for this single step. The latter corresponds to each particle being updated once, on average, in a single time unit. The equation for T_n is subject to the boundary conditions $T_0 = T_N = 0$; namely, if the population starts in a fixation state, the time to reach this state is 0. The result for the fixation time is (see appendix B)

$$T_n = E_n \sum_{m=1}^{N-1} Q_m - \sum_{m=1}^{n-1} Q_m, \tag{10}$$

where

$$Q_n \equiv \alpha_n + r_n \alpha_{n-1} + r_n r_{n-1} \alpha_{n-2} + \dots + r_n r_{n-1} \dots r_2 \alpha_1,$$

with $Q_0 = 0$, $r_n = b_n/a_n$, and $\alpha_n \equiv \delta t/a_n$.

It does not seem possible to reduce (10) to a compact form, but the main feature of this exact expression is that the fixation time scales exponentially in N and is nearly independent of n (or, equivalently, x), except for n close to 0 or to N (figure 4). The exponential dependence on N again arises because of the existence of an effective potential well, whose depth grows linearly with N , which draws the population toward the state $x = \frac{1}{2}$. The near independence of the fixation time on the initial condition is a consequence of the population being drawn toward the bottom of this potential well, where the concentrations of A and B are equal. As a result, the value of the fixation time for any initial value of x is close to the fixation time when the population starts from the bottom of the potential well at $x = \frac{1}{2}$.

It is possible, however, to obtain an analytical expression for the average fixation time by the WKB method [23–26]. The idea of this approach is that the probability distribution settles into a quasi-steady state that assumes an exponential large-deviation form. From equation (4), there is a slow leakage from this quasi-steady state to the fixation state whose rate, $\Gamma(N)$, is given by

$$\Gamma(N)\delta t = b_1 \tilde{P}_1 + a_{N-1} \tilde{P}_{N-1} = 2b_1 \tilde{P}_1, \tag{11}$$

i.e., the flux from states that are one step away from fixation to the fixation states. Here the tilde denotes the steady-state distribution and we also use the symmetry $n \leftrightarrow N - n$. We then identify the inverse of this leakage rate with the fixation time.

We obtain an approximate equation for the continuum probability distribution $\tilde{P}_n \rightarrow \tilde{P}(x)$ by setting the time derivative in the master equation (4) to zero, and writing $n \pm 1$ as $x \pm \delta x$ to give

$$a(x - \delta x)\tilde{P}(x - \delta x) + b(x + \delta x)\tilde{P}(x + \delta x) = [a(x) + b(x)]\tilde{P}(x). \tag{12}$$

We now assume that $\tilde{P}(x)$ has the exponential form $\tilde{P}(x) \sim e^{S(x)/\delta x} = e^{NS_0(x) + S_1(x) + \dots}$ and substitute this form into (12) to give (up to $O(1)$)

$$S_0(x) = \int^x dz \log \left[\frac{a(z)}{b(z)} \right], \quad S_1(x) = -\frac{1}{2} \log [a(x)b(x)]. \tag{13}$$

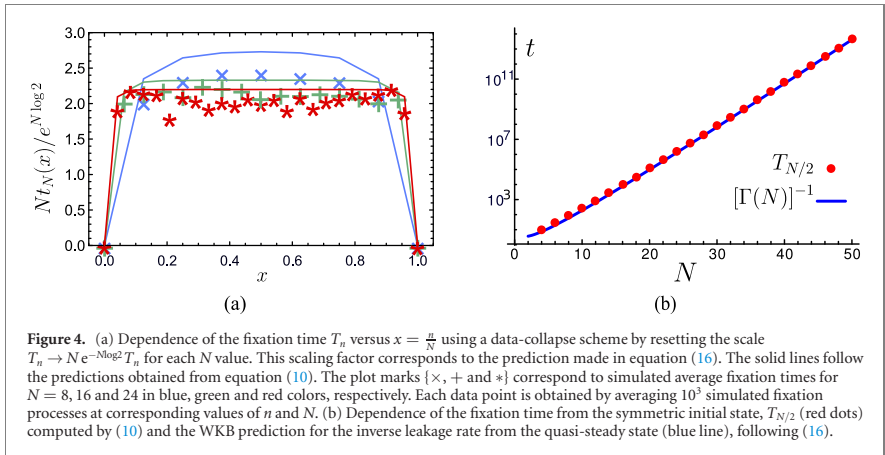


Figure 4. (a) Dependence of the fixation time T_n versus $x = \frac{n}{N}$ using a data-collapse scheme by resetting the scale $T_n \rightarrow N e^{-N \log 2} T_n$ for each N value. This scaling factor corresponds to the prediction made in equation (16). The solid lines follow the predictions obtained from equation (10). The plot marks $\{ \times, + \text{ and } * \}$ correspond to simulated average fixation times for $N = 8, 16$ and 24 in blue, green and red colors, respectively. Each data point is obtained by averaging 10^3 simulated fixation processes at corresponding values of n and N . (b) Dependence of the fixation time from the symmetric initial state, $T_{N/2}$ (red dots) computed by (10) and the WKB prediction for the inverse leakage rate from the quasi-steady state (blue line), following (16).

Now using (1a) and (1b) for $a(x)$ and $b(x)$, we have

$$\tilde{P}(x) \sim \frac{e^{N S_0(x)}}{x^{3/2}(1-x)^{3/2}}, \tag{14}$$

with $S_0(x) = -x \log x - (1-x) \log(1-x)$. Note that the action $S_0(x)$ is peaked at the quasi-state state $x = \frac{1}{2}$. We normalize $\tilde{P}(x)$ by using the Laplace method for $N \rightarrow \infty$ [34],

$$\int_0^1 \frac{e^{N S_0(x)}}{x^{3/2}(1-x)^{3/2}} dx \approx \sqrt{\frac{32}{N}} e^{N \log 2},$$

so that

$$\tilde{P}(x) \simeq \frac{N}{\sqrt{32\pi}} \frac{e^{N[-x \log x - (1-x) \log(1-x) - \log 2]}}{x^{3/2}(1-x)^{3/2}}. \tag{15}$$

We now compute the fixation rate Γ by substituting

$$\tilde{P}\left(\frac{1}{N}\right) \simeq \frac{N^3}{\sqrt{32\pi}} e^{1-N \log 2} \quad \text{and} \quad b\left(\frac{1}{N}\right) \simeq \frac{1}{N^2},$$

into equation (11) to give

$$\Gamma(N) = \frac{N e}{\sqrt{8\pi}} e^{-N \log 2}. \tag{16}$$

We now identify the inverse of this rate with the average fixation time. As shown in figure 4(b), this inverse rate accurately matches the simulation data for the fixation time.

5. Two-species cooperation with migration

We now incorporate migration into the dynamics, in which particles of either species migrate into the population at the same fixed rate λ , and each new particle replaces a randomly selected existing particle. Because migration is accompanied by replacement, the population size remains fixed, which is the physically most relevant case. Now the population is driven to a steady state rather than to fixation and we want to understand the nature of this steady state.

5.1. Probability distribution

For a population that consists of n A's and $(N-n)$ B's, suppose that the migrant is an A. With probability $\frac{1}{2}(1-x)$, the A migrant replaces a B and $n \rightarrow n+1$, while with probability $\frac{1}{2}x$, the A migrant replaces an A, and the composition of the population remains the same. Similar reasoning applies when the migrant is a B. As a result of a migration event, the average change in the number of A's is $\frac{1}{2}(1-x) - \frac{1}{2}x$. The rate equation for n now is (compare with equation (2))

$$\langle \dot{n} \rangle = N(1-\lambda) [x(1-x)(1-2x)] + \frac{1}{2} N \lambda (1-2x). \tag{17}$$

For $\lambda > 0$, $x = 0$ and $x = 1$ are no longer fixed points and only the remaining fixed point at $x = \frac{1}{2}$ is stable. In the absence of fluctuations, the population is thus driven to a steady-state distribution, $P_n(t \rightarrow \infty)$, that is peaked about $x = \frac{1}{2}$. Because there is no absorbing state in the stochastic dynamics, we might anticipate a similar behavior for $P_n(t \rightarrow \infty)$ when stochasticity is accounted for. We will show, however, that within the Fokker–Planck approximation the steady-state distribution can either be unimodal or trimodal in shape and the latter case corresponds to a steady state that is not truly steady.

The probability distribution P_n is now governed by the master equation

$$\begin{aligned} \dot{P}_n = & N(1 - \lambda) [a_{n-1}P_{n-1} + b_{n+1}P_{n+1} - (a_n + b_n)P_n] \\ & + N\lambda [c_{n-1}P_{n-1} + d_{n+1}P_{n+1} - (c_n + d_n)P_n], \end{aligned} \tag{18}$$

with hopping probabilities due to migration that are given by

$$c_n = \frac{1}{2} \left(1 - \frac{n}{N}\right) \quad d_n = \frac{1}{2} \frac{n}{N}. \tag{19}$$

We now determine the continuum probability distribution in the Fokker–Planck approximation. As we shall see, this continuum expression for the probability distribution matches simulation data quite well, thus justifying the Fokker–Planck approximation *ex post facto* as a way to probe steady-state properties.

In terms of $x = \frac{n}{N}$, $dx = \frac{1}{N}$, $P_n \rightarrow P(x)$, we expand (18) in a Taylor series up to second order. This gives the Fokker–Planck equation [28, 35]

$$\begin{aligned} P_t = & - \left\{ (1-2x) \left[(1-\lambda)x(1-x) + \frac{\lambda}{2} \right] P(x, t) \right\}_x + \frac{1}{2N} \left\{ \left[(1-\lambda)x(1-x) + \frac{\lambda}{2} \right] P(x, t) \right\}_{xx} \\ \equiv & - \{v(x)P(x, t)\}_x + \{D(x)P(x, t)\}_{xx}, \end{aligned} \tag{20}$$

where the subscripts denote partial derivatives.

The steady state is defined by solving this equation with the left-hand side set to zero. Integrating once gives $(DP)_x - vP = B$, where B is a constant. We determine the constant by evaluating this equation at the symmetry point $x = \frac{1}{2}$. Because the probability distribution is symmetric about $x = \frac{1}{2}$, $P_x(x = \frac{1}{2}) = 0$. Moreover, at $x = \frac{1}{2}$, $v = 0$ and $D_x = 0$, which implies that $B = 0$. Thus we only need to solve $(DP)_x - vP = 0$, whose solution is

$$\begin{aligned} P(x) = & C \exp \left\{ \int^x dy \frac{v(y) - D_y(y)}{D(y)} \right\} = C \exp \left\{ -\log D(x) + \int^x dy \frac{v(y)}{D(y)} \right\} \\ = & \frac{C}{D(x)} \exp \left\{ \int^x dy 2N(1 - 2y) \right\} \\ = & C' \left[\frac{1}{(1-\lambda)x(1-x) + \frac{\lambda}{2}} \right] e^{2Nx(1-x)}, \end{aligned} \tag{21}$$

where the constant C' is determined by normalization.

For $\lambda \rightarrow 0$, $P(x)$ is concentrated near $x = 0$ and near $x = 1$; these peaks correspond to the near-fixation states. Because of rare fluctuations, however, the population stochastically switches between states where almost all particles are of type A to states where almost all particles are of type B. Naively, one therefore anticipates that the steady-state distribution should be bimodal, with a peak at each of the two near-fixation states. Unexpectedly, there always remains a peak at $x = \frac{1}{2}$ (which may be vanishingly small), so that the this distribution is *trimodal* in the small- λ regime. As λ increases beyond a critical value, the steady-state distribution undergoes a trimodal to unimodal transition (figure 5).


For fixed N , we determine the transition between trimodality and unimodality by finding the point(s) where $P'(x) = 0$. This calculation gives, after straightforward algebra,

$$P'(x) \propto (1 - 2x) e^{2Nx(1-x)} \left[2N - \frac{(\lambda - 1)}{D(x)^2} \right],$$

where again $D(x)$ is the diffusion coefficient defined by equation (20). The leading factor of $1 - 2x$ equals 0 at $x = \frac{1}{2}$ and corresponds to the extremum in the distribution at $x = \frac{1}{2}$.

However, there are additional extrema at the points where the factor in the square brackets equals 0. To determine these extrema, we first determine the zero of this factor at $x = 0, 1$. Thus we have the condition $2N = (1 - \lambda)/(\lambda/2)^2$. Since we will find that $\lambda \ll 1$, we also neglect λ compared to 1 to give

$$\lambda_c = \sqrt{\frac{2}{N}}. \tag{22}$$



For $\lambda < \lambda_c$, the distribution $P(x)$ has three extrema. To find the location of the two secondary extrema for $\lambda \lesssim \lambda_c$, it is convenient to now use the variable $y = x - \frac{1}{2}$ with $y \rightarrow \pm \frac{1}{2}$, which corresponds to x close to zero or to 1. Now the condition that the factor in the square brackets equals zero gives

$$2N \simeq \frac{1}{(2\epsilon + \lambda/2)^2},$$

from which we obtain

$$\epsilon \simeq \frac{1}{4}(\lambda_c - \lambda) \quad \text{for } \lambda < \lambda_c. \tag{23}$$

In the regime $\lambda < \lambda_c$, the distribution $P(x)$ is necessarily trimodal because the point $x = \frac{1}{2}$ is always a local maximum of $P(x)$. To verify this point, we compute the second derivative of $P(x)$ at $x = \frac{1}{2}$,

$$P''(x) = e^{2N\epsilon(1-x)} \left[(1 - 2x)^2 - 2 + \frac{2(1 - \lambda)}{D(x)^3} D'(x) \right],$$

which is indeed negative at $x = \frac{1}{2}$.


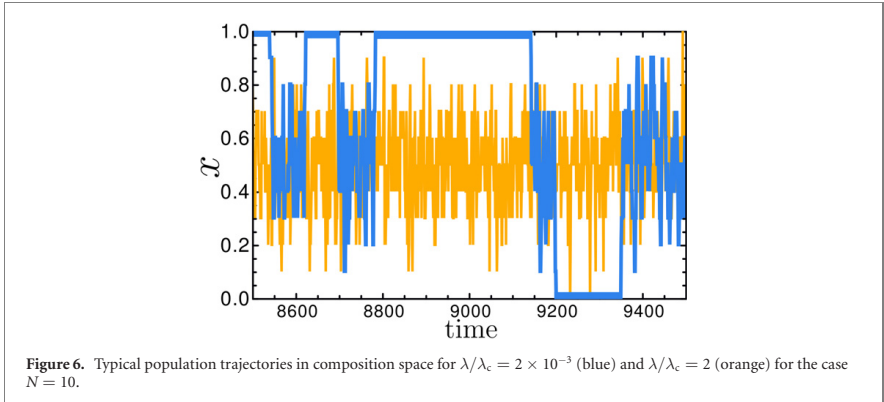
5.2. Macroscopic fluctuations in the steady state

An intriguing feature of two-species cooperation with migration is that the steady state is not strictly steady, especially when λ is small (figure 6). For $\lambda \ll \lambda_c$, substantial time ranges exist during which little or no immigration occurs. During these periods, the population will tend to approach one of the fixation states. Even if the population does reach a state of all A's or all B's, immigration eventually drives the population towards an equal number of A's and B's. The competing effects of cooperation and immigration therefore cause the population to wander stochastically from one near-fixation state to the other, with returns to the equal-concentration point $x = \frac{1}{2}$ controlled by the immigration rate (figure 6). Related phenomenology occurs in noisy voter models [36–39]. The extended time periods during which the population is close to fixation corresponds to the large weight in the secondary peaks of the probability distribution in figure 5. In contrast, when the immigration rate is much larger than λ_c , the rapid inflow of equal numbers of A's and B's ensures that the population contains roughly equal numbers of both species.

One way to quantify these fluctuations is by the *return time* RT , which we define as the time interval between successive points where $x = \frac{1}{2}$. As suggested by figure 6, this return time will be short for $\lambda \gg \lambda_c$ and become long for $\lambda \ll \lambda_c$. The latter behavior is a harbinger of large composition fluctuations in the population. We can again use the backward Kolmogorov approach to determine the λ dependence of the return time. Let τ_n denote the average time to reach the balanced state of equal numbers of A's and B's when starting from a state where the number of A's equals $n > \frac{1}{2}N$. Then

$$RT = \frac{1}{N} + \tau_{1+N/2}. \tag{24}$$

That is, starting from the balanced state, the average return time to this state is the average time for a single event, $\frac{1}{N}$, where the population now consists of $\frac{N}{2} + 1$ A's and $\frac{N}{2} - 1$ B's, plus the average time to reach the balanced state when starting from this minimally imbalanced state.



In analogy with equation (9) for the fixation time, the time τ_n satisfies the recursion

$$\begin{aligned} \tau_n = (1 - \lambda) [a_n \tau_{n+1} + b_n \tau_{n-1} + (1 - a_n - b_n) \tau_n] \\ + \lambda [c_n \tau_{n+1} + d_n \tau_{n-1} + (1 - c_n - d_n) \tau_n] + \delta t, \end{aligned} \tag{25a}$$

with a_n, b_n given by (1a) and (1b), and c_n, d_n given by (19). We can rewrite this recursion in the canonical form of equation (9):

$$\tau_n = a'_n \tau_{n+1} + b'_n \tau_{n-1} + (1 - a'_n - b'_n) \tau_n + \delta t, \tag{25b}$$

where $a'_n = (1 - \lambda)a_n + \lambda c_n$ and $b'_n = (1 - \lambda)b_n + \lambda d_n$. This recursion is valid for $\frac{1}{2}N < n < N$, while for $n = N$ the first term in (25b) is absent. This missing term acts as the effective reflecting boundary condition for $n = N$. At $n = \frac{1}{2}N$, the appropriate boundary condition is $\tau_{N/2} = 0$; namely, the balanced state corresponds to the end of the process. The result for τ_n for arbitrary n is given in (D.5), while the return time RT itself is given in (D.4) (see appendix D for details). Figure 7 shows that RT scales as λ^{-1} for $\lambda \rightarrow 0$. This behavior is the source of the long-lived temporal fluctuations in the composition of the population, as illustrated in figure 6.

6. Discussion

Much of the literature on populations of multiple cooperating species has focused on the continuous limit. Such a population is driven to an attractor state in which there are equal concentrations of each species. However, when finite-population stochasticity is incorporated, the true attractors of the dynamics of an isolated population are instead the fixation states, in which only one species exists. Because stochastic effects are relevant

in real systems, a discrete approach that incorporates this stochasticity is necessary to describe the dynamics in a faithful way.

Within the discrete approach, we determined the probability for a finite population to reach a given fixation state as a function of the initial condition, as well as the time to reach fixation. The behaviors of these two quantities reflect the effective bias that drives the system to the quasi-steady state of equal concentrations of the two species. Namely, the fixation probability is nearly independent of the initial condition and the fixation time scales exponentially with population size N . As a consequence of this exponential dependence, fixation is not observable in a laboratory or in a simulational time scale for any reasonable population size.

It is worth mentioning that the statistical features of two-species cooperation share similarities with the vacillating voter model [40], despite their different microscopic update rules. In the vacillating voter model, agents (voters) can be in one of two voting states and their agreement or disagreement is influenced by the state of yet another neighbor. Here ‘vacillation’ refers to the possibility that a voter does not adopt the state of a randomly chosen neighbor, as in the standard voter model, but may adopt the state of this additional neighbor. The properties of this decision process drives the population toward a zero-magnetization state. This state is equivalent to the attractor in two-species cooperation, where both species are equally represented.

When migration into the system can also occur, the population now ostensibly reaches a steady state. An unanticipated feature of this steady state for small migration rate λ is that this state is not genuinely steady, because there are long-term stochastic fluctuations that drive the population from one near-fixation state to the opposite near-fixation state, with the population spending long time periods in these near-fixation states. Such macroscopic wanderings are reflected in the steady-state abundance distribution, which is strongly peaked at these near-fixation states for sufficiently small λ . The time scale associated with these fluctuations increases rapidly as the migration rate decreases. Thus observations of a cooperative system have to be sufficiently long to incorporate many such wanderings so as to ensure that true average behavior is actually probed. We found that the return time—the time between successive instants where the number of A’s and B’s are equal—allows us to quantify these temporal fluctuations in a precise way.

This observation of large fluctuations also has important ramifications for populations that consist of more than two cooperating species. Depending the immigration rate, the population size N , and the number of distinct species S , the number of species that are actually present at any given time could be much less than S . Thus a *typical* steady state could have a very different character than the *average* steady state that is predicted by the time-independent density distribution. Moreover, a multispecies population will also exhibit large fluctuations in the actual composition of the species that are present. This intriguing issue has recently been investigated in the context of multispecies Lotka–Volterra models [41, 42]. The backward Kolmogorov equation offers the possibility of obtaining new insights into these large fluctuations because of the relative simplicity of neutral models of cooperating species.

The predictions presented in our study can also be tested in experimental settings that are based on microfluidic chambers, where small populations of cells can be maintained so that long-term monitoring can be performed. These small-scale devices offer a unique opportunity to explore the impact of population size and validate the approximations presented in this work. Both well-mixed populations, as well as two- and three-dimensional spatial populations, could be maintained in constant numbers over time. Future extensions of our model, such as including cell death or environmental noise, would be helpful to design experimental protocols and explore the conditions required to maintain stable cooperative cell assemblies [43, 44].

Acknowledgments

The authors thank Deepak Bhat and Jacopo Grilli for fruitful discussions, and Eric A Blair for inspiring ideas. JP and RS were supported by the Botín Foundation and the Spanish Ministry of Economy and Competitiveness through Grant FIS2015-67616-P MINEICO/AEI/FEDER. JP is also supported by ‘María de Maeztú’ fellowship MDM-2014-0370-17-2. SRs research was supported in part by NSF Grant DMR-1910736. JP and RS thank the hospitality of the Santa Fe Institute, where this project began.

Appendix A. The fixation probability

We want to solve equation (7) for the fixation probability:

$$E_n = a_n E_{n+1} + b_n E_{n-1} + (1 - a_n - b_n) E_n.$$

This calculation is standard (see, e.g., [45]) and we provide it here so that our presentation is self contained. It is convenient to rewrite the above equation as $a_n(E_{n+1} - E_n) = b_n(E_n - E_{n-1})$, and then define $u_n \equiv E_n - E_{n-1}$

and $r_n \equiv b_n/a_n$ to recast it as the following first-order recursion for the u_n :

$$u_n = r_{n-1}u_{n-1} = r_{n-1}r_{n-2}r_{n-3} \dots r_1u_1,$$

with $u_1 = E_1 - E_0 = E_1$. We now define $R_n = \prod_{m=1}^n r_m$ so that the equation for u_n becomes

$$u_n = R_{n-1}u_1 = R_{n-1}E_1. \tag{A.1}$$

Since the u_n are successive differences of the E_n , we determine E_n by summing the u_n . Thus

$$E_n = \sum_{m=1}^n u_m = \sum_{m=1}^n R_{m-1}E_1 = \sum_{m=0}^{n-1} R_m E_1, \tag{A.2}$$

where we need to define $R_0 = 1$ for consistency. We now determine the unknown E_1 by using the boundary condition $E_N = 1$ in (A.2) to give $E_N = \sum_{m=0}^{N-1} R_m E_1 = 1$. Having found E_1 , the general solution for E_n is

$$E_n = \sum_{m=0}^{n-1} R_m \bigg/ \sum_{m=0}^{N-1} R_m. \tag{A.3}$$

To simplify the above expression, we start with a_n and b_n defined in equation (1):

$$a_n = \frac{n}{N} \left(1 - \frac{n}{N}\right)^2 \quad b_n = \left(\frac{n}{N}\right)^2 \left(1 - \frac{n}{N}\right),$$

so that $r_n = b_n/a_n = n/(N - n)$. Then

$$R_n = r_1 r_2 \dots r_n = \frac{n!(N - n - 1)!}{(N - 1)!} = \left[\binom{N - 1}{n} \right]^{-1}.$$

Substituting this representation for R_n in (A.3) gives equation (8).

For completeness, we also give the continuum solution for the fixation probability. We take the continuum limit of equation (7) by letting $(n \pm 1)/N \rightarrow x \pm dx$, with $dx = \frac{1}{N}$, and then expanding this equation to second order in dx . This gives

$$E'' + 2N(1 - 2x)E' = 0, \tag{A.4}$$

where the prime denotes differentiation with respect to x . This equation is subject to the boundary conditions $E(0) = 0$ and $E(1) = 1$. As in the discrete formulation, the first condition corresponds to the impossibility of reaching a population of all A's if the initial state contains no A's, while the second condition corresponds to the initial state coinciding with the desired final state. The solution to (A.4), subject to the given boundary conditions is

$$E(x) = \frac{\int_0^x du e^{2N(u^2-u)}}{\int_0^1 du e^{2N(u^2-u)}} = \frac{1}{2} \left[1 + \frac{\operatorname{erfi}(\sqrt{2N}(x - \frac{1}{2}))}{\operatorname{erfi}(\sqrt{N/2})} \right], \tag{A.5}$$

where erfi is the imaginary error function. This expression from the continuum approximation agrees well with the exact discrete result (8) (see figure 3). However, the continuum approach is no longer accurate for the fixation time (see the next section).

Appendix B. The fixation time

We now solve the recursions (9) for the fixation time:

$$T_n = a_n T_{n+1} + b_n T_{n-1} + (1 - a_n - b_n) T_n + \delta t.$$

Following the same steps that led to (A.1), we obtain, for the difference $v_n \equiv T_n - T_{n-1}$,

$$v_n = r_{n-1}v_{n-1} - \alpha_{n-1}, \tag{B.1}$$

where $\alpha_n \equiv \delta t/a_n$. Notice that $v_1 = T_1 - T_0 = T_1 \equiv R_0 T_1$.

We develop the recursion (B.1) to give

$$\begin{aligned} v_n &= r_{n-1}r_{n-2}v_{n-2} - \alpha_{n-1} - r_{n-1}\alpha_{n-2} \\ &= r_{n-1}r_{n-2}r_{n-3}v_{n-3} - \alpha_{n-1} - r_{n-1}\alpha_{n-2} - r_{n-1}r_{n-2}\alpha_{n-3} \\ &\quad \vdots \end{aligned}$$

continuing this development to the end leads to

$$v_n = R_{n-1}v_1 - Q_{n-1} = R_{n-1}T_1 - Q_{n-1}, \tag{B.2}$$

where

$$Q_n \equiv \alpha_n + r_n\alpha_{n-1} + r_n r_{n-1}\alpha_{n-2} + \dots + r_n r_{n-1} \dots r_2\alpha_1,$$

and $Q_0 = 0$ by virtue of (B.2).

Finally, we sum the v_n to obtain the fixation time:

$$T_n = \sum_{m=1}^n v_m = \sum_{m=0}^{n-1} R_m T_1 - \sum_{m=1}^{n-1} Q_m. \tag{B.3}$$

To eliminate the unknown T_1 , we use the boundary condition $T_N = 0$ to give

$$T_1 = \sum_{m=1}^{N-1} Q_m / \sum_{m=0}^{N-1} R_m.$$

Substituting this in (B.3) and noting that $\sum_{m=0}^{N-1} R_m / \sum_{m=0}^{N-1} R_m$ is just the fixation probability E_n , we obtain the result quoted in equation (10).

Appendix C. WKB approximation

A comprehensive review of the application of the WKB method for large deviations in stochastic populations can be found in [26]. In this section we summarize the basic steps to reach equation (13). We make the ansatz $\bar{P}(x) \sim \exp \{NS_0(x) + S_1(x) + O(N^{-1})\}$, and expand to linear order in δx to give

$$\bar{P}(x \pm \delta x) \simeq \bar{P}(x) \exp \left\{ \pm S'_0 + \delta x \left(\frac{S''_0}{2} \pm S'_1 \right) \right\} \simeq \bar{P}(x) e^{\pm S_0} \left[1 + \delta x \left(\frac{S''_0}{2} \pm S'_1 \right) \right],$$

with $\delta x = 1/N$. We now define $\Lambda \equiv e^{S_0}$ and substitute this into (12) to obtain

$$(a - a'\delta x) \Lambda^{-1} \left[1 + \delta x \left(\frac{S''_0}{2} - S'_1 \right) \right] + (b + b'\delta x) \Lambda \left[1 + \delta x \left(\frac{S''_0}{2} + S'_1 \right) \right] = a + b,$$

which can be separated into terms of $O(1)$ and terms of $O(\delta x)$. For the former, we have

$$a\Lambda^{-1} + b\Lambda = a + b, \tag{C.1}$$

which has the two solutions $\Lambda_0 = 1$ and $\Lambda = \frac{a}{b}$. The Λ_0 solution corresponds to $S'_0 = 0$, and the resulting constant can be absorbed by the normalization condition on $\bar{P}(x)$. The second solution is

$$S'_0 = \log \left(\frac{a}{b} \right), \tag{C.2}$$

which, after integration, results in the first expression in (13). For the $O(\delta x)$ terms we must solve

$$-a'\Lambda^{-1} + a\Lambda^{-1} \left(\frac{S''_0}{2} - S'_1 \right) + b'\Lambda + b\Lambda \left(\frac{S''_0}{2} + S'_1 \right) = 0, \tag{C.3}$$

which, after substitution of $\Lambda = \frac{a}{b}$, yields

$$b \left(\frac{S''_0}{2} - S'_1 \right) + a \left(\frac{S''_0}{2} + S'_1 \right) = b \frac{a'}{a} - a \frac{b'}{b}. \tag{C.4}$$

On the other hand, $S''_0 = \frac{a'}{a} - \frac{b'}{b}$, which reduces the previous equation to

$$(a - b)S'_1 = \frac{1}{2} \left[b \frac{a'}{a} - a \frac{b'}{b} - (a' - b') \right] = -\frac{1}{2} \left(\frac{a'}{a} + \frac{b'}{b} \right) (a - b). \tag{C.5}$$

The first term in the brackets on the right-hand side corresponds to the derivative of $\log(ab)$. Hence, after integration, we obtain the second expression in (13).

Appendix D. The return time

We want to solve the recursion for the return time τ_n , defined as the time for the population to first reach the state with equal numbers of A's and B's when the initial state contains $n > \frac{1}{2}N$ A's. The state with $n = N$ plays the role of an effective reflecting boundary condition. The system of equations that we wish to solve is (25b):

$$\tau_n = a'_n \tau_{n+1} + b'_n \tau_{n-1} + (1 - a'_n - b'_n) \tau_n + \delta t, \tag{D.1a}$$

for $1 + N/2 \leq n < N$, while for $n = N$ the appropriate equation is

$$\tau_N = b'_N \tau_{N-1} + (1 - a'_N - b'_N) \tau_N + \delta t. \tag{D.1b}$$

Using the fact that $a'_N = 0$, this last equation can be rewritten as

$$\tau_N - \tau_{N-1} \equiv v_N = \frac{\delta t}{b'_N}. \tag{D.2}$$

The remaining $\frac{N}{2} - 1$ equations (D.1a) are of the same type as (9) for the fixation time. Thus the solution for τ_n has the same form as (B.3):

$$\tau_n = \sum_{m=N/2}^{n-1} R'_m \tau_{1+N/2} - \sum_{m=1+N/2}^{n-1} Q'_m, \tag{D.3}$$

where

$$R'_n = \prod_{m=1+N/2}^n \frac{b'_m}{a'_m} \equiv \prod_{m=1+N/2}^n r'_m$$

and

$$Q'_n = \alpha'_n + r'_n \alpha'_{n-1} + r'_n r'_{n-1} \alpha'_{n-2} + \dots + r'_n r'_{n-1} \dots r'_{2+N/2} \alpha'_{1+N/2},$$

with $\alpha'_n = \delta t/a'_n$ and $Q'_{N/2} = 0$.

To eliminate the unknown $\tau_{1+N/2}$, we now write (D.3) for the special cases of $n = N$ and $n = N - 1$:

$$\tau_N = \sum_{m=N/2}^{N-1} R'_m \tau_{1+N/2} - \sum_{m=1+N/2}^{N-1} Q'_m \quad \tau_{N-1} = \sum_{m=N/2}^{N-2} R'_m \tau_{1+N/2} - \sum_{m=1+N/2}^{N-2} Q'_m.$$

Their difference is

$$\tau_N - \tau_{N-1} = R'_{N-1} \tau_{1+N/2} - Q'_{N-1} = \frac{\delta t}{b'_N},$$

so that $\tau_{1+N/2}$ is given by

$$\tau_{1+N/2} = \frac{Q'_{N-1} + \delta t/b'_N}{R'_{N-1}}. \tag{D.4}$$

Substituting this expression for $\tau_{1+N/2}$ in (D.3) gives the average time to reach the balanced state of equal numbers of A's and B's when starting from a population that contains n A's with a reflecting boundary condition at $n = N$:

$$\tau_n = \sum_{m=N/2}^{n-1} R'_m \frac{Q'_{N-1} + \delta t/b'_N}{R'_{N-1}} - \sum_{m=1+N/2}^{n-1} Q'_m. \tag{D.5}$$

What we want, however, is, the return time, defined as the average time to start at the balanced state and first return to this state. This is $RT = \frac{1}{N} + \tau_{1+N/2}$.

ORCID iDs

Jordi Piñero  <https://orcid.org/0000-0002-4183-3733>

S Redner  <https://orcid.org/0000-0001-5906-4904>

Ricard Solé  <https://orcid.org/0000-0001-6974-1008>

References

- [1] Goel N S, Maitra S C and Montroll E W 1971 *Rev. Mod. Phys.* **43** 231
- [2] Murray J D 2001 *Mathematical Biology: I. An Introduction* (New York: Springer)
- [3] May R M 2001 *Stability and Complexity in Model Ecosystems* (Princeton, NJ: Princeton University Press)

- [4] May R M, Levin S A and Sugihara G 2008 *Nature* **451** 893–4
- [5] Bronstein J L 2015 *Mutualism* (Oxford: Oxford University Press)
- [6] Perc M 2016 *Phys. Lett. A* **380** 2803–8
- [7] Perc M, Jordan J J, Rand D G, Wang Z, Boccaletti S and Szolnoki A 2017 *Phys. Rep.* **687** 1–51
- [8] Mitri S, Clarke E and Foster K R 2016 *ISME J.* **10** 1471–82
- [9] Nadell C D, Drescher K and Foster K R 2016 *Nat. Rev. Microbiol.* **14** 589–600
- [10] Müller M J I, Neugeboren B I, Nelson D R and Murray A W 2014 *Proc. Natl Acad. Sci.* **111** 1037–42
- [11] Shou W, Ram S and Vilar J M G 2007 *Proc. Natl Acad. Sci.* **104** 1877–82
- [12] Amor D R, Montañez R, Duran-Nebreda S and Solé R 2017 *PLoS Comput. Biol.* **13** e1005689
- [13] Amor D R and Dal Bello M 2019 *Life* **9** 22
- [14] Nowak M A 2006 *Evolutionary Dynamics: Exploring the Equations of Life* (Cambridge, MA: Harvard University Press)
- [15] Antal T and Scheuring I 2006 *Bull. Math. Biol.* **68** 1923–44
- [16] Altrock P M and Traulsen A 2009 *New J. Phys.* **11** 013012
- [17] Black A J, Traulsen A and Galla T 2012 *Phys. Rev. Lett.* **109** 028101
- [18] Nadell C D, Foster K R and Xavier J B 2010 *PLoS Comput. Biol.* **6** e1000716
- [19] Rakoff-Nahoum S, Foster K R and Comstock L E 2016 *Nature* **533** 255–9
- [20] Foster K R, Schluter J, Coyte K Z and Rakoff-Nahoum S 2017 *Nature* **548** 43–51
- [21] Friedman J and Gore J 2017 *Curr. Opin. Syst. Biol.* **1** 114–21
- [22] Vega N M and Gore J 2018 *Curr. Opin. Syst. Biol.* **45** 195–202
- [23] Elgart V and Kamenev A 2004 *Phys. Rev. E* **70** 041106
- [24] Kessler D A and Shnerb N M 2007 *J. Stat. Phys.* **127** 861–86
- [25] Assaf M and Meerson B 2010 *Phys. Rev. E* **81** 021116
- [26] Assaf M and Meerson B 2017 *J. Phys. A: Math. Theor.* **50** 263001
- [27] Krapivsky P L, Redner S and Ben-Naim E 2010 *A Kinetic View of Statistical Physics* (New York: Cambridge University Press)
- [28] Van Kampen N G 1992 *Stochastic Processes in Physics and Chemistry* vol 1 (Amsterdam: Elsevier)
- [29] Redner S 2001 *A Guide to First-Passage Processes* (New York: Cambridge University Press)
- [30] Lee S et al 2012 *Mathematics Stack Exchange* <https://math.stackexchange.com/questions/151441/calculate-sums-of-inverses-of-binomial-coefficients>
- [31] Moran P A P et al 1962 *The Statistical Processes of Evolutionary Theory* (Oxford: Oxford University Press)
- [32] Kimura M 1983 *The Neutral Theory of Molecular Evolution* (Cambridge: Cambridge University Press)
- [33] Ewens W J 2012 *Mathematical Population Genetics I: Theoretical Introduction* (New York: Springer)
- [34] Bender C M and Orszag S A 1999 *Advanced Mathematical Methods for Scientists and Engineers I: Asymptotic Methods and Perturbation Theory* (New York: Springer)
- [35] Gardiner C W et al 1985 *Handbook of Stochastic Methods* 3rd edn (Berlin: Springer)
- [36] Fichthorn K, Gulari E and Ziff R 1989 *Phys. Rev. Lett.* **63** 1527
- [37] Considine D, Redner S and Takayasu H 1989 *Phys. Rev. Lett.* **63** 2857
- [38] Carro A, Toral R and Miguel M S 2016 *Sci. Rep.* **6** 24775
- [39] Herrerías-Azcué F and Galla T 2019 *Phys. Rev. E* **100** 022304
- [40] Lambiotte R and Redner S 2007 *J. Stat. Mech.* **L10001**
- [41] Bunin G 2017 *Phys. Rev. E* **95** 042414
- [42] Pearce M T, Agarwala A and Fisher D S 2020 *Proc. Natl Acad. Sci. USA* **117** 14572–83
- [43] Ferry M S, Razinkov I A and Hasty J 2011 *Methods Enzymol.* **497** 295–372
- [44] Luke C S, Selimkhanov J, Baumgart L, Cohen S E, Golden S S, Cookson N A and Hasty J 2016 *ACS Synth. Biol.* **5** 8–14
- [45] Karlin S and Taylor H M 1975 *A First Course in Stochastic Processes* (San Diego, CA: Academic)

Review



Cite this article: Piñero J, Solé R. 2019

Statistical physics of liquid brains.

Phil. Trans. R. Soc. B **374**: 20180376.

<http://dx.doi.org/10.1098/rstb.2018.0376>

Accepted: 7 March 2019

One contribution of 15 to a theme issue
'Liquid brains, solid brains: How distributed
cognitive architectures process information'.

Subject Areas:

theoretical biology, evolution, cognition,
biophysics, immunology

Keywords:

brains, collective intelligence, criticality,
phase transitions, evolution

Authors for correspondence:

Jordi Piñero

e-mail: jordi.pinero@upf.edu

Ricard Solé

e-mail: ricard.sole@upf.edu

Statistical physics of liquid brains

Jordi Piñero^{1,2} and Ricard Solé^{1,2,3}

¹ICREA-Complex Systems Lab, Universitat Pompeu Fabra, 08003 Barcelona, Spain

²Institut de Biologia Evolutiva (CSIC-UPF), Pg Marítim Barceloneta, 37, 08003 Barcelona, Spain

³Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

JP, 0000-0002-4183-3733; RS, 0000-0001-6974-1008

Liquid neural networks (or 'liquid brains') are a widespread class of cognitive living networks characterized by a common feature: the agents (ants or immune cells, for example) move in space. Thus, no fixed, long-term agent-agent connections are maintained, in contrast with standard neural systems. How is this class of systems capable of displaying cognitive abilities, from learning to decision-making? In this paper, the collective dynamics, memory and learning properties of liquid brains is explored under the perspective of statistical physics. Using a comparative approach, we review the generic properties of three large classes of systems, namely: standard neural networks (solid brains), ant colonies and the immune system. It is shown that, despite their intrinsic physical differences, these systems share key properties with standard neural systems in terms of formal descriptions, but strongly depart in other ways. On one hand, the attractors found in liquid brains are not always based on connection weights but instead on population abundances. However, some liquid systems use fluctuations in ways similar to those found in cortical networks, suggesting a relevant role for criticality as a way of rapidly reacting to external signals.

This article is part of the theme issue 'Liquid brains, solid brains: How distributed cognitive architectures process information'.

1. Introduction

As pointed out by physicist John Hopfield, biology is different from physics in one fundamental way: biological systems perform computations [1]. Within the context of evolution, a crucial ingredient for the emergence of biological complexity required the development of information-processing systems at multiple scales [2]. Adaptation to a dynamic environment deeply benefited from non-genetic processes that allowed response mechanisms to short-term changes. Thus, biological computation is an intrinsic part of our current understanding of cell phenotypes [3] and, not surprisingly, the molecular webs of interactions connecting genes, proteins and metabolites have often been represented in terms of computations [4].

Once fast-responding molecular signalling mechanisms were in place, a whole range of possibilities became available: individuals could not only respond to environmental cues, but they could also start to interact with other individuals prompting a higher-order cognitive network [5]. Such transition took place in a diverse range of ways. It included the development of the first brain-like structures [6,7] as well as societies formed by relatively simple agents (ants, termites or bees) capable of performing complex cognitive actions at the collective level [8,9]. Ant colonies have been compared to brains as both exhibit emergent collective phenomena (dynamical and structural patterns of organization and behaviour that cannot be reduced to the properties of single ants) and display cognition on a large scale beyond that of the individual components [10,11]. These two examples represent two distinguishable large classes of networks. Along with ant colonies, immune systems (ISs) also share traits characteristic of the metazoan nerve nets yet they strongly depart from them in the fluid embodiment and nature of cell-cell interactions.

The previous three examples are displayed in figure 1. Here coupled neurons (figure 1*a*), interacting ants (figure 1*b*) or immune cells responding to novel challenges (figure 1*c*) are shown, along with minimal representations of the underlying networks (figure 1*d-f*). The classical picture of a neural network involves a topological structure (a graph) with neurons occupying the nodes and interneuronal

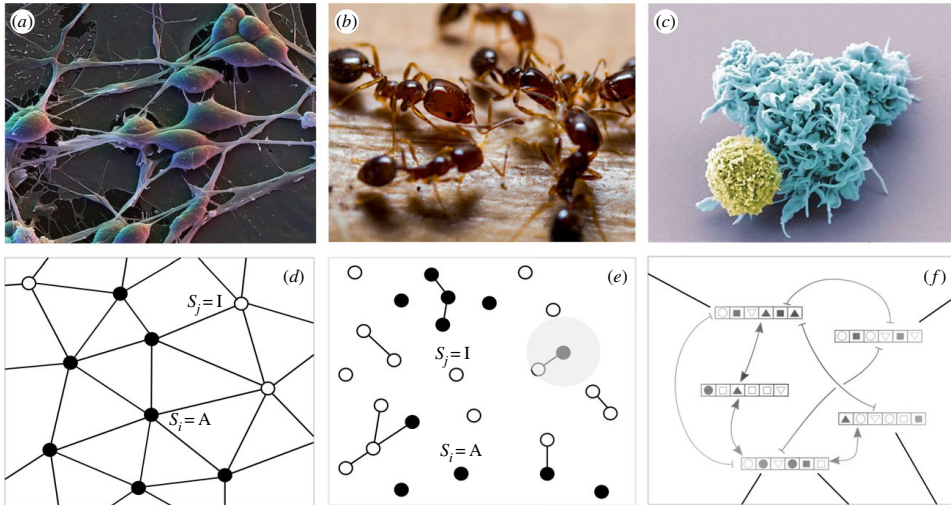


Figure 1. Network interactions in liquid versus solid brains. The three case studies analysed in this paper are shown, with examples of the agents involved in each case. (a) Standard neural networks involve spatially localized cells connected through synaptic weights. In contrast with this architecture, liquid brains, including (b) ant colonies and (c) the IS include mobile agents (or cell subsets) interacting in space and time with no fixed pairwise weights. A schematic of each case study is outlined in the row below. Standard neural networks are defined in terms of connected excitable elements that can be roughly classified in active (firing) and inactive (quiescent) neurons, here indicated as filled and open circles, respectively (d). The wiring matrix remains basically the same in terms of topology (who is connected with whom) but will be modified in strength due to experience. By contrast, ant colonies must be represented by disconnected graphs (e) where interactions are possible within a given spatial range, here indicated by means of the grey circle. The IS allows several representations of the interactions, but in many cases it is the molecular interaction between epitopes (strings of symbols in (f)) that truly represents the underlying liquid brain dynamics.

links becoming the edges (figure 1d). Two types of nodes are shown, open and closed, associated to inactive or active neurons, respectively. This is akin to simple interactive models between active (on) and inactive (off) neurons. In this approximation, one treats these agents as small magnets the orientation of which depends on their corresponding neighbours via some non-trivial interaction rules. In general, physicists dub these kinds of systems a spin-glass [12]. This concept lies at the core of what follows here.

Ant colonies also involve collectives of interacting individuals that can be modeled as on/off activity - e.g. engaged in a task or not. However, the physical location of such mobile agents change over time: the colony is 'liquid'. Thus, interactions are now extended to a spatiotemporal embedding, with agents interacting at a certain location in one instant and then moving on to other locations to continue interacting with each other. This particularity will constrain the system's adaptive properties in a non-trivial manner. Within the liquid realm, we can still characterize two paradigms: given their relative mobility and signal transmissivity (see below) insect colonies are strongly affected by the locality of their interactions, whereas ISs are highly mobile, such that a well-mixed approach might accurately represent their overall dynamics.

Other types of organisms, such as the slime mould *Physarum*, solve some classes of optimization problems by using a different form of fluid organization [13], although in this case there is no neural substrate. Systems of this kind are able to solve minimization problems on a network [14].

Upon the transition to multicellularity, cell types capable of sensing and responding to signals appeared and permitted the emergence of a novel class of systems: webs of connected cells.

These expanded the landscape of computations, including processing the information in non-trivial (aneural) ways [2]. Such primitive networks provided a reliable way of dealing with complex decisions, integrating and storing memory and creating the conditions for increasing behavioural complexity. Simple organisms such as hydra and planarian flatworms provide good illustrations of the early steps in this direction [7]. To some extent, all these systems can be modelled as networks of neurons that are connected in a stable way over time. Each pair of connected cells will remain linked over a given time scale and changes will take place at the level of the type and strength of the connection. Theoretical work has shown that cognitive tasks performed by these 'solid' brains (simple and complex) such as pattern recognition, associative memory or language processing can be properly described. But what about liquid brains?

In this paper, we review several models of both ant colony and IS dynamics based on a neural network perspective and compare them with previous studies on 'solid' brain models. In table 1, we summarize some general qualitative properties of the three classes of systems explored here, as well as others that we found relevant. The list is not exhaustive and involves generic descriptors that inevitably ignore the broad diversity of sizes, organization levels and ecological contexts. Several key components of each potential candidate, including size, age, context or developmental trajectories have some influence in the degree of robustness, memory potential or wiring patterns. All these factors make this basic table a tentative one. Nevertheless, it also highlights the commonalities that we consider relevant to our presentation.

Some key examples are worth mentioning. The label 'liquid' is used to describe a physical state that ignores spatial

Table 1. Comparative properties in liquid versus solid brains. This table summarizes a broad set of properties that are usually attributed to neural systems (solid brains) and here compared to those reported from two relevant examples of liquid brains, namely the IS and insect (mostly ant) colonies. While the way computations are performed is a parallel process in all systems, all also exhibit some degree of specialization, which can be understood as a modularity or a division of labour (DOL). This first is observable in vertebrate brains while the later is a characteristic allocation of tasks that can occur either in societies with different morphological castes or in monomorphic ones. Similarly, we label the learning and memory properties in terms of a simple, network-related set of properties. In most cases studied here, the memory potential of an ant colony is related to short-term phenomena tied to the production of a pheromone field, but long-term memories have also been reported at the individual level. In all these examples, we indicate by asterisk (*) those attributes that are not well established or have been found in some case studies, and that will benefit from a theory of liquid brains.

	brain	insect colonies	IS
computational nature	distributed/modular	distributed/DOL	distributed/DOL
reliability under agent loss	high	high	high
connectivity	hard-wired yet plastic	liquid but spatially constrained (*)	liquid (*)
memory and learning	synaptic	population-based and synaptic	population-based (Burnet) and synaptic (Jerne)
regeneration potential	low	high	high
externalities	peripheral nervous system and technology (tools)	nests and agriculture	no
weighted interactions	Hebbian	pheromone-mediated	antibody-mediated
dynamical state	critical	critical (*)	critical (*)

structuring, such as lymph nodes in the IS or the nest structure of ant or termite colonies. Some of these features cannot be taken as absolute indicators since they are strongly influenced by life styles, size or behavioural context. The neural network of a hydra or a planarian flatworm are simple and small and might not display the modularity found in more complex neural agents, but nevertheless they display spatially stable networks of neurons, which are reliable under cell loss. Other relevant features (which are not included in table 1) such as the self/non-self discrimination problem will be amply discussed later on. In the following sections, we summarize several types of models used to represent and understand the dynamics of the three case studies discussed here. By using them, we aim at enhancing the universal elements shared by these liquid systems while tracing a theoretical framework to study them.

2. Solid brains

Standard neural networks (NN), from cell cultures to brains, have received great attention since the 1950s. An especially successful approach has been based on the use of statistical physics as a robust formalism capable of capturing the collective properties exhibited by neural masses [15]. Both in statistical physics as well as in logic models of NN, neurons are replaced by a toy model representing only the minimal features exhibited by real cells. The intricate structure of physiological neurons is ignored and replaced by a formal object devoid of any specific traits associated to cellular or molecular biological mechanisms. Similarly, the way connections and propagation of activity occurs is mapped into a simple graph. Despite all these oversimplifications, NN theory (also known as connectionism) has been capable of explaining the nature and relevance of collective phenomena involved in a broad range of areas, from learning in small metazoans to more complex phenomena related to human cognition [16,17].

We use here the term ‘brains’ in a generic way too: it will refer to ensembles of interconnected neurons (or neural-like

elements). Over the years the field has been growing in multiple directions, but a special turning point is the classical paper by Hopfield [18] where the basis for a statistical physics description of neural networks emerged and largely marked the development of this class of systems. Such a physics perspective provided the basis for the understanding of their global properties out from the underlying microscopic description. Importantly, it also provided a systematic approach to identify the presence of different phases associated with the presence or lack of memory as well as dynamical states separating different types of activity. In this way, the physics of phase transitions [19–21] became a cornerstone to our understanding of neural networks.

The simplest, canonical model is based on an assembly of two-state agents description [22,23]. These are denoted as $S_i(t) \in \{0, 1\}$ or $S_i(t) \in \{-1, 1\}$ (with $i = 1, \dots, N$). Agents are connected to each other through fixed synaptic links (figure 1a): each element sends to and receives a signal from another. Connectivity is represented by a matrix $J_{ik} \in \mathbb{R}$. The system is modelled by a dynamical set of equations:

$$S_i(t+1) = \Theta \left(\sum_{j=1}^N J_{ij} S_j(t) - \theta_i \right), \quad (2.1)$$

where $\Theta(z) = 1$ for $z > 0$ and zero otherwise. The scalar θ_i is a threshold value. The so called external field, $h_i = \sum_j J_{ij} S_j(t)$, weights the total input of S_i . It is worth noting that the same class of threshold model used to describe the dynamics of NN has been used to approach the dynamics of gene regulatory networks (GRN) [24, pp. 411–521, 25,26].

(a) Attractor dynamics in recurrent neural networks

A general treatment of these systems involves a high-dimensional problem and a wide range of dynamical behaviours. However, an illustration of the potential of NN as a way of solving computational problems in a distributed manner is provided by the Hopfield model [18,27]. This consists of a fully connected neural network described by the dynamical

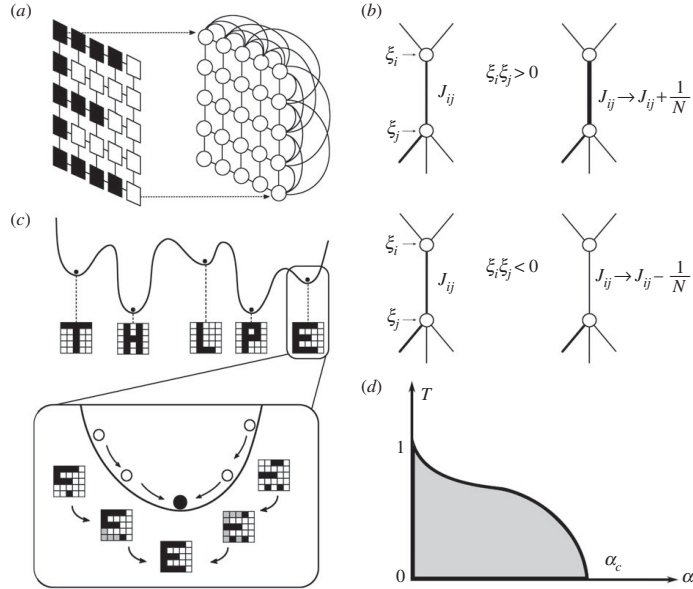


Figure 2. Distributed computation in neural networks. Using a very simple set of rules, an NN model can store and retrieve memories in a robust manner. In the Hopfield's model, a massively connected set of neurons (a) with symmetric connections obeying Hebb's rule (b) will display such properties. In (b), a pair of formal neurons is shown receiving inputs $\xi_i, \xi_j \in \{-1, 1\}$ from a given memory state or pattern ξ^μ . If they are identical, i.e. $\xi_i = \xi_j$, their connection is increased (in both directions). Otherwise, J_{ij} is decreased. (c) Network dynamics makes the system's state flow to energy minima, thus recovering the desired memory state. The model exhibits remarkable reliability against connection loss. In (d), we show how reliable is memory retrieval against stochastic thermal variability. Parameter α is a relative measure of memory capacity. The critical value $\alpha_c \simeq 0.138$ separates the two phases: memory reliability (shaded area) and unreliability (blank area). This transition occurs sharply. Note that this critical value is specific for Hopfield nets; different interaction rules would yield different limitations to memory capacity.

equations (2.1) with $\theta_i = 0$. Hopfield's model assumes no self-connection ($J_{ii} = 0$) and symmetry, i.e. $J_{ij} = J_{ji}$. It can be shown that the model only displays single-point equilibrium (attractors), i.e. asymptotically, the trained network will tend to a stable configuration where all elements remain in a given state (figure 2a–c). Additionally, Hopfield's model allows the network to store a number p of 'memories' (patterns) defined as a set of vectors $\xi_\mu = (\xi_1^\mu, \dots, \xi_i^\mu, \dots, \xi_N^\mu)$, $\mu = 1 \dots p$. The storage process takes place within a 'training phase' where they are presented to the network in such a way that each neuron S_i adopts the memory state, i.e. $S_i = \xi_i$ and all its synaptic weights J_{ij} are updated (starting from $J_{ij} = 0$ at time zero) following the so-called Hebb's rule, which is summarized in figure 2b. In a nutshell, correlated inputs increase weights, whereas uncorrelated ones decrease them. It can be shown [28] that the memory states ξ_μ are, in fact, the minima of a (high-dimensional) energy function, namely:

$$\mathcal{H}((\xi_i^t)) = -\frac{1}{2} \sum_{ij} J_{ij} S_i S_j, \quad J_{ij} = \frac{1}{N} \sum_{\mu=0}^p \xi_i^\mu \xi_j^\mu \quad (2.2)$$

and initial conditions close to a minimum will evolve towards it. This is also outlined in figure 2c where we represent such multiple minima. In summary, the Hopfield model is a dynamical process of memory retrieval: stored patterns are recovered by a purely dynamical process. Extensions to this approach come by introducing thermal noise for the $\{S_i\}$ degrees of freedom. Usually, this is obtained via a temperature T that accounts for stochastic thermal variations (and, more generally, for noise). Each time we choose a neuron, the

probability of changing to (or remaining in) state $S_i = +1$ is a saturating function of the local field, namely:

$$P[S_i(t+1) = +1 | h_i(t)] = \phi\left(\frac{1}{T} \sum_j J_{ij} S_j(t)\right) \quad (2.3)$$

with T defining a temperature and $\phi(x)$ a function such that $\phi(0) = 0$ and $\phi(x) \rightarrow \pm 1$ for $x \rightarrow \pm \infty$. Temperature is not just an additional attribute, as it actually provides a powerful mechanism to escape from local minima. By using a stochastic transition rule, it is possible to move to lower-energy states from a given, suboptimal (usually non-memory) state. In this context, a measure of memory capacity is introduced as $\alpha \equiv p/N$, where p here corresponds to the number of well-stored patterns. A phase-transition diagram captures the overall system behaviour, depicted in figure 2d. The shaded region represents states where the system is capable of retaining the memory patterns, while, for the blank region, these are lost due to noise. An abrupt transition separates these two regimes.

The previous model is an illustration of how cognitive functions can be understood in terms of a system of connected neurons. Here synaptic weights are modified in such a way that the resulting attractor dynamics allows associative memory to be the consequence of a relaxation towards energy minima. Only steady states are thus allowed. However, as discussed in the next section, a different picture emerges when we look at the actual dynamical patterns exhibited by neural tissues.

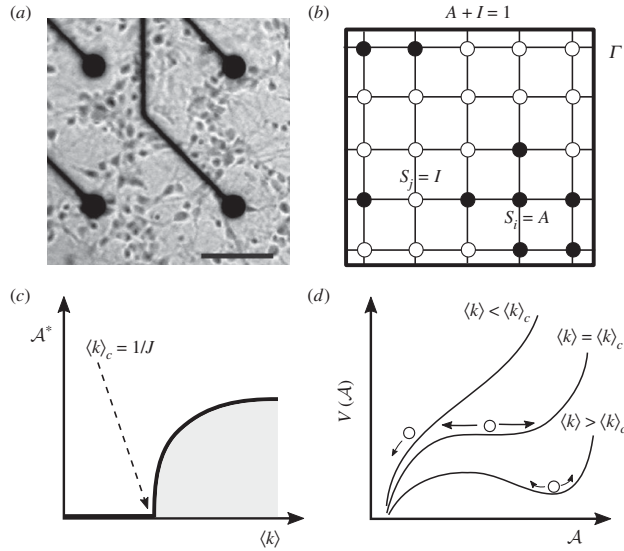


Figure 3. Phase transitions in neural dynamics. In a simple version of large-scale dynamics of neural tissues, (a) tissues (such as brain cortex) can be represented as a network of connected neighbouring areas that are connected with excitatory links (adapted from Eckmann *et al.* [34]). A toy model of this (b) could be represented as a lattice of neural elements connected as a grid with all elements linked to four elements in a homogeneous fashion. The analysis of this system reveals a phase transition from zero activity to high-activity by crossing a critical value of average connections at $\langle k \rangle_c = 1/J$ (c). A potential function can be obtained where the two phases are revealed as stable states of $V(A)$ (d). Here, large fluctuations show clear dominance around the critical point.

(b) Critical dynamics in cortical networks

If we think of an idealized graph such as the one described in figure 1d, two classes of nodes can be defined: either inactive or active. Active nodes are formed by firing neurons whose excitability can be propagated to the nearest inactive areas [29]. As a result, excitation waves can move across whole areas. This would be a requirement to maintain integration in a dynamical fashion [30]. This analysis gives rise to more complex types of attractors instead of local stable points.

The minimal model that can describe the propagation or activity is based on a contagion scenario where inactive nodes can become active if they are connected to active nodes. Moreover, an active node can spontaneously decay. At the smallest scale, this is similar to the threshold dynamics described above. The simplest case to consider is a homogeneous model where all connections are similar, capable of propagating excitability with $J_{ij} = J$ and an average connectivity $\langle k \rangle$. It can be shown that the large-scale (coarse-grained) dynamics for this homogeneous case can be defined by the equation [29]

$$\frac{dA}{dt} = f(A) = -\frac{1}{\tau}A + \frac{J(k)}{\tau}A(1-A), \quad (2.4)$$

where τ is a characteristic time decay. A specially relevant observation is that neural systems exhibit critical behaviour [31–33]. Two main classes of dynamical behaviour can occur. This can be shown using the fixed points, i.e. those A^* such that $(dA/dt)_{A^*} = 0$. Two states are obtained. One is the trivial, inactive phase where no activity propagates: $A_0^* = 0$. The second phase is associated with the second fixed point, namely:

$$A_1^* = 1 - \frac{1}{J(k)}, \quad (2.5)$$

which is properly defined (i.e. $A_1^* \geq 0$) provided that $J(k) \geq 1$. A critical point separating the two phases is thus achieved for $J(k) = 1$. For a given J value, the critical connectivity is given by $\langle k \rangle_c = 1/J$.

In figure 3, two important diagrams are shown that summarize the basic phenomena resulting from the previous model. One is the so-called bifurcation diagram [35] where the stable states A_0^* , A_1^* are plotted against the average connectivity $\langle k \rangle$, with a marked change occurring at criticality (figure 3c). Additionally, we also display (figure 3d) the potential function $V(A)$ [36], defined as

$$V(A) = -\int f(A) dA, \quad (2.6)$$

such that the dynamics derive from it, i.e. $dA/dt = -dV(A)/dA$. The minima (maxima) of the potential correspond to stable (unstable) fixed points. As we approach criticality, the potential function becomes increasingly flatter. What is the impact of this flatness in the activity? In general, shallow potentials are associated with higher time variability and fluctuations diverge close to criticality. To show this, we can use a linear stability analysis taking the state $A(t) = A_k^* + \delta A$, i.e. a small deviation δA from a fixed point A_k^* , and plugging it into the original equation for $A(t)$. On a first approximation, it can be shown that

$$\frac{d\delta A}{dt} = \left(\frac{\partial f(A)}{\partial A} \right)_{A_k^*} \delta A = \lambda_k^* \delta A, \quad (2.7)$$

where λ_k^* is a scalar to be evaluated at each fixed point. The resulting equation for fluctuations is linear. Thus, close to A_k^* , we expect a growth of fluctuations following an exponential

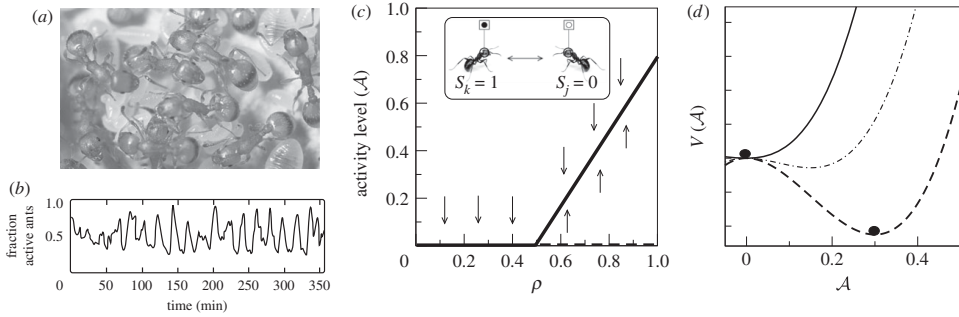


Figure 4. Ant colonies as excitable neural nets. In some ant species, such as those belonging to the genus *Leptothorax* (a), oscillations in activity have been recorded (b) revealing a collective synchronization phenomenon (both adapted from Solé [36]). This phenomenon can be described as an excitable neural system, where ants (inset of c) are reduced to a Boolean representation with active and inactive individuals. (c) As the density of ants ρ increases, a phase change occurs at a critical density, separating inactive from active colonies. (d) Potential function associated with the dynamics of these colonies: for densities larger (lower) than ρ_c , a well-defined minimum is displayed. Closer to criticality, this potential becomes flatter and allows wide fluctuations to occur.

growth or decay. For the inactive phase,¹ we have

$$\delta A(t) = (\delta A)_0 e^{\lambda_0 t} = (\delta A)_0 \exp\left(\frac{-1}{\tau}(1 - J\langle k \rangle)t\right). \quad (2.8)$$

As we can see, the system will return to the fixed point (when $\langle k \rangle < 1/J$) at a rate given by λ_0 . As we get close to criticality, the exponent gets smaller, the relaxation time rapidly increases. If the previous result is written in terms of a relaxation time $T(J, \langle k \rangle)$, i.e. $\delta A(t) \sim \exp(-t/T(J, \langle k \rangle))$ we have

$$T(J, \langle k \rangle) \sim \frac{1}{1 - J\langle k \rangle}, \quad (2.9)$$

which rapidly diverges as $J\langle k \rangle \rightarrow 1$. The divergence predicted by this simple model is confirmed by the analysis of the fluctuations found in neural systems.

The two previous models explore some essential components of neural complexity. Both deal with collective behaviour and exhibit special regions of parameter spaces that separate different phases. Phase transitions are of central importance within statistical physics, and provide a powerful framework to capture how microscopic interactions translate into system-level patterns and processes [36,37]. Their importance within our context becomes manifest as qualitative changes in collective behaviour are typically caused by phase transition phenomena often associated with the density of individuals or the signals they use to communicate. How these systems behave close to transition points turns out to be a key issue, as it provides understanding about how emergent phenomena occur.

3. Liquid brains

(a) Ant colony dynamics

Social insects, including ants and termites among other groups, amount to about the same biomass as humans on the Earth [38]. With an evolutionary history spanning around a hundred million years, eusocial colonies have deeply engineered the environment and dominated the terrestrial biosphere [39]. In trying to attach biological fitness, insect colonies appear to behave as superorganisms: it is the colony as a whole that plays an evolutionary role, rather than its individual agents (ants). Across the biosphere, we encounter both monomorphic and polymorphic ant colonies, the latter exhibiting

physiological–anatomical differences within a given colony. However, it is estimated that 80% of ant species are monomorphic. The rest of species (polymorphic) can generally include two or three different castes. Here onwards, we will focus our study on monomorphic ant species.

On the other hand, various estimates state that the behavioural repertoire of ant colonies ranges from 20 to 45 different individual-ant behaviours [8, pp. 180–200]. In order to shift from a given state to another and adapt to any given environmental circumstance, ants use chemical signals called pheromones. Different ant species use different sets of pheromones, some secrete only one type of molecule and others use up to 20.² Thus, information is processed in a two-level fashion: mobile agents (ants) interacting with a set of diffusive molecules (pheromones). Ants continuously detect the pheromone concentrations and, upon integrating this information, produce an internal image that affects their behavioural state. Moreover, ant states prompt the secretion of one (or more) pheromones thus reshaping their concentration values. This coalescence of signalling back and forth allows the whole colony to access global states where functions are achieved by means of its underlying network of interactions. Information is stored and processed through this ‘liquid brain’ to give rise to various large-scale collective behaviours. In the following examples, we will review several theoretical approaches to modelling ant colony dynamics and compare them with standard NN model efforts.

(i) Ant colonies as excitable neural nets

One of the simplest illustrations of the neural-like modelling of insect colony dynamics is provided by the emergent synchronization displayed by some small colonies of the genus *Leptothorax* ([40,41]; see also [9], ch. 6). In a nutshell, it has been observed that the colony-level activity displayed by their nests exhibits a remarkable bursting pattern (figure 4a,b) that exhibits a periodic component [42]. This means that ants can be active or inactive and the total number of active individuals changes in such a way that at times no ant in the colony is active while the synchronization events are linked to an almost fully active colony. These bursts have been found in other species [43] and result from the propagation of activity carried by moving individuals that can activate dormant ants in ways similar to those

found in epidemic models [36,44,45]. Synchronization of neural masses is in fact a major research field within neuroscience [46], and it has been shown to pervade a wide range of functional traits and behavioural patterns. Is there something similar taking place in ant colonies?

This problem provides a simple example of a fluid network. Here the description level of individuals and their interactions is limited to a Boolean set of variables $\Sigma = \{0, 1\}$ associated with the inactive (motionless) and active (moving) states, respectively (see inset of figure 4c). An NN model here is thus limited to a coarse-grained representation of ants. Such a model was suggested in Solé *et al.* [40] under the assumption that individuals can be described as an underlying continuous variable $S_i \in [0, 1]$ (with $i = 1, \dots, N$) which changes over time following a dynamical equation:

$$S_i(t+1) = \theta \left(\sum_{j \in \Gamma_i} J(\eta_j, \eta_i) S_j(t) \right), \quad (3.1)$$

which strongly resembles the familiar form of standard NN. However, a rapid inspection reveals a fundamental difference: here the matrix $J(\eta_j, \eta_i)$ is state-dependent. In other words, its value is a function of the specific pair of agents that interact at a given time step. Specifically, we partition the activity interval $[0, 1]$ into two domains associated with the active/inactive observables, i.e. $\eta_i = \theta [S_i - \theta]$. Thus, the interaction matrix will include only four possible pairs,

$$J(\eta_j, \eta_i) = \begin{bmatrix} J_{00} & J_{01} \\ J_{10} & J_{11} \end{bmatrix},$$

where $J \geq 0$. Once activity decreases below the threshold θ , the ant becomes inactive and stops moving. Otherwise, it moves around as a random walker (unless constrained by other ants occupying the nearest lattice sites). Here ants are assumed to move on a discrete two-dimensional lattice Ω and interactions occur in a strictly local manner, only affecting the set of nearest neighbouring positions Γ_i of S_i . Finally, an inactive ant (with $\eta < \theta$) can become active spontaneously (achieving a state $S_0 > \theta$) with probability p_a . A common feature of these matrices is the presence of coupling terms connecting active and inactive individuals, as expected from an excitable system where activity can be propagated among agents. It is important to notice that the collective synchronization does *not* result from the coupling of individuals' internal clock. Instead, single virtual ants behave randomly. The dynamics of single elements will be described by

$$S_i(t+1) = \theta (gJ_{11}S_i(t)). \quad (3.2)$$

A simple case can be solved, namely when the coupling is small and activity remains small (which is consistent with observation). If we choose $\theta(x) = \tanh x$, then we may use linear approximation $\tanh(gz) \approx gz$ which admits a solution to the previous equation. If, initially, an ant is activated to a level S_0 , then $S(t) = S_0(gJ)^t$, which is a decaying function of time. If an activation term is also introduced (i.e. active ants can activate inactive ones), then a coarse-grained model can be defined in probabilistic terms. Let us label as N_a the number of active ants. This number will change over time as a consequence of both interactions and decay. The efficiency of activation events will be proportional to gJ , assuming the

previous linear approximation. Hereafter, we will indicate by N and ρ the total number and density of ants, respectively.

If $\mathcal{A}(\mathbf{x}, t)$ indicates the probability density of active ants at a given point of our two-dimensional lattice $\mathbf{x} \in \Omega$, then it can be written as: $\mathcal{A}(\mathbf{x}, t) = P[S_{\mathbf{x}}(t) = 1]$. The activity density will evolve following a master equation according to the previous rules:

$$\frac{d\mathcal{A}(\mathbf{x}, t)}{dt} = gJ \sum_{\langle \mathbf{u} \rangle} P[S_{\mathbf{x}}(t) = 0 \cap S_{\mathbf{u}}(t) = 1] - \alpha \mathcal{A}(\mathbf{x}, t), \quad (3.3)$$

where $\langle \mathbf{u} \rangle$ indicates sum over the set of q nearest neighbours, $P[S_{\mathbf{x}} = 0 \cap S_{\mathbf{u}} = 1]$ is the probability of having a pair of nearest ants in different states.

The previous equation is exact, but its computation would require knowledge of the probabilities associated with the interactions between nearest sites. Several methods can be used to solve this model with different levels of approximation. Here we will consider the simplest one, commonly known as a mean field theory, which is based on suppressing the spatial correlation between nearest sites. This is done by assuming that the system is in fact well mixed and thus all sites are neighbours or, in mathematical terms, $q = \Omega$. If this is the case, we can use the total population

$$\rho(t) = \sum_{\mathbf{x}} \rho(\mathbf{x}, t). \quad (3.4)$$

By summing on both sides of the previous master equation, and ignoring correlations between active and inactive neighbours, it can be shown that the global dynamics can be described as

$$\frac{d\mathcal{A}}{dt} = gJ\mathcal{A}(\rho - \mathcal{A}) - \alpha\mathcal{A}, \quad (3.5)$$

and this equation can be studied as a deterministic model of ant colonies displaying excitable dynamics. The model has two equilibrium points, namely $\mathcal{A}_0^* = 0$ (no activity spreads) and $\mathcal{A}_1^* = \rho - \alpha/gJ$, associated with persistent propagation. The previous equation is similar to those used in epidemic dynamics [47] associated with a population composed by two classes of individuals (infected and susceptible). Using the density of ants as a control parameter, these two phases are separated by a critical point $\rho_c = \alpha/gJ$. The global behaviour of this model is summarized in figure 4c where the bifurcation diagram for this system is shown. Above ρ_c an active phase is present, whereas an inactive one is found for $\rho < \rho_c$.

In this system, the potential function $V(\mathcal{A})$ is

$$V(\mathcal{A}) = - \int (gJ\mathcal{A}(\rho - \mathcal{A}) - \alpha\mathcal{A}) d\mathcal{A} \quad (3.6)$$

and is displayed in figure 4d, where we show three examples of its behaviour for different density values. As we already discussed within the context of brain criticality, here too the transition between phases as density is changed involves a shallow potential function, indicating that wide fluctuations should be expected to occur. One remarkable observation from *Leptothorax* colonies is that they seem to be poised close to the critical density [48] at density levels where theory predicts that maximum information and behavioural diversity is achieved [49,50]. As discussed above within the context of neural tissues, criticality provides a source of fast response and optimal information processing.

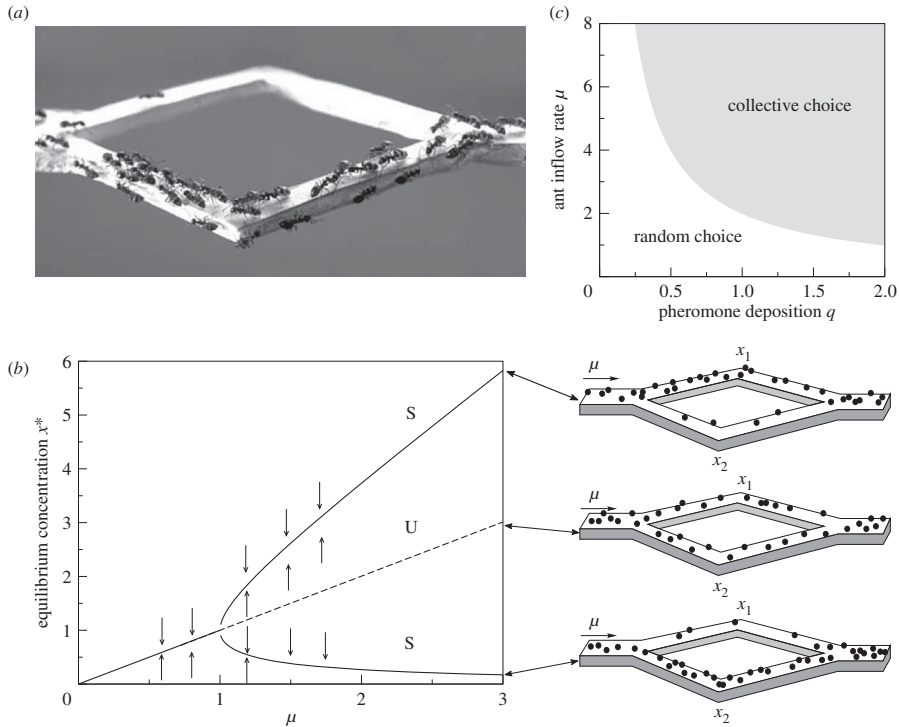


Figure 5. Collective decision-making. (a) A two-path experiment allows testing of the mechanisms by which emergent decision-making occurs. The photograph shows an example of a colony that has made a collective decision, as shown by the preferential use of the shortest path. (b) The mathematical analysis of the model associated with this phenomenon shows that two alternative solutions exist associated with the preferential choice of one branch, along with a third one where both branches are used. (c) The parameter space for the simple symmetric case.

The key message provided by this example is that a commonality with other excitable neural systems exists: a universal property is the use of critical points to perform cognitive tasks. Being poised close to critical states provides a natural way of amplifying input signals while remaining most of the time in a low-fluctuation state [51]. Such a compromise makes sense as a way of displaying optimal information while reducing the cost of the system's state. Is there a well-defined function that can be associated with this? The answer is yes. By using self-synchronized patterns of activity a task may be fulfilled more effectively than with non-synchronized activity, at the same average level of activity per individual [52,53].

(ii) Collective decision-making and symmetry breaking in ant colonies

The next case study involves a classic example of how fluid brains solve a well-defined optimization problem. Specifically, a given ant colony searching on a given spatial landscape needs to discriminate between different available food sources [54–56]. The decision-making rules involved here have inspired major applications within computer science [57–60].

Consider the determination of the shortest path to a single source, to be chosen between two alternatives. This problem can be easily implemented in the laboratory, using a two-bridge set-up (figure 5a). Here the ant nest would be located to the left side and ants would walk through the two-bridge

array to reach a food source located on the right side. The two branches can be identical or instead have different lengths. The problem to be solved here is which one is the shortest. Moreover, if one is chosen from a symmetric case (equal paths) how is this choice made. Once again, the solution cannot be found at the individual level: colony-level processes need to be in place to make the right decision.

This problem can be understood in terms of statistical mechanics where ants are described as Boolean, spin-like variables and the pheromone acts as a field [61]. Here the alternative paths are mathematically mapped to a ferromagnet and the spin-pheromone coupling is described accordingly. In order to introduce noise, explicit use of a temperature is made. In order to present the key ideas under a macroscopic picture, let us here explore a simplified version that can be easily explored analytically.

Ants can use quorum-sensing mechanisms as a way of creating and (responding to) pheromone fields thus generating a large-scale chemical field that allows them to properly make their decision. Initially, ants will walk on both bridges, choosing randomly their branch. We should expect at this point an equal number of ants on each branch, i.e. $\rho_1 = \rho_2$. However, once an ant has found the food source, it releases a pheromone as it returns to the nest. Other ants will detect the released signal, which helps ants to decide where to move, releasing further pheromones and amplifying the previous mark. The pheromone trail also evaporates, and evaporation will be

more effective in the longer trail, where more surface is available. As a result, the shortest path is more likely to be used, and is eventually chosen. Ants have computed the shortest path. A model describing this experiment can be defined as follows. If ρ_1 and ρ_2 indicate the concentrations of trail pheromone in each branch, their dynamics [62] is given by a pair of equations for the pheromone fields:

$$\frac{d\rho_k}{dt} = \mu q_k P_k(\rho_1, \rho_2) - \nu \rho_k, \quad (3.7)$$

with $k = 1, 2$. Here μ is the rate of ants entering each branch, q_i the rate of pheromone production at the i th branch and ν is the rate of evaporation. The functions $P_i(\rho_1, \rho_2)$ can now be understood as probabilities of choosing a bridge depending on the pheromone concentrations. These probabilities are well described by a nonlinear, threshold response function [63,64]:

$$P_i(\rho_1, \rho_2) = \frac{(\rho_i + K)^2}{\Theta(\rho_1, \rho_2)}, \quad (3.8)$$

where $\Theta(\rho_1, \rho_2) = \sum_{j=1,2} (\rho_j + K)^2$ and $i = 1, 2$. The parameter K gives the likelihood of choosing a path free of pheromones ($\rho_i = 0$).

This is a general model that incorporates attributes associated with each branch. But an interesting scenario arises when one considers the symmetric case where $q_1 = q_2 = q$. For this situation, the previous set of equations reduces to

$$\frac{d\rho_i}{dt} = \mu q \frac{(\rho_i + K)^2}{\Theta(\rho_1, \rho_2)} - \nu \rho_i. \quad (3.9)$$

Here there is no true optimal choice: both branches are equal. Now, although the obvious expectation is a similar distribution of ants in each branch, this is not what is observed. We would easily conclude that ants would choose both paths and that individuals will equally walk in both branches. However, what is typically seen is that the symmetry is broken in favour of one of the two branches. Why is this the case? This phenomenon illustrates a very important class of phase transition: the so called symmetry breaking process. Despite the symmetry of the system, amplification of initial fluctuations leads to the formation of a dominant pheromone trail that is used by all ants once established.

The fixed points associated with this system are obtained from $d\rho_i/dt = 0$. One possible solution to this system is the symmetric state $\rho_1^* = \rho_2^* = \rho^*$ (associated with equal use of both branches), ants equally distributed in both branches. For this special case, we have $P_i(\rho_1^*, \rho_2^*) = P(\rho^*) = \mu q/2$ and thus we only need to solve a single equation $d\rho^*/dt = \mu q/2 - \nu \rho^*$, which gives a fixed point $\rho^* = \mu q/(2\nu)$. This is the symmetric state to be broken. The second scenario corresponds to the choice of one of the branches ($\rho_1^* \neq \rho_2^*$). Since $\rho_1 + \rho_2 = 2\rho^* = \mu q/\nu$, we see that

$$\left(\frac{\mu q}{\nu} - \rho_i^*\right)(\rho_i^* + K)^2 = \rho_i^* \left(\frac{\mu q}{\nu} - \rho_i^* + K\right)^2, \quad (3.10)$$

after some algebra, this gives the new fixed points $\rho_{\pm}^* = (\rho_{1\pm}^*, \rho_{2\pm}^*)$ and $\rho_{\pm}^* = (\rho_{1\mp}^*, \rho_{2\pm}^*)$ with

$$\rho_{i\pm}^* = \frac{\mu q}{2\nu} + \left[\left(\frac{\mu q}{2\nu}\right)^2 \pm K^2 \right]^{1/2}. \quad (3.11)$$

This pair of fixed points will exist provided that $\mu q/2\nu > K$, which allows the derivation of a critical line (figure 5c)

$$\mu_c = \frac{2K\nu}{q}, \quad (3.12)$$

indicating that there is a minimal rate of ants entering the bridges required to observe the symmetry breaking phenomena. For $\mu > \mu_c$, the symmetric state becomes unstable (figure 5b,c) while the two other solutions can be equally likely. Below this value, the only fixed point is the symmetric case with identical flows of ants in each branch. This symmetric model can be generalized to (more interesting) asymmetric scenarios where the two potential choices are different (see [55] and references therein) either because the food sources are of different sizes or because the paths are of different lengths and the shortest path needs to be chosen. This symmetry breaking phenomenon has also been observed in the ant colony panic responses [65], army ant trails [54] and optimal group formation [66]. A specially interesting proposal concerning the phenomenon of symmetry breaking in ants was made in Bonabeau [67], where it was suggested that flexible behaviour leading to efficient decisions is more likely to occur close to critical points.

(iii) Task allocation in ant colonies as a parallel distributed process

In the previous example, we considered a set of agents described as binary variables, thus ignoring the combinatorial complexity that should be expected from an insect equipped with a brain. Moreover, it is clear that the active/inactive dichotomy hides a repertoire of potential activities that can be carried out by individuals, associated with the set of tasks needed to maintain the colony. DOL is in fact one of the most important and widespread phenomenon in nature, and very common in social groups [68]. It has been shown that the dynamics of subsets of individuals performing specific tasks within colonies is an emergent phenomenon [10]. In this scenario, a colony that needs to perform a given set of tasks under given environmental conditions (and respond to changes in flexible ways) must be capable of sensing its internal state using some kind of distributed information processing.

Inspired by the dynamics of harvester ants, Gordon *et al.* [69] proposed a neural network model of task allocation where individual ants are represented by a sequence of Boolean variables instead of a single ON-OFF description. Observations from extensive field work on harvester ants (*Pogonomyrmex*) show that members of an ant colony perform a variety tasks outside the nest, such as foraging and nest maintenance work. Remarkably, this is a monomorphic species, i.e. individuals exhibit identical phenotypes. The number of ants actively performing each task changes over time due to task switching as well as the presence of inactive workers [70]. As discussed in Gordon [11], interactions among ants involve physical contact. This allows sensing the state of other nest-mates to create a network of information exchanges. Experimental perturbation of the number of ants performing a given task triggers changes in the numbers of individuals performing other tasks. Importantly, this switching dynamics is a consequence of the microscopic, local ant-ant interactions. The attractors associated with normal and perturbed conditions are thus a collective-level outcome of individual interactions.

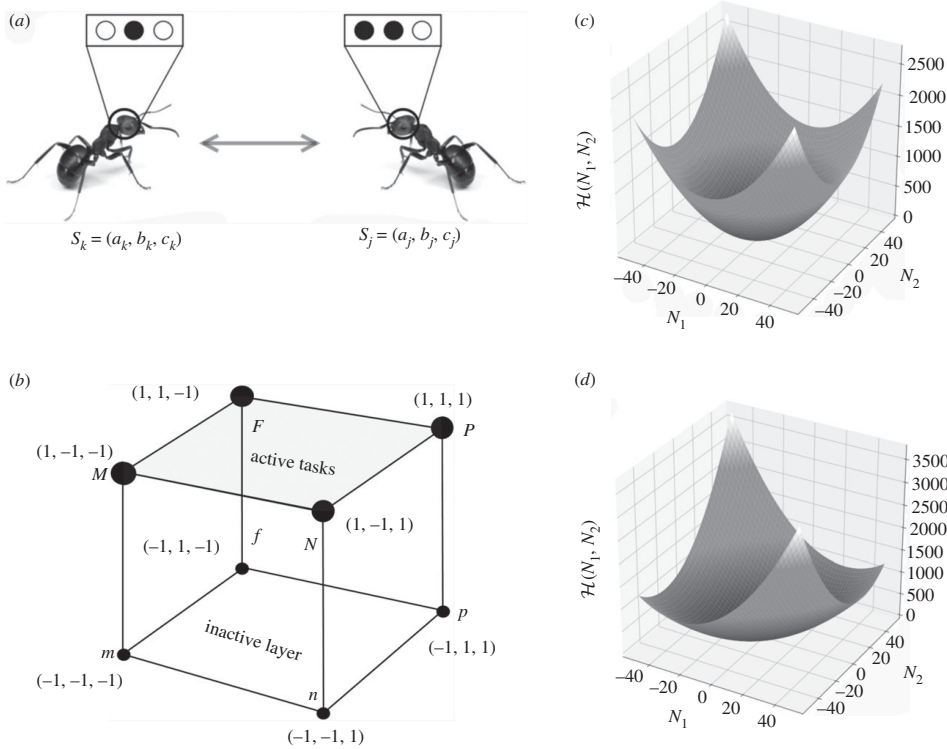


Figure 6. Neural network model of task allocation in ant colonies. The dynamics of harvester ants in Gordon *et al.* [69] can be described in terms of virtual ants (a) each carrying a 3-spin internal description, with changes taking place by means of direct pairwise interactions. The total state space is a three-dimensional Boolean cube (b) where we indicate active (observable) tasks in the top of the cube while a lower layer of inactive states is formed by a flip in the first spin (negative for inactive ants). The model exhibits an attractor dynamics with an associated potential (energy) function. (c,d) The potential function is easily found for a two-task system for the specific values of parameters $\alpha = 1$ and $\beta = 0.1$ (c) and $\beta = 0.5$ (d).

In their model, Gordon and co-workers consider a set of four observed tasks, namely: patrollers, foragers, nest maintenance and midden workers, displayed by harvester ants. Additionally, individuals can become inactive (as reported in ant colonies, see previous section). Since each type of ant performing any of the four tasks can become inactive, the model assumes that eight possible vectors can represent the available space state which can be covered by an internal state of three binary variables [69]. Specifically, ants are described now as 3-spin vectors $S_k = (S_k^1, S_k^2, S_k^3)$. In their original paper, they use the notation $P =$ active patroller, $F =$ active forager, $N =$ active nest maintenance worker and $M =$ active midden worker. The lower case versions (p, f, n, m) would indicate inactive versions of the previous vectors. The space of possible internal states is indicated in figure 6a,b. These are represented as vertices of a Boolean cube, where all states are, respectively, indicated as strings of +1 and -1 values.

The simplest approach for this problem is to assume that the different components of the internal state act independently, with different associated weight matrices. In this way, we would have

$$S_j^\mu(t+1) = \Theta(h_j^\mu(t)) = \Theta\left(\sum_k J_{jk}^\mu S_k^\mu(t)\right). \quad (3.13)$$

The (internal) state of S_j^μ will remain stable after one interaction, provided that $S_j^\mu h_j^\mu > 0$. An energy function is defined accordingly as follows:

$$\mathcal{H}((S_k^\mu, J_{ij}^\mu)) = -\frac{1}{2} \sum_\mu \sum_{ij} J_{ij}^\mu S_i^\mu S_j^\mu. \quad (3.14)$$

In the macroscopic realm, the observable state is the number of ants performing each task from the repertoire. It is then desirable to have a description where the energy minimization is defined in terms of the set $\{n_k\}$. Thus, the energy function now reads as

$$\mathcal{H}(\{n_k, \Gamma_{ij}\}) = -\frac{1}{2} \sum_{ij} \Gamma_{ij} n_i n_j, \quad (3.15)$$

with a new set of parameters $\{\Gamma_{ij}\}$ that depend on the microscopic couplings and can be derived from the initial matrices [69]. This energy function allows a description of the system's equilibrium states (attractors) as a high-dimensional surface the minima of which correspond to the task allocation solutions. This form is consistent with a reaction-based dynamics where pairwise interactions among classes of individuals conditions the global dynamics. As a simple illustration of this idea, let us consider a two-state/two-task case, where it is not difficult

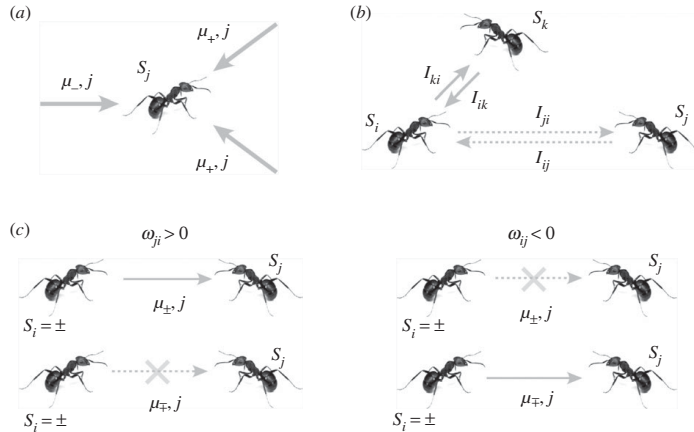


Figure 7. Collective communication dynamics in ant colonies. In (a), we display an agent i and a set of messages reaching it within time τ , all addressed to i while some carrying the $+$ order others the $-$ order. These messages will be integrated according to equation (3.18). On the other hand, (b) shows how interactions via message sending depends on the frequency (or intensity) of messaging between agents, I . Notice that I values decay with distance. Finally, the way that orders are sent by senders (c) depends on yet another set of couplings $\{\omega_{ij} \in \{-, +\}\}$, which determine whether a $+$ or a $-$ order will be dumped into the system depending on the actual state of the sender $S_i = \pm$. Schematically, the arrow connecting sender and receptor is blocked (crossed out) for anticorrelated correlation between coupling ω_{ij} and sender state S_i .

to show that the energy function will correspond to

$$\mathcal{H}(N_1, N_2) = -\frac{1}{2} \left(\sum_{s_i=+1} S_i h_i + \sum_{s_i=-1} S_i h_i \right) \quad (3.16)$$

$$= -\frac{1}{2} (\alpha N_1^2 + \alpha N_2^2 - 2\beta N_1 N_2), \quad (3.17)$$

where we use $\Gamma_{11} = \Gamma_{22} = \alpha$ and $\Gamma_{12} = \Gamma_{21} = \beta$. We can easily recognize in this solution the elliptic paraboloid, displaying a single minimum. In figure 6c, we show an almost symmetric energy surface for $\alpha = 1$, $\beta = 0.1$, whereas a less symmetric case is displayed in figure 6d, where $\beta = 0.5$. In the latter, the coupling between the two tasks creates an elongated valley that would allow for more population fluctuations.

(iv) Collective dynamics of communicating populations

Insect colonies use different organic molecules (pheromones) to transmit signals and process information at a colony level. It is safe to assume that evolution has imprinted on ants a hard-wired pheromone-based detection physiology that generates an internal image of the local environment for each individual ant; however, such a picture is incomplete when confronted with the full complexity of the colony. It is the cobweb of diffusing pheromone signals and ants acting as rewiring agents that confers on the colony its true evolutionary potency. Individual ants are relegated to acting merely as cogs within the macroscopical system [38]. This multiple-scale interrelation is the object of study of the present model by Mikhailov [71].

Imagine a colony of ants individually labelled as $i = 1, \dots, N$. Now, introduce two-state variables for each ant as $S_i \in \{-1, +1\}$, $\forall i$. Thus, vector $S = (S_1, \dots, S_N)$ characterizes the full configuration of the system. In this model, ants are again acting as neural agents, but they are also able to send out and receive messages into and from the colony. A message is

encoded in a pheromone cocktail, and ants continuously secrete it. To simplify the system's dynamics, we will consider that a message is fully described with two labels, namely $\mu_{\sigma,j} = (\sigma, j)$, where $\sigma \in \{-1, 1\}$ and j corresponds to the address tag. In other words, message $\mu_{\sigma,j}$ delivers information σ to the j th ant.

$$S_i(t + \tau) = \text{sign}(m(+, i) - m(-, i)). \quad (3.18)$$

Let us introduce a correspondence matrix, $\{\omega_{ij}\}$, with each of the $N(N-1)$ elements of the former taking values $\{-, +\}$. The function of this matrix is to determine whether a signal will be sent or not in a time interval τ . The way it works is depicted in figure 7c. If $\omega_{ij} = +$, the sender ant, i , will send a message $\mu_{\pm,j}$ to ant j only if $S_i = \pm$, whereas for $\omega_{ij} = -$, the message will be anticorrelated with the state of i , i.e. a message $\mu_{\pm,j}$ is sent only if $S_i = \mp$. In simpler terms, the correspondence matrix distinguished two channels of information transfer: correlated ($\omega = +$) or anticorrelated ($\omega = -$) message and sender state. On the other hand, we define a frequency distribution, I_{ij} , as the number of messages per unit time τ that ant j is sending to ant i (figure 7b). Within a spatial context, it is clear that $I_{ij} = I(|i-j|)$, where $|\cdot|$ is the distance between two ants.

Let us then consider the dynamics of the messages present in the system with labels (σ, i) ,

$$\frac{dm(\sigma, i)}{dt} = -\delta m(\sigma, i) + \frac{1}{2} \sum_j I_{ij} (1 + \sigma \omega_{ij} S_j), \quad (3.19)$$

where we have dubbed δ the message decay rate. Therefore, under the stationary regime, we expect

$$m(\sigma, i) = \frac{1}{2\delta} \sum_j I_{ij} (1 + \sigma \omega_{ij} S_j), \quad (3.20)$$

which, combined with (3.18), leads to

$$S_i(t + \tau) = \text{sign} \left(\frac{1}{\delta} \sum_j I_{ij} \omega_{ij} S_j \right). \quad (3.21)$$

Notice that (3.21) is equivalent to the Hopfield model (2.1), provided that $\omega_{ij} = \omega_{ji}$. Thus, patterns can be stored in a similar manner by following a Hebbian approach by associating:

$$\frac{1}{\delta} I_{ij} \omega_{ij} \rightarrow J_{ij} = \frac{1}{N} \sum_{\mu=0}^p \xi_i^{\mu} \xi_j^{\mu}, \quad (3.22)$$

where, as in §2a, $\{\xi_i^{\mu}\}$ will correspond to the agent states of $\mu = 1, \dots, p$ different stored patterns. Although limitations to capacity will also apply here, perhaps more interestingly, other constraints will arise too, namely:

- (1) Agent-to-agent distance dependence on the signal intensity, $I_{ij} = I(|i - j|)$, which should take the form of a monotonically decreasing function. Effectively, this leads to a *diluted* network, i.e. every agent does not connect to every other agent.
- (2) Environmental noise: signal loss due to fluctuations of the information channel. This can be formalized as thermal noise, which has also been discussed in §2a.
- (3) Cost-efficiency effects: the address-message system devised here carries with it a large cost on the senders to produce the necessary chemical repertoire so that the signal is well-transmitted with minimal error.

The previous models, with all the simplifications they contain, provide a range of examples of how to approach collective behavioural patterns in ant colonies. The first two examples include amplification processes with two rather different outcomes. In one case, time-dependent fluctuations are observed and are connected to high information transfer. Additionally, the suggested criticality provides a source for low energy use together with a rapid response to perturbations. The second deals with solving a short-path (optimization) task which requires exploiting the well-known phenomenon of symmetry-breaking in order to redirect the population towards the shortest branch. In both cases, the state of the system is described by population values where individual ants have no other identity than either being active/inactive or located in one of the two possible paths.

The two other examples provide an increasingly more defined state for ants. When dealing with a number of tasks that require a distribution of individuals performing them, ants are represented by a Boolean vector and interactions occur following a threshold function (as in standard neural nets) but with a state-dependent choice of links. In other words, in contrast with neurons in solid brains, there is no predefined weight matrix but a state-dependent one. This is in fact an especially important difference, along with the fact that the attractors are given by populations of Boolean strings and any ant can end up in one of them. The fourth model makes a strong attempt to get closer to Hopfield's picture, but is less realistic in terms of modelling real case scenarios. In order to achieve that result, the model requires incorporating chemical mediators along with the correspondence matrix.

An important message from the previous examples is that both liquid and solid brains share common descriptions (at some level) such as the presence of threshold-like phenomena. They also achieve attractor states but they are smaller in number, and more degenerate in the liquid state than the solid one, even if full connectivity is considered. In the Hopfield model, for example, a very large number of attractors are present, but each neuron will have a specific set of connections. In the next section, we explore the second class of liquid systems as defined by the dynamics of cells in the IS. In this case, the functional constraints are associated with detection and response to information involving signal cascades mediated by interacting sets of cells.

(b) The immune system as a liquid brain

The IS consists of a myriad of chemical compounds (e.g. antibodies, cytokines) and multiple cell lines (B cells and T cells or lymphocytes, macrophages, etc.) aggregated into a multi-component complex system. The essential purpose of the IS is to detect external and malicious agents (antigens) such as viruses, bacteria or cancerous cells, and prompt an appropriate reaction (antigen neutralization or tolerance). At the same time, it must be able to distinguish the latter from internal signals (the self). As such, the IS must be capable of processing, storing and manipulating large amounts of information [72].

The map of interactions of the IS can be depicted as an interwoven web of signalling and response functions between all its agents. Unravelling a full picture of the IS is beyond the scope of this work. For the purpose of our discussion, we will focus on the three core elements that significantly shape the IS architecture: T cells, B cells and antibodies (Ab). Lymphocytes have specific enzymes on their membranes that store a molecular compound that has been randomly generated during its maturation process. This compound binds to specific fragments of proteins (epitopes) coming from an antigen (often through an antigen presenting cell, APC), hence prompting an internal cascade of reactions that activate the lymphocyte. The collection of receptors of a given lymphocyte clone-line is dubbed an idio-type.

Upon detection, B cells (aided by helper T cells) will proliferate thus generating copies of the same receptor structure, while secreting large concentrations of its specific antibody. In summary, the clonal expansion theory [73] states that, since the generated clones share their idio-type, successive binding to the antigen will be triggered and an amplification process will lead to an immune response [74].

On the other hand, a more systemic approach to the IS reveals an underlying network of idiotypes that excite or inhibit one another through the same detection/reaction mechanisms as with antigens. This phenomenon is known as an idio-type cascade: an initial perturbation (antigen) activates a series of idiotypes filling the system with their corresponding antibodies (Ab1), which, in turn, are detected through another set of idiotypes thus prompting a second batch of antibodies (Ab2), and so on and so forth. This observation suggests a network scheme where each node is associated with an idio-type and each link will correspond to an interaction between any two idiotypes (see later figure 9a–c).

Idio-type cascades were first observed and theorized by Jerne [75] and have since spurred a scientific debate between the allopoietic/autopoietic (reductionist/systemic) approaches to

the IS [76–78]. While Burnet’s theory provides some mechanisms for how the IS generates its idiotypic repertoire capable of self/non-self discrimination, Jerne’s network approach complements this process and shows how a distributed computation concatenated to clonal theory might give rise to crucial information-processing aspects of the immune response.

In this section, we will study some fundamental aspects of the IS as a liquid brain. We will begin by looking at the size of the IS and how it is constrained by its fundamental function of antigen detection and discrimination. Then we will study how the IS is capable of storing information at a network level, discuss how it makes use of its idiotypic landscape structure to naturally reproduce a reliable self/non-self classification, and briefly comment on the implications of such a systems-view to the IS.

(i) Simple constraints for the probability of detection

Early studies of the IS showed that epitope reactivity for a generic lymphocyte (B cell or T cell) is of the order 10^{-5} ; in other words, the probability that a random epitope binds to the surface of a lymphocyte is given by $p \approx 10^{-5}$ [74]. This begs the question: why wouldn’t the IS organize such that $p \sim 1$?

In Percus *et al.* [79], a simple argument was put forward to show that the fact we observe such values of p might be related to the problem of self/non-self recognition, which strongly constrains the way the IS is assembled.

Consider the following definitions: n is the total number of expressed antibody receptors in the IS repertoire, N is the number of foreign epitopes for a given environment and N' denotes the number of self epitopes, or epitopes derived from cells belonging to the organism. Thus, the goal of the IS is to properly distinguish the foreign epitopes while avoiding an immune response to the self-originated ones. Let us denote by $P(N, N'; n)$ the probability that the repertoire of size n is able to properly detect N foreign epitopes and not detect N' self epitopes. Note that the probability of non-recognition of a random epitope for a single lymphocyte is given by $q = 1 - p$. Hence,

$$\begin{cases} q^n; & \text{probability of } n \text{ consecutive non-recognitions} \\ 1 - q^n; & \text{probability of at least one recognition.} \end{cases}$$

Therefore, we may now compute

$$P(N, N'; n) = (1 - q^n)^N (q^n)^{N'} \quad (3.23)$$

The goal is to maximize equation (3.23). This is easily done by maximizing the log $P(N, N'; n)$, which leads to an optimal value for q ,

$$q = \left(1 + \frac{N}{N'}\right)^{-1/n} \approx 1 - \frac{1}{n} \log\left(1 + \frac{N}{N'}\right), \quad (3.24)$$

where we expanded the previous expression using $1/n \ll 1$. Notice that we can now write

$$p \approx \frac{1}{n} \log\left(1 + \frac{N}{N'}\right). \quad (3.25)$$

Inman (see [74, pp. 1226–1227] and references therein) estimated the size of foreign epitopes to be of the order $N \sim 10^{16}$, while, for the human genome, if we approximate the number of self epitopes per protein to be of about 10, then $N' \sim 10^6$. The IS repertoire (total number of idiotypes) can also be approximated by $n \sim 10^7$ [80], which, by equation

(3.25), yields a prediction of $p \sim 2 \times 10^{-6}$. This is smaller than empirically obtained values of p , which are of the order $p_{\text{exp}} \sim 10^{-5}$. This might imply that the IS is operating at a non-optimal stage, reacting more often than necessary. This suggests that further mechanisms must be at play in optimizing the immune response. Notice that, due to the logarithmic nature of (3.25), possible miscalculations of N or N' will not entail substantial deviations for the predicted value of the detection probability p . For more details on this approach, see Perelson & Oster [81, pp. 656–657 and references therein] and Percus *et al.* [79].

(ii) Percolation thresholds in the immune system

After Jerne’s discovery of idiotypic cascades, novel ideas were put forward in trying to understand the organizational principles of the IS as a network. Perelson [77] introduced a simple model of the idiotypic cascading phenomenon. Given a repertoire of n idiotypes (i.e. n different types of antibodies), and assuming that paratopes and epitopes can be thought of as bit-strings of size L (figure 8c), then we will consider that an antibody can detect (bind to) a given string if the number of matched pairs of the ordered paratope–epitope interaction exceeds a threshold value, $\theta < n$. As we will see, this readily imposes strong bounds on the system performance.

Recall that, under Jerne’s paradigm, antibodies are now capable of matching with other antibody types and concatenate into an idiotypic cascade. Thus, we can infer that, for a high threshold value (low reactivity), fewer antibodies will be matching, but also fewer antibodies will be able to detect and react to a given antigen. On the other hand, the reverse is also true: for low values of θ (high reactivity), antibodies will be triggered altogether, as the matching probability is expected to increase. Therefore, it is interesting to study what type of structure will emerge from this simplified model.

Suppose that a given antibody is physically connected to a number of antibodies z , i.e. it will encounter up to z other antibody types but might or might not bind to them. Now, the probability that any pair of antibodies do match is denoted by p , which, by definition, will depend on θ (see below). Thus, given an initial perturbation into the system (such as antigen exposure) then an idiotypic cascade is triggered, where idiotypes react to each other. Such a process will look like a Bethe lattice of degree z (figure 8a). Denote by $\mathcal{A}(i)$ the number of activated antibodies at the i th layer of the tree, then it is easy to show that:

$$\mathcal{A}(i+1) = p(z-1)\mathcal{A}(i), \quad (3.26)$$

which implies that there will be a characteristic probability value $p = p_c = (z-1)^{-1}$, at which the network becomes connected, exhibiting a percolation phase transition [36]. For values of $p > p_c$, the network is fully connected, while for $p < p_c$, any initial perturbation will eventually die out (figure 8a,b).

For the IS one can argue that $z \sim n$ in other words, the system is sufficiently fluid and the coarse number of elements is sufficiently large so that any physical interaction can occur. This sets a value on the critical threshold at $p_c \sim n^{-1}$. On the other hand, one can compute $p = p(\theta)$ by assuming that each bit, of the L -sized strings, is generated by a coin toss. Then, the probability of having two strings with sufficient




Figure 8. Percolation in immune networks. Idiotypic cascades take place at a network level in the IS. (a) A critical percolation cascading on a Bethe lattice of degree $z = 3$. Concentric circles delimit successive layers of the cascade. (b) The percolation probability depends on the matching threshold θ . At low threshold values, the system is highly connected, allowing deep penetration across layers, while for high θ , the matching probability decays abruptly, leading to a phase of low connectivity with small-sized cascades. Right in the interface, we have the percolation point. (c) Two strings (epitope-paratope) of length $L = 10$ with seven matching pairs and three non-matching pairs. For example, if threshold $\theta = 5$, this particular pair of strings would react, whereas for high fidelity matching ($\theta = 8$), the pair would not connect.

complementary bit-to-bit values is

$$p = \sum_{k=\theta}^L \binom{L}{k} \left(\frac{1}{2}\right)^k \left(\frac{1}{2}\right)^{L-k} = \frac{1}{2^L} \sum_{k=\theta}^L \binom{L}{k}, \tag{3.27}$$

which is plotted in figure 8b. We observe a sudden transition from low to high reactivity at around $\theta \sim L/2$. In fact, as $L \rightarrow \infty$, then $p(\theta) \rightarrow 1 - \theta(L/2)$.

Both n and p have been independently measured ([77, pp. 19–20] and references therein). The repertoire size is estimated to be of the order $n \sim 10^6$, while $p \sim 10^{-5}$. Hence, the IS operates in the post-critical regime, where connectivity is high and large cascading events are common.

(iii) Information storage in immune networks

In the search for a clear understanding of how ISs store and process information, optimization arguments as above do not suffice under the light of Jerne’s theory of idiotypic networks. Initial attempts to describe how information is distributed over the network connecting different idiotypes were put forward by De Boer, Hogeweg, Weisbuch and Perelson (see [74, pp. 1229–1258], and references therein). Here, we will briefly summarize a minimal model by Parisi [78] that involves Hopfield-like NN and imposes global limits

on the pattern recognition processes that a distributed network of idiotypes must follow.

Consider the set of antibody binary concentrations $\{c_i(t) \in \{0, 1\}\}$, for $i = 1, \dots, N$, with N the total number of antibodies of a healthy human IS (around $10^6 - 10^7$). To all effects and purposes, ‘antibodies’ and ‘idiotypes’ are interchangeable from here onwards. Next, we model idiotypic interaction networks, by imposing a dynamical process of idiotypic concentrations in the same spirit as (2.1):

$$c_i(t + \tau) = \theta \left(\sum_j J_{ij} c_j(t) \right). \tag{3.28}$$

Now, the interactions between different idiotypes are mediated by $\{J_{ij}\}$, for which we consider the following properties:

- (i) $J_{ij} = 0$, i.e. no idiotypic self-interaction is allowed, which is the case for paratope–epitope complementarity matching;
- (ii) $J_{ij} = J_{ji}$, which is a simplification of the Onsager affinity relations between idiotypes,³ $\log|J_{ij}| = \log|J_{ji}|$;
- (iii) $J_{ij} = U(-1, +1), \forall i \neq j$.

Condition (iii) states that the values of the off-diagonal elements of J_{ij} are taken from the uniform distribution between $[-1, 1]$. These approximations allow for a derivation of overall limits of distributed storage of information. The system is now described as a spin glass [19,20,82].

Stable solutions for this particular problem turn out to be fully characterized by an average number of pre-assigned concentrations, M . In other words, a generic initial configuration of concentrations will inevitably flow into a stable state by switching concentration values on and off until a pre-assigned configuration of concentration levels is reached. These global stable states act as memory basins similarly to how memory is stored in the aforementioned NN models. Naturally, $M < N$, thus, we can define $\alpha < 1$ such that $M = \alpha N$.

Spin glass theory [82] predicts that, for $N \gg 1$, out of the total 2^N possible binary states of the system, and for conditions (i)–(iii), a total of $2^{\alpha N}$ patterns can be stored, with $\lambda \sim 0.3$. Withal, we can now try to understand the relation between λ and parameter α .

Let us consider the probability (p_m) of randomly choosing a ‘memorized state’ out of all the possible configurations or, simply, $p_m = 2^{\alpha N} / 2^N = 2^{-(1-\alpha)N}$. However, because only M preassigned concentrations are required to fully describe an attractor, we then expect a number of compatible solutions per stable state. Thus, let us compute the average number of compatible solutions per attractor as

$$\left(\begin{array}{c} \text{av. no.} \\ \text{of solutions} \end{array} \right) = p_m \times \left(\begin{array}{c} \text{degeneracy} \\ \text{of config.} \end{array} \right) \tag{3.29}$$

$$= 2^{-(1-\alpha)N} \times 2^{N-M} \tag{3.30}$$

$$= 2^{(\alpha-\alpha)M}. \tag{3.31}$$

Notice that the average number of solutions will be greater or equal to one iff $\alpha < \lambda \sim 0.3$. Essentially, this imposes a bound in M . In other words, if we denote $\alpha_c = \lambda$, then for $M > \alpha_c N$, no equilibrium states are found. Thus, $M_c \equiv \alpha_c N$ is the maximum number of pre-assigned antibody concentrations such that the dynamics imposed by (3.28) flow into well-defined stored patterns. This effectively constrains the memory content that an idiotypic interaction web is able to store.

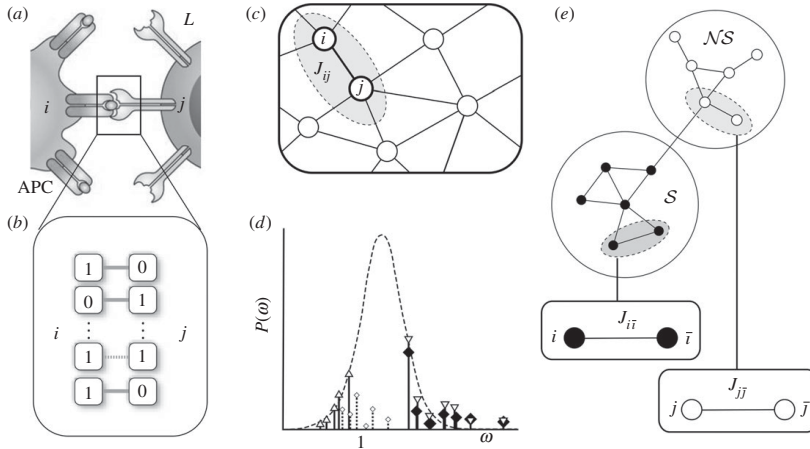


Figure 9. The IS as a liquid brain. (a) An interaction between an APC carrying a fragment of an antigen and presenting to a lymphocyte (L). (b) Upon matching, the lymphocyte will react by secreting antibodies with the corresponding matching code, thus flooding the system with its idiotypic information and prompting an idiotypic cascade. (c) A representation of the subjacent idiotypic network operating across the IS. This network is actually self-organized into two major blocks (e) of heavily influential (darker region) and weakly influential (lighter region) nodes. Such an effect can be computationally studied by looking at the strength distribution (d), $P(\omega)$, noting that picking a random node from the right/left (strong/weak) (i/j) ends of the spectrum, and then looking at its corresponding next neighbours strength (\bar{i}/\bar{j}), they typically fall under the same category, i.e. strong/weak nodes connect to strong/weak nodes. This suggests a network-like mechanism for tackling the self (S)/non-self (NS) classification problem (ω -axis is depicted in logarithmic scale). Strong nodes are responsible for self-addressed Ab, and vice versa. Part (d) is adapted from Barra & Agliari [76, pp. 15–16].

Parisi's contributions fuelled the statistical physics and spin-glass approach to the analysis of the IS (see sections below), but also suggested how complex interactions among a large number of elements yield strong constraints in the feasibility of IS information processing. This has major implications in that it shows how selective pressures pushing for a reliable IS must go beyond the purely genetic component involved in epitope/idiotope generation. Under this higher-level picture, the capacity for learning and reliably retrieving information from the immune network is limited by fundamental statistical attributes. In the next sections, we will study further the structure of these idiotypic networks and explore how other computational properties can emerge from such a systems analysis.

(iv) Idiotypic networks as liquid neural nets

In what remains of this section, we will outline a model by Barra & Agliari (BA) [76] based on statistical physics of a well-mixed/liquid neural web representing Jerne's idiotypic network. Let us assume:

- (i) A given clone idiootype is fully characterized by a string of L bits. All idiotypes are of the same size.
- (ii) Each string is obtained from successive, independent coin-tosses with values $\{0, 1\}$.
- (iii) The number of cells of a clone-type is sufficiently large so that potential idiotypic interactions are always carried out with their respective intensity values.

Assumptions (i)–(ii) are sensible first approximations to the biological processes the IS undergoes during maturation [72]. On the other hand, a sufficiently high number of lymphocytes per idiootype is not realistic under the light of clonal expansion

theory. However, the goal of the BA model is to figure out the overall implications of having an idiotypic network description.⁴

Let us construct an idiootope space $Y_L \equiv \{\xi \in \{0, 1\}^L\}$ spanning all possible strings with bit-size L . Indexes $i, j, \dots \in \{1, \dots, N\}$, with N corresponding to the total number of different clone types in the IS. *A priori*, a complete repertoire would seem to scale as $N \sim 2^L$, however, as we will see, the network constraints will give rise to another scaling behaviour between the repertoire size and epitope/paratope length.

Next, we construct the network following a simple model of chemical complementarity. As usual, let us define a complementarity function:

$$\begin{aligned} \kappa_{ij} &:= \frac{1}{L} \sum_{\mu=1}^L [\xi_{\mu}^i (1 - \xi_{\mu}^j) + \xi_{\mu}^j (1 - \xi_{\mu}^i)] \\ &= \frac{1}{L} (|\xi^i| + |\xi^j| - 2\xi^i \cdot \xi^j), \end{aligned} \quad (3.32)$$

with $|A| \equiv \sum_{\mu=1}^L A_{\mu}$. This accounts for the total number of complementary inputs between idiotypes i and j . For example, suppose $L = 5$, then for $\xi^i = (10101)$ and $\xi^j = (01011)$, $\kappa_{ij} = \frac{3}{5}$ (figure 9b). In turn, this allows the construction of a chemical affinity function

$$\begin{aligned} f_{\beta L}: Y \times Y &\rightarrow \mathbb{R} \\ (\xi^i, \xi^j) &\mapsto f_{\beta L}(\xi^i, \xi^j) = \beta \kappa_{ij} - (1 - \kappa_{ij}), \end{aligned} \quad (3.33)$$

which is defined as a balance between repulsion and attraction effects of anti-complementary and complementary bit-pairs, modulated by trade-off parameter $\beta \geq 0$. Thus, it will be bounded as $-1 \leq f_{\beta L} \leq \beta$, distinguishing two

interactive regimes for each pair of idiotypes:

$$f_{\beta,L} = \begin{cases} < 0 & \text{repulsive regime} \\ > 0 & \text{attractive regime.} \end{cases} \quad (3.34)$$

Following these precepts, let us outline how the unweighted network of idiotypic interactions will unfold. The IS can be arguably approximated as a well-mixed system. This means that, following (iii), any possible physical interaction (B cell/T cell or APC/T cell, etc.) occurs at a sufficiently high rate so that we need only to account for their internal affinity structure. Let us then define $p_{\beta,L}$ as the probability that two generic idiotypes display a matching interaction. Consider the following:

- (1) The idiotypic strings, $\{\xi^i\}$, are extracted by a successive L random coin-tosses with equal probability for $\{0, 1\}$ values, i.e. $p_0 = p_1 = 1/2$.
- (2) The complementarity κ_{ij} and affinity $f_{\beta,L}$ functions fully regulate the interactions. In particular, we define a link between two generic idiotypes (ξ^i, ξ^j) iff $f_{\beta,L}(\xi^i, \xi^j) > 0$, i.e. if the pair lays on the attractive regime.

In general, the probability for any two idiotypes to produce a complementarity value, κ , is $P(\kappa) = \binom{L}{\kappa}/2^L$. Now, owing to assumption (2) and (3.33), then

$$p_{\beta,L} = \mathcal{P}\left(\bigcup_{f_{\beta,L}>0} \kappa\right) = \sum_{\kappa=\lfloor L/(\beta+1) \rfloor+1}^L P(\kappa). \quad (3.35)$$

Since N is the total number of different idiotypes, the emergent network picture will be described by an Erdős–Renyi (ER) graph with degree distribution

$$\phi_{\beta,L}(k) = \binom{N}{k} p_{\beta,L}^k (1 - p_{\beta,L})^{N-k}, \quad (3.36)$$

the mean value of which corresponds to $\langle k \rangle \approx p_{\beta,L} N$. ER networks display a percolation point at which the system acquires a giant connected component [36]. Typically, this occurs at $\langle k \rangle = 1$, associated with $p_c := 1/N$. Next, we explore what regime we should expect the idiotypic network to be in and how this reflects on the IS's repertoire capacity.

For finite values of L , the shape of the function $p_{\beta,L}$ as a function of β is that of a transfer function. Recall that the trade-off parameter β separates the favourably repulsive regime ($\beta < 1$) from the favourably attractive one ($\beta > 1$), $\beta = 1$ corresponding to the symmetric case. Now, if chains (epitopes/idiotypes) are considered to be large, then an amplification process occurs depending on the favourably repulsive/attractive regimes determined by the value of β . Such amplification is reflected on the switch-like behaviour of the connection probability. On the other hand, since the percolation threshold will be of the order of $1/N$, even if the system is repulsively favoured ($\beta < 1$), it can still easily become fully connected.

Now, consider the three elements that are now coming together: probability of connection, $p_{\beta,L}$; number of idiotypes (or different clones), N ; and the average number of connections per idiotypic, $\langle k \rangle$. While the probability of connection is purely a result of the internal chemical interactions, $\langle k \rangle$ is a defining feature of our network. Yet experimental data shows a value of $N \sim 10^{18}$, and a connectivity between idiotypes in mature ISs

of $p_{\beta,L} \sim 3 - 5\%$ [83]. This means that we should expect a densely connected network of around $\langle k \rangle \sim 10^{12}$.

Moreover, the fact that $p_{\beta,L}$ is so low suggests that the system operates at the repulsive regime. That being, we may now compute the relation between epitope size L and number of idiotypes in the BA model by using (3.35) and $\beta < 1$. This results in a scaling relation $N \sim \sqrt{L}e^{\gamma L}$, with $\gamma < 1$, as opposed to the bit-by-bit repertoire size, which would grow as 2^L [76, pp. 5–13].

Hitherto, we have been able to characterize idiotypic networks using only basic assumptions for chemical affinity, which has led to a dense ER. But what kind of computations is this system able to perform? And how does the IS use its autopoeietic features to distinguish the self/non-self? To provide an answer to these questions, we ought to look at a fine-grained version of the idiotypic network and inquire into how its interaction intensities are distributed over the net.

(v) Stewart–Varela–Coutinho theory

In their seminal papers Stewart, Varela and Coutinho [84,85] showed that a network systems approach to the idiotypic webs described by Jerne actually displayed two major interactive blocks: a highly connected (strongly interacting) module and a loosely connected (low interactive) one. Such an observation suggested that each module's activity could correlate to the self and non-self reactions of the IS. More specifically, the strongly connected module acts as an auto-regulated dynamical subsystem that is continuously activated and auto-inhibited; this would correspond to a tolerance response, thus associated with the self (S) stimuli, in a healthy IS. On the other hand, the low interactive module shows a basal activity in the system, but under the presence of a stimulus it will be activated, thus prompting a neutralization response. The latter module is then associated to the non-self (NS) stimuli. Hence, through this network structural property the vast repertoire of the IS is capable of sorting out the self/non-self.

Although this two-block structure could appear to be the result of an intricate evolutionary process, by following the BA model, a twofold assembly akin to Stewart, Varela and Coutinho's system is shown to emerge for free. This would suggest a generative mechanism capable of explaining the underlying self/non-self modular structure independently of adaptive drives. Let us briefly explore how this phenomenon takes place at the weighted network level.

(vi) Weighted idiotypic networks and mirror types

The affinity function $f_{\beta,L}$ works as a representation of the chemical reactions that take place on the cell surfaces, then, following a simple extension to the concept of interaction, connection matrix $J_{ij}(\beta, L)$ can be defined as

$$J_{ij}(\beta, L) \sim \Theta[f_{\beta,L}(\xi^i, \xi^j)] e^{f_{\beta,L}(\xi^i, \xi^j)}. \quad (3.37)$$

Notice how we still impose a lower threshold of connectance by setting $J_{ij} = 0$ for $f_{\beta,L} \leq 0$, which keeps the previous network picture, while turning on the matrix values smoothly in the attractive regime. The exact values for all the J_{ij} will depend on each realization of the stochastically generated network of idiotypes $\{\xi^i\}$, thus, it will be necessary to normalize each interaction parameter over all the space of possible networks [76, pp. 13–20].

Once the interaction intensities are in place, one can look at the total strength for each idiotypic (node) as

$$\omega_i(\beta, L) = \sum_{j|i_i > 0} J_{ij}(\beta, L). \quad (3.38)$$

Heuristically, J_{ij} values characterize the robustness of a given $i - j$ interaction, while ω_i measures how influential idiotypic i is relative to the whole network. Now, consider the weight frequency distribution $P(\omega)$, which can be shown to be well approximated by a normal distribution [76, pp. 13–20]. If we select an idiotypic i and look at its first-order neighbours that inhibit i , namely mirror- i idiotypes (or simply \bar{i}), then it is possible to study how the system self-classifies these pairs into two major classes (figure 9e): strong and weakly interacting pairs. The fact that the interaction is symmetrically strong/weak for each pair is a consequence of chemical complementarity in the affinity function.

However, this simple realization turns out to be an extremely powerful tool to resolve the self/non-self distinction. In summary:

- The $P(\omega)$ degree distribution separates the two regimes of strong/weak influential nodes (figure 9d). The weak nodes (blank triangles) happen to have weak mirror types (blank squares) whereas the strongly interacting nodes (reversed blank triangles) have mirror types (black squares) that also are highly interactive nodes. This mechanism gives rise to the S/VNS modules.
- The strong block is hypothesized to account for the self-directed antibodies, while the weak module acts as a basal signal only activated by the presence of non-self antibodies (triggered by external antigens).
- Establishment of robust memories occurs more effectively at the weakly interacting block, as relative variations in the affinity values will produce more durable configuration changes in this network module.
- Autoimmune diseases can now be understood as deviations from the two-module structure, where strongly interacting circuits (responsible for self-addressed antibodies) may deviate towards lower weighted regions of the spectrum, thus triggering auto-immune response.

Hence, a natural mechanism for fundamental computational questions such as the self/non-self identification is derived from first principles. These are constructed under the assumption that the interaction time scales are small compared with the global observational time scales, while, on the other hand the ability of the ‘neural agents’ (idiotypes) to rapidly propagate throughout the environment ultimately allows the characterization of the idiotypic network as a biologically meaningful system. Thus, the IS appears to be a limit case scenario for ‘liquid brains’, where it is precisely the high levels of agent mobility that give rise to its capacity to solve classification problems.

This realization leads to novel questions: are there size limitations for ISs and their performance in terms of physical embodiment? What are the consequences of these constraints to the self/non-self distinguishability? Or, in general, can different ‘ISs’ exist across multiple scales?

(vii) B cell and T cell interactions as a glassy system

Immune systems seem to conceal within a concoction of networks the nodes (antibodies and lymphocytes) of which can be active or inactive in relation to their

presence or concentration levels. New approaches based on lymphocyte–lymphocyte interactions have shed light onto the intrinsic glassy mechanisms behind IS information processing [86,87].

The previous analysis has only focused on the presence–absence of Ab types in an idiotypic context, yet, underlying the dynamics of idiotypic cascades, B cell clone types are behind the production (or lack thereof) of the antibodies. This will entail the global state of the IS and, as discussed above, allow for self/non-self distinguishability or memory storage. On the other hand, T cells, helpers and suppressors, will promote and inhibit the proliferation of B cells. From a network perspective, this can be understood as a core bipartite network with T cells on one layer and B cells on the other. However, this internal network will contain both excitatory and inhibitory interactions. Such a rich dynamical picture is what ultimately confers the IS its glassy behaviour.

Under this framework, more profound questions can be formally pursued. In Agliari *et al.* [86], the authors showed how a spin-glass approach to the coupling between lymphocytes leads to different scenarios. In these, the overall clonal expansion extent and the ratio between the number of B cells and the number of T cells (helpers and suppressors) strongly define the regions of stable memory storage and inability to retrieve memories. This has major implications for complex diseases such as HIV. Here, T-cell levels are depleted by the virus and the system is driven into the spin-glass states where the IS fails to operate reliably [86].

Finally, this provides a yet deeper analogy between liquid brains and solid neural networks (see Hopfield model in §1). In the latter, the memory capacity α measures the relative number of patterns stored in a network of a given size, but in the IS this seems to be constrained by the ratio in gross numbers between the two type of lymphocytes. Nonetheless, there is yet another fundamental distinction between the two systems. Hopfield nets operate by interpreting a pattern at a time (provided an input, multiple outputs can be reliably retrieved), but the IS computational task works in a different way; it needs to carry out parallel tasks (as the organism is exposed to multiple pathogens at once), but its decision boils down to a single bit of information: immune reaction or tolerance. This imposes severe restrictions on both the idiotypic and the lymphocyte networks of interactions (see [88] for an excellent investigation into this matter).

4. Discussion

The emergence of cognition in our biosphere has been marked by several key events that allowed the evolution of special classes of cell phenotypes along with ways of wiring them together. Nerve cells and nerve nets pervade the revolution towards new life forms capable of dealing with non-genetic information in complex ways. But the basic ingredients for the emergence of complex forms of information processing have appeared multiple times at different scales and in different evolutionary contexts [2]. Neural-like processing systems have evolved as specialized organs but also as communities of moving agents. In both cases, agent–agent interactions involve some sort of recognition, internal communication coding and stimuli thresholds that decide if changes are made. As shown in previous sections, simple models can capture relevant phenomena associated with both classes of systems.

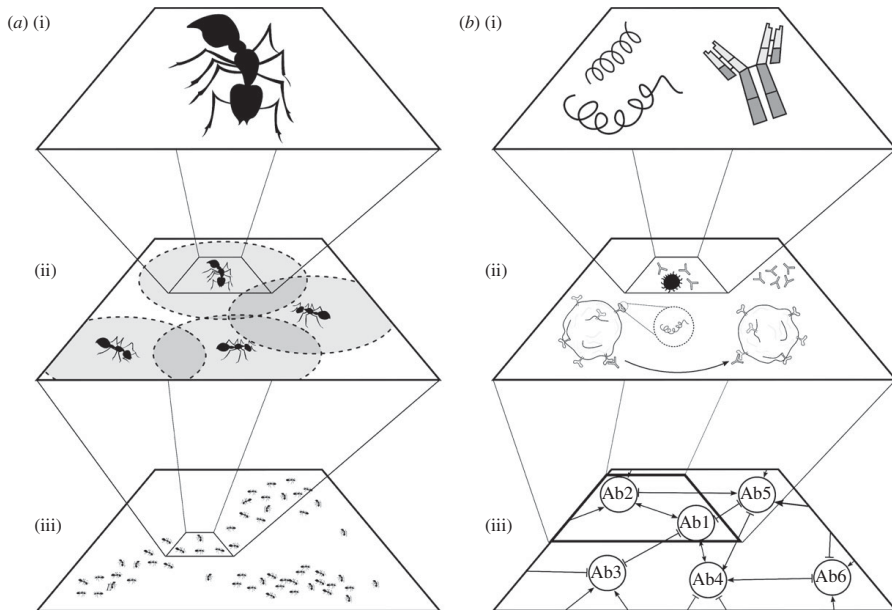


Figure 10. Multiscale dynamics in liquid brains. As occurs with many other complex systems, each example of liquid brains involves several scales of description. (a) Ant colonies perform diverse functionalities, such as collective foraging (a_{iii}) on a colony-level basis. At a smaller scale, pairwise interactions among ants take place (a_{ii}). Such interactions are localized and, thus, constrained by spatio-temporal properties such as agent mobility or density. At the top of this hierarchy (a_i), we encounter single ants as a system. These agents will be defined by a set of rules that drive their behaviour at this minimal scale. (b) A similar scheme can be made for the IS. Scales now involve the idiotypic (or antibody type) network (b_{iii}), where information is processed, for instance, at the self/non-self discrimination level (see above). As we zoom in, we encounter the cellular-scale interactions level (b_{ii}), which are also associated to the simple-matching recognition dynamics. Finally, yet another level of complexity is reached at the description of the IS elementary agents (b_i): viruses, paratopes, epitopes and surface receptors.

These two classes of networks share emergence as a major feature. Memory, learning or decision-making are grounded in a set of bottom-up phenomena where emergent properties arise from individual, microscopic interactions (figure 10). Collective phenomena and a physics approach becomes a natural common field from where to extract universal features. These emergent traits can be attractor basins associated with memory states or efficient task allocation, but can also be phase transitions due to the presence of critical connectivities, or even criticality itself, enabling rapid and efficient information transfer.

What is missing from our previous models? An important piece of complexity that has been ignored in this discussion is the internal complexity of the agents. This is not necessarily a limitation. When dealing with complex systems, we purposely ignore unnecessary detailed descriptions of the system in order to render the problem solvable and provide true understanding. The level of simplification is imposed by the kind of question being addressed. As in the Hopfield model and other classical neural network approaches, neuronal complexity is reduced to the minimal description. The IS too is a rather sophisticated system, exhibiting a considerable diversity of cell types and interactions. Ants, on the other hand, are not simple strings of Boolean bits. Individual cognition exists: ants carry actual brains inside their heads, even if small ones [89].

As pointed out by some authors [90] even if these brains are orders of magnitude smaller than ours, they exhibit some types of cognitive skills. How important are they compared to colony-level cognition? This is an open question that will

require further attention. It might be the case that increases in the cognitive complexity of ant colonies is accompanied by reductions in an individual's cognition. Such a trade-off has been explored in other contexts, like the evolution of multicellularity [91], described by the term 'complexity drain'. Previous work used coupled discrete maps to suggest that collective computation might also display this phenomenon [92]. Further multiscale models of liquid brains should be developed to properly address this question.

What other systems can be described as liquid brains? As mentioned before, computations arising from gene–gene regulatory links within cells have been studied using similar formal schemes, from Boolean tables to threshold networks [93,94]. Several observations also suggest that gene networks might be critical [95–98].

A final comment needs to be made concerning the physical phases used to present our classification between liquid and solid. The use of the term 'liquid' to label the classes of systems discussed here is only partially appropriate. Particularly in relation to ants, their collective dynamics is more appropriately described as 'active matter': ants (as well as bacterial and robotic swarms) need to be understood as interacting self-propelled robots [99]. Here too the statistical physics approach has played a key role in understanding coordinated behaviour and its transitions. Once again, in spite of considerable differences, deep analogies exist between classical equilibrium statistical physics systems and those made of active units. Understanding the cognitive complexity of liquid brains and its limits can provide deep insights into the evolution of information-

processing, computational systems grounded in living structures.

Data accessibility. This article has no additional data.

Authors' contributions. All authors built and analysed the mathematical models. All authors gave final approval for publication.

Competing interests. We have no competing interests.

Funding. This study was supported by the Botin Foundation, by Banco Santander through its Santander Universities Global Division, FIS2015-67616-P and by the Santa Fe Institute. J.P. acknowledges support from 'María de Maeztú' fellowship MDM-2014-0370-17-2. This work has also counted on the support of Secretaria d'Universitats i Recerca del Departament d'Economia i Coneixement de la Generalitat de Catalunya.

Acknowledgements. The authors want to thank Luís Seoane, Sergi Valverde, Victor Maull and the other members of the Complex Systems Laboratory for stimulating discussions. Special thanks to Bob Merriman for his inspiring ideas. We also thank the Santa Fe

Institute for hosting the Working Group on 'Liquid brains, solid brains'.

Endnotes

¹The same results are obtained when the active phase is used, since the two points just exchange their stability.

²A trade-off between polymorphism and pheromone repertory is evidenced, as caste differentiation already segregates behavioural states in a decisive way.

³Considering antisymmetric interactions leads to chaotic behaviour for the time-dependent dynamics, which is arguably not a good description of the IS as Ab concentrations would be observed to behave randomly.

⁴Assumption (iii) acquires more relevance in the full BA model as the coupling between the actual lymphocyte activity and the subjacent idiotypic network is studied. In this paper, though, we concentrate on the network-like features.

References

- Hopfield JJ. 1994 Physics, computation and why biology looks so different. *J. Theor. Biol.* **171**, 53–60. (doi:10.1006/jtbi.1994.1211)
- Baluška F, Levin M. 2016 On having no head: cognition throughout biological systems. *Front. Psychol.* **7**, 902. (doi:10.3389/fpsyg.2016.00902)
- Benenson Y. 2012 Biomolecular computing systems: principles, progress and potential. *Nat. Rev. Genet.* **13**, 455–468. (doi:10.1038/nrg3197)
- Bray D. 1990 Intracellular signalling as a parallel distributed process. *J. Theor. Biol.* **143**, 215–231. (doi:10.1016/S0022-5193(05)80268-1)
- Jablonka E, Lamb MJ. 2006 The evolution of information in the major transitions. *J. Theor. Biol.* **239**, 236–246. (doi:10.1016/j.jtbi.2005.08.038)
- Rose S. 2006 *The future of the brain*. New York, NY: Oxford University Press.
- Pagán O. 2018 *Strange survivors: how organisms attack and defend in the game of life*. Dallas, TX: BenBella Books.
- Oster GF, Wilson EO. 1978 *Caste and ecology in the social insects*. Princeton, NJ: Princeton University Press.
- Solé R, Goodwin B. 2000 *Signs of life: how complexity pervades biology*. New York, NY: Basic Books.
- Gordon DM. 1999 *Ants at work: how insect society is organized*. New York, NY: W.W. Norton & Co.
- Gordon DM. 2010 *Ant encounters*. Princeton, NJ: Princeton University Press.
- Stein DL, Newman CM. 2013 *Spin glasses and complexity*. Princeton, NJ: Princeton University Press.
- Tero A, Kobayashi R, Nakagaki T. 2007 A mathematical model for adaptive transport network in path finding by true slime mold. *J. Theor. Biol.* **244**, 553–564. (doi:10.1016/j.jtbi.2006.07.015)
- Adamatzky A. 2010 *Physarum machines: computers from slime mould*. Singapore: World Scientific.
- Deco G, Jirsa VK, Robinson PA, Breakspear M, Friston K. 2008 The dynamic brain: from spiking neurons to neural masses and cortical fields. *PLoS Comp. Biol.* **4**, e1000092. (doi:10.1371/journal.pcbi.1000092)
- Forrest S. 1990 Emergent computation: self-organizing, collective, and cooperative phenomena in natural and artificial computing networks. *Physica D* **42**, 1–11. (doi:10.1016/0167-2789(90)90063-U)
- Farmer JD. 1990 A rosetta stone for connectionism. *Physica D* **42**, 153–187. (doi:10.1016/0167-2789(90)90072-W)
- Hopfield JJ. 1982 Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl Acad. Sci. USA* **79**, 2554–2558. (doi:10.1073/pnas.79.8.2554)
- Amit DJ, Gutfreund H, Sompolinsky H. 1985 Spin-glass models of neural networks. *Phys. Rev. A* **32**, 1007–1018. (doi:10.1103/PhysRevA.32.1007)
- Sompolinsky H. 1988 Statistical mechanics of neural networks. *Phys. Today* **41**, 70–80. (doi:10.1063/1.881142)
- Haken H. 1991 *Synergetic computers and cognition: a top-down approach to neural nets*. Berlin, Germany: Springer.
- McCulloch WS, Pitts W. 1943 A logical calculus of the ideas immanent in nervous activity. *Bull. Math. Biophys.* **5**, 115–133. (doi:10.1007/BF02478259)
- Rashevsky N. 1960 *Mathematical biophysics: physico-mathematical foundations of biology*, 3rd edn. New York, NY: Dover P., Inc.
- Kauffman SA. 1993 *The origins of order: self-organization and selection in evolution*. New York, NY: Oxford University Press.
- Bornholdt S. 2005 Less is more in modelling large genetic networks. *Science* **310**, 449–451. (doi:10.1126/science.1119959)
- Bornholdt S. 2008 Boolean network models of cellular regulation: prospects and limitations. *J. R. Soc. Interface* **5**, S85–S94. (doi:10.1098/rsif.2008.0132.focus)
- Peretto P. 1992 *An introduction to the modeling of neural networks*. Cambridge, UK: Cambridge University Press.
- Hertz J, Krogh A, Palmer RG. 1991 *Introduction to the theory of neural computation*, vol. I. Redwood City, CA: Addison-Wesley.
- Hesse J, Gross T. 2014 Self-organized criticality as a fundamental property of neural systems. *Front. Neurosci.* **8**, 166.
- Muñoz MA. 2018 Colloquium: criticality and dynamical scaling in living systems. *Rev. Mod. Phys.* **90**, 031001. (doi:10.1103/RevModPhys.90.031001)
- Chialvo DR. 2004 Critical brain networks. *Physica A* **340**, 756–765. (doi:10.1016/j.physa.2004.05.064)
- Chialvo DR. 2010 Emergent complex neural dynamics. *Nat. Phys.* **6**, 744–750. (doi:10.1038/nphys1803)
- Plenz D, Niebur E, Schuster HG. 2014 *Criticality in neural systems*. Weinheim, Germany: Wiley-VCH.
- Eckmann JP, Feinerman O, Gruendlinger L, Moses E, Soriano J, Tlusty T. 2007 The physics of living neural networks. *Phys. Rep.* **449**, 54–76. (doi:10.1016/j.physrep.2007.02.014)
- Strogatz SH. 1994 *Nonlinear dynamics and chaos*. Cambridge, MA: Westview Press.
- Solé RV. 2011 *Phase transitions*. Princeton, NJ: Princeton University Press.
- Goldenfeld N. 1992 *Lectures on phase transitions and the renormalization group*. Cambridge, MA: Westview Press.
- Wilson EO. 2012 *The social conquest of earth*. New York, NY: W.W. Norton & Co.
- Wilson EO, Hölldobler B. 2005 Eusociality: origin and consequences. *Proc. Natl Acad. Sci. USA* **102**, 13 367–13 371. (doi:10.1073/pnas.0505858102)
- Solé RV, Miramontes O, Goodwin BC. 1993 Oscillations and chaos in ant societies. *J. Theor. Biol.* **161**, 343–357. (doi:10.1006/jtbi.1993.1060)
- Miramontes O, Solé RV, Goodwin BC. 1993 Collective behaviour of random-activated mobile cellular automata. *Physica D* **63**, 145–160. (doi:10.1016/0167-2789(93)90152-Q)
- Cole BJ. 1991 Short-term activity cycles in ants: generation of periodicity by worker interaction. *Am. Nat.* **137**, 244–259. (doi:10.1086/285156)

43. Hölldobler B, Wilson EO. 1990 *The ants*. Cambridge, MA: Harvard University Press.
44. Goss S, Deneubourg JL. 1988 Autocatalysis as a source of synchronised rhythmic activity in social insects. *Insectes Soc.* **35**, 310–315. (doi:10.1007/BF02224063)
45. Bonabeau E, Theraulaz G, Deneubourg JL. 1998 The synchronization of recruitment-based activities in ants. *BioSystems* **45**, 195–211. (doi:10.1016/S0303-2647(98)00004-5)
46. Buszaki G, Draguhn A. 2004 Neuronal oscillations in cortical networks. *Science* **304**, 1926–1929. (doi:10.1126/science.1099745)
47. Murray JD. 1989 *Mathematical biology*, pp. 610–650. New York, NY: Springer.
48. Miramontes O. 1995 Order-disorder transitions in the behavior of ant societies. *Complexity* **1**, 56–60. (doi:10.1002/cplx.v1.3)
49. Solé RV, Miramontes O. 1995 Information at the edge of chaos in fluid neural networks. *Physica D* **80**, 171–180. (doi:10.1016/0167-2789(95)90075-6)
50. Miramontes O, DeSouza O. 1996 The nonlinear dynamics of survival and social facilitation in termites. *J. Theor. Biol.* **181**, 373–380. (doi:10.1006/jtbi.1996.0138)
51. Mora T, Bialek W. 2011 Are biological systems poised at criticality? *J. Stat. Mech.* **144**, 268–302. (doi:10.1007/s10955-011-0229-4)
52. Delgado J, Solé RV. 1997 Noise induced transitions in fluid neural networks. *Phys. Lett. A* **229**, 183–189. (doi:10.1016/S0375-9601(97)00184-9)
53. Delgado J, Solé RV. 2000 Self-synchronization and task fulfilment in ant colonies. *J. Theor. Biol.* **205**, 433–441. (doi:10.1006/jtbi.2000.2077)
54. Deneubourg JL, Goss S. 1989 Collective patterns and decision-making. *Ethol. Ecol. Evol.* **1**, 295–311. (doi:10.1080/08927014.1989.9525500)
55. Detrain C, Deneubourg JL. 2006 Self-organized structures in a superorganism: do ants 'behave' like molecules? *Phys. Life Rev.* **3**, 162–187. (doi:10.1016/j.plrev.2006.07.001)
56. Garnier S, Gautrais J, Theraulaz G. 2007 The biological principles of swarm intelligence. *Swarm Intell.* **1**, 3–31. (doi:10.1007/s11721-007-0004-y)
57. Bonabeau E, Dorigo M, Theraulaz G. 1999 *Swarm intelligence: from natural to artificial systems*. New York, NY: Oxford University Press.
58. Dorigo M, Maniezze V, Colonna A. 1996 The ant system: optimization by a colony of cooperating agents. *IEEE Trans. Syst. Man Cybern. B* **26**, 1–13. (doi:10.1109/3477.484436)
59. Dorigo M, Birattari M, Stützle T. 2006 Ant colony optimization. *IEEE Comp. Intell. Mag.* **1**, 28–39. (doi:10.1109/MCI.2006.329691)
60. Marshall JAR, Bogacz R, Dornhaus A, Planqué R, Kovacs T, Franks NR. 2009 Optimal decision-making in brains and social insects. *J. R. Soc. Interface* **6**, 1065–1074. (doi:10.1098/rsif.2008.0511)
61. Millonas MM. 1993 Swarms, phase transition, and collective intelligence; and a nonequilibrium statistical field theory of swarms and other spatially extended complex systems. Santa Fe Institute Working Papers 93-06-039.
62. Nicolis SC, Deneubourg JL. 1999 Emergent patterns and food recruitment in ants: an analytical study. *J. Theor. Biol.* **198**, 575–592. (doi:10.1006/jtbi.1999.0934)
63. Beekers R, Deneubourg JL, Goss S. 1992 Trails and U-turns in the selection of a path by the ant *Lasius niger*. *J. Theor. Biol.* **159**, 397–415. (doi:10.1016/S0022-5193(05)80686-1)
64. Deneubourg JL, Aron S, Goss S, Pasteels JM. 1990 The self-organizing exploratory pattern of the Argentine ant. *J. Insect. Behav.* **3**, 159–168. (doi:10.1007/BF01417909)
65. Altshuler E, Ramos O, Núñez Y, Fernández J, Batista-Leyva AJ, Noda C. 2005 Symmetry breaking in escaping ants. *Am. Nat.* **166**, 643–649. (doi:10.1086/498139)
66. Amé JM, Halloy J, Rivault C, Detrain C, Deneubourg JL. 2006 Collegial decision making based on social amplification leads to optimal group formation. *Proc. Natl Acad. Sci. USA* **103**, 5835–5840. (doi:10.1073/pnas.0507877103)
67. Bonabeau E. 1996 Marginally stable swarms are flexible and efficient. *J. Phys. I France* **6**, 309–324. (doi:10.1051/jpl:1996151)
68. Duarte A, Weissing FJ, Pen I, Keller L. 2011 An evolutionary perspective on self-organized division of labor in social insects. *Annu. Rev. Ecol. Syst.* **42**, 91–110. (doi:10.1146/annurev-ecolsys-102710-145017)
69. Gordon DM, Goodwin BC, Trainor LEH. 1992 A parallel distributed model of the behaviour of ant colonies. *J. Theor. Biol.* **156**, 293–307. (doi:10.1016/S0022-5193(05)80677-0)
70. Gordon DM. 1986 The dynamics of the daily round of the harvester ant colony (*Pogonomyrmex barbatus*). *Anim. Behav.* **34**, 1402–1419. (doi:10.1016/S0003-3472(86)80211-1)
71. Mikhailov AS. 1993 Collective dynamics in models of communicating populations. In *Interdisciplinary approaches to nonlinear complex systems*, pp. 89–108. Berlin, Germany: Springer.
72. Delves PJ, Roitt IM. 2000 The immune system. *N. Engl. J. Med.* **343**, 37–49. (doi:10.1056/NEJM200007063430107)
73. Burnet FM. 1959 *The clonal selection theory of acquired immunity*. Nashville, TN: Vanderbilt University Press.
74. Perelson AS, Weisbuch G. 1997 Immunology for physicists. *Rev. Mod. Phys.* **69**, 1219–1267. (doi:10.1103/RevModPhys.69.1219)
75. Jerne NK. 1974 Towards a network theory of the immune system. *Ann. Inst. Pasteur Immunol.* **125C**, 373–389.
76. Barra A, Agliari E. 2010 A statistical mechanics approach to autopoietic immune networks. *J. Stat. Mech. Theory Exp.* **2010**, P07004. (doi:10.1088/1742-5468/2010/07/P07004)
77. Perelson AS. 1989 Immune network theory. *Immunol. Rev.* **110**, 5–36. (doi:10.1111/imr.1989.110.issue-1)
78. Parisi G. 1990 A simple model for the immune network. *Proc. Natl Acad. Sci. USA* **87**, 429–433. (doi:10.1073/pnas.87.1.429)
79. Percus JK, Percus OE, Perelson AS. 1993 Predicting the size of T-cell receptor and antibody combining region from consideration of efficient self-nonself discrimination. *Proc. Natl Acad. Sci. USA* **90**, 1691–1695. (doi:10.1073/pnas.90.5.1691)
80. Behn U. 2007 Idiotypic networks: toward a renaissance? *Immunol. Rev.* **216**, 142–152. (doi:10.1111/imr.2007.216.issue-1)
81. Perelson AS, Oster GF. 1979 Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination. *J. Theor. Biol.* **81**, 645–670. (doi:10.1016/0022-5193(79)90275-3)
82. Mezard M, Parisi G, Virasoro MA. 1987 *Spin glass theory and beyond: an introduction to the replica method and its applications*. Singapore: World Scientific.
83. Detours V, Sulzer B, Perelson A. 1996 Size and connectivity of the idiotypic network are independent of the discreteness of the affinity distribution. *J. Theor. Biol.* **183**, 409–416. (doi:10.1006/jtbi.1996.0231)
84. Stewart J, Varela FJ, Coutinho A. 1989 The relationship between connectivity and tolerance as revealed by computer simulation of the immune network: some lessons for an understanding of autoimmunity. *J. Autoimmun.* **2**, 15–23. (doi:10.1016/0896-8411(89)90113-3)
85. Varela FJ, Coutinho A. 1991 Second generation immune networks. *Immunol. Today* **12**, 159–166. (doi:10.1016/S0167-5699(05)80046-5)
86. Agliari E, Barra A, Guerra F, Moauro F. 2011 A thermodynamic perspective of the immune capabilities. *J. Theor. Biol.* **287**, 48–63. (doi:10.1016/j.jtbi.2011.07.027)
87. Agliari E, Barra A, Del Ferraro G, Guerra F, Tantari D. 2015 Energy in self-directed B lymphocytes: a statistical mechanics perspective. *J. Theor. Biol.* **375**, 21–31. (doi:10.1016/j.jtbi.2014.05.006)
88. Agliari E, Annibale A, Barra A, Coolen ACC, Tantari D. 2017 Retrieving infinite numbers of patterns in a spin-glass model of immune networks. *Europhys. Lett.* **117**, 28003. (doi:10.1209/0295-5075/117/28003)
89. Webb B. 2012 Cognition in insects. *Phil. Trans. R. Soc. B* **367**, 2715–2722. (doi:10.1098/rstb.2012.0218)
90. Dornhaus A, Franks NR. 2008 Individual and collective cognition in ants and other insects (Hymenoptera: Formicidae). *Myrmecol. News* **11**, 215–226.
91. McShea DW. 2002 A complexity drain on cells in the evolution of multicellularity. *Evolution* **56**, 441–452. (doi:10.1111/evo.2002.56.issue-3)
92. Delgado J, Solé RV. 1997 Collective-induced computation. *Phys. Rev. E* **55**, 2338–2344. (doi:10.1103/PhysRevE.55.2338)
93. Kauffman SA. 1984 Emergent properties in random complex automata. *Physica D* **10**, 145–156. (doi:10.1016/0167-2789(84)90257-4)
94. Karlebach G, Shamir R. 2008 Modelling and analysis of gene regulatory networks. *Nat. Rev. Mol. Cell. Biol.* **9**, 770–780. (doi:10.1038/nrm2503)

95. Serra R, Villani M, Graudenzi A, Kauffman SA. 2007 Why a simple model of genetic regulatory networks describes the distribution of avalanches in gene expression data. *J. Theor. Biol.* **246**, 449-460. (doi:10.1016/j.jtbi.2007.01.012)
96. Balleza E, Alvarez-Buylla ER, Chaos A, Kauffman S, Shmulevich I, Aldana M. 2008 Critical dynamics in genetic regulatory networks: examples from four kingdoms. *PLoS ONE* **3**, e2456. (doi:10.1371/journal.pone.0002456)
97. Torres-Sosa C, Huang S, Aldana M. 2012 Criticality is an emergent property of genetic networks that exhibit evolvability. *PLoS Comp. Biol.* **8**, 1002669. (doi:10.1371/journal.pcbi.1002669)
98. Daniels BC, Kim H, Moore D, Zhou S, Smith HB, Karas B, Kauffman SA, Walker SI. 2018 Criticality distinguishes the ensemble of biological regulatory networks. *Phys. Rev. Lett.* **12**, 138102. (doi:10.1103/PhysRevLett.121.138102)
99. Vicsek T, Zafeiris A. 2012 Collective motion. *Phys. Rep.* **517**, 71–140. (doi:10.1016/j.physrep.2012.03.004)

Appendix

A1. Other publications

In this appendix, we give the references of publications that appeared during the development of this PhD project co-authored by the author of this document which are not part of this PhD Thesis.

- Hernández-Navarro, L. & Piñero, J. (2022). Exact face-landing probabilities for bouncing objects: Edge probability in the coin toss and the three-sided die problem. *Phys. Rev. E*, **105**(2), L022201. DOI: 10.1103/PhysRevE.105.L022201
- Duran-Nebreda, S., Pla, J., Vidiella, B., Piñero, J. Conde-Pueyo, N., & Solé, R. (2021). Synthetic lateral inhibition in periodic pattern forming microbial colonies. *ACS Synth.*, **10**(2), 277-285. DOI: 10.1021/acssynbio.0c00318
- Sardanyés, J., Piñero, J., & Solé, R. (2019). Habitat loss-induced tipping points in metapopulations with facilitation. *Popul.*, **61**(4), 436-449. DOI: 10.1002/1438-390X.12020
- Bhat, D., Piñero, J., & Redner, S. (2019). Fixation in fluctuating populations. *JSTAT*, **2019**(6), 063501. DOI: 10.1088/1742-5468/ab1ddd
- Valverde, S., Piñero, J., Corominas-Murtra, B., Montoya, J., Joppa, L., & Solé, R. (2018). The architecture of mutualistic networks as an evolutionary spandrel. *Nat. Ecol. Evol.*, **2**(1), 94-99. DOI: 10.1038/s41559-017-0383-4

A2. The Harvesting Guesser Model

In this section, we showcase the potential of our results concerning the universal bounds on energy harvesting systems discussed in the previous sections. We do so by examining an example called the *two-state guesser*.




Figure 7.1: **Guesser model.** The left shows the system (guesser and environment) with all internal microstates connected by a clique via its baseline mechanisms. We show that this system is equivalent to a two-state system (right) where the guesser only needs to be in the *correct* (1) state, i.e. $x = y$ in order to harvest energy from the environment, and so the system is optimized via some control mechanisms (dashed-lined transitions). If the agent fails to guess the state of the environment (0), then no energy is harvested.

Consider system with a configuration as shown in the left of Figure 7.1. The system is separated into two subsystems, namely the agent (state space X) with n possible states and an environment (state space Y) also with n possible states. Suppose the baseline mechanisms connect all the states of the each subsystem separately via a transition matrix that forms a clique with equal rates, i.e. $R_{x'x}^X = k^X, \forall x, x' \in X$, and $R_{y'y}^Y = k^Y, \forall y, y' \in Y$. We note that, since baseline processes are fully independent of the state of the environment and vice-versa, then the agent-and-environment independency constraint is held.

Our main premise is that, whenever $x = y$, then θ units of energy (in units of $k_B T = 1$) per unit time are transferred to the work reservoir. Thus, the agent is dubbed a *guesser* since it requires to ‘guess’ the state the environment in order to harvest energy from it. For simplicity, we will assume that the system operates in a dissipativeless manner ($q = 0$) and

that no coarse-graining is involved ($s = 0$). The energy pay-off is given by the vector

$$d_{xy} = D\delta_{xy}, \quad (7.1)$$

with D a constant in units of energy per unit time. We also assume that all possible control mechanisms applicable to optimize the efficiency of energy harvesting for the guesser are *local*, i.e., the condition of ‘local control’ holds.

Variable reduction We begin by showing that the $X \times Y$ state space can be reduced to a 2 state system. Firstly, observe that given the above definitions, we can rewrite:

$$R_{x'x}^X = k^X(1 - \delta_{x'x}) - nk^X\delta_{x'x}; \quad R_{y'y}^Y = k^Y(1 - \delta_{y'y}) - nk^Y\delta_{y'y}. \quad (7.2)$$

Thus, the system’s baseline dynamics follow:

$$\partial_t^R p_{xy} = k^X p_x + k^Y p_y - nK p_{xy}, \quad (7.3)$$

where p_x and p_y indicate respective marginal distributions and $K := k^X + k^Y$. Since our example corresponds to a system with a fluctuating environment in which local control and agent-and-environment independence conditions being held, then we have $p^Y = \pi^Y$ (see Section 2.1 of the SM in [PSK23] for a proof). In particular, owing to the symmetries of R^Y , $\pi_y = 1/n$, for all $y \in Y$. Then,

$$p_x = \sum_y p_{x|y} p_y = \frac{1}{n} \sum_y p_{x|y} = \frac{1}{n} = \pi_x,$$

where the latter equality is derived immediately from the symmetries R^X . This can be readily interpreted as having no special $x \in X$ or $y \in Y$ states; only correlation between these two affect the computation of $\Delta \dot{W}^*$

and p^* . Thus, the joint distribution can effectively be given for a single parameter, the *correctness* probability, namely

$$p_{xy} = \begin{cases} \frac{\chi}{n} & x = y \\ \frac{1-\chi}{n(n-1)} & x \neq y \end{cases}, \quad (7.4)$$

with $\chi := \sum_{x,y:x=y} p_{xy}$. This effectively reduces the problem into a two-state system, with only two possible states, namely *correct* ($x = y \Leftrightarrow C = 1$), with probability χ , and *incorrect* ($x \neq y \Leftrightarrow C = 0$), with probability $1 - \chi$. Shifting to this reduced version of the problem implies that now equation (7.3) reduces to:

$$\partial_t^R p_{xy} = \frac{K}{n} - nKp_{xy} \xrightarrow{\sum_{x,y:x=y}} \partial_t^R \chi = K(1 - n\chi). \quad (7.5)$$

For simplicity, we will redefine constant D given in (7.1) as:

$$D = \frac{Kn}{\beta} \theta. \quad (7.6)$$

Maximum power increase and optimal distribution. Using (2.15), we study each term applied to the problem outlined above. First, we derive:

$$\begin{aligned} -\partial_t^R \mathcal{S}(p) &= \frac{K}{n} \ln \left(\prod_{xy} p_{xy} \right) + KnS(p) \\ &= -Kn \left(\chi - \frac{1}{n} \right) \ln \left[\frac{\chi(n-1)}{1-\chi} \right], \end{aligned} \quad (7.7)$$

where we used (7.4) in the last equality. On the other hand, using (7.6), for the second term in (2.15) we obtain:

$$\langle p - \pi, d \rangle = D \left(\chi - \frac{1}{n} \right). \quad (7.8)$$

Combining (7.7), (7.8) and redefinition (7.6), it is possible to arrive at the following optimization problem for χ ,

$$\Delta \dot{W}^* = \max_{\chi} K n \beta^{-1} \mathcal{F}_n(\chi; \theta), \quad (7.9)$$

with

$$\mathcal{F}_n(\chi; \theta) := \left(\chi - \frac{1}{n} \right) \left\{ \theta - \ln \left[\frac{\chi(n-1)}{1-\chi} \right] \right\},$$

and solution for χ^* given, in general, by:

$$\chi_n^*(\theta) = \operatorname{argmax}_{\chi} \mathcal{F}_n(\chi; \theta). \quad (7.10)$$

Next, we will outline a method to solve (7.9) in the LR, FE and M regimes.

Implicit Solution Albeit simple to state, the optimization problem (7.9) does not admit a close form solution. However, an implicit form is easy to derive by setting

$$\left(\frac{\partial \mathcal{F}_n}{\partial \chi} \right)_{\theta} = 0 \Rightarrow \theta(\chi^*) = \frac{\chi^* - \frac{1}{n}}{\chi^*(1-\chi^*)} + \ln \left[\frac{\chi^*(n-1)}{1-\chi^*} \right]. \quad (7.11)$$

Inverting (7.11) is not possible in closed-form, but inverse plots can be studied for different ranges of θ and different values for N , as shown in Figure 7.2. Notice that, for large number of states $n \gg 1$, then (7.11) can be approximated by

$$\theta - \ln n = \frac{1}{(1-\chi^*)} + \ln \left[\frac{\chi^*}{1-\chi^*} \right], \quad (7.12)$$

which can be solved exactly using the Lambert \mathcal{W} function as

$$\chi_{n \gg 1}^*(\theta) \simeq \frac{\mathcal{W} \left(\frac{e^{\theta-1}}{n} \right)}{1 + \mathcal{W} \left(\frac{e^{\theta-1}}{n} \right)}. \quad (7.13)$$

This solution becomes exact for large n as shown using dashed lines in Figure 7.2. As expected, the larger the space of states, the more energetic influx is required for the system to shift into the correct state. Shifting the values of $\theta \rightarrow \theta' = \theta - \ln n$ collapses the behavior of χ^* to a universal curve for large n (inset of the same figure).




Figure 7.2: **Optimal correctness.** Optimal correctness probability (χ^*) as a function of the power yield (θ) for $n = 10, 10^2, 10^4$. The inset shows how a shift in the driving power scale, $\theta \rightarrow \theta' = \theta - \ln(n - 1)$, collapses all curves at the large n limit.

Linear Response Regime By hypothesis, in the LR approximation, $\chi^* \approx 1/n$, i.e. p_{xy} is close to equidistribution. Recall that, in general, this occurs when the energy influx is sufficiently small, then entropic equilibration forces dominate the energetic trade-off in (2.15). Expanding (7.11) around $\chi^* \sim 1/n$ and inverting for $\chi^*(\theta)$, which in turn can be expanded around $\theta \sim 0$, yields

$$\chi_{\text{LR}}^*(\theta) = \frac{1}{n} \left[1 + \frac{n-1}{2n} \theta + \mathcal{O}(\theta^2) \right]. \quad (7.14)$$

Accordingly, we compute the optimal value for the reduced lagrangian \mathcal{F}_n , which gives

$$\mathcal{F}_n^{\text{LR}*}(\theta) = \frac{\theta}{n} \left[\left(\frac{n-1}{n} \right) \frac{\theta}{2} + \mathcal{O}(\theta^2) \right]. \quad (7.15)$$

This result is shown in dashed lines the left plot of Figure 7.3.

Far-from-Equilibrium Regime Similarly, the system is at the FE regime when the agent's state is highly correlated with the environment's, i.e.




Figure 7.3: **Guesser's maximum power increase.** Curves (in log-log scales) of the normalized optimal target for values of $n = 10, 10^2, 10^3$. Dashed/straight lines correspond to the LR and FE approximations given by (7.15) and (7.19), respectively. The inset shows the same plot after shifting the reduced coupling $\theta \rightarrow \theta' = \theta - \ln(n - 1)$.

$\chi^* \sim 1$. This occurs when the energy influx is high, since under this circumstances the entropic contribution in (2.15) becomes of second order, and the agent is capable of spending most of the time in the correct state.

In order to solve for this regime, let us consider an approximate form for \mathcal{F}_N , namely

$$\mathcal{F}_n(\chi, \theta) \simeq \mathcal{F}_N^{\text{FE}}(\chi, \theta) := - \left(1 - \frac{1}{n}\right) \ln \left[\frac{\chi(n-1)}{1-\chi} \right] + \theta \left(\chi - \frac{1}{n} \right), \quad (7.16)$$

where we have used $\chi \sim 1$ in the prefactor of the entropic contribution. This allows for a direct solution to the optimization problem, i.e.

$$\left(\frac{\partial \mathcal{F}_N^{\text{FE}}}{\partial \chi} \right)_{\theta} = 0 \Rightarrow \chi_{\text{FE}}^*(\theta) = \frac{1}{2} \left[1 + \sqrt{1 - \frac{4(n-1)}{n\theta}} \right], \quad (7.17)$$

where the negative root is discarded since it contradicts $\chi^* \sim 1$. Using a large θ expansion, we write an asymptotic expression for the true $\chi^*(\theta)$ at the high information regime as:

$$\chi_{\text{FE}}^*(\theta) = 1 - \left(1 - \frac{1}{N}\right) \theta^{-1} + \mathcal{O}(\theta^{-2}). \quad (7.18)$$

Equivalently, finding the maximum reduced lagrangian value as a power-series of θ^{-1} can be achieved by simple substitution of (7.18) into (7.16), yielding

$$\mathcal{F}_N^{\text{FE}*}(\theta) = \left(1 - \frac{1}{N}\right) \left[\theta - \ln(eN\theta) + \left(1 - \frac{1}{N}\right) \frac{1}{\theta} + \mathcal{O}(\theta^{-2}) \right], \quad (7.19)$$

which is too shown in straight lines the left plot of Figure 7.3.

Controlled solution In this section, we study the thermodynamics a particular instance of the system governed by both baseline and control processes, i.e. $\bar{R} = R + R'$, where R' refers to the control mechanism transition rates. Introduce:

$$R_{x'x}^Y \Big|_{y'=y} = [k' \delta_{x'y} + k' e^{-\phi} (1 - \delta_{x'y})] (1 - \delta_{x'x}), \quad (7.20)$$

such that $R'_{(x'y')(xy)} = \delta_{y'y} R_{x'x}^Y$. Heuristically, this rate matrix defines a *guessing* process that always shifts the agent's state $x \rightarrow x'$ either to the 'correct' state (i.e. $x' = y$) or works away from it due to thermal fluctuations. Notice the above obeys local detailed balance:

$$\frac{R_{yx}^Y}{R_{xy}^Y} \Big|_{x \neq y} = e^{\phi}.$$

The thermodynamic interpretation of the parameter can be put in terms such that $\phi = -\beta w' > 0$ with $w' < 0$ the work done *by* the agent (hence the negative sign) to tilt the transition from $x \neq y \rightarrow x = y$. In other

words, $w' < 0$ represents an expenditure of energy from the agent that is required in order to proceed with the state-matching mechanism.

Let us study the stationary distribution,

$$0 = \partial_t \bar{\pi}_{xy} = k^X \sum_{y'} (\bar{\pi}_{xy'} - \bar{\pi}_{xy}) + k^Y \sum_{x'} (\bar{\pi}_{x'y} - \bar{\pi}_{xy}) + k' \sum_{x'} (\bar{\pi}_{x'y} - e^{-\phi} \bar{\pi}_{xy}) (1 - \delta_{x'y}) \delta_{xy}, \quad (7.21)$$

Owing to symmetry, both marginal probabilities are equidistributed, $\bar{\pi}_x = \bar{\pi}_y = 1/n$. This allows for a simplification as before:

$$\bar{\pi}_{xy} = \begin{cases} \frac{\chi^{\text{st}}}{n} & , x = y \\ \frac{1 - \chi^{\text{st}}}{n(n-1)} & , x \neq y \end{cases}; \quad (7.22)$$

where χ^{st} now stands for the agent's stationary probability of sitting at the same state as the environment's, provided both baseline and control mechanisms. Thus, we obtain χ^{st} after marginalizing (7.21) over states $x = y$, which gives

$$\chi^{\text{st}} = \frac{\frac{1}{n} + \rho}{1 + \rho[1 + (n-1)e^{-\phi}]}, \quad (7.23)$$

where we have defined the adimensional coupling:

$$\rho := \frac{k'}{Kn} = \frac{k'}{(k^X + k^Y)n}.$$

Increase of power yield: finding the average increase power flux is now a straightforward computation using the definition, $\Delta \dot{W}(\bar{\pi}) = \langle \bar{\pi}, w + w' \rangle - \langle \pi, w \rangle$, and (7.23), which yields $\Delta \dot{W}(\bar{\pi}) = KN\beta^{-1} \bar{\mathcal{F}}_n(\theta, \phi, \rho)$, where

$$\bar{\mathcal{F}}_n(\theta, \phi, \rho) = \left(\chi^{\text{st}} - \frac{1}{n} \right) \theta - \rho \phi [1 - (1 + e^{-\phi}) \chi^{\text{st}}]. \quad (7.24)$$

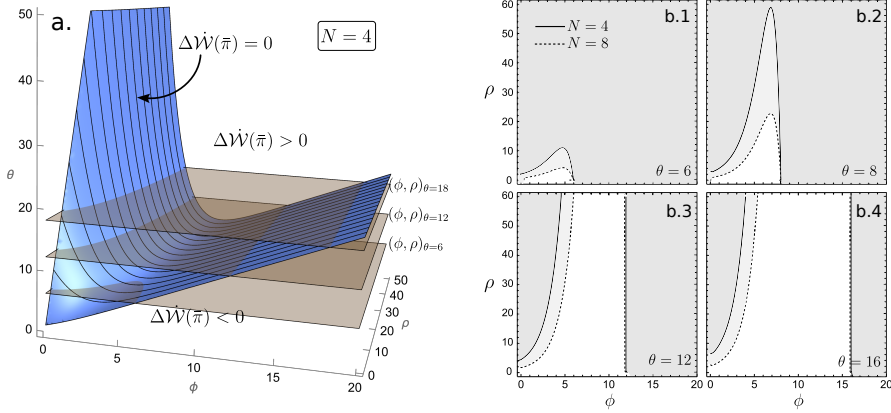


Figure 7.4: Guesser phases. (a) Shows (in blue) the sheet for which $\Delta\dot{W}(\bar{\pi}) = 0$. Above/below this sheet the net power difference for the controlled agent is positive/negative. The orange transparent planes show how different sections of the three-dimensional functions (each at values of $\theta = 6, 12, 18$) separate the space of parameters $\{\phi, \rho\}_\theta$ into phases of effective/ineffective control ($\Delta\dot{W}(\bar{\pi}) > 0/\Delta\dot{W}(\bar{\pi}) < 0$). This plot is done for $N = 4$. Subfigures (b.1), (b.2), (b.3) and (b.4) show a projection of the latter for four theta values, $\theta = 6, 8, 12$ and 16 , respectively. Straight (dashed) lines separate effective from ineffective control regimes for $n = 4$ ($n = 8$).

Here, χ^{st} is left as the implicit function of $\{n; \theta, \phi, \rho\}$ given by (7.23). **Efficiency Patterns:** having readily computed the maximum power yield, (7.9), we now may define the efficiency of any controlled system with respect to the optimal performance as

$$\eta = \frac{\Delta\dot{W}(\bar{\pi})}{\Delta\dot{W}^*}. \quad (7.25)$$

where η will depend on both the baseline and control parameters, as well as system size (n). Pertaining to our current analysis, $\eta_n(\theta, \phi, \rho) = \bar{\mathcal{F}}_n(\theta, \phi, \rho)/\mathcal{F}_n(\theta)$. The resulting patterns are numerically obtained, depicted and described in Figure 7.4 and Figure 7.5. As shown, for specific regions of the parameter space, the system performs at zero or nega-

tive efficiencies, i.e. $\Delta\dot{W}(\bar{\pi}) < 0$, while, for the same parameter values, $\Delta\dot{W}^* \geq 0$. This means that the choice of control mechanisms can perform better or poorly depending on the parameter regions. In Figure 7.4, the region of reduced control parameters $\{\rho, \phi\}$ for which $\eta < 0$ sits under the sheet (given by condition $\eta_n(\theta, \phi, \rho) = 0$). Here parameter θ parameter is drawn along the vertical axis. Figs. 7.4b.1-b.4 show slices of constant θ for different N values. Here, shaded regions display negative efficiency. As expected, such regions grow with N . This is because the larger the search space for the agent, the more energy it is required to pump out energy to account for its entropic cost.

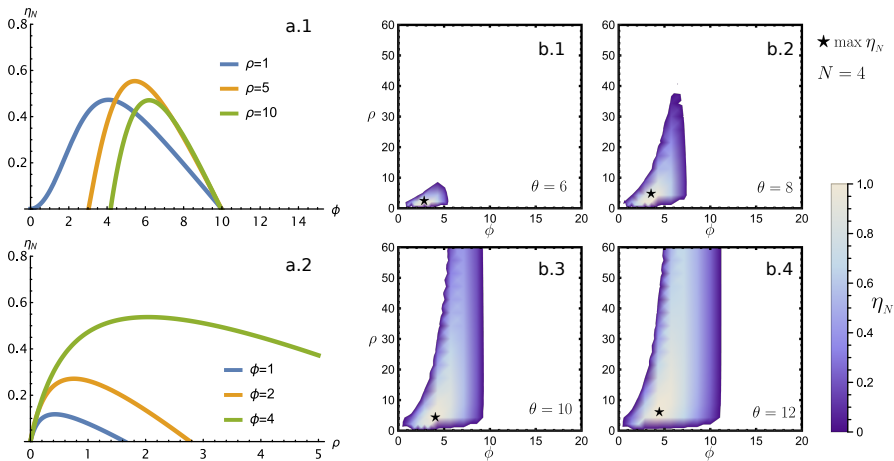


Figure 7.5: Guesser efficiency patterns. Subplots (a.1)-(a.2) show the non-monotonic behavior of $\eta_n(\theta = 10, \phi, \rho)$, for $n = 10$. (a.1) displays three curves for η_N against ϕ , each fixing $\rho = 1, 5, 10$; (a.2) displays three η_n against ρ , each fixing $\phi = 1, 2, 4$. Note that there are regions where the efficiency becomes negative, i.e., an ineffective controller ($\Delta\dot{W}(\bar{\pi}) < 0 \Leftrightarrow \eta_N < 0$). Both subplots show a maximum efficiency value. Subfigures (b.1)–(b.4) are density maps of efficiency values $\eta_n > 0$ for $n = 4$ with $\theta = 6, 8, 10, 12$, respectively. Blank spaces are inefficient ($\eta_n < 0$) regions. The maximum efficiency values attained in each case are computed numerically and shown on each plot with the symbol \star .

Another interesting result is the presence of a global efficiency maxima given our example of a control process. Figure 7.5 shows different efficiency curves and density plots for different values of θ sweeping around the values of (ϕ, ρ) . In all cases, a maxima is reached at some particular value (shown in the figure) that sits within the effective region.

Bibliography

- [ABMM16] Armen E Allahverdyan, Sanasar G Babajanyan, NH Martirosyan, and AV Melkikh. Adaptive heat engine. *Physical Review Letters*, 117(3):030601, 2016.
- [Ada10] Andrew Adamatzky. *Physarum machines: computers from slime mould*, volume 74. World Scientific, 2010.
- [AEM06] David Alonso, Rampal S Etienne, and Alan J McKane. The merits of neutral theory. *Trends in ecology & evolution*, 21(8):451–457, 2006.
- [Alb89] Pere Alberch. The logic of monsters: evidence for internal constraint in development and evolution. *Geobios*, 22:21–57, 1989.
- [AM17] Michael Assaf and Baruch Meerson. Wkb theory of large deviations in stochastic populations. *Journal of Physics A: Mathematical and Theoretical*, 50(26):263001, 2017.
- [And88a] Philip W Anderson. Spin glass i: A scaling law rescued. *Physics Today*, 41(1):9–11, 1988.
- [And88b] Philip W Anderson. Spin glass ii: Is there a phase transition? *Physics Today*, 41(3):9–11, 1988.
- [And88c] Philip W Anderson. Spin glass iii: Theory raises its head. *Physics Today*, 41(6):9, 1988.

- [And88d] Philip W Anderson. Spin glass iv: Glimmerings of trouble. *Physics Today*, 41(9):9–11, 1988.
- [And89a] Philip W Anderson. Spin glass v: Real power brought to bear. *Physics Today*, 42(7):9–11, 1989.
- [And89b] Philip W Anderson. Spin glass vi: Spin glass as cornucopia. *Physics Today*, 42(9):9–11, 1989.
- [And90] Philip W Anderson. Spin glass vii: Spin glass as a paradigm. *Physics Today*, 43(3):9–11, 1990.
- [Arb12] Michael A Arbib. *Brains, machines, and mathematics*. Springer Science & Business Media, 2012.
- [AS06] Camille Stephan-Otto Attolini and Peter F Stadler. Evolving towards the hypercycle: A spatial model of molecular evolution. *Physica D: Nonlinear Phenomena*, 217(2):134–141, 2006.
- [ASG⁺16] Sandro Azaele, Samir Suweis, Jacopo Grilli, Igor Volkov, Jayanth R Banavar, and Amos Maritan. Statistical mechanics of ecological systems: Neutral theory and beyond. *Reviews of Modern Physics*, 88(3):035003, 2016.
- [AW13] Armen E Allahverdyan and QA Wang. Adaptive machine and its thermodynamic costs. *Physical Review E*, 87(3):032139, 2013.
- [BA10] Adriano Barra and Elena Agliari. A statistical mechanics approach to autopoietic immune networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2010(07):P07004, 2010.
- [Bak13] Per Bak. *How nature works: the science of self-organized criticality*. Springer Science & Business Media, 2013.

- [BAP⁺12] Antoine Bérut, Artak Arakelyan, Artyom Petrosyan, Sergio Ciliberto, Raoul Dillenschneider, and Eric Lutz. Experimental verification of Landauer's principle linking information and thermodynamics. *Nature*, 483(7388):187–189, 2012.
- [Bau22] Daniel Baumann. *Cosmology*. Cambridge University Press, 2022.
- [BBN96] Maarten C Boerlijst, Sebastian Bonhoeffer, and Martin A Nowak. Viral quasi-species and recombination. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 263(1376):1577–1584, 1996.
- [BCLP16] Anthony Bartolotta, Sean M Carroll, Stefan Leichenauer, and Jason Pollack. Bayesian second law of thermodynamics. *Physical Review E*, 94(2):022102, 2016.
- [Beg08] John M Beggs. The criticality hypothesis: how local cortical networks might optimize information processing. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 366(1864):329–343, 2008.
- [Ben73] Charles H Bennett. Logical reversibility of computation. *IBM journal of Research and Development*, 17(6):525–532, 1973.
- [BH95] Maarten C Boerlijst and Pauline Hogeweg. Spatial gradients enhance persistence of hypercycles. *Physica D: Nonlinear Phenomena*, 88(1):29–39, 1995.
- [BL16] František Baluška and Michael Levin. On having no head: cognition throughout biological systems. *Frontiers in psychology*, 7:902, 2016.

- [BM14] György Buzsáki and Kenji Mizuseki. The log-dynamic brain: how skewed distributions affect network operations. *Nature Reviews Neuroscience*, 15(4):264–278, 2014.
- [BMV12] Vincenzo Bonifaci, Kurt Mehlhorn, and Girish Varma. Physarum can compute shortest paths. *Journal of theoretical biology*, 309:121–133, 2012.
- [Boo54] George Boole. *An investigation of the laws of thought: on which are founded the mathematical theories of logic and probabilities*. Walton and Maberly, 1854.
- [BP95] Per Bak and Maya Paczuski. Complexity, contingency, and criticality. *Proceedings of the National Academy of Sciences*, 92(15):6689–6696, 1995.
- [Bro15] Judith L Bronstein. *Mutualism*. Oxford University Press, USA, 2015.
- [BSR01] Madeleine Beekman, David JT Sumpter, and Francis LW Ratnieks. Phase transition between disordered and ordered foraging in pharaoh’s ants. *Proceedings of the National Academy of Sciences*, 98(17):9703–9706, 2001.
- [BTW88] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality. *Physical review A*, 38(1):364, 1988.
- [Cas99] Ted J Case. Illustrated guide to theoretical ecology. *Ecology*, 80(8):2848–2848, 1999.
- [CDF⁺20] Scott Camazine, Jean-Louis Deneubourg, Nigel R Franks, James Sneyd, Guy Theraula, and Eric Bonabeau. Self-organization in biological systems. In *Self-Organization in Biological Systems*. Princeton university press, 2020.

- [Chi04] Dante R Chialvo. Critical brain networks. *Physica A: Statistical Mechanics and its Applications*, 340(4):756–765, 2004.
- [Chi10] Dante R Chialvo. Emergent complex neural dynamics. *Nature physics*, 6(10):744–750, 2010.
- [CM05] Kim Christensen and Nicholas R Moloney. *Complexity and criticality*, volume 1. World Scientific Publishing Company, 2005.
- [CP99] B Jack Copeland and Diane Proudfoot. Alan turing’s forgotten ideas in computer science. *Scientific American*, 280(4):98–103, 1999.
- [Cro98] Gavin E Crooks. Nonequilibrium measurements of free energy differences for microscopically reversible markovian systems. *Journal of Statistical Physics*, 90:1481–1487, 1998.
- [Cro99] Gavin E Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Physical Review E*, 60(3):2721, 1999.
- [Cro00] Gavin E Crooks. Path-ensemble averages in systems driven far from equilibrium. *Physical review E*, 61(3):2361, 2000.
- [CT91] Thomas M Cover and Joy A Thomas. Information theory and statistics. *Elements of information theory*, 1(1):279–335, 1991.
- [Dar59] Charles Darwin. *On the origin of species*. Arcturus Publishing Lt. (2015), 1859.
- [Dar92] Charles Darwin. *The formation of vegetable mould, through the action of worms: with observations on their habits*. J. Murray, 1892.

- [dB88] RJ de Boer. Symmetric idiotypic networks: connectance and switching, stability, and suppression. In *Theoretical Immunology: the proceedings of the theoretical immunology workshop, Part Two*, volume 3, page 265. Addison-Wesley Publishing Company, 1988.
- [DB16] Steven Diamond and Stephen Boyd. Cvxpy: A python-embedded modeling language for convex optimization. *The Journal of Machine Learning Research*, 17(1):2909–2913, 2016.
- [DBH89] Rob J De Boer and Pauline Hogeweg. Stability of symmetric idiotypic networks: a critique of hoffmann’s analysis. *Bulletin of mathematical biology*, 51(2):217–222, 1989.
- [DBKP90] RJ De Boer, IG Kevrekidis, and AS Perelson. A simple idiotypic network model with complex dynamics. *Chemical engineering science*, 45(8):2375–2382, 1990.
- [DD05] Christian De Duve. *Singularities: landmarks on the pathways of life*. Cambridge University Press, 2005.
- [DJR⁺08] Gustavo Deco, Viktor K Jirsa, Peter A Robinson, Michael Breakspear, and Karl Friston. The dynamic brain: from spiking neurons to neural masses and cortical fields. *PLoS computational biology*, 4(8):e1000092, 2008.
- [DLBS10] Audrey Dussutour, Tanya Latty, Madeleine Beekman, and Stephen J Simpson. Amoeboid organism solves complex nutritional challenges. *Proceedings of the National Academy of Sciences*, 107(10):4607–4611, 2010.
- [DNB19] Salva Duran-Nebreda and George W Bassel. Plant behaviour in response to the environment: information processing in the solid state. *Philosophical Transactions of the Royal Society B*, 374(1774):20180370, 2019.

- [DNBM⁺16] Salva Duran-Nebreda, Adriano Bonforti, Raúl Montañez, Sergi Valverde, and Ricard Solé. Emergence of proto-organisms from bistable stochastic differentiation and adhesion. *Journal of The Royal Society Interface*, 13(117):20160108, 2016.
- [DNS15] Salva Duran-Nebreda and Ricard Solé. Emergence of multicellularity in a model of cell growth, death and aggregation under size-dependent selection. *Journal of the Royal Society Interface*, 12(102):20140982, 2015.
- [DW58] Charles Darwin and Alfred Wallace. On the tendency of species to form varieties; and on the perpetuation of varieties and species by natural means of selection. *Zoological Journal of the Linnean Society*, 3(9):45–62, 1858.
- [EFG⁺07] Jean-Pierre Eckmann, Ofer Feinerman, Leor Grunendlinger, Elisha Moses, Jordi Soriano, and Tsvi Tlusty. The physics of living neural networks. *Physics Reports*, 449(1-3):54–76, 2007.
- [Eig13] Manfred Eigen. *From strange simplicity to complex familiarity: a treatise on matter, information, life and thought*. OUP Oxford, 2013.
- [EL21] Mo R Ebrahimkhani and Michael Levin. Synthetic living machines: A new window on life. *Isience*, 24(5):102505, 2021.
- [Eng13] Jeremy L England. Statistical physics of self-replication. *The Journal of chemical physics*, 139(12):09B623_1, 2013.
- [ES77] Manfred Eigen and Peter Schuster. A principle of natural self-organization: Part a: Emergence of the hypercycle. *Naturwissenschaften*, 64:541–565, 1977.

- [ES78a] Manfred Eigen and Peter Schuster. The hypercycle: A principle of natural self-organization part b: The abstract hypercycle. *Naturwissenschaften*, 65:7–41, 1978.
- [ES78b] Manfred Eigen and Peter Schuster. The hypercycle: a principle of natural self-organization part c: The realistic hypercycle. *The Science of Nature*, 65(7):341–369, 1978.
- [ES79] Manfred Eigen and Peter Schuster. *The hypercycle: a principle of natural self-organization*. Springer Science & Business Media, 1979.
- [ES82] Manfred Eigen and Peter Schuster. Stages of emerging life: five principles of early organization. *Journal of molecular evolution*, 19:47–61, 1982.
- [ES12] Manfred Eigen and Peter Schuster. *The hypercycle: a principle of natural self-organization*. Springer Science & Business Media, 2012.
- [Far90] J Doyne Farmer. A rosetta stone for connectionism. *Physica D: Nonlinear Phenomena*, 42(1-3):153–187, 1990.
- [GGT92] Deborah M Gordon, Brian C Goodwin, and Lynn EH Trainor. A parallel distributed model of the behaviour of ant colonies. *Journal of theoretical Biology*, 156(3):293–307, 1992.
- [Gor99] Deborah M Gordon. *Ants at work: how an insect society is organized*. Simon and Schuster, 1999.
- [Gou90] Stephen Jay Gould. *Wonderful life: the Burgess Shale and the nature of history*. WW Norton & Company, 1990.
- [Gou02] Stephen Jay Gould. *The structure of evolutionary theory*. Harvard University Press, Cambridge, MA, USA., 2002.

- [GST12] Scott F Gilbert, Jan Sapp, and Alfred I Tauber. A symbiotic view of life: we have never been individuals. *The Quarterly review of biology*, 87(4):325–341, 2012.
- [Hat99] Takahiro Hatano. Jarzynski equality for the transitions between nonequilibrium steady states. *Physical Review E*, 60(5):R5017, 1999.
- [HCM23] Rudolf Hanel and Bernat Corominas-Murtra. The typical set and entropy in stochastic systems with arbitrary phase space growth. *Entropy*, 25(2):350, 2023.
- [HG14] Janina Hesse and Thilo Gross. Self-organized criticality as a fundamental property of neural systems. *Frontiers in systems neuroscience*, 8:166, 2014.
- [Hil13] Terrell L Hill. *Free energy transduction and biochemical cycle kinetics*. Courier Corporation, 2013.
- [HL15] Paul G Higgs and Niles Lehman. The rna world: molecular cooperation at the origins of life. *Nature Reviews Genetics*, 16(1):7–17, 2015.
- [Hop82] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.
- [Hop84] John E Hopcroft. Turing machines. *Scientific American*, 250(5):86–E9, 1984.
- [Hop94] JJ Hopfield. Physics, computation, and why biology looks so different. *Journal of Theoretical Biology*, 171(1):53–60, 1994.
- [HS88] Josef Hofbauer and Karl Sigmund. The theory of evolution and dynamical systems. mathematical aspects of

selection. *CAMBRIDGE UNIVERSITY PRESS, NEW YORK, NY(USA). 1988., 1988.*

- [HS01] Takahiro Hatano and Shin-ichi Sasa. Steady-state thermodynamics of langevin systems. *Physical review letters*, 86(16):3463, 2001.
- [Hub11] Stephen P Hubbell. The unified neutral theory of biodiversity and biogeography (mpb-32). In *The Unified Neutral Theory of Biodiversity and Biogeography (MPB-32)*. Princeton University Press, 2011.
- [Hut65] G Evelyn Hutchinson. *The ecological theater and the evolutionary play*. Yale University Press, 1965.
- [ID20] Sosuke Ito and Andreas Dechant. Stochastic time evolution, information geometry, and the cramér-rao bound. *Physical Review X*, 10(2):021056, 2020.
- [IDC⁺21] Shamreen Iram, Emily Dolson, Joshua Chiel, Julia Pelesko, Nikhil Krishnan, Özenç Güngör, Benjamin Kuznets-Speck, Sebastian Deffner, Efe Ilker, Jacob G Scott, et al. Controlling the speed and trajectory of evolution with counterdiabatic driving. *Nature Physics*, 17(1):135–142, 2021.
- [IMS91] EM Izhikevich, AS Mikhailov, and NA Sveshnikov. Memory, learning and neuromediators. *BioSystems*, 25(4):219–229, 1991.
- [Ito18] Sosuke Ito. Stochastic thermodynamic interpretation of information geometry. *Physical review letters*, 121(3):030605, 2018.
- [Jar97] Christopher Jarzynski. Equilibrium free-energy differences from nonequilibrium measurements: A master-equation approach. *Physical Review E*, 56(5):5018, 1997.

- [Jay65] Edwin T Jaynes. Gibbs vs boltzmann entropies. *American Journal of Physics*, 33(5):391–398, 1965.
- [Jer74] Niels K Jerne. Towards a network theory of the immune system. *Ann. Immunol.*, 125:373–389, 1974.
- [JL06] Eva Jablonka and Marion J Lamb. The evolution of information in the major transitions. *Journal of theoretical biology*, 239(2):236–246, 2006.
- [JL14] Eva Jablonka and Marion J Lamb. *Evolution in four dimensions, revised edition: Genetic, epigenetic, behavioral, and symbolic variation in the history of life*. MIT press, 2014.
- [Kau69a] Stuart Kauffman. Homeostasis and differentiation in random genetic control networks. *Nature*, 224(5215):177–178, 1969.
- [Kau69b] Stuart A Kauffman. Metabolic stability and epigenesis in randomly constructed genetic nets. *Journal of theoretical biology*, 22(3):437–467, 1969.
- [Kau71] Stuart A Kauffman. Cellular homeostasis, epigenesis and replication in randomly aggregated macromolecular systems. 1971.
- [Kau86] Stuart A Kauffman. Autocatalytic sets of proteins. *Journal of theoretical biology*, 119(1):1–24, 1986.
- [Kau91] Stuart A Kauffman. Antichaos and adaptation. *Scientific American*, 265(2):78–85, 1991.
- [Kau93] Stuart A Kauffman. *The origins of order: Self-organization and selection in evolution*. Oxford University Press, USA, 1993.

- [Kau00] Stuart A Kauffman. *Investigations*. Oxford University Press, 2000.
- [KBLB20] Sam Kriegman, Douglas Blackiston, Michael Levin, and Josh Bongard. A scalable pipeline for designing reconfigurable organisms. *Proceedings of the National Academy of Sciences*, 117(4):1853–1859, 2020.
- [Kim83] Motoo Kimura. *The neutral theory of molecular evolution*. Cambridge University Press, 1983.
- [KK+95] Stuart Kauffman, Stuart A Kauffman, et al. *At home in the universe: The search for laws of self-organization and complexity*. Oxford University Press, USA, 1995.
- [KL87] Stuart Kauffman and Simon Levin. Towards a general theory of adaptive walks on rugged landscapes. *Journal of theoretical Biology*, 128(1):11–45, 1987.
- [KRBN10] Pavel L Krapivsky, Sidney Redner, and Eli Ben-Naim. *A kinetic view of statistical physics*. Cambridge University Press, 2010.
- [Lan61] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM journal of research and development*, 5(3):183–191, 1961.
- [Lan10] Nick Lane. *Life ascending: the ten great inventions of evolution*. Profile books, 2010.
- [Lan15] Nick Lane. *The vital question: energy, evolution, and the origins of complex life*. WW Norton & Company, 2015.
- [Lan22] Nick Lane. *Transformer: The Deep Chemistry of Life and Death*. Profile Books, 2022.
- [LD20] Michael Levin and Daniel C Dennett. Cognition all the way down. *Aeon Essays*, 2020.

- [LKR11] Yonatan Loewenstein, Annerose Kuras, and Simon Rumpel. Multiplicative dynamics underlie the emergence of the log-normal distribution of spine sizes in the neocortex in vivo. *Journal of Neuroscience*, 31(26):9481–9488, 2011.
- [Lli02] Rodolfo R Llinás. *I of the vortex: From neurons to self*. MIT press, 2002.
- [Lot20] Alfred J Lotka. Analytical note on certain rhythmic relations in organic systems. *Proceedings of the National Academy of Sciences*, 6(7):410–415, 1920.
- [Lot22] Alfred J Lotka. Contribution to the energetics of evolution. *Proceedings of the National Academy of Sciences*, 8(6):147–151, 1922.
- [Lot56] Alfred J Lotka. *Elements of mathematical biology*. Dover Publications, 1956.
- [LR90] Harvey S Leff and Andrew F Rex. Maxwell’s demon. *Entropy, information, computing*, 1990.
- [LR02] Harvey Leff and Andrew F Rex. *Maxwell’s Demon 2 Entropy, Classical and Quantum Information, Computing*. CRC Press, 2002.
- [Mag88] Anne E Magurran. *Ecological diversity and its measurement*. Princeton university press, 1988.
- [Mag03] Anne E Magurran. *Measuring Biological Diversity*. John Wiley & Sons, 2003.
- [Mar81] Lynn Margulis. *Symbiosis in cell evolution: Life and its environment on the early earth*. W.H Freeman and Co., 1981.

- [MAS00] Alan McKane, David Alonso, and Ricard V Solé. Mean-field stochastic theory for species-rich assembled communities. *Physical Review E*, 62(6):8466, 2000.
- [MAS04] Alan J McKane, David Alonso, and Ricard V Solé. Analytic solution of hubbell’s model of local community dynamics. *Theoretical Population Biology*, 65(1):67–73, 2004.
- [Max72] J Clerk Maxwell. Theory of heat. *Astronomical register*, 10:107–107, 1872.
- [May19] Robert M May. *Stability and complexity in model ecosystems*. Princeton university press, 2019.
- [MB11] Thierry Mora and William Bialek. Are biological systems poised at criticality? *Journal of Statistical Physics*, 144:268–302, 2011.
- [MJ12] Dibyendu Mandal and Christopher Jarzynski. Work and information processing in a solvable model of maxwell’s demon. *Proceedings of the National Academy of Sciences*, 109(29):11641–11645, 2012.
- [MNNM14] Melanie JI Müller, Beverly I Neugeboren, David R Nelson, and Andrew W Murray. Genetic drift opposes mutualism during spatial population expansion. *Proceedings of the National Academy of Sciences*, 111(3):1037–1042, 2014.
- [Mor68] Harold J Morowitz. Energy flow in biology; biological organization as a problem in thermal physics. 1968.
- [Mor03] Simon Conway Morris. *Life’s solution: inevitable humans in a lonely universe*. Cambridge university press, 2003.

- [Mor04] Harold J Morowitz. *The emergence of everything: How the world became complex*. Oxford University Press, USA, 2004.
- [Mor13] Harold J Morowitz. *Entropy for biologists: an introduction to thermodynamics*. Elsevier, 2013.
- [MP43] Warren S McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5:115–133, 1943.
- [MSS95] John Maynard-Smith and Eörs Szathmáry. *The major transitions in evolution*. WH Freeman Spektrum Oxford, UK., 1995.
- [Mun18] Miguel A Munoz. Colloquium: Criticality and dynamical scaling in living systems. *Reviews of Modern Physics*, 90(3):031001, 2018.
- [MW01] Robert H MacArthur and Edward O Wilson. *The theory of island biogeography*, volume 1. Princeton university press, 2001.
- [NAS22] National Academy of Sciences NAS. *Physics of Life*. National Academic Press, 2022.
- [NGM91] Sean Nee, Richard D Gregory, and Robert M May. Core and satellite species: theory and artefacts. *Oikos*, pages 83–87, 1991.
- [NP89] Gregoire Nicolis and Ilya Prigogine. Exploring complexity an introduction. 1989.
- [NYT00] Toshiyuki Nakagaki, Hiroyasu Yamada, and Ágota Tóth. Maze-solving by an amoeboid organism. *Nature*, 407(6803):470–470, 2000.

- [OTW17] Thomas E Ouldridge and Pieter Rein Ten Wolde. Fundamental costs in the production and destruction of persistent polymer copies. *Physical review letters*, 118(15):158103, 2017.
- [OVS19] Aina Ollé-Vila and Ricard Solé. Cellular heterogeneity results from indirect effects under metabolic tradeoffs. *Royal Society Open Science*, 6(9):190281, 2019.
- [Par90] Giorgio Parisi. A simple model for the immune network. *Proceedings of the National Academy of Sciences*, 87(1):429–433, 1990.
- [Pel11] Luca Peliti. Statistical mechanics in a nutshell. In *Statistical Mechanics in a Nutshell*. Princeton University Press, 2011.
- [PHS15] Juan MR Parrondo, Jordan M Horowitz, and Takahiro Sagawa. Thermodynamics of information. *Nature physics*, 11(2):131–139, 2015.
- [PJ04] Natasha Paul and Gerald F Joyce. Minimal self-replicating systems. *Current opinion in chemical biology*, 8(6):634–639, 2004.
- [PME16] Nikolay Perunov, Robert A Marsland, and Jeremy L England. Statistical physics of adaptation. *Physical Review X*, 6(2):021036, 2016.
- [PP21] Luca Peliti and Simone Pigolotti. *Stochastic Thermodynamics: An Introduction*. Princeton University Press, 2021.
- [PPP92] JK Percus, Ora E Percus, and Alan S Perelson. Probability of self-nonsel self discrimination. In *Theoretical and Experimental Insights into Immunology*, pages 63–70. Springer, 1992.

- [PPP93] Jerome K Percus, Ora E Percus, and Alan S Perelson. Predicting the size of the t-cell receptor and antibody combining region from consideration of efficient self-nonsel self discrimination. *Proceedings of the National Academy of Sciences*, 90(5):1691–1695, 1993.
- [PRS22] Jordi Piñero, Sidney Redner, and Ricard Solé. Fixation and fluctuations in two-species cooperation. *Journal of Physics: Complexity*, 3(1):015011, 2022.
- [PS18] Jordi Piñero and Ricard Solé. Nonequilibrium entropic bounds for darwinian replicators. *Entropy*, 20(2):98, 2018.
- [PS19] Jordi Piñero and Ricard Solé. Statistical physics of liquid brains. *Philosophical Transactions of the Royal Society B*, 374(1774):20180376, 2019.
- [PSK23] Jordi Piñero, Ricard Solé, and Artemy Kolchinsky. Universal bounds and thermodynamic tradeoffs in nonequilibrium energy harvesting. *arXiv preprint arXiv:2303.04975*, 2023.
- [PUZ20] Giorgio Parisi, Pierfrancesco Urbani, and Francesco Zamponi. *Theory of simple glasses: exact solutions in infinite dimensions*. Cambridge University Press, 2020.
- [PW97] Alan S Perelson and Gérard Weisbuch. Immunology for physicists. *Reviews of modern physics*, 69(4):1219, 1997.
- [PWMM04] David Peak, Jevin D West, Susanna M Messinger, and Keith A Mott. Evidence for complex, collective dynamics and emergent, distributed computation in plants. *Proceedings of the National Academy of Sciences*, 101(4):918–922, 2004.

- [Ras60] Nicolas Rashevsky. *Mathematical Biophysics: physico-mathematical foundations of biology. Vol 1 and 2*. Dover, USA, 1960.
- [Rau66] David M Raup. Geometric analysis of shell coiling: general problems. *Journal of paleontology*, pages 1178–1190, 1966.
- [Ric06] Robert E Ricklefs. The unified neutral theory of biodiversity: do the numbers add up? *Ecology*, 87(6):1424–1431, 2006.
- [RWB43] Arturo Rosenblueth, Norbert Wiener, and Julian Bigelow. Behavior, purpose and teleology. *Philosophy of science*, 10(1):18–24, 1943.
- [Sac20] Andrea Sacchetti. Francesco carlini: Kepler’s equation and the asymptotic solution to singular differential equations. *Historia Mathematica*, 53:1–32, 2020.
- [Sag67] Lynn Sagan. On the origin of mitosing cells. *Journal of theoretical biology*, 14(3):225–IN6, 1967.
- [Sag14] Takahiro Sagawa. Thermodynamic and logical reversibilities revisited. *Journal of Statistical Mechanics: Theory and Experiment*, 2014(3):P03025, 2014.
- [SAM02] Ricard V Solé, David Alonso, and Alan McKane. Self-organized instability in complex ecosystems. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1421):667–681, 2002.
- [SB12] Ricard Solé and Jordi Bascompte. Self-organization in complex ecosystems.(mpb-42). In *Self-Organization in Complex Ecosystems.(MPB-42)*. Princeton University Press, 2012.

- [Sch44] Erwin Schrodinger. *What is life? The physical aspect of the living cell*. 1944.
- [Sch11] Peter Schuster. The mathematics of darwins theory of evolution: 1859 and 150 years later. *The Mathematics of Darwins Legacy*, pages 27–66, 2011.
- [Sch16] Peter Schuster. Some mechanistic requirements for major transitions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 371(1701):20150439, 2016.
- [Sei08] Udo Seifert. Stochastic thermodynamics: principles and perspectives. *The European Physical Journal B*, 64:423–431, 2008.
- [Sei12] Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on progress in physics*, 75(12):126001, 2012.
- [SEMSRM17] Ben Shirt-Ediss, Sara Murillo-Sánchez, and Kepa Ruiz-Mirazo. Framing major prebiotic transitions as stages of protocell development: three challenges for origins-of-life research. *Beilstein journal of organic chemistry*, 13(1):1388–1395, 2017.
- [SET98] Frank Schweitzer, Werner Ebeling, and Benno Tilch. Complex motion of brownian particles with energy depots. *Physical Review Letters*, 80(23):5044, 1998.
- [Set21] James P Sethna. *Statistical mechanics: entropy, order parameters, and complexity*, volume 14. Oxford University Press, USA, 2021.
- [SFS18] Naoto Shiraishi, Ken Funo, and Keiji Saito. Speed limit for classical stochastic processes. *Physical review letters*, 121(7):070601, 2018.

- [SG00] Ricard V Solé and Brian C Goodwin. *Signs of life: How complexity pervades biology*. Basic books New York, 2000.
- [Sha48] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- [Sim49] Edward H Simpson. Measurement of diversity. *nature*, 163(4148):688–688, 1949.
- [SM95] Ricard V Solé and Octavio Miramontes. Information at the edge of chaos in fluid neural networks. *Physica D: Nonlinear Phenomena*, 80(1-2):171–180, 1995.
- [SM16] Eric Smith and Harold J Morowitz. *The origin and nature of life on earth: the emergence of the fourth geosphere*. Cambridge University Press, 2016.
- [SMCL⁺96] Ricard V Solé, Susanna Manrubia Cuevas, Bartolo Luque, Jordi Delgado, and Jordi Bascompte. Phase transitions and complex systems: Simple, nonlinear models capture complex systems at the edge of chaos. 1996.
- [SMDN15] Ricard V Solé, Raúl Montañez, and Salva Duran-Nebreda. Synthetic circuit designs for earth terraformation. *Biology Direct*, 10(1):1–10, 2015.
- [SMF19] Ricard Solé, Melanie Moses, and Stephanie Forrest. *Liquid brains, solid brains*, 2019.
- [SMG93] Ricard V Solé, Octavio Miramontes, and Brian C Goodwin. Oscillations and chaos in ant societies. *Journal of theoretical Biology*, 161(3):343–357, 1993.
- [Smi99] John Maynard Smith. The idea of information in biology. *The Quarterly Review of Biology*, 74(4):395–400, 1999.

- [Smi00] John Maynard Smith. The concept of information in biology. *Philosophy of science*, 67(2):177–194, 2000.
- [Smo92] Lee Smolin. Did the universe evolve? *Classical and Quantum Gravity*, 9(1):173–190, 1992.
- [Smo04] Lee Smolin. Cosmological natural selection as the explanation for the complexity of the universe. *Physica A: Statistical Mechanics and Its Applications*, 340(4):705–713, 2004.
- [SMS95] Eörs Szathmáry and John Maynard-Smith. The major evolutionary transitions. *Nature*, 374(6519):227–232, 1995.
- [SN13] Daniel L Stein and Charles M Newman. *Spin glasses and complexity*, volume 4. Princeton University Press, 2013.
- [Sol11] Ricard Solé. *Phase transitions*, volume 3. Princeton University Press, 2011.
- [Soy01] Valery N Soyfer. The consequences of political dictatorship for russian science. *Nature Reviews Genetics*, 2(9):723–729, 2001.
- [SP13] Woodrow L Shew and Dietmar Plenz. The functional benefits of criticality in the cortex. *The neuroscientist*, 19(1):88–100, 2013.
- [SS79] Montgomery Slatkin and John Maynard Smith. Models of coevolution. *The Quarterly Review of Biology*, 54(3):233–263, 1979.
- [SS82] Jörg Swetina and Peter Schuster. Self-replication with errors: A model for polynucleotide replication. *Biophysical chemistry*, 16(4):329–345, 1982.

- [SS97] Eörs Szathmáry and John Maynard Smith. From replicators to reproducers: the first major transitions leading to life. *Journal of theoretical biology*, 187(4):555–571, 1997.
- [SS01] István Scheuring and Eörs Szathmáry. Survival of replicators with parabolic growth tendency and exponential decay. *Journal of Theoretical Biology*, 212(1):99–105, 2001.
- [SS06] Josep Sardanyés and Ricard V Solé. Bifurcations and phase transitions in spatially extended two-member hypercycles. *Journal of Theoretical Biology*, 243(4):468–482, 2006.
- [SS07] Josep Sardanyés and Ricard V Solé. Spatio-temporal dynamics in simple asymmetric hypercycles under weak parasitic coupling. *Physica D: Nonlinear Phenomena*, 231(2):116–129, 2007.
- [SS13] Luís F Seoane and Ricard V Solé. A multiobjective optimization approach to statistical mechanics. *arXiv preprint arXiv:1310.6372*, 2013.
- [SS18] Luís F Seoane and Ricard V Solé. Information theory, predictability and the emergence of complex life. *Royal Society open science*, 5(2):172221, 2018.
- [SSE21] Ricard Solé, Josep Sardanyés, and Santiago F Elena. Phase transitions in virology. *Reports on Progress in Physics*, 84(11):115901, 2021.
- [SW49] Claude E Shannon and Warren Weaver. A mathematical model of communication. *Urbana, IL: University of Illinois Press*, 11:11–20, 1949.

- [Sza06] Eörs Szathmáry. The origin of replicators and reproducers. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 361(1474):1761–1776, 2006.
- [Sza15] Eörs Szathmáry. Toward major evolutionary transitions theory 2.0. *Proceedings of the National Academy of Sciences*, 112(33):10104–10111, 2015.
- [Szi29] Leo Szilard. Über die entropieverminderung in einem thermodynamischen system bei eingriffen intelligenter wesen. *Zeitschrift für Physik*, 53(11):840–856, 1929.
- [TBD98] Guy Theraulaz, Eric Bonabeau, and Jean-Louis Deneubourg. The origin of nest complexity in social insects. *Complexity*, 3(6):15–25, 1998.
- [Tho17] J Arthur Thomson. On growth and form, 1917.
- [TKN07] Atsushi Tero, Ryo Kobayashi, and Toshiyuki Nakagaki. A mathematical model for adaptive transport network in path finding by true slime mold. *Journal of theoretical biology*, 244(4):553–564, 2007.
- [Ton12] David Tong. Statistical physics. *Lecture Notes*, 2012.
- [TSU+10] Shoichi Toyabe, Takahiro Sagawa, Masahito Ueda, Eiro Muneyuki, and Masaki Sano. Experimental demonstration of information-to-energy conversion and validation of the generalized jarzynski equality. *Nature physics*, 6(12):988–992, 2010.
- [Tur36] Alan Mathison Turing. On computable numbers, with an application to the entscheidungsproblem. *J. of Math*, 58(345-363):5, 1936.
- [Tur11] J Scott Turner. Termites as models of swarm cognition. *Swarm intelligence*, 5:19–43, 2011.

- [Tur19] J Scott Turner. Homeostasis as a fundamental principle for a coherent theory of brains. *Philosophical Transactions of the Royal Society B*, 374(1774):20180373, 2019.
- [VBHM03] Igor Volkov, Jayanth R Banavar, Stephen P Hubbell, and Amos Maritan. Neutral theory and relative species abundance in ecology. *Nature*, 424(6952):1035–1037, 2003.
- [VCMP+09] Sergi Valverde, Bernat Corominas-Murtra, Andrea Perna, Pascale Kuntz, Guy Theraulaz, and Ricard V Solé. Percolation in insect nest networks: evidence for optimal wiring. *Physical Review E*, 79(6):066106, 2009.
- [vK86] Günter von Kiedrowski. A self-replicating hexadeoxynucleotide. *Angewandte Chemie International Edition in English*, 25(10):932–935, 1986.
- [VNK12] John Von Neumann and Ray Kurzweil. *The computer and the brain*. Yale university press, 2012.
- [Vol26] V Volterra. Fluctuations in the abundance of species considered mathematically. *Nature*, 118:558–560, 1926.
- [Vol83] Mikhail Vladimirovich Volkenstein. *Biophysics*. Mir Publishers Moscow, 1983.
- [Vol13a] Mikhail Vladimirovich Volkenstein. *General Biophysics: Volume I*, volume 1. Academic Press, 2013.
- [Vol13b] Mikhail Vladimirovich Volkenstein. *General Biophysics: Volume II*, volume 2. Academic Press, 2013.
- [VWKK22] Vitaly Vanchurin, Yuri I Wolf, Eugene V Koonin, and Mikhail I Katsnelson. Thermodynamics of evolution and the origin of life. *Proceedings of the National Academy of Sciences*, 119(6):e2120042119, 2022.

- [VZ12] Tamás Vicsek and Anna Zafeiris. Collective motion. *Physics reports*, 517(3-4):71–140, 2012.
- [Wad14] Conrad Hal Waddington. *The strategy of the genes*. Routledge, 2014.
- [Wag00] Jorge Wagensberg. Complexity versus uncertainty: The question of staying alive. *Biology and philosophy*, 15:493–508, 2000.
- [Wie19] Norbert Wiener. *Cybernetics or Control and Communication in the Animal and the Machine*. MIT press, 2019.
- [Wil07] David M Wilkinson. *Fundamental processes in ecology: an earth systems approach*. OUP Oxford, 2007.
- [Yeo92] Julia M Yeomans. *Statistical mechanics of phase transitions*. Clarendon Press, 1992.
- [Zim20] Vincent Zimmern. Why brain criticality is clinically relevant: a scoping review. *Frontiers in neural circuits*, 14:54, 2020.
- [ZO87] Wojciech S Zielinski and Leslie E Orgel. Autocatalytic synthesis of a tetranucleotide analogue. *Nature*, 327(6120):346–347, 1987.
- [Zur89] Wojciech H Zurek. Thermodynamic cost of computation, algorithmic complexity and the information metric. *Nature*, 341(6238):119–124, 1989.
- [Zur90] Wojciech H Zurek. Complexity, entropy and the physics of information. *Santa Fe Institute Studies in the Sciences of Complexity*, 1990.

