




UNIVERSITÀ
DI PAVIA

DIPARTIMENTO DI FISICA

Corso di Laurea Magistrale in Scienze Fisiche

**Coevolutionary Games on Structured
Populations:
A Statistical Mechanics Approach**

*Giochi coevolutivi su popolazioni strutturate:
un approccio di meccanica statistica*

Tesi di Laurea di:
Dario Demetri

Relatore:
Prof. Giacomo Livan

Correlatori:
Dott. Giacomo Frigerio
Dott. Federico Maria Quetti

A.A. 2025–2026

Abstract

Cooperation in systems of many agents is a collective state whose stability cannot be understood from individual incentives alone. This phenomenon becomes even more subtle when the incentives themselves are not fixed, but are reshaped by the population's behaviour through the decision-making of individuals. In such systems, strategies and games coevolve: the population determines the strategic environment that will in turn govern its future evolution. Previous work introduced this mechanism in well-mixed populations, where each pair of individuals interacts with the same probability, and showed that endogenous feedback can generate chimera games, namely stable regimes that would be unstable or impossible in corresponding scenarios where incentives remain fixed over time.

This thesis develops a statistical mechanics treatment of this problem by studying the endogenous feedback model on structured populations, where interactions are local and topologically constrained. The results show that chimera-game regimes survive both on square lattices and on complex network topologies. On lattices, the dynamics displays two distinct time scales: the global cooperation level relaxes relatively quickly, while cooperative clusters continue to reorganise over much longer times. The analysis also reveals sharp changes in cluster morphology depending on the game reached by the feedback dynamics. On complex networks, clustering plays a central role: it can sustain cooperation in regimes where low-clustering networks lead to full defection, while also generating higher-order local correlations that reduce the accuracy of the pair-approximation prediction.

Sommario

La cooperazione in sistemi composti da molti agenti è uno stato collettivo la cui stabilità non può essere compresa a partire dai soli incentivi individuali. Questo fenomeno diventa ancora più sottile quando gli incentivi stessi non sono fissati, ma vengono modificati dal comportamento della popolazione attraverso le decisioni dei singoli individui. In questi sistemi, strategie e giochi coevolvono: la popolazione determina l'ambiente strategico che a sua volta governerà la sua evoluzione futura. Un lavoro precedente ha introdotto questo meccanismo in popolazioni ben mescolate, dove ogni coppia di individui interagisce con la stessa probabilità, mostrando che il feedback endogeno può generare giochi chimera, cioè regimi stabili che sarebbero instabili o impossibili in scenari corrispondenti in cui gli incentivi restano fissati nel tempo.

Questa tesi sviluppa un trattamento di questo problema basato sulla meccanica statistica, studiando il modello con feedback endogeno su popolazioni strutturate, dove le interazioni sono locali e vincolate da una topologia. I risultati mostrano che i regimi chimera sopravvivono sia su reticoli quadrati sia su reti complesse. Sui reticoli, la dinamica mostra due scale temporali distinte: il livello globale di cooperazione rilassa relativamente rapidamente, mentre i cluster cooperativi continuano a riorganizzarsi su tempi molto più lunghi. L'analisi rivela inoltre cambiamenti netti nella morfologia dei cluster a seconda del gioco raggiunto dalla dinamica di feedback. Sulle reti complesse, il clustering svolge un ruolo centrale: può sostenere la cooperazione in regimi in cui reti a basso clustering portano alla defezione completa, generando al tempo stesso correlazioni locali di ordine superiore che riducono l'accuratezza della predizione basata sull'approssimazione di coppia.

Contents

Introduction	1
1 Game Theory and the Emergence of Cooperation	5
1.1 Complex Systems and Emergence	5
1.2 The Problem of Cooperation	6
1.2.1 The Prisoner's Dilemma	7
1.2.2 The Emergence of Cooperation	8
1.3 Foundations of Game Theory	10
1.3.1 Normal-Form Games	10
1.3.2 Classical Solution Concepts and Evolutionary Equilibria	12
1.4 Evolutionary Game Theory	16
1.4.1 Finite and Infinite Population Dynamics	16
1.4.2 The Replicator Equation	17
1.4.3 Nash Equilibria and Evolutionary Stability	18
1.4.4 Replicator Dynamics in Well-Mixed Populations	19
2 Coevolutionary Game Theory and the Endogenous Feedback Model	28
2.1 Main Classes of Coevolutionary Dynamics	29
2.1.1 Network Coevolution	29
2.1.2 Payoff-Matrix Evolution	31
2.1.3 Behavioural-Rule and Trait Evolution	34
2.2 The Endogenous Feedback Model	36
2.2.1 Model Formulation	37
2.2.2 Well-Mixed Dynamics	39
2.2.3 Emergence of Chimera Games	41
3 Evolutionary Games on Structured Populations	51
3.1 Complex Networks as Models of Structured Populations	52
3.2 Analytical Approaches on Structured Populations	56
3.2.1 The Mean-Field Approximation	56

3.2.2	The Pair Approximation	57
3.2.3	The Replicator Equation on Graphs	61
3.3	Evolutionary Dynamics on Square Lattices	67
3.3.1	Simulation Framework	67
3.3.2	Von Neumann Neighbourhood	69
3.3.3	Moore Neighbourhood	77
3.3.4	Self-Interaction	84
3.4	Evolutionary Dynamics on Complex Networks	87
3.4.1	Erdős–Rényi Networks	87
3.4.2	Watts–Strogatz Networks	91
3.4.3	Barabási–Albert Networks	94
4	The Endogenous Feedback Model on Square Lattices	104
4.1	Chimera Games on Square Lattices	104
4.1.1	Simulation Framework	105
4.1.2	Von Neumann Neighbourhood	106
4.1.3	Moore Neighbourhood	114
4.1.4	Self-Interaction	118
4.2	Cluster Dynamics	118
4.2.1	Fast Global Convergence and Slow Spatial Reorganisation . .	120
4.2.2	Quantifying Cluster Structure	122
4.2.3	Sharp Change in Cluster Morphology	123
5	The Endogenous Feedback Model on Complex Networks	128
5.1	Chimera Games on Complex Networks	128
5.1.1	Erdős–Rényi Networks	129
5.1.2	Watts–Strogatz Networks	131
5.1.3	Barabási–Albert Networks	132
5.2	Clustering and the Breakdown of Pair Approximation	137
5.2.1	The Degree-Preserving Rewiring Procedure	138
5.2.2	Effect of Clustering on Cooperation and Analytical Predictions	139
	Conclusions	143
	Bibliography	146
	Appendix A Endogenous Feedback Simulations in Well-Mixed Populations	151
A.1	Prisoner’s Dilemma as G_2	151
A.2	Stag Hunt as G_2	152

A.3	Snowdrift as G_2	155
Appendix B Endogenous Feedback Simulations on Square Lattices		160
B.1	Von Neumann Neighbourhood	160
B.1.1	Prisoner's Dilemma as G_2	160
B.1.2	Stag Hunt as G_2	161
B.1.3	Snowdrift as G_2	165
B.2	Moore Neighbourhood	166
Appendix C Endogenous Feedback Simulations on Complex Networks		174
C.1	Erdős–Rényi Networks	174
C.2	Watts–Strogatz Networks	174
C.3	Barabási–Albert Networks	175

Introduction

Cooperation in systems of many agents is a collective state whose stability cannot be understood from individual incentives alone [1, 2]. This phenomenon becomes even more subtle when the incentives themselves are not fixed, but are reshaped by the population’s behaviour through the decision-making of individuals. In such systems, strategies and games coevolve: the population determines the strategic environment that will in turn govern its future evolution. Previous work introduced this mechanism in well-mixed populations [3], where each pair of individuals interacts with the same probability, and showed that *endogenous feedback* can generate *chimera games*, namely stable regimes that would be unstable or impossible in corresponding scenarios where incentives remain fixed over time. The aim of this thesis is to investigate whether such regimes survive when interactions are local and structured, and how spatial correlations and network clustering reshape their stability.

The thesis lies at the intersection of statistical physics, evolutionary game theory, network science and complex systems. Classical *game theory* provides the mathematical framework to study strategic interactions among a small number of decision-makers, in order to identify optimal strategies under given rules of interaction [4]. Since its early development, game theory has shown a strong predictive power and has been empirically tested in several domains, ranging from economics and social coordination to evolutionary biology [5, 6, 7].

When one moves from isolated strategic encounters to populations of many interacting agents whose strategies evolve in time, the focus shifts from individual optimal behaviour to collective dynamics. The problem therefore naturally requires a statistical-mechanics perspective: macroscopic outcomes arise from microscopic update rules, stochasticity, local interactions, correlations, and network structure [8, 9, 10]. *Evolutionary game theory* is the mathematical formulation that focuses on determining long-term strategic outcomes from repeated local interactions. The need for this formulation is motivated, for instance, by evolutionary games on lattices, where local interactions can generate phase-transition-like behaviour at the macroscopic level of cooperation [11].

From a physicist’s perspective, the problem of cooperation is especially compelling because strategically interacting populations display many hallmarks of complex systems: nonlinear dynamics, feedback loops, multiscale organisation and emergent order. Cooperation can therefore be studied as a macroscopic collective state generated by microscopic interaction rules and by the topology of the interaction network. This viewpoint motivates the use of tools such as mean-field and pair approximation methods, master and Fokker–Planck equations, finite-size scaling analysis, and

quasistationary methods [12, 13, 14, 15, 16, 17].

In many real-world systems, however, not only strategies evolve over time. Agents may also modify the conditions under which interactions take place: they can reshape their social ties, move in space, or affect environmental variables that determine the rewards, or payoffs, associated with each strategic outcome. The feedback between strategies and interaction environments is formalised within the framework of *coevolutionary games*, where the population and the rules of interaction evolve together [18, 19]. This framework can describe a wide range of adaptive systems, including ecological interactions, financial markets, and urban mobility systems, where individual behaviour and the interaction landscape mutually reshape one another [20, 21, 22].

The thesis focuses on a particular class of coevolutionary mechanisms in which the strategic incentives change endogenously with the state of the population. These incentives are encoded in the payoff matrix, which specifies the rewards associated with the possible outcomes of an interaction. The model studied was developed in a recent work [3] carried out within the research group of Prof. Giacomo Livan at the University of Pavia, in collaboration with the group of Prof. Vito Latora at Queen Mary University of London. In this endogenous feedback framework, the payoff matrix depends on the current fraction of cooperators, so that the population reshapes the strategic incentives that determine its future evolution. This mechanism is relevant, for instance, in online communities, where positive interactions may make cooperative engagement increasingly rewarding [23], or in financial markets, where speculative bubbles may amplify incentives for opportunistic exit and free-riding [24].

This feedback mechanism can generate so-called chimera games: regimes in which stable cooperation arises despite being incompatible with the predictions of standard fixed-game dynamics. A possible real-world analogue of this chimera-like mechanism is liquidity provision in financial markets [25]. Providing liquidity can be interpreted as a cooperation-like behaviour, since it makes trading easier and improves market functioning. At the same time, the liquidity created by market participants can also create opportunities for better-informed traders to exploit liquidity providers. The coexistence of liquidity provision and opportunistic exploitation may therefore appear puzzling in a fixed-game interpretation, but becomes more natural when incentives are allowed to be reshaped endogenously by collective market behaviour.

The original framework developed in the paper focuses on *well-mixed populations*, where each agent is equally likely to interact with any other agent and spatial structure is neglected. This thesis studies whether chimera-game regimes survive in *structured populations*, where interactions are local and topologically constrained, as in *spatial lattices* and *complex networks* [26, 27, 28]. In structured populations, the global cooperation level is no longer sufficient to characterise the state of the system: two populations with the same fraction of cooperators can display different spatial arrangements, leading to different evolutionary stability.

The thesis has three main aims. The first is to test whether chimera-game regimes emerge when moving from a highly idealised well-mixed setting to a more realistic framework of structured populations. The second is to analyse how cooperators organise in space, since their spatial arrangement can determine whether cooperation

resists or collapses under the invasion of defectors. The third is to study the role of clustering in complex networks, since it can affect both the final cooperation level reached by the system and the validity of the analytical prediction based on pair approximation.

From a methodological point of view, the thesis provides a systematic statistical-mechanics analysis of these systems. This is relevant since, although evolutionary and coevolutionary games have been widely studied, a unified statistical-mechanics framework for coevolutionary games on structured populations is still largely missing. The work combines Monte Carlo simulations with analytical approaches such as mean-field theory and pair approximation. Network modelling is used to represent structured populations, including square lattices and different classes of complex networks, such as Erdős–Rényi, Watts–Strogatz and Barabási–Albert networks. The numerical analysis includes systematic parameter scans, repeated independent runs, and averages over different network realisations, in order to obtain robust estimates of cooperation levels and fluctuations. Due to the computational cost of the simulations, part of the work has been carried out using INFN computing clusters.

The numerical simulations presented in this thesis were carried out by the author using Python code developed for this work. The same framework was also used to reproduce and re-analyse benchmark results from the literature on evolutionary and coevolutionary game theory, in order to validate the implementation and ensure consistency across the different models considered. Some analytical derivations, benchmark results or figures are adapted from the literature. In those cases, this is explicitly indicated in the text or in the corresponding figure captions. The original numerical results and illustrations produced for this thesis are indicated accordingly. All the code used to run the simulations is openly available in the GitHub repository:

<https://github.com/demetridario-gif/Spatial-Coevolutionary-Games->

This makes the numerical analysis transparent and reproducible, and allows the simulations to be extended or adapted for future studies.

The thesis is organised as follows. Chapter 1 introduces the theoretical framework for this work. After discussing complex systems and emergence, the focus shifts to the problem of cooperation, introduced through the Prisoner’s Dilemma as the canonical example of a social dilemma. Game theory is then presented as the mathematical framework for strategic interactions, before moving to evolutionary game theory and introducing the replicator equation. Finally, the well-mixed population case is analysed in detail by comparing the analytical predictions of the replicator equation with the results of Monte Carlo simulations. Chapter 2 extends this framework to coevolutionary game theory, where strategies evolve together with additional dynamical variables. After reviewing several coevolutionary mechanisms, the endogenous feedback model is introduced and its behaviour in well-mixed populations is analysed, with particular attention to the emergence of chimera-game regimes.

Chapter 3 marks the transition from well-mixed to structured populations, the main focus of this thesis. Interaction structure is introduced through spatial lattices and complex networks, and the main definitions and structural properties of networks are presented. The analytical tools used to study evolutionary dynamics on graphs are

then discussed, including mean-field methods and pair approximation. In particular, the replicator equation on regular graphs in the weak-selection regime is derived. Finally, evolutionary games are analysed on square lattices and on Erdős–Rényi, Watts–Strogatz and Barabási–Albert networks, providing the reference behaviour against which the endogenous feedback model is later compared.

The original research contribution of the thesis starts in Chapter 4, where the endogenous feedback model is studied on square lattices. The chapter investigates whether chimera-game regimes emerge in spatially structured populations and then analyses the formation, motion and reorganisation of cooperative clusters. Particular attention is devoted to the relation between the global cooperation level and the spatial organisation of cooperative domains, as well as to how cluster morphology changes when the game reached by the dynamics is varied.

Chapter 5 studies the endogenous-feedback model on complex networks, considering Erdős–Rényi, Watts–Strogatz and Barabási–Albert topologies. It investigates whether chimera-game regimes emerge when interactions are topologically constrained beyond regular lattices, and analyses how network structure affects the resulting cooperation level. The chapter then focuses specifically on clustering. In order to isolate its effect from that of the degree distribution, a degree-preserving rewiring procedure is used to vary the clustering coefficient while keeping the degree of every node fixed. This makes it possible to analyse how clustering affects the final cooperation level and how it contributes to the breakdown of the pair-approximation prediction.

The thesis concludes with a discussion of the main results obtained, their scientific relevance, and the future research perspectives opened by this work.

Chapter 1

Game Theory and the Emergence of Cooperation

In this chapter, we introduce the theoretical background for our work. We first discuss complex systems and emergence, describing how macroscopic collective behaviour can arise from local interactions among many individual units. Within this framework, sustained cooperation can be interpreted as a paradigmatic example of collective behaviour emerging at the population scale from individual interactions. We introduce the problem of cooperation through the Prisoner’s Dilemma, the canonical example of a social dilemma, where individual incentives conflict with collective welfare. We then discuss mechanisms that can support cooperation, focusing in particular on repeated interactions and reciprocity. After motivating the problem in this way, we introduce game theory as the mathematical framework used to study strategic interactions among decision-making agents. In this framework, an interaction is strategic when the payoff obtained by each agent depends not only on its own action, but also on the actions chosen by the other agents. In its classical formulation, game theory usually considers fixed strategic situations, in which the available actions and payoff structure are specified in advance. Finally, the chapter shifts from this classical, static formulation to evolutionary game theory, where strategies spread or disappear according to their relative success in the population. In this context, we introduce the replicator equation as the fundamental deterministic description of evolutionary dynamics. We then analyse the main two-strategy games and compare the analytical predictions with stochastic numerical simulations.

1.1 Complex Systems and Emergence

The starting point of *complex systems science* is the observation that many systems in nature and society cannot be understood solely by analysing their components in isolation. This perspective is closely related to statistical physics, where macroscopic properties are often described in terms of collective variables emerging from the interaction of many microscopic degrees of freedom [29]. Building on this idea, and following the conceptual framework introduced by Anderson in his article “More is Different” [30], we can describe a *complex system* as

*“A large collection of **units interacting locally** with each other, giving rise — without explicit coordination — to **emergent macroscopic behaviour** that may be difficult to predict from the behaviour of individual units.”*

The central idea behind Anderson’s principle is that “more is different”: increasing the number of interacting elements does not simply produce a larger version of the same system. Instead, collective interactions may generate qualitatively new macroscopic properties and dynamical behaviours that are absent at the microscopic level. This is analogous to what happens in many-body physics, where phenomena such as magnetic ordering, critical behaviour and phase transitions cannot be understood from the properties of a single particle alone [31, 32]. Examples of *emergent phenomena* are abundant in nature and society. Flocking behaviour in birds emerges from local alignment interactions among individuals and can be modelled in terms of collective motion of self-propelled particles [33]. Traffic jams may spontaneously arise even in the absence of accidents or external obstacles, as a consequence of nonlinear interactions among vehicles [34]. Collective opinion formation may emerge from repeated social interactions and has been studied using tools inspired by statistical physics [35]. Similarly, financial markets display collective fluctuations, correlations and scaling properties [36], while ecosystems and biological populations often display complex collective dynamics generated by the interactions among many individual units [37]. From a physical perspective, complex systems are particularly interesting because they frequently display phenomena analogous to those encountered in statistical mechanics, including collective ordering, metastability, criticality and phase transitions. Understanding how macroscopic order emerges from microscopic interactions therefore constitutes one of the central goals of complex systems science.

1.2 The Problem of Cooperation

Building on the notion of emergence, cooperation can be interpreted as a paradigmatic example of collective behaviour arising from local interactions. Cooperative structures are observed in many biological, social and economic systems, ranging from microbial communities [38], animal societies [2], human institutions [39], collective decision-making [6], and economic exchange [40]. From a theoretical perspective, however, the persistence of cooperation is not straightforward. In many interaction scenarios, individuals may obtain a larger immediate payoff by behaving selfishly than by cooperating. Thus, if agents are driven only by individual payoff maximisation, or if selection favours strategies with higher short-term fitness, one might expect non-cooperative behaviour to spread. This prediction contrasts with the widespread occurrence of cooperation in real systems. The central problem is therefore to understand which mechanisms allow cooperative behaviour to emerge and persist despite incentives to defect. This tension between individual advantage and collective welfare is formalised by the concept of a *social dilemma*. A social dilemma is a strategic situation in which actions that are advantageous from the perspective of individual agents may lead to outcomes that are suboptimal for the population as a whole. In other words, individual incentives and collective opti-

ality are not aligned [41]. From a physical perspective, this makes cooperation an interesting problem because it can be formulated as the emergence of a macroscopic collective state from many microscopic interactions. At the microscopic level, each individual interaction is governed by local incentives; at the macroscopic level, the population may display collective outcomes such as widespread cooperation, widespread defection, or coexistence between the two behaviours. The problem is therefore analogous in spirit to many problems in statistical physics, where one studies how global order, correlations and collective patterns arise from the interaction of many elementary units [29, 32].

1.2.1 The Prisoner's Dilemma

The *Prisoner's Dilemma* is the paradigmatic model of a social dilemma and represents one of the central frameworks for studying cooperation in game theory and evolutionary dynamics. The name originates from a classical thought experiment involving two prisoners accused of a crime and interrogated separately. Each prisoner has two possible actions. The first is to remain silent, thereby cooperating with the accomplice. The second is to confess and betray the accomplice. In the language of game theory, remaining silent is identified with *cooperation*, denoted by C , while confessing is identified with *defection*, denoted by D . If both prisoners remain silent, they cooperate with one another and both receive a moderate sentence. If one prisoner confesses while the other remains silent, the confessing prisoner defects against the cooperator and is rewarded with freedom, or with a strongly reduced sentence. The prisoner who remained silent is instead exploited and receives the harshest punishment. Finally, if both prisoners confess, both defect: each avoids being the only exploited prisoner, but both receive an intermediate sentence, which is worse than mutual silence. As will be discussed in the following section, the mathematical formalisation of this situation is provided by game theory, and more specifically by a symmetric two-player game with two available strategies. The strategic structure can therefore be represented by the two actions *cooperation*, C , and *defection*, D . The payoff received by each player depends on the action chosen by both players. If both players cooperate, each receives the *reward* for mutual cooperation, denoted by R . If both defect, each receives the *punishment* for mutual defection, denoted by P . If one player defects while the other cooperates, the defector receives the *temptation* payoff T , while the cooperator receives the *sucker's* payoff S . The *payoff matrix* can therefore be written as

$$\begin{array}{c|cc} & C & D \\ \hline C & (R, R) & (S, T) \\ D & (T, S) & (P, P) \end{array} \quad (1.1)$$

where the first entry in each pair corresponds to the payoff of the row player and the second entry corresponds to the payoff of the column player. The Prisoner's Dilemma is defined by the payoff ordering

$$T > R > P > S.$$

This ranking captures the essence of the dilemma. The largest payoff is obtained by defecting against a cooperator, since $T > R$. Mutual cooperation is preferable to

mutual defection, since $R > P$. However, cooperating against a defector produces the worst possible outcome, since $S < P$. From the perspective of an isolated rational individual, defection is the dominant strategy. As a matter of fact, if the other player cooperates, defecting gives T instead of R , and is therefore advantageous. If the other player defects, defecting gives P instead of S , and is again preferable. Consequently, independently of the opponent's behaviour, defection always yields a larger individual payoff. Two rational players are therefore driven toward mutual defection, producing the outcome (D, D) in which both players receive the payoff P . The *dilemma* lies in the fact that both players would have obtained a larger payoff through mutual cooperation, since $R > P$. The Prisoner's Dilemma therefore represents a paradigmatic example of a situation in which *individual rationality leads to collective irrationality*. Each player chooses the individually optimal strategy, yet the resulting collective outcome is globally suboptimal. For this reason, the Prisoner's Dilemma has become one of the most influential models in the theoretical study of cooperation and evolutionary game dynamics [42].

1.2.2 The Emergence of Cooperation

The theoretical difficulty posed by cooperation can now be stated more clearly. If defection is individually advantageous as in the Prisoner's Dilemma, why is cooperation so frequently observed? Why do biological organisms, human societies and artificial systems often display cooperative behaviour instead of collapsing into purely selfish exploitation? This question is particularly important because cooperation is not an exceptional phenomenon, but appears across many domains and scales. In biology, cooperative behaviour is observed in microbial communities [38], animal groups and social insects [43, 2]. In human societies, cooperation is essential for economic exchange [40], institutions [39], scientific collaboration [44], and collective decision-making [6]. In ecological and technological systems, global order often emerges from local interactions among many units. Cooperation therefore appears as a widespread collective phenomenon, even though simple game-theoretic reasoning often predicts defection.

The problem of cooperation is a problem of emergence. At the microscopic level, individual agents may follow local rules and respond to local incentives. At the macroscopic level, however, the population may develop cooperative structures, stable patterns or collective regimes. Understanding how cooperative behaviour can emerge despite self-interested incentives is one of the central problems in the study of cooperation dynamics. So far, the discussion has focused on one-shot interactions, where players meet only once and the consequences of their choices are limited to the immediate payoff of that encounter. However, many interactions in biological, social and economic systems are not isolated events. Individuals often meet repeatedly, and when future encounters are expected, the logic of the Prisoner's Dilemma changes substantially [45]. In a one-shot Prisoner's Dilemma, defection is individually rational because actions have no future consequences: a player can exploit the opponent and then leave the interaction. In a repeated setting, instead, present actions can affect future payoffs. Defection can be punished in subsequent rounds, while cooperation can be rewarded by continued cooperation. Repetition therefore

makes reciprocity possible. This mechanism is commonly referred to as the “shadow of the future”. When future interactions are sufficiently important, the short-term gain obtained by exploiting the opponent may be outweighed by the long-term benefit of maintaining a cooperative relationship. Cooperation can therefore become advantageous not because individuals stop pursuing their own interest, but because self-interest is extended over a longer time horizon.

A landmark contribution to this mechanism was provided by Robert Axelrod in his book *The Evolution of Cooperation* [46]. Axelrod organised a computer tournament in which different strategies for the iterated Prisoner’s Dilemma were submitted and made to compete against one another. Each strategy specified how a player should act over repeated interactions, depending on the previous history of play. The most successful strategy was TIT FOR TAT. Its rule is simple: cooperate in the first round, and then repeat the opponent’s previous move. Thus, cooperation is met with cooperation, while defection is answered by defection. Its success relied on four main properties: it is *nice*, because it never defects first; *retaliatory*, because it responds to defection; *forgiving*, because it returns to cooperation once the opponent does; and *clear*, because its behaviour is easy to understand. The main lesson of Axelrod’s tournament is that successful cooperation does not require unconditional altruism. Pure cooperation can be exploited, while purely aggressive behaviour prevents stable cooperative relations from forming. Conditional cooperation, instead, can be robust: it starts cooperatively, punishes exploitation, and allows cooperation to be restored after conflict. Axelrod’s work therefore showed that reciprocity can provide a mechanism for cooperation among self-interested agents. When individuals interact repeatedly, remember previous behaviour, and condition their future actions on past interactions, cooperative strategies may perform well and persist within the population [1].

Axelrod also used historical examples to show that reciprocal cooperation can emerge even under highly adverse conditions. One of the most striking cases is the *live-and-let-live* system that developed in some sectors of trench warfare during the First World War [47]. Although soldiers on opposite sides of the front were formal enemies, local and tacit forms of restraint sometimes emerged between opposing units. The mechanism was not based on sympathy, explicit agreement or institutional enforcement. Rather, it arose from repeated interaction under stable local conditions. The same units often faced one another for extended periods of time, so each side could condition its behaviour on the previous actions of the other. Escalation by one side could be met with retaliation, whereas mutual restraint reduced the risks and costs faced by both. This situation can be interpreted through the logic of the iterated Prisoner’s Dilemma. Mutual restraint plays the role of cooperation, while aggressive escalation corresponds to defection. Since interactions were repeated and retaliation was possible, reciprocal behaviour could emerge even in a context defined by conflict and limited trust. The importance of the live-and-let-live system is that it shows how cooperation does not necessarily require friendship, moral agreement or centralised coordination. Under appropriate conditions, repeated interaction alone can generate local patterns of cooperation among self-interested agents.

1.3 Foundations of Game Theory

Game theory provides the mathematical framework for studying strategic interactions among agents, namely situations in which the payoff of each agent depends not only on its own action, but also on the actions chosen by the others. In this section, we first discuss normal-form games and then introduce the main equilibrium concepts used throughout the thesis. We follow the approach of Leyton-Brown and Shoham [48].

1.3.1 Normal-Form Games

We begin by introducing the notion of finite normal-form games.

Definition 1.1. A finite **normal-form game** is a triple $\mathcal{G} = (N, S, u)$, where

- $N = \{1, \dots, n\}$ is a finite set of players;
- $S = S_1 \times \dots \times S_n$ is a non-empty set of available actions, with S_i the finite strategy set of player i . Each vector $s = (s_1, \dots, s_n) \in S$ is called a *strategy profile*;
- $u = (u_1, \dots, u_n)$ is a vector of payoff functions, with $u_i : S \rightarrow \mathbb{R}$ assigning a payoff to player i for each strategy profile.

For a two-player game with strategy sets S_1 and S_2 , the normal-form representation can be written as a *payoff matrix* of size $|S_1| \times |S_2|$. Rows correspond to strategies in S_1 , columns to strategies in S_2 , and each entry contains the payoff pair $(u_1(s_1, s_2), u_2(s_1, s_2))$. In the two-strategy case $S_1 = S_2 = \{C, D\}$ considered above, this reduces to a 2×2 payoff matrix, as in Eq. (1.1). A class of particular interest for this thesis is given by symmetric games. Symmetry means that both players have the same strategy set and that payoffs depend only on the strategies played, not on the identity of the players.

Definition 1.2. A normal-form game $\mathcal{G} = (N, S, u)$ is **symmetric** if all players have the same strategy set, namely $S_1 = \dots = S_n$, and payoffs are invariant under permutations of player labels. More precisely, for any strategy profile $s = (s_1, \dots, s_n)$ and for any permutation π of the player set N , one has

$$u_i(s_1, \dots, s_n) = u_{\pi(i)}(s_{\pi^{-1}(1)}, \dots, s_{\pi^{-1}(n)}).$$

In the two-player case, this condition reduces to

$$u_1(a, b) = u_2(b, a),$$

for every pair of strategies a, b . In this thesis, we focus mainly on symmetric two-strategy games with strategies cooperation, C , and defection, D . In this scenario, the payoff matrix can be written as

$$A = \begin{pmatrix} R & S \\ T & P \end{pmatrix}, \quad (1.2)$$

where rows refer to the strategy of the first player and columns to the strategy of the opponent. The entries have the following interpretation:

- R is the reward for mutual cooperation;
- S is the sucker's payoff of a cooperator facing a defector;
- T is the temptation payoff of a defector facing a cooperator;
- P is the punishment for mutual defection.

Canonical Two-Strategy Games

Although the Prisoner's Dilemma introduced in Sect. 1.2.1 is the most famous example, it belongs to a broader class of two-player games describing different forms of conflict between individual incentives and collective welfare. These games are commonly classified according to the relative ordering of the four payoffs T , R , P and S , which determine the strategic structure of the interaction and the evolutionary dynamics of cooperation.

Snowdrift Game (SD) The *Snowdrift Game*, also known as the *Chicken Game*, describes a situation in which unilateral cooperation is preferable to mutual defection. In this scenario, the payoff ranking is

$$T > R > S > P.$$

The intuition behind the name is the following. Two drivers are trapped on opposite sides of a snowdrift. Each can either leave the car and shovel the snow (cooperate) or remain inside the car (defect). If both cooperate, they share the effort required to clear the road and both obtain the reward R . If one cooperates while the other defects, the defector obtains the largest payoff T , since it benefits from the other player's effort without paying any cost, while the cooperator receives the sucker's payoff S . Finally, if both defect and remain in their cars, neither driver can pass through the snowdrift and both receive the lowest payoff P . Unlike in the Prisoner's Dilemma, cooperating against a defector is still preferable to mutual defection. In the snowdrift interpretation, shovelling alone is costly, but it is better than both drivers remaining inactive and staying trapped. Indeed, the payoff of unilateral cooperation, S , is larger than the payoff of mutual defection, P , namely $S > P$.

Stag Hunt Game (SH) The *Stag Hunt Game* represents a coordination problem in which cooperation is collectively optimal but risky. The canonical payoff ordering is

$$R > T > P > S.$$

The name derives from Rousseau's classical example of two hunters choosing between hunting a stag together or hunting hares individually. Hunting a stag requires cooperation between both hunters but yields the largest payoff R . Hunting a hare

instead provides a smaller but guaranteed individual payoff T . If one player cooperates while the other defects, the cooperator fails to catch the stag and receives the lowest payoff S , whereas the defector successfully hunts the hare and receives T . If both defect, both hunt hares individually and obtain the intermediate payoff P . The dilemma therefore lies in the trade-off between a safe but suboptimal strategy and a risky but potentially more rewarding cooperative outcome. Unlike the Prisoner's Dilemma, cooperation is not opposed by direct selfish incentives, but rather by strategic uncertainty and coordination risk.

Harmony Game (HG) The *Harmony Game* represents the opposite scenario of the Prisoner's Dilemma. The payoff ordering is

$$R > T > S > P.$$

In this case, cooperation is always advantageous, independently of the opponent's action. Mutual cooperation yields the largest payoff, while defection never provides a strategic advantage over cooperation. This situation can be illustrated by two workers performing a common task. If both cooperate, the task is completed efficiently and both obtain the highest reward R . If one defects while the other cooperates, the defector still receives a smaller benefit T , but cooperation remains preferable from the individual perspective. Unlike the previous games, the Harmony Game does not represent a true social dilemma because there is no conflict between individual and collective interests. Cooperation is simultaneously the individually optimal and collectively optimal strategy. For this reason, the Harmony game is often interpreted as the symmetric counterpart of the Prisoner's Dilemma, in which the payoff structure favours cooperation rather than defection.

Together, these games provide a general framework for studying cooperation, competition, coordination and strategic conflict, and will play a key role throughout this thesis.

1.3.2 Classical Solution Concepts and Evolutionary Equilibria

Game theory uses different solution concepts to characterise outcomes of strategic interactions. In this section, we focus on three different notions: Pareto efficiency, Nash equilibrium and evolutionarily stable strategies. They capture different forms of stability or optimality.

Pareto Efficiency

Pareto efficiency is a criterion for collective optimality. It compares outcomes according to whether some players can be made better off without making any other player worse off.

Definition 1.3. Let $\mathcal{G} = (N, S, u)$ be a normal-form game. Given two strategy profiles $s, s' \in S$, we say that s' **Pareto dominates** s if

$$u_i(s') \geq u_i(s) \quad \forall i \in N,$$

and

$$u_j(s') > u_j(s) \quad \text{for at least one } j \in N. \quad (1.3)$$

Definition 1.4. A strategy profile $s \in S$ is **Pareto efficient** if there exists no $s' \in S$ that Pareto dominates s .

Thus, a strategy profile is Pareto efficient if there exists no alternative strategy profile that makes at least one player better off without making at least one other player worse off. Pareto efficiency is therefore a criterion of collective efficiency, rather than a prediction of which strategy profile rational players will choose.

Nash Equilibrium

Nash equilibrium formalises strategic stability. An outcome is stable in this sense if no player can increase its payoff by changing strategy unilaterally, given the strategies of the other players.

Let $s = (s_1, \dots, s_n) \in S$ be a strategy profile. For player i , we denote by

$$s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$$

the strategies chosen by all players except i . The global profile can therefore be written as

$$s = (s_i, s_{-i}).$$

Definition 1.5. A strategy $s_i^* \in S_i$ is a **best response** to s_{-i} if

$$u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i}) \quad \forall s_i \in S_i.$$

Definition 1.6. A strategy profile $s^* = (s_1^*, \dots, s_n^*) \in S$ is a **Nash equilibrium** if each player's strategy is a best response to the strategies of the others, namely if

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \quad \forall s_i \in S_i, \quad \forall i \in N.$$

At a Nash equilibrium, no unilateral deviation is profitable. The concept therefore captures individual strategic stability, but not necessarily collective efficiency.

Definition 1.7. A Nash equilibrium $s^* = (s_1^*, \dots, s_n^*)$ is a **strict Nash equilibrium** if every unilateral deviation gives a strictly lower payoff, namely if

$$u_i(s_i^*, s_{-i}^*) > u_i(s_i, s_{-i}^*) \quad \forall s_i \neq s_i^*, \quad \forall i \in N.$$

Thus, strict Nash equilibrium is a stronger notion of strategic stability: each player has a unique incentive to keep its equilibrium strategy, given the strategies of the others.

The Prisoner's Dilemma illustrates the distinction between strategic stability and collective efficiency. In this game, the payoff ordering is $T > R > P > S$. The outcome (D, D) is a Nash equilibrium because, when the other player defects, switching from D to C decreases the payoff from P to S . However, (D, D) is not Pareto efficient, since both players would obtain a higher payoff under mutual cooperation: $R > P$. Conversely, (C, C) is Pareto efficient, but it is not a Nash equilibrium: if one player cooperates, the other can increase its payoff by defecting, since $T > R$. Thus, the Prisoner's Dilemma shows that strategic stability and collective optimality need not coincide.

Evolutionarily Stable Strategies

The concept of an evolutionarily stable strategy, or ESS, was introduced by Maynard Smith and Price to describe strategies that are robust against invasion by rare mutants [49]. It is defined for symmetric games and provides an evolutionary refinement of Nash equilibrium.

Consider a resident population in which almost all individuals play strategy A . Let $A' \neq A$ be a rare mutant strategy. The resident strategy A is evolutionarily stable if, when the mutant is sufficiently rare, the resident obtains a strictly higher expected payoff than the mutant.

Definition 1.8. Consider a symmetric two-player normal-form game and let A be a strategy. The strategy A is an **evolutionarily stable strategy** if, for every alternative strategy $A' \neq A$, there exists $\bar{\varepsilon} > 0$ such that, for all $0 < \varepsilon < \bar{\varepsilon}$,

$$u(A, (1 - \varepsilon)A + \varepsilon A') > u(A', (1 - \varepsilon)A + \varepsilon A').$$

Here $(1 - \varepsilon)A + \varepsilon A'$ denotes a population in which a fraction $1 - \varepsilon$ plays the resident strategy A , while a small fraction ε plays the mutant strategy A' . Thus, $u(A, (1 - \varepsilon)A + \varepsilon A')$ is the expected payoff of the resident strategy against the mixed population, whereas $u(A', (1 - \varepsilon)A + \varepsilon A')$ is the expected payoff of the mutant strategy against the same population. By linearity of expectation, the ESS condition is equivalent to

$$(1 - \varepsilon)u(A, A) + \varepsilon u(A, A') > (1 - \varepsilon)u(A', A) + \varepsilon u(A', A').$$

Since the inequality is required only for sufficiently small ε , this is equivalent to requiring that, for every $A' \neq A$, either

$$u(A, A) > u(A', A),$$

or

$$u(A, A) = u(A', A) \quad \text{and} \quad u(A, A') > u(A', A').$$

Definition 1.9. A strategy A is a **weak evolutionarily stable strategy** if, for every alternative strategy $A' \neq A$, either

$$u(A, A) > u(A', A),$$

or

$$u(A, A) = u(A', A) \quad \text{and} \quad u(A, A') \geq u(A', A').$$

The strict ESS condition requires the resident strategy to outperform any rare mutant strategy. In the weak definition, the mutant cannot perform better than the resident, but it may perform equally well in the neutral case.

For the symmetric two-strategy game, whose payoff matrix is given by Eq. (1.2), consider a resident population in which almost all individuals play cooperation C , while a small fraction ε plays defection D . The expected payoff of a cooperator




Figure 1.1: Relationship between Nash equilibria, strict symmetric Nash equilibria and evolutionarily stable strategies. Every evolutionarily stable strategy is a Nash equilibrium, but not every Nash equilibrium is evolutionarily stable. Conversely, every strict symmetric Nash equilibrium is evolutionarily stable, but not every evolutionarily stable strategy is necessarily a strict Nash equilibrium. *Illustration by the author.*

against this population is $\Pi_C = (1 - \varepsilon)R + \varepsilon S$, whereas the expected payoff of a defector is $\Pi_D = (1 - \varepsilon)T + \varepsilon P$. Cooperation is resistant to invasion by defectors if

$$(1 - \varepsilon)R + \varepsilon S > (1 - \varepsilon)T + \varepsilon P.$$

For sufficiently small ε , this condition reduces to

$$R > T.$$

If the leading-order terms are equal, namely if $R = T$, then the next-order condition is

$$S > P.$$

Therefore, cooperation C is an evolutionarily stable strategy if $R > T$, or, in the neutral case $R = T$, if $S > P$. Analogously, defection D is evolutionarily stable if $P > S$, or, in the neutral case $P = S$, if $T > R$. In the standard Prisoner's Dilemma, defection is the dominant strategy and satisfies the ESS condition, since a rare cooperator facing a population of defectors obtains S , while defectors obtain P , with $P > S$. The relationship between Nash equilibrium and evolutionary stability is therefore one of refinement. From the definition, every ESS must be a Nash equilibrium: if a resident strategy A is evolutionarily stable, then no alternative strategy $A' \neq A$ can obtain a higher payoff against A . Hence, A is a best response to itself. However, the converse is not true in general: a Nash equilibrium may fail to be evolutionarily stable if it can be invaded by a rare mutant strategy. Every strict symmetric Nash equilibrium is instead an ESS, because the resident strategy obtains a strictly higher payoff against itself than any alternative strategy does. Figure 1.1 illustrates this inclusion structure.

1.4 Evolutionary Game Theory

The previous sections introduced cooperation as a strategic problem arising from the tension between individual incentives and collective outcomes. While classical game theory analyses this tension at the level of fixed strategic interactions, evolutionary game theory studies how behaviours change in frequency within a population. Strategies are therefore treated as dynamical entities: behaviours, rules, traits or phenotypes whose abundance changes according to their success, relative to alternative strategies. In biological contexts, this success is naturally interpreted in terms of *reproductive fitness*. In social or economic contexts, the same mathematical framework can describe imitation, learning, cultural transmission, or the diffusion of behavioural rules. In this sense, the payoff obtained by individuals adopting a given strategy is interpreted as the fitness of that strategy. Evolutionary dynamics then describe how the distribution of strategies changes as a consequence of payoff differences within the population. The central question is therefore no longer only whether cooperation is rational for an isolated player, but whether cooperative strategies can invade, persist, and remain stable under selection. Evolutionary game theory connects the microscopic level of strategic interactions with the macroscopic evolution of the population, providing a natural bridge between game theory, population dynamics and the study of emergent behaviour.

1.4.1 Finite and Infinite Population Dynamics

Evolutionary dynamics can be formulated at different levels of description, depending in particular on population size and on the role played by stochastic fluctuations. A first important distinction is between finite-population models and infinite-population models. In finite-population models, the number of individuals is finite and, in many classical formulations, it is kept fixed in order to focus on changes in strategy composition rather than on demographic growth. This is the case, for instance, in the *Moran process*, originally introduced in population genetics by Moran [50]. In its simplest form, a population of fixed size N contains individuals of two types. At each update step, one individual is selected for reproduction, typically with probability proportional to its fitness, and one individual is selected for death. The offspring of the reproducing individual then replaces the individual selected for death. Since one reproduction event and one death event occur together, the total population size remains constant, while the number of individuals adopting each strategy may change. A schematic illustration of one elementary update of the Moran process is shown in Fig. 1.2. As a consequence, the dynamics is intrinsically *stochastic*. Even if one strategy has a selective advantage, random fluctuations may still lead to its extinction. Conversely, a disadvantageous mutant may sometimes spread and eventually fixate simply because of demographic noise. In the absence of mutation, the process eventually reaches an absorbing state in which one type has fixed in the population and the other has gone extinct. The relevant quantities are therefore fixation probabilities, fixation times and the effect of stochastic drift. Finite-population models are especially important when population size is small or when fluctuations cannot be neglected, as in microbial populations, small animal




Figure 1.2: Schematic illustration of one elementary update in the Moran process. One individual is selected for reproduction (marked R) and one individual is selected for death (marked D). The offspring of the reproducing individual replaces the individual selected for death, so that the total population size remains constant while the composition of the population may change. *Illustration by the author.*

groups, early invasion events by rare mutants, or stochastic agent-based simulations. In infinite populations, by contrast, the state of the system is described by continuous strategy frequencies. Random fluctuations are averaged out at the macroscopic level, and the evolutionary dynamics can be described *deterministically* through differential equations. This limit provides a clean analytical baseline for studying the direction of selection, the existence of equilibria and their stability. In this thesis, we will primarily focus on the infinite-population scenario, using deterministic equations as the analytical baseline for the evolutionary dynamics. However, the two treatments are connected. Indeed, Traulsen, Claussen and Hauert showed that the replicator equation can be obtained as the deterministic large-population limit of a generalised Moran process, with finite-size corrections accounting for the stochastic effects present in finite populations [10].

1.4.2 The Replicator Equation

In the infinite-population scenario, the state of the population is described by continuous strategy frequencies, and the stochastic effects associated with individual update events become negligible at the macroscopic level. The evolutionary dynamics can therefore be represented by deterministic differential equations. The fundamental equation describing frequency-dependent selection in this setting is the *replicator equation*.

Consider a population composed of n strategies, and let $x_i(t)$ denote the frequency of strategy i at time t , satisfying the normalisation condition $\sum_{i=1}^n x_i(t) = 1$. The interaction between strategies is described by a payoff matrix A , an $n \times n$ matrix whose entries a_{ij} , $i, j = 1, \dots, n$ represent the payoff obtained by an individual using strategy i when interacting with an individual using strategy j . The expected payoff, or fitness, of strategy i is

$$f_i = \sum_{j=1}^n x_j a_{ij} = (A\mathbf{x})_i,$$

while the average fitness of the population is

$$\phi = \sum_{i=1}^n x_i f_i = \mathbf{x}^T A \mathbf{x}.$$

The replicator equation is then given by

$$\dot{x}_i = x_i (f_i - \phi), \quad i = 1, \dots, n. \quad (1.4)$$

Equation (1.4) formalises a simple selection principle. In biological applications, fitness measures reproductive success: strategies associated with higher fitness become more represented in the population. In evolutionary game theory, the payoff obtained from interactions is therefore interpreted as a contribution to fitness. Strategies whose payoff, or fitness, is larger than the population average increase in frequency, whereas strategies whose payoff is smaller than the average decrease. In other words, evolutionary change is driven by relative success. What matters is not the absolute value of the payoff, but whether a strategy performs better or worse than the average strategy in the population. The replicator equation provides the standard deterministic description of evolutionary dynamics in well-mixed populations. In this setting, individuals are not constrained by spatial or network structure: each player can interact with any other player, and the probability of encountering strategy j is simply its global frequency x_j . Equivalently, the local environment experienced by each individual coincides with the average composition of the whole population. From a network perspective, this corresponds to a *complete graph*, where every individual is connected to every other individual. The replicator equation is therefore a *mean-field description*: it neglects local correlations, clustering and the specific topology of interactions.

Stability Concepts Before analysing the qualitative behaviour of the replicator equation, it is useful to recall the basic notions of stability for deterministic dynamical systems. Consider a system $\dot{\mathbf{x}} = F(\mathbf{x})$. A point $\hat{\mathbf{x}}$ is a rest point if $F(\hat{\mathbf{x}}) = 0$. It is *Lyapunov stable* if, for every $\varepsilon > 0$, there exists $\delta > 0$ such that

$$\|\mathbf{x}(0) - \hat{\mathbf{x}}\| < \delta \quad \Rightarrow \quad \|\mathbf{x}(t) - \hat{\mathbf{x}}\| < \varepsilon \quad \forall t \geq 0.$$

Thus, trajectories starting sufficiently close to $\hat{\mathbf{x}}$ remain close to it for all future times. The rest point is *asymptotically stable* if it is Lyapunov stable and, in addition, there exists $\delta > 0$ such that

$$\|\mathbf{x}(0) - \hat{\mathbf{x}}\| < \delta \quad \Rightarrow \quad \lim_{t \rightarrow \infty} \mathbf{x}(t) = \hat{\mathbf{x}}.$$

1.4.3 Nash Equilibria and Evolutionary Stability

In the context of evolutionary game dynamics, Nash equilibria are closely related to rest points of the replicator equation. In the following, we use the standard definitions and results presented by Hofbauer and Sigmund [51].

For a symmetric game with payoff matrix A , a population state $\hat{\mathbf{x}}$ is a Nash equilibrium if no alternative population state $\mathbf{x}$ obtains a higher payoff against $\hat{\mathbf{x}}$ than $\hat{\mathbf{x}}$ obtains against itself:

$$\mathbf{x}^T A \hat{\mathbf{x}} \leq \hat{\mathbf{x}}^T A \hat{\mathbf{x}}.$$

This means that no mutant strategy distribution can outperform the resident population $\hat{\mathbf{x}}$. From the definition, it follows that every Nash equilibrium is a rest

point of the replicator equation. Indeed, if $\hat{x}_i > 0$, then all strategies present in the population must have payoff equal to the average payoff,

$$(A\hat{\mathbf{x}})_i = \hat{\mathbf{x}}^T A\hat{\mathbf{x}},$$

and therefore

$$\dot{x}_i = \hat{x}_i \left[(A\hat{\mathbf{x}})_i - \hat{\mathbf{x}}^T A\hat{\mathbf{x}} \right] = 0.$$

The converse is not true in general: a rest point only requires all strategies present in the population to have the same payoff, whereas a Nash equilibrium also requires that absent strategies do not obtain a higher payoff against the resident population. However, dynamical stability restores a connection with Nash equilibrium. In particular, as shown in [51], every Lyapunov stable rest point of the replicator dynamics is a Nash equilibrium; since asymptotic stability implies Lyapunov stability, every asymptotically stable rest point is also a Nash equilibrium.

The notion of evolutionary stability can also be expressed in this context. A state $\hat{\mathbf{x}}$ is an *evolutionarily stable state* if, for every $\mathbf{x} \neq \hat{\mathbf{x}}$ sufficiently close to $\hat{\mathbf{x}}$,

$$\hat{\mathbf{x}}^T A\mathbf{x} > \mathbf{x}^T A\mathbf{x}.$$

Thus, when the population is slightly perturbed, the resident state $\hat{\mathbf{x}}$ obtains a higher payoff against the perturbed population than the perturbation obtains against itself. Evolutionary stability is therefore stronger than the Nash condition and provides a dynamical stability criterion: every evolutionarily stable state is an asymptotically stable rest point of the replicator dynamics as shown in [51].

1.4.4 Replicator Dynamics in Well-Mixed Populations

At this point, we specialise the replicator dynamics to the case of two strategies: cooperation, denoted by C , and defection, denoted by D .

The Dilemma Space As introduced in the previous section, a symmetric two-player game can be written in the normal form

$$G = \begin{pmatrix} R & S \\ T & P \end{pmatrix},$$

where R is the reward for mutual cooperation, S is the sucker's payoff obtained by a cooperator facing a defector, T is the temptation payoff obtained by a defector facing a cooperator, and P is the punishment for mutual defection. In the following, we adopt a normalised parametrisation of this payoff matrix. Since the replicator equation depends on payoff differences, adding the same constant to all payoffs does not change the dynamics, while multiplying all payoffs by a positive constant only rescales time. Therefore, provided $R > P$, the payoff matrix can be normalised without loss of generality by setting

$$R = 1, \quad P = 0.$$

Equivalently, this corresponds to applying a positive affine transformation to the original payoff matrix. Henceforth, for convenience, we adopt the following parametrisation of the payoff matrix, introduced in [52] and later extended in [53],

$$G(D_g, D_r) = \begin{pmatrix} 1 & -D_r \\ 1 + D_g & 0 \end{pmatrix}. \quad (1.5)$$

where $D_g := T - R$ and $D_r := P - S$ are the *gamble-intending* and *risk-averting* dilemma strength parameters, respectively. The parameter D_g measures the incentive to defect against a cooperator, while D_r quantifies the disadvantage experienced by cooperators when facing defectors. Different combinations of the parameters $D_g, D_r \in [-1, 1]$ correspond to different games. In the (D_g, D_r) plane, also referred to as the *dilemma space*, this domain forms a square that is naturally partitioned into four regions, as shown in Fig. 1.3. The signs of D_g and D_r determine the strategic structure of the game. Since $D_g = T - R$ and $D_r = P - S$, the condition $D_g > 0$ means that defecting against a cooperator is more profitable than mutual cooperation, while $D_g < 0$ means that mutual cooperation gives a larger payoff than unilateral defection. Similarly, $D_r > 0$ means that cooperating against a defector is worse than mutual defection, while $D_r < 0$ means that unilateral cooperation is still preferable to mutual defection. The axes correspond to degenerate boundary cases. If $D_g = 0$, then $T = R$, so the temptation payoff is equal to the reward for mutual cooperation. If $D_r = 0$, then $P = S$, so the payoff of mutual defection is equal to the sucker's payoff. At the origin, $D_g = D_r = 0$, both equalities hold: $T = R$ and $P = S$. Therefore, the axes do not correspond to strict game classes, but to transition lines between the four canonical regions. The four open quadrants correspond to the canonical games: Prisoner's Dilemma for $D_g > 0$ and $D_r > 0$, Snowdrift Game for $D_g > 0$ and $D_r < 0$, Stag Hunt Game for $D_g < 0$ and $D_r > 0$, and Harmony Game for $D_g < 0$ and $D_r < 0$.

Analytical Characterisation We now derive the analytical equilibria of the four game classes in the dilemma space. We call ρ the frequency of cooperators in the population. The fitness of cooperators and defectors are given by their expected payoffs against the current population, that is,

$$\begin{aligned} f_C &= \Pi_C = \rho - D_r(1 - \rho), \\ f_D &= \Pi_D = (1 + D_g)\rho. \end{aligned}$$

The average payoff in the population is

$$\langle \Pi \rangle = \rho \Pi_C + (1 - \rho) \Pi_D.$$

In this setting, the replicator equation for cooperators reads

$$\dot{\rho} = \rho(\Pi_C - \langle \Pi \rangle) = \rho(1 - \rho)\Delta\Pi(\rho) = \rho(1 - \rho)[\rho(D_r - D_g) - D_r],$$

where $\Delta\Pi(\rho) = \Pi_C(\rho) - \Pi_D(\rho)$ is the payoff difference between cooperators and defectors. The *stationary points* are obtained from $\dot{\rho} = 0$, which yields $\rho = 0$, $\rho =$




Figure 1.3: Classification of symmetric two-strategy games in the (D_g, D_r) plane. The parameter $D_g = T - R$ measures the temptation to defect against a cooperator: $D_g > 0$ means that unilateral defection is more profitable than mutual cooperation, whereas $D_g < 0$ means that mutual cooperation gives a higher payoff than unilateral defection. The parameter $D_r = P - S$ measures the risk of cooperating against a defector: $D_r > 0$ means that being exploited as a cooperator is worse than mutual defection, whereas $D_r < 0$ means that unilateral cooperation is still preferable to mutual defection. The four open quadrants correspond to the canonical game classes: *Prisoner's Dilemma* (PD), *Snowdrift* (SD), *Stag Hunt* (SH), and *Harmony Game* (HG). The axes $D_g = 0$ and $D_r = 0$ are boundary cases where the corresponding payoff inequalities become equalities. *Illustration by the author.*

1, and the interior fixed point, whenever it exists,

$$\rho^* = \frac{D_r}{D_r - D_g}.$$

The boundary points $\rho = 0$ and $\rho = 1$ are always stationary because of the prefactor $\rho(1 - \rho)$: if the population is entirely composed of one strategy, the frequency of cooperators cannot change under the replicator equation. However, as discussed in Sect. 1.4.3, being a stationary point does not by itself imply being a Nash equilibrium. The connection with Nash equilibrium is recovered through stability. Indeed, as discussed in the previous section, every asymptotically stable stationary point is also Lyapunov stable and therefore corresponds to a Nash equilibrium. The *stability* of the fixed points is determined by the sign of the *Jacobian*

$$J(\rho) = \frac{d\dot{\rho}}{d\rho},$$

Game class	Parameters	Stable equilibrium
Prisoner's Dilemma (PD)	$D_g > 0, D_r > 0$	Full defection: $\rho = 0$
Snowdrift Game (SD)	$D_g > 0, D_r < 0$	Mixed: $\rho^* \in (0, 1)$
Stag Hunt (SH)	$D_g < 0, D_r > 0$	Bistability: $\rho = 0$ or $\rho = 1$
Harmony Game (HG)	$D_g < 0, D_r < 0$	Full cooperation: $\rho = 1$

Table 1.1: Classification of symmetric two-strategy games in the (D_g, D_r) plane and corresponding stable outcomes under the well-mixed replicator dynamics. The signs of $D_g = T - R$ and $D_r = P - S$ determine the game class. In the Prisoner's Dilemma, full defection is stable; in the Snowdrift Game, the stable state is an interior coexistence equilibrium; in the Stag Hunt Game, both full defection and full cooperation are locally stable, leading to bistability; in the Harmony Game, full cooperation is stable.

which is given by

$$J(\rho) = (1 - \rho)\Delta\Pi(\rho) - \rho\Delta\Pi(\rho) + \rho(1 - \rho)\frac{d\Delta\Pi(\rho)}{d\rho}.$$

A fixed point is asymptotically stable when $J(\rho) < 0$ and unstable when $J(\rho) > 0$. We proceed by evaluating the Jacobian at the three fixed points. We obtain that:

- (i) $\rho = 0$ (full defection) is asymptotically stable when $D_r > 0$. Hence, in this region, full defection is also a Nash equilibrium.
- (ii) $\rho = 1$ (full cooperation) is asymptotically stable when $D_g < 0$. Hence, in this region, full cooperation is also a Nash equilibrium.
- (iii) The interior equilibrium $\rho^* \in (0, 1)$ is asymptotically stable when $D_g > 0, D_r < 0$, and unstable when $D_g < 0, D_r > 0$. When asymptotically stable, this interior equilibrium is also a Nash equilibrium.

Recalling the classification of social dilemmas in the (D_g, D_r) plane as per Fig. 1.3, we can now characterise the evolutionary dynamics in each of the four quadrants. The resulting analysis is summarised in Table 1.1.

To fully characterise the equilibrium structure in the dilemma space, we used the analytical conditions derived above to construct a deterministic prediction of the final cooperation level across the whole (D_g, D_r) plane. In particular, we scanned the square domain $D_g, D_r \in [-1, 1]$ using a fine grid of 1001×1001 equally spaced parameter pairs. For each pair (D_g, D_r) , we computed the stationary points of the replicator equation and determined their stability from the Jacobian. This procedure assigns to each point of the dilemma space the stable equilibrium predicted by the well-mixed replicator dynamics. Thus, the analytical prediction is $\rho = 0$ in the Prisoner's Dilemma region, $\rho = 1$ in the Harmony Game region, and $\rho = \rho^*$ in the Snowdrift region. The Stag Hunt region requires an additional step, because the dynamics is bistable. In this case, both boundary equilibria $\rho = 0$ and $\rho = 1$ are stable, while the interior equilibrium ρ^* is unstable and separates their basins of attraction. Since the stochastic simulations discussed below are initialised at

$\rho_0 = 0.5$, we assigned the analytical outcome in the Stag Hunt region by comparing ρ_0 with the unstable threshold ρ^* . If $\rho_0 > \rho^*$, the predicted final state is full cooperation, $\rho = 1$; if $\rho_0 < \rho^*$, the predicted final state is full defection, $\rho = 0$. This gives an analytical heatmap of the expected equilibrium cooperation level across the dilemma space, shown in Fig. 1.4. This deterministic prediction will be used as a reference for comparison with the finite-population stochastic simulations discussed below.




Figure 1.4: Analytical prediction of the asymptotic cooperation level in the (D_g, D_r) dilemma space under the well-mixed replicator dynamics. The square domain $D_g, D_r \in [-1, 1]$ was sampled on a 1001×1001 grid. For each parameter pair, the stationary points of the two-strategy replicator equation were computed and classified through the sign of the Jacobian. The prediction assigns full defection, $\rho = 0$, in the Prisoner's Dilemma region, full cooperation, $\rho = 1$, in the Harmony Game region, and the stable coexistence equilibrium, $\rho = \rho^*$, in the Snowdrift region. In the Stag Hunt region, where the dynamics is bistable, the final state depends on the basin of attraction. Since the simulations used for comparison are initialised at $\rho_0 = 0.5$, the analytical outcome is assigned by comparing ρ_0 with the unstable interior threshold ρ^* : if $\rho_0 > \rho^*$ the predicted state is full cooperation, whereas if $\rho_0 < \rho^*$ the predicted state is full defection. This heatmap therefore provides the deterministic well-mixed reference against which finite-population stochastic simulations are compared. *Original illustration by the author.*

Numerical Simulations At this point, we perform numerical simulations to verify the results predicted by the replicator equation obtained above. Throughout the simulations, strategy updating is implemented through the *Fermi updating rule*,

first introduced in the context of evolutionary games by Szabó and Tóke in [11]. At each elementary update, a focal player X compares its payoff Π_X with the payoff Π_Y of another player Y , and adopts the strategy of Y with probability

$$p = \frac{1}{1 + \exp[-\beta(\Pi_Y - \Pi_X)]}, \quad (1.6)$$

where β is the *intensity of selection*. This rule is closely related to *Glauber dynamics* in statistical mechanics, originally introduced for the stochastic dynamics of Ising spin systems [54]. In kinetic spin systems, transition probabilities depend sigmoidally on energy differences; analogously, in evolutionary games, strategy adoption depends sigmoidally on payoff differences. The parameter β plays the role of an *inverse temperature*. For small β , selection is weak and strategy adoption is strongly stochastic, with $p \simeq 1/2$ even when payoff differences are present. For large β , selection is strong and the rule approaches deterministic imitation of the better performing strategy. Thus, the Fermi rule incorporates microscopic noise, modelling imperfect information, bounded rationality, and fluctuations not explicitly controlled in the system. Figure 1.5 shows the Fermi adoption probability as a function of the payoff difference $\Delta\Pi = \Pi_Y - \Pi_X$ for different values of the selection intensity β . For small β , payoff differences have only a weak effect on strategy adoption, corresponding to weak selection and high stochasticity. For large β , the response becomes increasingly sharp, approaching deterministic imitation of the better-performing strategy. The system evolves through an asynchronous strategy-




Figure 1.5: Fermi update probability as a function of the payoff difference $\Delta\Pi = \Pi_Y - \Pi_X$ for different values of the selection intensity β . Small values of β correspond to weak selection and strong stochasticity, whereas large values of β make the update rule increasingly deterministic. *Illustration by the author.*

update process based on reciprocal interactions. At each elementary update step, a focal player X and a model player Y are selected uniformly at random from the population. Each of them plays the game against $k = 4$ randomly chosen individuals and collects a total payoff, denoted by Π_X and Π_Y , respectively. The focal player X

then adopts the strategy of the model player Y according to the Fermi update rule in Eq. (1.6). We considered a population of size $N = 1000$, while the selection intensity was fixed to $\beta = 0.2$. Time is measured in Monte Carlo sweeps. One Monte Carlo sweep corresponds to N elementary update attempts, so that each individual has, on average, one opportunity to revise its strategy per sweep. We selected four representative parameter pairs (D_g, D_r) , one for each quadrant of the dilemma space, corresponding to the Prisoner's Dilemma, the Snowdrift Game, the Stag Hunt and the Harmony Game. For each game, simulations were initialised with nine equally spaced initial cooperation levels, $\rho_0 = 0.1, 0.2, \dots, 0.9$. For every initial condition, we performed 20 independent realisations and computed the average cooperation level ρ as a function of Monte Carlo sweeps. The results are reported in Fig. 1.6. The stochastic simulations reproduce the predictions of the analytical fixed-point analysis summarised in Table 1.1. In each panel, the horizontal red reference line indicates the corresponding prediction obtained from the replicator equation. The agreement between simulations and analytical predictions is excellent: cooperation vanishes in the Prisoner's Dilemma, the Snowdrift Game converges to a mixed equilibrium, the Stag Hunt displays initial-condition-dependent bistability, and the Harmony Game converges to full cooperation.

To fully characterise the evolutionary dynamics in the dilemma space, we also performed a systematic numerical scan of the whole (D_g, D_r) plane. Specifically, we considered $21 \times 21 = 441$ equally spaced parameter pairs (D_g, D_r) in the square domain $D_g, D_r \in [-1, 1]$. Each pair corresponds to a different two-strategy game in the dilemma space. For each game, we used the same asynchronous Fermi update protocol introduced above. The initial fraction of cooperators was fixed to $\rho_0 = 0.5$, and for each parameter pair we performed 20 independent realisations. The simulations were run until the cooperation level reached a stationary regime. To estimate the final cooperation level, we computed the average value of $\rho(t)$ over the last 20% of the simulation time. This procedure yields an estimate of the asymptotic cooperation level associated with each point of the dilemma space. The resulting values were represented as a heatmap in the (D_g, D_r) plane, shown in Fig. 1.7. This representation allows one to visualise the global organisation of cooperation across the four game classes and to compare the numerical outcomes with the analytical predictions of the replicator equation. Overall, the numerical results agree with the analytical prediction in Fig. 1.4, recovering the same organisation of cooperation across the dilemma space. The small differences between the two heatmaps are due to the fact that the analytical prediction is obtained from the deterministic infinite-population replicator equation on a fine 1001×1001 grid, whereas the numerical results come from finite-population stochastic simulations performed, for computational reasons, on a coarser 21×21 grid and averaged over a finite number of realisations.

This chapter introduced the game-theoretic and evolutionary-game framework used to study the emergence of cooperation. In the models discussed so far, the payoff matrix, the interaction structure and the update rule were kept fixed, while only the strategy composition of the population changed over time.

The next chapter moves from this evolutionary-game setting to a coevolutionary one. In coevolutionary games, strategies evolve together with at least one additional component of the system, such as the interaction network, the payoff matrix,

Evolutionary Dynamics of the Social Dilemmas




Figure 1.6: Well-mixed stochastic simulations for the four canonical two-strategy game classes. Each panel shows the average cooperation level ρ as a function of Monte Carlo sweeps for one representative parameter pair (D_g, D_r) , with the specific values reported in the title of the corresponding panel. The four panels correspond respectively to the Prisoner's Dilemma, Snowdrift Game, Stag Hunt and Harmony Game. Simulations were performed with population size $N = 1000$, selection intensity $\beta = 0.2$, and Fermi strategy updating. For each game, the dynamics was initialised from nine equally spaced initial cooperation levels, $\rho_0 = 0.1, 0.2, \dots, 0.9$, and averaged over 20 independent realisations. The horizontal red reference line indicates the analytical prediction obtained from the well-mixed replicator equation. Different panels are shown over different Monte Carlo sweep windows. The excellent agreement between simulations and analytical predictions confirms the validity of the replicator-equation fixed-point analysis in the well-mixed setting. *Simulation by the author.*

an environmental variable or an individual trait. This transition prepares the introduction of the endogenous-feedback model, in which the game itself changes as a function of the cooperation level. This mechanism constitutes the central coevolutionary framework studied in this thesis.




Figure 1.7: Final cooperation level obtained from well-mixed stochastic simulations across the (D_g, D_r) dilemma space. The heatmap reports the estimated stationary cooperation level for 21×21 equally spaced parameter pairs in the domain $D_g, D_r \in [-1, 1]$. For each parameter pair, the population was initialised at $\rho_0 = 0.5$ and evolved according to the asynchronous Fermi update rule. The reported value is obtained by averaging $\rho(t)$ over 20 independent realisations and then over the last 20% of the simulation time after convergence. The heatmap provides the finite-population stochastic counterpart of the deterministic well-mixed prediction shown in Fig. 1.4. The overall structure agrees with the replicator-equation prediction, while small discrepancies are due to finite-size fluctuations, stochastic updating, the finite number of realisations, and the coarser 21×21 parameter grid used in the simulations. *Simulation by the author.*

Chapter 2

Coevolutionary Game Theory and the Endogenous Feedback Model

In the previous chapter, evolutionary game theory was introduced as a dynamical extension of classical game theory. In this framework, the focus is no longer only on the strategic choices of isolated rational agents, but on the evolution of strategy frequencies in a population. The standard evolutionary setting assumes that individuals interact according to a fixed payoff matrix and that the success of each strategy determines whether it increases or decreases in frequency over time. This framework provides a natural description of frequency-dependent selection, however, it relies on important simplifying assumptions: the payoff matrix, the interaction structure and the update rule are usually treated as externally specified and constant in time. Only the strategies are allowed to evolve. In many biological, ecological and social systems this assumption is restrictive. Individuals may modify their social ties, move in space, alter their influence, change their learning behaviour, or affect the environmental conditions that determine the payoffs of the game. This motivates the transition from evolutionary game theory to *coevolutionary game theory*. Broadly, a coevolutionary game is an evolutionary game in which strategies evolve together with at least one additional dynamical variable. The term therefore does not refer to a single specific model, but to a broad class of models in which strategic evolution is coupled to the evolution of some other property of the system [18].

In this chapter, we first discuss the main forms of coevolutionary dynamics. We then focus on a specific class of models based on *endogenous feedback*, in which the payoff matrix is not fixed externally but changes as a function of the state of the population. The endogenous-feedback framework used in this thesis was introduced in [3]. This framework leads to the emergence of so-called *chimera games*, stable intermediate-cooperation regimes that are not predicted by the corresponding fixed-game dynamics. These regimes are one of the main focuses of this thesis and will be introduced in detail in this chapter.

2.1 Main Classes of Coevolutionary Dynamics

Coevolutionary games can be organised into several broad classes, depending on which component of the system evolves together with the strategies. This coupling is especially relevant for the study of cooperation, because mechanisms that are absent in fixed games may emerge once strategies are allowed to coevolve with other dynamical components of the system. For example, cooperators may avoid exploitation by breaking links with defectors; a resource depleted by defectors may alter the payoff matrix and make cooperation advantageous; or successful individuals may become more influential and therefore transmit their strategies more efficiently. In this section, we discuss three representative classes of coevolutionary dynamics:

1. *network coevolution*, where the interaction structure changes over time;
2. *payoff-matrix evolution*, where the payoffs or the environmental conditions determining the payoffs change over time;
3. *behavioural-rule and trait coevolution*, where individual properties such as update rules, influence, mobility or teaching activity evolve together with strategies.

This classification is not exhaustive, and the boundaries between the three classes are not always sharp. A more detailed discussion of coevolutionary mechanisms in evolutionary games can be found in the review by Perc and Szolnoki [18].

2.1.1 Network Coevolution

One of the most studied forms of coevolutionary dynamics concerns the simultaneous evolution of strategies and interaction networks. Here, we use only the basic idea of a network as a structure of interactions among individuals; a systematic introduction to complex networks and their mathematical properties will be provided in Chapter 3. In standard evolutionary games on graphs, the network is fixed: vertices represent players, edges represent possible interactions, and only strategies change over time. In network coevolution, by contrast, the interaction structure becomes itself a dynamical variable. The general feedback loop can be written schematically as

strategies $\longrightarrow$ network structure $\longrightarrow$ future interactions $\longrightarrow$ strategies.

Thus, the network is no longer a passive substrate for evolutionary dynamics. Instead, it is reshaped by the strategic state of the population and, in turn, modifies the future evolutionary process. This idea is particularly natural in social and biological systems, where individuals are often not forced to maintain the same interaction partners indefinitely. Unfavourable connections may be broken, new partners may be selected, and successful interactions may be reinforced. In the case of cooperation, this can allow cooperators to avoid persistent exploitation by defectors and to form more assortative interaction patterns. However, as emphasised by Perc and Szolnoki in [18], coevolutionary rules must be introduced carefully: if the rule is designed to




Figure 2.1: Representative interaction-updating rules in network coevolution. The figure illustrates three ways in which the interaction network can change together with the strategies of the players. Cooperators are shown in green and defectors in black. In Type A, link creation and deletion depend directly on the strategies of the connected players: different types of links, such as CC , CD and DD links, can have different creation and deletion rates. Links marked in red are removed, while links marked in blue are newly created. In Type B, rewiring is driven by payoff-based partner evaluation: players may cut adverse or unprofitable ties and replace them with more favourable connections. In Type C, network adaptation is coupled to strategy adoption: when a player adopts the strategy of a neighbour, the link to the strategy donor is preserved, while some other links may be removed. These three rules exemplify how the network can become an active dynamical variable rather than a fixed substrate for evolutionary dynamics. *Adapted from [18].*

favour cooperators directly, the resulting promotion of cooperation may be artificial rather than emergent. Following the classification proposed by Perc and Szolnoki, three representative classes of interaction-updating rules can be distinguished, as illustrated in Fig. 2.1.

- **Type A: strategy-dependent link dynamics** In the first class, the probability that a link is created or deleted depends directly on the strategies of the connected players. For instance, links between cooperators may be longer lived than links involving defectors, because mutual cooperation yields benefits to both endpoints, whereas defector-defector interactions are unprofitable for both players. This mechanism is exemplified by active-linking models, where the creation and deletion rates of CC , CD and DD links are strategy-dependent [55].
- **Type B: payoff-dependent rewiring** In the second class, link adaptation depends on the payoff obtained from a given interaction. A player may evaluate an existing connection and replace it if an alternative partner is expected to

yield a higher payoff. In this case, rewiring is not determined solely by the strategies at the endpoints, but by the performance associated with the link. Such rules represent adaptive partner choice and can generate heterogeneous networks in which cooperation is supported by the selective removal of adverse ties.

- **Type C: strategy-update-induced rewiring** In the third class, changes in the interaction network are coupled directly to strategy adoption. For example, when a player adopts the strategy of a neighbour, the adoption event may trigger the deletion of some of its previous links, while the link to the strategy donor is preserved. In this case, network evolution is not an independent process, but is tied to the microscopic events through which strategies spread.

A central role is played by the relative time scale of network adaptation and strategy updating. If links are rewired much faster than strategies change, individuals do not effectively evolve on the initial network, experiencing instead an averaged interaction structure generated by the rapid link dynamics. As a consequence, the microscopic payoff matrix may remain unchanged, while the average payoffs experienced by the strategies are modified. In this sense, fast network coevolution can induce an effective transformation of the game. Likewise, in the limit of very slow rewiring, the model reduces to standard evolutionary dynamics on a fixed network. If, instead, network and strategy dynamics occur on comparable time scales, the outcome depends on their coupled evolution and may interpolate between the behaviour of static networks and that of rapidly adapting systems. Network coevolution therefore provides a mechanism through which cooperation may be promoted without changing the microscopic payoff matrix directly, but by changing the pattern of interactions among individuals.

2.1.2 Payoff-Matrix Evolution

A second major class of coevolutionary games is obtained when the payoff matrix itself is allowed to evolve. In standard evolutionary game theory, the payoff matrix is fixed: strategies evolve, but the incentives defining the game remain unchanged. In payoff-matrix evolution, this assumption is relaxed. The strategic environment becomes dynamic, so that the payoffs associated with cooperation and defection may change over time. In general, one may write

$$A = A(t),$$

or, more specifically,

$$A = A(z(t)),$$

where $z(t)$ is a dynamical variable that modifies the entries of the payoff matrix. This variable may represent an environmental state [19], the abundance of a resource [20], an institutional condition [39], a social norm [56], or any collective quantity capable of changing the incentives faced by individuals. In this sense, the payoff matrix can be interpreted as a dynamical strategic environment. The essential point is that

the game is no longer an externally fixed object: it becomes part of the dynamical system.

A particularly important example is provided by *environmental feedback games*. In these models, the payoff matrix depends on an additional variable representing the environment, which is itself modified by the strategies present in the population. This produces a feedback loop of the form

$$\text{strategies} \longrightarrow \text{environment} \longrightarrow \text{payoff matrix} \longrightarrow \text{strategies}.$$

The first influential formulation of this idea was introduced by Weitz et al. in the context of a tragedy of the commons with game-environment feedback [19]. Their model introduces a coupled strategy-environment dynamical system, in which the frequency of cooperators, denoted by x , and the environmental state, denoted by n , coevolve. The system is therefore two-dimensional and can be written as

$$\begin{aligned}\varepsilon \dot{x} &= x(1-x) [r_1(x, A(n)) - r_2(x, A(n))], \\ \dot{n} &= n(1-n)f(x).\end{aligned}$$

Here x denotes the fraction of individuals adopting strategy 1, interpreted as cooperation, while $1-x$ is the fraction adopting strategy 2, interpreted as defection. The variable $n \in [0, 1]$ describes the environmental state. Depending on the specific application, n may represent the concentration of a vital nutrient, the biomass of a biological species, the availability of a common resource, the level of pollution, or the quality of a social or institutional environment. Equation (2.1.2) is a replicator equation with an environment-dependent payoff matrix. The quantities $r_1(x, A(n))$ and $r_2(x, A(n))$ are the expected payoffs, or fitnesses, of the two strategies when the game is defined by the current environmental state. If $r_1(x, A(n)) > r_2(x, A(n))$, then strategy 1 increases in frequency; if $r_1(x, A(n)) < r_2(x, A(n))$, it decreases. The factor $x(1-x)$ ensures that the boundary states $x = 0$ and $x = 1$ are invariant. Equation (2.1.2) describes the evolution of the environment. The function $f(x)$ represents the feedback of the population on the environment. Its sign determines the direction of environmental change. If $f(x) > 0$, the environmental state increases, corresponding for instance to resource regeneration or environmental enhancement. If $f(x) < 0$, the environmental state decreases, corresponding to depletion or degradation. The factor $n(1-n)$ confines the environmental variable to the interval $[0, 1]$ by making both boundaries invariant. Finally, the parameter ε controls the relative speed of strategy dynamics and environmental dynamics.

The key novelty of this model is that the payoff matrix is environment-dependent. Let the two strategies be cooperation C and defection D , with rows denoting the strategy of the focal player and columns denoting the strategy of the opponent. We use the payoff notation introduced in Sect. 1.2.1, where T , R , P and S denote the temptation, reward, punishment and sucker's payoff, respectively. Environmental feedback is introduced by allowing the payoff matrix to interpolate between two limiting games:

$$A(n) = (1-n) \begin{pmatrix} T & P \\ R & S \end{pmatrix} + n \begin{pmatrix} R & S \\ T & P \end{pmatrix}.$$

Here $n \in [0, 1]$ denotes the environmental state. When $n = 0$, the population plays the first game; when $n = 1$, it plays the second game; intermediate values of n

correspond to intermediate payoff structures. The two limiting matrices represent opposite strategic regimes. Assuming the usual ordering $T > R > P > S$, the matrix obtained for $n = 1$ is the Prisoner's Dilemma payoff matrix discussed in Sect. 1.2.1, where defection is favoured. The matrix obtained for $n = 0$ instead corresponds to the Harmony Game discussed in Sect. 1.3.1, where cooperation is favoured. Thus, the environmental variable controls a continuous transition between a cooperation-favouring regime and a defection-favouring regime. This construction captures the logic of a tragedy of the commons. When the environment is replete, defection is individually advantageous because individuals can exploit the available resource. As defectors become more common, the environment is progressively degraded. In the depleted environment, cooperation becomes advantageous because cooperative behaviour contributes to environmental recovery. As cooperation spreads, the environment improves again, but this restored state creates renewed incentives for defection. The system can therefore cycle between cooperative and defective regimes, and between depleted and replete environmental states. The main result of Weitz et al. is precisely the emergence of an oscillating tragedy of the commons. Instead of converging monotonically to a static defective state, the system may display persistent oscillations generated by the feedback between strategies and environment. In the symmetric version of the model, the dynamics exhibit closed orbits in the interior of the (x, n) state space. In more general asymmetric cases, oscillations may approach a heteroclinic cycle connecting boundary states, as the population repeatedly moves through the sequence: defection dominates, the environment deteriorates, cooperation becomes favoured, the environment recovers, and defection becomes favoured again. This model is important because it shows that changing the payoff matrix through environmental feedback can generate qualitatively new dynamics that are absent in fixed-game replicator dynamics. In a fixed Prisoner's Dilemma, for example, defection simply dominates. In a feedback-evolving game, by contrast, the payoff structure changes as a consequence of the accumulated actions of the population, and this can sustain recurrent transitions between cooperation and defection.

The framework of Weitz et al. was later generalised by Tilman, Plotkin and Akçay, who developed a broader theory of the so-called *eco-evolutionary games* [20]. Their starting point is the observation that in many real systems the environment is not merely changed by strategies, but also has its own intrinsic dynamics. For example, a biological population may grow logistically in the absence of harvesting; pollution may decay even without behavioural change; nutrient levels may be affected both by organisms and by abiotic processes. In their general formulation, the fitness of strategy i depends both on the population state and on the environmental state:

$$\pi_i = \pi_i(\mathbf{x}, n).$$

The strategy frequencies evolve according to a replicator equation,

$$\dot{x}_i = \varepsilon_3 x_i [\pi_i(\mathbf{x}, n) - \phi(\mathbf{x}, n)],$$

where ϕ is the average fitness of the population and ε_3 sets the time scale of strategy dynamics. The environmental state expressed by the variable n evolves according to

$$\dot{n} = \varepsilon_1 f(n) - \varepsilon_2 h(\mathbf{x}, n),$$

where $f(n)$ represents the intrinsic environmental dynamics and $h(\mathbf{x}, n)$ represents the impact of the strategies on the environment. The parameters ε_1 , ε_2 and ε_3 control the relative time scales of intrinsic environmental change, strategy-induced environmental change and strategy updating. Differently from the Weitz model, where the environmental state changes only as a direct consequence of the strategic composition, in the Tilman et al. framework, by contrast, the environment can also renew, decay or undergo tipping point dynamics independently of the current strategic state. For linear two-strategy eco-evolutionary games, Tilman et al. write the environment-dependent payoff matrix as

$$\Pi(n) = (1 - n) \begin{pmatrix} R_0 & S_0 \\ T_0 & P_0 \end{pmatrix} + n \begin{pmatrix} R_1 & S_1 \\ T_1 & P_1 \end{pmatrix}.$$

Here $n = 0$ and $n = 1$ correspond to two environmental extremes, such as poor and rich resource states. The two matrices specify the game played in each of these environmental states. This representation makes it possible to analyse how the system moves through a continuum of games as the environment changes. A central contribution of Tilman et al. is the classification of dynamical outcomes in terms of a small number of interpretable parameters. These parameters measure the incentives to lead or follow behavioural change in rich or poor environments. For example, in a resource-harvesting interpretation, one can distinguish the incentive to be the first individual to switch to a high-impact strategy in a rich environment, the incentive to follow others in doing so, and the corresponding incentives to switch back to a low-impact strategy in a poor environment. These “lead” and “follow” incentives, together with the relative time scale of environmental and strategic change, determine whether the system approaches a monomorphic state, a mixed equilibrium, a bistable regime or persistent cycles. The main conceptual difference between the two approaches is therefore the role of the environment. In Weitz et al., the environment acts as a feedback variable driven directly by the strategies. In Tilman et al., the environment is a dynamical system in its own right, with intrinsic growth, decay or tipping-point behaviour, in addition to being affected by the strategies. Consequently, Tilman et al. show that the Weitz model can be interpreted as a limiting case of a broader class of eco-evolutionary games. For the purposes of this thesis, these models provide the main background for understanding payoff-matrix evolution. They show that once the payoff matrix is allowed to depend on a dynamical variable, evolutionary games can display fixed points, bistability, coexistence, oscillations and tragedy-of-the-commons cycles that are impossible in standard fixed-game dynamics.

2.1.3 Behavioural-Rule and Trait Evolution

Coevolutionary dynamics may also act on individual properties that determine how strategies spread. In many standard evolutionary games, all players are assumed to be identical except for their strategy. They use the same update rule, have the same influence on their neighbours, and respond to payoff differences in the same way. Behavioural-rule and trait coevolution relax this assumption by allowing individual characteristics to evolve together with strategies. In this class of models, the

additional evolving variable is not the network or the payoff matrix, but a property of the agents. Examples include teaching activity [57], reproduction capability [58], mobility [59], age [60], reputation [61], aspiration level [62], risk attitude [63], or the intensity of selection [64]. The general feedback loop can be written as

$$\text{strategies} \longrightarrow \begin{array}{c} \text{individual traits} \\ \text{or update rules} \end{array} \longrightarrow \text{strategy transmission} \longrightarrow \text{strategies}.$$

One of the most studied examples is the coevolution of teaching activity. In standard pairwise comparison rules, a player with a higher payoff may transmit its strategy to a neighbour. Teaching activity introduces heterogeneity in this transmission ability. A player x is assigned a parameter w_x , measuring how effectively it can spread its strategy. The Fermi update rule can then be modified as

$$W(s_x \rightarrow s_y) = w_x \frac{1}{1 + \exp[(\Pi_y - \Pi_x)/K]},$$

where K controls the uncertainty of strategy adoption and w_x is the teaching activity of player x . If w_x is fixed, teaching activity represents quenched heterogeneity. In a coevolutionary version, instead, w_x changes during the evolutionary process. A simple rule, introduced by Szolnoki and Perc [57], assumes that whenever a player successfully transmits its strategy, its teaching activity increases:

$$w_x \rightarrow w_x + \Delta w.$$

This implements a “successful become more influential” mechanism. The resulting dynamics may generate a heterogeneous distribution of influence, even if all players are initially identical. The relevance of this mechanism lies in the fact that heterogeneity emerges endogenously. Cooperation is not promoted by imposing influential cooperators from the beginning; rather, influence is produced by the evolutionary process itself. In this sense, teaching-activity coevolution provides a microscopic mechanism for the self-organisation of influential individuals. Such individuals play a role analogous to hubs in heterogeneous networks, because their strategies can affect many others through an enhanced transmission capability.

A second example is the coevolution of the strategy update rule itself. In the Fermi rule introduced in Eq. (1.6), the parameter controlling selection intensity is usually fixed. If the adoption probability is written as

$$p = \frac{1}{1 + \exp[-\beta(\Pi_Y - \Pi_X)]},$$

then β determines how strongly payoff differences bias imitation. In behavioural-rule coevolution, one may allow this parameter to vary between individuals and evolve over time:

$$\beta = \beta_i(t).$$

In this case, selection acts not only on strategies, but also on the way in which strategies are updated. The system can therefore select an effective level of stochasticity, rationality or responsiveness to payoff differences [64].

Mobility provides another important example. In spatial games, individuals are often fixed on a lattice or network. If mobility is introduced, agents may move

in response to payoffs, local neighbourhood composition, or other cues. Movement modifies who interacts with whom, thereby changing the future evolutionary dynamics. Random mobility may either support or inhibit cooperation, whereas success-driven migration can favour cooperation by allowing individuals to move toward more favourable local environments [59].

Age can also be treated as a coevolving trait. In models where age affects teaching activity, older individuals may have a higher probability of transmitting their strategy. If adopting a new strategy resets the age of a player, then strategy transmission and demographic renewal become coupled. This can generate asymmetric effects on the propagation of cooperation and defection, because newborn players have low influence and cannot immediately transmit their newly acquired strategy efficiently [60].

Other traits, such as reputation [61], aspiration level [62], and risk attitude [63], can be incorporated in a similar way. Reputation may affect how likely a player is to be imitated or chosen as an interaction partner. Aspiration levels determine whether players are satisfied with their current payoff and therefore whether they are likely to revise their strategy. Risk attitudes modify how uncertain payoffs are evaluated. In each case, the trait influences the evolutionary dynamics and may itself be modified by the outcome of the dynamics.

The common feature of these models is that individual heterogeneity is not merely imposed as a static background condition. Instead, it is generated and transformed by the same evolutionary process that changes the strategies. Behavioural-rule and trait coevolution therefore provide a general framework for studying how differences in influence, mobility, learning behaviour or responsiveness can emerge endogenously and feed back into the evolution of cooperation.

2.2 The Endogenous Feedback Model

In this section, we introduce the Endogenous Feedback Model (EFM), which constitutes the central coevolutionary framework considered in this thesis and will be used in the following chapters. We first formulate the model and highlight its differences with respect to the coevolutionary mechanisms discussed in the previous section. We then analyse the model using the replicator equation, which provides analytical predictions for the possible stationary states and their stability. Finally, we compare these predictions with Monte Carlo simulations and show how endogenous feedback can generate the so-called *chimera games*, which represent one of the main focuses of this thesis.

The model, together with the notion of chimera games, was introduced in [3]. The analytical framework and the main qualitative results discussed below are based on that work, while the Monte Carlo simulations reported in this thesis were reproduced and, in some cases, extended by the author.

2.2.1 Model Formulation

The models discussed in the previous section show that the payoff matrix of an evolutionary game need not be fixed. In eco-evolutionary models, the game is usually written as a function of an explicit environmental variable, such as a resource level or ecological state. The payoff matrix may take the form

$$G(n(t)) = [1 - n(t)]G_1 + n(t)G_2,$$

where $n(t) \in [0, 1]$ represents the state of the environment. The environmental variable has its own dynamics and is affected by the population composition. The resulting system is therefore typically two-dimensional, involving both the strategy variable and the environmental variable.

The *Endogenous Feedback Model* (EFM) first introduced in [3] and considered in this thesis follows a different principle. The game is not mediated by an explicit environmental variable. Instead, the payoff matrix depends directly on the current, or possibly past, state of the population. In the case of social dilemmas, the relevant collective variable is the cooperation level $\rho(t)$, defined as the fraction of individuals adopting the cooperative strategy at time t . The game played at time t is therefore shaped by the same population state that it subsequently acts upon. This gives rise to a feedback loop of the form

$$\rho(t) \longrightarrow G(t) \longrightarrow \rho(t),$$

rather than the eco-evolutionary loop

$$\rho(t) \longrightarrow n(t) \longrightarrow G(n) \longrightarrow \rho(t).$$

This distinction is important. In eco-evolutionary models, the environment may represent a concrete resource, such as biomass, nutrient concentration, pollution level, or the state of a common pool. In the endogenous feedback model, by contrast, the “environment” is not represented by an independent resource variable. The relevant quantity is the incentive structure itself, which is generated by the collective behaviour of the population. This makes the model particularly suitable for systems where strategic incentives emerge endogenously from social, economic, or institutional interactions. Examples include reputation systems in online platforms, where cooperative behaviour modifies future visibility and trust [65], or financial markets, where collective expectations reshape the incentives for liquidity provision, speculation, and opportunistic exit [66]. In such cases, the game does not evolve because an external resource is depleted or replenished: it evolves because the population collectively modifies the strategic conditions under which future decisions are made. Consider two games G_1 and G_2 , written in the dilemma-strength parametrisation introduced in Sect. 1.4.4:

$$G_i = G(D_g^i, D_r^i) = \begin{pmatrix} 1 & -D_r^i \\ 1 + D_g^i & 0 \end{pmatrix}, \quad i = 1, 2.$$

Here G_1 and G_2 represent two baseline strategic environments. They may belong to the same dilemma class or to different regions of the (D_g, D_r) plane (see Fig. 1.3). In

the endogenous feedback model, the game played at time t is a convex combination of these two reference games:

$$G(t) = [1 - f(\rho(t - \tau))] G_1 + f(\rho(t - \tau)) G_2. \quad (2.1)$$

Here, $f : [0, 1] \rightarrow [0, 1]$ is a continuous and increasing function that controls how the cooperation level modifies the payoff matrix. Different choices of the feedback function f correspond to different ways in which the population state affects the payoff matrix. The simplest choice is the linear feedback

$$f(\rho) = \rho.$$

More generally, one may consider nonlinear feedback functions, for example

$$f(\rho) = \rho^\lambda, \quad \lambda > 0,$$

which allow one to model different response profiles of the incentive structure to changes in the cooperation level. Values of λ modify how strongly the game reacts to low or high levels of cooperation. The parameter $\tau \geq 0$ captures the *delay* in the system dynamics. It represents the time required for the incentive structure to react to the current state of the population. When $\tau = 0$, the feedback is *instantaneous*. When $\tau > 0$, the game played at time t depends on the cooperation level observed at the previous time $t - \tau$.

However, in this thesis, we restrict attention to the linear and instantaneous case, namely

$$f(\rho) = \rho, \quad \tau = 0.$$

With this choice, Eq. (2.1) reduces to

$$G(t) = [1 - \rho(t)] G_1 + \rho(t) G_2. \quad (2.2)$$

This choice isolates the fundamental mechanism of endogenous feedback while keeping the dynamics analytically transparent. It also provides the simplest setting in which chimera games can already emerge. A detailed analysis of nonlinear feedbacks and delayed feedbacks can be found in [3].

The choice of the two reference games G_1 and G_2 is not unique. In principle, they can be chosen from any two regions of the dilemma space, depending on the feedback mechanism one wants to model. A particularly relevant case is obtained by letting the feedback interpolate between a cooperation-favouring Harmony Game and a defection-favouring Prisoner's Dilemma. For instance, if G_1 is chosen as a Harmony Game and G_2 as a Prisoner's Dilemma, then increasing cooperation moves the effective payoff matrix from G_1 toward G_2 . In other words, a more cooperative population progressively generates stronger incentives for defection: unilateral defection becomes more rewarding, while cooperation becomes more vulnerable to exploitation. This creates a negative feedback between collective cooperation and the incentives supporting it. Conversely, if the direction of the interpolation is reversed, increasing cooperation can shift the game from a Prisoner's Dilemma toward a Harmony Game, producing a feedback that reinforces cooperative behaviour. Eq. (2.1) captures the defining feature of the model: the population does not evolve

inside a fixed game. Instead, the game itself changes as the population changes. The strategic composition determines the payoff matrix, while the payoff matrix determines the future evolution of the strategic composition as illustrated in Fig. 2.2.




Figure 2.2: Schematic representation of the endogenous feedback model. The payoff matrix is not fixed externally, but changes as a function of the cooperation level ρ . The figure illustrates two possible feedback directions between the population state and the game being played. In the left panel, the reference games are chosen so that G_1 is a cooperation-favouring Harmony Game and G_2 is a defection-favouring Prisoner's Dilemma. As cooperation increases, the effective game moves from G_1 toward G_2 : unilateral defection becomes more rewarding and cooperation becomes more vulnerable to exploitation. In the right panel, the direction is reversed: increasing cooperation shifts the effective game from a Prisoner's Dilemma toward a Harmony Game, thereby favouring mutual cooperation. In both cases, the social dilemma coevolves with the population composition through the feedback loop $\rho(t) \rightarrow G(t) \rightarrow \rho(t)$. Adapted from [3].

2.2.2 Well-Mixed Dynamics

We now analyse the endogenous-feedback model in the well-mixed setting. This provides the simplest baseline in which the feedback mechanism can be studied before introducing any spatial or network structure. We first describe the finite-population Monte Carlo implementation, emphasising how the payoff matrix is updated during the dynamics. We then derive the corresponding replicator-equation predictions, which provide the analytical reference for the stationary cooperation levels and their stability.

Numerical Setup We investigate the endogenous-feedback model through Monte Carlo simulations. The payoff matrix is updated during the dynamics. For a simulation initialised at cooperation level ρ_0 , the initial game is

$$G(\rho_0) = [1 - \rho_0] G_1 + \rho_0 G_2.$$

At later times, the current cooperation level determines the current game:

$$G(t) = [1 - \rho(t)] G_1 + \rho(t) G_2.$$

Thus, strategy updates modify $\rho(t)$, which in turn modifies the payoff matrix used in subsequent interactions. Simulations are performed in a well-mixed population of size $N = 1000$. At each elementary update step, a focal player X and a model player Y are selected uniformly at random. Each player interacts with $k = 4$ randomly chosen individuals and accumulates a total payoff computed using the current payoff matrix $G(t)$. The focal player then adopts the strategy of the model player according to the Fermi rule considered in Eq. (1.6). The selection intensity is fixed to $\beta = 0.2$. Time is measured in Monte Carlo sweeps. One Monte Carlo sweep corresponds to N elementary update attempts, so that each individual has, on average, one opportunity to revise its strategy per sweep. For each pair of reference games G_1 and G_2 , simulations are initialised from three cooperation levels, $\rho_0 = 0.1, 0.5, 0.9$. For every initial condition, we perform 20 independent realisations and compute the average cooperation level $\rho(t)$ as a function of Monte Carlo sweeps.

Replicator Equation Analysis In the well-mixed scenario, the endogenous-feedback model can be analysed through the replicator equation according to Eq. (1.4). Denoting by $\rho(t)$ the fraction of cooperators, the dynamics reads, explicitly retaining the time dependence,

$$\frac{d\rho(t)}{dt} = \rho(t)(1 - \rho(t))\Delta\Pi(\rho(t)), \quad (2.3)$$

where

$$\Delta\Pi(\rho) = \Pi_C(\rho) - \Pi_D(\rho)$$

is the payoff difference between cooperators and defectors. For a fixed game in the parametrisation introduced in Eq. (1.5), the payoff difference is

$$\Delta\Pi(\rho) = \rho(D_r - D_g) - D_r. \quad (2.4)$$

In the endogenous-feedback model, however, the game itself depends on ρ . Using the linear instantaneous feedback introduced in Eq. (2.2), the effective game is

$$G(\rho) = (1 - \rho)G_1 + \rho G_2.$$

Equivalently, the dilemma-strength parameters evolve as

$$D_r(\rho) = (1 - \rho)D_r^1 + \rho D_r^2, \quad D_g(\rho) = (1 - \rho)D_g^1 + \rho D_g^2.$$

Substituting these expressions into Eq. (2.4) gives

$$\Delta\Pi(\rho) = \rho^2 \left[(D_r^2 - D_r^1) + (D_g^1 - D_g^2) \right] + \rho (2D_r^1 - D_r^2 - D_g^1) - D_r^1. \quad (2.5)$$

The stationary points of Eq. (2.3) are

$$\rho = 0, \quad \rho = 1,$$

and any internal solution $\rho^* \in (0, 1)$ satisfying

$$\Delta\Pi(\rho^*) = 0. \quad (2.6)$$

The local stability of a stationary point is determined by the sign of the Jacobian

$$J(\rho) = \frac{d\dot{\rho}}{d\rho}.$$

For an internal stationary point, using $\Delta\Pi(\rho^*) = 0$, this reduces to

$$J(\rho^*) = \rho^*(1 - \rho^*) \left. \frac{d\Delta\Pi}{d\rho} \right|_{\rho=\rho^*}.$$

Hence, an internal equilibrium is locally asymptotically stable when $J(\rho^*) < 0$ and unstable when $J(\rho^*) > 0$. At equilibrium, the system settles on the effective game

$$G_{\text{eq}} = (1 - \rho^*)G_1 + \rho^*G_2,$$

with corresponding parameters

$$D_r^{\text{eq}} = D_r(\rho^*), \quad D_g^{\text{eq}} = D_g(\rho^*).$$

Observe that for an internal equilibrium, Eq. (2.6) is equivalent to the usual mixed-equilibrium condition for the effective game,

$$\rho^* = \frac{D_r^{\text{eq}}}{D_r^{\text{eq}} - D_g^{\text{eq}}}.$$

2.2.3 Emergence of Chimera Games

In this subsection, we use the replicator-equation analysis developed above to study the emergence of chimera games in two representative cases. We first consider the case in which the final reference game G_2 is a Prisoner's Dilemma, and then the complementary case in which G_2 is a Harmony Game. For each case, we compare the analytical predictions of the replicator equation with Monte Carlo simulations.

Prisoner's Dilemma as G_2 . We first consider the case in which the final reference game is a Prisoner's Dilemma, with parameters $(D_g^2, D_r^2) = (1, 1)$. Substituting this choice into Eq. (2.5) gives

$$\Delta\Pi(\rho) = \rho^2 (D_g^1 - D_r^1) + \rho (2D_r^1 - D_g^1 - 1) - D_r^1.$$

In this case, an internal equilibrium exists for $D_r^1 < 0$, namely when G_1 lies in the Harmony or Snowdrift regions. One can realise that this equilibrium is asymptotically stable. Therefore, the system can converge to a nontrivial coexistence level $0 < \rho^* < 1$.

We now compare these replicator predictions with stochastic simulations. We fix G_2 to be a Prisoner's Dilemma and consider four representative choices of the initial game G_1 :

$$(D_g^1, D_r^1) \in \{(-0.92, 0.25), (-0.8, -0.3), (-0.3, -0.8), (0.7, -0.7)\}.$$

These cases span different regions of the dilemma space and allow us to compare conventional outcomes with feedback-induced anomalous regimes. A key difference with respect to static evolutionary games is that, in the endogenous-feedback model, the game being played is not fixed. At each time t , the effective game is given by Eq. (2.2). As a consequence, the evolution can be represented not only through the time series of the cooperation level $\rho(t)$, but also as a trajectory in the (D_g, D_r) dilemma space. Since the interpolation is linear, the trajectory is constrained to the straight segment joining G_1 and G_2 (see Fig. 2.3). For each choice of G_1 , simulations were initialised from three different cooperation levels, $\rho_0 = 0.1, 0.5, 0.9$. For each initial condition, the results were averaged over 20 independent realisations. Fig. 2.3 shows the corresponding trajectories in the dilemma space. The fixed game G_2 is indicated by a filled purple square, while the four choices of G_1 are indicated by open squares. For each initial condition, the starting game $G(\rho_0) = [1 - \rho_0]G_1 + \rho_0 G_2$ is marked by a coloured dot. The black lines represent the trajectories followed by the effective game in the (D_g, D_r) dilemma space as the cooperation level $\rho(t)$ changes during the endogenous-feedback dynamics. The red crosses indicate the equilibrium games reached in the stochastic simulations. For each initial condition, the final cooperation level is estimated by averaging $\rho(t)$ over the last 20% of the simulation time after convergence; this value determines the final position of the effective game in the (D_g, D_r) dilemma space. Fig. 2.4 shows the average cooperation level as a function of Monte Carlo sweeps for the same four choices of G_1 . Each panel corresponds to one initial game, and the curves refer to the three initial conditions $\rho_0 = 0.1, 0.5, 0.9$. The horizontal red line indicates the stable equilibrium cooperation level predicted by the replicator equation (2.3). Overall, the agreement between simulations and replicator predictions is excellent, confirming that the mean-field analysis captures the main features of the feedback-driven dynamics in the well-mixed setting. The most interesting case is

$$(D_g^1, D_r^1) = (-0.8, -0.3),$$

for which G_1 lies in the Harmony Game region. In this case, all three initial conditions converge to the same intermediate cooperation level $0 < \rho^* < 1$, in agreement with the stable internal equilibrium predicted by the replicator equation. The corresponding equilibrium game G_{eq} , however, lies in the *Stag Hunt* region of the dilemma space. Under standard evolutionary dynamics with a fixed payoff matrix, a Stag Hunt game admits bistability: the only stable outcomes are full cooperation and full defection, while the interior equilibrium is unstable. Here, by contrast, the endogenous feedback stabilizes a mixed equilibrium. This is precisely the defining feature of a *chimera game*: the population converges to a stable intermediate cooperation level satisfying the Nash condition of the equilibrium game, even though the same equilibrium game would not support such a stable mixed state under standard static evolutionary dynamics. The mechanism can be understood from the position of G_1 and G_2 in the dilemma space. For the pair $(D_g^1, D_r^1) = (-0.8, -0.3)$, the initial reference game G_1 lies in the Harmony Game region, whereas G_2 is a Prisoner's Dilemma. Therefore, when cooperation increases, the effective game is shifted toward G_2 , where defection is favoured; this tends to reduce cooperation. Conversely, when cooperation decreases, the effective game is shifted back toward G_1 , where cooperation is favoured; this tends to restore cooperation. The stable




Figure 2.3: Trajectories in the (D_g, D_r) dilemma space for the endogenous-feedback model with G_2 fixed to a Prisoner's Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$, and four representative choices of the initial reference game G_1 : $(D_g^1, D_r^1) = (-0.92, 0.25), (-0.8, -0.3), (-0.3, -0.8), (0.7, -0.7)$. The fixed reference game G_2 is indicated by a filled purple square, while the four choices of G_1 are indicated by open squares. For each G_1 , simulations are initialised from three cooperation levels, $\rho_0 = 0.1, 0.5, 0.9$, shown by coloured dots. Since the payoff matrix evolves according to the linear endogenous-feedback rule $G(t) = [1 - \rho(t)]G_1 + \rho(t)G_2$, the effective game follows a trajectory in the dilemma space as the cooperation level $\rho(t)$ changes. The black lines represent the trajectories followed by the effective game during the stochastic simulations. The red crosses indicate the equilibrium games reached in the simulations. For each initial condition, the final cooperation level is estimated by averaging $\rho(t)$ over the last 20% of the simulation time after convergence; this value determines the final position of the effective game in the dilemma space. When different initial conditions converge to the same stationary state, the corresponding red crosses overlap. The case $(D_g^1, D_r^1) = (-0.8, -0.3)$ displays the chimera regime: the trajectory starts from a Harmony Game, moves toward the Prisoner's Dilemma, and reaches an equilibrium game in the Stag Hunt region while stabilising an intermediate cooperation level. *Simulation by the author.*

intermediate state is therefore produced by this negative feedback mechanism: the system is pushed away from both full cooperation and full defection and settles at

Cooperation Behaviour over Time




Figure 2.4: Time evolution of the average cooperation level $\rho(t)$ for the same endogenous-feedback simulations shown in Fig. 2.3. The second reference game is fixed to a Prisoner's Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$, while the four panels correspond to the four representative choices of G_1 : $(D_g^1, D_r^1) = (-0.92, 0.25), (-0.8, -0.3), (-0.3, -0.8), (0.7, -0.7)$. In each panel, the curves show the average cooperation level as a function of Monte Carlo sweeps for the three initial conditions $\rho_0 = 0.1, 0.5, 0.9$, averaged over 20 independent realisations. The horizontal red line marks the stable equilibrium cooperation level predicted by the replicator equation, Eq. (2.3). Different panels are displayed over different Monte Carlo sweep windows. The agreement between the stochastic simulations and the analytical prediction is excellent. In particular, the case $(D_g^1, D_r^1) = (-0.8, -0.3)$ exhibits a chimera regime: endogenous feedback stabilises an intermediate cooperation level even though the corresponding equilibrium game lies in the Stag Hunt region, where standard fixed-game dynamics would predict bistability rather than stable coexistence. *Simulation by the author.*

a nontrivial cooperation level. To make the mechanism more explicit, Fig. 2.5 compares the standard fixed-game picture with the endogenous-feedback dynamics. The figure highlights why the case $(D_g^1, D_r^1) = (-0.8, -0.3)$ should be interpreted as a chimera game.

To test the robustness of these observations more systematically, we extended the analysis to a larger set of reference games G_1 , while keeping $G_2 = G(D_g^2, D_r^2) = G(1, 1)$, corresponding to a Prisoner's Dilemma. The selected games G_1 and the corresponding trajectories in the dilemma space are shown in Fig. 2.6. The games G_1 were chosen along two lines, namely $D_g^1 = -0.9$ and $D_r^1 = -0.9$. This choice provides a clear geometric scan of the dilemma space while keeping the selected games away from the boundary axes $D_g = 0$ and $D_r = 0$, where the classification of game

classes becomes degenerate. For each pair (G_1, G_2) , we used the same simulation protocol described above, with initial cooperation levels $\rho_0 = 0.1, 0.5, 0.9$. The trajectories in the dilemma space were constructed as in Fig. 2.3. The equilibrium game reached by the system was estimated by averaging the cooperation level over the final 20% of the simulation time, after convergence. The full cooperation time series for all selected pairs and initial conditions, together with the corresponding replicator predictions, are reported in Appendix A. The systematic scan confirms the excellent agreement between stochastic simulations and the analytical predictions of the replicator equation. When G_1 lies in the Stag Hunt region, namely for $D_g^1 < 0$ and $D_r^1 > 0$, the dynamics converges to the trivial absorbing state $\rho = 0$. By contrast, when G_1 lies in the Harmony Game region along the branch $D_g^1 = -0.9$, the dynamics typically converges to a stable intermediate cooperation level. In these cases, the equilibrium game reached by the system lies in the Stag Hunt region, while the observed cooperation level is a stable interior equilibrium, therefore producing chimera game regimes. This systematic scan therefore provides a well-mixed benchmark for the analysis developed in the following chapters. In particular, the same set of reference games G_1 will be used when studying endogenous feedback in structured populations, allowing us to compare directly how population structure modifies the feedback-induced regimes observed in the well-mixed case.

Harmony Game as G_2 . We now consider the complementary case in which the final reference game is a Harmony Game,

$$G_2 = G(D_g^2, D_r^2) = G(-1, -1).$$

Substituting this choice into Eq. (2.5) gives

$$\Delta\Pi(\rho) = \rho^2 (D_g^1 - D_r^1) + \rho (2D_r^1 - D_g^1 + 1) - D_r^1.$$

In this case, an internal equilibrium exists for $D_r^1 > 0$, namely when G_1 lies in the Prisoner's Dilemma or Stag Hunt regions. However, this internal equilibrium is unstable and separates the basins of attraction of full defection and full cooperation. Therefore, the final outcome depends on the initial cooperation level: if $\rho_0 > \rho^*$, the population converges to full cooperation, whereas if $\rho_0 < \rho^*$, it converges to full defection. In this scenario, increasing cooperation shifts the effective game toward stronger incentives in favour of cooperation. The feedback is therefore self-reinforcing: as cooperation grows, the strategic environment becomes progressively more favourable to further cooperation. The corresponding trajectories in the dilemma space are shown in Fig. 2.7. The corresponding time evolution of the cooperation level is shown in Fig. 2.8. The horizontal red line in each panel reports the equilibrium cooperation level predicted by the replicator equation, allowing a direct comparison between the stochastic simulations and the analytical dynamics. The results confirm the replicator-equation prediction. When G_1 lies in a region with $D_r^1 > 0$, namely in the Prisoner's Dilemma or Stag Hunt region, the internal equilibrium ρ^* is unstable and acts as a threshold: initial conditions above it converge to full cooperation, while initial conditions below it converge to full defection. When $D_r^1 < 0$, namely when G_1 lies in the Snowdrift region, no physical internal

equilibrium exists and the population converges to full cooperation. Therefore, when G_2 is a Harmony Game, endogenous feedback can reshape the basins of attraction, but it does not generate stable interior coexistence. Unlike the case in which G_2 is a Prisoner's Dilemma, no chimera regime emerges here.

The remaining cases in which the game G_2 is chosen as a Stag Hunt or as a Snowdrift Game are reported in Appendix A.2 and Appendix A.3, respectively.

This chapter introduced the coevolutionary-game framework. We first discussed the main coevolutionary mechanisms studied in the literature, including the evolution of interaction networks, environmental feedbacks, behavioural rules and individual traits. We then focused on the endogenous-feedback model, which constitutes the central coevolutionary mechanism studied in this thesis. In this model, the payoff matrix is no longer fixed, but changes with the state of the population through the cooperation level $\rho(t)$. This creates a feedback loop between the strategies adopted by the agents and the strategic incentives that shape their future evolution. In the well-mixed setting, this mechanism was shown to generate chimera-game regimes, namely stable regimes that would be unstable or impossible in the corresponding fixed-game dynamics.

The analysis carried out so far, however, is based on a highly idealised well-mixed description, in which each pair of individuals interacts with the same probability. The next step is therefore to move towards more realistic settings with an explicit network structure, where interactions are local and topologically constrained. The next chapter introduces structured populations, described by spatial lattices and complex networks, and develops the analytical and numerical tools needed to study evolutionary dynamics on graphs.




Figure 2.5: Schematic explanation of the emergence of chimera games in the endogenous-feedback model. Top panels: in the absence of feedback, the game is fixed and corresponds to a single point in the (D_g, D_r) dilemma space. In the specific example considered here, the equilibrium game lies in the Stag Hunt region. Under standard evolutionary dynamics with a fixed payoff matrix, a Stag Hunt game displays bistability: full cooperation and full defection are stable, whereas the interior equilibrium is unstable. Bottom panels: in the endogenous-feedback model, the game is not fixed but evolves according to Eq. (2.2). Since the feedback is linear, the effective game moves along the segment connecting G_1 and G_2 , with its instantaneous position determined by the current cooperation level $\rho(t)$. At stationarity, the population reaches ρ^* and the game settles at $G^* = (1 - \rho^*)G_1 + \rho^*G_2$. The figure shows that distinct dilemmas drive the dynamics toward the equilibrium game G^* in the Stag Hunt region. Although ρ^* satisfies the Nash condition for G^* , endogenous feedback stabilises this intermediate cooperation level, whereas the same Stag Hunt game under fixed-game dynamics would display bistability rather than stable coexistence. We refer to these hybrid regimes as *chimera games*: the final equilibrium game belongs to one class of social dilemmas, but the stability of the observed outcome is inherited from the full feedback-driven dynamics rather than from the static dynamics of that game alone. Endogenous feedback therefore makes possible stable strategic outcomes that would be forbidden in conventional evolutionary games with a fixed payoff matrix. *Adapted from [3].*




Figure 2.6: Systematic scan of the endogenous-feedback model with G_2 fixed to a Prisoner's Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$. The figure shows the trajectories in the (D_g, D_r) dilemma space obtained for the selected reference games G_1 . The games G_1 are chosen along the lines $D_g^1 = -0.9$ and $D_r^1 = -0.9$, avoiding the boundary axes $D_g = 0$ and $D_r = 0$ where the classification of game classes is degenerate. The fixed game G_2 is indicated by a filled purple square, while the selected games G_1 are indicated by open squares. For each pair (G_1, G_2) , simulations are initialised from $\rho_0 = 0.1, 0.5, 0.9$, and the corresponding initial games $G(\rho_0) = [1 - \rho_0]G_1 + \rho_0 G_2$ lie on the segment connecting G_1 to G_2 . The black trajectories show the evolution of the effective game in the dilemma space, while red crosses mark the equilibrium games reached in the simulations. The equilibrium position is computed from the final cooperation level, estimated by averaging $\rho(t)$ over the last 20% of the simulation time after convergence. The scan shows that Stag Hunt initial games G_1 lead to full defection, whereas Harmony-Game choices of G_1 along the $D_g^1 = -0.9$ branch generate stable intermediate cooperation levels, corresponding to chimera-game regimes. *Original simulation by the author.*




Figure 2.7: Trajectories in the (D_g, D_r) dilemma space for the endogenous-feedback model with G_2 fixed to a Harmony Game, $G_2 = G(D_g^2, D_r^2) = G(-1, -1)$. The fixed reference game G_2 is indicated by a filled purple square, while the different choices of G_1 are indicated by open squares. For each choice of G_1 , simulations are initialised from three cooperation levels, $\rho_0 = 0.1, 0.5, 0.9$, shown by coloured dots. Since the payoff matrix evolves according to the linear feedback rule $G(t) = [1 - \rho(t)]G_1 + \rho(t)G_2$, the effective game follows a trajectory in the dilemma space as the cooperation level $\rho(t)$ changes. The black lines represent the trajectories followed by the effective game during the simulations. The red crosses indicate the equilibrium games reached in the stochastic simulations. For each initial condition, the final cooperation level is estimated by averaging $\rho(t)$ over the last 20% of the simulation time after convergence; this value determines the final position of the effective game in the dilemma space. In this case, the trajectories do not lead to stable interior coexistence: the final outcomes are absorbing states, corresponding either to full cooperation or to full defection. *Simulation by the author.*




Figure 2.8: Time evolution of the average cooperation level $\rho(t)$ for the same endogenous-feedback simulations shown in Fig. 2.7. The second reference game is fixed to a Harmony Game, $G_2 = G(D_g^2, D_r^2) = G(-1, -1)$, while the four panels correspond to four representative choices of G_1 . In each panel, the curves show the average cooperation level as a function of Monte Carlo sweeps for the three initial conditions $\rho_0 = 0.1, 0.5, 0.9$, averaged over 20 independent realisations. When an internal equilibrium exists, the horizontal red line marks the corresponding value of ρ^* predicted by the replicator equation. In this case, however, the internal equilibrium is unstable and acts as a threshold separating the basins of attraction of full defection and full cooperation. Different panels are displayed over different Monte Carlo sweep windows. The agreement between stochastic simulations and replicator predictions is excellent. Unlike the Prisoner’s-Dilemma-as- G_2 case, no stable interior coexistence is observed: endogenous feedback reshapes the basins of attraction, but the asymptotic outcome remains absorbing. *Simulation by the author.*

Chapter 3

Evolutionary Games on Structured Populations

The previous chapter introduced coevolutionary game theory and the Endogenous Feedback Model, in which the payoff matrix is no longer fixed but changes as a function of the cooperation level. The analysis was first developed in the well-mixed setting, where each individual effectively interacts with a random sample of the whole population. This provides a natural baseline and allows for a straightforward analytical treatment, but it neglects one of the most important features of many biological, social and economic systems: interactions are rarely global and homogeneous. Individuals usually interact with a limited set of neighbours, determined by spatial proximity, social ties or the topology of an underlying network. This motivates the transition from well-mixed populations to structured populations. In a well-mixed system, cooperators are uniformly exposed to defectors, and local organisation plays no role. In a structured population, by contrast, the identity of interaction partners matters. Cooperators may form clusters, interact repeatedly with other cooperators, and thereby reduce their exposure to exploitation. Population structure can therefore modify the effective selective pressures acting on cooperation and defection. This idea is closely related to *network reciprocity*, one of the main mechanisms for the evolution of cooperation identified by Nowak [2]. Network reciprocity arises when the population is not well mixed, so that individuals interact locally rather than with the population as a whole. In this case, cooperative behaviour can be supported by spatial or network correlations: clusters of cooperators can help one another and resist invasion by defectors more effectively than isolated cooperators in a well-mixed population. The study of evolutionary games on structured populations therefore shifts the focus from only asking which game is being played to also asking *who interacts with whom*. Indeed, the same payoff matrix can generate different macroscopic outcomes depending on the interaction structure.

In this chapter, we present complex networks as mathematical models of structured populations, introducing the basic definitions and the main structural properties that will be relevant in the rest of the thesis, such as degree, clustering, path length and degree heterogeneity. We then discuss the main analytical approaches used to study evolutionary dynamics on structured populations. In particular, we introduce

the mean-field approximation and its limitations, before moving to the pair approximation. This framework allows us to derive the replicator equation on regular graphs, which will serve as the fundamental analytical benchmark for the numerical simulations performed throughout the thesis. After developing this theoretical background, we analyse evolutionary games on structured populations. We first consider square lattices and then move to complex network topologies, focusing on Erdős–Rényi, Watts–Strogatz and Barabási–Albert networks. For each case, we compare the stochastic simulation results with the analytical predictions, examine how the outcomes differ from the well-mixed case, and identify the regimes in which the analytical approximation succeeds or breaks down.

This chapter therefore provides both the theoretical and numerical foundation for the following analysis of the endogenous-feedback model on structured populations. The central question addressed in the next chapters will be whether the feedback-induced regimes identified in the well-mixed case, including chimera games, survive or change when individuals interact on lattices and complex networks.

3.1 Complex Networks as Models of Structured Populations

In this section, we introduce complex networks as mathematical models of structured populations. We present the basic graph-theoretical definitions and the main structural properties relevant for evolutionary dynamics, such as degree, path length, clustering, assortativity and degree heterogeneity. These concepts will be used throughout the chapter to compare different interaction structures.

Structured populations require a mathematical representation of who interacts with whom. Complex networks provide a natural framework for this purpose, because they allow one to describe a system in terms of its components and the interactions among them [27, 28]. A complex network is mathematically represented by a graph.

Definition 3.1. A **graph** is a pair $\mathcal{G} = (V, E)$, where $V = \{1, \dots, N\}$ is the set of vertices, or nodes, and E is the set of edges, or links.

In this representation, nodes correspond to the elementary units of the system, while edges encode the interactions between them. In the context of game theory, nodes represent individuals in the population, and edges represent possible strategic interactions. Edges can encode different types of relations depending on whether the interaction is symmetric or asymmetric. This leads to the distinction between directed and undirected graphs.

Definition 3.2. A graph is **directed** if its edges have an orientation, so that a link from node i to node j does not necessarily imply a link from node j to node i . A graph is **undirected** if its edges have no orientation, so that

$$(i, j) \in E \iff (j, i) \in E.$$




Figure 3.1: Examples of simple undirected graphs and their corresponding adjacency matrices. The path graph contains only nearest-neighbour connections along a chain; the triangle-with-tail graph contains a local cycle; the complete graph contains all possible links between nodes. *Illustration by the author.*

Directed graphs are useful to represent asymmetric relations, such as citation networks, food webs or follower networks. Since the evolutionary games considered in this thesis are based on reciprocal interactions, our focus will be on undirected graphs. Once the set of nodes and edges has been specified, the connectivity pattern of the graph can be represented algebraically through the adjacency matrix.

Definition 3.3. Given a graph $\mathcal{G} = (V, E)$ with N nodes, its **adjacency matrix** is the matrix $A \in \{0, 1\}^{N \times N}$ whose entries are defined as

$$A_{ij} = \begin{cases} 1, & \text{if nodes } i \text{ and } j \text{ are connected,} \\ 0, & \text{otherwise.} \end{cases}$$

Observe that for an undirected graph, the adjacency matrix is symmetric, namely, $A_{ij} = A_{ji}$. If self-interactions are excluded, then $A_{ii} = 0 \quad \forall i \in V$, that is, the diagonal entries of the adjacency matrix are zero. Fig. 3.1 shows three simple examples of undirected graphs and their corresponding adjacency matrices. A first important local property of a graph is the number of connections incident to each node. This quantity is called the degree of the node and measures how many direct interaction partners the node has.

Definition 3.4. The **degree** k_i of node i is the number of neighbours connected to it. For an undirected graph, it is given by

$$k_i = \sum_{j=1}^N A_{ij}.$$

At the level of the whole network, one can summarise the typical number of connections per node through the average degree.

Definition 3.5. The **average degree** of a graph is

$$\langle k \rangle = \frac{1}{N} \sum_{i=1}^N k_i.$$

If the graph is undirected and has $M = |E|$ edges, then

$$\langle k \rangle = \frac{2M}{N},$$

because each edge contributes to the degree of two nodes.

At one extreme, a graph may contain all possible connections between its nodes. This leads to the notion of a complete graph.

Definition 3.6. A **complete graph** is a graph in which every node is connected to all the other nodes. Equivalently, for an undirected graph with no self-loops, it contains the maximum possible number of edges,

$$M_{\max} = \frac{N(N-1)}{2}.$$

From the perspective of evolutionary game theory, a complete graph corresponds to an all-to-all interaction structure and provides a graph representation of the well-mixed population. Most real networks, however, are *sparse*: the number of actual links is much smaller than the maximum possible number of links. Equivalently,

$$\langle k \rangle \ll N.$$

Sparsity implies that most pairs of nodes are not directly connected. It is therefore important to describe how nodes can be reached indirectly through sequences of links. This leads to the notions of path and distance.

Definition 3.7. A **path** between two nodes i and j is a sequence of adjacent edges connecting i to j . The **length** of a path is the number of edges it contains. The **distance** between two nodes i and j , denoted by $d(i, j)$, is the length of the shortest path connecting them.

A basic global property of a graph is whether all nodes belong to the same connected structure.

Definition 3.8. A graph is **connected** if, for every pair of nodes $i, j \in V$, there exists at least one path connecting i and j .

Once distances between nodes are defined, one can characterise the overall size of the network in terms of the longest shortest path.

Definition 3.9. For a connected graph, the **diameter** is the largest distance between any pair of nodes:

$$D = \max_{i, j \in V} d(i, j).$$

While the diameter captures the largest separation in the network, the average path length measures the typical separation between nodes.

Definition 3.10. For a connected graph, the **average path length** is

$$\ell = \frac{1}{N(N-1)} \sum_{i \neq j} d(i, j).$$

Small values of ℓ indicate that information, perturbations or strategies can spread across the system through few intermediate steps.

Another important local property of a network is the tendency of neighbours of the same node to be connected to one another, a property is captured by the clustering coefficient.

Definition 3.11. Let node i have degree $k_i \geq 2$, and let e_i be the number of edges that are actually present among the neighbours of i . The **local clustering coefficient** of node i is defined as

$$C_i = \frac{e_i}{k_i(k_i - 1)/2} = \frac{2e_i}{k_i(k_i - 1)}.$$

For nodes with $k_i < 2$, one usually sets $C_i = 0$.

The local clustering coefficient measures the fraction of possible edges among the neighbours of a node that are actually present. Equivalently, it measures how many triangles are formed around that node. Observe that, if A is the adjacency matrix of a undirected graph, the number of triangles in which node i participates is

$$\tau_i = \frac{(A^3)_{ii}}{2}.$$

The factor 2 appears because each triangle containing node i can be traversed in two directions. Consequently, the total number of triangles in the graph is

$$\tau = \frac{\text{Tr}(A^3)}{6},$$

since each triangle is counted once for each of its three nodes and in two possible directions. At the global level, clustering can be summarised by averaging the local clustering coefficient over all nodes.

Definition 3.12. The **average clustering coefficient** of a graph is

$$C = \frac{1}{N} \sum_{i=1}^N C_i.$$

High clustering indicates the presence of locally cohesive groups and many triangular motifs. In evolutionary games, this is relevant because local clustering can affect the formation and stability of cooperative clusters. More generally, local organisation is not limited to triangles. One may also study higher-order cycles and motifs, such as squares, pentagons or larger closed structures. These quantities are related to higher powers of the adjacency matrix. For example, the diagonal elements of A^4 count closed walks of length four starting and ending at the same node. However, powers of the adjacency matrix count all closed walks, including degenerate ones involving repeated edges or backtracking. Therefore, extracting the number of simple cycles, such as squares or pentagons, requires additional corrections.

3.2 Analytical Approaches on Structured Populations

In this section, we introduce the analytical approaches used in this thesis to study structured populations. We first discuss the mean-field approximation, highlighting its assumptions and limitations. We then introduce the pair approximation, which represents the first analytical correction beyond mean-field theory by retaining information about neighbouring pairs and local correlations. The treatment will be restricted to regular graphs, namely graphs in which all nodes have the same degree k . Finally, we use pair approximation to obtain the replicator equation on regular graphs in the limit of weak selection, which will serve as the main analytical benchmark for the structured-population simulations discussed later on.

3.2.1 The Mean-Field Approximation

The replicator equation in the well-mixed setting has already been introduced in Sect. 1.4.4. Here, we recall it in order to highlight its interpretation as a *mean-field approximation*. This perspective is important because it clarifies which assumptions are made when the population is treated as well mixed, and therefore what changes when one moves to structured populations.

Consider a population composed of n strategies, and let $x_i(t)$ denote the frequency of strategy i at time t , with $\sum_{i=1}^n x_i(t) = 1$. Let A be the payoff matrix, which entries a_{ij} represent the payoff obtained by an individual using strategy i when interacting with an individual using strategy j . In the *mean-field approximation*, local interaction structure is averaged out. Each individual is assumed to interact with a representative sample of the whole population, so that its local neighbourhood has the same expected composition as the global population. Therefore, the probability of meeting a player using strategy j is simply the global frequency x_j . Under this assumption, the expected payoff, or fitness, of strategy i is

$$f_i = \sum_{j=1}^n a_{ij}x_j = (A\mathbf{x})_i,$$

while the average fitness of the population is

$$\phi = \sum_{i=1}^n x_i f_i = \mathbf{x}^T A \mathbf{x}.$$

The corresponding replicator equation for each strategy is

$$\dot{x}_i = x_i(f_i - \phi), \quad i = 1, \dots, n. \quad (3.1)$$

Equation (3.1) describes selection at the level of global strategy frequencies. Strategies whose fitness is larger than the population average increase in frequency, whereas strategies whose fitness is smaller than the average decrease. The crucial point is that the local interaction environment of each individual is replaced by the average composition of the whole population. Thus, the model does not distinguish between

local and global strategy frequencies. From the perspective of network theory, the mean-field approximation corresponds to a complete graph, or fully connected interaction structure. If the population has N individuals, every individual can interact with every other individual. Thus, no local neighbourhood is privileged: each player samples the whole population, and the expected composition of its interaction partners coincides with the global frequency vector $\mathbf{x}$.

More generally, the limitation of the mean-field approximation can be expressed in terms of conditional neighbourhood probabilities. Let $q_{i|j}$ denote the probability that a neighbour of an individual using strategy j uses strategy i . Mean-field theory assumes that the local composition around any player coincides with the global population composition, namely

$$q_{i|j} = x_i.$$

This equality holds under random mixing, where the strategy of a neighbour is independent of the strategy of the focal individual. In structured populations, however, local correlations typically imply

$$q_{i|j} \neq x_i.$$

The mean-field approximation therefore neglects spatial correlations, network topology, clustering and local fluctuations in the composition of neighbourhoods. Moving to structured populations means relaxing precisely this assumption. Individuals no longer interact with a representative sample of the whole population, but only with the neighbours specified by an underlying graph. As a result, local strategy frequencies may differ from global frequencies, and the replicator equation may no longer provide an accurate description of the evolutionary dynamics. The aim of pair approximation is precisely to retain this first level of local information by keeping track of the conditional probabilities $q_{i|j}$.

3.2.2 The Pair Approximation

The mean-field approximation describes the population only through the global strategy frequencies x_i . This is appropriate when interactions are well mixed, because the local neighbourhood of each individual is assumed to have the same composition as the whole population. In structured populations, however, individuals interact only with their neighbours, and local strategy frequencies may differ from global frequencies. Pair approximation is the first step beyond mean-field theory: instead of tracking only the densities of single strategies, it also tracks the densities of neighbouring pairs. In this subsection we describe its mathematical formulation following the approach of [12].

Let $\mathcal{S}$ denote the set of sites of a finite but large population, and assume that each site has exactly k neighbours. The set of links between neighbouring sites is denoted by $\mathcal{L} \subseteq \mathcal{S} \times \mathcal{S}$. Thus, two sites $x, y \in \mathcal{S}$ form a pair if $xy \in \mathcal{L}$. Each site $x \in \mathcal{S}$ is in a state

$$\sigma_x \in \Omega,$$

where Ω is the set of possible states or strategies. The state of the whole population is denoted by

$$\sigma = (\sigma_x)_{x \in \mathcal{S}},$$




Figure 3.2: Local environment of a neighbouring pair xy . The central edge xy identifies the pair whose state may change through a local update event. The neighbours connected to x form the left environment E_{xy}^x , while the neighbours connected to y form the right environment E_{xy}^y . The full local environment of the pair is the union of these two sets, $E_{xy} = E_{xy}^x \cup E_{xy}^y$. Adapted from [12].

while the state of a pair xy is denoted by

$$\sigma_x \sigma_y.$$

The microscopic dynamics changes the state of sites through local update events. In the general pair-dynamics formalism, these events can be written as pair events of the form

$$\sigma_x \sigma_y \longrightarrow \sigma'_x \sigma'_y,$$

or, equivalently,

$$ab \longrightarrow kl,$$

where $a, b, k, l \in \Omega$. This notation also includes single-site events as a special case, for example when only one of the two states changes. Each event has an associated rate

$$r_\sigma(ab \rightarrow kl),$$

which may depend on the current state of the population σ . Since interactions are local, this dependence is usually restricted to the immediate neighbourhood of the pair. Denoting by E_{xy} the local environment of the pair xy , one writes

$$r_\sigma(ab \rightarrow kl) = r(ab \rightarrow kl \mid \sigma_{E_{xy}}).$$

An example of a local environment is illustrated in Fig. 3.2: the pair xy has an environment on the x side and an environment on the y side, and the union of these two sets determines the entire environment on which the local transition rate may depend.

We now introduce pair frequencies. Let x_a denote the frequency of sites in state a , and let x_{ab} denote the frequency of neighbouring pairs in state ab . From these

quantities one can define the conditional probability

$$q_{i|a} = \frac{x_{ia}}{x_a},$$

which gives the probability that a neighbour of an a -site is in state i . More generally, when the local environment of a pair is considered, we write

$$q_{i|ab}$$

for the probability that an additional neighbour of the a -site in an ab pair is in state i . Similarly, $q_{j|ba}$ denotes the probability that an additional neighbour of the b -site in an ab pair is in state j . The dynamics of pair frequencies can be derived by a gain–loss bookkeeping argument. Consider a pair in state ab . Its frequency changes when microscopic events either create or destroy pairs of this type. There are two ways in which this can happen. First, a pair can be created or destroyed directly by an event acting on the pair itself:

$$ij \rightarrow ab \quad \text{or} \quad ab \rightarrow kl.$$

Second, a pair can be created or destroyed indirectly by an event occurring in its neighbourhood. For example, if an event changes an adjacent pair in such a way that the state of one neighbour of an a -site changes from i to b , then an ai pair becomes an ab pair. This distinction between direct and indirect effects is illustrated in Fig. 3.3. Let $\bar{r}_\sigma(\cdot)$ denote the average rate of an event over the relevant local environments in the current population state σ . The rate of change of the pair frequency x_{ab} is obtained by summing all events that destroy ab pairs and all events that create them. The loss terms are the following. First, an ab pair can be directly transformed into another pair kl . Second, an event in the neighbourhood of the a -site can destroy the ab pair. Since the a -site has $k-1$ neighbours other than the b -site, this contribution is weighted by $(k-1)q_{i|ab}$. Third, an analogous neighbourhood event can occur on the b -side, weighted by $(k-1)q_{j|ba}$. The gain terms are analogous. An ab pair can be created directly from a source pair ij . It can also be created indirectly when an event occurring next to a pair jb or ai changes one of the neighbouring states so as to produce an ab pair. Putting these contributions together gives the pair-dynamics master equation

$$\begin{aligned} \frac{dx_{ab}}{dt} = & -x_{ab} \sum_{k,l} \left[\bar{r}_\sigma(ab \rightarrow kl) + \sum_i (k-1)q_{i|ab} \bar{r}_\sigma(ia \rightarrow kl \mid iab) \right. \\ & \left. + \sum_j (k-1)q_{j|ba} \bar{r}_\sigma(bj \rightarrow kl \mid abj) \right] \\ & + \sum_{i,j} \left[x_{ij} \bar{r}_\sigma(ij \rightarrow ab) + x_{jb} \sum_k (k-1)q_{i|jb} \bar{r}_\sigma(ij \rightarrow ka \mid ijb) \right. \\ & \left. + x_{ai} \sum_l (k-1)q_{j|ia} \bar{r}_\sigma(ij \rightarrow bl \mid aij) \right]. \end{aligned}$$

Equation (3.2.2) shows why pair dynamics is more informative than mean-field theory. The equation does not only depend on the global frequencies x_a , but also on




Figure 3.3: Direct and indirect contributions to the dynamics of pair frequencies. (a) A pair of type ab can be created directly when a microscopic event transforms a source pair ij into ab . (b) It can also be created indirectly when an event occurs in the neighbourhood of an existing pair, for example when an event next to an ai pair changes the state i into b , thereby turning the pair ai into ab . This bookkeeping is the basis of pair-dynamics equations, where the evolution of pair frequencies depends both on events acting on the pair itself and on events acting in its local neighbourhood. Adapted from [12].

pair frequencies and on conditional probabilities describing the local neighbourhood of pairs. At the same time, it also shows the main difficulty: the dynamics of pairs is not closed at the pair level. Terms such as $q_{i|ab}$ depend on triplet information, because they condition not only on one site but on an entire pair. The purpose of *pair approximation* is to close this hierarchy by expressing higher-order local quantities in terms of pair-level quantities. The standard closure assumes that, once the state of the nearest neighbour is known, the additional information carried by more distant sites can be neglected. Mathematically, one writes

$$q_{i|j\ell} \simeq q_{i|j}. \quad (3.2)$$

This approximation means that the probability of finding an i -strategist next to a j -strategist is assumed not to depend explicitly on the state ℓ of a second-order neighbour. The local structure is therefore described through pair correlations, while higher-order correlations are approximated from them.

In ecological applications, the pair events may correspond to birth, death or movement processes: birth creates an occupied site next to an already occupied one, death changes an occupied site into an empty site, and movement swaps the states of neighbouring sites. In evolutionary games, the same formalism is used with strategy-update events instead. These update rules, such as birth–death, death–birth, imitation and pairwise comparison updating, will be specified explicitly in the next subsection, where pair approximation is used to derive the replicator equation on graphs.

3.2.3 The Replicator Equation on Graphs

We now apply the pair approximation method discussed in the previous subsection to derive the replicator equation on graphs. We follow the approach of [13], which derives an effective replicator equation for evolutionary games on regular graphs by combining weak selection with pair approximation.

We consider *regular graphs* of degree k , so that each individual has exactly k neighbours, and we restrict to the case $k > 2$. Each individual accumulates a payoff Π by interacting with all its neighbours. The fitness of an individual is given by

$$W = 1 - w + w\Pi,$$

where $w \in [0, 1]$ is the intensity of selection. Throughout this derivation, we work in the *weak-selection regime*, that is,

$$w \ll 1.$$

We consider four updating rules for the evolutionary dynamics:

1. **Birth–death** (BD) updating. An individual is selected for reproduction from the whole population with probability proportional to fitness. Its offspring then replaces one randomly chosen neighbour.
2. **Death–birth** (DB) updating. A randomly selected individual is removed, and its neighbours compete to occupy the vacant site with probability proportional to fitness.
3. **Imitation** (IM) updating. A randomly selected individual is given the opportunity to revise its strategy. It may retain its current strategy or adopt the strategy of one of its neighbours with probability proportional to fitness.
4. **Pairwise comparison** (PC) updating. A focal individual compares its payoff Π_X with that of a randomly chosen neighbour, say Π_Y , and adopts the neighbour's strategy with probability given by the Fermi updating rule introduced in Eq. (1.6),

$$p_{X \leftarrow Y} = \frac{1}{1 + \exp[-\beta(\Pi_Y - \Pi_X)]}.$$

For $\beta \rightarrow 0$, updating becomes almost independent of payoff differences, whereas for $\beta \rightarrow \infty$ it approaches deterministic imitation of the better-performing neighbour.

These four updating rules are illustrated in Fig. 3.4. In BD and DB, reproduction and replacement occur in different orders; in IM, a focal player updates by possibly copying one of its neighbours; in PC, strategy revision is based on payoff comparison through the Fermi rule. As will be discussed below, in the present weak-selection setting on regular graphs the PC rule leads to the same effective first-order correction as BD updating. For this reason, the final replicator equation on graphs can be written in terms of three distinct cases: BD (equivalently PC), DB, and IM. Let A be the payoff matrix, which entries a_{ij} represent the payoff obtained by an individual




Figure 3.4: Microscopic updating rules considered in the derivation of the replicator equation on graphs. Panel (a) shows *birth-death* (BD) updating: an individual is first selected for reproduction, and its offspring then replaces a randomly chosen neighbour. Panel (b) shows *death-birth* (DB) updating: a focal individual is first removed, and its neighbours then compete to occupy the vacant site. Panel (c) shows *imitation* (IM) updating: a focal individual is selected for strategy revision and may imitate one of its neighbours. Panel (d) shows *pairwise comparison* (PC) updating: a focal individual compares its payoff with that of a neighbour and may adopt the neighbour's strategy according to the Fermi rule. *Adapted from [13].*

using strategy i when interacting with an individual using strategy j . Let x_i denote the global frequency of strategy i . In the well-mixed scenario, the payoff of strategy i is

$$f_i = \sum_{j=1}^n x_j a_{ij} = (A\mathbf{x})_i,$$

while the population average payoff is

$$\phi = \sum_{i=1}^n x_i f_i = \mathbf{x}^T A \mathbf{x}.$$

The effect of graph structure enters through an additional local competition term. More precisely, one introduces an $n \times n$ matrix B , whose entries b_{ij} depend on the

updating rule given by:

$$\begin{aligned}
 \text{BD/PC:} \quad b_{ij} &= \frac{a_{ii} + a_{ij} - a_{ji} - a_{jj}}{k-2}, \quad i, j = 1, \dots, n. \\
 \text{DB:} \quad b_{ij} &= \frac{(k+1)a_{ii} + a_{ij} - a_{ji} - (k+1)a_{jj}}{(k+1)(k-2)}, \quad i, j = 1, \dots, n. \\
 \text{IM:} \quad b_{ij} &= \frac{(k+3)a_{ii} + 3a_{ij} - 3a_{ji} - (k+3)a_{jj}}{(k+3)(k-2)}, \quad i, j = 1, \dots, n.
 \end{aligned} \tag{3.3}$$

Let us further introduce the quantities

$$g_i = \sum_{j=1}^n x_j b_{ij}, \quad i = 1, \dots, n. \tag{3.4}$$

Using the pair-approximation framework developed in the previous subsection, the resulting *replicator equation on graphs* is

$$\dot{x}_i = x_i (f_i + g_i - \phi), \quad i = 1, \dots, n. \tag{3.5}$$

Thus, the structured-population dynamics differs from the standard replicator equation only by the additional term g_i , which captures the effect of local competition induced by graph structure and by the updating rule. The derivation of Eq. (3.5) is provided below for birth–death updating.

Observe that the matrix B is antisymmetric, namely,

$$b_{ij} = -b_{ji}, \quad b_{ii} = 0,$$

and therefore

$$\sum_{i=1}^n x_i g_i = \sum_{i,j=1}^n x_i x_j b_{ij} = 0.$$

As a consequence, the average fitness of the population remains

$$\dot{\phi} = \sum_{i=1}^n x_i \dot{f}_i,$$

exactly as in the standard well-mixed replicator equation. Equation (3.5) can also be written as

$$\dot{x}_i = x_i \left[\sum_{j=1}^n x_j (a_{ij} + b_{ij}) - \phi \right].$$

This form shows that, under weak selection and pair approximation, moving from a well-mixed population to a regular graph is equivalent to a transformation of the payoff matrix,

$$A \longrightarrow A + B,$$

where A is the original payoff matrix and B is the graph-induced correction matrix determined by the update rule.

Derivation of the Replicator Equation on Graphs (BD Updating) We now outline the derivation of Eq. (3.5) for BD updating. The derivations of the other two update rules, DB and IM follow the same logic and can be found in [13]. Under birth–death updating, an individual is selected for reproduction with probability proportional to its fitness, and its offspring replaces a randomly chosen neighbour. On a graph, limited dispersal induces local correlations: individuals using the same strategy tend to be more clustered than they would be under random mixing. Therefore, as discussed in Sect. 3.2.2, we introduce the conditional probability $q_{i|j}$, defined as the probability that a neighbour of a j -strategist uses strategy i . The dynamics of the local frequencies $q_{i|j}$ is obtained from the pair-frequency dynamics. Recall that $q_{i|j} = \frac{x_{ij}}{x_j}$ where x_{ij} is the frequency of neighbouring pairs in which an i -strategist is adjacent to a j -strategist, while x_j is the global frequency of strategy j . Under weak selection, the global frequencies x_j vary only at order w , whereas local correlations relax on the faster $\mathcal{O}(1)$ time scale. Hence, on the fast time scale, the global frequencies can be treated as approximately constant, and one obtains

$$\dot{q}_{i|j} = \frac{\dot{x}_{ij}}{x_j} + \mathcal{O}(w).$$

To compute $\dot{x}_{ij}$, we use a gain–loss argument under BD updating. Each update replaces one node and therefore destroys and creates the k ordered pairs incident to that node. The loss of $i \rightarrow j$ pairs is proportional to the expected number of i -neighbours of a j -node, namely $kq_{i|j}$. The creation term has two contributions. First, the copying event creates one pair between the offspring and the reproducing individual; this contribution is captured by the Kronecker delta δ_{ij} , where $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ otherwise. Second, the remaining $k - 1$ neighbours of the replaced node contribute through local correlations. Using the pair-approximation introduced in Eq. (3.2), two-step correlations are approximated as $q_{i|\ell j} \simeq q_{i|\ell}$, which gives the term $(k-1) \sum_{\ell=1}^n q_{i|\ell} q_{\ell|j}$. After normalisation, the local-frequency dynamics becomes

$$\dot{q}_{i|j} = \frac{2}{k} \left[\delta_{ij} + (k-1) \sum_{\ell=1}^n q_{i|\ell} q_{\ell|j} - kq_{i|j} \right] + \mathcal{O}(w), \quad (3.6)$$

where δ_{ij} is the Kronecker delta. Setting $\dot{q}_{i|j} = 0$ in Eq. (3.6) and using the consistency identity $q_{i|j}x_j = q_{j|i}x_i$, we obtain the equilibrium local frequencies

$$q_{i|j}^* = \frac{(k-2)x_i + \delta_{ij}}{k-1}.$$

This implies that

$$q_{i|i}^* > x_i, \quad q_{i|j}^* < x_i \quad (j \neq i).$$

Thus, individuals adopting strategy i are surrounded by a higher proportion of i -neighbours than predicted by the global frequency x_i , whereas individuals using any other strategy have a lower proportion of i -neighbours than expected under random mixing. This expresses the local assortment induced by graph structure.

Having determined the equilibrium local frequencies, we can derive the dynamics of the global strategy frequencies. For convenience, we write $q_{i|j}^*$ simply as $q_{i|j}$. We introduce the notation $(i; k_1, \dots, k_n)$ -player to denote a player using strategy i

who has k_1 neighbours using strategy 1, k_2 neighbours using strategy 2, and so on up to k_n neighbours using strategy n . Here, the quantities $k_1, \dots, k_n$ denote the composition of the neighbourhood and satisfy $\sum_{\ell=1}^n k_\ell = k$, where k is the degree of the regular graph. Let us consider one elementary step of BD updating. The number of i -strategists increases by one if two events occur. First, an $(i; k_1, \dots, k_n)$ -player is selected for reproduction. Second, its offspring replaces one of its neighbours who does not use strategy i . The probability that an $(i; k_1, \dots, k_n)$ -player is selected for reproduction is

$$\frac{x_i \left(\frac{k!}{k_1! \dots k_n!} q_{1|i}^{k_1} \dots q_{n|i}^{k_n} \right) W(i; k_1, \dots, k_n)}{\bar{W}},$$

where $W(i; k_1, \dots, k_n)$ denotes the fitness of the player and $\bar{W}$ is the average fitness in the population. The fitness is given by

$$W(i; k_1, \dots, k_n) = 1 - w + w \sum_{\ell=1}^n k_\ell a_{i\ell}.$$

The second event, namely replacing a neighbour who does not use strategy i , occurs with probability $1 - \frac{k_i}{k}$. Conversely, the number of i -strategists decreases by one when two events occur. First, a $(j; k_1, \dots, k_n)$ -player with $j \neq i$ is selected for reproduction. Second, its offspring replaces one of its neighbours using strategy i . The probability that a $(j; k_1, \dots, k_n)$ -player is selected for reproduction is

$$\frac{x_j \left(\frac{k!}{k_1! \dots k_n!} q_{1|j}^{k_1} \dots q_{n|j}^{k_n} \right) W(j; k_1, \dots, k_n)}{\bar{W}},$$

where

$$W(j; k_1, \dots, k_n) = 1 - w + w \sum_{\ell=1}^n k_\ell a_{j\ell}.$$

The probability that the offspring replaces a neighbour using strategy i is $\frac{k_i}{k}$. From these gain and loss terms one can compute the expected change in the global frequency of strategy i , denoted by $E[\Delta x_i]$, during one elementary update step of duration Δt . In the infinite-population limit, we obtain the deterministic dynamics

$$\dot{x}_i = \frac{E[\Delta x_i]}{\Delta t}.$$

Carrying out the calculation and expanding to first order in the selection intensity w yields

$$\dot{x}_i \simeq w \frac{(k-2)^2}{k-1} x_i (f_i + g_i - \phi).$$

Here

$$f_i = \sum_{j=1}^n x_j a_{ij}, \quad \phi = \sum_{i=1}^n x_i f_i = \sum_{i,j=1}^n x_i x_j a_{ij},$$

and

$$g_i = \sum_{j=1}^n x_j b_{ij},$$

with

$$b_{ij} = \frac{a_{ii} + a_{ij} - a_{ji} - a_{jj}}{k - 2}.$$

The prefactor $w \frac{(k-2)^2}{k-1}$ only rescales time. Neglecting this constant factor, we obtain

$$\dot{x}_i = x_i (f_i + g_i - \phi),$$

which is the replicator equation on graphs for BD updating.

Replicator Equation on Graphs for Social Dilemmas We now specialise Eq. (3.5) to the case of two strategies, cooperation and defection, using the payoff parametrisation introduced in Sect. 1.4.4. In particular, we consider the payoff matrix

$$G(D_g, D_r) = \begin{pmatrix} 1 & -D_r \\ 1 + D_g & 0 \end{pmatrix},$$

where the first row and column correspond to cooperation and the second row and column to defection. Here we focus on the BD updating rule since, on regular graphs and in the weak-selection limit, it yields the same effective graph correction as the pairwise comparison update considered throughout the thesis. Let ρ denote the frequency of cooperators. The replicator equation on graphs becomes

$$\dot{\rho} = \rho(1 - \rho) [(\Pi_C + g_C) - (\Pi_D + g_D)],$$

where Π_C and Π_D are the well-mixed payoffs of cooperators and defectors as per Eq. 1.4.4, while g_C and g_D are the graph-induced correction terms defined in Eq. (3.4). The correction matrix B is given by Eq. (3.3). Using the parametrisation above, one obtains

$$b_{CD} = -\frac{D_r + D_g}{k - 2}, \quad b_{DC} = \frac{D_r + D_g}{k - 2},$$

while

$$b_{CC} = b_{DD} = 0.$$

Therefore,

$$\begin{aligned} g_C &= (1 - \rho)b_{CD} = -\frac{(1 - \rho)(D_r + D_g)}{k - 2}, \\ g_D &= \rho b_{DC} = \frac{\rho(D_r + D_g)}{k - 2}. \end{aligned}$$

Using the payoff difference already derived in Sect. 1.4.4,

$$\Pi_C - \Pi_D = \rho(D_r - D_g) - D_r,$$

we obtain

$$\dot{\rho} = \rho(1 - \rho) \left[\rho(D_r - D_g) - D_r - \frac{D_r + D_g}{k - 2} \right]. \quad (3.7)$$

As in the well-mixed case, the boundary states

$$\rho = 0, \quad \rho = 1$$

are stationary points of the dynamics. In addition, an internal stationary point may exist. Setting the term in square brackets in Eq. (3.7) equal to zero gives

$$\rho^* = \frac{D_r + \frac{D_r + D_g}{k-2}}{D_r - D_g}.$$

This equilibrium is admissible only if

$$0 < \frac{D_r + \frac{D_r + D_g}{k-2}}{D_r - D_g} < 1.$$

Thus, in contrast with the well-mixed case, the existence and the value of internal equilibria depend not only on the dilemma strength parameters D_g and D_r , but also on the degree k of the graph. Therefore, the local competition term induced by pair approximation changes the equilibrium structure of the dynamics.

3.3 Evolutionary Dynamics on Square Lattices

In this section, we consider the square lattice as the first structured population model for evolutionary dynamics beyond the well-mixed approximation. We first introduce the simulation framework used throughout this section, specifying the lattice geometry, the update rule and the neighbourhood structures considered. In particular, we distinguish between the von Neumann neighbourhood and the Moore neighbourhood. We then perform numerical simulations and compare the results with the analytical predictions obtained from the replicator equation on graphs. Finally, we analyse the effect of adding self-interaction to the system.

3.3.1 Simulation Framework

We consider a square lattice of linear size $L = 200$, corresponding to a population of

$$N = L^2 = 40000$$

individuals. Each individual occupies one site of the lattice, and each site is occupied by exactly one player. Thus, there are no empty sites. Each player can adopt one of two strategies, cooperation C or defection D . Periodic boundary conditions are imposed in both spatial directions, so that the lattice has the topology of a torus and boundary effects are removed. An example of a random initial configuration is shown in Fig. 3.5. In this example, the initial fraction of cooperators is $\rho_0 = 0.5$. Cooperators are shown in black, while defectors are shown in white.

Players interact locally with their neighbours on the lattice. Depending on the simulation setting, we consider either a von Neumann neighbourhood, with degree $k = 4$, or a Moore neighbourhood, with degree $k = 8$, as illustrated in Fig. 3.6. In the von Neumann case, each player interacts with its four nearest neighbours in the horizontal and vertical directions. In the Moore case, each player also interacts with the four diagonal neighbours. In addition, we also consider the possible inclusion of




Figure 3.5: Example of a random initial configuration on a square lattice with periodic boundary conditions. The lattice has linear size $L = 200$, corresponding to $N = 40000$ individuals. Each site is occupied by one player, which can be either a cooperator or a defector. Cooperators are represented in black, while defectors are represented in white. The configuration shown here is initialised with cooperation level $\rho_0 = 0.5$. *Illustration by the author.*

self-interaction, in which each player interacts once with itself in addition to interacting with its neighbours. The game played in each interaction is the two-strategy game parametrised by the dilemma-strength parameters D_g and D_r , introduced in Sect. 1.4.4:

$$G(D_g, D_r) = \begin{pmatrix} 1 & -D_r \\ 1 + D_g & 0 \end{pmatrix}.$$

The first row and column correspond to cooperation, while the second row and column correspond to defection. The evolutionary dynamics follows pairwise comparison with Fermi updating, as introduced in Eq. (1.6). At each elementary update, a focal player X is selected uniformly at random, and one of its neighbours Y is chosen uniformly at random. Both players accumulate their payoffs by interacting with all their neighbours in the chosen neighbourhood. If self-interaction is included, each player also interacts once with itself. The focal player X then adopts the strategy of Y with probability

$$p_{X \leftarrow Y} = \frac{1}{1 + \exp[-\beta(\Pi_Y - \Pi_X)]},$$

where Π_X and Π_Y are the total payoffs of the two players and β is the selection intensity. Time is measured in Monte Carlo sweeps, which correspond to N elementary update attempts, so that each individual has, as expected value, one opportunity to revise its strategy per sweep. For each parameter setting, we perform 10 independent runs, starting from independently generated initial configurations. The




Figure 3.6: Schematic representation of the two neighbourhood structures considered on the square lattice. In both panels, the central site represents the focal player, while the highlighted surrounding sites indicate the players with whom it interacts. **Left:** von Neumann neighbourhood, in which the focal player interacts with its four nearest neighbours in the horizontal and vertical directions, corresponding to degree $k = 4$. **Right:** Moore neighbourhood, in which the focal player interacts with all eight surrounding neighbours, including the four diagonal ones, corresponding to degree $k = 8$. *Illustration by the author.*

cooperation level $\rho(t)$ is recorded as a function of Monte Carlo sweeps and the results are averaged over the independent runs. Simulations are run for several sweeps until convergence of the cooperation level to a stationary regime.

3.3.2 Von Neumann Neighbourhood

We first consider evolutionary dynamics on a square lattice with von Neumann neighbourhood, corresponding to degree $k = 4$, without self-interaction. In the simulations, the intensity of selection in the Fermi update rule is fixed to $\beta = 0.2$. Although this value does not strictly place the stochastic dynamics in the weak-selection regime assumed in the derivation of the replicator equation on graphs, it is sufficiently small for the dynamics to remain qualitatively comparable with the analytical prediction. We therefore use the graph replicator equation derived in Sect. 3.2.3 as a theoretical benchmark. For a von Neumann neighbourhood, the degree of the graph is $k = 4$. Substituting this value into Eq. (3.7) gives

$$\dot{\rho} = \rho(1 - \rho) \left[\rho(D_r - D_g) - \frac{3D_r + D_g}{2} \right]. \quad (3.8)$$

This equation predicts two boundary stationary states,

$$\rho = 0, \quad \rho = 1,$$

corresponding respectively to full defection and full cooperation. In addition, an internal stationary point may exist. Setting the term in square brackets in Eq. (3.8)

equal to zero gives

$$\rho^* = \frac{3D_r + D_g}{2(D_r - D_g)}.$$

This fixed point is admissible only if it lies in the physical interval, namely

$$0 < \frac{3D_r + D_g}{2(D_r - D_g)} < 1.$$

The stability of the stationary states is determined by the sign of the Jacobian $J(\rho) = \frac{d\dot{\rho}}{d\rho}$. Evaluating it at the fixed points gives the following stability conditions:

- $\rho = 0$ is stable if

$$3D_r + D_g > 0.$$

- $\rho = 1$ is stable if

$$D_r + 3D_g < 0.$$

- the internal equilibrium $\rho^* \in (0, 1)$, when admissible, is stable if

$$D_r - D_g < 0.$$

These conditions show explicitly how the graph correction modifies the well-mixed stability structure. In the well-mixed case, the stability of the boundary equilibria is controlled by the signs of D_r and D_g . For the von Neumann graph approximation, these boundaries are shifted to the oblique lines

$$3D_r + D_g = 0, \quad D_r + 3D_g = 0.$$

To visualise the analytical prediction across the whole dilemma space, we constructed a deterministic heatmap of the equilibrium cooperation level predicted by Eq. (3.8). We scanned the domain $D_g, D_r \in [-1, 1]$ using a fine grid of 1001×1001 equally spaced parameter pairs. For each pair (D_g, D_r) , we computed the stationary points of Eq. (3.8) and determined their stability from the Jacobian. If a stable internal equilibrium was present, the analytical outcome was assigned as $\rho = \rho^*$. If only one boundary equilibrium was stable, the corresponding value $\rho = 0$ or $\rho = 1$ was assigned. In bistable regions, where both boundary equilibria are stable and the internal equilibrium is unstable, the predicted outcome was assigned by comparing the initial condition used in the simulations, $\rho_0 = 0.5$, with the unstable threshold ρ^* . If $\rho_0 > \rho^*$, the predicted final state is full cooperation, whereas if $\rho_0 < \rho^*$, the predicted final state is full defection. The resulting analytical prediction is shown in Fig. 3.7. Compared with the well-mixed prediction shown in Fig. 1.4, the Prisoner's Dilemma, Stag Hunt and Harmony Game regions remain unchanged. The main differences occur in the Snowdrift region, where the graph correction shifts the location and admissibility of the stable internal equilibrium. To highlight the effect of graph structure more directly, we compare the asymptotic cooperation level predicted by the graph replicator equation with the corresponding well-mixed prediction. We define

$$\Delta\rho_\infty = \rho_\infty^{\text{graph}} - \rho_\infty^{\text{WM}}, \quad (3.9)$$




Figure 3.7: Analytical prediction of the equilibrium cooperation level in the (D_g, D_r) dilemma space obtained from the replicator equation on graphs for a von Neumann neighbourhood, corresponding to degree $k = 4$. The prediction is computed from Eq. (3.8) on a 1001×1001 grid in the domain $D_g, D_r \in [-1, 1]$. For each parameter pair, the stationary points are computed and their stability is determined from the Jacobian. The boundary equilibrium $\rho = 0$ is stable when $3D_r + D_g > 0$, while $\rho = 1$ is stable when $D_r + 3D_g < 0$. An internal equilibrium $\rho^* = (3D_r + D_g)/[2(D_r - D_g)]$ is admissible when it lies in $(0, 1)$ and is stable when $D_r - D_g < 0$. In bistable regions, the final analytical outcome is assigned by comparing the simulation initial condition $\rho_0 = 0.5$ with the unstable threshold ρ^* . *Original illustration by the author.*

where $\rho_\infty^{\text{graph}}$ is the equilibrium cooperation level predicted by Eq. (3.8), while ρ_∞^{WM} is the corresponding equilibrium cooperation level predicted by the well-mixed replicator equation. The resulting difference heatmap is shown in Fig. 3.8. The difference heatmap shows that the graph correction mainly affects the Snowdrift region. When both the well-mixed and graph equations predict a stable internal equilibrium, their difference is

$$\rho_\infty^{\text{graph}} - \rho_\infty^{\text{WM}} = \frac{D_r + D_g}{2(D_r - D_g)}.$$

In the Snowdrift region, where $D_g > 0$ and $D_r < 0$, the denominator is negative. Therefore, the graph replicator equation predicts a higher equilibrium cooperation level than the well-mixed equation when $D_r < -D_g$, and a lower equilibrium cooperation level when $D_r > -D_g$. The graph correction can also make the internal equilibrium non-admissible in part of the Snowdrift region. In particular, for

$$3D_r + D_g > 0,$$




Figure 3.8: Difference between the analytical prediction of the graph replicator equation for a von Neumann neighbourhood and the corresponding well-mixed replicator prediction. The plotted quantity is $\Delta\rho_\infty = \rho_\infty^{\text{graph}} - \rho_\infty^{\text{WM}}$ defined in Eq. (3.9). Positive values indicate that the graph replicator equation predicts a higher equilibrium cooperation level than the well-mixed approximation, while negative values indicate a lower predicted cooperation level. The main differences are concentrated in the Snowdrift region, where the stable internal equilibrium is shifted by the graph-induced correction term. *Original illustration by the author.*

the internal equilibrium lies outside the physical interval and the analytical prediction collapses to full defection. Conversely, for

$$D_r + 3D_g < 0,$$

the internal equilibrium is again non-admissible, and the stable outcome is full cooperation. This shows that the graph correction can modify the equilibrium structure predicted by the well-mixed approximation, both by shifting coexistence levels and by replacing coexistence with absorbing states in part of the Snowdrift region.

We now compare the analytical predictions of the replicator equation on graphs for the von Neumann neighbourhood with the corresponding stochastic simulations on the square lattice. The simulation setting is the one described in Sect. 3.3.1, and is therefore not repeated here in full. In particular, we consider a square lattice with periodic boundary conditions, no self-interaction, and selection intensity $\beta = 0.2$. Figure 3.9 reports the time evolution of the cooperation level for four representative games in the dilemma space, namely

$$(D_g, D_r) = (0.2, 0.2), (-0.5, 0.8), (0.5, -0.3), (0.5, -0.8).$$

These four parameter pairs sample different regions of the dilemma space and allow one to compare directly the simulation outcomes with the predictions of the graph

replicator equation for $k = 4$ given by Eq. (3.7), as well as with the well-mixed replicator prediction discussed in Sect. 1.4.4. As can be seen from the figure, the




Figure 3.9: Time evolution of the cooperation level ρ for Monte Carlo simulations on a square lattice with von Neumann neighbourhood ($k = 4$), periodic boundary conditions, and no self-interaction. The four panels correspond to the games $(D_g, D_r) = (0.2, 0.2)$, $(-0.5, 0.8)$, $(0.5, -0.3)$ and $(0.5, -0.8)$, respectively. In each panel, the dynamics is shown for three initial cooperation levels, $\rho_0 = 0.1, 0.5, 0.9$, and results are averaged over 10 independent realisations. The solid red horizontal line marks the stable equilibrium cooperation level predicted by the replicator equation on graphs for the von Neumann neighbourhood, while the black dashed horizontal line marks the corresponding well-mixed prediction. The selection intensity is $\beta = 0.2$, and time is measured in Monte Carlo sweeps. Different panels are shown over different Monte Carlo sweep windows. The figure shows very good agreement between simulations and pair approximation in the Prisoner’s Dilemma and Stag Hunt cases, while visible discrepancies emerge in the Snowdrift region. *Original simulation by the author.*

agreement between simulations and the pair-approximation prediction is very good in the Prisoner’s Dilemma and Stag Hunt cases, corresponding to panels (a) and (b). By contrast, visible deviations emerge in the Snowdrift region, corresponding to panels (c) and (d). In particular, the discrepancy is especially evident in panel (c), corresponding to $(D_g, D_r) = (0.5, -0.3)$. In these cases, the prediction based on the replicator equation on graphs is clearly more accurate than the well-mixed one, but it still does not fully capture the actual lattice dynamics. This reflects the fact that pair approximation retains only pair-level correlations, whereas on the square lattice higher-order spatial correlations, cluster geometry and interface effects may play a

significant role. Understanding this breakdown of pair approximation constitutes one of the main focuses of the thesis and will be analysed more deeply in the following chapters. The spatial arrangement of individuals on the square lattice also makes it possible to visualise the evolutionary dynamics directly in configuration space. In particular, the lattice representation provides a clear picture of how clusters of cooperators and defectors form, evolve and eventually disappear. As an illustrative example, in Fig. 3.10 we report a sequence of snapshots for the case corresponding to panel (a) of Fig. 3.9, namely

$$(D_g, D_r) = (0.2, 0.2).$$

Black sites represent cooperators, whereas white sites represent defectors. The sequence shows the progressive erosion of cooperative domains and the eventual convergence to full defection. This final state is in agreement with the prediction of the replicator equation on graphs, which predicts the stable stationary state $\rho = 0$. To characterise the evolutionary dynamics over the whole dilemma space, we performed a systematic numerical scan of the (D_g, D_r) plane. As in the well-mixed case, we considered a 21×21 grid of equally spaced parameter pairs in the square domain $D_g, D_r \in [-1, 1]$. For each parameter pair, we ran 10 independent simulations, all initialised at

$$\rho_0 = 0.5,$$

and estimated the asymptotic cooperation level by averaging $\rho(t)$ over the last 20% of the simulation time, after convergence had been reached. The resulting heatmap is shown in Fig. 3.11. The numerical heatmap should be compared with the analytical prediction of the replicator equation on graphs for $k = 4$, shown in Fig. 3.7. The two figures show a good overall agreement and reproduce the same global organisation of the dilemma space. In particular, the main qualitative structure of the four game classes is correctly captured by the pair approximation. To visualise the deviations more clearly, in Fig. 3.12 we plot the difference

$$\rho_\infty^{\text{sim}} - \rho_\infty^{\text{pred}},$$

where ρ_∞^{sim} is the stationary cooperation level measured in the simulations and $\rho_\infty^{\text{pred}}$ is the corresponding prediction of the graph replicator equation. This figure shows that the agreement is essentially perfect in the Prisoner's Dilemma and Harmony Game regions, and also in most of the Stag Hunt region. A visible exception appears along the secondary diagonal

$$D_g = -D_r$$

within the Stag Hunt region. Along this line, the pair approximation predicts an unstable internal threshold at

$$\rho^* = 0.5.$$

Since the simulations are initialised at $\rho_0 = 0.5$, the system starts exactly on the unstable separatrix. The analytical prediction is therefore marginal, and small perturbations select either absorbing state. In the Monte Carlo simulations, the system more often converges to full cooperation. However, the most significant discrepancies are observed in the Snowdrift region. Here, the sign of $D_r + D_g$ separates two distinct behaviours. When

$$D_r + D_g > 0,$$


Square-lattice snapshots for $D_g = 0.2, D_r = 0.2$


Figure 3.10: Snapshots of the stochastic dynamics on a square lattice with von Neumann neighbourhood ($k = 4$), periodic boundary conditions, no self-interaction, and selection intensity $\beta = 0.2$, for the game $(D_g, D_r) = (0.2, 0.2)$, corresponding to panel (a) of Fig. 3.9. Black sites denote cooperators, whereas white sites denote defectors. The snapshots are shown at the initial configuration (0 sweeps) and after 20, 50, 100, 150, and 200 Monte Carlo sweeps. The initial condition is a random configuration with cooperation level $\rho_0 = 0.5$. The figure illustrates the progressive disappearance of cooperators and the convergence to full defection, in agreement with the analytical prediction of the replicator equation on graphs for the von Neumann neighbourhood. *Simulation by the author.*

one has

$$D_g > -D_r = S.$$

Thus, the temptation to defect against a cooperator, measured by $D_g = T - R$, is relatively strong compared with the payoff obtained by a cooperator facing a defector. Defectors located at the boundaries of cooperative clusters can therefore exploit neighbouring cooperators efficiently. The pair approximation already predicts a reduction of cooperation in this region. However, on the square lattice this effect is amplified by spatial structures that are not fully captured by pair-level correlations. As a result, one observes

$$\rho_\infty^{\text{sim}} < \rho_\infty^{\text{pred}}.$$

By contrast, when

$$D_r + D_g < 0,$$




Figure 3.11: Stationary cooperation level obtained from stochastic simulations on a square lattice with von Neumann neighbourhood ($k = 4$), periodic boundary conditions, and no self-interaction. The heatmap reports the average asymptotic cooperation level for 21×21 equally spaced parameter pairs (D_g, D_r) in the domain $D_g, D_r \in [-1, 1]$. For each parameter pair, the initial condition is $\rho_0 = 0.5$, results are averaged over 10 independent realisations, and the final cooperation level is estimated by averaging $\rho(t)$ over the last 20% of the simulation time after convergence. The selection intensity is $\beta = 0.2$. *Original simulation by the author.*

one has

$$D_g < -D_r = S.$$

Thus, the payoff obtained by a cooperator facing a defector is relatively large compared with the temptation to defect against a cooperator. Cooperative clusters therefore become more robust, and the square lattice amplifies their persistence beyond what is captured by pair approximation. In this case, one finds

$$\rho_\infty^{\text{sim}} > \rho_\infty^{\text{pred}}.$$

For completeness, we also compare the simulation results with the well-mixed analytical predictions. To this end, in Fig. 3.13 we report the heatmap of the difference

$$\rho_\infty^{\text{sim}} - \rho_\infty^{\text{WM}},$$

where ρ_∞^{WM} denotes the stationary cooperation level predicted by the well-mixed replicator equation. As expected, the discrepancies are more pronounced than in the




Figure 3.12: Difference between the stationary cooperation level measured in stochastic simulations and the analytical prediction of the graph replicator equation for the von Neumann neighbourhood: $\rho_{\infty}^{\text{sim}} - \rho_{\infty}^{\text{pred}}$. Positive values indicate parameter regions in which the square-lattice simulations yield more cooperation than predicted by pair approximation, whereas negative values indicate regions in which the simulations yield less cooperation. The agreement is essentially exact in the Prisoner's Dilemma and Harmony Game regions, and also in most of the Stag Hunt region, while the main discrepancies are concentrated in the Snowdrift region. These deviations reveal the effect of higher-order spatial correlations and interface structures that are not captured by pair approximation. *Original illustration by the author.*

comparison with pair approximation. This confirms that the graph replicator equation provides a substantially better analytical description of the lattice dynamics than the well-mixed theory. In particular, pair approximation captures the leading effect of local spatial correlations, whereas the well-mixed description neglects them entirely.

3.3.3 Moore Neighbourhood

We now consider the case of a Moore neighbourhood, corresponding to degree $k = 8$, without self-interaction. The numerical simulations are performed with the same pairwise comparison update rule used in the von Neumann case, with selection intensity $\beta = 0.2$. As before, this value does not correspond to the strict weak-selection limit, but it is sufficiently small for the stochastic dynamics to be compared qualitatively with the analytical prediction obtained from the replicator equation on




Figure 3.13: Difference between the stationary cooperation level measured in the Monte Carlo simulations and the prediction of the well-mixed replicator equation: $\rho_{\infty}^{\text{sim}} - \rho_{\infty}^{\text{WM}}$. Positive values indicate parameter regions in which the square-lattice simulations yield more cooperation than predicted by the well-mixed dynamics, whereas negative values indicate regions in which they yield less cooperation. Compared with Fig. 3.12, the discrepancies are clearly larger, showing that pair approximation provides a more accurate description of the square-lattice dynamics. This reflects the fact that pair approximation captures local pair correlations, while the well-mixed theory neglects spatial structure altogether. *Original illustration by the author.*

graphs. For a Moore neighbourhood, each player interacts with its eight surrounding neighbours: the four horizontal and vertical nearest neighbours, together with the four diagonal neighbours. Therefore, the degree entering the graph replicator equation is $k = 8$. Substituting this value into Eq. (3.7) gives

$$\dot{\rho} = \rho(1 - \rho) \left[\rho(D_r - D_g) - \frac{7D_r + D_g}{6} \right]. \quad (3.10)$$

This equation has the two boundary stationary states

$$\rho = 0, \quad \rho = 1,$$

corresponding respectively to full defection and full cooperation. In addition, an internal stationary point may exist. Setting the term in square brackets in Eq. (3.10) equal to zero gives

$$\rho^* = \frac{7D_r + D_g}{6(D_r - D_g)}.$$

This fixed point is admissible only if

$$0 < \frac{7D_r + D_g}{6(D_r - D_g)} < 1.$$

The stability of the stationary states is obtained from the sign of the Jacobian $J(\rho) = \frac{d\dot{\rho}}{d\rho}$. Evaluating it at the fixed points gives the following conditions:

- $\rho = 0$ is stable if

$$7D_r + D_g > 0.$$

- $\rho = 1$ is stable if

$$D_r + 7D_g < 0.$$

- the internal equilibrium $\rho^* \in (0, 1)$, when admissible, is stable if

$$D_r - D_g < 0.$$

To visualise the analytical prediction across the dilemma space, we constructed a deterministic heatmap of the asymptotic cooperation level predicted by Eq. (3.10). As in the von Neumann case, we scanned the domain $D_g, D_r \in [-1, 1]$ using a grid of 1001×1001 equally spaced parameter pairs. For each pair (D_g, D_r) , we computed the stationary points of Eq. (3.10) and determined their stability from the Jacobian. If a stable internal equilibrium was present, the predicted asymptotic cooperation level was assigned as $\rho_\infty = \rho^*$. If only one boundary equilibrium was stable, the outcome was assigned as $\rho_\infty = 0$ or $\rho_\infty = 1$. In bistable regions, where both boundary equilibria are stable and the internal equilibrium is unstable, the predicted outcome was assigned by comparing the initial condition used in the simulations, $\rho_0 = 0.5$, with the unstable threshold ρ^* . The resulting prediction is shown in Fig. 3.14. As expected, the Moore-neighbourhood prediction is closer to the well-mixed prediction than the von Neumann case. Qualitatively, this is because each player interacts with a larger local sample of the population. Increasing the number of neighbours makes the local interaction environment less sensitive to individual pair correlations and more similar to the average composition of the population. This is also reflected analytically by the correction term in Eq. (3.7), which is proportional to $\frac{1}{k-2}$. Thus, increasing the degree k reduces the strength of the graph-induced correction.

To further verify the approach toward the well-mixed limit as the number of interaction partners increases, we also considered an extended Moore neighbourhood. In this setting, each player interacts not only with the eight nearest surrounding sites, but with all sites contained in the 5×5 square centred on the focal player, excluding the focal site itself. The number of interaction partners is therefore $k = 5^2 - 1 = 24$. Substituting $k = 24$ into Eq. (3.7) gives

$$\dot{\rho} = \rho(1 - \rho) \left[\rho(D_r - D_g) - \frac{23D_r + D_g}{22} \right]. \quad (3.11)$$

The internal fixed point is

$$\rho^* = \frac{23D_r + D_g}{22(D_r - D_g)},$$




Figure 3.14: Analytical prediction of the asymptotic cooperation level in the (D_g, D_r) dilemma space obtained from the replicator equation on graphs for a Moore neighbourhood, corresponding to degree $k = 8$. The prediction is computed from Eq. (3.10) on a 1001×1001 grid in the domain $D_g, D_r \in [-1, 1]$. For each parameter pair, the stationary points are computed and their stability is determined from the Jacobian. The boundary equilibrium $\rho = 0$ is stable when $7D_r + D_g > 0$, while $\rho = 1$ is stable when $D_r + 7D_g < 0$. An internal equilibrium $\rho^* = (7D_r + D_g)/[6(D_r - D_g)]$ is admissible when it lies in $(0, 1)$ and is stable when $D_r - D_g < 0$. In bistable regions, the final analytical outcome is assigned by comparing the simulation initial condition $\rho_0 = 0.5$ with the unstable threshold ρ^* . Compared with the von Neumann case, the prediction is closer to the well-mixed result because the larger neighbourhood reduces the graph-induced correction term. *Original illustration by the author.*

which is admissible if

$$0 < \frac{23D_r + D_g}{22(D_r - D_g)} < 1.$$

The corresponding stability conditions are

- $\rho = 0$ is stable if

$$23D_r + D_g > 0;$$

- $\rho = 1$ is stable if

$$D_r + 23D_g < 0;$$

- the internal equilibrium $\rho^* \in (0, 1)$, when admissible, is stable if

$$D_r - D_g < 0.$$

The analytical heatmap obtained for the extended Moore neighbourhood is shown in Fig. 3.15. The same 1001×1001 grid and the same stability-based assignment procedure described above are used. As expected, the prediction moves even closer




Figure 3.15: Analytical prediction of the asymptotic cooperation level in the (D_g, D_r) dilemma space obtained from the replicator equation on graphs for an extended Moore neighbourhood, corresponding to degree $k = 24$. In this neighbourhood, each player interacts with all sites in the 5×5 square centred on the focal site, excluding the focal player itself. The prediction is computed from Eq. (3.11) on a 1001×1001 grid in the domain $D_g, D_r \in [-1, 1]$. The boundary equilibrium $\rho = 0$ is stable when $23D_r + D_g > 0$, while $\rho = 1$ is stable when $D_r + 23D_g < 0$. An internal equilibrium $\rho^* = (23D_r + D_g)/[22(D_r - D_g)]$ is admissible when it lies in $(0, 1)$ and is stable when $D_r - D_g < 0$. In bistable regions, the final analytical outcome is assigned by comparing the simulation initial condition $\rho_0 = 0.5$ with the unstable threshold ρ^* . The heatmap shows that increasing the neighbourhood size further reduces the graph-induced correction and brings the analytical prediction closer to the well-mixed limit shown in Fig. 1.4. *Original illustration by the author.*

to the well-mixed case (Fig. 1.4). Indeed, for $k = 24$ the graph correction is proportional to $1/(k - 2) = 1/22$, which is much smaller than the correction for the von Neumann neighbourhood, $1/2$, and for the Moore neighbourhood, $1/6$. The stability boundaries

$$23D_r + D_g = 0, \quad D_r + 23D_g = 0$$

are therefore very close to the well-mixed axes $D_r = 0$ and $D_g = 0$.

To complement the analytical characterisation discussed above, we now compare the predictions of the replicator equation on graphs for the Moore neighbourhood with the corresponding Monte Carlo simulations on the square lattice. The simulation

framework is the one introduced in Sect. 3.3.1, and is therefore not repeated here in full. In particular, we consider a square lattice with periodic boundary conditions, Moore neighbourhood ($k = 8$), no self-interaction, and selection intensity $\beta = 0.2$. To characterise the stationary behaviour over the whole dilemma space, we performed a systematic numerical scan using $21 \times 21 = 441$ equally spaced parameter pairs (D_g, D_r) in the square domain $D_g, D_r \in [-1, 1]$. For each pair, the dynamics was initialised at $\rho_0 = 0.5$, and the results were averaged over 10 independent realisations. As in the previous cases, the stationary cooperation level ρ_∞^{sim} was estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. The resulting heatmap is shown in Fig. 3.16. As in the von Neumann case, the com-




Figure 3.16: Stationary cooperation level ρ_∞^{sim} obtained from stochastic simulations on a square lattice with Moore neighbourhood ($k = 8$), periodic boundary conditions, no self-interaction, and selection intensity $\beta = 0.2$. The heatmap reports the average final cooperation level over 21×21 equally spaced parameter pairs (D_g, D_r) in the domain $D_g, D_r \in [-1, 1]$. For each pair, the dynamics is initialised at $\rho_0 = 0.5$, averaged over 10 independent realisations, and the stationary value is estimated by averaging $\rho(t)$ over the final 20% of the simulation time. The figure provides the numerical counterpart of the analytical prediction heatmap obtained from the replicator equation on graphs for the Moore neighbourhood shown in Fig. 3.14. *Original simulation by the author.*

parison with the graph replicator equation reveals non-negligible discrepancies in specific regions of the parameter space. To visualise these discrepancies, in Fig. 3.17 we report the difference

$$\rho_\infty^{\text{sim}} - \rho_\infty^{\text{pred}},$$

where $\rho_\infty^{\text{pred}}$ denotes the equilibrium cooperation level predicted by the replicator equation on graphs for $k = 8$. The deviations are concentrated mainly in part of the Stag Hunt region and in the Snowdrift region. In the Stag Hunt region, a visible




Figure 3.17: Difference between the stationary cooperation level measured in the stochastic simulations and the analytical prediction of the replicator equation on graphs for the Moore neighbourhood, namely $\rho_\infty^{\text{sim}} - \rho_\infty^{\text{pred}}$. The simulation setting is the same as in Fig. 3.16. Positive values indicate that the simulations yield more cooperation than predicted by pair approximation, while negative values indicate less cooperation. The main discrepancies are concentrated in part of the Stag Hunt region, close to the line $D_r = -D_g$, and in the Snowdrift region, where higher-order spatial correlations and interface effects generate systematic deviations from pair-level theory. *Original illustration by the author.*

discrepancy appears close to the line

$$D_r = -D_g.$$

Along this line, the analytical threshold is close to the initial condition $\rho_0 = 0.5$, making the deterministic prediction particularly sensitive. For some parameter pairs just above this diagonal, the stochastic dynamics on the lattice converges to full cooperation, $\rho_\infty^{\text{sim}} = 1$, whereas the pair-approximation prediction gives full defection, $\rho_\infty^{\text{pred}} = 0$. This indicates that finite-size fluctuations and spatial correlations can shift the effective basin of attraction with respect to the deterministic pair-approximation prediction. In the Snowdrift region, the discrepancies have the same qualitative interpretation discussed for the von Neumann neighbourhood. Here, the sign of $D_r + D_g$ separates two distinct regimes. When

$$D_r + D_g > 0,$$

one has

$$D_g > -D_r = S.$$

Thus, the temptation to defect against a cooperator, measured by $D_g = T - R$, is relatively strong compared with the payoff obtained by a cooperator facing a defector. Defectors located at the boundaries of cooperative domains can therefore exploit neighbouring cooperators efficiently. The pair approximation already predicts a reduction of cooperation in this region. However, on the square lattice this effect is amplified by spatial structures that are not fully captured by pair-level correlations. As a result, one observes

$$\rho_{\infty}^{\text{sim}} < \rho_{\infty}^{\text{pred}}.$$

By contrast, when

$$D_r + D_g < 0,$$

one has

$$D_g < -D_r = S.$$

Thus, the payoff obtained by a cooperator facing a defector is relatively large compared with the temptation to defect against a cooperator. Cooperators at the boundaries of cooperative domains are therefore not strongly penalised when interacting with defectors, and cooperative domains become more robust. The pair approximation already predicts an increase of cooperation in this region. However, also in this case, the square lattice amplifies the effect through spatial structures that are not fully captured by pair-level correlations. Consequently, one finds

$$\rho_{\infty}^{\text{sim}} > \rho_{\infty}^{\text{pred}}.$$

Thus, also for the Moore neighbourhood, the pair approximation does not fully reproduce the Monte Carlo dynamics on the square lattice. The discrepancies highlight the role of spatial structures, interfaces and higher-order local correlations that are not captured by a closure based only on pair-level information.

3.3.4 Self-Interaction

In this subsection, we investigate the effect of *self-interaction*, namely the possibility that each player interacts not only with its neighbours, but also once with itself. In addition, within this framework we reproduce the classical results of Szabó and Tóke on the spatial Prisoner's Dilemma with self-interaction [11]. This comparison provides a useful validation of the Monte Carlo simulation framework used throughout this thesis.

We consider a square lattice with periodic boundary conditions and von Neumann neighbourhood, so that each player interacts with its four nearest neighbours and, in addition, once with itself. The strategy-update dynamics is the same as in the previous subsections and follows the Fermi rule with selection intensity β . The effect of self-interaction can be understood directly from the payoff matrix

$$G(D_g, D_r) = \begin{pmatrix} 1 & -D_r \\ 1 + D_g & 0 \end{pmatrix}.$$

If self-interaction is included, a cooperator receives an additional payoff equal to $G_{CC} = 1$, whereas a defector receives an additional payoff equal to $G_{DD} = 0$. Thus, self-interaction systematically favours cooperators over defectors, since it adds a positive contribution to the payoff of cooperators while not providing any analogous benefit to defectors. One therefore expects an overall increase in the equilibrium cooperation level. To quantify this effect, we performed the usual scan of the dilemma space using $21 \times 21 = 441$ equally spaced parameter pairs $(D_g, D_r) \in [-1, 1] \times [-1, 1]$. For each pair, the dynamics was initialised at $\rho_0 = 0.5$, simulated for 10 independent realisations, and the stationary cooperation level was estimated by averaging $\rho(t)$ over the final 20% of the simulation window after convergence. The resulting heatmap is shown in Fig. 3.18. As expected, self-interaction leads to an overall in-




Figure 3.18: Stationary cooperation level ρ_∞ obtained from Monte Carlo simulations on a square lattice with periodic boundary conditions, von Neumann neighbourhood ($k = 4$), and self-interaction. The heatmap reports the average final cooperation level for 21×21 equally spaced parameter pairs (D_g, D_r) in the domain $D_g, D_r \in [-1, 1]$. For each pair, the dynamics is initialised at $\rho_0 = 0.5$, evolved according to the Fermi update rule with selection intensity $\beta = 0.2$, and averaged over 10 independent realisations. The stationary value is estimated by averaging $\rho(t)$ over the final 20% of the simulation time. Compared with the corresponding case without self-interaction, cooperation is enhanced across a broad portion of the dilemma space. *Original simulation by the author.*

crease in the observed cooperation level across large portions of the dilemma space. This confirms that allowing players to interact with themselves makes cooperative strategies more resilient, precisely because cooperators gain an additional positive self-contribution whereas defectors do not.

At this point, we reproduce the classical results of Szabó and Tóke [11], who reported

an interesting phase-transition-like behaviour in the spatial Prisoner's Dilemma on a square lattice. This comparison is useful both as a validation of our numerical implementation and as a reference point for interpreting the role of self-interaction. We consider the same setting as in their original work, namely a square lattice with von Neumann neighbourhood ($k = 4$), self-interaction, and Fermi updating. Following [11], we restrict attention to the *weak Prisoner's Dilemma* by fixing $D_r = 0$, and varying $D_g \in [0, 1]$. In our parametrisation, the payoff matrix becomes

$$G(D_g, 0) = \begin{pmatrix} 1 & 0 \\ 1 + D_g & 0 \end{pmatrix},$$

which is equivalent to the parametrisation used by Szabó and Tóke, where the temptation parameter is denoted by b and satisfies

$$b = 1 + D_g.$$

For each value of D_g , we performed stochastic simulations starting from the initial condition $\rho_0 = 0.5$, and measured the asymptotic cooperation level ρ by averaging over the final 20% of the simulation time. The results are shown in Fig. 3.19. We recover the same qualitative behaviour reported in [11]. As D_g increases, the




Figure 3.19: Reproduction of the results of Szabó and Tóke [11] for the spatial Prisoner's Dilemma on a square lattice with von Neumann neighbourhood ($k = 4$) and self-interaction. The asymptotic cooperation level ρ is shown as a function of D_g , with $D_r = 0$, corresponding in our parametrisation to the weak Prisoner's Dilemma $G(D_g, 0) = \begin{pmatrix} 1 & 0 \\ 1 + D_g & 0 \end{pmatrix}$. Monte Carlo simulations are performed using the Fermi update rule with selection intensity $\beta = 10$, starting from $\rho_0 = 0.5$. For each value of D_g , the final cooperation level is estimated by averaging $\rho(t)$ over the final 20% of the simulation time. The figure reproduces the phase-transition-like behaviour reported in the original work, with a transition from full cooperation to coexistence and then to full defection as D_g increases. *Simulation by the author.*

stationary cooperation level decreases and the system undergoes a phase-transition-like change between three regimes: full cooperation, coexistence of cooperators and defectors, and full defection. In particular, below the first transition point the

population converges to full cooperation, while above a second threshold cooperators disappear and the system reaches full defection. Between these two thresholds, cooperators and defectors coexist at a nontrivial stationary level. This agreement provides an important consistency check for our implementation.

3.4 Evolutionary Dynamics on Complex Networks

In this section, we study evolutionary dynamics on complex networks. In particular, we consider three network models: Erdős–Rényi random graphs, Watts–Strogatz small-world networks and Barabási–Albert scale-free networks. These models allow us to investigate how different structural properties, such as randomness, clustering, short path lengths and degree heterogeneity, affect the evolution of cooperation. For each network class, we first succinctly discuss their main structural properties. We then perform Monte Carlo simulations on the corresponding networks and whenever appropriate, the numerical results are compared with the analytical prediction of the replicator equation on regular graphs derived in Sect. 3.2.3. This comparison allows us to assess when pair approximation provides an accurate description of the dynamics and when network-specific structures lead to deviations from the analytical benchmark.

3.4.1 Erdős–Rényi Networks

We begin with Erdős–Rényi random networks, which provide the simplest model of a random interaction structure. In this model, links are placed independently between pairs of nodes with a fixed probability. As a result, the network has no imposed spatial organisation or preferential pattern of connections: before the graph is realised, all pairs of nodes are statistically equivalent.

Definition 3.13. An **Erdős–Rényi random graph** $G(N, p)$ is a graph with N nodes in which each possible undirected edge between two distinct nodes is present independently with probability p . Therefore, for every pair of nodes $i, j \in V$, with $i \neq j$,

$$\mathbb{P}((i, j) \in E) = p.$$

The parameter p controls the density of the network. Since an undirected graph with N nodes has at most $N(N - 1)/2$ possible edges, the expected number of links is

$$\langle M \rangle = p \frac{N(N - 1)}{2}.$$

The corresponding expected average degree is

$$\langle k \rangle = p(N - 1).$$

Equivalently, if one wants to generate an Erdős–Rényi network with a given expected average degree $\langle k \rangle$, one can choose

$$p = \frac{\langle k \rangle}{N - 1}.$$

Different realisations generated with the same parameters N and p are not identical: both the edge set and the total number of links fluctuate from one realisation to another. This is illustrated in Fig. 3.20. The degree distribution of an Erdős–Rényi




Figure 3.20: Three realisations of Erdős–Rényi random graphs generated with the same parameters, $N = 12$ and $p = 1/6$. Each possible edge is included independently with probability p , so different realisations generally have different edge sets and may also have different numbers of links. This illustrates the intrinsically probabilistic nature of the Erdős–Rényi network construction. *Illustration by the author.*

graph is binomial:

$$P(k) = \binom{N-1}{k} p^k (1-p)^{N-1-k}.$$

In the sparse large-network limit, where N is large and $\langle k \rangle$ remains finite, this distribution is well approximated by a Poisson distribution:

$$P(k) \simeq e^{-\langle k \rangle} \frac{\langle k \rangle^k}{k!}.$$

Thus, Erdős–Rényi networks are relatively homogeneous: most nodes have degree close to the average degree, and very high-degree nodes are exponentially rare. Figure 3.21 shows an example of this behaviour for a large sparse Erdős–Rényi network. The degree distribution is concentrated around the average degree and follows closely the Poisson distribution.

The clustering coefficient of an Erdős–Rényi graph is low. Indeed, since edges are placed independently with probability p , the probability that two neighbours of a node are connected is approximately

$$C_{\text{ER}} \simeq p = \frac{\langle k \rangle}{N-1}.$$

For sparse networks, this quantity tends to zero as N increases. Therefore, Erdős–Rényi graphs do not display strong local clustering. At the same time, typical distances are short. It turns out that the average path length scales approximately as

$$\ell_{\text{ER}} \sim \frac{\log N}{\log \langle k \rangle}.$$

We now present the results obtained for evolutionary dynamics on Erdős–Rényi networks. Simulations are performed on large populations, with $N = 40000$, so that




Figure 3.21: Degree distribution of an Erdős–Rényi graph with $N = 1000$ nodes and expected average degree $\langle k \rangle = 8$. The bars show the empirical degree distribution of one network realisation, while the red curve shows the Poisson prediction $P(k) = e^{-\langle k \rangle} \langle k \rangle^k / k!$. The agreement illustrates the Poisson-like degree distribution of Erdős–Rényi networks in the sparse large-network limit. *Illustration by the author.*

finite-size effects are reduced and the dynamics can be compared meaningfully with the analytical predictions discussed above. In order to compare these results with those obtained on the square lattice, we choose the link probability p so that the expected average degree of the Erdős–Rényi network is

$$\langle k \rangle = p(N - 1) = 4.$$

Thus, the typical number of interaction partners is the same as in the von Neumann neighbourhood considered on the square lattice. As analytical benchmark, we use the replicator equation on regular graphs for degree $k = 4$, given in Eq. (3.7), together with the corresponding equilibrium conditions derived for this case. For each pair (D_g, D_r) in the dilemma space, we perform 10 independent simulations and average the results. Strategy updating follows the Fermi rule, with selection intensity fixed to $\beta = 0.2$. All simulations are initialised from $\rho_0 = 0.5$, and are run until the cooperation level reaches a stationary regime. The final cooperation level is then estimated by averaging $\rho(t)$ over the last 20% of the simulation time window. A further modelling choice concerns how payoffs are computed on the network. In principle, one may either use accumulated payoff or normalise the payoff by the degree of the node. We tested both possibilities and found that the results differ only slightly. Since the analytical treatment developed above is naturally formulated in terms of average payoff, in the following we report only the results obtained using degree-normalised payoffs.

To characterise the evolutionary outcomes across the dilemma space, we perform a systematic numerical scan over a 21×21 grid of parameter pairs $(D_g, D_r) \in [-1, 1] \times [-1, 1]$. The resulting heatmap is shown in Fig. 3.22

To assess the quality of the analytical approximation, we then compare the simula-




Figure 3.22: Final cooperation level on Erdős–Rényi networks across the dilemma space. The heatmap reports the average stationary cooperation level obtained from Monte Carlo simulations on Erdős–Rényi networks with $N = 40000$ nodes and expected average degree $\langle k \rangle = 4$. The link probability is chosen accordingly as $p = \langle k \rangle / (N - 1)$. Simulations are performed using the Fermi update rule with selection intensity $\beta = 0.2$, starting from the initial condition $\rho_0 = 0.5$. For each pair (D_g, D_r) on a 21×21 grid in the square domain $[-1, 1] \times [-1, 1]$, the result is averaged over 10 independent realisations. The final cooperation level is computed by averaging $\rho(t)$ over the last 20% of the simulation time window after convergence. Payoffs are normalised by degree. *Original simulation by the author.*

tion results with the prediction of the replicator equation on graphs. This is done by plotting the heatmap of the difference between the simulated stationary cooperation level and the analytical prediction,

$$\rho_{\infty}^{\text{sim}} - \rho_{\infty}^{\text{graph}},$$

shown in Fig. 3.23. As can be seen from the figure, the agreement between the analytical prediction and the stochastic dynamics on Erdős–Rényi networks is excellent. The only visible discrepancies are concentrated in the Stag Hunt region, along the usual line $D_r = -D_g$. This is precisely the boundary at which the bistable structure becomes most sensitive to fluctuations and to the initial condition. In this region, the analytical prediction is close to the threshold separating the basins of attraction of full cooperation and full defection, so small stochastic fluctuations or small variations in the network realisation may push the system toward one absorbing state or the other. Outside this narrow boundary region, pair approximation captures the




Figure 3.23: Difference between Monte Carlo simulation results on Erdős–Rényi networks and the analytical prediction of the replicator equation on graphs. The heatmap shows $\rho_{\infty}^{\text{sim}} - \rho_{\infty}^{\text{graph}}$, where $\rho_{\infty}^{\text{sim}}$ is the stationary cooperation level measured in the simulations and $\rho_{\infty}^{\text{graph}}$ is the corresponding prediction of the replicator equation on graphs for degree $k = 4$, as given by Eq. (3.7). The simulation setting is the same as in Fig. 3.22: Erdős–Rényi networks with $N = 40000$ and expected average degree $\langle k \rangle = 4$, Fermi updating with $\beta = 0.2$, initial condition $\rho_0 = 0.5$, payoff normalisation by degree, and averaging over 10 independent realisations. The final cooperation level is computed over the last 20% of the simulation time window. *Original illustration by the author.*

dynamics very accurately. This result is consistent with the structural properties of Erdős–Rényi networks. Indeed, these networks are sparse and locally tree-like, with low clustering and relatively weak higher-order correlations. As a consequence, pair-level information is already sufficient to describe the relevant local structure, and there is no strong need to include higher-order clusters beyond pairs. In this sense, the excellent agreement observed here indicates that, for Erdős–Rényi networks, the pair approximation provides an accurate analytical description of the evolutionary dynamics.

3.4.2 Watts–Strogatz Networks

The second class of complex networks considered in this thesis is the Watts–Strogatz small-world network. This model was introduced to capture a structural property observed in many real systems: networks can be highly clustered at the local level

while still having short typical distances between nodes [67]. Examples include social acquaintance networks, collaboration networks, neural networks and technological infrastructures. In such systems, individuals or units often interact within local groups, but a few long-range connections can strongly reduce the distance between otherwise distant parts of the network. The central idea of the Watts–Strogatz model is to interpolate between a regular lattice and a random network.

Definition 3.14. A **Watts–Strogatz network** is constructed from a regular ring lattice with N nodes and degree k , where each node is initially connected to its k nearest neighbours. Each edge is then rewired with probability p_{rew} by keeping one endpoint fixed and connecting the other endpoint to a randomly chosen node, avoiding self-loops and multiple edges.

The parameter p_{rew} controls the amount of randomness in the network. When $p_{\text{rew}} = 0$, the network is a regular ring lattice. When $p_{\text{rew}} = 1$, all eligible links have been rewired, and the network becomes much closer to a random graph. The most interesting regime is the intermediate one, where

$$0 < p_{\text{rew}} \ll 1.$$

In this regime, short average path length and high clustering coexist. This is the small-world regime. Figure 3.24 illustrates the construction of the model for three different values of the rewiring probability. The small-world effect can be quantified by studying the average path length and the clustering coefficient as functions of the rewiring probability. For a regular lattice, the clustering coefficient is high, but the average path length is large because distant nodes can be reached only through many local steps. When a small fraction of edges is rewired, the average path length drops rapidly, while the clustering coefficient remains close to its regular-lattice value. This behaviour is shown in Fig. 3.25.

The degree distribution of Watts–Strogatz networks is narrow. Since the model starts from a regular graph, most nodes keep a degree close to the initial value k . Rewiring introduces fluctuations in the degrees, but it does not generate the broad degree heterogeneity or hubs typical of scale-free networks. An example of the degree distribution for the network size and rewiring probability used in the simulations is shown in Fig. 3.26.

We now turn to the results obtained from the Monte Carlo simulations on Watts–Strogatz networks. In this case, we considered populations of size $N = 40000$ and generated the interaction network using the Watts–Strogatz model with rewiring probability $p = 0.1$, starting from a regular ring lattice with degree $k = 4$. This choice ensures an average degree equal to that used in the previous analyses, allowing a direct comparison with the square-lattice and Erdős–Rényi cases. The evolutionary dynamics is the same as described above. In particular, we used the Fermi update rule with selection intensity $\beta = 0.2$, performed 10 independent realisations for each parameter pair, and measured the final cooperation level by averaging $\rho(t)$ over the last 20% of the simulation time after convergence. As in the Erdős–Rényi case, we considered the average payoff, since using total payoff leads to qualitatively identical results, while the analytical derivation of the replicator equation on graphs is naturally closer to an average-payoff description.




Figure 3.24: Construction of Watts–Strogatz networks. Panel (a) shows the regular ring lattice obtained for $p_{\text{rew}} = 0$, where each node is connected to its nearest and next-nearest neighbours. Panel (b) shows the small-world regime, here represented for $p_{\text{rew}} = 0.1$: most local links are preserved, while a small number of rewired long-range connections creates shortcuts across the network. Panel (c) shows the randomised regime obtained for $p_{\text{rew}} = 1$, where all eligible links have been rewired. The figure illustrates how the Watts–Strogatz model interpolates between an ordered lattice and a random network. *Adapted from [26].*

To characterise the behaviour over the whole dilemma space, we performed the usual scan over a 21×21 grid of parameter pairs $(D_g, D_r) \in [-1, 1]^2$, starting from the initial condition $\rho_0 = 0.5$. The resulting heatmap of the asymptotic cooperation level is shown in Fig. 3.27. As an analytical benchmark, we use the replicator equation on regular graphs, Eq. (3.7), with degree $k = 4$. To assess the agreement between simulations and analytical predictions, we also report in Fig. 3.28 the heatmap of the difference

$$\rho_{\infty}^{\text{sim}} - \rho_{\infty}^{\text{graph}},$$

where $\rho_{\infty}^{\text{graph}}$ denotes the equilibrium cooperation level predicted by the replicator equation on graphs. As can be seen from the figure, the main discrepancies are again concentrated in the Snowdrift region. This behaviour is qualitatively similar to what was observed on the square lattice. Compared with the Erdős–Rényi case, where the agreement with pair approximation was excellent, visible deviations now reappear. This suggests that the breakdown of pair approximation is mainly associated with local loops, clustering and higher-order neighbourhood correlations. Pair approximation captures pair-level correlations, but it cannot fully describe the additional local structure present in clustered networks.




Figure 3.25: Small-world behaviour in the Watts–Strogatz model. The figure shows the average path length $\ell(p_{\text{rew}})$ and the clustering coefficient $C(p_{\text{rew}})$ as functions of the rewiring probability, normalised by their values at $p_{\text{rew}} = 0$. For small values of p_{rew} , the average path length decreases rapidly, while the clustering coefficient remains relatively high. This identifies the small-world regime, where short typical distances coexist with strong local clustering. *Illustration by the author.*




Figure 3.26: Degree distribution of a Watts–Strogatz network with $N = 40000$ nodes, initial degree $k = 4$, and rewiring probability $p_{\text{rew}} = 0.1$. The bars show the sampled degree distribution, while the red dashed line indicates the average degree. The distribution remains concentrated around the initial degree, showing that the network is relatively homogeneous and does not display the broad degree heterogeneity characteristic of scale-free networks. *Illustration by the author.*

3.4.3 Barabási–Albert Networks

The third class of complex networks considered is the Barabási–Albert network. This model was introduced to describe networks whose degree distribution is highly heterogeneous and characterised by the presence of hubs, namely nodes with a very large




Figure 3.27: Asymptotic cooperation level $\rho_{\infty}^{\text{sim}}$ obtained from Monte Carlo simulations on Watts–Strogatz networks across the dilemma space. The network consists of $N = 40000$ nodes and is generated from a regular ring lattice with degree $k = 4$, using rewiring probability $p = 0.1$. For each parameter pair (D_g, D_r) in a 21×21 grid over the domain $[-1, 1] \times [-1, 1]$, simulations were initialised at $\rho_0 = 0.5$ and repeated over 10 independent realisations. The evolutionary dynamics follows the Fermi update rule with selection intensity $\beta = 0.2$, and payoffs are normalised by degree. The final cooperation level is estimated by averaging $\rho(t)$ over the last 20% of the simulation time after convergence. *Original simulation by the author.*

number of connections [68]. Examples of systems in which hub-like structures are important include the World Wide Web, citation networks, collaboration networks, Internet topologies and biological interaction networks [26]. The Barabási–Albert model is based on two mechanisms: growth and preferential attachment. Growth means that the network is not generated all at once, but expands progressively as new nodes are added. Preferential attachment means that new nodes are more likely to connect to nodes that already have many links. In this way, highly connected nodes tend to become even more connected over time.

Definition 3.15. A **Barabási–Albert network** is constructed by starting from a small connected seed graph with m_0 nodes. At each time step, a new node is added to the network and connected to $m \leq m_0$ existing nodes. The probability that the new node connects to an existing node i is proportional to the degree k_i of that node:

$$\Pi_i = \frac{k_i}{\sum_j k_j}.$$




Figure 3.28: Heatmap of the difference between the asymptotic cooperation level obtained in Monte Carlo simulations and the analytical prediction of the replicator equation on graphs, $\rho_\infty^{\text{sim}} - \rho_\infty^{\text{graph}}$, for Watts–Strogatz networks. The analytical benchmark is given by Eq. (3.7) with degree $k = 4$. Simulation parameters are the same as in Fig. 3.27: $N = 40000$, rewiring probability $p = 0.1$, initial condition $\rho_0 = 0.5$, 10 independent realisations per parameter pair, Fermi update rule with selection intensity $\beta = 0.2$, and payoffs normalised by degree. The figure shows that the most visible discrepancies are concentrated in the Snowdrift region, indicating that on Watts–Strogatz networks pair approximation does not fully capture the higher-order local correlations induced by clustering and short loops. *Original illustration by the author.*

The parameter m controls the number of links introduced by each new node. Since every new node creates m edges, the average degree of a large Barabási–Albert network is approximately

$$\langle k \rangle \simeq 2m.$$

Thus, choosing $m = 2$ gives an average degree close to 4, making this network comparable with the Erdős–Rényi and Watts–Strogatz cases considered above. The defining structural property of the Barabási–Albert model is its power-law degree distribution. In the large-network limit, it turns out that (see [26])

$$P(k) \sim k^{-3}.$$

This implies that most nodes have small degree, while a small number of nodes become highly connected hubs. This is in sharp contrast with Erdős–Rényi or

Watts–Strogatz networks, where the degree distribution is narrowly concentrated around the average degree. The term *scale-free* refers precisely to this power-law behaviour. A power-law distribution does not have a characteristic degree scale: under a rescaling of the degree variable, its functional form remains unchanged up to a multiplicative factor. Indeed, if

$$P(k) \sim k^{-\gamma},$$

then

$$P(ak) \sim a^{-\gamma} P(k).$$

For this reason, there is no typical degree that fully characterises the network. Instead, connectivity is distributed across many scales, from many low-degree nodes to a few highly connected hubs. Fig. 3.29 compares an Erdős–Rényi network with a Barabási–Albert network with similar average degree. While the Erdős–Rényi network appears relatively homogeneous, the Barabási–Albert network displays a more heterogeneous structure, with some nodes acting as hubs. The difference between

Homogeneous and Heterogeneous Random Networks




Figure 3.29: Visual comparison between an Erdős–Rényi network and a Barabási–Albert network with the same number of nodes and comparable average degree. Both networks have $N = 40$ nodes and expected average degree close to $\langle k \rangle \simeq 4$. Panel (a) shows an Erdős–Rényi network generated with connection probability $p = \langle k \rangle / (N - 1)$, where links are placed independently between pairs of nodes. Panel (b) shows a Barabási–Albert network generated by preferential attachment with $m = 2$. Node size is proportional to degree. The Erdős–Rényi network appears relatively homogeneous, whereas the Barabási–Albert network already displays stronger degree heterogeneity, with a few nodes acquiring more connections than the rest of the network. *Illustration by the author.*

the two models becomes even clearer when one compares their degree distributions. Figure 3.30 shows the degree distribution of a Barabási–Albert network and of an Erdős–Rényi network generated with the same parameters used in the Monte Carlo simulations discussed below. Both networks have size $N = 40000$ and average degree approximately equal to $\langle k \rangle \simeq 4$. The Erdős–Rényi distribution is concentrated




Figure 3.30: Degree distributions of a Barabási–Albert network and an Erdős–Rényi network with $N = 40000$ nodes and comparable average degree $\langle k \rangle \simeq 4$. The Erdős–Rényi network is generated with connection probability $p = \langle k \rangle / (N - 1)$, while the Barabási–Albert network is generated by preferential attachment with $m = 2$. The plot is shown on logarithmic axes. The Erdős–Rényi degree distribution is concentrated around the average degree, reflecting the relative homogeneity of random graphs. By contrast, the Barabási–Albert distribution has a broad tail, indicating the presence of highly connected hubs. The dashed reference line shows the scaling $P(k) \sim k^{-3}$ expected for the Barabási–Albert model in the large-network limit. *Illustration by the author.*

around $\langle k \rangle$, while the Barabási–Albert distribution has a heavy tail that is compatible with a power-law decay. Barabási–Albert networks also have short typical distances, because hubs act as central connectors between different parts of the graph. At the same time, the original model has relatively low clustering. In particular, the average clustering coefficient is small and decreases with the network size, approaching zero in the large- N limit. Their main structural feature is instead degree heterogeneity, namely the coexistence of many low-degree nodes with a few highly connected hubs. In the context of evolutionary dynamics, this heterogeneity is especially important. In homogeneous networks, most individuals have a similar number of interaction partners. In Barabási–Albert networks, instead, hubs can interact with many more individuals than typical nodes. Therefore, the strategy adopted by a hub may have a disproportionately large influence on the spreading and persistence of cooperation. A cooperative hub can support cooperation in its neighbourhood, while a defective hub can exploit many cooperators. For this reason, Barabási–Albert networks provide a useful setting to study how heterogeneous connectivity affects evolutionary games.

We now present the results obtained on Barabási–Albert networks. In the simulations, we considered a population of size $N = 40000$. The networks are generated

using the mechanism described above with $m = 2$, so that the average degree is approximately

$$\langle k \rangle \simeq 2m = 4.$$

This choice allows a direct comparison with the Erdős–Rényi, Watts–Strogatz and von Neumann lattice cases considered above. In Barabási–Albert networks, however, the choice between accumulated and normalised payoff is much more important than in homogeneous networks. Indeed, Erdős–Rényi and Watts–Strogatz networks with narrow degree distributions are relatively homogeneous, so normalising or not normalising the payoff by degree produces only limited differences. By contrast, Barabási–Albert networks are strongly heterogeneous and contain hubs. If payoffs are accumulated, high-degree nodes collect contributions from many more interactions than low-degree nodes, and therefore the evolutionary dynamics becomes strongly degree-dependent. For this reason, the results obtained with average payoff and total payoff must be analysed separately.

We first consider the average-payoff convention, in analogy with the previous network cases. In this setting, the payoff accumulated by each player is divided by its degree. This removes the direct advantage due only to having many neighbours and allows us to compare the simulations with the analytical benchmark given by the replicator equation on graphs. Strictly speaking, the replicator equation on graphs derived in Sect. 3.2.3 applies to regular graphs, where all nodes have the same degree. A Barabási–Albert network is not regular. Nevertheless, when payoffs are normalised by degree, the explicit effect of degree heterogeneity on payoff accumulation is strongly reduced. For this reason, we use Eq. (3.7) with $k = 4$ as an analytical benchmark. We scan the dilemma space using a 21×21 grid of parameter pairs $(D_g, D_r) \in [-1, 1] \times [-1, 1]$. For each pair, simulations are initialised at $\rho_0 = 0.5$, and the dynamics follows the Fermi update rule with selection intensity $\beta = 0.2$. For each parameter pair, we perform 10 independent realisations and estimate the final cooperation level by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. The resulting heatmap is shown in Fig. 3.31. To compare simulations with the analytical benchmark, in Fig. 3.32 we also report the difference $\rho_\infty^{\text{sim}} - \rho_\infty^{\text{graph}}$, where $\rho_\infty^{\text{graph}}$ is the prediction of the graph replicator equation with $k = 4$. The agreement is excellent and is comparable to the Erdős–Rényi case. This suggests that, once payoffs are normalised by degree, the heterogeneity of the Barabási–Albert network does not substantially reduce the predictive power of the pair-approximation benchmark.

We now consider the total-payoff convention in the setting introduced by Santos and Pacheco, where scale-free networks were shown to promote cooperation in evolutionary games [69, 70, 71]. Reproducing this behaviour is also useful as a validation of the Monte Carlo simulation framework used in this thesis. In the Santos–Pacheco setting, all directly connected pairs play one round of the game, and payoffs are accumulated over all neighbours. Strategies are then updated synchronously. When player x updates, one of its neighbours y is chosen at random. If $P_y > P_x$, the strategy of y replaces that of x with probability

$$p_{x \leftarrow y} = \frac{P_y - P_x}{k_{>\Delta}}, \quad (3.12)$$




Figure 3.31: Asymptotic cooperation level $\rho_{\infty}^{\text{sim}}$ obtained from stochastic simulations on Barabási–Albert networks with average payoff. The networks have $N = 40000$ nodes and are generated with attachment parameter $m = 2$, giving average degree $\langle k \rangle \simeq 4$. The heatmap reports the final cooperation level over a 21×21 grid of parameter pairs (D_g, D_r) in the domain $[-1, 1] \times [-1, 1]$. For each pair, simulations are initialised at $\rho_0 = 0.5$, evolved according to the Fermi update rule with selection intensity $\beta = 0.2$, and averaged over 10 independent realisations. Payoffs are normalised by degree, and the final cooperation level is computed by averaging $\rho(t)$ over the last 20% of the simulation time after convergence. *Original simulation by the author.*

where

$$k_{>} = \max(k_x, k_y),$$

and

$$\Delta = \max(T, 1) - \min(S, 0)$$

normalises the transition probability. In our parametrisation,

$$T = 1 + D_g, \quad S = -D_r.$$

We repeated the dilemma-space scan on Barabási–Albert networks using total payoff. The resulting heatmap is shown in Fig. 3.33. Compared with the average-payoff case, the total-payoff dynamics produces a marked increase in cooperation across the dilemma space. This reproduces the cooperation-enhancing effect of scale-free networks reported by Santos and Pacheco. The effect can be understood as a consequence of degree heterogeneity. When payoffs are accumulated, hubs interact with




Figure 3.32: Difference between the asymptotic cooperation level measured in stochastic simulations on Barabási–Albert networks with average payoff and the analytical prediction of the graph replicator equation: $\rho_{\infty}^{\text{sim}} - \rho_{\infty}^{\text{graph}}$. The analytical benchmark is Eq. (3.7) with degree $k = 4$. The simulation setting is the same as in Fig. 3.31: $N = 40000$, $m = 2$, $\rho_0 = 0.5$, Fermi updating with $\beta = 0.2$, average payoff and 10 independent realisations per parameter pair. The small magnitude of the discrepancy shows that, once payoffs are normalised by degree, the pair-approximation benchmark remains highly predictive despite the heterogeneous degree distribution of the Barabási–Albert network. *Original illustration by the author.*

many more neighbours than typical nodes and can therefore acquire a strong evolutionary advantage. If a hub is occupied by a cooperator and is surrounded by sufficiently many cooperators, it accumulates a large payoff through repeated cooperative interactions. Since many neighbouring nodes compare themselves with the hub, the cooperative strategy can then spread efficiently through the network. This creates a positive feedback: cooperation around the hub increases the hub’s payoff, and the hub in turn promotes the spread of cooperation. By contrast, a defective hub may initially gain a high payoff by exploiting neighbouring cooperators, but this advantage is self-limiting. If defection spreads to its neighbours, the hub becomes surrounded by defectors and its payoff decreases toward the mutual-defection payoff. For this reason, the dynamics tends to stabilise cooperation on hubs, which then enhances the spread of cooperation across the network.

We also verified that replacing the synchronous Santos–Pacheco update with the asynchronous Fermi rule used throughout this thesis leaves the results essentially unchanged.

This chapter established the theoretical and numerical baseline for studying evolu-

Simulated Equilibrium Cooperation (BA with total payoff)




Figure 3.33: Asymptotic cooperation level ρ_∞^{sim} obtained from Monte Carlo simulations on Barabási–Albert networks with total payoff, following the Santos–Pacheco setting [69, 70, 71]. The networks have $N = 40000$ nodes and are generated with attachment parameter $m = 2$, giving average degree $\langle k \rangle \simeq 4$. The heatmap reports the final cooperation level over a 21×21 grid of parameter pairs (D_g, D_r) in the domain $[-1, 1] \times [-1, 1]$. For each pair, simulations are initialised at $\rho_0 = 0.5$ and averaged over 10 independent realisations. Payoffs are accumulated over all neighbours without degree normalisation, and strategies are updated according to the Santos–Pacheco rule given by Eq. (3.12). The figure shows a strong enhancement of cooperation compared with the average-payoff case. *Simulation by the author.*

tionary games on structured populations. We first introduced the main definitions of network theory and discussed how complex networks can be used to represent interaction structures. We then presented the analytical tools used to study evolutionary dynamics on graphs. In particular, we showed how pair approximation provides the first correction beyond mean-field theory by retaining information on local pair correlations. This led to the replicator equation on regular graphs, which was used throughout the chapter as the main analytical benchmark.

The first structured interaction setting considered was the square lattice. For the von Neumann neighbourhood, corresponding to degree $k = 4$, the graph replicator equation provided a good description of the Monte Carlo dynamics in most of the dilemma space. The agreement was especially good in the Prisoner’s Dilemma, Harmony Game and most of the Stag Hunt regions. The main discrepancies appeared in the Snowdrift region, where the square lattice produced systematic deviations from the pair-approximation prediction. These deviations showed that pair-level information is not always sufficient to describe spatial evolutionary dynamics: higher-order correlations, interface effects and the geometry of cooperative domains can significantly affect the final cooperation level. The study of self-interaction showed that

allowing each player to interact with itself favours cooperation, since cooperators receive an additional positive payoff contribution. In this setting, the simulations also reproduced the classical behaviour of the spatial Prisoner's Dilemma reported by Szabó and Tóke, providing a useful validation of the numerical implementation.

We then moved from square lattices to complex networks. Erdős–Rényi networks provided the simplest random interaction structure. Since these networks are sparse, weakly clustered and locally close to tree-like graphs, the agreement between Monte Carlo simulations and the pair-approximation benchmark was excellent. The only relevant deviations appeared near the bistable boundary in the Stag Hunt region, where the dynamics is particularly sensitive to fluctuations and to the initial condition.

Watts–Strogatz networks displayed a different behaviour. Although their degree distribution is still relatively homogeneous, their local clustering and short loops led to discrepancies similar to those observed on square lattices. In particular, the largest deviations appeared again in the Snowdrift region, highlighting the fact that clustering and higher-order local correlations can reduce the accuracy of pair approximation.

Finally, Barabási–Albert networks allowed us to analyse the role of degree heterogeneity. With degree-normalised payoffs, the agreement with the pair-approximation benchmark remained very good, suggesting that the explicit effect of heterogeneity vanishes when payoffs are normalised by degree. By contrast, with accumulated payoffs, the dynamics changed substantially: hubs acquired a disproportionate payoff advantage and cooperation was strongly enhanced, reproducing the cooperation-promoting effect of scale-free networks reported in the literature.

Overall, the results of this chapter provide the reference behaviour against which the endogenous-feedback model will be compared. The analysis showed when pair approximation gives an accurate prediction and when it breaks down because of spatial correlations, clustering, interface effects or degree heterogeneity. This baseline is essential for the following chapters. Starting from the next chapter, the payoff matrix will no longer be fixed, but will change endogenously with the cooperation level. The key question will then be whether chimera-game regimes survive when the same structured populations considered here are coupled to endogenous feedback, and how spatial and network structure reshape the resulting cooperative dynamics.

Chapter 4

The Endogenous Feedback Model on Square Lattices

In this chapter, we study the endogenous-feedback model introduced in Chapter 2 on a square lattice, thereby embedding the population in a spatial structure in which each individual interacts only with a finite set of local neighbours. In the well-mixed setting analysed previously, endogenous feedback was shown to generate so-called chimera games. These are regimes in which the dynamical outcome of the feedback model does not correspond to the standard fixed-game behaviour associated with the equilibrium game reached by the system. The first aim here is to investigate whether analogous chimera-game regimes also emerge on square lattices. At the same time, we ask whether spatial structure changes their location, stability or qualitative nature relative to the well-mixed case. The square lattice also makes it possible to study aspects of the dynamics that cannot be observed in a well-mixed population. In particular, the spatial arrangement of cooperators and defectors allows us to analyse the formation, motion and reorganisation of cooperative clusters. For this reason, the second part of the chapter focuses on the spatial dynamics of cooperators on the lattice. We show that the system can display a separation between the fast convergence of the global cooperation level and the much slower spatial reorganisation of cooperative clusters. We then quantify the structure of these clusters and discuss the sharp change in their morphology as the game reached by the endogenous-feedback dynamics varies.

4.1 Chimera Games on Square Lattices

In this section, we study the endogenous-feedback model on square lattices. We verify whether the chimera-game regimes identified in the well-mixed setting persist under local, spatially structured interactions. This is done by comparing the endogenous-feedback results with the corresponding fixed-game dynamics on the same spatial structure. In addition, the outcomes of the Monte Carlo simulations are compared with the analytical predictions provided by the graph-replicator equation.

As described in Sect. 2.2.1, in the endogenous-feedback setting the payoff matrix is

not fixed, but evolves with the state of the population. Given two reference games G_1 and G_2 , the game played at time t is determined by the current cooperation level $\rho(t)$:

$$G(\rho(t)) = (1 - \rho(t))G_1 + \rho(t)G_2. \quad (4.1)$$

Equivalently, omitting the explicit time dependence, the dilemma-strength parameters become functions of the cooperation level:

$$D_r(\rho) = (1 - \rho)D_r^1 + \rho D_r^2, \quad D_g(\rho) = (1 - \rho)D_g^1 + \rho D_g^2. \quad (4.2)$$

As an analytical benchmark, we use the replicator equation for regular graphs in the weak selection regime derived in Sect. 3.2.3. For a fixed game with dilemma-strength parameters D_g and D_r , the equation reads

$$\dot{\rho} = \rho(1 - \rho) \left[\rho(D_r - D_g) - D_r - \frac{D_r + D_g}{k - 2} \right], \quad (4.3)$$

where k is the degree of the graph. Substituting the expressions for D_g and D_r given by Eqs. (4.2) into Eq. (4.3), we obtain the graph-replicator equation for the endogenous-feedback model:

$$\dot{\rho} = \rho(1 - \rho) \left[\rho(D_r(\rho) - D_g(\rho)) - D_r(\rho) - \frac{D_r(\rho) + D_g(\rho)}{k - 2} \right]. \quad (4.4)$$

For notational compactness, we define $c_k = \frac{1}{k-2}$. After expanding the terms in Eq. (4.4), we obtain

$$\begin{aligned} \dot{\rho} = \rho(1 - \rho) & \left[\rho^2 (D_r^2 + D_g^1 - D_r^1 - D_g^2) + \rho \left((2 + c_k)D_r^1 - (1 + c_k)D_r^2 \right. \right. \\ & \left. \left. - (1 - c_k)D_g^1 - c_k D_g^2 \right) - (1 + c_k)D_r^1 - c_k D_g^1 \right]. \end{aligned} \quad (4.5)$$

This equation provides the analytical prediction used in this section. It should be interpreted as a pair-approximation benchmark for regular graphs in the weak-selection regime.

4.1.1 Simulation Framework

As in Sect. 3.3, we consider a square lattice of linear size $L = 200$, corresponding to a population of $N = L^2 = 40000$ individuals. Each individual occupies one site of the lattice, and each site is occupied by exactly one player. Thus, there are no empty sites. Each player can adopt one of two strategies, cooperation denoted as C or defection denoted as D . Periodic boundary conditions are imposed in both spatial directions, so that the lattice has the topology of a torus and boundary effects are removed. In contrast with the fixed-game setting considered in Sect. 3.3, here we consider the endogenous-feedback model as described above. The payoff matrix is updated during the dynamics. For a simulation initialised at cooperation level ρ_0 , the initial game is

$$G(\rho_0) = (1 - \rho_0)G_1 + \rho_0 G_2.$$

At later times, the current cooperation level determines the current game:

$$G(t) = (1 - \rho(t))G_1 + \rho(t)G_2,$$

played by the population. Thus, strategy updates modify $\rho(t)$, which in turn modifies the payoff matrix used in subsequent interactions. In the Monte Carlo simulations, this feedback loop is implemented by updating the effective game according to the current value of $\rho(t)$. At each elementary update, a focal player X is selected uniformly at random. Then, one of its neighbours Y is chosen uniformly at random. The neighbourhood can be either the von Neumann neighbourhood, with degree $k = 4$, or the Moore neighbourhood, with degree $k = 8$, as illustrated in Fig. 3.6. The two selected players accumulate their payoffs by playing the current game $G(t)$ against all their neighbours and, when self-interaction is included, also against themselves. Once the payoffs Π_X and Π_Y have been computed, the focal player X adopts the strategy of player Y according to the Fermi rule introduced in Eq. (1.6). The selection intensity is fixed to $\beta = 0.2$. Time is measured in Monte Carlo sweeps. For each pair of reference games G_1 and G_2 , we run several independent Monte Carlo simulations for different initial cooperation levels ρ_0 . The trajectories are averaged over independent runs, and the asymptotic cooperation level ρ_∞ is estimated by averaging $\rho(t)$ over the final part of the simulation after convergence.

4.1.2 Von Neumann Neighbourhood

We first consider a square lattice with von Neumann neighbourhood, without self-interaction. In this case, each player interacts with its four nearest neighbours, so that the degree of the graph is $k = 4$. Substituting $k = 4$ into Eq. (4.5), we obtain

$$\begin{aligned} \dot{\rho} = \rho(1 - \rho) & \left[\rho^2 (D_r^2 + D_g^1 - D_r^1 - D_g^2) \right. \\ & \left. + \frac{\rho}{2} (5D_r^1 - 3D_r^2 - D_g^1 - D_g^2) - \frac{1}{2} (3D_r^1 + D_g^1) \right]. \end{aligned} \quad (4.6)$$

This equation always admits the two boundary stationary points $\rho = 0$ and $\rho = 1$. In addition, up to two internal stationary points may exist. These are obtained by setting the expression inside the square brackets of Eq. (4.6) equal to zero and retaining only the solutions $\rho^* \in (0, 1)$. The local stability of stationary points is determined by the sign of the Jacobian $J(\rho) = \frac{d\dot{\rho}}{d\rho}$. Since for an internal stationary point $\rho^* \in (0, 1)$, the expression inside the square brackets of Eq. (4.6) vanishes, the Jacobian reduces to

$$J(\rho^*) = \rho^*(1 - \rho^*) \frac{dQ(\rho)}{d\rho} \Big|_{\rho=\rho^*},$$

where $Q(\rho)$ denotes the expression inside the square brackets of Eq. (4.6). Since $\rho^*(1 - \rho^*) > 0$, the stability of an internal fixed point is determined only by the sign of $dQ/d\rho$ evaluated at ρ^* . It turns out that the full-defection state is stable if $3D_r^1 + D_g^1 > 0$. This condition depends only on the parameters of G_1 . This is expected because, when $\rho = 0$, the effective game played by the population is

$G(0) = G_1$. Therefore, full defection can be stable only if the first reference game is a game in which full defection is stable, consistently with the fixed-game condition derived in Sect. 3.3.2. Similarly, the full-cooperation state is stable if $D_r^2 + 3D_g^2 < 0$. This condition depends only on the parameters of G_2 . Indeed, when $\rho = 1$, the effective game is $G(1) = G_2$. Therefore, full cooperation can be stable only if the second reference game is a game in which full cooperation is stable. Finally, an internal stationary point $\rho^* \in (0, 1)$ is stable if

$$2(D_r^2 + D_g^1 - D_r^1 - D_g^2)\rho^* + \frac{1}{2}(5D_r^1 - 3D_r^2 - D_g^1 - D_g^2) < 0.$$

Prisoner's Dilemma as G_2 . We first specialise the analysis to the case in which the reference game associated with full cooperation, G_2 , is a Prisoner's Dilemma. This choice induces a feedback dynamics in which increasing cooperation shifts the effective game towards a more defective regime. Indeed, according to Eq. (4.1), as ρ increases, the game $G(\rho)$ moves in the direction of G_2 . As a representative case, we fix $D_r^2 = D_g^2 = 1$, corresponding to a Prisoner's Dilemma in the top-right region of the dilemma space. Substituting these values into Eq. (4.6), the graph-replicator equation becomes

$$\dot{\rho} = \rho(1 - \rho) \left[\rho^2(D_g^1 - D_r^1) + \rho \left(\frac{5}{2}D_r^1 - \frac{1}{2}D_g^1 - 2 \right) - \frac{3}{2}D_r^1 - \frac{1}{2}D_g^1 \right]. \quad (4.7)$$

The stability analysis shows that $\rho = 0$ is stable if $3D_r^1 + D_g^1 > 0$, while $\rho = 1$ is always unstable. An internal equilibrium $\rho^* \in (0, 1)$ is stable if

$$2\rho^*(D_g^1 - D_r^1) + \frac{5}{2}D_r^1 - \frac{1}{2}D_g^1 - 2 < 0.$$

At this point, we compare the analytical predictions with Monte Carlo simulations. The simulation framework is the one described in Sect. 4.1.1. In order to make the comparison with the well-mixed scenario as direct as possible, we consider the same choices of G_1 analysed in Sect. 2.2.3. In particular, we take

$$(D_g^1, D_r^1) \in \{(-0.92, 0.25), (-0.8, -0.3), (-0.3, -0.8), (0.7, -0.7)\}.$$

For each pair, we initialise the system at three different cooperation levels, $\rho_0 \in \{0.1, 0.5, 0.9\}$. For each initial condition, 10 independent Monte Carlo simulations are performed and the results are averaged. The simulations are run until convergence to a stationary regime. Figure 4.1 shows the corresponding trajectories in the dilemma space. The second reference game G_2 is represented by a purple square, while the different choices of G_1 are represented by white squares. The three initial conditions are shown as coloured points. The black curves represent the trajectories of the effective game $G(\rho(t))$ during the dynamics, and the red crosses mark the final positions reached once the stationary regime has been attained. In order to assess whether chimera-game regimes are present also on square lattices, the asymptotic outcomes of the endogenous-feedback dynamics must be compared with the corresponding fixed-game dynamics on the same spatial structure. The relevant benchmark is therefore not the well-mixed fixed-game case, but the structured




Figure 4.1: Trajectories of the effective game $G(\rho(t))$ in the dilemma space for the endogenous-feedback model on a square lattice with von Neumann neighbourhood. The second reference game is fixed to a Prisoner's Dilemma, $G_2 = G(1, 1)$, and is represented by a purple square. The first reference games G_1 are represented by white squares and correspond to $(D_g^1, D_r^1) = (-0.92, 0.25)$, $(-0.8, -0.3)$, $(-0.3, -0.8)$ and $(0.7, -0.7)$. For each pair, simulations are initialised at $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$, shown as coloured points. The black curves show the trajectories followed by the effective game during the dynamics, while the red crosses indicate the final positions reached after convergence. Simulations are performed on a square lattice with linear size $L = 200$, periodic boundary conditions, no self-interaction and Fermi updating with $\beta = 0.2$. For each initial condition, 10 independent Monte Carlo simulations are performed and the results are averaged. The final position is computed from the asymptotic cooperation level, estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. *Original simulation by the author.*

fixed-game scenario discussed in Sect. 3.3.2, and in particular the results shown in Fig. 3.11. With respect to this benchmark, the results show that chimera-game regimes are also present on square lattices. In particular, this occurs for the two parameter pairs $(D_g^1, D_r^1) = (-0.8, -0.3)$ and $(D_g^1, D_r^1) = (-0.92, 0.25)$. In both cases, the endogenous-feedback dynamics stabilises an intermediate cooperation level while the final effective game lies in the Stag Hunt region. This is not a standard fixed-game outcome on the von Neumann lattice, where Stag Hunt games are bistable and converge to one of the two absorbing states. Hence, the observed intermediate

equilibria correspond to chimera-game regimes generated by the endogenous feedback mechanism. The first case, $(D_g^1, D_r^1) = (-0.8, -0.3)$, is consistent with the well-mixed analysis of Sect. 2.2.3, where the same pair already produced a chimera-game regime. The second case, $(D_g^1, D_r^1) = (-0.92, 0.25)$, is more specific to the spatial setting. In the well-mixed model, this pair led to full defection, whereas on the square lattice it reaches a positive stationary cooperation level.

The evolution of the global cooperation level over time is reported in Fig. 4.2. The four panels correspond to the same choices of G_1 shown in Fig. 4.1. In each panel, the three curves show the averaged trajectories obtained from the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. The horizontal red line indicates the stable stationary level predicted by the analytical benchmark given by Eq. (4.7). The

Cooperation Behaviour over Time on the Square Lattice




Figure 4.2: Time evolution of the cooperation level for the endogenous-feedback model on a square lattice with von Neumann neighbourhood, with $G_2 = G(1, 1)$ fixed as a Prisoner's Dilemma. The four panels correspond to the same choices of the first reference game G_1 shown in Fig. 4.1, namely $(D_g^1, D_r^1) = (-0.92, 0.25)$, $(-0.8, -0.3)$, $(-0.3, -0.8)$ and $(0.7, -0.7)$. In each panel, the curves show the global cooperation level over sweeps obtained from the three initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. The horizontal red line marks the stable stationary cooperation level predicted by the graph-replicator equation in Eq. (4.7). Simulations are performed on a square lattice with linear size $L = 200$, periodic boundary conditions, no self-interaction and Fermi updating with $\beta = 0.2$. For each initial condition, results are averaged over 10 independent Monte Carlo simulations. Different panels are shown over different Monte Carlo sweep windows. *Original simulation by the author.*

agreement between Monte Carlo simulations and the analytical prediction depends

on the region of the dilemma space in which G_1 is located. When G_1 is a Harmony Game, the agreement is excellent. When G_1 is a Snowdrift Game, small deviations appear, but the analytical prediction still captures the qualitative behaviour of the dynamics. By contrast, the largest discrepancy occurs when G_1 is a Stag Hunt. In this case, the lattice structure allows cooperators to survive more effectively than predicted by the pair-approximation benchmark. As discussed in greater detail below, this discrepancy is related to the formation of spatial cooperative domains. These domains protect cooperators from invasion by defectors and are not fully captured by the analytical approximation.

To test the robustness of these observations more systematically, we extend the analysis to a larger set of reference games G_1 , while keeping $G_2 = G(D_g^2, D_r^2) = G(1, 1)$. The selected games G_1 are chosen along two lines in the dilemma space, namely $D_g^1 = -0.9$ and $D_r^1 = -0.9$. This is the same systematic scan performed in the well-mixed scenario studied in Sect. 2.2.3. This choice provides a clear geometric scan of the dilemma space while keeping the selected games away from the boundary axes $D_g = 0$ and $D_r = 0$, where the classification of game classes becomes degenerate. The aim is to assess how the introduction of spatial structure modifies the dynamics of the system. For each pair (G_1, G_2) , we use the simulation protocol described in Sect. 4.1.1, with initial cooperation levels $\rho_0 \in \{0.1, 0.5, 0.9\}$. For each initial condition, 10 independent Monte Carlo simulations are performed and averaged. The equilibrium game reached by the system is estimated by averaging the final cooperation level over the last 20% of the simulation time, after convergence. Figure 4.3 shows the corresponding trajectories in the dilemma space. As before, G_2 is represented by a filled purple square, while the selected games G_1 are represented by open squares. The black curves show the trajectories of the effective game $G(\rho(t))$, and the red crosses mark the final games reached by the Monte Carlo simulations. In addition, cyan crosses indicate the stationary games predicted by the graph-replicator equation in Eq. (4.7). This allows a direct visual comparison between the spatial simulations and the analytical pair-approximation benchmark. The scan shows that chimera regimes appear for several choices of the first reference game G_1 . In particular, as in the well-mixed case, they appear for the Harmony-Game choices of G_1 along the branch $D_g^1 = -0.9$ with $D_r^1 > -0.9$. In these cases, the endogenous-feedback dynamics converges to stable intermediate cooperation levels, even though the final effective game lies in the Stag Hunt region where the corresponding fixed game would not produce the same outcome as shown in Sect. 3.3.2. A significant difference with respect to the well-mixed case appears for the Stag Hunt choices of G_1 . In the well-mixed scan reported in Fig. 2.6, the trajectories starting from Stag Hunt reference games converged to full defection. By contrast, on the square lattice, most of the corresponding simulations reach a final cooperation level significantly different from zero. This shows that spatial structure can promote the *emergence of cooperation* in parameter regions where the well-mixed model produces full defection.

The comparison with the cyan crosses also shows where the analytical prediction is accurate and where it breaks down. The largest discrepancies between Monte Carlo simulations and the graph-replicator predictions occur when G_1 is a Stag Hunt. This is consistent with the behaviour already observed in Fig. 4.2. The reason is that the




Figure 4.3: Systematic scan of the endogenous-feedback model on a square lattice with von Neumann neighbourhood, with G_2 fixed to a Prisoner's Dilemma, $G_2 = G(1, 1)$. The figure shows the trajectories in the (D_g, D_r) dilemma space obtained for selected reference games G_1 . The games G_1 are chosen along the lines $D_g^1 = -0.9$ and $D_r^1 = -0.9$, avoiding the boundary axes $D_g = 0$ and $D_r = 0$ where the classification of game classes is degenerate. The fixed game G_2 is indicated by a filled purple square, while the selected games G_1 are indicated by open squares. For each pair (G_1, G_2) , simulations are initialised from $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. The corresponding initial games lie on the segment connecting G_1 to G_2 . The black trajectories show the evolution of the effective game $G(\rho(t))$ in the dilemma space. Red crosses mark the equilibrium games reached in the Monte Carlo simulations, while cyan crosses mark the equilibrium games predicted by the graph-replicator equation given by Eq. (4.7). Simulations are performed on a square lattice with linear size $L = 200$, periodic boundary conditions, no self-interaction and Fermi updating with $\beta = 0.2$. For each initial condition, results are averaged over 10 independent Monte Carlo simulations. The final cooperation level is estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. *Original simulation by the author.*

square lattice allows cooperators to form spatial clusters, which can protect them from invasion by defectors. This clustering effect is not fully captured by the pair-approximation benchmark and will be analysed in more detail later on. The full cooperation time-series plots for all values of G_1 considered in Fig. 4.3 are reported

in Appendix B.1.1. To visualise the comparison between Monte Carlo simulations and analytical predictions more directly, we also plot the final cooperation level as a function of the varying parameter along the two scan lines. Figure 4.4 reports the comparison along the line $D_g^1 = -0.9$, where D_r^1 is varied, while Fig. 4.5 gives the corresponding comparison along the line $D_r^1 = -0.9$, where D_g^1 is varied. Blue circles represent the final cooperation level obtained from Monte Carlo simulations along with the corresponding standard deviation. Orange squares represent stable analytical equilibria, while green squares represent unstable analytical equilibria.




Figure 4.4: Comparison between Monte Carlo simulations and analytical predictions for the endogenous-feedback model on a square lattice with von Neumann neighbourhood and $G_2 = G(1, 1)$ fixed as a Prisoner's Dilemma. The figure shows the final cooperation level along the line $D_g^1 = -0.9$ as D_r^1 is varied. Blue circles indicate the asymptotic cooperation level obtained from Monte Carlo simulations, estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. Error bars represent the standard deviation of the cooperation level computed over the same final 20% time window; they are barely visible because they are very small. Orange squares indicate stable equilibria predicted by the graph-replicator equation, while green squares indicate unstable analytical equilibria. The analytical equilibria are obtained from Eq. (4.7). Simulations are performed with $\rho_0 = 0.5$, $N = 40000$, periodic boundary conditions, no self-interaction and Fermi updating with $\beta = 0.2$. The largest discrepancies appear in the Stag Hunt region, where spatial clustering allows cooperators to survive more effectively than predicted by the pair-approximation benchmark. *Original illustration by the author.*

The comparison confirms the qualitative picture obtained from the dilemma space trajectories. The analytical benchmark captures the behaviour well when the first reference game is a Harmony Game, and remains qualitatively accurate for several Snowdrift choices. However, the largest deviations appear in the Stag Hunt region. In this region, Monte Carlo simulations on the square lattice often reach cooperation levels significantly different from zero even when the analytical approximation




Figure 4.5: Comparison between Monte Carlo simulations and analytical predictions for the endogenous-feedback model on a square lattice with von Neumann neighbourhood and $G_2 = G(1, 1)$ fixed as a Prisoner's Dilemma. The figure shows the final cooperation level along the line $D_r^1 = -0.9$ as D_g^1 is varied. Blue circles indicate the asymptotic cooperation level obtained from Monte Carlo simulations, estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. Error bars represent the standard deviation of the cooperation level computed over the same final 20% time window; they are barely visible because they are very small. Orange squares indicate stable equilibria predicted by the graph-replicator equation, while green squares indicate unstable analytical equilibria. The analytical equilibria are obtained from Eq. (4.7). Simulations are performed with $\rho_0 = 0.5$, $N = 40000$, periodic boundary conditions, no self-interaction and Fermi updating with $\beta = 0.2$. Along this scan, the main discrepancy appears in the Snowdrift region, corresponding to $D_g^1 > 0$. However, the disagreement is small, and the analytical benchmark still captures the qualitative behaviour of the dynamics. *Original illustration by the author.*

predicts full defection or a much lower cooperation level. This provides further evidence that the discrepancy is driven by spatial clustering and by the formation of cooperative domains, which allow cooperators to survive locally despite the invasion pressure from defectors.

Harmony Game as G_2 . At this point, we analyse the scenario in which the reference game associated with full cooperation, G_2 , is a Harmony Game. As a representative choice, we consider $G_2 = G(D_g^2, D_r^2) = G(-1, -1)$, corresponding to the bottom-left corner of the social dilemma plane. Substituting $D_g^2 = D_r^2 = -1$ into Eq. (4.5), the replicator equation becomes

$$\dot{\rho} = \rho(1 - \rho) \left[\rho^2(D_g^1 - D_r^1) + \rho \left(\frac{5}{2}D_r^1 - \frac{1}{2}D_g^1 + 2 \right) - \frac{1}{2}(3D_r^1 + D_g^1) \right]. \quad (4.8)$$

Therefore, besides the two boundary stationary points $\rho = 0$ and $\rho = 1$, the dynamics may admit up to two internal stationary points. These are obtained by setting the expression inside the square brackets of Eq. (4.8) equal to zero and retaining only the solutions $\rho^* \in (0, 1)$. The local stability of a stationary point is determined by the Jacobian $J(\rho) = \frac{d\dot{\rho}}{d\rho}$. It turns out that $\rho = 0$ is stable if $3D_r^1 + D_g^1 > 0$, whereas $\rho = 1$ is always stable. Finally, an interior stationary point ρ^* , when present, is stable if

$$2\rho^*(D_g^1 - D_r^1) + \frac{5}{2}D_r^1 - \frac{1}{2}D_g^1 + 2 < 0.$$

We use the same four choices of G_1 considered in the well-mixed case of Sect. 2.2.3:

$$(D_g^1, D_r^1) \in \{(-0.65, 0.75), (0.30, 0.80), (0.80, 0.30), (0.70, -0.70)\}.$$

This allows us to isolate the effect of spatial structure on the endogenous-feedback dynamics. The corresponding trajectories in the dilemma space are shown in Fig. 4.6, while the associated time evolution of the cooperation level is reported in Fig. 4.7.

The two figures should be compared with the corresponding well-mixed results reported in Sect. 2.2.3. As in the well-mixed case, the system eventually converges to absorbing states, namely full cooperation or full defection in agreement with the analytical prediction of Eq. (4.8). Thus, unlike the case in which G_2 is a Prisoner's Dilemma, here the spatial structure of the square lattice does not lead to qualitatively different asymptotic outcomes.

4.1.3 Moore Neighbourhood

We now consider the case of a square lattice with Moore neighbourhood, corresponding to degree $k = 8$, while keeping $G_2 = G(D_g^2, D_r^2) = G(1, 1)$, that is, a Prisoner's Dilemma as the reference game associated with full cooperation. Substituting $k = 8$ into Eq. (4.5), the replicator equation on graphs becomes

$$\dot{\rho} = \rho(1 - \rho) \left[\rho^2 (D_g^1 - D_r^1) + \rho \left(\frac{13}{6}D_r^1 - \frac{5}{6}D_g^1 - \frac{4}{3} \right) - \left(\frac{7}{6}D_r^1 + \frac{1}{6}D_g^1 \right) \right]. \quad (4.9)$$

Thus, the dynamics admits the two boundary stationary points $\rho = 0$ and $\rho = 1$, together with up to two internal stationary points. These are obtained by setting the expression inside the square brackets of Eq. (4.9) equal to zero and retaining only the solutions $\rho^* \in (0, 1)$. The stability conditions are the following:

- $\rho = 0$ is stable if

$$7D_r^1 + D_g^1 > 0;$$

- $\rho = 1$ is always unstable;
- an interior equilibrium $\rho^* \in (0, 1)$, when present, is stable if

$$2\rho^* (D_g^1 - D_r^1) + \frac{13}{6}D_r^1 - \frac{5}{6}D_g^1 - \frac{4}{3} < 0.$$




Figure 4.6: Trajectories in the (D_g, D_r) dilemma space for the endogenous-feedback model on a square lattice with von Neumann neighbourhood, with $G_2 = G(D_g^2, D_r^2) = G(-1, -1)$ fixed as a Harmony Game. The four representative choices of the reference game G_1 are $(D_g^1, D_r^1) = (-0.65, 0.75)$, $(0.30, 0.80)$, $(0.80, 0.30)$ and $(0.70, -0.70)$, chosen to match the corresponding well-mixed analysis of Sect. 2.2.3. The game G_2 is indicated by a filled purple square, while the four choices of G_1 are indicated by open squares. For each G_1 , Monte Carlo simulations are initialised from three cooperation levels, $\rho_0 = 0.1, 0.5, 0.9$, shown by coloured dots. Since the payoff matrix evolves according to the endogenous-feedback rule $G(t) = [1 - \rho(t)]G_1 + \rho(t)G_2$, changes in the cooperation level induce a trajectory of the effective game in the dilemma space. These trajectories are represented by black lines. Red crosses indicate the final equilibrium games reached from each initial condition. For each initial condition, 10 independent Monte Carlo simulations are performed and averaged. The final position is obtained from the asymptotic cooperation level, estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. The corresponding time evolution of the cooperation level, together with the analytical prediction, is reported in Fig. 4.7. In this case, the final outcomes are absorbing, in agreement with the analytical predictions and with the corresponding well-mixed case shown in Fig. 2.7. *Original simulation by the author.*

As in the previous subsection, the full-defection equilibrium depends only on the parameters of G_1 , because when $\rho = 0$ the effective game coincides with G_1 . By contrast, full cooperation is always unstable, consistently with the fact that G_2 is a Prisoner's Dilemma.

At this point, we compare these analytical predictions with Monte Carlo simulations on the square lattice. The simulation framework is the one described in Sect. 4.1.1. To systematically explore the dilemma space and ensure a consistent comparison with the well-mixed and von Neumann lattice results, we consider the same set of reference games G_1 as in the previous cases. In particular, we vary G_1 along the two lines $D_g^1 = -0.9$ and $D_r^1 = -0.9$, while keeping the target game fixed at


Cooperation Behaviour over Time on the Square Lattice: HG as G_2 

Figure 4.7: Time evolution of the cooperation level ρ for the endogenous-feedback model on a square lattice with von Neumann neighbourhood, with $G_2 = G(D_g^2, D_r^2) = G(-1, -1)$ fixed as a Harmony Game. The four panels correspond to the four choices $(D_g^1, D_r^1) = (-0.65, 0.75)$, $(0.30, 0.80)$, $(0.80, 0.30)$ and $(0.70, -0.70)$, respectively. In each panel, the dynamics is shown for three initial cooperation levels, $\rho_0 = 0.1, 0.5, 0.9$, and the results are averaged over 10 independent Monte Carlo simulations. The red horizontal line marks the stable equilibrium cooperation level predicted by the replicator equation Eq. (4.8). Simulations are performed on a square lattice of size $L = 200$ with periodic boundary conditions, no self-interaction and Fermi updating with selection intensity $\beta = 0.2$. Different panels are shown over different Monte Carlo sweep windows. In this case, the final outcomes are absorbing, in agreement with the analytical predictions and with the corresponding well-mixed case shown in Fig. 2.8. *Original simulation by the author.*

$G_2 = G(1, 1)$ throughout the analysis. The resulting trajectories in the dilemma space are shown in Fig. 4.8. The filled purple square marks the fixed game G_2 , the open squares mark the selected values of G_1 , and the coloured dots denote the initial games corresponding to $\rho_0 = 0.1, 0.5$ and 0.9 . The black curves represent the trajectories followed by the effective game during the Monte Carlo simulations, while the red crosses indicate the equilibrium games reached numerically after convergence. In addition, blue crosses mark the analytical equilibrium games predicted by the replicator equation on graphs. The results for the endogenous-feedback dynamics on the Moore neighbourhood must be compared with the corresponding fixed-game results discussed in Subsection 3.3.3. As in the von Neumann case, the introduction of endogenous feedback gives rise to chimera-game regimes also on the Moore neighbourhood. In particular, for all the reference games tested along the line $D_g^1 = -0.9$




Figure 4.8: Systematic scan of the endogenous-feedback model on a square lattice with Moore neighbourhood ($k = 8$), with the target game fixed at $G_2 = G(D_g^2, D_r^2) = G(1, 1)$, corresponding to a Prisoner's Dilemma. The reference games G_1 are selected along the two lines $D_g^1 = -0.9$ and $D_r^1 = -0.9$, providing a geometric scan of the dilemma space while avoiding the boundary axes. The filled purple square indicates the fixed game G_2 , whereas the open squares indicate the selected values of G_1 . For each pair (G_1, G_2) and each initial condition $\rho_0 = 0.1, 0.5, 0.9$, we performed 10 independent Monte Carlo simulations and averaged the results. The coloured dots indicate the initial conditions, while the black curves show the corresponding trajectories followed by the effective game in the dilemma space. The red crosses indicate the equilibrium games reached numerically, obtained from the average final cooperation level computed over the last 20% of the simulation time after convergence. The blue crosses indicate the analytical predictions obtained by Eq. (4.9). The figure shows that chimera-game regimes also emerge in the Moore case, while the largest discrepancies between numerical simulations and analytical predictions occur for Stag Hunt choices of G_1 . *Original simulation by the author.*

with $D_r^1 > -0.9$, the system converges to stable nontrivial cooperation levels, even though the corresponding fixed-game dynamics on the same structured population would not display the same asymptotic behaviour. This confirms that the emergence of chimera-game outcomes is not specific to the von Neumann interaction structure, but persists when the local neighbourhood is enlarged to include diagonal neighbours. At the same time, the comparison with the analytical predictions shows the same qualitative limitation already observed in the von Neumann case: the largest

simulation–theory discrepancies occur when G_1 lies in the Stag Hunt region. The corresponding time series of the cooperation level for the different reference games, together with the associated analytical predictions, are reported in Appendix B.2. For the same set of games, the direct comparison between the final cooperation level obtained numerically and the one predicted analytically, analogous to Figs. 4.5 and 4.4, is also reported in Appendix B.2. Comparing these results with the corresponding von Neumann plots shows that the agreement with the pair-approximation predictions is slightly worse in the Moore case. This is consistent with the fact that, as discussed in the following chapter, the inclusion of diagonal interactions increases clustering effects, which are not fully captured by a pair-level analytical description.

4.1.4 Self-Interaction

We then analyse the effect of self-interaction on the endogenous-feedback dynamics on the square lattice. In this setting, each player interacts not only with its neighbours, but also with itself when payoffs are computed. Recalling the payoff parametrisation introduced in Eq. (1.5),

$$G(D_g, D_r) = \begin{pmatrix} 1 & -D_r \\ 1 + D_g & 0 \end{pmatrix},$$

we see that a cooperator therefore receives an additional payoff contribution equal to 1, while a defector receives an additional payoff contribution equal to 0. Self-interaction thus introduces an intrinsic advantage for cooperators. We focus on the case in which the second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$, and we repeat the same systematic scan considered above, with the first reference games G_1 chosen along the two lines $D_g^1 = -0.9$ and $D_r^1 = -0.9$. The aim is to compare the resulting trajectories with the corresponding case without self-interaction, shown in Fig. 4.3. The trajectories in the dilemma space are reported in Fig. 4.9. The game G_2 is represented by a filled purple square, while the selected games G_1 are represented by open squares. The black curves show the trajectories followed by the effective game $G(\rho(t))$, and the red crosses indicate the final equilibrium games reached in the Monte Carlo simulations. Compared with the case without self-interaction, the final equilibrium points are systematically shifted towards G_2 . Since the effective game is defined as $G(\rho) = (1 - \rho)G_1 + \rho G_2$, a shift towards G_2 corresponds to a larger final value of ρ . Therefore, self-interaction shifts the dynamics towards more cooperative outcomes. This effect is consistent with the payoff structure of the model discussed above. As shown in the figure, self-interaction can sustain non-zero cooperation levels even when the final effective game remains in the Prisoner’s Dilemma region, where the corresponding fixed-game dynamics without self-interaction would lead to full defection.

4.2 Cluster Dynamics

In this section, we study the cluster dynamics generated by the endogenous-feedback model on the square lattice. Thanks to the spatial structure of the lattice, we can




Figure 4.9: Systematic scan of the endogenous-feedback model on a square lattice with von Neumann neighbourhood and self-interaction, with $G_2 = G(D_g^2, D_r^2) = G(1, 1)$ fixed as a Prisoner's Dilemma. The selected reference games G_1 are chosen along the two lines $D_g^1 = -0.9$ and $D_r^1 = -0.9$. The filled purple square indicates the fixed game G_2 , whereas the open squares indicate the selected values of G_1 . For each pair (G_1, G_2) , Monte Carlo simulations are initialised at $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$, shown by coloured dots. For each initial condition, 10 independent Monte Carlo simulations are performed and averaged. The black curves represent the trajectories followed by the effective game in the dilemma space, while the red crosses indicate the equilibrium games reached numerically, computed from the final cooperation level averaged over the last 20% of the simulation time after convergence. Simulations are performed on a square lattice with linear size $L = 200$, periodic boundary conditions, von Neumann neighbourhood, self-interaction and Fermi updating with $\beta = 0.2$. Compared with the corresponding case without self-interaction, the equilibrium points are systematically shifted towards G_2 , showing that self-interaction increases the final level of cooperation across the scanned parameter region. *Original simulation by the author.*

analyse not only the global level of cooperation, but also the spatial organisation of cooperators into clusters. We first show that, in some parameter regimes, the system displays two dynamical processes occurring on clearly separated time scales: a fast convergence of the global cooperation level and a much slower spatial reorganisation of cooperative clusters. We then introduce quantitative measures to characterise the structure of these clusters. Finally, we study how the compactness ratio, which

measures the degree of spatial aggregation of cooperative domains, varies as a function of the effective game played by the population. This reveals a sharp change in the spatial organisation of cooperators, from compact cooperative domains to more fragmented configurations.

4.2.1 Fast Global Convergence and Slow Spatial Reorganisation

We investigate the spatial organisation underlying the asymptotic cooperation levels discussed in the previous section. We focus on a representative endogenous-feedback simulation on a square lattice with von Neumann neighbourhood. The second reference game is fixed as a Prisoner's Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$, while the first reference game is chosen as $G_1 = G(D_g^1, D_r^1) = G(-0.9, 0.1)$. The lattice is initialised with cooperation level $\rho_0 = 0.5$, and the simulation is run for 50000 Monte Carlo sweeps. Figure 4.10 shows a sequence of lattice snapshots taken at different Monte Carlo sweeps. Cooperators are represented in black, while defectors are represented in white. The corresponding evolution of the global cooperation level is reported in Fig. 4.11. The comparison between Figs. 4.10 and 4.11 shows




Figure 4.10: Snapshots of the square lattice during the endogenous-feedback dynamics with von Neumann neighbourhood. The reference games are $G_1 = G(-0.9, 0.1)$ and $G_2 = G(1, 1)$, with G_2 fixed as a Prisoner's Dilemma. The lattice is initialised at $\rho_0 = 0.5$ and the simulation run for 50000 Monte Carlo sweeps. Cooperators are represented in black, while defectors are represented in white. The snapshots correspond to Monte Carlo sweeps 0, 50, 200, 500, 1000, 5000, 10000, 20000, 35000 and 50000. The first row reports the first five snapshots, while the second row reports the remaining five. The figure shows that, after the global cooperation level has already approached a stationary value after a few sweeps, the spatial configuration continues to evolve through a slow rearrangement of cooperative clusters, eventually leading to the formation of a single large cooperative cluster. *Original simulation by the author.*

that two distinct dynamics occur on different time scales. The global cooperation




Figure 4.11: Global cooperation level $\rho(t)$ for the same simulation shown in Fig. 4.10. The dynamics is run on a square lattice with von Neumann neighbourhood, with $G_1 = G(-0.9, 0.1)$ and $G_2 = G(1, 1)$, where G_2 is fixed as a Prisoner's Dilemma. The lattice is initialised at $\rho_0 = 0.5$ and the simulation is run for 50000 Monte Carlo sweeps. The main panel shows the full time evolution, while the inset zooms in on the first 300 sweeps and highlights the initial transient. The global cooperation level rapidly relaxes towards an approximately stationary value, well before the spatial rearrangement displayed in Fig. 4.10 is complete. This indicates that macroscopic stationarity in $\rho(t)$ can coexist with a continuing microscopic reorganisation of the spatial configuration. *Original simulation by the author.*

level rapidly approaches an approximately stationary value, whereas the spatial configuration continues to evolve for much longer. Indeed, macroscopic stationarity does not imply spatial stationarity. Even when $\rho(t)$ remains nearly constant, local strategy updates still rearrange the microscopic configuration, causing cooperative clusters to move, merge and reorganise across the lattice. In the representative simulation considered here, this slow spatial reorganisation eventually leads to the formation of a single large cooperative cluster. The qualitative mechanism can be understood by looking at the trajectory of the effective game in the dilemma space. As shown in Fig. 4.3, the equilibrium game lies in the Stag Hunt region. Since Stag Hunt is a coordination game, cooperators are favoured when they are surrounded by other cooperators. The spatial dynamics therefore promotes the formation of compact cooperative domains, which reduce the cooperator-defector interface and protect cooperators from invasion by defectors. This phenomenon is analysed more quantitatively in the following subsection, where we introduce measures of cluster structure and study the mechanism behind this spatial reorganisation.

4.2.2 Quantifying Cluster Structure

To quantify the morphology of the spatial configurations, we use a measure based on the total interface between cooperators and defectors. At each recorded snapshot, the square lattice is represented by a binary field $s_{m,n}$, where $s_{m,n} = 1$ if the site with lattice coordinates (m, n) is occupied by a cooperator and $s_{m,n} = 0$ otherwise.

Definition 4.1. The cooperative area is defined as

$$A = \sum_{m,n} s_{m,n}.$$

The **cooperative-defective interface** is measured as the number of nearest-neighbour links connecting a cooperator to a defector. With periodic boundary conditions, it can be written as

$$P = \sum_{m,n} [\mathbb{I}(s_{m,n} \neq s_{m+1,n}) + \mathbb{I}(s_{m,n} \neq s_{m,n+1})],$$

where all indices are understood modulo L , with L the linear size of the lattice. This implements periodic boundary conditions. The function $\mathbb{I}(\cdot)$ is the indicator function, defined as

$$\mathbb{I}(s_{m,n} \neq s_{p,q}) = \begin{cases} 1, & \text{if } s_{m,n} \neq s_{p,q}, \\ 0, & \text{if } s_{m,n} = s_{p,q}. \end{cases}$$

In this expression, the terms $s_{m+1,n}$ and $s_{m,n+1}$ denote the right and upper nearest neighbours of site (m, n) , respectively. Therefore, each horizontal and vertical nearest-neighbour pair is counted once. Thus, P measures the total length of the discrete cooperator–defector interface.

At this point, we combine the cooperative area and the cooperative-cluster perimeter into a single quantity that measures how compact or fragmented the spatial distribution of cooperators is.

Definition 4.2. The **compactness ratio** χ is defined as

$$\chi = \frac{P^2}{A}.$$

For a fixed cooperative area A , smaller values of χ correspond to more compact or more perimeter-efficient cooperative domains, whereas larger values indicate more elongated, fragmented, or irregular spatial structures. In this sense, χ can be interpreted as an inverse measure of compactness. This interpretation is illustrated in Fig. 4.12. The three configurations have the same cooperative area, but progressively larger perimeters. The compact square has the smallest value of χ , the elongated rectangle has an intermediate value, and the fragmented configuration has a much larger value since the total cooperator–defector interface is substantially larger. Observe that a continuum normalisation such as $P^2/(4\pi A)$ is not used here.


Compactness ratio on a 200×200 square lattice

Figure 4.12: Illustration of the compactness ratio $\chi = P^2/A$ on a 200×200 square lattice like the one used in the simulations. The three configurations have the same cooperative area $A = 100$, but different lattice perimeters P . Compact domains have a smaller interface relative to their area and therefore a lower value of χ , whereas elongated or fragmented configurations have larger cooperator–defector interfaces and higher values of χ . In the figure, only a central window of the lattice is displayed. *Original illustration by the author.*

The reason is that the relevant geometry is not continuous and isotropic. Boundaries are made of horizontal and vertical lattice edges, so the natural compact reference shapes are discrete square-like or rectangular domains rather than circles. Therefore, the unnormalised quantity P^2/A is more appropriate for comparing the relative morphology of configurations generated on a square lattice. There is, however, an additional point to consider. Since the simulations use periodic boundary conditions, the lattice has the topology of a torus. As a result, the perimeter is affected not only by the local shape of the boundary, but also by the global topology of the cooperative domain. A sufficiently extended cluster may reduce its perimeter by winding around the lattice. This effect is illustrated in Fig. 4.13, where a wrapping strip has a smaller interface than a non-wrapping domain with the same area. For this reason, χ should be interpreted as a global interface-based descriptor of morphology. It is sensitive to local compactness, fragmentation, elongation, and also to perimeter reductions induced by toroidal wrapping. A low value of χ therefore does not necessarily imply that a cooperative cluster is compact in the usual Euclidean sense. It may also indicate that the cluster has organised into a perimeter-efficient configuration allowed by the periodic geometry of the lattice.

4.2.3 Sharp Change in Cluster Morphology

In this subsection, we use the compactness ratio introduced in Sect. 4.2.2 to analyse the spatial structure of the cooperative clusters at equilibrium.

We focus again on the endogenous-feedback model on a square lattice with von Neumann neighbourhood, with the second reference game fixed as a Prisoner’s Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$. The parameter pairs considered here are selected from the systematic scan shown in Fig. 4.3. In particular, we consider reference games

Effect of toroidal wrapping on the lattice perimeter




Figure 4.13: Illustration of the effect of periodic boundary conditions on the compactness ratio. The two configurations have the same cooperative area $A = 20000$, but different spatial organisation. The left panel shows a non-wrapping rectangle, while the right panel shows a strip that spans the full horizontal direction of the lattice. Since the system has toroidal geometry, the left and right boundaries are identified, so the wrapping strip has only two interfaces and therefore a smaller perimeter. This shows that, on a torus, a lower value of $\chi = P^2/A$ may result not only from local compactness, but also from topological wrapping. *Original illustration by the author.*

G_1 along the two scan lines $D_g^1 = -0.9$ and $D_r^1 = -0.9$, excluding the boundary axes. More explicitly, we consider

$$D_g^1 = -0.9, D_r^1 \in \{-0.9, -0.8, -0.7, -0.6, -0.5, -0.4, -0.3, -0.2, -0.1, 0.1, 0.2, 0.3\},$$

and

$$D_r^1 = -0.9, D_g^1 \in \{-0.8, -0.7, -0.6, -0.5, -0.4, -0.3, -0.2, -0.1, 0.1, 0.2, 0.3\}.$$

These choices span the relevant part of the dilemma space already explored in the main text and allow us to analyse how cluster formation changes across different equilibrium games. For each pair (G_1, G_2) , the lattice is initialised with $\rho_0 = 0.5$, and the simulation is run for 50000 Monte Carlo sweeps. This large simulation time is chosen to ensure that the system has enough time not only to reach an approximately stationary cooperation level, but also to allow the slower spatial reorganisation of cooperative clusters discussed in Sect. 4.2.1. In all cases considered below, we verify that the observables used to quantify cluster structure fluctuate around a stationary value during the final part of the simulation. Fig. 4.14 shows how the equilibrium compactness ratio varies with the effective dilemma parameter reached by the dynamics. For each simulation, the compactness ratio reported is computed according to Def. 4.2. The value shown on the vertical axis is obtained by averaging this quantity over the final 20% of the simulation time. The corresponding error bar is the standard deviation computed over the same final time

window. The value of D_g shown on the horizontal axis is obtained in the same way: for each simulation, we compute the effective value $D_g(\rho(t))$ during the dynamics and average it over the final 20% of the simulation time. The horizontal error bars are the standard deviations of $D_g(\rho(t))$ over the same final time window. At the same time, Fig. 4.15 shows that, for the representative parameter pairs dis-




Figure 4.14: Equilibrium compactness ratio as a function of the equilibrium effective value of D_g for the endogenous-feedback model on a square lattice with von Neumann neighbourhood. The compactness ratio is computed as P^2/A , where A is the area of the cooperative cluster and P is the length of the cooperative-defective interface. The second reference game is fixed as a Prisoner's Dilemma, $G_2 = G(1, 1)$, and the selected values of G_1 are taken from the systematic scan shown in Fig. 4.3. Each point corresponds to one simulation initialised at $\rho_0 = 0.5$ and run for 50000 Monte Carlo sweeps. The value of the compactness ratio shown on the vertical axis is obtained by averaging $P(t)^2/A(t)$ over the final 20% of the simulation time. The value of D_g shown on the horizontal axis is obtained by averaging the effective value $D_g(\rho(t))$ over the same final time window. Error bars represent the corresponding standard deviations over this final 20% time window. The sharp increase in P^2/A indicates a sharp change in cluster morphology, from compact cooperative domains to more fragmented spatial configurations. *Original simulation by the author.*

played, the compactness ratio reaches a stationary fluctuating regime within the simulation time. This justifies the use of the final 20% time window to estimate it. At the same time, the figure shows that the stationary compactness ratio varies strongly across different choices of G_1 . It is useful to compare these values with the random initial condition. At $\rho_0 = 0.5$, each cooperator has, on average, two cooperative neighbours and two defective neighbours considering a von Neumann neighbourhood. Therefore, for a random initial configuration, the cooperative area is approximately $A_0 \simeq \frac{N}{2}$, while the cooperative-defective interface has expected




Figure 4.15: Time evolution of the compactness ratio for representative parameter pairs from the same set used in Fig. 4.14. The simulations are performed on a square lattice with von Neumann neighbourhood, with $G_2 = G(1,1)$ fixed as a Prisoner's Dilemma, initial condition $\rho_0 = 0.5$, and total duration of 50000 Monte Carlo sweeps. Only a subset of the parameter pairs is shown. The displayed curves correspond to $(D_g^1, D_r^1) = (-0.9, -0.4), (-0.9, -0.6), (-0.9, -0.8), (-0.8, -0.9), (-0.7, -0.9), (-0.5, -0.9), (-0.3, -0.9), (-0.1, -0.9), (-0.9, 0.3)$ and $(0.3, -0.9)$. The curves show that the compactness ratio reaches a fluctuating stationary regime, while its stationary value depends strongly on the equilibrium game reached by the endogenous-feedback dynamics. *Original simulation by the author.*

length $P_0 \simeq N$. For the compactness ratio used here, this gives the rough baseline $\frac{P_0^2}{A_0} \simeq \frac{N^2}{N/2} = 2N$. Since in our simulations $N = 40000$, the random initial condition corresponds approximately to $\frac{P_0^2}{A_0} \simeq 80000$. This estimate provides a useful reference scale and is consistent with the initial value of the compactness ratio observed in Fig. 4.15. The behaviour shown in Fig. 4.14 reveals a sharp change in cluster morphology. The ratio P^2/A acts as a structural indicator of the spatial organisation of cooperators. It remains small when cooperators form compact domains with a relatively short interface, and it increases sharply when cooperators become more fragmented and mixed with defectors. The mechanism behind this behaviour can be understood by comparing Fig. 4.14 with the trajectories in the dilemma space shown in Fig. 4.3. When the equilibrium effective game has negative effective D_g , the population plays a game in the Stag Hunt region. Stag Hunt games are coordination games. In this regime, cooperators benefit from being surrounded by other cooperators. As a consequence, cooperators tend to aggregate into compact domains. This reduces the length of the interface between cooperators and defectors and protects the cooperative cluster from exploitation. By contrast, as the equilibrium value of D_g^* moves towards the Snowdrift region, the strategic structure becomes closer to an anti-coordination game. In a Snowdrift-like regime, coexistence is favoured by

mixed local neighbourhoods rather than by compact cooperative domains. Cooperators no longer need to minimise their interface with defectors in the same way as in the Stag Hunt region. Therefore, cluster formation becomes less pronounced, and the compactness ratio takes larger values. The sharp change observed in the compactness ratio can therefore be interpreted as a crossover between two spatial regimes. In one regime, cooperators reorganise into compact cooperative clusters with a relatively small interface. In the other, cooperators remain more fragmented and mixed with defectors, producing a larger interface relative to their area.

The results of this chapter show that the endogenous-feedback mechanism can sustain chimera-game regimes on square lattices, where interactions are local and spatially organised. Compared with the fixed-game baseline developed in Chapter 3, the feedback between the cooperation level and the payoff matrix changes the evolutionary dynamics substantially, allowing stable coexistence regimes to survive in parameter regions where they would not be expected from the corresponding fixed-game dynamics. The spatial nature of the lattice also made it possible to analyse the organisation of cooperators in terms of clusters. The simulations showed that the global cooperation level and the spatial configuration do not necessarily evolve on the same time scale: while $\rho(t)$ can reach a quasi-stationary value relatively quickly, cooperative clusters may continue to reorganise over much longer times. Moreover, different effective games produced different cluster morphologies, ranging from fragmented cooperative regions to more compact domains able to resist defector invasion.

The square lattice, however, is a regular and spatially homogeneous structure. The next chapter therefore moves to complex networks, where interactions are still local but no longer constrained by a regular geometry. This makes it possible to investigate whether chimera-game regimes also survive on more general network topologies, and how structural features such as randomness, degree heterogeneity and clustering affect the final cooperation level and the accuracy of the analytical prediction.

Chapter 5

The Endogenous Feedback Model on Complex Networks

In this chapter, we investigate the endogenous-feedback model on complex networks. In the previous chapters, chimera-game regimes were observed in well-mixed populations and on square lattices. The natural question is whether these regimes also emerge when interactions take place on more general network topologies. This question is important because complex networks introduce structural features that are absent, or only weakly represented, in regular lattices. In particular, degree heterogeneity, short path lengths, and local clustering can affect the spreading of strategies and the stability of cooperation. We therefore investigate whether chimera-game regimes survive despite the presence of an underlying network topology. When they do, we analyse how network structure affects the final cooperation level and the accuracy of the analytical prediction. We first analyse the dynamics on Erdős–Rényi, Watts–Strogatz, and Barabási–Albert networks, comparing the Monte Carlo simulations with the analytical predictions obtained from pair approximation. We then focus on the role of clustering. By varying the amount of clustering while keeping the degree of every node fixed, we isolate its effect on the final cooperation level and show how clustering can lead to a breakdown of the pair-approximation prediction.

5.1 Chimera Games on Complex Networks

In this section, we investigate the emergence of chimera-game regimes on the complex network topologies introduced in Chapter 3, namely, Erdős–Rényi, Watts–Strogatz, and Barabási–Albert networks. We consider the endogenous-feedback model studied in the previous chapters. Given two reference games G_1 and G_2 , the payoff matrix at time t is

$$G(\rho(t)) = (1 - \rho(t))G_1 + \rho(t)G_2,$$

where $\rho(t)$ is the current cooperation level. Throughout this section, we focus on the case in which the second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. The selected games G_1 are chosen along two lines in the dilemma space, $D_g^1 = -0.9$ and $D_r^1 = -0.9$. This choice provides a geometric scan of the dilemma space while avoiding the boundary axes $D_g = 0$ and $D_r = 0$, where the classification

of game classes is degenerate. For each network topology, we perform Monte Carlo simulations of the endogenous-feedback dynamics and compare the results with the analytical predictions provided by the graph-replicator equation for regular graphs in the weak-selection regime. Unless otherwise stated, the numerical protocol used in this section is the same for all network topologies. For each selected pair (G_1, G_2) and each initial condition $\rho_0 \in \{0.1, 0.5, 0.9\}$, we generate 5 independent network realisations and perform 5 independent Monte Carlo simulations on each realisation. The reported results are therefore averaged over 25 runs, accounting both for the randomness of the network generation and for the stochasticity of the evolutionary dynamics. As in the previous chapters, the final cooperation level is estimated by averaging $\rho(t)$ over the final 20% of the simulation time window, after convergence to a stationary regime.

5.1.1 Erdős–Rényi Networks

We first consider Erdős–Rényi networks, whose main structural properties have been discussed in Sect. 3.4.1. To compare these results with those obtained on the square lattice of linear size $L = 200$, we fix the network size to $N = 40000$ and choose the link probability p so that the expected average degree is $\langle k \rangle = p(N - 1) = 4$. Thus, the typical number of interaction partners is the same as in the von Neumann neighbourhood considered on the square lattice. As an analytical benchmark, we use the graph-replicator equation obtained for $G_2 = G(1, 1)$ on a regular graph of degree $k = 4$, given by Eq. (4.7). A modelling choice concerns how payoffs are computed. One can either use accumulated payoffs, in which case high-degree nodes naturally receive contributions from more interactions, or degree-normalised payoffs, in which the payoff of each node is divided by its degree. We tested both prescriptions and found that the results are practically unchanged. Since the analytical treatment is naturally formulated in terms of average payoff, in what follows we focus on the degree-normalised case. Strategy updating follows the Fermi rule given by Eq. (1.6), with selection intensity $\beta = 0.8$. The simulations follow the numerical protocol described at the beginning of this section.

Figure 5.1 shows that chimera-game regimes are also present on Erdős–Rényi networks. In particular, along the line $D_g^1 = -0.9$, for the parameter range $-0.9 < D_r^1 < 0.3$, the effective game reached at equilibrium lies in the Stag Hunt region. In the corresponding fixed-game dynamics on Erdős–Rényi networks, discussed in Sect. 3.4.1, this region leads to full defection. By contrast, in the endogenous-feedback model the population sustains a non-zero level of cooperation. This shows that the chimera-game behaviour is not restricted to well-mixed populations or square lattices, but also emerges on random network topologies.

As we can see, the agreement between Monte Carlo simulations and the analytical prediction is excellent for all the parameter pairs considered. This indicates that, in the Erdős–Rényi case with degree-normalised payoffs and expected average degree $\langle k \rangle = 4$, the graph-replicator equation captures the macroscopic endogenous-feedback dynamics very accurately. The origin of this agreement, and its possible breakdown for other network structures, will be analysed in the following section. Observe that small residual cooperation levels may appear in parameter regions




Figure 5.1: Systematic scan of the endogenous-feedback model on Erdős–Rényi networks, with G_2 fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. The selected games G_1 are chosen along the lines $D_g^1 = -0.9$ and $D_r^1 = -0.9$ in the dilemma space. The fixed game G_2 is represented by a filled purple square, while the selected games G_1 are represented by open squares. The black curves show the trajectories of the effective game $G(\rho(t))$ in the dilemma space. Red crosses mark the final games reached by the Monte Carlo simulations, while cyan crosses indicate the stationary games predicted by the graph-replicator equation given by Eq. (4.7). Simulations are performed on Erdős–Rényi networks with $N = 40000$ nodes and expected average degree $\langle k \rangle = 4$, using degree-normalised payoffs and Fermi updating with $\beta = 0.8$. For each pair (G_1, G_2) , simulations are initialised from $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. The final cooperation level is estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. The figure shows that chimera-game regimes are also present on Erdős–Rényi networks and that the analytical predictions excellently agree with the Monte Carlo simulations. *Original simulation by the author.*

where the dynamics otherwise converges to full defection, for example along the line $D_g^1 = -0.9$ for sufficiently large positive values of D_r^1 . This effect is due to the presence of isolated nodes, and possibly small disconnected components, in the Erdős–Rényi network. If isolated nodes are initialised as cooperators, they have no interaction partners and therefore cannot change strategy. More generally, disconnected components evolve independently from the rest of the network and can

therefore preserve small local cooperation levels. As a result, the simulated cooperation level may not converge exactly to zero, and small discrepancies with the analytical prediction can also appear in regimes with intermediate cooperation. The corresponding time series of the cooperation level are reported in Appendix C.

5.1.2 Watts–Strogatz Networks

We then consider Watts–Strogatz networks, whose main structural properties have been discussed in Sect. 3.4.2. As in the Erdős–Rényi case, we fix the network size to $N = 40000$ and choose the average degree to be $\langle k \rangle = 4$. More precisely, we start from a regular ring in which each node is connected to its two nearest neighbours on each side, and rewire the links with probability $p_{\text{rew}} = 0.1$. As analytical benchmark, we use the same graph-replicator equation adopted in the previous cases, namely Eq. (4.7), which corresponds to a regular graph of degree $k = 4$ with $G_2 = G(1, 1)$. As in the Erdős–Rényi case, the results obtained with total and degree-normalised payoffs are similar. In the following, we report only the degree-normalised case. Strategy updating follows the Fermi rule with selection intensity $\beta = 0.8$. The simulations follow the numerical protocol described at the beginning of this section.

The corresponding systematic scan is shown in Fig. 5.2. Chimera-game regimes are again observed. In particular, along the line $D_g^1 = -0.9$, they appear for $-0.9 < D_r^1 < 0.8$, since the equilibrium game reached by the population lies in the Stag Hunt region while the endogenous-feedback dynamics sustains a non-zero level of cooperation. This differs from the corresponding fixed-game dynamics on Watts–Strogatz networks, discussed in Sect. 3.4.2, where Stag Hunt games lead to absorbing outcomes.

Compared with the Erdős–Rényi case, however, visible discrepancies now emerge between the Monte Carlo simulations and the analytical prediction. The agreement remains good when G_1 is a Harmony Game, and is still reasonably accurate when G_1 is a Snowdrift Game. By contrast, the most relevant discrepancies occur when G_1 lies in the Stag Hunt region. In some of these cases, the analytical prediction yields a final cooperation level equal to zero, whereas the simulations reach a stationary cooperation level significantly different from zero. As in the square-lattice case, this discrepancy is driven by clustering and by the formation of cooperative domains, which allow cooperators to survive locally despite the invasion pressure exerted by defectors. This effect will be analysed more carefully in the following section. The full time evolution of the cooperation level corresponding to the different choices of G_1 shown in Fig. 5.2 is reported in Appendix C.

To visualise the comparison between Monte Carlo simulations and analytical predictions more directly, we also plot the final cooperation level as a function of the varying parameter along the two scan lines. Figure 5.3 reports the comparison along the line $D_g^1 = -0.9$, where D_r^1 is varied, while Fig. 5.4 shows the corresponding comparison along the line $D_r^1 = -0.9$, where D_g^1 is varied. Blue circles represent the final cooperation level obtained from Monte Carlo simulations, along with the corresponding standard deviation. Orange squares represent stable analytical equilibria, while green squares represent unstable analytical equilibria.




Figure 5.2: Systematic scan of the endogenous-feedback model on Watts–Strogatz networks, with G_2 fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. The selected games G_1 are chosen along the lines $D_g^1 = -0.9$ and $D_r^1 = -0.9$ in the dilemma space. The fixed game G_2 is represented by a filled purple square, while the selected games G_1 are represented by open squares. The black curves show the trajectories of the effective game $G(\rho(t))$ in the dilemma space. Red crosses mark the final games reached by the Monte Carlo simulations, while cyan crosses indicate the stationary games predicted by the graph-replicator equation given by Eq. (4.7). Simulations are performed on Watts–Strogatz networks with $N = 40000$ nodes, average degree $\langle k \rangle = 4$, rewiring probability $p_{\text{rew}} = 0.1$, degree-normalised payoffs, and Fermi updating with $\beta = 0.8$. For each pair (G_1, G_2) , simulations are initialised from $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. The final cooperation level is estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. The figure shows the presence of chimera-game regimes and reveals visible discrepancies with the analytical prediction, especially when G_1 lies in the Stag Hunt region. *Original simulation by the author.*

5.1.3 Barabási–Albert Networks

Finally, we consider Barabási–Albert networks, whose main structural properties have been discussed in Sect. 3.4.3. As in the previous cases, the network size is fixed to $N = 40000$. The networks are generated with attachment parameter $m = 2$, so that the average degree is approximately $\langle k \rangle \simeq 2m = 4$. In contrast with




Figure 5.3: Comparison between Monte Carlo simulations and analytical predictions for the endogenous-feedback model on Watts–Strogatz networks with $G_2 = G(1, 1)$ fixed as a Prisoner’s Dilemma. The figure shows the final cooperation level along the line $D_g^1 = -0.9$ as D_r^1 is varied. Blue circles indicate the asymptotic cooperation level obtained from Monte Carlo simulations. Error bars represent the standard deviation of the final cooperation level across the 25 realisations considered (5 network generations and 5 Monte Carlo simulations for each generation), where each final cooperation level is estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. Orange squares indicate stable analytical equilibria, while green squares indicate unstable ones, obtained from Eq. (4.7). Simulations are performed with $\rho_0 = 0.5$, $N = 40000$, rewiring probability $p_{\text{rew}} = 0.1$, average degree $\langle k \rangle = 4$, degree-normalised payoffs and Fermi updating with $\beta = 0.8$. The largest discrepancies appear in the Stag Hunt region, where clustering allows cooperators to survive more effectively than predicted by the pair-approximation benchmark. *Original illustration by the author.*

Erdős–Rényi and Watts–Strogatz networks, Barabási–Albert networks are strongly heterogeneous. A small number of nodes, the hubs, have a degree much larger than that of typical nodes. For this reason, the choice between degree-normalised and accumulated payoffs becomes dynamically relevant and it can substantially affect the evolutionary dynamics. We therefore analyse both cases separately.

We first consider degree-normalised payoffs. In this case, the payoff of each node is divided by its degree, so that hubs do not acquire a direct payoff advantage simply because they have more interaction partners. As an analytical benchmark, we use the same graph-replicator equation adopted in the previous sections, namely Eq. (4.7). Strategy updating follows the Fermi rule with selection intensity $\beta = 0.8$. The simulations follow the numerical protocol described at the beginning of this section.

Figure 5.5 shows that, in the degree-normalised case, chimera-game regimes are again observed. In particular, along the line $D_g^1 = -0.9$, they appear for $-0.9 < D_r^1 < 0.3$. In this range, the effective game reached at equilibrium lies in the Stag Hunt region, while the endogenous-feedback dynamics sustains a non-zero level of




Figure 5.4: Comparison between Monte Carlo simulations and analytical predictions for the endogenous-feedback model on Watts–Strogatz networks with $G_2 = G(1, 1)$ fixed as a Prisoner’s Dilemma. The figure shows the final cooperation level along the line $D_r^1 = -0.9$ as D_g^1 is varied. Blue circles indicate the asymptotic cooperation level obtained from Monte Carlo simulations. Error bars represent the standard deviation of the final cooperation level across the 25 realisations considered (5 network generations and 5 Monte Carlo simulations for each generation), where each final cooperation level is estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. Orange squares indicate stable analytical equilibria, while green squares indicate unstable ones, obtained from Eq. (4.7). Simulations are performed with $\rho_0 = 0.5$, $N = 40000$, rewiring probability $p_{\text{rew}} = 0.1$, average degree $\langle k \rangle = 4$, degree-normalised payoffs and Fermi updating with $\beta = 0.8$. Along this scan, the main discrepancy appears in the Snowdrift region, corresponding to $D_g^1 > 0$. However, the disagreement remains limited, and the analytical benchmark still captures the qualitative behaviour of the dynamics. *Original illustration by the author.*

cooperation. The agreement between Monte Carlo simulations and the analytical prediction is excellent. This indicates that, once payoffs are normalised by degree, the heterogeneity of the Barabási–Albert network does not significantly affect the validity of the pair-approximation benchmark. The full time evolution of the cooperation level corresponding to the different choices of G_1 shown in Fig. 5.5 is reported in Appendix C.

We now turn to the accumulated-payoff case, where payoffs are no longer divided by the degree of the nodes. As a consequence, high-degree nodes accumulate larger payoffs since they take part in more interactions. In this case, Eq. (4.7) no longer provides an appropriate analytical benchmark. That equation is derived for regular graphs, where all nodes have the same degree and therefore have comparable structural weight. This assumption is not valid on strongly heterogeneous networks such as Barabási–Albert networks, where hubs can dominate the evolutionary dynamics. For this reason, we do not include the analytical prediction in the dilemma-space plot. Indeed, the results shown in Fig. 5.6 are both quantitatively and qualitatively different from the degree-normalised case discussed above. As we can see, overall,




Figure 5.5: Systematic scan of the endogenous-feedback model on Barabási-Albert networks with degree-normalised payoffs, with G_2 fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. The selected games G_1 are chosen along the lines $D_g^1 = -0.9$ and $D_r^1 = -0.9$ in the dilemma space. The fixed game G_2 is represented by a filled purple square, while the selected games G_1 are represented by open squares. The black curves show the trajectories of the effective game $G(\rho(t))$ in the dilemma space. Red crosses mark the final games reached by the Monte Carlo simulations, while cyan crosses indicate the stationary games predicted by the graph-replicator equation given by Eq. (4.7). Simulations are performed on Barabási-Albert networks with $N = 40000$ nodes and attachment parameter $m = 2$, corresponding to average degree $\langle k \rangle \simeq 4$. Payoffs are degree-normalised and strategies are updated according to the Fermi rule with $\beta = 0.8$. For each pair (G_1, G_2) , simulations are initialised from $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. The final cooperation level is estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. The figure shows that, when payoffs are normalised by degree, chimera-game regimes are present and the analytical prediction remains in excellent agreement with the Monte Carlo simulations. *Original simulation by the author.*

the final cooperation level is considerably higher. This is consistent with the role of hubs in heterogeneous networks, discussed in Sect. 3.4.3. When payoffs are accumulated, hubs interact with many more neighbours than typical nodes and can therefore acquire a strong evolutionary advantage. If a hub is occupied by a cooperator and is surrounded by sufficiently many cooperators, it accumulates a large




Figure 5.6: Systematic scan of the endogenous-feedback model on Barabási–Albert networks with accumulated payoffs, with G_2 fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. The selected games G_1 are chosen along the lines $D_g^1 = -0.9$ and $D_r^1 = -0.9$ in the dilemma space. The fixed game G_2 is represented by a filled purple square, while the selected games G_1 are represented by open squares. The black curves show the trajectories of the effective game $G(\rho(t))$ in the dilemma space, and red crosses mark the final games reached by the Monte Carlo simulations. No analytical prediction is shown. Simulations are performed on Barabási–Albert networks with $N = 40000$ nodes and attachment parameter $m = 2$, corresponding to average degree $\langle k \rangle \simeq 4$. Payoffs are accumulated over all neighbours, without degree normalisation, and strategies are updated according to the Fermi rule with $\beta = 0.8$. For each pair (G_1, G_2) , simulations are initialised from $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. The final cooperation level is estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. The figure shows a strong enhancement of cooperation with respect to the degree-normalised case and to the other network topologies considered above. Several chimera-game outcomes are observed. The most significant ones occur in the Prisoner’s Dilemma region, where cooperation survives at a non-trivial level despite the effective game favouring defection. *Original simulation by the author.*

payoff through repeated cooperative interactions. Since many neighbouring nodes compare themselves with the hub, the cooperative strategy can then spread efficiently through the network. This creates a positive feedback: cooperation around the hub increases the hub’s payoff, and the hub in turn promotes the spread of coop-

eration. By contrast, a defective hub may initially gain a high payoff by exploiting neighbouring cooperators, but this advantage is self-limiting. If defection spreads to its neighbours, the hub becomes surrounded by defectors and its payoff decreases. For this reason, the dynamics tends to stabilise cooperation on hubs, which then enhances the spread of cooperation across the network.

To assess the presence of chimera-game regimes in this setting, the endogenous-feedback results should be compared with the fixed-game dynamics on Barabási–Albert networks with accumulated payoffs, reported in Fig. 3.33. This comparison shows that chimera-game behaviour is also present in this case, but with a different character from that observed for the previous network topologies. For some choices of G_1 in the Stag Hunt region, the endogenous-feedback dynamics reaches an intermediate level of cooperation, whereas the corresponding fixed-game dynamics would lead to full cooperation. A more distinctive effect appears for several choices of G_1 in the Harmony Game region and for some choices in the Snowdrift region. In these cases, the effective game reached at equilibrium lies in the Prisoner’s Dilemma region, but the population still sustains a non-zero level of cooperation. In the corresponding fixed-game dynamics, these Prisoner’s Dilemma games would lead to full defection. Hence, the endogenous-feedback mechanism extends the possibility of observing cooperation in the Prisoner’s Dilemma region, even in parameter regimes where defection is strongly favoured. In this sense, the feedback dynamics pushes cooperative behaviour beyond what is observed in the fixed-game Barabási–Albert case with accumulated payoffs. Finally, observe that, unlike in the degree-normalised case and in the previous network topologies, the final cooperation level is not always robust with respect to the initial condition. Different values of ρ_0 can lead to different stationary cooperation levels and, consequently, to different final effective games in the dilemma space. The full time evolution of the cooperation level for the pairs shown in Fig. 5.6 is reported in Appendix C.

5.2 Clustering and the Breakdown of Pair Approximation

Starting from Chapter 3, we have seen that the outcome of evolutionary games on structured populations depends on the underlying interaction network. In the fixed-game case, different network topologies can lead to different final cooperation levels, and the agreement between Monte Carlo simulations and the analytical prediction based on pair approximation is not equally accurate for all structures. These effects become much more pronounced when the endogenous-feedback model is introduced. A recurring feature of the cases in which these discrepancies are most visible is the presence of clustering. This is the case for square lattices and Watts–Strogatz networks, where neighbours of the same node are more likely to be connected to each other. Clustering can therefore alter the expected cooperation level by changing the local structure of interactions and the correlations between neighbouring strategies. The aim of this section is therefore to analyse the role of clustering in the evolutionary dynamics. To this end, we use a degree-preserving rewiring procedure, which changes the local structure of the network while keeping the degree of every node

fixed, as shown below. In this way, variations in the dynamics can be attributed to changes in clustering rather than to changes in the degree distribution. We then use this procedure to study how clustering shapes the final cooperation level and how it affects the validity of the pair-approximation prediction used throughout the thesis.

5.2.1 The Degree-Preserving Rewiring Procedure

The degree-preserving rewiring procedure consists in swapping pairs of edges while keeping the degree of every node unchanged. At each step, two links are selected at random. If the selected links are (i, j) and (m, n) , the operation is

$$(i, j), (m, n) \longrightarrow (i, n), (m, j).$$

In other words, the two original edges are swapped: node i is disconnected from node j and connected to node n , while node m is disconnected from node n and connected to node j . Therefore, this operation changes the local connectivity pattern, but it does not modify the degree of any node. Indeed, each of the four nodes involved in the swap loses one neighbour and gains one new neighbour. A schematic example of a single swapping step is shown in Fig. 5.7. By repeating this operation many times




Figure 5.7: Schematic illustration of one step of the degree-preserving edge-swapping procedure. **Left:** the original network contains the two edges (i, j) and (m, n) . **Right:** the two selected edges are removed and replaced by (i, n) and (m, j) . This operation preserves the degree of all nodes involved, since each node loses one link and gains one link. Repeating this procedure over the whole network progressively modifies the local connectivity pattern while leaving the degree sequence unchanged. *Illustration by the author.*

over the whole network, one progressively randomises the original structure while preserving the degree sequence. In this way, global quantities related to node degree are left unchanged, whereas local structural properties, most notably the clustering coefficient, can vary substantially.

In our work, we consider a regular ring lattice as the initial network. The network has $N = 40000$ nodes, and each node is connected to its first and second nearest neighbours along the ring. The initial network is therefore regular with degree

$k = 4$. This choice provides a highly clustered reference structure from which networks with progressively lower clustering can be generated through the procedure described above. To quantify the intensity of the rewiring, we introduce a parameter α , defined as the number of attempted edge swaps per link. If M is the total number of links in the network, the number of attempted swaps is $n_{\text{swap}}(\alpha) = \lfloor \alpha M \rfloor$, where $\lfloor x \rfloor$ denotes the integer part of x . Thus, $\alpha = 0$ corresponds to the original regular ring lattice, while larger values of α correspond to increasingly randomised networks. Since every accepted swap preserves all node degrees, all networks generated by this procedure remain 4-regular. By increasing α , the short loops and local ordering of the initial ring lattice are progressively destroyed. As a consequence, the clustering coefficient decreases, while the degree sequence remains exactly fixed. This makes the procedure particularly useful for isolating the effect of clustering on the evolutionary dynamics.

To build the family of networks used to isolate the role of clustering, we apply the degree-preserving edge-swapping procedure for several values of the rewiring intensity α . These values are chosen so as to sample both the initial high-clustering regime and the low-clustering plateau reached after extensive randomisation. For each value of α , we generate 5 independent network realisations and compute the corresponding average clustering coefficient. Fig. 5.8 shows the behaviour of the clustering coefficient as a function of α . The figure shows that this procedure provides a controlled way to tune the clustering coefficient. For small values of α , the network retains a high level of local clustering. As α increases, local triangles are progressively destroyed and the clustering coefficient decreases. For large values of α , the clustering coefficient reaches a low plateau, corresponding to a randomised regular network with clustering coefficient close to zero. This family of networks is therefore well suited to isolate the role of clustering. The degree of every node is fixed throughout the procedure, while the clustering coefficient is systematically varied.

5.2.2 Effect of Clustering on Cooperation and Analytical Predictions

We now analyse how clustering affects both the final cooperation level reached by the system and the accuracy of the analytical prediction based on pair approximation. As discussed in the previous sections, the effect of network structure becomes particularly relevant in the endogenous-feedback model. For this reason, we focus here on this case. We consider the family of networks generated from a regular ring lattice through the degree-preserving rewiring procedure described in Sect. 5.2.1. The rewiring intensities α are the same as in Fig. 5.8. They span a range from high to low average clustering, allowing us to study the effect of clustering while keeping the degree of every node fixed. The second reference game is fixed as a Prisoner's Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$. For the first reference game, we consider Stag Hunt games along the line $D_g^1 = -0.9$, namely

$$D_g^1 = -0.9, \quad D_r^1 \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}.$$




Figure 5.8: Average clustering coefficient as a function of the rewiring intensity α for degree-preserving rewired networks. The initial network is a regular ring lattice with $N = 40000$ nodes and degree $k = 4$, where each node is connected to its first and second nearest neighbours. For each value of α , 5 independent network realisations are generated. Points indicate the mean clustering coefficient across realisations, while error bars represent the corresponding standard deviation. The error bars are practically invisible because the standard deviations across network realisations are very small. Increasing α progressively randomises the initial network and reduces the clustering coefficient, while the degree of every node remains fixed. *Original simulation by the author.*

The simulations are initialised with cooperation level $\rho_0 = 0.5$. This choice is sufficient for the present analysis since the final cooperation level is robust with respect to different initial conditions. For each value of α , we generate 5 independent network realisations. For each network realisation and each pair (G_1, G_2) , we perform 5 independent Monte Carlo simulations. Thus, for each value of α and each pair (G_1, G_2) , the results are averaged over 25 simulations. The final cooperation level is estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence.

Figure 5.9 reports the final cooperation level as a function of the average clustering coefficient. Each curve corresponds to a different choice of G_1 , while the dashed horizontal line of the same colour indicates the corresponding analytical prediction obtained from the replicator equation for regular graph given by Eq. (4.7). Horizontal error bars represent the standard deviation of the clustering coefficient across network realisations, while vertical error bars represent the standard deviation of $\rho(t)$ computed over the same final 20% time window used to estimate the mean final cooperation level.

In this controlled setting, Fig. 5.9 shows that clustering has a significant effect on the final cooperation level reached by the system. The effect is especially clear for the pairs $(D_g^1, D_r^1) = (-0.9, 0.4)$ and $(D_g^1, D_r^1) = (-0.9, 0.5)$. For low values




Figure 5.9: Final cooperation level as a function of the average clustering coefficient for the endogenous-feedback model on degree-preserving rewired regular networks. The initial network is a regular ring lattice with $N = 40000$ nodes and degree $k = 4$. The second reference game is fixed as a Prisoner's Dilemma, $G_2 = G(1, 1)$, while the first reference game is chosen in the Stag Hunt region along the line $D_g^1 = -0.9$, with $D_r^1 \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$. Each curve corresponds to one value of D_r^1 . For each value of the rewiring intensity α , 5 independent network realisations are generated, and 5 independent Monte Carlo simulations are performed on each realisation. The simulations are initialised at $\rho_0 = 0.5$, and the final cooperation level is estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. Dashed horizontal lines indicate the corresponding analytical predictions obtained from the graph-replicator equation in Eq. (4.7). Horizontal error bars represent the standard deviation of the clustering coefficient across network realisations, while vertical error bars represent the standard deviation of $\rho(t)$ computed over the same final 20% time window used to estimate the mean final cooperation level. The figure shows that the final cooperation level increases progressively with the clustering coefficient. At the same time, the discrepancy between Monte Carlo simulations and analytical predictions becomes larger as clustering increases, showing the progressive breakdown of the pair-approximation description. *Original simulation by the author.*

of the clustering coefficient, the system reaches full defection, in agreement with the analytical prediction. However, as clustering increases, cooperation appears above a threshold-like value of the clustering coefficient and remains stable. For sufficiently clustered networks, the final cooperation level reaches values of the order of 20%. This behaviour is consistent with the mechanism discussed in Sect. 4.2.3: local clustering allows cooperators to form spatially organised cooperative domains, which can protect them from invasion by defectors. More generally, all curves display an increasing trend as the clustering coefficient grows. This indicates that, in the

parameter region considered here, clustering favours the survival of cooperation. The effect is not only quantitative. For some values of G_1 , increasing clustering changes the qualitative outcome of the dynamics, from full defection to stable coexistence between cooperators and defectors.

The same figure also shows how clustering affects the accuracy of the analytical prediction obtained from pair approximation. For small clustering coefficients, the Monte Carlo simulations are close to the analytical prediction. This is expected because highly rewired regular networks are locally closer to random regular graphs, for which the pair-approximation description is more appropriate. As the clustering coefficient increases, however, the discrepancy between simulations and analytical predictions becomes larger. Thus, clustering is one of the structural mechanisms responsible for the breakdown of the pair approximation. The clearest examples are $(D_g^1, D_r^1) = (-0.9, 0.4)$ and $(-0.9, 0.5)$. At low clustering, both the simulations and the analytical benchmark predict full defection. As clustering increases, however, the simulations converge to a stable non-zero cooperation level, whereas the pair approximation still predicts full defection. The disagreement is therefore qualitative: clustering allows cooperation to persist in regimes where the analytical approximation predicts its extinction. This result confirms that the pair approximation does not fully capture the higher-order local correlations generated by clustered networks. In particular, triangles and short loops favour the formation of cooperative domains, whose stability depends on spatial correlations beyond pair-level information. Consequently, the analytical benchmark progressively breaks down as clustering increases.

The results of this chapter complete the analysis of the endogenous-feedback model on structured populations. They show that chimera-game regimes can survive not only on regular lattices, but also on complex network topologies. At the same time, the comparison with the pair-approximation benchmark reveals that network structure can strongly affect both the final cooperation level and the accuracy of the analytical prediction. A central role was played by clustering. By using a degree-preserving rewiring procedure, we isolated the effect of clustering from that of the degree distribution. The results showed that increasing clustering can sustain cooperation in regimes where low-clustering networks lead to full defection. At the same time, the higher-order local correlations generated by clustering are not captured by pair approximation and therefore lead to a progressive breakdown of the analytical prediction. These results are brought together and discussed in the conclusions.

Conclusions

This thesis investigated the emergence of cooperation, one of the fundamental problems in the study of interacting populations and complex systems. The phenomenon was approached from a statistical-mechanics perspective, treating populations of interacting agents as nonequilibrium systems whose macroscopic regularities emerge from microscopic update rules, stochasticity, local interactions and network structure. Within this framework, evolutionary game theory and coevolutionary game theory were interpreted as complementary descriptions of adaptive many-agent systems, in which structure, fluctuations and feedback jointly determine the collective behaviour of the population.

The first main contribution of the thesis was to show that chimera-game regimes survive in structured populations, where interactions are local and topologically constrained. When the endogenous-feedback model was implemented on square lattices and complex networks, the system was still able to reach stable coexistence between cooperators and defectors, even in parameter regions where the corresponding fixed-game dynamics would not predict such stability. This suggests that chimera-game regimes are not an artefact of the well-mixed setting, but can survive the transition to more realistic interaction structures.

The second main contribution concerned the spatial dynamics of cooperative clusters, namely domains of neighbouring cooperators. The analysis revealed the presence of two distinct dynamical time scales. On a relatively short time scale, the global fraction of cooperators approached a quasi-stationary value. On a longer time scale, however, the spatial organisation of the population could continue to evolve through the coarsening, merging or fragmentation of cooperative regions. This behaviour is reminiscent of phase-ordering dynamics in spin systems: the global density may become nearly stationary while the spatial configuration continues to reorganise through the motion and coarsening of domains [72].

The analysis also revealed sharp changes in cluster morphology across the different games reached by the feedback dynamics. In some regimes, cooperators remained fragmented into small and vulnerable domains, whereas in others they organised into compact domains able to resist invasion by defectors. This provides a possible rule of thumb for inferring the underlying effective interactions from the emerging spatial patterns. Compact cooperative domains are naturally associated with coordination-like dynamics, as in Stag Hunt games, whereas more fragmented cooperative structures are indicative of anti-coordination dynamics, as in Snowdrift games. This connection between spatial morphology and effective game type may be particularly relevant for experimental evolution in microbial populations, where

the spatial organisation of cooperative strains can often be observed directly [73].

The third main contribution concerned the role of clustering, understood here as a graph-theoretic property of the interaction network. Using a degree-preserving rewiring procedure, we isolated the effect of clustering from that of the degree distribution. This made it possible to show that clustering is not a secondary structural detail, but a mechanism that can qualitatively change the outcome of the endogenous-feedback dynamics. In the parameter regimes analysed, low-clustering networks could drive the system to full defection. Increasing clustering, while keeping the degree of every node fixed, allowed cooperation to emerge and reach finite stationary levels. Thus, clustering can sustain cooperation in regimes where low-clustering networks with the same degree sequence cannot.

At the same time, clustering was responsible for the breakdown of the pair approximation prediction. For sufficiently randomised networks, the analytical prediction based on pair approximation remained accurate. As clustering increased, however, the discrepancy between simulations and theory grew systematically. This shows that the failure of the analytical description is not simply due to the degree distribution, which was kept fixed, but to the higher-order local correlations generated by clustered network structure. The degree-preserving rewiring analysis therefore identifies clustering as a key structural mechanism controlling both the level of cooperation and the validity of the analytical approximation.

Overall, the thesis contributes to the statistical physics of adaptive many-agent systems by showing how endogenous feedback, local interactions, stochastic dynamics and network structure combine to produce cooperation. The results highlight a rich phenomenology in structured coevolutionary games, including the survival of chimera-game regimes, the separation between global density relaxation and spatial reorganisation, the emergence of distinct cooperative cluster morphologies, and the role of clustering in both sustaining cooperation and inducing the breakdown of pair approximation.

The work carried out in this thesis is expected to lead to two research articles. The first will report the emergence of chimera-game regimes in structured populations, showing that the regimes identified in the well-mixed setting of [3] survive when the population is described by more realistic interaction structures, where contacts are local and topologically constrained. The second will focus on the spatial dynamics of cooperative clusters, with particular attention to time-scale separation, domain coarsening and the relation between cluster morphology and the game reached by the feedback dynamics.

Several directions remain open for future research. A first direction is the analytical characterisation of the two dynamical time scales observed in the spatial dynamics, together with a more systematic theory of cluster morphology. Such a theory would help explain why the global cooperation level can reach stationarity before the spatial configuration has completed its reorganisation. It would also make it possible to relate the stability of cooperation to measurable spatial quantities, such as domain size, interface density, cluster-size distributions and coarsening rates. In this way, one could clarify how the morphology of cooperative domains depends on the game reached by the feedback dynamics and how this morphology controls the resistance of cooperation to defector invasion.

A second direction is the study of local adaptive games, in which the endogenous feedback does not depend on the global fraction of cooperators, but on the local fraction of cooperators in the neighbourhood of each agent. This would make the game itself spatially heterogeneous, allowing different regions of the same population to experience different strategic environments. A third possible direction is the study of stochastic annealed feedback, where the local game is sampled probabilistically from the local state of the population rather than being determined by a deterministic interpolation.

More broadly, this thesis suggests that a statistical-mechanics approach to coevolutionary games can reveal a rich phenomenology in endogenous-feedback models, including coexistence regimes, time-scale separation, changes in spatial organisation and the breakdown of analytical approximations. At present, however, a general statistical-mechanics theory of coevolutionary games is still largely missing. Developing such a framework would be an important step towards a more systematic understanding of adaptive populations in which strategies, interaction structure and strategic incentives coevolve.

Bibliography

- [1] Robert Axelrod and William D. Hamilton. “The Evolution of Cooperation”. In: *Science* 211.4489 (1981), pp. 1390–1396. DOI: [10.1126/science.7466396](https://doi.org/10.1126/science.7466396).
- [2] Martin A. Nowak. “Five Rules for the Evolution of Cooperation”. In: *Science* 314.5805 (2006), pp. 1560–1563. DOI: [10.1126/science.1133755](https://doi.org/10.1126/science.1133755).
- [3] Federico Maria Quetti et al. “Endogenous Feedback Reshapes the Stability of Cooperation”. In: *arXiv preprint arXiv:2506.20603* (2025). DOI: [10.48550/arXiv.2506.20603](https://doi.org/10.48550/arXiv.2506.20603).
- [4] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. Cambridge, MA: MIT Press, 1994.
- [5] Colin F. Camerer. *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton: Princeton University Press, 2003.
- [6] Ernst Fehr and Simon Gächter. “Altruistic Punishment in Humans”. In: *Nature* 415.6868 (2002), pp. 137–140. DOI: [10.1038/415137a](https://doi.org/10.1038/415137a).
- [7] John Maynard Smith. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press, 1982.
- [8] Martin A Nowak. *Evolutionary Dynamics: Exploring the Equations of Life*. Harvard University Press, 2006.
- [9] György Szabó and Gábor Fáth. “Evolutionary Games on Graphs”. In: *Physics Reports* 446.4–6 (2007), pp. 97–216. DOI: [10.1016/j.physrep.2007.04.004](https://doi.org/10.1016/j.physrep.2007.04.004).
- [10] Arne Traulsen, Jens Christian Claussen, and Christoph Hauert. “Coevolutionary Dynamics: From Finite to Infinite Populations”. In: *Physical Review Letters* 95.23 (2005), p. 238701. DOI: [10.1103/PhysRevLett.95.238701](https://doi.org/10.1103/PhysRevLett.95.238701).
- [11] György Szabó and Csaba T öke. “Evolutionary prisoner’s dilemma game on a square lattice”. In: *Phys. Rev. E* 58 (1998), pp. 69–73. DOI: [10.1103/PhysRevE.58.69](https://doi.org/10.1103/PhysRevE.58.69).
- [12] Minus van Baalen. “Pair Approximations for Different Spatial Geometries”. In: *The Geometry of Ecological Interactions: Simplifying Spatial Complexity*. Ed. by Ulf Dieckmann, Richard Law, and Johan A. J. Metz. Cambridge Studies in Adaptive Dynamics. Cambridge: Cambridge University Press, 2000, pp. 359–387.
- [13] Hisashi Ohtsuki and Martin A. Nowak. “The replicator equation on graphs”. In: *Journal of Theoretical Biology* 243.1 (2006), pp. 86–97. DOI: [10.1016/j.jtbi.2006.06.004](https://doi.org/10.1016/j.jtbi.2006.06.004).

- [14] N. G. van Kampen. *Stochastic Processes in Physics and Chemistry*. 3rd ed. Amsterdam: North-Holland, 2007.
- [15] Hannes Risken. *The Fokker–Planck Equation: Methods of Solution and Applications*. 2nd ed. Berlin: Springer, 1996.
- [16] Michael N. Barber. “Finite-Size Scaling”. In: *Phase Transitions and Critical Phenomena*. Ed. by Cyril Domb and Joel L. Lebowitz. Vol. 8. London: Academic Press, 1983, pp. 145–266.
- [17] M. M. de Oliveira and Ronald Dickman. “How to Simulate the Quasistationary State”. In: *Physical Review E* 71.1 (2005), p. 016129. DOI: [10.1103/PhysRevE.71.016129](https://doi.org/10.1103/PhysRevE.71.016129).
- [18] Matjaž Perc and Attila Szolnoki. “Coevolutionary Games: A Mini Review”. In: *BioSystems* 99.2 (2010), pp. 109–125. DOI: [10.1016/j.biosystems.2009.10.003](https://doi.org/10.1016/j.biosystems.2009.10.003).
- [19] Joshua S. Weitz et al. “An Oscillating Tragedy of the Commons in Replicator Dynamics with Game-Environment Feedback”. In: *Proceedings of the National Academy of Sciences* 113.47 (2016), E7518–E7525. DOI: [10.1073/pnas.1604096113](https://doi.org/10.1073/pnas.1604096113).
- [20] Andrew R. Tilman, Joshua B. Plotkin, and Erol Akçay. “Evolutionary Games with Environmental Feedbacks”. In: *Nature Communications* 11 (2020), p. 915. DOI: [10.1038/s41467-020-14531-6](https://doi.org/10.1038/s41467-020-14531-6).
- [21] Daniel Friedman. “Towards Evolutionary Game Models of Financial Markets”. In: *Quantitative Finance* 1.1 (2001), pp. 177–185. DOI: [10.1080/713665544](https://doi.org/10.1080/713665544).
- [22] Marta C. González, César A. Hidalgo, and Albert-László Barabási. “Understanding Individual Human Mobility Patterns”. In: *Nature* 453 (2008), pp. 779–782. DOI: [10.1038/nature06958](https://doi.org/10.1038/nature06958).
- [23] Peter Kollock. “The Economies of Online Cooperation: Gifts and Public Goods in Cyberspace”. In: *Communities in Cyberspace*. Ed. by Marc A. Smith and Peter Kollock. London: Routledge, 1999, pp. 220–239.
- [24] Taisei Kaizoji. “Speculative Bubbles and Crashes in Stock Markets: An Interacting-Agent Model of Speculative Activity”. In: *Physica A: Statistical Mechanics and its Applications* 287.3–4 (2000), pp. 493–506. DOI: [10.1016/S0378-4371\(00\)00388-5](https://doi.org/10.1016/S0378-4371(00)00388-5).
- [25] Mario Bellia. “High-Frequency Market Making: Liquidity Provision, Adverse Selection, and Competition”. In: *Adverse Selection, and Competition* (2017). November 15, 2017. DOI: [10.2139/ssrn.3074313](https://doi.org/10.2139/ssrn.3074313).
- [26] Albert-László Barabási. *Network Science*. Cambridge: Cambridge University Press, 2016.
- [27] Mark E. J. Newman. *Networks: An Introduction*. Oxford: Oxford University Press, 2010.
- [28] Vito Latora, Vincenzo Nicosia, and Giovanni Russo. *Complex Networks: Principles, Methods and Applications*. Cambridge: Cambridge University Press, 2017.

-
- [29] R. K. Pathria and Paul D. Beale. *Statistical Mechanics*. 3rd ed. Amsterdam: Elsevier/Academic Press, 2011.
 - [30] Philip W. Anderson. “More Is Different”. In: *Science* 177.4047 (1972), pp. 393–396. DOI: [10.1126/science.177.4047.393](https://doi.org/10.1126/science.177.4047.393).
 - [31] P. M. Chaikin and T. C. Lubensky. *Principles of Condensed Matter Physics*. Cambridge: Cambridge University Press, 1995.
 - [32] Nigel Goldenfeld. *Lectures on Phase Transitions and the Renormalization Group*. Reading, MA: Addison-Wesley, 1992.
 - [33] Tamás Vicsek et al. “Novel Type of Phase Transition in a System of Self-Driven Particles”. In: *Physical Review Letters* 75.6 (1995), pp. 1226–1229. DOI: [10.1103/PhysRevLett.75.1226](https://doi.org/10.1103/PhysRevLett.75.1226).
 - [34] Dirk Helbing. “Traffic and Related Self-Driven Many-Particle Systems”. In: *Reviews of Modern Physics* 73.4 (2001), pp. 1067–1141. DOI: [10.1103/RevModPhys.73.1067](https://doi.org/10.1103/RevModPhys.73.1067).
 - [35] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. “Statistical Physics of Social Dynamics”. In: *Reviews of Modern Physics* 81.2 (2009), pp. 591–646. DOI: [10.1103/RevModPhys.81.591](https://doi.org/10.1103/RevModPhys.81.591).
 - [36] Rosario N. Mantegna and H. Eugene Stanley. *An Introduction to Econophysics: Correlations and Complexity in Finance*. Cambridge: Cambridge University Press, 2000.
 - [37] Robert M. May. *Stability and Complexity in Model Ecosystems*. Princeton: Princeton University Press, 1973.
 - [38] Stuart A. West et al. “Social Evolution Theory for Microorganisms”. In: *Nature Reviews Microbiology* 4.8 (2006), pp. 597–607. DOI: [10.1038/nrmicro1461](https://doi.org/10.1038/nrmicro1461).
 - [39] Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press, 1990.
 - [40] Samuel Bowles and Herbert Gintis. *A Cooperative Species: Human Reciprocity and Its Evolution*. Princeton: Princeton University Press, 2011.
 - [41] Robyn M. Dawes. “Social Dilemmas”. In: *Annual Review of Psychology* 31 (1980), pp. 169–193. DOI: [10.1146/annurev.ps.31.020180.001125](https://doi.org/10.1146/annurev.ps.31.020180.001125).
 - [42] Anatol Rapoport and Albert M. Chammah. *Prisoner’s Dilemma: A Study In Conflict And Cooperation*. University of Michigan Press, 1965.
 - [43] W. D. Hamilton. “The Genetical Evolution of Social Behaviour. I”. In: *Journal of Theoretical Biology* 7.1 (1964), pp. 1–16. DOI: [10.1016/0022-5193\(64\)90038-4](https://doi.org/10.1016/0022-5193(64)90038-4).
 - [44] Stefan Wuchty, Benjamin F. Jones, and Brian Uzzi. “The Increasing Dominance of Teams in Production of Knowledge”. In: *Science* 316.5827 (2007), pp. 1036–1039. DOI: [10.1126/science.1136099](https://doi.org/10.1126/science.1136099).
 - [45] Robert L. Trivers. “The Evolution of Reciprocal Altruism”. In: *The Quarterly Review of Biology* 46.1 (1971), pp. 35–57. DOI: [10.1086/406755](https://doi.org/10.1086/406755).
 - [46] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.

-
- [47] Tony Ashworth. *Trench Warfare, 1914–1918: The Live and Let Live System*. London: Macmillan, 1980.
 - [48] Kevin Leyton-Brown and Yoav Shoham. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Synthesis Lectures on Artificial Intelligence and Machine Learning. San Rafael, CA: Morgan & Claypool Publishers, 2008.
 - [49] John Maynard Smith and George R. Price. “The Logic of Animal Conflict”. In: *Nature* 246.5427 (1973), pp. 15–18. DOI: [10.1038/246015a0](https://doi.org/10.1038/246015a0).
 - [50] P. A. P. Moran. “Random Processes in Genetics”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 54.1 (1958), pp. 60–71. DOI: [10.1017/S0305004100033193](https://doi.org/10.1017/S0305004100033193).
 - [51] Josef Hofbauer and Karl Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge: Cambridge University Press, 1998.
 - [52] Jun Tanimoto and Hiroki Sagara. “Relationship between dilemma occurrence and the existence of a weakly dominant strategy in a two-player symmetric game”. In: *BioSystems* 90.1 (2007), pp. 105–114. DOI: [10.1016/j.biosystems.2006.07.005](https://doi.org/10.1016/j.biosystems.2006.07.005).
 - [53] Zhen Wang et al. “Universal scaling for the dilemma strength in evolutionary games”. In: *Physics of Life Reviews* 14 (2015), pp. 1–30. DOI: [10.1016/j.plrev.2015.04.033](https://doi.org/10.1016/j.plrev.2015.04.033).
 - [54] Roy J. Glauber. “Time-Dependent Statistics of the Ising Model”. In: *Journal of Mathematical Physics* 4.2 (1963), pp. 294–307. DOI: [10.1063/1.1703954](https://doi.org/10.1063/1.1703954).
 - [55] Jorge M. Pacheco, Arne Traulsen, and Martin A. Nowak. “Active Linking in Evolutionary Games”. In: *Journal of Theoretical Biology* 243.3 (2006), pp. 437–443. DOI: [10.1016/j.jtbi.2006.06.027](https://doi.org/10.1016/j.jtbi.2006.06.027).
 - [56] Alessandro Tavoni, Maja Schlüter, and Simon Levin. “The Survival of the Conformist: Social Pressure and Renewable Resource Management”. In: *Journal of Theoretical Biology* 299 (2012), pp. 152–161. DOI: [10.1016/j.jtbi.2011.07.003](https://doi.org/10.1016/j.jtbi.2011.07.003).
 - [57] Attila Szolnoki and György Szabó. “Cooperation Enhanced by Inhomogeneous Activity of Teaching for Evolutionary Prisoner’s Dilemma Games”. In: *Europhysics Letters* 77.3 (2007), p. 30004. DOI: [10.1209/0295-5075/77/30004](https://doi.org/10.1209/0295-5075/77/30004).
 - [58] Zhi-Xi Wu, Zhihai Rong, and Petter Holme. “Diversity of Reproduction Time Scale Promotes Cooperation in Spatial Prisoner’s Dilemma Games”. In: *Physical Review E* 80.3 (2009), p. 036106. DOI: [10.1103/PhysRevE.80.036106](https://doi.org/10.1103/PhysRevE.80.036106).
 - [59] Luo-Luo Jiang, Matjaž Perc, and Bing-Hong Wang. “Role of Adaptive Migration in Promoting Cooperation in Spatial Games”. In: *Physical Review E* 81.3 (2010), p. 036108. DOI: [10.1103/PhysRevE.81.036108](https://doi.org/10.1103/PhysRevE.81.036108).
 - [60] Attila Szolnoki et al. “Impact of Aging on the Evolution of Cooperation in the Spatial Prisoner’s Dilemma Game”. In: *Physical Review E* 80.2 (2009), p. 021901. DOI: [10.1103/PhysRevE.80.021901](https://doi.org/10.1103/PhysRevE.80.021901).

-
- [61] Martin A. Nowak and Karl Sigmund. “Evolution of Indirect Reciprocity by Image Scoring”. In: *Nature* 393.6685 (1998), pp. 573–577. DOI: [10.1038/31225](https://doi.org/10.1038/31225).
 - [62] Jun Zhang et al. “Promotion of Cooperation in Aspiration-Based Spatial Prisoner’s Dilemma Game”. In: *Physica A: Statistical Mechanics and Its Applications* 390.12 (2011), pp. 2258–2266. DOI: [10.1016/j.physa.2011.02.045](https://doi.org/10.1016/j.physa.2011.02.045).
 - [63] Weijun Zeng, Minqiang Li, and Fuzan Chen. “Cooperation in the Evolutionary Iterated Prisoner’s Dilemma Game with Risk Attitude Adaptation”. In: *Applied Soft Computing* 44 (2016), pp. 122–132. DOI: [10.1016/j.asoc.2016.03.025](https://doi.org/10.1016/j.asoc.2016.03.025).
 - [64] Alex McAvoy, Arvind Rao, and Christoph Hauert. “Intriguing Effects of Selection Intensity on the Evolution of Prosocial Behaviors”. In: *PLoS Computational Biology* 17.11 (2021), e1009611. DOI: [10.1371/journal.pcbi.1009611](https://doi.org/10.1371/journal.pcbi.1009611).
 - [65] Paul Resnick et al. “Reputation Systems”. In: *Communications of the ACM* 43.12 (2000), pp. 45–48. DOI: [10.1145/355112.355122](https://doi.org/10.1145/355112.355122).
 - [66] Markus K. Brunnermeier and Lasse Heje Pedersen. “Market Liquidity and Funding Liquidity”. In: *The Review of Financial Studies* 22.6 (2009), pp. 2201–2238. DOI: [10.1093/rfs/hhn098](https://doi.org/10.1093/rfs/hhn098).
 - [67] Duncan J. Watts and Steven H. Strogatz. “Collective Dynamics of Small-World Networks”. In: *Nature* 393 (1998), pp. 440–442.
 - [68] Albert-László Barabási and Réka Albert. “Emergence of Scaling in Random Networks”. In: *Science* 286.5439 (1999), pp. 509–512.
 - [69] Francisco C. Santos and Jorge M. Pacheco. “Scale-Free Networks Provide a Unifying Framework for the Emergence of Cooperation”. In: *Physical Review Letters* 95.9 (2005), p. 098104. DOI: [10.1103/PhysRevLett.95.098104](https://doi.org/10.1103/PhysRevLett.95.098104).
 - [70] Francisco C. Santos, Jorge F. Rodrigues, and Jorge M. Pacheco. “Graph Topology Plays a Determinant Role in the Evolution of Cooperation”. In: *Proceedings of the Royal Society B: Biological Sciences* 273.1582 (2006), pp. 51–55. DOI: [10.1098/rspb.2005.3272](https://doi.org/10.1098/rspb.2005.3272).
 - [71] Francisco C. Santos, Jorge M. Pacheco, and Tom Lenaerts. “Evolutionary Dynamics of Social Dilemmas in Structured Heterogeneous Populations”. In: *Proceedings of the National Academy of Sciences of the United States of America* 103.9 (2006), pp. 3490–3494. DOI: [10.1073/pnas.0508201103](https://doi.org/10.1073/pnas.0508201103).
 - [72] Alan J. Bray. “Theory of Phase-Ordering Kinetics”. In: *Advances in Physics* 43.3 (1994), pp. 357–459. DOI: [10.1080/00018739400101505](https://doi.org/10.1080/00018739400101505).
 - [73] Carey D. Nadell, Knut Drescher, and Kevin R. Foster. “Spatial Structure, Cooperation and Competition in Biofilms”. In: *Nature Reviews Microbiology* 14.9 (2016), pp. 589–600. DOI: [10.1038/nrmicro.2016.84](https://doi.org/10.1038/nrmicro.2016.84).

Appendix A

Endogenous Feedback Simulations in Well-Mixed Populations

In this appendix, we report additional well-mixed simulations for the endogenous-feedback model considered in Chapter 2. These simulations complement the results discussed in the main text by showing further representative choices of the reference games G_1 and G_2 .

A.1 Prisoner's Dilemma as G_2

In this section, we analyse the scenario in which the second reference game is fixed to a Prisoner's Dilemma, namely

$$G_2 = G(D_g^2, D_r^2) = G(1, 1).$$

These simulations complement the systematic scan shown in Fig. 2.6. In particular, we report three representative cases, each corresponding to a different choice of the first reference game G_1 : a Harmony Game, a Snowdrift Game, and a Stag Hunt Game. The corresponding results are shown in Figs. A.1, A.2, and A.3, respectively. The simulation setting is the same as that used in the main text. The population is well mixed and has size $N = 1000$. At each elementary update, a focal player X and a model player Y are selected uniformly at random. Each player interacts with $k = 4$ randomly chosen individuals and accumulates a total payoff computed from the current endogenous payoff matrix

$$G(t) = [1 - \rho(t)]G_1 + \rho(t)G_2.$$

The focal player then adopts the strategy of the model player according to the Fermi rule in Eq. (1.6), with selection intensity fixed to $\beta = 0.2$. For each pair (G_1, G_2) , simulations were initialised from three cooperation levels, $\rho_0 = 0.1, 0.5, 0.9$, and averaged over 20 independent realisations. Time is measured in Monte Carlo sweeps, where one sweep corresponds to N elementary update attempts. In each figure, the horizontal red line indicates the equilibrium cooperation level predicted by the replicator equation. Overall, the three representative cases show excellent agreement between the stochastic simulations and the analytical predictions obtained from the replicator equation.




Figure A.1: Well-mixed endogenous-feedback simulations with G_2 fixed to a Prisoner's Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$, and G_1 chosen in the Harmony Game region. The curves show the average cooperation level $\rho(t)$ as a function of Monte Carlo sweeps for the three initial cooperation levels $\rho_0 = 0.1, 0.5, 0.9$, averaged over 20 independent realisations. The horizontal red line indicates the equilibrium cooperation level predicted by the replicator equation. Different panels are displayed over different Monte Carlo sweep windows. The excellent agreement between the stochastic simulations and the analytical prediction confirms that the replicator equation accurately captures the asymptotic behaviour of the well-mixed endogenous-feedback dynamics. *Original simulation by the author.*

A.2 Stag Hunt as G_2

We now consider the case in which the final reference game is a Stag Hunt, namely

$$G_2 = G(D_g^2, D_r^2) = G(-1, 1).$$


Cooperation Behaviour over Time: SD as G_1


Figure A.2: Well-mixed endogenous-feedback simulations with G_2 fixed to a Prisoner's Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$, and G_1 chosen in the Snowdrift Game region. The curves show the average cooperation level $\rho(t)$ as a function of Monte Carlo sweeps for the three initial cooperation levels $\rho_0 = 0.1, 0.5, 0.9$, averaged over 20 independent realisations. The horizontal red line reports the equilibrium cooperation level predicted by the replicator equation. Different panels are displayed over different Monte Carlo sweep windows. The excellent agreement between the stochastic simulations and the analytical prediction confirms that the replicator equation accurately captures the asymptotic behaviour of the well-mixed endogenous-feedback dynamics. *Original simulation by the author.*

Substituting this choice into Eq. (2.5) gives

$$\Delta\Pi(\rho) = \rho^2 (2 + D_g^1 - D_r^1) + \rho (2D_r^1 - D_g^1 - 1) - D_r^1. \quad (\text{A.1})$$

A first observation follows from the boundary values of the payoff difference:

$$\Delta\Pi(0) = -D_r^1, \quad \Delta\Pi(1) = 1.$$

Hence, when $D_r^1 > 0$, namely when G_1 lies in the Stag Hunt or Prisoner's Dilemma region, there exists a unique internal equilibrium $\rho^* \in (0, 1)$. This equilibrium is unstable and therefore separates the basins of attraction of full defection and full cooperation, leading to bistability. The region $D_r^1 < 0$, corresponding to the


Cooperation Behaviour over Time: SH as G_1


Figure A.3: Well-mixed endogenous-feedback simulations with G_2 fixed to a Prisoner's Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$, and G_1 chosen in the Stag Hunt region. The curves show the average cooperation level $\rho(t)$ as a function of Monte Carlo sweeps for the three initial cooperation levels $\rho_0 = 0.1, 0.5, 0.9$, averaged over 20 independent realisations. The horizontal red line indicates the cooperation level predicted by the replicator equation. Different panels are displayed over different Monte Carlo sweep windows. The excellent agreement between stochastic simulations and analytical prediction confirms that, in this region, the endogenous feedback toward the Prisoner's Dilemma does not stabilise cooperation. The dynamics converges to the absorbing state of full defection, consistently with the well-mixed replicator analysis. *Original simulation by the author.*

Harmony Game and part of the Snowdrift region, displays a richer solution structure. In this case, internal equilibria exist only when the discriminant of Eq. (A.1) is non-negative, namely when

$$(D_g^1)^2 + 2D_g^1 + 4D_r^1 + 1 \geq 0. \quad (\text{A.2})$$

If this condition is not satisfied, the system displays full cooperation, since $\Delta\Pi(\rho) > 0$ throughout the physical interval $\rho \in [0, 1]$. If instead Eq. (A.2) is satisfied, two internal equilibria exist. These have opposite stability, with the smaller one being stable and the larger one unstable. In this case, the dynamics displays a peculiar bistability between a stable intermediate cooperation level and the absorbing state

$\rho = 1$. In this case, we can realise the nature of chimera games in the portion of parameter space satisfying

$$D_r^1 < 0, \\ (D_g^1)^2 + 2D_g^1 + 4D_r^1 + 1 > 0,$$

and

$$-D_r^1 > D_g^1.$$

In this regime, the system converges to a stable intermediate cooperation level even though the equilibrium game belongs to the SH class that, under standard fixed-game dynamics, would not support a stable coexistence state.

We now compare these analytical predictions with stochastic simulations. We selected four representative values of the first reference game,

$$(D_g^1, D_r^1) = (-0.6, -0.4), (0.3, -0.7), (0.7, -0.3), (0.6, 0.4),$$

which span different regions of the dilemma space. The simulation setting is the same as in the previous appendix section: a well-mixed population of size $N = 1000$, Fermi update rule with selection intensity $\beta = 0.2$, $k = 4$ interactions per player at each elementary update, and three initial cooperation levels $\rho_0 = 0.1, 0.5, 0.9$, with results averaged over 20 independent realisations. The trajectories followed by the effective game in the dilemma space are shown in Fig. A.4. The corresponding time evolution of the cooperation level is shown in Fig. A.5, where the horizontal red line indicates the analytical prediction obtained from the replicator equation. Overall, the Stag-Hunt-as- G_2 scenario shows excellent agreement between the Monte Carlo simulations and the analytical predictions obtained from the replicator equation. The stochastic trajectories converge toward the equilibrium levels indicated by the horizontal red reference lines, confirming that the well-mixed replicator analysis accurately captures the asymptotic behaviour of the endogenous-feedback dynamics.

A.3 Snowdrift as G_2

We now consider the case in which the final reference game is a Snowdrift Game, namely

$$G_2 = G(D_g^2, D_r^2) = G(1, -1).$$

Substituting this choice into Eq. (2.5) gives

$$\Delta\Pi(\rho) = \rho^2 (D_g^1 - D_r^1 - 2) + \rho (2D_r^1 - D_g^1 + 1) - D_r^1. \quad (\text{A.3})$$

In this case,

$$\Delta\Pi(0) = -D_r^1, \quad \Delta\Pi(1) = -1.$$

Hence, when $D_r^1 < 0$, namely when G_1 lies in the Harmony Game or Snowdrift region, there exists a unique internal equilibrium $\rho^* \in (0, 1)$. This equilibrium is stable, leading to a nontrivial coexistence state. The corresponding equilibrium game G^* lies in the Snowdrift region, so this behaviour is fully consistent with standard fixed-game evolutionary predictions. The region $D_r^1 > 0$, corresponding




Figure A.4: Trajectories in the (D_g, D_r) dilemma space for the endogenous-feedback model with G_2 fixed to a Stag Hunt, $G_2 = G(D_g^2, D_r^2) = G(-1, 1)$, and four representative choices of G_1 , namely $(D_g^1, D_r^1) = (-0.6, -0.4)$, $(0.3, -0.7)$, $(0.7, -0.3)$, and $(0.6, 0.4)$. The fixed reference game G_2 is indicated by a filled purple square, while the selected games G_1 are indicated by open squares. For each pair (G_1, G_2) , simulations are initialised from $\rho_0 = 0.1, 0.5, 0.9$, and the corresponding initial games lie on the segment connecting G_1 to G_2 . The black lines represent the trajectories followed by the effective game during the dynamics. The red crosses indicate the equilibrium games reached in the simulations, obtained from the final cooperation level estimated by averaging $\rho(t)$ over the last 20% of the simulation time after convergence. The corresponding cooperation time series are shown in Fig. A.5. The figure shows that this scenario is generally characterised by absorbing outcomes, but also contains a richer region in which intermediate equilibria may appear. Overall, the agreement with the analytical predictions is excellent. *Simulation by the author.*

to the Prisoner's Dilemma and part of the Stag Hunt region, has a richer solution structure. Internal equilibria exist only when the discriminant of Eq. (A.3) is non-negative, namely when

$$(D_g^1)^2 - 2D_g^1 - 4D_r^1 + 1 \geq 0. \quad (\text{A.4})$$

If this condition is not satisfied, the system displays full defection, since $\Delta\Pi(\rho) < 0$ throughout the physical interval $\rho \in [0, 1]$. If instead Eq. (A.4) is satisfied, two internal equilibria exist. These have opposite stability, with the larger one being

Cooperation Behaviour over Time




Figure A.5: Time evolution of the average cooperation level $\rho(t)$ for the same well-mixed endogenous-feedback simulations shown in Fig. A.4. The second reference game is fixed to a Stag Hunt, $G_2 = G(D_g^2, D_r^2) = G(-1, 1)$, while the four panels correspond to the four representative choices of $G_1 = (D_g^1, D_r^1)$ considered in the text. In each panel, the curves show the average cooperation level as a function of Monte Carlo sweeps for the three initial conditions $\rho_0 = 0.1, 0.5, 0.9$, averaged over 20 independent realisations. The horizontal red line indicates the equilibrium cooperation level predicted by the replicator equation. Different panels are displayed over different Monte Carlo sweep windows. Overall, the agreement with the analytical predictions is excellent. *Simulation by the author.*

stable. The system therefore displays a peculiar bistability between the absorbing state $\rho = 0$ and a stable intermediate cooperation level. Unlike the Stag-Hunt-as- G_2 case, no chimera games arise here. Whenever a stable internal equilibrium exists, the corresponding equilibrium game G^* lies in the Snowdrift region, so the observed coexistence state remains consistent with standard fixed-game evolutionary dynamics.

We now compare these analytical predictions with stochastic simulations. We selected four representative values of the first reference game,

$$(D_g^1, D_r^1) = (-0.6, -0.4), (-0.7, 0.3), (-0.3, 0.7), (0.6, 0.4),$$

again spanning different regions of the dilemma space. The simulation setting is the same as above: a well-mixed population of size $N = 1000$, Fermi update rule

with selection intensity $\beta = 0.2$, $k = 4$ interactions per player, and three initial cooperation levels $\rho_0 = 0.1, 0.5, 0.9$, with averages over 20 independent realisations. The trajectories in the dilemma space are shown in Fig. A.6. The corresponding time evolution of the cooperation level is shown in Fig. A.7, where the horizontal red line indicates the analytical prediction obtained from the replicator equation. The




Figure A.6: Trajectories in the (D_g, D_r) dilemma space for the endogenous-feedback model with G_2 fixed to a Snowdrift Game, $G_2 = G(D_g^2, D_r^2) = G(1, -1)$, and four representative choices of G_1 , namely $(D_g^1, D_r^1) = (-0.6, -0.4)$, $(-0.7, 0.3)$, $(-0.3, 0.7)$, and $(0.6, 0.4)$. The fixed reference game G_2 is indicated by a filled purple square, while the selected games G_1 are indicated by open squares. For each pair (G_1, G_2) , simulations are initialised from $\rho_0 = 0.1, 0.5, 0.9$, and the corresponding initial games lie on the segment connecting G_1 to G_2 . The black lines represent the trajectories followed by the effective game during the dynamics. The red crosses indicate the equilibrium games reached in the simulations, obtained from the final cooperation level estimated by averaging $\rho(t)$ over the last 20% of the simulation time after convergence. The corresponding cooperation time series are shown in Fig. A.7. The figure shows that this scenario can lead both to absorbing states and to nontrivial intermediate equilibria. *Simulation by the author.*

results again show excellent agreement between the Monte Carlo simulations and the analytical predictions of the replicator equation. In the Snowdrift-as- G_2 scenario, the feedback dynamics can lead either to absorbing states or to stable intermediate

Cooperation Behaviour over Time




Figure A.7: Time evolution of the average cooperation level $\rho(t)$ for the same well-mixed endogenous-feedback simulations shown in Fig. A.6. The second reference game is fixed to a Snowdrift Game, $G_2 = G(D_g^2, D_r^2) = G(1, -1)$, while the four panels correspond to the four representative choices of $G_1 = (D_g^1, D_r^1)$ considered in the text. In each panel, the curves show the average cooperation level as a function of Monte Carlo sweeps for the three initial conditions $\rho_0 = 0.1, 0.5, 0.9$, averaged over 20 independent realisations. The horizontal red line indicates the equilibrium cooperation level predicted by the replicator equation. Different panels are displayed over different Monte Carlo sweep windows. Overall, the numerical results are in excellent agreement with the analytical prediction. *Simulation by the author.*

cooperation levels, depending on the position of G_1 in the dilemma space. However, unlike the Stag-Hunt-as- G_2 case, the stable nontrivial equilibria correspond to equilibrium games lying in the Snowdrift region. Therefore, the observed coexistence is consistent with the standard fixed-game prediction for a Snowdrift Game. In this sense, endogenous feedback modifies the transient trajectory and the basin structure, but it does not generate chimera-game regimes in this scenario.

Appendix B

Endogenous Feedback Simulations on Square Lattices

In this appendix, we report additional Monte Carlo simulations for the endogenous-feedback model on square lattices studied in Chapter 4. These simulations complement the results discussed in the main text by showing further representative choices of the reference games G_1 and G_2 , together with different neighbourhood structures.

B.1 Von Neumann Neighbourhood

In this section, we report the additional simulations performed on a square lattice with von Neumann neighbourhood. In this case, each player interacts with its four nearest neighbours, so that the degree of the interaction graph is $k = 4$. We consider three choices of the second reference game G_2 : a Prisoner's Dilemma, a Stag Hunt, and a Snowdrift Game.

B.1.1 Prisoner's Dilemma as G_2

We first consider the case in which the second reference game is fixed as a Prisoner's Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$. Figures B.1, B.2, and B.3 show the evolution of cooperation over time associated with the systematic scan in Fig. 4.3. The selected games G_1 are grouped by region of the dilemma space: Stag Hunt, Harmony Game, and Snowdrift Game, respectively. As can be seen from the three figures, the agreement between Monte Carlo simulations and the analytical prediction given by Eq. (4.7) varies across game classes. It is excellent when G_1 is a Harmony Game, remains qualitatively good with only small discrepancies when G_1 is a Snowdrift Game, and becomes much weaker when G_1 lies in the Stag Hunt region. In the latter case, simulations sustain higher cooperation levels than predicted by the graph-replicator equation, indicating that spatial clustering favours the survival of cooperators beyond the pair-approximation prediction.


Cooperation Behaviour over Time on the Square Lattice: SH as G_1


Figure B.1: Time evolution of the cooperation level ρ for the endogenous-feedback model on a square lattice with von Neumann neighbourhood, with $G_2 = G(1, 1)$ fixed as a Prisoner's Dilemma, for all selected choices of G_1 belonging to the Stag Hunt region in the systematic scan reported in Fig. 4.3. Each panel corresponds to a different choice of G_1 . In every panel, the dynamics is shown for $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$, averaged over 10 independent Monte Carlo simulations. The red horizontal line marks the stable equilibrium cooperation level predicted by Eq. (4.7). Different panels are shown over different Monte Carlo sweep windows. In this case, significant discrepancies are observed between the analytical prediction and the Monte Carlo simulations, indicating that the graph-replicator equation does not fully capture the effect of local spatial structure when G_1 belongs to the Stag Hunt region. *Original simulation by the author.*

B.1.2 Stag Hunt as G_2

We now report the additional square-lattice results obtained when the second reference game is fixed as a Stag Hunt, $G_2 = G(D_g^2, D_r^2) = G(-1, 1)$. Substituting $D_g^2 = -1$ and $D_r^2 = 1$ into Eq. (4.6) gives

$$\dot{\rho} = \rho(1 - \rho) \left[\rho^2(2 + D_g^1 - D_r^1) + \rho \left(\frac{5}{2}D_r^1 - \frac{1}{2}D_g^1 - 1 \right) - \frac{3}{2}D_r^1 - \frac{1}{2}D_g^1 \right]. \quad (\text{B.1})$$

This equation admits the two boundary stationary states $\rho = 0$ and $\rho = 1$. Direct evaluation of the Jacobian shows that the full-defection equilibrium $\rho = 0$ is stable


Cooperation Behaviour over Time on the Square Lattice: HG as G_1


Figure B.2: Time evolution of the cooperation level ρ for the endogenous-feedback model on a square lattice with von Neumann neighbourhood, with $G_2 = G(1, 1)$ fixed as a Prisoner's Dilemma, for all selected choices of G_1 belonging to the Harmony Game region in the systematic scan reported in Fig. 4.3. Each panel corresponds to a different choice of G_1 . In every panel, the dynamics is shown for $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$, averaged over 10 independent Monte Carlo simulations. The red horizontal line marks the stable equilibrium cooperation level predicted by Eq. (4.7). Different panels are shown over different Monte Carlo sweep windows. For this class of initial games, the agreement between simulations and analytical predictions is excellent. *Original simulation by the author.*

if $3D_r^1 + D_g^1 > 0$, whereas the full-cooperation equilibrium $\rho = 1$ is always stable. In addition to the two boundary stationary states, internal equilibria may also exist. They are obtained by setting the square bracket in Eq. (B.1) equal to zero. By considering the Jacobian, it turns out that an internal equilibrium $\rho^* \in (0, 1)$ is stable if

$$2\rho^*(2 + D_g^1 - D_r^1) + \frac{5}{2}D_r^1 - \frac{1}{2}D_g^1 - 1 < 0.$$


Cooperation Behaviour over Time on the Square Lattice: SD as G_1


Figure B.3: Time evolution of the cooperation level ρ for the endogenous-feedback model on a square lattice with von Neumann neighbourhood, with $G_2 = G(1, 1)$ fixed as a Prisoner's Dilemma, for all selected choices of G_1 belonging to the Snowdrift region in the systematic scan reported in Fig. 4.3. Each panel corresponds to a different choice of G_1 . In every panel, the dynamics is shown for $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$, averaged over 10 independent Monte Carlo simulations. The red horizontal line marks the stable equilibrium cooperation level predicted by Eq. (4.7). Different panels are shown over different Monte Carlo sweep windows. The analytical benchmark captures the qualitative behaviour, although small discrepancies may arise for some choices of G_1 . *Original simulation by the author.*

To allow a direct comparison with the corresponding well-mixed results reported in Sect. A.2, we use the same four representative choices of the reference game G_1 ,

$$(D_g^1, D_r^1) \in \{(-0.6, -0.4), (0.3, -0.7), (0.7, -0.3), (0.6, 0.4)\}.$$

Figures B.4 and B.5 show the square-lattice counterparts of Figs. A.4 and A.5, respectively. The first figure reports the trajectories in the dilemma space, while the second reports the corresponding time evolution of the cooperation level together with the analytical prediction. As can be seen by comparing Figs. B.4 and B.5 with the corresponding well-mixed results in Figs. A.4 and A.5, the square-lattice dynamics gives the same qualitative outcome for most representative choices of G_1 . The only relevant difference concerns $(D_g^1, D_r^1) = (0.7, -0.3)$. In the well-mixed case, this pair led to a non-zero final cooperation level for $\rho_0 = 0.1$ and $\rho_0 = 0.5$,




Figure B.4: Trajectories in the (D_g, D_r) dilemma space for the endogenous-feedback model on a square lattice with von Neumann neighbourhood, with $G_2 = G(-1, 1)$ fixed as a Stag Hunt. The four representative choices of G_1 are the same as those used in the corresponding well-mixed analysis reported in Fig. A.4, namely $(D_g^1, D_r^1) = (-0.6, -0.4)$, $(0.3, -0.7)$, $(0.7, -0.3)$ and $(0.6, 0.4)$. The filled purple square indicates G_2 , while open squares indicate the four choices of G_1 . For each pair (G_1, G_2) , simulations are initialised from $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$, shown by coloured dots. For each initial condition, 10 independent Monte Carlo simulations are performed and averaged. Black lines denote the trajectories followed by the effective game during the Monte Carlo simulations, while red crosses mark the final equilibrium games reached in the simulations. The final position is computed from the asymptotic cooperation level estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. Compared with the well-mixed case, the only qualitative difference concerns $(D_g^1, D_r^1) = (0.7, -0.3)$, for which the square-lattice dynamics converges to full defection, while an intermediate level of cooperation $\rho \in (0, 1)$ is reached in the well-mixed scenario. *Original simulation by the author.*

whereas on the square lattice the system converges to full defection. This reduction in cooperation is already captured by the pair-approximation-based replicator equation, which predicts a lower equilibrium cooperation level than the well-mixed model. However, pair approximation does not capture the full spatial structure of the square lattice. In this case, the lattice geometry further amplifies the reduction of cooperation, driving the system to $\rho = 0$. As can be seen from the figures, this same parameter pair is also the only case in which a clear discrepancy is observed between the Monte Carlo simulations and the analytical prediction.


Cooperation Behaviour over Time on the Square Lattice: SH as G_2


Figure B.5: Time evolution of the cooperation level ρ for the endogenous-feedback model on a square lattice with von Neumann neighbourhood, with $G_2 = G(-1, 1)$ fixed as a Stag Hunt. The four panels correspond to the same representative choices of G_1 used in the well-mixed simulations shown in Fig. A.5. In each panel, the dynamics is shown for $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$, and results are averaged over 10 independent Monte Carlo simulations. The red horizontal line marks the stable equilibrium cooperation level predicted by Eq. (B.1). Different panels are shown over different Monte Carlo sweep windows. The comparison with the well-mixed case shows that the main difference occurs for $(D_g^1, D_r^1) = (0.7, -0.3)$: while the well-mixed dynamics produced a positive final cooperation level for $\rho_0 = 0.1$ and $\rho_0 = 0.5$, the square-lattice simulations converge to $\rho = 0$. This tendency is already anticipated by the pair-approximation benchmark, which predicts a lower cooperation level than in the well-mixed case, but the actual square-lattice structure amplifies the effect. This parameter pair is also the only clear discrepancy between the Monte Carlo simulations and the analytical prediction. *Original simulation by the author.*

B.1.3 Snowdrift as G_2

Finally, we report the additional square-lattice results obtained when the second reference game is fixed as a Snowdrift Game, $G_2 = G(D_g^2, D_r^2) = G(1, -1)$. Substituting $D_g^2 = 1$ and $D_r^2 = -1$ into Eq. (4.6) gives

$$\dot{\rho} = \rho(1 - \rho) \left[\rho^2(D_g^1 - D_r^1 - 2) + \rho \left(\frac{5}{2}D_r^1 - \frac{1}{2}D_g^1 + 1 \right) - \frac{3}{2}D_r^1 - \frac{1}{2}D_g^1 \right]. \quad (\text{B.2})$$

This equation admits the two boundary stationary states $\rho = 0$ and $\rho = 1$. Direct evaluation of the Jacobian shows that the full-defection equilibrium $\rho = 0$ is stable

if

$$3D_r^1 + D_g^1 > 0,$$

whereas the full-cooperation equilibrium $\rho = 1$ is always unstable. In addition to the two boundary stationary states, internal equilibria may also exist. They are obtained by setting the square bracket in Eq. (B.2) equal to zero. By considering the Jacobian, it turns out that an internal equilibrium $\rho^* \in (0, 1)$ is stable if

$$2\rho^*(D_g^1 - D_r^1 - 2) + \frac{5}{2}D_r^1 - \frac{1}{2}D_g^1 + 1 < 0.$$

To allow a direct comparison with the corresponding well-mixed results reported in Sect. A.3, we use the same four representative choices of the reference game G_1 ,

$$(D_g^1, D_r^1) \in \{(-0.6, -0.4), (-0.7, 0.3), (-0.3, 0.7), (0.6, 0.4)\}.$$

Figures B.6 and B.7 show the square-lattice counterparts of Figs. A.6 and A.7, respectively. The first figure reports the trajectories in the dilemma space, while the second reports the corresponding time evolution of the cooperation level together with the analytical prediction. As can be seen by comparing Figs. B.6 and B.7 with the corresponding well-mixed results in Figs. A.6 and A.7, the square-lattice dynamics gives the same qualitative outcome for all representative choices of G_1 considered here. In particular, the system can converge either to absorbing states or to stable intermediate cooperation levels, depending on the position of G_1 and on the initial cooperation level. However, unlike the case discussed in Sect. B.1.2, the introduction of the square-lattice structure does not change the qualitative nature of the asymptotic outcomes with respect to the well-mixed model. The pair-approximation-based replicator equation captures the Monte Carlo results well, although small quantitative differences may remain because pair approximation does not fully encode the spatial correlations generated by the square lattice.

B.2 Moore Neighbourhood

In this section, we report the additional simulations performed on a square lattice with Moore neighbourhood. In this case, each player interacts with the eight surrounding sites, so that the degree of the interaction graph is $k = 8$. These results complement the Moore-neighbourhood analysis presented in the main text in Sect. 4.1.3.

Here we focus only on the case in which the second reference game is fixed as a Prisoner's Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$. The following figures complement the systematic scan performed in Fig. 4.8. Figures B.8, B.9, and B.10 report the cooperation time series for the selected values of G_1 used in that scan, grouped according to the region of the dilemma space in which G_1 lies: Stag Hunt, Harmony Game, and Snowdrift Game, respectively. Figures B.11 and B.12 provide a direct comparison between the asymptotic cooperation levels obtained from Monte Carlo simulations and the equilibria predicted by Eq. (4.9). The first comparison refers to the scan line $D_g^1 = -0.9$, where D_r^1 is varied, while the second refers to the scan line $D_r^1 = -0.9$, where D_g^1 is varied. As can be seen from the figures,




Figure B.6: Trajectories in the (D_g, D_r) dilemma space for the endogenous-feedback model on a square lattice with von Neumann neighbourhood, with $G_2 = G(1, -1)$ fixed as a Snowdrift Game. The four representative choices of G_1 are the same as those used in the corresponding well-mixed analysis reported in Fig. A.6, namely $(D_g^1, D_r^1) = (-0.6, -0.4)$, $(-0.7, 0.3)$, $(-0.3, 0.7)$ and $(0.6, 0.4)$. The filled purple square indicates G_2 , while open squares indicate the four choices of G_1 . For each pair (G_1, G_2) , simulations are initialised from $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$, shown by coloured dots. For each initial condition, 10 independent Monte Carlo simulations are performed and averaged. Black lines denote the trajectories followed by the effective game during the Monte Carlo simulations, while red crosses mark the final equilibrium games reached in the simulations. The final position is computed from the asymptotic cooperation level estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. Compared with the corresponding well-mixed case, the square-lattice dynamics leads to the same qualitative asymptotic outcomes for the representative choices considered here. *Original simulation by the author.*

the agreement between Monte Carlo simulations and the pair-approximation-based analytical benchmark is generally very good when G_1 belongs to the Harmony Game region, remains satisfactory for most Snowdrift choices, and becomes less accurate when G_1 lies in the Stag Hunt region.


Cooperation Behaviour over Time on the Square Lattice: SD as G_2


Figure B.7: Time evolution of the cooperation level ρ for the endogenous-feedback model on a square lattice with von Neumann neighbourhood, with $G_2 = G(1, -1)$ fixed as a Snowdrift Game. The four panels correspond to the same representative choices of G_1 used in the well-mixed simulations shown in Fig. A.7. In each panel, the dynamics is shown for $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$, and results are averaged over 10 independent Monte Carlo simulations. The red horizontal line marks the stable equilibrium cooperation level predicted by Eq. (B.2). Different panels are shown over different Monte Carlo sweep windows. The comparison with the well-mixed case shows that the square-lattice dynamics does not produce qualitatively different asymptotic outcomes in this scenario. The analytical prediction captures the observed behaviour well, with only minor quantitative deviations. *Original simulation by the author.*


Cooperation Behaviour over Time on the Moore Lattice: SH as G_1


Figure B.8: Time evolution of the cooperation level ρ for the endogenous-feedback model on a square lattice with Moore neighbourhood, with $G_2 = G(1, 1)$ fixed as a Prisoner's Dilemma, for the selected choices of G_1 belonging to the Stag Hunt region in the systematic scan reported in Fig. 4.8. Each panel corresponds to a different choice of G_1 . In each panel, the dynamics is shown for $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$, averaged over 10 independent Monte Carlo simulations. The red horizontal line marks the stable equilibrium cooperation level predicted by Eq. (4.9). Simulations are performed on a square lattice with linear size $L = 200$, periodic boundary conditions, Moore neighbourhood, no self-interaction and Fermi updating with $\beta = 0.2$. Different panels are shown over different Monte Carlo sweep windows. The figure shows significant discrepancies between Monte Carlo simulations and analytical predictions for this class of reference games. *Original simulation by the author.*


Cooperation Behaviour over Time on the Moore Lattice: HG as G_1


Figure B.9: Time evolution of the cooperation level ρ for the endogenous-feedback model on a square lattice with Moore neighbourhood, with $G_2 = G(1, 1)$ fixed as a Prisoner's Dilemma, for the selected choices of G_1 belonging to the Harmony Game region in the systematic scan reported in Fig. 4.8. Each panel corresponds to a different choice of G_1 . In each panel, the dynamics is shown for $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$, averaged over 10 independent Monte Carlo simulations. The red horizontal line marks the stable equilibrium cooperation level predicted by Eq. (4.9). Simulations are performed on a square lattice with linear size $L = 200$, periodic boundary conditions, Moore neighbourhood, no self-interaction and Fermi updating with $\beta = 0.2$. Different panels are shown over different Monte Carlo sweep windows. For this class of reference games, the agreement between Monte Carlo simulations and analytical predictions is generally very good. *Original simulation by the author.*


Cooperation Behaviour over Time on the Moore Lattice: SD as G_1


Figure B.10: Time evolution of the cooperation level ρ for the endogenous-feedback model on a square lattice with Moore neighbourhood, with $G_2 = G(1, 1)$ fixed as a Prisoner's Dilemma, for the selected choices of G_1 belonging to the Snowdrift region in the systematic scan reported in Fig. 4.8. Each panel corresponds to a different choice of G_1 . In each panel, the dynamics is shown for $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$, averaged over 10 independent Monte Carlo simulations. The red horizontal line marks the stable equilibrium cooperation level predicted by Eq. (4.9). Simulations are performed on a square lattice with linear size $L = 200$, periodic boundary conditions, Moore neighbourhood, no self-interaction and Fermi updating with $\beta = 0.2$. Different panels are shown over different Monte Carlo sweep windows. The analytical benchmark captures the qualitative behaviour well, although moderate deviations may appear for some choices of G_1 . *Original simulation by the author.*




Figure B.11: Comparison between Monte Carlo simulations and analytical predictions for the endogenous-feedback model on a square lattice with Moore neighbourhood and $G_2 = G(1, 1)$ fixed as a Prisoner's Dilemma. The figure shows the final cooperation level along the line $D_g^1 = -0.9$, as D_r^1 is varied. Blue circles indicate the asymptotic cooperation level obtained from Monte Carlo simulations, estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. Error bars represent the standard deviation of the cooperation level computed over the same final 20% time window; they are barely visible because they are very small. Orange squares indicate stable equilibria predicted by Eq. (4.9), while green squares indicate unstable analytical equilibria. Simulations are performed with $\rho_0 = 0.5$, linear size $L = 200$, periodic boundary conditions, Moore neighbourhood, no self-interaction and Fermi updating with $\beta = 0.2$. The largest deviations between simulations and analytical predictions occur in the Stag Hunt region. *Original illustration by the author.*




Figure B.12: Comparison between Monte Carlo simulations and analytical predictions for the endogenous-feedback model on a square lattice with Moore neighbourhood and $G_2 = G(1, 1)$ fixed as a Prisoner's Dilemma. The figure shows the final cooperation level along the line $D_r^1 = -0.9$, as D_g^1 is varied. Blue circles indicate the asymptotic cooperation level obtained from Monte Carlo simulations, estimated by averaging $\rho(t)$ over the final 20% of the simulation time after convergence. Error bars represent the standard deviation of the cooperation level computed over the same final 20% time window; they are barely visible because they are very small. Orange squares indicate stable equilibria predicted by Eq. (4.9), while green squares indicate unstable analytical equilibria. Simulations are performed with $\rho_0 = 0.5$, linear size $L = 200$, periodic boundary conditions, Moore neighbourhood, no self-interaction and Fermi updating with $\beta = 0.2$. Along this scan, a limited discrepancy appears in the Snowdrift region, corresponding to $D_g^1 > 0$. *Original illustration by the author.*

Appendix C

Endogenous Feedback Simulations on Complex Networks

In this appendix, we report the full time evolution of the cooperation level for the endogenous feedback simulations on complex networks discussed in Sect. 5.1. These results complement the dilemma-space trajectories shown in the main text. For each network topology, the time series are grouped according to the region of the dilemma space in which the first reference game G_1 is chosen, while the second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(D_g^2, D_r^2) = G(1, 1)$. The simulation settings are the same as those used for the corresponding figures in the main text.

C.1 Erdős–Rényi Networks

This section reports the time evolution of the cooperation level for the endogenous feedback model on Erdős–Rényi networks. The simulations correspond to the dilemma space trajectories shown in Fig. 5.1. The networks have $N = 40000$ nodes and expected average degree $\langle k \rangle = 4$. Payoffs are degree-normalised and strategy updating follows the Fermi rule with $\beta = 0.8$. The time series are shown in Figs. C.1–C.3, for choices of G_1 in the Stag Hunt, Harmony Game and Snowdrift regions, respectively.

C.2 Watts–Strogatz Networks

This section reports the time evolution of the cooperation level for the endogenous-feedback model on Watts–Strogatz networks. The simulations correspond to the dilemma-space trajectories shown in Fig. 5.2. The networks have $N = 40000$ nodes, average degree $\langle k \rangle = 4$ and rewiring probability $p_{\text{rew}} = 0.1$. Payoffs are degree-normalised and strategy updating follows the Fermi rule with $\beta = 0.8$. The time series are shown in Figs. C.4–C.6, for choices of G_1 in the Stag Hunt, Harmony Game and Snowdrift regions, respectively.


Cooperation Behaviour over Time on Erdős–Rényi Networks: SH as G_1


Figure C.1: Time evolution of the cooperation level for the endogenous-feedback model on Erdős–Rényi networks, for choices of G_1 in the Stag Hunt region. The second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. Each panel corresponds to a different choice of G_1 , and the curves show the evolution of $\rho(t)$ for the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. Simulations are performed on Erdős–Rényi networks with $N = 40000$ nodes and expected average degree $\langle k \rangle = 4$, using degree-normalised payoffs and Fermi updating with $\beta = 0.8$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. Solid red lines indicate the corresponding analytical predictions obtained from the pair approximation on regular graphs in the weak-selection regime, given by Eq. (4.7). Different panels are shown over different Monte Carlo sweep windows. The excellent agreement between the Monte Carlo simulations and the analytical predictions shows that the pair approximation accurately captures the dynamics in this case. *Original simulation by the author.*

C.3 Barabási–Albert Networks

This section reports the time evolution of the cooperation level on Barabási–Albert networks for the model with endogenous feedback. The simulations correspond to the trajectories in the dilemma space shown in Fig. 5.5 and Fig. 5.6. The networks have $N = 40000$ nodes and attachment parameter $m = 2$, corresponding to average degree $\langle k \rangle \simeq 4$. Strategy updating follows the Fermi rule with $\beta = 0.8$. As discussed


Cooperation Behaviour over Time on Erdős–Rényi Networks: HG as G_1


Figure C.2: Time evolution of the cooperation level for the endogenous-feedback model on Erdős–Rényi networks, for choices of G_1 in the Harmony Game region. The second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. Each panel corresponds to a different choice of G_1 , and the curves show the evolution of $\rho(t)$ for the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. Simulations are performed on Erdős–Rényi networks with $N = 40000$ nodes and expected average degree $\langle k \rangle = 4$, using degree-normalised payoffs and Fermi updating with $\beta = 0.8$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. Solid red lines indicate the corresponding analytical predictions obtained from the pair approximation on regular graphs in the weak-selection regime, given by Eq. (4.7). Different panels are shown over different Monte Carlo sweep windows. The excellent agreement between the Monte Carlo simulations and the analytical predictions shows that the pair approximation accurately captures the dynamics in this case. *Original simulation by the author.*

in the main text, the results on Barabási–Albert networks depend significantly on how payoffs are computed. For this reason, both cases are reported here. The time


Cooperation Behaviour over Time on Erdős–Rényi Networks: SD as G_1


Figure C.3: Time evolution of the cooperation level for the endogenous-feedback model on Erdős–Rényi networks, for choices of G_1 in the Snowdrift region. The second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. Each panel corresponds to a different choice of G_1 , and the curves show the evolution of $\rho(t)$ for the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. Simulations are performed on Erdős–Rényi networks with $N = 40000$ nodes and expected average degree $\langle k \rangle = 4$, using degree-normalised payoffs and Fermi updating with $\beta = 0.8$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. Solid red lines indicate the corresponding analytical predictions obtained from the pair approximation on regular graphs in the weak-selection regime, given by Eq. (4.7). Different panels are shown over different Monte Carlo sweep windows. The excellent agreement between the Monte Carlo simulations and the analytical predictions shows that the pair approximation accurately captures the dynamics in this case. *Original simulation by the author.*

series are shown in Figs. C.7–C.9 for the degree-normalised case and in Figs. C.10–C.12 for the accumulated-payoff case, with G_1 chosen in the Stag Hunt, Harmony Game and Snowdrift regions, respectively.


Cooperation Behaviour over Time on Watts--Strogatz Networks: SH as G_1


Figure C.4: Time evolution of the cooperation level for the endogenous-feedback model on Watts–Strogatz networks, for choices of G_1 in the Stag Hunt region. The second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. Each panel corresponds to a different choice of G_1 , and the curves show the evolution of $\rho(t)$ for the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. Simulations are performed on Watts–Strogatz networks with $N = 40000$ nodes, average degree $\langle k \rangle = 4$ and rewiring probability $p_{\text{rew}} = 0.1$, using degree-normalised payoffs and Fermi updating with $\beta = 0.8$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. Solid red lines indicate the corresponding analytical predictions obtained from the pair approximation on regular graphs in the weak-selection regime, given by Eq. (4.7). Different panels are shown over different Monte Carlo sweep windows. The comparison reveals a strong discrepancy between the Monte Carlo simulations and the analytical predictions, showing that the pair approximation does not accurately capture the dynamics in this case. *Original simulation by the author.*


Cooperation Behaviour over Time on Watts--Strogatz Networks: HG as G_1


Figure C.5: Time evolution of the cooperation level for the endogenous-feedback model on Watts–Strogatz networks, for choices of G_1 in the Harmony Game region. The second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. Each panel corresponds to a different choice of G_1 , and the curves show the evolution of $\rho(t)$ for the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. Simulations are performed on Watts–Strogatz networks with $N = 40000$ nodes, average degree $\langle k \rangle = 4$ and rewiring probability $p_{\text{rew}} = 0.1$, using degree-normalised payoffs and Fermi updating with $\beta = 0.8$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. Solid red lines indicate the corresponding analytical predictions obtained from the pair approximation on regular graphs in the weak-selection regime, given by Eq. (4.7). Different panels are shown over different Monte Carlo sweep windows. The agreement between the Monte Carlo simulations and the analytical predictions is almost perfect in this case. *Original simulation by the author.*


Cooperation Behaviour over Time on Watts--Strogatz Networks: SD as G_1


Figure C.6: Time evolution of the cooperation level for the endogenous-feedback model on Watts–Strogatz networks, for choices of G_1 in the Snowdrift region. The second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. Each panel corresponds to a different choice of G_1 , and the curves show the evolution of $\rho(t)$ for the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. Simulations are performed on Watts–Strogatz networks with $N = 40000$ nodes, average degree $\langle k \rangle = 4$ and rewiring probability $p_{\text{rew}} = 0.1$, using degree-normalised payoffs and Fermi updating with $\beta = 0.8$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. Solid red lines indicate the corresponding analytical predictions obtained from the pair approximation on regular graphs in the weak-selection regime, given by Eq. (4.7). Different panels are shown over different Monte Carlo sweep windows. The discrepancies between the Monte Carlo simulations and the analytical predictions are limited, and the analytical prediction captures the qualitative behaviour of the numerical dynamics. *Original simulation by the author.*


Cooperation Behaviour over Time: SH as G_1 on a Barabási--Albert Network


Figure C.7: Time evolution of the cooperation level for the endogenous-feedback model on Barabási–Albert networks with degree-normalised payoffs, for choices of G_1 in the Stag Hunt region. The second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. Each panel corresponds to a different choice of G_1 , and the curves show the evolution of $\rho(t)$ for the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. Simulations are performed on Barabási–Albert networks with $N = 40000$ nodes and attachment parameter $m = 2$, using degree-normalised payoffs and Fermi updating with $\beta = 0.8$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. Solid red lines indicate the corresponding analytical predictions obtained from the pair approximation on regular graphs in the weak-selection regime, given by Eq. (4.7). Different panels are shown over different Monte Carlo sweep windows. The agreement between the Monte Carlo simulations and the analytical predictions is excellent in this case. *Original simulation by the author.*


Cooperation Behaviour over Time: HG as G_1 on a Barabási–Albert Network


Figure C.8: Time evolution of the cooperation level for the endogenous-feedback model on Barabási–Albert networks with degree-normalised payoffs, for choices of G_1 in the Harmony Game region. The second reference game is fixed as a Prisoner's Dilemma, $G_2 = G(1,1)$. Each panel corresponds to a different choice of G_1 , and the curves show the evolution of $\rho(t)$ for the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. Simulations are performed on Barabási–Albert networks with $N = 40000$ nodes and attachment parameter $m = 2$, using degree-normalised payoffs and Fermi updating with $\beta = 0.8$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. Solid red lines indicate the corresponding analytical predictions obtained from the pair approximation on regular graphs in the weak-selection regime, given by Eq. (4.7). Different panels are shown over different Monte Carlo sweep windows. The agreement between the Monte Carlo simulations and the analytical predictions is excellent in this case. *Original simulation by the author.*


Cooperation Behaviour over Time: SD as G_1 on a Barabási--Albert Network


Figure C.9: Time evolution of the cooperation level for the endogenous-feedback model on Barabási–Albert networks with degree-normalised payoffs, for choices of G_1 in the Snowdrift region. The second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. Each panel corresponds to a different choice of G_1 , and the curves show the evolution of $\rho(t)$ for the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. Simulations are performed on Barabási–Albert networks with $N = 40000$ nodes and attachment parameter $m = 2$, using degree-normalised payoffs and Fermi updating with $\beta = 0.8$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. Solid red lines indicate the corresponding analytical predictions obtained from the pair approximation on regular graphs in the weak-selection regime, given by Eq. (4.7). Different panels are shown over different Monte Carlo sweep windows. The agreement between the Monte Carlo simulations and the analytical predictions is excellent in this case. *Original simulation by the author.*


Cooperation Behaviour over Time on a Barabási–Albert Network with Total Payoff: SH as G_1


Figure C.10: Time evolution of the cooperation level for the endogenous-feedback model on Barabási–Albert networks with accumulated payoffs, for choices of G_1 in the Stag Hunt region. The second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. Each panel corresponds to a different choice of G_1 , and the curves show the evolution of $\rho(t)$ for the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. Simulations are performed on Barabási–Albert networks with $N = 40000$ nodes and attachment parameter $m = 2$. Payoffs are accumulated over all neighbours, without degree normalisation, and strategies are updated according to the Fermi rule with $\beta = 0.8$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. Different panels are shown over different Monte Carlo sweep windows. Solid red lines indicate the analytical benchmark obtained from the pair approximation on regular graphs in the weak-selection regime, given by Eq. (4.7). In this accumulated-payoff case, this benchmark is shown only as a reference, because the assumptions of the approximation are not satisfied on heterogeneous Barabási–Albert networks. *Original simulation by the author.*


Cooperation Behaviour over Time on a Barabási–Albert Network with Total Payoff: HG as G_1


Figure C.11: Time evolution of the cooperation level for the endogenous-feedback model on Barabási–Albert networks with accumulated payoffs, for choices of G_1 in the Harmony Game region. The second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. Each panel corresponds to a different choice of G_1 , and the curves show the evolution of $\rho(t)$ for the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. Simulations are performed on Barabási–Albert networks with $N = 40000$ nodes and attachment parameter $m = 2$. Payoffs are accumulated over all neighbours, without degree normalisation, and strategies are updated according to the Fermi rule with $\beta = 0.8$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. Different panels are shown over different Monte Carlo sweep windows. Solid red lines indicate the analytical benchmark obtained from the pair approximation on regular graphs in the weak-selection regime, given by Eq. (4.7). In this accumulated-payoff case, this benchmark is shown only as a reference, because the assumptions of the approximation are not satisfied on heterogeneous Barabási–Albert networks. *Original simulation by the author.*


Cooperation Behaviour over Time on a Barabási–Albert Network with Total Payoff: SD as G_1


Figure C.12: Time evolution of the cooperation level for the endogenous-feedback model on Barabási–Albert networks with accumulated payoffs, for choices of G_1 in the Snowdrift region. The second reference game is fixed as a Prisoner’s Dilemma, $G_2 = G(1, 1)$. Each panel corresponds to a different choice of G_1 , and the curves show the evolution of $\rho(t)$ for the initial conditions $\rho_0 = 0.1$, $\rho_0 = 0.5$ and $\rho_0 = 0.9$. Simulations are performed on Barabási–Albert networks with $N = 40000$ nodes and attachment parameter $m = 2$. Payoffs are accumulated over all neighbours, without degree normalisation, and strategies are updated according to the Fermi rule with $\beta = 0.8$. For each initial condition, results are averaged over 5 independent network realisations, with 5 independent Monte Carlo simulations performed on each realisation. Different panels are shown over different Monte Carlo sweep windows. Solid red lines indicate the analytical benchmark obtained from the pair approximation on regular graphs in the weak-selection regime, given by Eq. (4.7). In this accumulated-payoff case, this benchmark is shown only as a reference, because the assumptions of the approximation are not satisfied on heterogeneous Barabási–Albert networks. *Original simulation by the author.*