

University of Massachusetts Boston

ScholarWorks at UMass Boston

Graduate Doctoral Dissertations

Doctoral Dissertations and Masters Theses

6-1-2013

Evolutionary Dynamics on Complex Networks

Swami Iyer

University of Massachusetts Boston

Follow this and additional works at: https://scholarworks.umb.edu/doctoral_dissertations



Part of the [Computer Sciences Commons](#)

Recommended Citation

Iyer, Swami, "Evolutionary Dynamics on Complex Networks" (2013). *Graduate Doctoral Dissertations*. 113.
https://scholarworks.umb.edu/doctoral_dissertations/113

This Open Access Dissertation is brought to you for free and open access by the Doctoral Dissertations and Masters Theses at ScholarWorks at UMass Boston. It has been accepted for inclusion in Graduate Doctoral Dissertations by an authorized administrator of ScholarWorks at UMass Boston. For more information, please contact scholarworks@umb.edu.

EVOLUTIONARY DYNAMICS ON COMPLEX NETWORKS

A Dissertation Presented

by

SWAMI IYER

Submitted to the Office of Graduate Studies,
University of Massachusetts Boston,
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

June 2013

Computer Science Program

© 2013 by Swami Iyer
All rights reserved

EVOLUTIONARY DYNAMICS ON COMPLEX NETWORKS

A Dissertation Presented

by

SWAMI IYER

Approved as to style and content by:

Timothy Killingback, Associate Professor of Mathematics
Chairperson of Committee

Dan A. Simovici, Professor of Computer Science
Member

Nurit Haspel, Assistant Professor of Computer Science
Member

Bala Sundaram, Professor of Physics
Member

Dan A. Simovici, Program Director
Computer Science Program

Peter Fejer, Chairperson
Computer Science Department

ABSTRACT

EVOLUTIONARY DYNAMICS ON COMPLEX NETWORKS

June 2013

Swami Iyer

B.E., VESIT, Bombay, India

M.S., University of Massachusetts Boston, USA

Ph.D., University of Massachusetts Boston, USA

Directed by Associate Professor of Mathematics Timothy Killingback

Many complex systems such as the Internet can be represented as networks, with vertices denoting the constituent components of the systems and edges denoting the patterns of interactions among the components. In this thesis, we are interested in how the structural properties of a network, such as its average degree, degree distribution, clustering, and homophily affect the processes that take place on it. In the first part of the thesis we focus on evolutionary game theory models for studying the evolution of cooperation in a population of predominantly selfish individuals. In the second part we turn our attention to an evolutionary model of disease dynamics and the impact of vaccination on the spread of infection. Throughout the thesis we use a network as an abstraction for a population, with vertices representing individuals in the population and edges specifying who can interact with whom. We analyze our models for a well-mixed population, i.e., an infinite population with random mixing, and compare the theoretical results with those obtained from computer simulations on model and empirical networks.

DEDICATION

To my aunts Subbalakshmi and Vasantha for nurturing me with care and affection, and my parents Inthumathi and Raghunathan and brother Shiva for all their support.

ACKNOWLEDGEMENTS

I am grateful to the department of Computer Science for providing me with financial support throughout the course of my study. I would like to thank my advisor, Timothy Killingback, for getting me interested in the research topics that underlie this thesis, for his guidance with the research, and for providing helpful comments on the thesis; the committee members: Dan Simovici, Bala Sundaram, and Nurit Haspel for all their support; the administrative staff: Allison Christiansen, Danielle Fontaine, Gemma Galecia, John Lewis, Maureen Title, and Sunera Amin for helping me out with my administrative needs; the system staff: Rick Martin, Leonard David, and Antonio Pera for helping me with my computational necessities; professors Ethan Bolker, Bill Campbell, Peter Fejer, Sheldon Kovitz, Maura Mast, Bob Morris, Alfred Noël, Carl Offner, Elizabeth O’Neil, Karen Ricciardi, Geza Schay, Dennis Wortman, and Catalin Zara from the Mathematics and Computer Science departments for all their support and encouragement; and my friends for making life at UMass Boston a fun, positive, and memorable experience. I am thankful to Jessica Gregory and Maryam Mahdavi, undergraduate students from the Mathematics department, for doing some of the initial work with me on the Disease Dynamics project, which is the topic of Chapter 6. I am grateful to Steve Revilak and David Weisman for perfecting the UMass Boston dissertation template in \LaTeX and thus saving the rest of us a lot of time and effort. I am also thankful to Sara Leahy for making sure that the thesis conforms to the UMass Boston dissertation standards. Finally, I would like to acknowledge the loving support of my family, and in particular, that of my aunts Subbalakshmi and Vasantha, my parents Inthumathi and Raghunathan, and my brother Shiva, without which I could not have carried out this work.

Swami Iyer; Boston; April, 2013

TABLE OF CONTENTS

DEDICATION	v
ACKNOWLEDGEMENTS	vi
LIST OF TABLES	x
LIST OF FIGURES	xii
Chapter	Page
1. INTRODUCTION	1
2. COMPLEX NETWORKS	19
Introduction	19
Examples of Complex Networks	20
Network Concepts	27
Network Models	40
Basic Properties of Select Empirical Networks	53
3. EVOLUTIONARY DYNAMICS	56
Introduction	56
Classical Game Theory	58
Evolutionary Game Theory	82
Replicator Dynamics	95
Adaptive Dynamics	111
Agent-Based Models	120
4. DISCRETE GAMES	127
Introduction	127
Assortative Interactions	128
Agent-Based Model	130
Prisoner's Dilemma	138

Chapter	Page
Snowdrift Game	154
Sculling Game	166
Discussion	179
Future Directions	182
5. CONTINUOUS GAMES	187
Introduction	187
Assortative Interactions	189
Agent-Based Model	190
Prisoner's Dilemma	198
Snowdrift Game	221
Tragedy of the Commons	243
Discussion	257
Future Directions	267
6. DISEASE DYNAMICS	270
Introduction	270
Models of Disease Dynamics	272
Simulation Results	296
Discussion	314
Future Directions	318
APPENDIX	
A. DYNAMICAL SYSTEMS	325
Introduction	325
One-Dimensional Dynamical Systems	325
Linear Systems	331
Higher Dimensional Dynamical Systems	335

Chapter	Page
B. SIMULATION SOFTWARE	343
Introduction	343
Obtaining the Software	343
Common Utilities	345
Simulating Discrete Games	349
Simulating Continuous Games	355
Simulating Disease Dynamics	364
REFERENCE LIST	372

LIST OF TABLES

Table		Page
1.	Worldwide mortality due to infectious diseases.	14
2.	Basic properties of select empirical networks.	55
3.	Effect of update rule on discrete-game dynamics on a complete network.	182
4.	Effect of update rule on discrete-game dynamics on a Barabási-Albert network.	183
5.	Effect of update rule on continuous-game dynamics on a complete network.	265
6.	Effect of update rule on continuous-game dynamics on a Barabási-Albert network.	266
7.	Values of R_0 for well-known infectious diseases.	273
A.8.	Classification of fixed points of $\dot{\mathbf{x}} = A\mathbf{x}$ according to the real and imaginary parts of the eigenvalues of A	335
B.9.	Some of the network topologies supported by the <code>network.py</code> module, along with the corresponding parameters.	345
B.10.	Parameters for simulating symmetric, two-strategy, pairwise games using our Python code.	350
B.11.	Parameters for simulating symmetric, two-strategy, pairwise games using our Java code.	353
B.12.	Parameters for simulating continuous games using our Python code.	356
B.13.	Parameters for simulating continuous games using our Java code.	360

Table	Page
B.14. Parameters for simulating disease dynamics using our Python code.	365
B.15. Parameters for simulating disease dynamics using our Java code.	369

LIST OF FIGURES

Figure	Page
1. The neural network of the worm <i>C. elegans</i>	28
2. A simple undirected network.	29
3. Degree distribution of the neural network of the worm <i>C. elegans</i>	37
4. 2-simplex S_2 and 3-simplex S_3	96
5. Phase line plots for a general symmetric, two-strategy, pairwise game.	102
6. Behavior of the full system 3.100.	108
7. Behavior of the full system 3.106.	111
8. Evolution of the distribution of persistence time x in a war of attrition game.	119
9. Sequential interactions and updates per generation in a discrete game.	132
10. Simultaneous interactions and updates per generation in a discrete game.	135
11. Phase line plots for the system 4.9 for different values of assortativity r	144
12. Evolution of the fraction x of cooperators in a symmetric, two-strategy, pairwise prisoner's dilemma game for different values of assortativity r	145
13. Variation of the long-term fraction x_∞ of cooperators with assortativity r and cost-to-benefit ratio ρ , on a complete network.	148

Figure	Page
14. Variation of the long-term fraction x_∞ of cooperators with average degree d and cost-to-benefit ratio ρ , on model networks. . .	149
15. Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on model power-law networks with clustering coefficient $C^{(2)}$	150
16. Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on model power-law networks with homophily coefficient h	151
17. Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on empirical networks.	153
18. Phase line plots for the system 4.14 for different values of assortativity r	157
19. Evolution of the fraction x of cooperators in a symmetric, two-strategy, pairwise snowdrift game, for different values of assortativity r	158
20. Variation of the long-term fraction x_∞ of cooperators with assortativity r and cost-to-benefit ratio ρ , on a complete network. .	160
21. Variation of the long-term fraction x_∞ of cooperators with average degree d and cost-to-benefit ratio ρ , on model networks. . .	162
22. Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on model power-law networks with clustering coefficient $C^{(2)}$	163
23. Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on model power-law networks with homophily coefficient h	164
24. Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on empirical networks.	165

Figure	Page
25. Phase line plots for the system 4.20 for different values of assortativity r	169
26. Evolution of the fraction x of cooperators in a symmetric, two-strategy, pairwise sculling game, for different values of assortativity r	170
27. Variation of the long-term fraction x_∞ of cooperators with assortativity r and cost-to-benefit ratio ρ , on a complete network.	173
28. Variation of the long-term fraction x_∞ of cooperators with average degree d and cost-to-benefit ratio ρ , on model networks.	174
29. Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on model power-law networks with clustering coefficient $C^{(2)}$	175
30. Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on model power-law networks with homophily coefficient h	176
31. Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on empirical networks.	178
32. Sequential interactions and updates in continuous games.	192
33. Simultaneous interactions and updates in continuous games.	195
34. Linear cost, linear benefit, and payoff functions for the CPD game.	202
35. Evolution of the distribution of strategies x in a CPD game with linear cost and benefit functions.	203
36. Quadratic cost, quadratic benefit, and payoff functions for the CPD game.	204

Figure	Page
37. Evolution of the distribution of strategies x in a CPD game with quadratic cost and benefit functions.	205
38. Variation of the long-term mean strategy \bar{x}_∞ with assortativity r and cost-to-benefit ratio ρ , on a complete network.	209
39. Variation of the long-term mean strategy \bar{x}_∞ with average degree d and cost-to-benefit ratio ρ , on model networks.	210
40. Variation of the long-term mean strategy \bar{x}_∞ with cost-to-benefit ratio ρ , on model power-law networks with clustering coefficient $C^{(2)}$	212
41. Variation of the long-term mean strategy \bar{x}_∞ with cost-to-benefit ratio ρ , on model power-law networks with homophily coefficient h	213
42. Variation of the long-term mean strategy \bar{x}_∞ with cost-to-benefit ratio ρ , on empirical networks.	214
43. Variation of the distribution of the long-term values x_∞ of strategies with assortativity r (average degree d) on complete (model) networks.	216
44. Evolution of the distribution of strategies x on model power-law networks with clustering coefficient $C^{(2)}$	217
45. Evolution of the distribution of strategies x on model power-law networks with homophily coefficient h	218
46. Evolution of the distribution of strategies x on empirical networks.	220
47. Quadratic cost, quadratic benefit, and payoff functions for the CSD game.	224
48. Evolution of the distribution of strategies x in a CSD game with quadratic cost and quadratic benefit functions.	226

Figure	Page
49. Variation of the distribution of the long-term values x_∞ of strategies with assortativity r (average degree d) on complete (model) networks.	229
50. Evolution of the distribution of strategies x on model power-law networks with average degree $d = 4$ and clustering coefficient $C^{(2)}$	230
51. Evolution of the distribution of strategies x on model power-law networks with average degree $d = 10$ and clustering coefficient $C^{(2)}$	231
52. Evolution of the distribution of strategies x on model power-law networks with average degree $d = 4$ and homophily coefficient h .	232
53. Evolution of the distribution of strategies x on model power-law networks with average degree $d = 10$ and homophily coefficient h .	233
54. Evolution of the distribution of strategies x on empirical networks.	236
55. Variation of the distribution of the long-term values x_∞ of strategies with assortativity r (average degree d) on complete (mode) networks.	237
56. Evolution of the distribution of strategies x on model power-law networks with average degree $d = 4$ and clustering coefficient $C^{(2)}$	238
57. Evolution of the distribution of strategies x on model power-law networks with average degree $d = 10$ and clustering coefficient $C^{(2)}$	239
58. Evolution of the distribution of strategies x on model power-law networks with average degree $d = 4$ and homophily coefficient h .	240
59. Evolution of the distribution of strategies x on model power-law networks with average degree $d = 10$ and homophily coefficient h .	241

Figure	Page
60. Evolution of the distribution of strategies x on empirical networks.	242
61. Quadratic cost, cubic benefit, and payoff functions for the CTOC game.	246
62. Evolution of the distribution of strategies x in a CTOC game with quadratic cost and cubic benefit functions.	248
63. Variation of the distribution of the long-term values x_∞ of strategies with assortativity r (average degree d) on complete (model) networks.	251
64. Evolution of the distribution of strategies x on model power-law networks with average degree $d = 4$ and clustering coefficient $C^{(2)}$	252
65. Evolution of the distribution of strategies x on model power-law networks with average degree $d = 10$ and clustering coefficient $C^{(2)}$	253
66. Evolution of the distribution of strategies x on model power-law networks with average degree $d = 4$ and homophily coefficient h .	254
67. Evolution of the distribution of strategies x on model power-law networks with average degree $d = 10$ and homophily coefficient h .	255
68. Evolution of the distribution of strategies x on empirical networks.	256
69. Allowed transitions between states in the SI model.	274
70. The trajectories for the SI model.	276
71. Allowed transitions between states in the SIR model.	278
72. The trajectories for SIR model.	279
73. Allowed transitions between states in the SIR model with births and deaths.	281

Figure	Page
74. The trajectories for SIR model with birth and deaths.	282
75. The R_0 curve for the SIR model with births and deaths.	285
76. Roulette-wheel selection of transmission probability for the S - vertex i from its neighborhood.	288
77. The R_0 curve for our model.	294
78. Disease dynamics on a complete network.	295
79. Evolution of mean virulence \bar{c} on a complete network	297
80. Evolution of mean virulence \bar{c} on random regular networks with varying average degree d	298
81. Evolution of mean virulence \bar{c} on Erdős-Rényi networks with varying average degree d	299
82. Evolution of mean virulence \bar{c} on Barabási-Albert networks with varying average degree d	300
83. Long-term value \bar{c}_∞ of mean virulence versus average degree d , for random regular, Erdős-Rényi, and Barabási-Albert networks.	301
84. Evolution of mean virulence \bar{c} on power-law networks with clustering coefficient $C^{(2)}$	302
85. Long-term value \bar{c}_∞ of mean virulence versus clustering coeffi- cient $C^{(2)}$, on power-law networks.	303
86. Evolution of mean virulence \bar{c} on power-law networks with ho- mophily coefficient h	304
87. Long-term value \bar{c}_∞ of mean virulence versus homophily coef- ficient h , on power-law networks.	305
88. Evolution of mean virulence \bar{c} on empirical networks.	306

Figure	Page
89. Long-term value \bar{c}_∞ of mean virulence for empirical networks.	307
90. Evolution of mean virulence \bar{c} on a Barabási-Albert network with vaccinations.	310
91. Long-term value of mean virulence \bar{c}_∞ versus the fraction ν of individuals vaccinated on a Barabási-Albert network.	311
92. Evolution of mean virulence \bar{c} on an empirical network with vaccinations.	312
93. Long-term value \bar{c}_∞ of mean virulence versus the fraction ν of individuals vaccinated on an empirical network.	313
94. Long-term value \bar{c}_∞ of mean virulence versus the number s of double-edge swaps per generation on power-law networks.	316
A.95. Phase line plot and phase potrait for the system A.2.	327
A.96. Phase line plot for $\dot{x} = f(x)$	328
A.97. Phase line plot and phase potrait for the system A.9.	331
A.98. Behavior of the system A.30 near fixed points and as a whole.	342
B.99. Evolution of the fraction x of cooperators in a discrete prisoner's dilemma game.	353
B.100. Evolution of the fraction x of cooperators in a discrete prisoner's dilemma game.	356
B.101. Evolution of the mean strategy \bar{x} and the distribution of strategies x in a continuous tragedy of the commons game.	360
B.102. Evolution of the mean strategy \bar{x} and the distribution of strategies x in a continuous tragedy of the commons game.	364

Figure	Page
B.103. Evolution of the fraction of susceptible s , infected x , and recovered r individuals, and evolution of the mean virulence \bar{c}	368
B.104. Evolution of the fraction of susceptible s , infected x , and recovered r individuals, and evolution of the mean virulence \bar{c}	371

CHAPTER 1

INTRODUCTION

It is not the strongest of the species that survives, nor the most intelligent that survives. It is the one that is the most adaptable to change.

- Charles Darwin

A network is a collection of vertices connected together in pairs by edges. There are many real-world systems that are composed of individual components linked together in some fashion and which can be thought of as networks. For example, we can consider the Internet as a (technological) network [CCG02, FFF99], with the vertices denoting computers and routers, and the edges denoting the physical connections between them; we can consider the patterns of interactions among individuals in a population to form a (social) network [Zac77], with the vertices denoting individuals in the population, and the edges denoting who interacts with whom.

Real-world networks are called “complex” networks, and the systems they represent are called “complex” systems. This is because: these systems/networks exhibit emergent phenomena; their whole is larger than the sum of their constituent parts; they can adapt to, evolve, and resist random removals of their components; and they exhibit complications on all scales possible within the physical constraints of the system [BBV08].

Scientists study networks for various reasons, one of which is to understand how the structure of a network affects its function [AB02, New03]. For example, the structure of the Internet affects the efficiency with which it transports data; and the structure of a social network affects how people learn, form opinions, and affect phenomena, such as the spread of a disease. Our goal in this thesis is to study how the structural aspects of a

network affect the long-term behavior of dynamical processes that take place on the network, such as the evolution of cooperative behavior, and the spreading of a disease in a population. We represent the population as a network, with vertices corresponding to the individuals in the population, and the edges corresponding to the interaction patterns among the individuals.

In Chapter 2 we provide a brief introduction to complex networks. We look at some examples of real-world networks, which are typically divided into four broad categories: technological networks, social networks, information networks, and biological networks. We describe key concepts from network theory used in the description and analysis of networks, restricting ourselves to those that are most pertinent to our work. Some of these concepts include: degree of a vertex which is the number of neighbors that the vertex has; average degree of a network which is the average value of the degree of all the vertices in the network; degree distribution which is the frequency distribution of the degree of vertices in the network; clustering coefficient of a network which measures the extent to which vertices in the network tend to cluster together; and homophily coefficient of a network which measures the tendency of the vertices in the network to associate with others whom they perceive as being (dis)similar to themselves. We consider several mathematical models for building networks that resemble the ones seen in the real world. Such model networks are incredibly valuable in our study because they allow us to freely change the network properties and learn about their effects on the processes taking place on the networks. The model networks we consider are: complete networks in which every vertex is connected to every other vertex; random regular networks [Wor99] in which every vertex is connected to a fixed number of other vertices; Erdős-Rényi networks [ER59, ER60, ER61] which have Poisson degree distributions; Barabási-Albert networks [BA99] which have power-law degree distributions; power-law networks [HK02] with clustering; and power-law networks [XS04] with homophily. Model networks do not

entirely match the properties of real-world networks. We would like to know if the subtle differences that exist between the model and real-world networks have an effect on the outcome of dynamical processes that take place on the networks. We conclude the chapter by listing the basic properties of select real-world networks that we will use—in addition to the model networks—as substrates for the dynamical processes that we consider in this thesis. The empirical networks we consider are: a symmetrized snapshot of the structure of the Internet at the level of autonomous systems, reconstructed from BGP tables posted by the University of Oregon Route Views Project; a network of coauthorships between scientists posting preprints on the High-Energy Theory E-Print Archive [New01]; network of coauthorships between scientists posting preprints on the Astrophysics E-Print Archive [New01]; a portion of the New Orleans Facebook network of user-to-user links [VMC09]; a snapshot of the Gnutella peer-to-peer file sharing network [LKF07, RF02]; and a network of protein-protein interactions in the yeast *S. cerevisiae* [JTA00, JMB01].

The interactions among individuals in the processes we consider are such that their outcome to an individual depends not only on the state or action of that individual, but also on the states or actions of that individual's neighbors, i.e., the individuals with whom the given individual interacts. For example, when an animal fights another animal over a territory, whether or not it uses a passive or an aggressive fighting strategy depends on the strategy used by its opponent [SP73, Smi74]; similarly, in an epidemic outbreak in a population, whether or not an individual is infected, and the extent to which the individual is infected, depends not only on the state (susceptible or immune) of the individual but also on the state (infected or not) of the individual's neighbors and extent to which they are infected [AM92]. Such interdependent interactions are best understood in the language of game theory.

Classical game theory [NMK07] provides a mathematical framework for describing and analyzing strategic interactions between rational individuals. Ideas from classical

game theory are applicable not only to contexts that are literally games, such as chess or poker, but also in settings that are not manifestly games [EK10], such as the pricing of a new product when other firms have similar new products, deciding how to bid in an auction, choosing a route on the Internet or through a transportation network, or deciding whether to adopt an aggressive or passive stance in international relations. The theory assumes that the interacting decision makers (players) are rational, i.e., they act so as to maximize their payoffs. It is also assumed that in reasoning about behavior in strategic games, each player actually succeeds in selecting the optimal strategies. In simple games played by experienced players, this is a reasonable assumption. But if the players are inexperienced and the games complex, then this assumption is often violated. It is also unreasonable to assume that the players will always behave rationally. Thus, individuals cannot always be expected to play games the way classical game theory predicts they should, and in fact they do not [GH01]. Furthermore, if we intend to apply the game-theoretic ideas to non-human organisms—and our models are indeed applicable in such settings—then the assumption about rational behavior must be relaxed.

A refinement of classical game theory, called evolutionary game theory [SP73, Smi74], incorporates ideas from Darwinian evolution [Dar69], and does not assume that the interacting individuals are rational. The theory has been enormously successful in evolutionary biology, and more recently, has also been extended to social settings [SN99]. In spite of the radical differences in how the two theories perceive strategic interactions, there are some deep connections between the two. In addition to providing a framework for describing situations in which individuals are not overtly making decisions, evolutionary game theory allows us to model the process by which frequencies of strategies change in a population. Once a model of the population dynamics is specified [Zee80], all of the standard concepts from dynamical systems [Str94, Ban05], such as attractors and stability, can be applied.

The models we propose in this thesis do not lend themselves to easy experimentation. We instead rely heavily on computer simulations to illustrate and verify our models. We propose agent-based models [Bon02] for in-silico emulation of how individuals in a population might behave if our models were implemented in reality. We use the agent-based models not only to verify the predictions made by our models, but also to simulate scenarios that are mathematically intractable.

We begin Chapter 3 with an introduction to classical game theory. We describe concepts such as rational behavior, games and their types, representation of games, strategies, payoffs, and ways (called solution concepts) to reason about optimal strategies in a game, including the all important Nash equilibrium and the related Nash's existence theorem. We introduce the three canonical two-player, two-strategy games: prisoner's dilemma, chicken, and coordination games, and identify the Nash equilibria in each. We conclude the discussion of classical game theory by highlighting some of its drawbacks.

The next topic we consider is that of evolutionary game theory. We revisit the concepts from classical game theory and discuss them in the light of Darwinian evolution. We describe the central concept of an evolutionarily stable strategy (ESS), and its relationship to the idea of Nash equilibrium from classical game theory. We reconsider the three canonical two-player, two-strategy games and identify the ESSs in each. We conclude the section on evolutionary game theory by discussing some of its shortcomings.

We next look at games from a dynamical systems standpoint. We discuss replicator dynamics [TJ78] and adaptive dynamics [GKM97, MB07] for analyzing discrete games with two or more strategies and continuous-strategy games respectively. We conclude the chapter by looking at agent-based models for simulating evolutionary dynamics on a computer.

In Chapter 4 we turn our attention to one of the two dynamical phenomena that this thesis proposes to study, namely the evolution of cooperative behavior in a population of selfish individuals. Cooperation—an act costly to individuals performing it, but beneficial to others—is essential for the formation of systems at all levels of complexity, from replicating molecules to multi-cellular organisms to human societies. Examples of cooperative behavior abound in the life, social, and information sciences. We sample a handful of them here.

Cooperation in the life sciences: assembly of earliest replicating molecules to form larger replicating entities capable of encoding more information [SS97]; integration of once free-living prokaryote ancestors of mitochondria and chloroplasts into eukaryotic cells [SS97]; differential production of intracellular products needed for replication in an RNA phage—a virus that reproduces inside a bacterium [TC99]; vampire bats sharing blood with one another [Wil90]; lions hunting and defending their territories in packs; predator inspection in guppies and sticklebacks [Mil87]; animals engaging in allogrooming [MH97]; and alarm calls in animals in response to danger [SCM80].

Cooperation in the social sciences: coordinating on social norms such as greeting customs, dress code, and driving regulations; political activism; consumption of performance-enhancement drugs in competitive sports; nuclear arms race; blood donation; evolution of nation states from hunter gatherer societies; evolution of language and culture; and dependence on natural resources such as fossil fuels, the global atmosphere, or the world's fisheries, as well as man-made resources such as social welfare [MSK06, Ost90].

Cooperation in the information sciences: correct implementation of the TCP protocol so as to avoid congestion in Internet traffic [LS08]; contributions towards free and open-source software development; working in groups on a software engineering project;

sharing files over a peer-to-peer network [RF02, SR04]; sharing Internet bandwidth [Ost90]; and adopting uniform technological standards such as encryption protocols.

In spite of its ubiquity, obtaining a satisfactory understanding of the origin and stability of cooperation is fundamentally difficult [Dug97, Fra98, Kol98, SS97]. From a purely sociological standpoint, the difficulty is obvious. An act of cooperation by an individual benefits another individual, but is costly to the one cooperating. Why should then individuals cooperate if they can reap the benefits of cooperation without bearing the costs of cooperating themselves? Explaining the evolutionary origins of cooperation is just as hard. Cooperation, from an evolutionary perspective, means that individuals forgo some of their reproductive potential to help one another. Selfish individuals on the other hand always have a higher fitness than cooperators, since they receive the benefits of cooperation without bearing the costs. Therefore, in a world of non-cooperators, a cooperative mutant would be eliminated by natural selection, and it is thus hard for cooperation to evolve in the first place. Furthermore, even if cooperation is established by some means, it cannot be stable, since selfish mutants can invade the population and increase in frequency until the cooperators have been eliminated.

Social interactions among individuals in a population can be studied using classical and evolutionary game theory, and various discrete and continuous games can be used to model cooperative behavior. It is, however, often very difficult to determine the payoffs to the interacting individuals and hence to determine with absolute certainty whether the payoff structure meets the requirements of any particular model. There are some exceptions. In the RNA phage described in [TC99], the payoff values can be measured with some precision. In case of individuals sharing Internet bandwidth, it is in principle possible to measure the payoffs to the individuals.

We study discrete forms of cooperative behavior, in which individuals either cooperate or defect, using the three canonical two-strategy games: prisoner's dilemma, chicken, and coordination games. We formulate the three games in terms of two strategies, namely cooperation and defection. The chicken and coordination games formulated in this manner are called the snowdrift and sculling games respectively. Analysis of the three games using classical and evolutionary game theory suggests that in the prisoner's dilemma game, the best thing for an individual to do is to defect. Absent any external mechanism, cooperation is thus not possible in the prisoner's dilemma game. In the snowdrift game, it is best to cooperate when the other individual defects, and vice versa. Cooperators can thus coexist with defectors in the snowdrift game. Finally, in the sculling game, the best thing for an individual to do is to imitate the opponent, i.e., cooperate (defect) if the opponent cooperates (defects).

Prisoner's dilemma has been extensively studied as a model for cooperative behavior [Dug97, Now06b, SN99]. There are five mechanisms that have been proposed for the evolution of cooperation: kin selection [Ham63], direct reciprocity [Axe06, NS92, NS93, Tri71], indirect reciprocity [NS98, NS05], network reciprocity [LHN05], and group selection [TN06]. Apart from [HD04], the snowdrift game has received very little attention as a metaphor for cooperation, even though the chicken game is quite extensively studied in the fields of political science, economics, and biology as models of conflict between two players [Rus59, SP73, SF07]. The sculling game formulation, to the best of our knowledge, is entirely new, although games in the coordination class, such as the stag-hunt game, have been studied in social science, as models for situations in which all parties can realize mutual gains, but only by making mutually consistent decisions [Sky03, SF07].

We propose assortative interactions [Gra90] as a mechanism for promoting cooperative behavior. Assortative interaction is closely related to the idea of kin selection

[Ham63], which was suggested by William Donald Hamilton, a British evolutionary biologist widely recognized as one of the greatest evolutionary theorists of the 20th century, as a possible mechanism for the evolution of altruistic behavior. In a well-mixed population—an infinitely large population with random mixing, an assortative interaction is an interaction among individuals of the same type, which in the context of games would mean individuals playing the same strategy. We quantify the level of assortative interactions among individuals using a single parameter $r \in [0, 1]$, called the degree of assortativity, and which is interpreted as the probability with which an individual interacts with another individual of its own type. In a population represented as a network, the structure of the network dictates who can interact with whom, and thus induces a degree of network assortativity in game interactions, which bears an inverse relationship to the average degree of the network. We are interested in how assortativity affects the fraction of cooperators in a population. As we shall see, assortativity is an essential ingredient for cooperation in the prisoner’s dilemma game, and in the case of the snowdrift and sculling games, it promotes higher fractions of cooperators.

Chapter 4 is organized as follows: we present a framework of replicator dynamics [TJ78], modified to take into account assortative interactions; we propose agent-based models for emulating the prisoner’s dilemma, snowdrift, and sculling games on a computer; we formulate the games as models for cooperation; we analyze the games for a well-mixed population using the modified framework of replicator dynamics; we present results of simulating the games on model and empirical networks and provide a general discussion of the results; and finally, we conclude the chapter with suggestions for future inquiry.

Cooperative behavior is not always discrete in nature. For example, when vampire bats share blood with one another [Wil90], or when individuals share files over a peer-to-peer network [RF02, SR04], they exhibit varying levels of cooperative behavior. Besides not

being able to take degrees of cooperative behavior into account, discrete game models suffer from another problem in that they consider the invasiveness and stability of fully developed, highly cooperative interactions, despite the fact that the gradual evolution of cooperation from an initially selfish state represents a more plausible evolutionary scenario. A framework which allows different degrees of cooperation is thus more natural. Once variable levels of cooperation are considered it becomes possible to study the crucial issue of how cooperation can evolve gradually from an initial state consisting of non-cooperative entities.

Situations in which individuals exhibit varying degrees of cooperative behavior can be described using continuous games, and this is the topic of Chapter 5. The level of cooperation or investment of an individual is that individual's strategy in the game. An investment is represented as a continuous variable, and the cost and benefit associated with an investment are represented as continuous functions of the investment. We consider three continuous games: the continuous prisoner's dilemma game (CPD) in which a cooperative investment made by an individual (donor) towards another individual (recipient) benefits the recipient but costs the donor; the continuous snowdrift game (CSD) in which the cooperative investment benefits both the donor and recipient but costs the donor; and the continuous tragedy of the commons (CTOC) game in which the cooperative investment—in this context an investment means the level of consumption of a limited common-pool resource, and cooperative investments mean modest levels of consumption—benefits the donor but costs both the donor and recipient.

In the CPD game, just as in the case of the discrete prisoner's dilemma game, absent any external mechanism, the investments made by individuals will gradually evolve to zero for any cost and benefit function. Various mechanisms have been proposed for the promotion and maintenance of cooperative investments in the CPD game. In [KDK99] the authors show that when the CPD game is played on two-dimensional spatial lattices,

investments evolve readily from very low levels to significant levels corresponding to cooperation. In [KD02] the authors consider the CPD game in a well-mixed population and demonstrate both analytically and by simulation, that payoff-based strategies, which is the idea that individuals invest more in cooperative interactions when they profit from these interactions, provide a natural explanation for the gradual evolution of cooperation from an initially noncooperative state and for the maintenance of cooperation thereafter. In [IKD04] the authors examine the effect of increasing the neighborhood size when the CPD game is played on a two-dimensional lattice, and one of their findings is that the mean-field limit of no cooperation is reached for a critical neighborhood size of about five neighbors.

In the CSD game, depending on the nature of the cost and benefit functions, different evolutionary outcomes are possible. In [DHK04] the authors study the CSD game in the context of a well-mixed population. They provide a complete classification of the adaptive dynamics for quadratic cost and benefit functions. Depending on the coefficients of these functions, the end state of the population can be: a state in which high and low investors can coexist—termed “tragedy of the commune” by the authors; a state in which every individual invests the same amount; or a state in which every individual makes either zero investment or the maximum possible investment.

In the classic tragedy of the commons [Har68], each individual benefits from consuming the common-pool resource, but the incurring costs are diluted and shared among all interacting individuals. As a result, the common resource is prone to exploitation and self-interest drives over-consumption of the resource to the detriment of all. In [KDH10] the authors consider the CTOC game in a well-mixed setting and demonstrate that the evolutionary outcomes based on a continuous range of consumption levels of common-pool resources are often strikingly different from the the classic tragedy of the commons. The authors reveal a second tragedy: not only is the common resource

overexploited, but selection may result in states in which high and low consumers stably coexist.

We propose, as we do for the discrete games, assortative interactions [Gra90] as a mechanism for promoting and maintaining cooperative investments. In a population represented as a network, the structure of the network induces a degree of network assortativity in game interactions. We are interested in how assortativity affects the level of investments made by individuals. As we shall see, assortative interactions provide a mechanism for the promotion and maintenance of cooperative investments.

Chapter 5 is organized as follows: we present the framework of adaptive dynamics [GKM97, MB07], modified to take into account assortative interactions; we propose agent-based models for emulating on a computer how individuals in a population might play the CPD, CSD, and CTOC games; we analyze the games for a well-mixed population for general cost and benefit functions using the modified framework of adaptive dynamics, and also consider a few specific cost and benefit functions for which the analysis is tractable; we present results of simulating the three games on model and empirical networks and provide a general discussion of the results; and finally, we conclude the chapter with suggestions for future inquiry.

In Chapter 6 we turn our attention to the other dynamical process on complex networks, namely the spreading of infectious diseases in a population. Humans, for a long time, have been battling with deadly diseases such as small pox, malaria, tuberculosis, cholera, and somewhat recently, with HIV/AIDS. Some of the most devastating pandemics—diseases that affect people over an extensive geographical area—witnessed by human history include: the Plague of Justinian (541-750) that killed between 50% and 60% of Europe's population [Bec]; the Black Death (1347-1352) that killed about 25 million individuals in Europe [Bec]; the introduction of smallpox, measles, and typhus to

the Americas by European explorers during the 15th and 16th centuries that caused millions of deaths among the the natives [DC96]; the first European influenza epidemic (1556-1560) that killed about 20% of the population [DC96]; smallpox that killed an estimated 60 million Europeans during the 18th century [His]; tuberculosis that killed about a quarter of the adult population of Europe in the 19th century [DP]; and the influenza of 1918 (or the Spanish Flu) that killed 25-50 million people [Lib].

The World Health Organization (WHO) collects information on global deaths by International Classification of Disease (ICD) code categories, a health care classification system that provides codes to classify diseases [Org]. Table 1 lists the top infectious disease killers which caused more than 100,000 deaths in 2002, with the 1993 data included for comparison [Org95, Org04]. The top three killers are HIV/AIDS, tuberculosis, and malaria. While the number of deaths due to nearly every disease have decreased, deaths due to HIV/AIDS have increased fourfold.

Modern technological societies have witnessed a new kind of virus, namely the computer virus—a computer program that can replicate itself and spread from one computer to another—which has disrupted productivity and caused billions of dollars in damage [KSC97, Kle07, LM01]. Some of the worst viruses that have wreaked havoc in the computer world include: *Melissa*, which attacked computers in 1999. The first fifty individuals in a user’s Microsoft Outlook contact list were sent an infected Microsoft Word document. The virus ran wild on corporate networks causing approximately \$1 billion in damages; *MSBlast*, which was launched in 2003 to exploit a vulnerability in the Microsoft operating system. The virus infected hundreds of thousands of computers and caused over \$2 billion in damages; and *LoveLetter*, which was introduced in 2010, targeted users through email, chat clients, and peer-to-peer networks. The virus attacked addresses in Microsoft Outlook and sent a malicious executable file attached to an email titled “I Love You”. The virus caused over \$5.5 billion in damages to computer users [Cro10].

Cause of Death	% of All Deaths	2002 Deaths ($\times 10^6$)	1993 Deaths ($\times 10^6$)	2002 Rank	1993 Rank
All infectious diseases	25.9%	14.7	16.4	-	-
Lower respiratory infections	6.9%	3.9	4.1	1	1
HIV/AIDS	4.9%	2.8	0.7	2	7
Diarrheal diseases	3.2%	1.8	3.0	3	2
Tuberculosis	2.7%	1.6	2.7	4	3
Malaria	2.2%	1.3	2.0	5	4
Measles	1.1%	0.6	1.1	6	5
Pertussis	0.5%	0.29	0.36	7	7
Tetanus	0.4%	0.21	0.15	8	12
Meningitis	0.3%	0.17	0.25	9	8
Syphilis	0.3%	0.16	0.19	10	11
Hepatitis B	0.2%	0.1	0.93	11	6
Tropical diseases	0.2%	0.13	0.53	12-17	9, 10, 16-18

Table 1: Worldwide mortality due to infectious diseases. Note: other causes of death include maternal and perinatal conditions (5.2%), nutritional deficiencies (0.9%), noncommunicable conditions (58.8%), and injuries (9.1%).

The epidemiological models that are used in the study of infectious diseases in humans also apply to computer viruses. In fact they go even further. They also apply to the study of information diffusion [KSC97, Kle07, LM01]. For example, the models could be used to understand why certain innovations end up being widely adopted, while others fizzle out; or how certain videos, such as “Gangnam Style”¹ managed to go viral on YouTube. Mathematical modeling of epidemics is thus an active area of research spanning multiple disciplines. Epidemiologists, computer scientists, and social scientists share a common interest in studying spreading phenomena and use similar models to describe the diffusion of viruses, knowledge, and innovation.

When a host individual contracts an infection, the pathogen responsible for the infection typically multiplies within the host while the immune system tries to combat it. The details of the interactions between the host and the pathogen is rather complicated. If we want to understand fully how diseases spread through populations we need to take all of these details into consideration, which is close to impossible. Fortunately, we can resort to simplified models that often provide good insight on the disease dynamics.

We first look at compartmental models from epidemiology, in which the within-host dynamics of a disease is represented as changes between a few disease states or compartments, such as susceptible (S), infected (I), recovered (R), and so forth [AM92, Bai75, DH00]. These models assume a fully-mixed population in which every individual has an equal chance of coming into contact with every other individual. The models are deterministic and one can write down a set of differential equations to describe the disease dynamics. The individual-level disease parameters such as the transmission rate of a disease are fixed in such models, and one is primarily interested in the change of the number of infected individuals in the population as a function of time. The main aim

¹As of January 6, 2013, the “Gangnam Style” music video (<http://www.youtube.com/watch?v=9bZkp7q19f0>) by Park Jae-sang (PSY) has been viewed over 1.13 billion times on YouTube. [Source: Wikipedia]

in these models is therefore to understand the properties of epidemics at equilibrium, the existence of a non-zero density of infected individuals, the presence or absence of a global outbreak, etc. We consider three different compartmental models: the susceptible-infected (SI) model, the susceptible-infected-recovered (SIR) model, and a variant of the SIR model which incorporates births and deaths. For each of the models we derive the expression for the basic reproductive number R_0 , which is the mean number of secondary cases a single infected case will cause in a population with no immunity to the disease. The basic reproductive number is an important metric that is used in the study of epidemics.

In reality, most diseases spread over networks of contacts between individuals. For example, contagious diseases and parasites are communicated when people come into close contact, and HIV and other sexually transmitted diseases (STDs) are communicated when people have sex. Such diseases are better understood using epidemiological models that take into account the connection patterns among the individuals in a population. The introduction of network structure into the models leads to important new results concerning the basic properties of disease spreading processes [AM92, BPV03, MA84, ML01, PV01b].

We propose a stochastic, discrete-time, evolutionary variant of the susceptible-infected-recovered (SIR) compartmental model to study disease dynamics on complex networks. The model is based on ideas from evolutionary game theory, and focuses on the evolutionary dynamics of disease-causing parasites—anything that lives and multiplies inside another organism called the “host” and usually causes harm. We assume that the host does not evolve on the time scale under consideration, which is a reasonable assumption, since parasites generally evolve much faster than their hosts. Unlike the equation-based epidemiological models, our model is individual based and focuses on how the disease-causing pathogen evolves in virulence [Ewa93, Now06a] in

response to the selection pressures acting on it. As a result, the individual-level disease parameters, such as the transmission probability and the virulence level of a pathogen within a host, vary with time. We derive the expression for R_0 for our model for a complete network. When the transmission probability is a saturating function of the virulence, the latter evolves so as to maximize the value of R_0 .

In addition to studying the evolution of virulence, we use our model to evaluate the effectiveness of different vaccination strategies in combating a disease [AM85, DB02, PV02b] spreading over model and empirical networks with power-law degree distribution. We consider three different strategies: vaccinating individuals in reverse order of their degrees; vaccinating random neighbors of randomly picked individuals; and vaccinating individuals completely at random.

We describe how our model can be implemented on a computer. We present results from agent-based simulations on model and empirical networks, and provide a general discussion of the results. We conclude the chapter with suggestions for future inquiry.

In addition to the six chapters, this thesis also includes two appendices. In Appendix A we provide a brief introduction to dynamical systems, a subject that deals with systems that evolve with time. Some of the concepts introduced in this appendix will be used in the analysis of the dynamical systems we study in this thesis. The appendix starts by looking at one-dimensional dynamical systems, which are relatively easy to understand, and then proceeds to look at dynamical systems in two or more dimensions.

In Appendix B we describe the software we have implemented for the in-silico exploration of the agent-based models we propose in Chapters 4, 5, and 6. We provide two different implementations of the models: one in Python and the other in Java². The appendix provides information on where to obtain the software from, what is contained in

²The simulation results presented in this thesis are from running the Java-based implementations of our models.

the distribution, and how to configure and run the software. The code is made available under the GNU General Public License.

CHAPTER 2

COMPLEX NETWORKS

The world of the quark has everything to do with a jaguar circling in the night.

- Murray Gell-Mann

2.1 Introduction

A *network*—also called a *graph* in mathematical literature—is a collection of vertices connected together in pairs by edges¹. There are many real-world systems of interest that are composed of individual components linked together in some manner and which can be thought of as networks. The patterns of connections in each such system can be represented as a network, with the components of the system forming the network’s vertices and the connections forming its edges. A network is thus an abstraction that captures the essentials of a system’s connection patterns.

Over the years, scientists have developed an extensive set of mathematical, computational, and statistical tools [New10] for analyzing, modeling, and understanding networks. Since these tools work with networks in their abstract form, they can in principle be applied to almost any system that can be represented as a network. Networks are therefore a general but powerful means of representing patterns of connections between the components of a system.

One of the reasons why scientists are interested in networks is because they would like to know the effects of the networks’ structure on their function. The structural properties

¹Computer scientists refer to vertices and edges as *nodes* and *links*, physicists call them *sites* and *bonds*, and sociologists use the terms *actors* and *ties*. We employ the terms network(s), vertex (vertices), and edge(s).

of a network can have huge effects on the dynamical processes that take place on the system that the network represents [AB02, New03]. For example, the patterns of connections between computers on the Internet affect the efficiency with which the network transports data; the connections in a social network affect how people learn, form opinions, and affect phenomena, such as the spread of a disease. Our goal in this thesis is to study how the structure of networks affects the long-term behavior of dynamical processes such as the evolution of cooperative behavior, and the spreading of infectious diseases in a population. We use a network to represent the population, with the vertices of the network corresponding to the individuals in the population, and the connection patterns in the network specifying who can interact with whom.

In this chapter we first look at some examples of real-world networks, and describe the key concepts from network theory used in the description and analysis of networks. We then explore a few mathematical models for building networks that resemble the ones seen in the real world. We conclude the chapter by listing the basic properties of select real-world networks that will play the role as substrates for the dynamical processes that we consider in Chapters 4, 5, and 6.

2.2 Examples of Complex Networks

Many real-world systems are referred to as “complex systems”, and by extension, their network representations, where possible, are referred to as “complex networks”. In order to understand why this is, we must clarify what we mean by “complex” and how it differs from something that is merely “complicated”. A complex system, such as the Internet or the World Wide Web, differs from a complicated system, such as a computer or an airplane in the following fundamental aspects [BBV08]:

1. Complex systems exhibit emergent phenomena in the sense that they are the spontaneous outcome of the interactions among the many constituent parts. In other words, complex systems are not engineered systems put in place according to a definite blueprint.
2. In a complex system, the whole is larger than the sum of its constituent parts. Thus decomposing the system and studying each subpart in isolation will not suffice for understanding the whole system and its dynamics.
3. Many complex systems can adapt to, evolve, and resist random removals of their components. Random removal of components of a complicated system, such as a computer, however, will rapidly lead to the malfunction or even complete failure of the system.
4. Finally, complex systems exhibit complications on all scales possible within the physical constraints of the system. In the networks representing such systems, these complications are statistically encoded in the heavy-tail distributions characterizing the network's structural properties. It is then impossible to define a typical scale in which an average description would be reliable. Heterogeneity, heavy-tail properties, and complex topological features, such as clustering, homophily, hierarchies, and communities are characteristic of a large number of real-world networks.

Complex networks occurring in the real world are divided [New10] into four broad categories: technological networks, social networks, information networks, and biological networks. This categorization is by no means rigorous and there is some overlap between them, i.e., some networks belong to more than one category. Nevertheless, the categorization is useful since networks belonging to the same category can often be analyzed using similar techniques. We list a few important examples from each category.

Technological Networks

Technological networks are the physical networks forming the backbone of modern technological societies. The most commonly cited example of a technological network is the Internet [CCG02, FFF99], the worldwide network of physical data connections between computers and related devices. The vertices in the Internet represent computers and routers, and the edges represent physical connections between them, such as optical fiber lines. The internet is an *undirected network* since the edges run in both directions.

The telephone network is an undirected network consisting of land lines and wireless lines that transmit telephone calls. The vertices in a telephone network represent individual subscribers, local telephone exchanges and long-distance offices, and the edges represent the local and trunk lines.

The power grid network [ASB00, WS98] is an undirected network of high-voltage transmission lines that transport electric power within and across countries. The vertices in a power grid network represent generating stations and switching substations, and the edges represent the high-voltage lines.

Other examples of technological networks include transportation networks [ASB00, SDC03], such as airline routes, road networks, and railway networks. In these typically undirected networks, the vertices represent the geographic locations and the edges represent the routes between them. For example, in road networks the vertices correspond to road intersections and the edges correspond to the roads.

Finally, we have distribution networks, such as oil and gas pipelines, water and sewage lines, and the routes used by post office, package delivery and cargo companies. What we call vertices and edges in these networks depends on the actual network. For example, in a network of gas pipelines [CBB09], the vertices are their intersections including pumping,

switching, and storage facilities, and the edges are the pipelines. Distribution networks are typically *directed*, i.e., the edges run in one direction only.

Social Networks

Social networks are networks in which the vertices represent people or groups of people, and the edges represent some form of social interaction between them. For example, the edges might represent friendship between individuals, exchange of goods or money, communication patterns, or sexual relationships.

One way of building social networks is by surveying individuals in a population. An example of such a network is the sexual contact network [LEA03], an undirected network in which the vertices represent individuals and the edges indicate that the individuals corresponding to the vertices they connect had sex with each other.

Social networks are also built from direct observations of interactions between individuals over a period of time. Examples of such networks include: the famous Zachary Karate Club network [Zac77], an undirected network in which the vertices represent the members of a karate club in a north American university and the edges represent the friendship patterns among them; and behavioral networks of animals, such as monkeys [FGD06, Sad72, SG84], kangaroos [Gra73], and dolphins [Lus03], which are directed networks with vertices representing animals and the edges indicating that the animals corresponding to the vertices they connect engage in social behaviors, such as mutual grooming or courting.

Yet another way of building social networks is from archival records. A well-known example is the undirected network of the ruling families of Florence in the fifteenth century [PA93]. In recent years, with the widespread availability of computer and online services, many networks have been constructed from electronic records of various types.

These include: email networks [EMB02, TWH05], which are directed networks with vertices corresponding to people and the edges between them corresponding to email messages from one person to another; and telephone call networks [OSH07], which are directed networks with vertices representing telephone numbers and the edges between them representing calls from one number to another. Recent years have also seen the rapid emergence of online social networks [BE12] such as *Facebook* and *Twitter*. Weblogs [AG05] are another source of online social networks which are directed networks with vertices representing the blogs and the edges representing links from one blog to another.

An important kind of network built from archived data is the affiliation network, an undirected network in which the individuals are connected via comembership of groups of some kind. The best known example of such a network is the network of collaboration of film actors [ASB00, BA99, WS98] in which the vertices are actors and the groups to which they belong are the casts of films. This network is the basis for a well-known parlor game called “six degrees of Kevin Bacon” in which one attempts to connect an actor via a chain of intermediate actors to Kevin Bacon. This is similar to the small-world experiments [Mil67, TM69] conducted by Stanley Milgram in which he concluded that the “geodesic distance” between vertices in a network is quite small² for most pair of vertices. A network similar in spirit is the affiliation network of mathematicians [Gro02] in which the vertices represent individuals who have directly or indirectly coauthored a paper with the famous Hungarian mathematician Paul Erdős, and the edges denote that the individuals corresponding to the vertices they connect have coauthored at least one paper. The length of the shortest path from any vertex to the vertex representing Paul Erdős denotes the “Erdős number” of the individual corresponding to that vertex; Albert

²This is the origin of the idea of “six degrees of separation”, the notion that there are only about six steps between any two people in the world.

Einstein, for example, has an Erdős number of two, and the author of this thesis has an Erdős number of four.

Information Networks

Information networks are man-made networks consisting of data items linked together in some fashion. The best known example of an information network is the World Wide Web [BAJ99, BAJ00, BKM00], a directed network in which the vertices are web pages and the edges are the hyperlinks that allow users to navigate from one page to another.

Another example of an information network is a citation network [Pri65], a directed network in which the vertices are papers and there is an edge from paper i to j if i cites j in its bibliography. There are citations besides those among academic papers. Two of particular importance are citations between patents and legal opinions. In patent (legal opinion) networks [CD06, FJS07, FJ08, LCS07], the vertices are patents (legal opinions) and the edges are citations of one patent (legal opinion) by another. Citation networks are in general *acyclic*, i.e., the networks do not contain closed-loop paths.

A peer-to-peer network [RF02, SR04] is an undirected information network in which the vertices are computers containing information in the form of files and the edges between them are virtual links—they exist only in software—established for the purposes of sharing the contents of those files.

A recommender network [CCK06, Gru08] is a type of information network extensively used in commerce. Such a network is undirected and is typically represented as a *bipartite network*, a network with two types of vertices, one representing the products and the other representing the individuals, with edges connecting individuals to the products they buy or like.

Another type of information network, also bipartite in form, is the undirected keyword index network, used for searching large bodies of data. In such a network, there are two types of vertices representing words and documents, and there is an edge from a vertex denoting a word to a vertex denoting a document if the word appears in that document.

Biological Networks

Networks are widely used in biology as a convenient representation of interaction patterns between biological elements. Biochemical networks [JTA00, JMB01] represent interaction patterns and control mechanisms within the cell. Examples of biochemical networks include: the metabolic network, a directed bipartite network in which the vertices represent the chemicals (metabolites) produced or consumed in a metabolic reaction and the edges indicate which metabolites are substrates (inputs) and which are products (outputs); the protein-protein interaction network, an undirected network in which the vertices represent proteins and the edges denote that the proteins they connect interact with each other; and the directed gene regulatory network in which the vertices are proteins or equivalently the genes that code for them and an edge from one gene to another indicates that the former regulates the expression of the latter.

Networks are also used in the study of the brain and the central nervous system in animals [WS98, WST86]. Such networks are called neural networks. The vertices in these directed networks represent neurons, and an edge between two vertices indicates that the output of one neuron is an input to the other.

Ecological networks [HBR96, Mar91] are networks of ecological interactions between species. An example of such a network is a food web, a directed network that represents which species prey on which other species in a given ecosystem. The vertices in the network correspond to species and the edges to predator-prey interactions. Food webs are

typically acyclic. Host-parasite networks [New03] are networks of parasitic relationships between organisms, such as the relationship between a large-bodied animals and the insects and micro organisms that live on or within them. The other class of ecological networks is that of mutualistic networks, meaning networks of mutually beneficial interactions between species. Examples of such networks include networks of plants and animals that pollinate them or disperse their seeds, and networks of ant species and the plants that they protect and eat. These networks are undirected, and typically bipartite.

2.3 Network Concepts

Once we have a representation of a system as a network, we want to analyze it and measure some of its key properties. In this thesis, we are interested in how the structural properties of networks shape the long-term behavior of dynamical processes that take place on the networks, such as the evolution of cooperative behavior, and the spreading of infectious diseases in a population.

Since the human eye is extraordinarily adept at figuring out patterns, the simplest way to analyze the structure of a network is to draw a picture of it³. Direct visualization of networks, however, is only useful for small networks, i.e., networks with a small number of vertices. Many of the networks of interest have hundreds of thousands or even millions of vertices, which means that visualization is not of much help in their analysis and we need to employ other techniques to determine their structural features. Fortunately, network theory has developed a plethora of concepts that we can use to deduce the structural properties⁴ of networks, even in cases where useful visualization is impossible.

³Gephi (<https://gephi.org>) is one of many free software tools that makes visualization of complex networks possible.

⁴Free software programs such as NetworkX (<http://networkx.lanl.gov/>) and iGraph (<http://igraph.sourceforge.net/>) implement algorithms to evaluate these properties for arbitrary networks.

In this section we describe the concepts that are most pertinent to the study of real-world networks, and to our thesis.

Networks and Their Representation

A *network* is a collection of vertices joined by edges. We denote a network by Γ , at times with a suitable suffix to indicate the type of the network. We denote the set of vertices by V and the set of edges by E . The number of vertices is denoted by n and the number of edges by m , i.e., $n = |V|$ and $m = |E|$. For example, the neural network Γ^{neural} of the worm *C. elegans* [WS98, WST86] has $n = 297$ vertices and $m = 2148$ edges; this network is small enough that it can be visualized (see Figure 1).




Figure 1: The neural network of the worm *C. elegans*. The network has 297 vertices and 2148 edges. The color of a vertex provides a visual indication of its degree, with blue (red) representing vertices with low (high) degrees. The visualization was produced using Gephi (<https://gephi.org/>).

An *undirected* network, such as Γ^{neural} ⁵, is one in which the edges are symmetric, i.e., if there is an edge from i to j then there is also an edge from j to i . On the other hand, a

⁵The neural network of the worm *C. elegans*, like all neural networks, is directed, but for our purposes in this section we ignore the directions of the edges and treat the network as undirected.

directed network is one in which each edge has a direction, pointing *from* one vertex *to* another.

The networks we consider in this thesis are *unweighted*, i.e., networks in which there are no weights associated with the edges; and *simple*, i.e., networks that have at most an edge between any pair of vertices, and no edges that connect the vertices to themselves. Γ^{neural} , for example, is a simple, unweighted network.

A *planar network* is a network in which the vertices can be arranged on a plane in such a manner that the edges do not cross.

There are various ways of representing networks. Consider a simple undirected network with n vertices labeled $1, \dots, n$. An example is shown in Figure 2. If we denote




Figure 2: A simple undirected network.

the edge between vertices i and j by (i, j) then the entire network can be specified by giving the value of n and a list of all the edges. For example, the network in Figure 2 has $n = 6$ and edges $(1, 2)$, $(1, 3)$, $(1, 5)$, $(1, 6)$, $(2, 6)$, and $(3, 4)$. Though such a representation is commonly used to store networks on computers, it is cumbersome for mathematical analysis.

A better representation of a network is the *adjacency matrix*. The adjacency matrix A of a simple network is the matrix with elements A_{ij} such that

$$A_{ij} = \begin{cases} 1 & \text{if there is an edge between vertices } i \text{ and } j, \\ 0 & \text{otherwise.} \end{cases} \quad (2.1)$$

For example, the adjacency matrix of the network in Figure 2 is

$$A = \begin{pmatrix} 0 & 1 & 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}. \quad (2.2)$$

There are two points to notice about the adjacency matrix for a simple undirected network such as this one. First, the diagonal matrix elements are all zero, and second, the matrix is symmetric. The adjacency matrix of a directed network, on the other hand, has the element $A_{ij} = 1$ if there is an edge from j to i , and is asymmetric.

Degree

The *degree* of a vertex in an undirected network is the number of vertices connected to it. We denote the degree of vertex i by k_i . The set of vertices connected to i is denoted by $N(i)$. Thus, $k_i = |N(i)|$. For an undirected network with n vertices the degree can be written in terms of the adjacency matrix as

$$k_i = \sum_{j=1}^n A_{ij}. \quad (2.3)$$

Every edge in an undirected network has two ends and if there are m edges in total then there are $2m$ ends of edges. But the number of ends of edges is also equal to the sum

of the degrees of all the vertices, so

$$2m = \sum_{i=1}^n k_i, \quad (2.4)$$

or

$$m = \frac{1}{2} \sum_{i=1}^n k_i = \frac{1}{2} \sum_{j=1}^n A_{ij}. \quad (2.5)$$

The *average degree* d of a vertex in an undirected network is

$$d = \frac{1}{n} \sum_{i=1}^n k_i, \quad (2.6)$$

and combining this with Equation 2.4 we get

$$d = \frac{2m}{n}. \quad (2.7)$$

For example, the average degree of Γ^{neural} is $d = 14.465$.

In a directed network, the *in-degree* is the number of incoming edges on a vertex and the *out-degree* is the number of out-going edges. Therefore, the in- and out-degrees of the vertex i can be written as

$$k_i^{\text{in}} = \sum_{j=1}^n A_{ij}, \quad k_i^{\text{out}} = \sum_{i=1}^n A_{ij}. \quad (2.8)$$

The number of edges m in a directed network is equal to the total number of in-coming ends of edges at all vertices, or equivalently to the total number of out-going ends of edges, so

$$m = \sum_{i=1}^n k_i^{\text{in}} = \sum_{j=1}^n k_j^{\text{out}} = \sum_{ij} A_{ij}. \quad (2.9)$$

Thus the average in-degree d_{in} and the average out-degree d_{out} of every directed network are equal:

$$d_{\text{in}} = \frac{1}{n} \sum_{i=1}^n k_i^{\text{in}} = \frac{1}{n} \sum_{j=1}^n k_j^{\text{out}} = d_{\text{out}} = d. \quad (2.10)$$

Networks in which all vertices have the same degree are called *regular networks*. A regular network in which all vertices have degree d are called d -regular. An example of a regular network is a periodic lattice, such as a square lattice in which every vertex has degree four.

An important concept in network analysis is that of *centrality*, which tells us something about the most important or central vertices in a network. There are many possible definitions of importance, and correspondingly many centrality measures [WF94]. We will, for the purposes of this thesis, only need *degree centrality*, which is simply the degree of a vertex. Although a simple measure, it can be very illuminating. In a social network, for example, individuals who have connections to many others might have more influence [Zac77]. A non-social network example is the use of citation counts in the evaluation of scientific papers [Pri65].

The maximum possible number of edges in a simple undirected network is $\binom{n}{2} = \frac{n(n-1)}{2}$. The *connectance* or *edge density* ρ of a network is the fraction of these edges that are actually present in the network:

$$\rho = \frac{m}{\binom{n}{2}} = \frac{2m}{n(n-1)} = \frac{d}{n-1}, \quad (2.11)$$

where we have made use of Equation 2.7. The density lies strictly in the range $0 \leq \rho \leq 1$. Most of the networks one is interested in are sufficiently large that Equation 2.11 can safely be approximated as $\rho = \frac{d}{n}$. A network for which the density ρ tends to a constant as $n \rightarrow \infty$ is said to be *dense*. On the other hand, a network in which $\rho \rightarrow 0$ as $n \rightarrow \infty$ is said to be *sparse*. Γ^{neural} is a sparse network with $\rho = 0.049$. The Internet, the World Wide Web, and most friendship networks, are also sparse [New10].

Paths and Components

A *path* in a network is any sequence of vertices such that every consecutive pair of vertices in the sequence is connected by an edge in the network. In other words, a path is a route across the network that runs from vertex to vertex along the edges of the network.

The *length* of a path in a network is the number of edges traversed along the path. A *geodesic path*, also called the *shortest path*, is a path between two vertices such that no shorter path exists. The length of a geodesic path is often called the *geodesic distance* or the *shortest distance*. In mathematical terms, the geodesic distance between vertices i and j is the smallest value of l such that $[A^l]_{ij} > 0$. It must be noted that there may be several shortest paths between a pair of vertices, so geodesic paths are not unique. The *diameter* of a network is the length of the longest geodesic path between any pair of vertices in the network for which a path actually exists. The diameter of a network is denoted by ℓ . Most real-world networks have a very small diameter [WS98] and are called “small-world” networks, by analogy with the small-world phenomenon popularly known as six degrees of separation [Mil67, TM69]. For Γ^{neural} the diameter is $\ell = 2.455$.

It is possible to have no path at all between a given pair of vertices in a network. The network in this case is divided into subgroups of vertices, with no edges between any two subgroups. A network of this kind is said to be disconnected. Conversely, if there is a path from every vertex in a network to every other vertex then the network is said to be *connected*. The subgroups of vertices in a disconnected network are called the *components* of the network. Note that a disconnected network has an infinite diameter. The number of components in a network is denoted by κ . A network is connected if and only if $\kappa = 1$. We denote the fractional size of the largest component (also known as the giant component) relative to the total number of vertices in a network by σ . Thus, σ has

the maximum value of one when $\kappa = 1$ and the least value when $\kappa = n$. For Γ^{neural} , we have $\kappa = 1$ and therefore $\sigma = 1$.

A *tree* is a connected, undirected network that contains no closed loops, and is thus *acyclic*. If a network consists of several components that are all trees, then the entire network is called a *forest*. The most important property of trees is that there is exactly one path between any pair of vertices. Another useful property of trees is that a tree of n vertices always has exactly $n - 1$ edges, and the converse is also true.

Clustering

Transitivity or *clustering* is a very important property of networks. Social networks, for example, exhibit high clustering [HK02]. In mathematics, a relation “ \bullet ” is said to be transitive if $a \bullet b$ and $b \bullet c$ together imply $a \bullet c$. An example of such a relation is the *equality* relation; if $a = b$ and $b = c$ then $a = c$, so “ $=$ ” is transitive. In a network, the simple relation “pairs of vertices connected by an edge” is transitive, meaning if vertex u is connected to vertex v and v is connected to vertex w , then u is also connected to w . In common parlance such a relation is described as “the friend of a friend is also a friend”.

We can quantify the level of transitivity in a network as follows: if u knows v and v knows w , then we have a path uvw of two edges in the network. If u also knows w , we say that the path is closed—it forms a loop of length three, or a triangle in the network. We define the *clustering coefficient* $C^{(1)}$ as the fraction of paths of length two in the network that are closed.

$$C^{(1)} = \frac{\text{number of closed paths of length two}}{\text{number of paths of length two}}. \quad (2.12)$$

There is an alternative way to calculate the clustering coefficient of a network. We can define the local clustering coefficient for a single vertex i as

$$C_i = \frac{\text{number of pairs of neighbors of } i \text{ that are linked to each other}}{\text{number of pairs of neighbors of } i}. \quad (2.13)$$

That is, to calculate C_i we go through all distinct pairs of vertices that are neighbors of i in the network, count the number of such pairs that are connected to each other, and divide by the total number of pairs, which is $\frac{k_i(k_i-1)}{2}$. The value C_i represents the average probability that a pair of i 's friends are friends of one another. These local clustering coefficients can be used to calculate the global clustering coefficient $C^{(2)}$ as

$$C^{(2)} = \frac{1}{n} \sum_{i=1}^n C_i. \quad (2.14)$$

It must be noted that while $C^{(1)}$ and $C^{(2)}$ both quantify how clustered a network is, they do not define the same quantity, but are two different forms of clustering coefficients. Also, $0 \leq C^{(1)}, C^{(2)} \leq 1$ and $C^{(1)} = C^{(2)} = 1$ implies perfect transitivity while $C^{(1)} = C^{(2)} = 0$ implies no closed triads, which happens for various network topologies, such as trees. A network is said to be highly clustered if its clustering coefficient has a much higher value relative to its edge density, i.e., $C^{(1)}, C^{(2)} \gg \rho$. For Γ^{neural} the clustering coefficient is $C^{(2)} = 0.292$.

Homophily

Vertices in a network sometimes exhibit a strong tendency to associate with others whom they perceive as being similar to themselves. This tendency is called *homophily*. The corresponding networks are called *homophilic networks*. One also encounters, albeit rarely, the tendency for vertices to associate with others that are unlike themselves. The corresponding networks in this case are called *heterophilic networks*⁶.

⁶In network theory, the terms *assortativity* and *disassortativity* are used instead of *homophily* and *heterophily*, and the corresponding networks are referred to as *assortative* and *disassortative* networks. We

We consider homophilic association in a network according to scalar characteristics, such as age or income. These are characteristics whose values come in a particular order. A special case of homophilic association according to a scalar quantity, and one of particular interest for us, is that of association by degree. In a network that shows homophilic association by degree, the high-degree vertices are preferentially connected to other high-degree vertices, and low to low, creating a *core/periphery* structure in the network. Conversely, we could have heterophilic association by degree meaning that the high-degree vertices tend to be connected with low-degree vertices and vice versa, creating a *star-like* structure in the network. Social networks generally tend to be homophilic, while technological and biological networks tend to be heterophilic [New02].

The homophily coefficient h of a network can be defined as

$$h = \frac{\sum_{ij} (A_{ij} - \frac{k_i k_j}{2m}) k_i k_j}{\sum_{ij} (k_i \delta_{ij} - \frac{k_i k_j}{2m}) k_i k_j}, \quad (2.15)$$

where k_i is the degree of vertex i in the network and δ_{ij} is the Kronecker delta. The homophily coefficient by degree is essentially the Pearson correlation coefficient of degree between pairs of linked vertices. For a maximally homophilic network $h = 1$, for a maximally heterophilic network $h = -1$, and for an uncorrelated network $h = 0$. For Γ^{neural} the value of the coefficient is $h = -0.163$ meaning that the network is heterophilic.

Degree Distribution

One of the most fundamental and defining properties of a network is its *degree distribution*, which is the frequency distribution of the degree of vertices in the network. If p_k is the fraction of vertices in the network that have degree k , the quantities p_k represent

refrain from using the terms from network theory in order to avoid any confusion with the notion of assortative interactions that we introduce in Chapters 4 and 5, which is unrelated to network homophily that is being introduced here.

the degree distribution of the network. The value p_k can also be thought of as the probability that a randomly chosen vertex in the network has degree k . Another construct containing the same information as the degree distribution is the *degree sequence*, which is the set $\{k_1, k_2, \dots, k_n\}$ of degrees of the n vertices in the network. It must be noted that a knowledge of the degree distribution (or degree sequence) does not, in most cases, give the complete picture of the network, since for most choices of vertex degrees, there is more than one network with those degrees.

It is often illuminating to look at a plot of the degree distribution of a large network as a function of k . It turns out that in many real-world networks, most of the vertices have




Figure 3: Degree distribution of the neural network of the worm *C. elegans*. (a) Linear scale. (b) Logarithmic scale.

low degree but there is a significant “tail” to the distribution corresponding to vertices with substantially high degree [ASB00, New05]. In statistical parlance, we say that the degree distribution is *right-skewed*. A well-connected vertex is called a *hub*. For example, the degree distribution (see Figure 3(a)) of Γ^{neural} has a significant tail, with several hubs with degrees ranging from 60 to 134, which are relatively large numbers given that the

network only has 297 vertices. Hubs in a network play a vital role in processes such as the spreading of a disease [PV01b]. Such networks are also robust to uniformly random removal of vertices, but susceptible to removal which targets the hubs [CNS00, CEB00]. In other words, the networks with right-skewed degree distributions are robust to “error”, but vulnerable to “attack”.

When plotted on a logarithmic scale, a right-skewed degree distribution follows, roughly speaking, a straight line. Figure 3(b) shows the degree distribution of Γ^{neural} on a logarithmic scale. In mathematical terms

$$\ln p_k = -\alpha \ln k + c, \quad (2.16)$$

where α and c are constants and $\alpha > 0$. Taking the exponential of both sides we have

$$p_k = Ck^{-\alpha}, \quad (2.17)$$

where $C = e^c$ is another constant. Distributions of this form, varying as a power of k , are called *power-laws*. The degree distributions of many real-world networks, roughly speaking, follow a power-law [BA99]. The constant α is known as the *exponent* of the power-law. Values in the range $2 \leq \alpha \leq 3$ are typical, although values slightly outside this range are possible and are observed occasionally [AB02, New03]. It must be noted that degree distributions of real-world networks differ from those of true power-law distributions in the small- k regime.

Networks with power-law distributions are sometimes called *scale-free networks*. There are also many networks that are not scale-free and which have degree distributions with non power-law forms [ASB00], but the scale-free ones are of particular interest because of their intriguing properties. Telling scale-free networks from ones that are not is not always easy. The simplest strategy is to look at a histogram of the degree distribution on a log-log plot to see if we have a straight line. There are problems with this approach and alternative methods have been proposed [CSN09].

Power-laws show up in a wide variety of places, not just in networks. For example, they arise in the frequency of use of words in human languages [Zip49]; and in the distribution of biological taxa [WY22]. A review of some mathematical properties of power-laws can be found in [New05]. One that is relevant to this thesis is the m th moment of the degree distribution, which is given by

$$\langle k^m \rangle = \sum_{k=0}^{\infty} k^m p_k. \quad (2.18)$$

Suppose we have a degree distribution p_k that has a power-law tail for $k \geq k_{min}$. Then

$$\langle k^m \rangle = \sum_{k=0}^{k_{min}-1} k^m p_k + C \sum_{k=k_{min}}^{\infty} k^{m-\alpha}. \quad (2.19)$$

Since a power-law is a slowly varying function of k , for large k , we can approximate the second term by an integral as follows:

$$\begin{aligned} \langle k^m \rangle &\approx \sum_{k=0}^{k_{min}-1} k^m p_k + C \int_{k_{min}}^{\infty} k^{m-\alpha} dk \\ &= \sum_{k=0}^{k_{min}-1} k^m p_k + \frac{C}{m-\alpha+1} \left[k^{m-\alpha+1} \right]_{k_{min}}^{\infty}. \end{aligned} \quad (2.20)$$

For a given value of α all moments will diverge for which $m \geq \alpha - 1$. Of particular interest is the second moment $\langle k^2 \rangle$, which is finite if and only if $\alpha > 3$. However, most real-world networks have values of α in the range $2 \leq \alpha \leq 3$, which means the second moment should diverge. An important consequence of this is that power-law networks are prone to spreading of infectious diseases [BPV03].

We can calculate the m th moment from the degree sequence as

$$\langle k^m \rangle = \frac{1}{n} \sum_{i=0}^n k_i^m, \quad (2.21)$$

and since all the k_i are finite, so is the sum. The m th moment is infinite for an arbitrarily large network. On finite, practical networks the divergence is replaced by large finite values which can be treated as infinite for most practical purposes.

2.4 Network Models

Although it is possible to obtain network representations of real-world systems, this is often an arduous enterprise. Some of the procedures used to build the networks we encountered in Section 2.2 are highlighted in [New10]. Even if we manage to obtain the network representation of a real-world system, it is only a *fixed* snapshot of the system. Its properties, such as the number of vertices, edges, average degree, degree distribution, etc., are fixed. The network is of little use if, for example, our goal is to study the effects of these properties on the dynamical processes that take place on the network, in which case we ought to be able to *change* the network's properties freely. We would like to be able to build *model networks* having properties desired by the modeler. In this section we consider a few such model networks for in-silico construction of real-world-like networks.

2.4.1 Complete Network

A complete or well-mixed network, which we denote by Γ_{WM} , is the simplest of all network models. We start with n vertices and establish an edge between every distinct pair (i, j) of vertices, where $i \neq j$. The network thus has $m = \binom{n}{2} = \frac{1}{2}n(n-1)$ edges. A complete network has maximum connectance, i.e., $\rho = 1$. Since every vertex is connected to exactly $n-1$ other vertices, the average degree d of the network is $n-1$. The degree distribution p_k of a complete network with n vertices is given by

$$p_k = \begin{cases} 1 & \text{if } k = n-1, \\ 0 & \text{otherwise.} \end{cases} \quad (2.22)$$

We are interested in complete networks because in the large- n limit they can be used to model a well-mixed population of individuals—an infinitely large population in which every individual can potentially interact with every other individual in the population. The

dynamical processes we study in Chapters 4, 5, and 6 can be analyzed completely if they are modeled for a well-mixed population, i.e., we can obtain analytical results that accurately predict the long-term behavior of the processes. One can then verify these predictions by simulating the model on a complete network. As we shall see in later chapters, a complete network serves as a reasonable approximation even for a population in which individuals interact not with *every* other individual, but with *sufficiently large* number of individuals.

2.4.2 Erdős-Rényi Network

A *random network* is a model network in which some specific set of parameters take fixed values, but are random in other aspects. One example, called the Erdős-Rényi model [ER59, ER60, ER61], denoted by $G(n, p)$, is one in which the number of vertices n and the probability p of there being edges between them are fixed. We take n vertices and place an edge between each distinct pair of vertices with independent probability p ⁷. We denote an Erdős-Rényi model network by Γ_{ER} .

$G(n, p)$ is an ensemble of networks with n vertices in which each simple network G having m edges appears with probability

$$P(G) = p^m (1 - p)^{\binom{n}{2} - m}. \quad (2.23)$$

The probability of drawing a network with m edges from the ensemble is

$$P(m) = \binom{\binom{n}{2}}{m} p^m (1 - p)^{\binom{n}{2} - m}. \quad (2.24)$$

⁷To see the network formation process in action visit:

<http://ccl.northwestern.edu/netlogo/models/GiantComponent>

The average value of m is

$$\langle m \rangle = \sum_{m=0}^{\binom{n}{2}} mP(m) = \binom{n}{2} p. \quad (2.25)$$

We can use this result to calculate the average degree d of the network as

$$d = \sum_{m=0}^{\binom{n}{2}} \frac{2m}{n} P(m) = \frac{2}{n} \binom{n}{2} p = (n-1)p. \quad (2.26)$$

The probability that a vertex is connected to specific k other vertices and not to any of the other vertices is $p^k(1-p)^{n-1-k}$. There are $\binom{n-1}{k}$ ways to choose those k vertices, and hence the total probability of being connected to exactly k other vertices is

$$p_k = \binom{n-1}{k} p^k (1-p)^{n-1-k}, \quad (2.27)$$

which is a binomial distribution. Thus, $G(n, p)$ has a binomial degree distribution.

In case of a large network $n \rightarrow \infty$, so $p = \frac{d}{n-1}$ will become vanishingly small.

Therefore

$$\begin{aligned} \ln[(1-p)^{n-1-k}] &= (n-1-k) \ln\left(1 - \frac{d}{n-1}\right) \\ &\approx -(n-1-k) \frac{d}{n-1} \quad (\text{using Taylor series expansion}) \\ &\approx -d \\ &= -d \text{ as } n \rightarrow \infty. \end{aligned} \quad (2.28)$$

Taking exponentials of both sides, we find that $(1-p)^{n-1-k} = e^{-d}$ in the large- n limit.

Also for large n we have

$$\binom{n-1}{k} = \frac{(n-1)!}{(n-1-k)!k!} \approx \frac{(n-1)^k}{k!}, \quad (2.29)$$

and Equation 2.27 becomes

$$p_k = \frac{(n-1)^k}{k!} p^k e^{-d} = \frac{(n-1)^k}{k!} \left(\frac{d}{n-1}\right)^k e^{-d} = e^{-d} \frac{d^k}{k!}, \quad (2.30)$$

which is a Poisson distribution. Thus, in the large- n limit, $G(n, p)$ has a Poisson degree distribution.

While the Erdős-Rényi model is able to emulate some of the aspects of real-world networks, such as the small-world effect and the existence of a giant component, it falls short in many ways. The most serious shortcoming of the model is that it is unable to explain the heavy-tail degree distribution that is so pervasive among real-world networks. Moreover, the network formation process is unrealistic and too random. It is hard to imagine that the components of a system come into existence all at once, and then at random establish connections with other components. The World Wide Web, for example, is not formed in this manner.

2.4.3 Random Regular Network

There is a large class of situations where the interaction network is determined by the geographical location of the participating individuals [SF07]. Biological games offer excellent examples of such situations. The typical structure in such games is two-dimensional, and can be modeled by a regular two-dimensional lattice. Here we discuss a model [Wor99] that generates random d -regular networks, i.e., random regular networks in which each vertex has degree d , using the following procedure:

1. Start with a set U of nd points (nd even) partitioned into n groups, each of size d .
2. A pair of vertices is a *suitable pair* if a connection between them neither creates a self-edge nor a multi-edge. Repeat the following until no suitable pair can be found: choose two random points i and j in U and if (i, j) is suitable, pair i and j and delete them from U .

3. Create a network G with an edge from r to s if and only if there is a pair containing points in the r th and s th groups. If G is regular, output it. Otherwise, return to step 1.

We denote a random regular network by Γ_{RR} . It must be noted that a random regular network is in general not planar. Since every vertex in the random regular network is connected to exactly d other vertices, the average degree of the network is d . The degree distribution p_k of a random regular network with n vertices is given by

$$p_k = \begin{cases} 1 & \text{if } k = d, \\ 0 & \text{otherwise.} \end{cases} \quad (2.31)$$

2.4.4 Barabási-Albert Network

In the models we have seen so far the parameters of the network, such as the number of vertices and edges or the degree sequence were fixed at the start, i.e., chosen by the modeler to have some desired values. The models we will consider next belong to the class of *generative models*, i.e., they model the mechanism by which networks are created. The networks produced by these models turn out to be similar to the networks we observe in the real world, suggesting that similar generative mechanisms are probably at work in the formation of real-world networks. The best-known example of a generative network model is the “preferential attachment” model [BA99] for the growth of networks with power-law degree distributions.

As we have seen before, many networks observed in the real world have degree distribution that approximately follow a power-law, at least in the tail of the distribution. Examples include the Internet, the World Wide Web, citation networks, and some social and biological networks [BA99]. A natural question to ask therefore is how might a network come to have such a distribution? The question was first considered by Derek

John de Solla Price [Pri76], who proposed a simple and elegant model of network formation that gives rise to power-law degree distribution.

Price was interested in power-law degree distributions seen in citation networks of scientific papers [Pri65]. He was inspired by the work of economist Herbert Simon, who noted the occurrence of power-laws in a variety of (non-network) economic data, such as people's personal wealth distributions [Sim55]. Simon proposed the "rich-get-richer" effect as an explanation for wealth distribution, and showed mathematically that it can give rise to a power-law distribution. Price adapted Simon's methods to the network context. Price called his model the *cumulative advantage* model although it is more often known today by the name "preferential attachment"⁸, which was coined by Barabási and Albert [BA99].

In Price's model of a citation network, we assume that papers are published continually, and that newly appearing papers cite previously existing ones. Let the average number of papers cited by a newly appearing paper be d , i.e., d is the average out-degree.

The crucial assumption in Price's model is that a newly appearing paper cites previous ones chosen at random with probability proportional to the number of citations those previous papers already have plus a positive constant a ⁹. The model is thus a simplified representation of the citation process, and builds a directed, acyclic network. Let $p_q(n)$ be the fraction of vertices in the network that have in-degree q when the network contains n vertices. Let us examine what happens when we add a single new vertex to the network.

Consider one of the citations made by the new vertex. In the model, the probability that the citation is to a particular other vertex i is proportional to $q_i + a$. Since the citation

⁸To see the network formation process in action visit:

<http://ccl.northwestern.edu/netlogo/models/PreferentialAttachment>

⁹The constant a gives each paper a number of "free" citations to get it started in the race.

in question has to be *some* paper, this probability must be normalized. In other words the probability must be

$$\frac{q_i + a}{\sum_i (q_i + a)} = \frac{q_i + a}{n\langle q \rangle + na} = \frac{q_i + a}{n(d + a)}. \quad (2.32)$$

Since $\langle q \rangle = \frac{1}{n} \sum_i q_i$ and the average out-degree of the network is d by definition and the average in- and out-degrees of a directed network are equal, so $\langle q \rangle = d$.

Since each new paper cites d others on average and since there are $np_q(n)$ vertices with degree q in our network, the expected number of new citations to all vertices with degree q is

$$np_q(n) \times d \times \frac{q + a}{n(d + a)} = \frac{d(q + a)}{d + a} p_q(n). \quad (2.33)$$

When we add a single new vertex to our network of n vertices, the number of vertices in the network with in-degree q increases by one for every vertex previously of in-degree $q - 1$ that receives a new citation, thereby becoming a vertex of in-degree q . From Equation 2.33 the expected number of such vertices is

$$\frac{d(q - 1 + a)}{d + a} p_{q-1}(n). \quad (2.34)$$

Similarly, we lose one vertex of in-degree q every time such a vertex receives a new citation, thereby becoming a vertex of in-degree $q + 1$. The expected number of such vertices receiving citations is

$$\frac{d(q + a)}{d + a} p_q(n). \quad (2.35)$$

The number of vertices with in-degree q in the network after the addition of a single vertex is $(n + 1)p_q(n + 1)$. We can thus write down the *master equation* for the evolution of the in-degree distribution as

$$(n + 1)p_q(n + 1) = np_q(n) + \frac{d(q - 1) + a}{d + a} p_{q-1}(n) - \frac{d(q + a)}{d + a} p_q(n). \quad (2.36)$$

The equation is valid for all values of q except $q = 0$. When $q = 0$ there are no vertices of lower degree that can gain an edge to become a vertex of degree zero. On the other hand, we can gain a vertex of degree zero whenever a new vertex is added, since by hypothesis papers have no citations when they are first published. So the approximate equation for $q = 0$ is

$$(n+1)p_0(n+1) = np_0(n) + 1 - \frac{da}{d+a}p_0(n). \quad (2.37)$$

Now let us consider the limit of large network size $n \rightarrow \infty$ and calculate the asymptotic form of the degree distribution in this limit. Taking the limit $n \rightarrow \infty$ and using the shorthand $p_q = p_q(\infty)$, Equations 2.36 and 2.37 become

$$p_q = \frac{q+a-1}{q+a+1+\frac{a}{d}}p_{q-1} \quad \text{for } q \geq 1 \quad (2.38)$$

and

$$p_0 = \frac{1+\frac{a}{d}}{a+1+\frac{a}{d}} \quad \text{for } q = 0. \quad (2.39)$$

Solving the recurrence relation given by Equations 2.38 and 2.39, we get the correct expression for general q as

$$p_q = \frac{(q+a-1)(q+a-2)\dots q}{(q+a+1+\frac{a}{d})\dots(a+2+\frac{a}{d})} \frac{1+\frac{a}{d}}{a+1+\frac{a}{d}}. \quad (2.40)$$

The Gamma function, $\Gamma(x) = \int_0^\infty t^{x-1}e^{-t}dt$, has a property that $\Gamma(x+1) = x\Gamma(x)$ for all $x > 0$, iterating which we get

$$\frac{\Gamma(x+n)}{\Gamma(x)} = (x+n-1)(x+n-2)\dots x. \quad (2.41)$$

Using this result in Equation 2.40 we can write

$$\begin{aligned} p_q &= \left(1 + \frac{q}{d}\right) \frac{\Gamma(q+a)\Gamma(q+1+\frac{a}{d})}{\Gamma(a)\Gamma(q+a+2+\frac{a}{d})} \\ &= \frac{\Gamma(q+a)\Gamma(2+\frac{a}{d})}{\Gamma(q+a+2+\frac{a}{d})} \frac{\Gamma(a+1+\frac{a}{d})}{\Gamma(a)\Gamma(1+\frac{a}{d})} \\ &= \frac{B(q+a, 2+\frac{a}{d})}{B(a, 1+\frac{a}{d})}, \end{aligned} \quad (2.42)$$

where $B(x, y) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}$ is Euler's beta function. Note that the Equation 2.42 is valid for $q \geq 0$.

For large q and fixed a and d , we can rewrite the beta function using Stirling's approximation for the gamma function

$$\Gamma(x) \approx \sqrt{2\pi} e^{-x} x^{x-\frac{1}{2}}, \quad (2.43)$$

thus giving us

$$B(x, y) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)} \approx \frac{e^{-x} x^{x-\frac{1}{2}}}{e^{-(x+y)} (x+y)^{x+y-\frac{1}{2}}} \Gamma(y). \quad (2.44)$$

But

$$(x+y)^{x+y-\frac{1}{2}} = x^{x+y-\frac{1}{2}} \left[1 + \frac{y}{x}\right]^{x+y-\frac{1}{2}} \approx x^{x+y-\frac{1}{2}} e^y. \quad (2.45)$$

Therefore,

$$B(x, y) \approx \frac{e^{-x} x^{x-\frac{1}{2}}}{e^{-(x+y)} x^{x+y-\frac{1}{2}} e^y} \Gamma(y) = x^{-y} \Gamma(y). \quad (2.46)$$

In other words, the beta function $B(x, y)$ falls off as a power-law for large values of x , with exponent y .

Applying this finding to 2.42 we find that for large values of q the degree distribution of the network goes as $p_q \propto (q+a)^{-\alpha}$ or simply

$$p_q \propto q^{-\alpha}, \quad (2.47)$$

when $q \gg a$, where the exponent α is

$$\alpha = 2 + \frac{a}{d}. \quad (2.48)$$

Thus Price's model for a citation network gives rise to a degree distribution with a power-law tail, very much in agreement with the degree distribution of real citation

networks. Note that $\alpha > 2$ since $a, d > 0$. For real citation networks, the exponent of the power-law is usually around $\alpha = 3$, which can be achieved by setting $a = d$.

The Barabási-Albert model [BA99] is similar to Price's model, though not identical, since it is a model of an undirected network. In this model, we again add vertices one by one to a growing network and each vertex connects to a suitably chosen set of previously existing vertices. The connections, however, are now undirected and the number of connections made by each vertex is exactly d , unlike Price's model where the number was required only to take an average value of d . This implies that d must be an integer. Connections are made to vertices with probability precisely proportional to the vertices' current degree. We will denote the degree of a vertex i by k_i . There are no vertices with $k < d$. The smallest degree in the network is always $k = d$. We next show that this is a special case of Price's model. Imagine that we give each edge added to the network a direction, running from the vertex just added to previously existing vertex it connects to. In this way we convert our network into a directed network in which each vertex has out-degree d . The total degree k_i of a vertex in the original undirected network is the sum of the vertex's in-degree and out-degree, which is $k_i = q_i + d$ where q_i is the in-degree as before.

But given that the probability of an edge attaching to a vertex is simply proportional to k_i , it is thus also proportional to $q_i + d$, which is the same as Price's model if we make the particular choice $a = d$. Thus the distribution of in-degree in this directed network is the same as for Price's model with $a = d$, which we find from Equation 2.42 as

$$p_q = \frac{B(q+d, 3)}{B(d, 2)}. \quad (2.49)$$

To get the distribution of the total degree we then simply replace $q + d$ by k to get

$$p_k = \begin{cases} \frac{B(k,3)}{B(d,2)} & \text{for } k \geq d \\ 0 & \text{for } k < d. \end{cases} \quad (2.50)$$

Simplifying Equation 2.50 using $B(x,y) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}$, we get

$$p_k = \frac{3d(d+1)}{k(k+1)(k+2)} \quad \text{for } k \geq d. \quad (2.51)$$

In the limit where k becomes large, we have

$$p_k \propto k^{-3}. \quad (2.52)$$

Thus the Barabási-Albert model generates a degree distribution with a power-law tail that always has an exponent $\alpha = 3$. This model is attractive, since it does not require the offset parameter a . Also, we can write the degree distribution without using special functions. However, the model can no longer match exponents observed in real networks, being restricted to just a simple exponent value $\alpha = 3$. We will denote a Barabási-Albert model network by Γ_{BA} .

The average degree of the network generated by Barabási-Albert model is $2d$. This is direct consequence of the fact that d edges are added when each vertex enters the network, and each edge contributes twice to the total degree.

A slight variation of the Barabási-Albert model, in which each new vertex entering the network attaches to a fixed number of existing vertices chosen uniformly at random, is called the *growing random network model* [CHK01]. This model produces networks with exponential degree distribution. Such networks do show up in the real world. For example, the electric power grid network of Southern California [ASB00] has a cumulative degree distribution that is exponential.

2.4.5 Power-Law Clustered Network

Many real-world networks, besides properties such as the small-world effect, and power-law degree distribution, also exhibit high clustering. The networks generated by the standard Barabási-Albert model do not have any clustering. Here we look at an extension [HK02] of the model that produces small-world networks with power-law degree distribution, and with high clustering. We denote these networks by Γ_{BA}^C . The standard Barabási-Albert model involves the following steps:

1. *Initial condition* - start with a network Γ_0 containing n_0 vertices and 0 edges.
2. *Growth* - at each time step t , add one vertex v to the current vertex set V_{t-1} to give V_t .
3. *Preferential attachment (PA)* - each of d edges of v is then attached to an existing vertex $w \in V_{t-1}$ with probability proportional to its degree, i.e., the probability for a vertex w to be attached to v is

$$p_w = \frac{k_w}{\sum_{v \in V_{t-1}} k_v}. \quad (2.53)$$

In the Barabási-Albert model, the growth step is iterated n times, and for each growth step the PA step is iterated d times for d edges of the newly added vertex v . In order to get high clustering the Barabási-Albert model algorithm is modified by adding the following additional step:

4. *Triad formation (TF)* - if an edge between v and w was added in the previous PA step, then add one more edge from v to a randomly chosen neighbor of w . If there remains no pair to connect, i.e., if all neighbors of w are already connected to v , do a PA step instead.

When a vertex v with d edges is added to the existing network, we first perform one PA step, and then perform a TF step with probability p_t or a PA step with probability

$1 - p_t$. The average number d_t of the TF triads per added vertex is then given by $d_t = (d - 1)p_t$, which is considered as the control parameter of the model. If $d_t = 0$ then the model reduces to the standard Barabási-Albert model.

2.4.6 Power-Law Homophilic Network

The existence of homophilic association among vertices in a network is an important property of real-world networks [New02]. Many social networks are homophilic, while many technological and biological networks are heterophilic. Such association tendencies in a network influence phenomena, such as the spreading of information or infectious diseases on the network.

Here we look at an extension [XS04] of the Barabási-Albert model that can generate undirected networks with homophilic or heterophilic association, and some clustering. We denote such networks by Γ_{BA}^h . The model is governed by a single parameter p . The model starts with a seed network built using the Barabási-Albert model and having the desired average degree. At each step two edges of the network are chosen at random, so that four vertices, in general with different degrees, are considered. The four vertices are ordered by their degrees. Then with probability p , the edges are rewired in such a way that one edge connects the two vertices with the larger degrees, and the other connects the two vertices with smaller degrees. With probability $1 - p$ they are randomly rewired. In the case when one or more of these new edges already exist in the new network, the step is discarded and a new pair of edges is selected. This avoids appearance of multiedges. The repeated application of the rewiring step leads to a homophilic version of the original network. One way of terminating the algorithm is when the desired level of homophily is reached, or when a predetermined number of rewiring steps have been performed. Note that the

rewiring procedure preserves the degree of every vertex, and thus does not alter the degree distribution of the seed network.

The model can also generate networks with heterophilic association. To achieve this, the rewiring step is modified so that with probability p , the highest-degree vertex is connected to the lowest-degree vertex and the remaining two vertices are connected to each other.

2.5 Basic Properties of Select Empirical Networks

Model networks do not entirely match the properties of real-world networks. We would like to know if the subtle differences that exist between the model and real-world networks have an effect on the outcome of dynamical processes that take place on the networks. Therefore, in Chapters 4, 5, and 6, where we consider various dynamical processes on networks, in addition to using model networks, we employ a selection of real-world networks¹⁰ as substrates for the processes.

We consider the following undirected real-world networks (see Table 2 for a listing of their basic properties):

- A symmetrized snapshot of the structure of the Internet at the level of autonomous systems, reconstructed from BGP tables posted by the University of Oregon Route Views Project, by Mark Newman. We denote this technological network by Γ^{internet} .
- A network of coauthorships between scientists who posted preprints on the High-Energy Theory E-Print Archive between January 1, 1995 and December 31, 1999 [New01]. We denote this social network by $\Gamma^{\text{hepcoauth}}$.

¹⁰Only the largest component of the networks are considered.

- A network of coauthorships between scientists who posted preprints on the Astrophysics E-Print Archive between January 1, 1995 and December 31, 1999. [New01]. We denote this social network by $\Gamma^{\text{astrocoauth}}$.
- The largest component of the New Orleans Facebook network of user-to-user links [VMC09]. We denote this social network by Γ^{fb} .
- A snapshot of the Gnutella peer-to-peer file sharing network from August 4, 2002 [LKF07, RF02]. We denote this information network by Γ^{p2p} .
- A network of protein-protein interactions in the yeast *S. cerevisiae* [JTA00, JMB01]. We denote this biological network by Γ^{protein} .

Γ	Category	n	m	ρ	d	κ	σ	ℓ	$C^{(2)}$	h
Γ^{internet}	Technological	22963	48436	0.0002	4.219	1	1.000	3.842	0.230	-0.198
$\Gamma^{\text{hepcoauth}}$	Social	5835	13815	0.0008	4.735	1	1.000	7.026	0.506	0.185
$\Gamma^{\text{astrocoauth}}$	Social	17903	197031	0.0012	22.011	1	1.000	4.194	0.633	0.201
Γ^{fb}	Social	13262	308420	0.0035	46.512	1	1.000	3.096	0.263	0.091
Γ^{p2p}	Information	10876	39994	0.0007	7.355	1	1.000	4.636	0.006	-0.013
Γ^{protein}	Biological	4941	8795	0.0007	3.560	1	1.000	11.409	0.079	0.025

Table 2: Basic properties of select empirical networks. These properties include: number of vertices n ; number of edges m ; connectance ρ ; average degree d ; number of connected components κ ; fractional size of the largest component σ ; diameter ℓ ; clustering coefficient $C^{(2)}$; and homophily coefficient h .

CHAPTER 3

EVOLUTIONARY DYNAMICS

Nothing in biology makes sense except in the light of evolution.

- Theodosius Dobzhansky

Another curious aspect of the theory of evolution is that everybody thinks he understands it.

- Jacques Monod

3.1 Introduction

The main objective of this thesis is to study the effects of network structure on the dynamics of evolutionary processes, i.e., how the structural properties of networks affect the long-term behavior of processes that take place on the networks. In subsequent chapters we propose models to describe processes such as evolution of cooperative behavior, and the spreading of infectious diseases in a population. We use a network to represent the population, with the vertices of the network corresponding to the individuals in the population, and the connection patterns in the network specifying who can interact with whom in the population.

The interactions among individuals in the dynamical processes we consider are such that their outcomes to an individual depend not only on the state or action of that individual, but also on the states or actions of the individual's neighbors, i.e., the individuals with whom the given individual interacts. For example, when an animal fights another animal over a territory, whether or not it uses a passive or an aggressive fighting strategy depends on the strategy used by its opponent [SP73, Smi74]; similarly, in an epidemic outbreak in a population, whether or not an individual is infected, and the extent

to which the individual is infected, depends not only on the state (susceptible or immune) of the individual but also on the state (infected or not) of the individual's neighbors and extent to which they are infected [AM92]. Such interdependent interactions are best understood in the language of *game theory*, which provides a mathematical framework for describing and analyzing strategic interactions.

Classical game theory¹ assumes that the interacting decision makers are rational, i.e., they act so as to maximize their payoffs. This assumption, at times, is untenable even when the decision makers are sentient human beings [GH01]. If we intend to apply game-theoretic ideas to non-human organisms—and our models are indeed applicable in such settings—then the assumption about rational behavior is not viable and must be relaxed. A refinement of the theory, called evolutionary game theory [SP73, Smi74], which incorporates ideas from Darwinian evolution [Dar69], does not assume the interacting individuals to be rational. The theory has been enormously successful in evolutionary biology, and more recently, has also been applied in social contexts [SN99]. Though classical and evolutionary game theory differ quite radically in how they perceive strategic interactions, i.e., games, there are some deep connections between the two theories.

Besides providing a framework for describing situations in which individuals are not overtly making decisions, evolutionary game theory offers us an additional bonus. It allows us to build an explicit model of the process by which the frequency of strategies change in a population and study the properties of the evolutionary dynamics within that model. Once a model of the population dynamics is specified [Zee80], all of the standard concepts from the theory of dynamical systems, such as attractors and stability, can be brought to bear.

¹The original formulation of game theory is termed “classical” to contrast it with its relatively modern incarnation called “evolutionary game theory”.

As much as we would like, it is not always easy to validate the results of our models by conducting experiments in the real world. We instead rely heavily on computer simulations to illustrate and verify our models. We propose agent-based models [Bon02] that can be simulated on a computer, and which emulate how individuals in a population might behave if our models were implemented in reality. We use agent-based simulations not only to verify the predictions made by our models, but also to simulate scenarios for which mathematical analysis is either very difficult, or simply impossible. We thus rely on an in-silico exploration of the results of our models.

In this chapter we first provide a brief introduction to classical and evolutionary game theories. We then look at games from a dynamical systems standpoint which will help us analyze games with both discrete and continuous strategies. Finally, we look at agent-based models for simulating evolutionary dynamics on computers.

3.2 Classical Game Theory

Classical game theory originated in the 1920s, when John von Neumann, a brilliant Hungarian mathematician, wanted to develop a scientific approach to bluffing in the game of poker. He observed that playing poker involves a strategic component that other parlor games, such as chess, do not. In most ordinary games the players know everything that they could know about what has happened in the game so far. In poker, however, the players only possess information about their own cards and know nothing about the other players' cards, and this gives the players the chance to misrepresent their strengths and weaknesses. Von Neumann realized that a science of how to best act in these situations would apply to a huge swath of decision-making problems in society. In 1944, von Neumann and economist Oskar Morgenstern wrote a classic book [NMK07] on game theory in which they developed, for the first time, a systematic theory of games.

Around 1950, a flurry of research in classical game theory was carried out in places such as Princeton University and the RAND corporation, a nonprofit global policy think tank. One of the most important contributions came from the American mathematician John Nash, who invented a simple but powerful concept, now called the “Nash equilibrium”², which is used to analyze the outcome of strategic interactions among decision makers.

Ideas from classical game theory are applicable not only in contexts that are literally games, such as chess or poker, but also in settings that are not usually considered games [EK10]. Examples include: pricing of a new product when other firms have similar new products; deciding how to bid in an auction; choosing a route on the Internet or through a transportation network; deciding whether to adopt an aggressive or passive stance in international relations; and choosing whether to use performance-enhancing drugs in a professional sport.

Since classical game theory crucially depends on ideas from *rational choice theory*, a framework for understanding and formally modeling social and economic behavior, we will need to explore some of those ideas before we get into the details of classical game theory proper.

Suppose we are faced with the problem of decision making. One approach to the problem is to determine the desired outcome and then to behave in a way that leads to that result. But this is not always possible. An alternative approach is to list the courses of action that are available and pick the outcome that maximizes something of value. The second approach is called “making an optimal decision”. Rational choice theory provides a mathematical framework for studying such decision problems. It is important to realize that these decision problems are *non interactive*, i.e., the outcome to the individual making

²Nash’s PhD thesis led to a single-page paper [Nas50] in the Proceedings of the National Academy of Sciences in 1950, which earned him the 1994 Nobel Prize for economics.

a decision depends on that individual's decision alone and not on those made by other individuals who are also confronted with the same decision problem. For example, an individual's preference for a particular flavor of ice cream does not depend on what other individuals might prefer.

A choice of behavior in a decision problem is called an *action*. The set of alternative actions available to an individual is denoted by A , which can either be discrete, such as $\{a_1, a_2, \dots\}$, or continuous, such as the unit interval $[0, 1]$.

A *payoff* is a function $\pi : A \rightarrow \mathbb{R}$ that associates a numerical value to every action $a \in A$. An action a^* is an *optimal action* if

$$\pi(a^*) \geq \pi(a) \quad \forall a \in A, \tag{3.1}$$

or, equivalently

$$a^* \in \operatorname{argmax}_{a \in A} \pi(a). \tag{3.2}$$

Theorem 3.1.³ The optimal action is unchanged if payoffs are altered by an *affine transformation*, i.e., a transformation that changes payoffs $\pi(a)$ into payoffs $\pi'(a)$ according to the rule

$$\pi'(a) = \alpha\pi(a) + \beta, \tag{3.3}$$

where α and β are constants independent of a , and $\alpha > 0$.

The set of possible outcomes is denoted by $\Omega = \{\omega_1, \omega_2, \dots\}$. We write $\omega_i \succ \omega_j$ if an individual *strictly* prefers ω_i over ω_j for some i, j ; $\omega_i \sim \omega_j$ if an individual is *indifferent* between the two; and $\omega_i \succeq \omega_j$ if an individual either prefers ω_i to ω_j or is indifferent between the two. An individual is called *rational* if her preferences for outcomes satisfy the following conditions:

³The proofs of theorems we state in this chapter, unless specified otherwise, can be found in [Web07].

1. (Completeness) Either $\omega_i \succeq \omega_j$ or $\omega_j \succeq \omega_i$; and
2. (Transitivity) If $\omega_i \succeq \omega_j$ and $\omega_j \succeq \omega_k$, then $\omega_i \succeq \omega_k$.

An individual has a personal *utility function* $u(\omega)$ that gives her the utility for any outcome ω . The outcome may be numeric, such as, an amount of money or number of vacation days, or something less tangible, such as degree of happiness. The utility function assigns a number to the outcome and encapsulates everything about an outcome that is important to the individual. Formally, the utility function is a function $u : \Omega \rightarrow \mathbb{R}$ such that

$$\begin{aligned} u(\omega_i) > u(\omega_j) &\Leftrightarrow \omega_i \succ \omega_j, \text{ and} \\ u(\omega_i) = u(\omega_j) &\Leftrightarrow \omega_i \sim \omega_j. \end{aligned} \tag{3.4}$$

An immediate consequence of this definition is that a rational individual should seek to maximize her utility. The relation between the utility function and the payoff function is straightforward. If choosing an action a produces an outcome $\omega(a)$, then $\pi(a) = u(\omega(a))$.

We only consider decision problems with *ordinal payoffs*, i.e., problems in which individuals make decisions only based on the relative orderings of the payoffs, and not their absolute differences. For example, an individual deciding between actions a_1 and a_2 , will prefer a_1 over a_2 as long as $\pi(a_1) > \pi(a_2)$, no matter how small the difference may be.

Individuals might also randomize their behavior. A general *randomized* behavior β is given by a vector of probabilities $p(a)$ with which each available action $a \in A$ is chosen and

$$\sum_{a \in A} p(a) = 1. \tag{3.5}$$

The set of all possible randomizing behaviors for a given decision problem is denoted by B , and the payoff for using a behavior β is given by

$$\pi(\beta) = \sum_{a \in A} p(a)\pi(a). \quad (3.6)$$

An *optimal behavior* β^* is one for which

$$\pi(\beta^*) \geq \pi(\beta) \quad \forall \beta \in B, \quad (3.7)$$

or, equivalently

$$\beta^* \in \operatorname{argmax}_{\beta \in B} \pi(\beta). \quad (3.8)$$

The *support* of a behavior β is the set $A(\beta) \subseteq A$ of all the actions for which β specifies $p(a) > 0$, i.e., $A(\beta) = \{a | a \in A \text{ and } p(a) > 0\}$.

Theorem 3.2. If β^* is an optimal behavior with support A^* , then $\pi(a) = \pi(\beta^*) \forall a \in A^*$.

A consequence of Theorem 3.2 is that if a randomizing behavior is optimal, then two or more actions are optimal as well. Therefore, in a non-interactive decision problem, randomization is not necessary to achieve maximal payoff, except may be as a tie breaker in choosing between two or more equally acceptable actions.

So far we have been concerned with non-interactive decision problems in which the payoff to an individual only depends on her own action and not on the actions of other individuals around her. We now shift our attention to *interactive decision problems*, i.e., decision problems involving multiple individuals in which the outcome to an individual depends not only on her own decisions but also on the other individuals' decisions. We illustrate this with a couple of examples.

Consider two firms A and B that manufacture a similar product. Any pricing decision that each of the two companies makes will affect not only its own profits but also the

profits of the other company. If, for example, A decides to lower the price of its product, some of B 's customers will switch to using A 's product. As a result, A 's profits will increase while B 's profits will drop, thus forcing B to also lower the price of its product. Thus the action of one company depends on the action of the other company.

As another example, consider two candidates A and B from different political parties, contesting in an election. Any reform that either of the candidates promises to bring about if elected to office, will affect not only her chances of winning the election but also the other candidate's chances. For example, if A promises to increase educational spending, then some of B 's followers who care about education will gravitate towards supporting A , thus forcing B to respond in kind. Here again, the action of one depends on that of the other.

In each of the two examples, we would like to predict the optimal behavior. In the case of the two companies, we would like to know the price that each company should charge for its product that is its best response to the price charged by the other company. In the case of the election candidates, we would like to know the reforms that each candidate must optimally promise. Classical game theory makes such predictions possible.

Classical game theory is a mathematical framework for dealing with situations, called games, in which rational—as defined by rational choice theory—decision makers, called players, interact with one another, and in which each participant's satisfaction with the outcome depends not only on her own action but also on the actions taken by everyone else.

A *strategy* is a rule for choosing an action at every point that a decision might have to be made. A *pure strategy* s is one in which there is no randomization. The finite set of all pure strategies is denoted by S . A *mixed strategy* σ specifies the probability $p(s)$ with which each of the pure strategies $s \in S$ is used. Suppose the set of n pure strategies is

$S = \{s_1, \dots, s_n\}$. Then a mixed strategy σ can be represented as a vector of probabilities

$$\sigma = (p(s_1), \dots, p(s_n)). \quad (3.9)$$

A pure strategy can then be represented as a vector in which all the entries are zero except one. For example, the strategy s_2 can be represented as

$$s_2 = (0, 1, 0, \dots, 0). \quad (3.10)$$

And a mixed strategy can be represented as a linear combination of pure strategies, i.e.,

$$\sigma = \sum_{s \in S} p(s)s. \quad (3.11)$$

The set of all mixed strategies is denoted by Σ . The *support* of a mixed strategy σ is the set $S(\sigma) \subseteq S$ of all the pure strategies for which σ specifies $p(s) > 0$.

Consider a game with n players indexed by $i \in \{1, \dots, n\}$. Let S_1, \dots, S_n denote the pure-strategy sets of the n players, and let $\Sigma_1, \dots, \Sigma_n$ denote their mixed strategy sets. The *pure-strategy profile* \hat{s} is a tuple of pure strategies given by $\hat{s} = (s_1, \dots, s_n)$, where $s_i \in S_i$, and it denotes the pure strategies played by each of the n players. Note that $\hat{s} \in \hat{S} = S_1 \times \dots \times S_n$. We use \hat{s}_{-i} to denote the pure-strategy profile that excludes the strategy s_i of player i , i.e., $\hat{s}_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n) \in \hat{S}_{-i} = S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_n$. Similarly, the *mixed-strategy profile* $\hat{\sigma}$ is a tuple of mixed strategies given by $\hat{\sigma} = (\sigma_1, \dots, \sigma_n)$, where $\sigma_i \in \Sigma_i$, and it denotes the mixed strategies played by each of the n players. Note that $\hat{\sigma} \in \hat{\Sigma} = \Sigma_1 \times \dots \times \Sigma_n$. We use $\hat{\Sigma}_{-i}$ to denote the mixed-strategy profile that excludes the strategy σ_i of player i , i.e., $\hat{\sigma}_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n) \in \hat{\Sigma}_{-i} = \Sigma_1 \times \dots \times \Sigma_{i-1} \times \Sigma_{i+1} \times \dots \times \Sigma_n$.

For each choice of strategies, each player receives a *payoff* that depends on the strategies used by everyone. Thus, the payoff to player i is given by

$$\pi_i : S_1 \times \dots \times S_n \rightarrow \mathbb{R}. \quad (3.12)$$

We will use the abbreviation $\pi_i(s_i; \hat{s}_{-i})$ for $\pi_i(s_1, \dots, s_n)$ when $n > 2$. If the players use mixed strategies $\sigma_1, \dots, \sigma_n$, then the payoff to player i is given by

$$\pi_i(\sigma_1, \dots, \sigma_n) = \sum_{s_1 \in \mathcal{S}_1} \sum_{s_2 \in \mathcal{S}_2} \dots \sum_{s_n \in \mathcal{S}_n} p_1(s_1) p_2(s_2) \dots \pi_i(s_1, s_2, \dots, s_n), \quad (3.13)$$

where $p_i(s)$ is the probability of choosing strategy s by player i . We will use the abbreviation $\pi_i(\sigma_i; \hat{\sigma}_{-i})$ for $\pi_i(\sigma_1, \dots, \sigma_n)$ when $n > 2$.

In order to reason about the behavior of players in a game, we assume that everything that a player cares about is summarized in the player's payoff. We also assume that each player chooses a strategy that maximizes her own payoff given her beliefs about the strategies used by the other players, and that each player succeeds in selecting the optimal strategy.

In a game of *perfect information*, a player knows everything about the structure of the game. She knows her own list of possible strategies, who the other players are, the strategies available to them, and what is their payoff for any choice of strategies. An example of such a game is chess. Contrarily, a game of *imperfect information* is one in which not all the information is available to every player. Poker is an example of such a game. A *static game* is one in which a single decision is made by each player, and each player has no knowledge of the decision made by the other players before making her own decision. Such games are also referred to as *simultaneous games* because any actual order in which the decisions are made is irrelevant. The popular children's game of rock-paper-scissors (also known as Rochambeau) is an example of a static game⁴. There are situations in which decisions are made at various times, with at least some of the earlier choices being public knowledge when the later decisions are being made. Such games are called *dynamic games*. Tic-tac-toe is an example of a dynamic game. In this thesis we only deal with static games with perfect information.

⁴The rock-paper-scissors game has interesting parallels in nature [KRF02, SL96].

A static game is usually represented in *normal* or *strategic* form. A finite n -player normal-form game consists of

1. a set of n players indexed by $i \in \{1, \dots, n\}$;
2. a set S_i of pure strategies assigned to each player i ; and
3. a payoff $\pi_i(s_i; \hat{s}_{-i})$ assigned to player i for every possible combination of pure strategies used by all players.

A natural way of representing the payoffs in a normal-form game with two players is using an $m \times n$ matrix where m and n are the number of pure strategies in the strategy sets of the two players. We illustrate this with an example.

Example 3.1. (Simplified poker) Consider a simplified game of poker played between two players with a deck of just two cards: an Ace and a King. Each player puts two dollars into the pot and one of them (say 1) is nominated to be the *dealer* and the other (say 2) to be the *recipient*. Player 1 deals herself a card at random and after seeing it can either fold, or raise by putting two more dollars into the pot. In the former case Player 2 wins, and in the latter case, Player 2 can either fold, in which case she loses, or can put two more dollars into the pot and demand to see the remaining card, winning if it is Ace and losing if it is King. The winner gets the contents of the pot.

The strategies for the dealer are: 1) always raise; 2) raise with Ace and fold with King; 3) fold with Ace and raise with King; and 4) always fold. The strategies for the recipient are: 1) stay if Player 1 raises; and 2) fold if Player 1 raises. We assume that Player 1 is equally likely to deal herself an Ace or a King.

We therefore have a game with two players: Player 1 and Player 2. The pure-strategy set for Player 1 is $S_1 = \{1, 2, 3, 4\}$, and that for Player 2 is $S_2 = \{1, 2\}$. The payoffs π are

given by the 4×2 matrix

$$\pi = \begin{matrix} & \begin{matrix} 1 & 2 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \\ 4 \end{matrix} & \begin{pmatrix} (0,0) & (2,-2) \\ (1,-1) & (0,0) \\ (-3,3) & (0,0) \\ (-2,2) & (-2,2) \end{pmatrix} \end{matrix}. \quad (3.14)$$

The matrix in Equation 3.14 is called a *bimatrix*. Player 1, called row-player, chooses a strategy i along the rows and Player 2, called column-player, chooses a strategy j along the columns. The pair of numbers (π_1, π_2) at (i, j) denote the payoffs to Player 1 and Player 2 respectively. For example, if Player 1 chooses the strategy 1 (always raise) and Player 2 chooses strategy 2 (fold if Player 1 raises) then Player 1 wins and Player 2 loses the two dollars that Player 2 put into the pot. \square

The game in Example 3.1 is an *asymmetric* game because there is an asymmetry in the roles of the two players; for example, $\pi_1(1,2) = 2 \neq \pi_2(2,1) = -1$. The payoffs in such two-player games are usually represented using bimatrices. When a game does not have such role asymmetries, it is called a *symmetric* game; in such games the players have the same strategy sets and their payoffs are symmetric. A two-player symmetric game can be represented using a regular matrix in which the elements are single numbers instead of pairs of numbers⁵.

A two-player normal-form game is a *constant-sum game* if there exists a constant c such that $\pi_1(s,t) + \pi_2(s,t) = c \forall s \in S_1, t \in S_2$. A *zero-sum game* is a special case of a constant sum game, with $c = 0$. The game in Example 3.1 is a zero-sum game. These are games of pure conflict, i.e., in such games one player's gain must come at the expense of the other player.

⁵We will use a bimatrix even for symmetric games.

Example 3.2. (Rock-paper-scissors game) The popular children's game of rock-paper-scissors provides yet another example of a two-player zero-sum game. Two children simultaneously make the shape of the item with one of their hands: rock (R) beats scissors (S); scissors beats paper (P); and paper beats rock. If both players choose the same item, then the game is a draw.

We therefore have a two-player game in which both players have the same pure-strategy set $S = \{R, P, S\}$, and the payoffs π are given by the 3×3 bimatrix

$$\pi = \begin{matrix} & \begin{matrix} R & P & S \end{matrix} \\ \begin{matrix} R \\ P \\ S \end{matrix} & \begin{pmatrix} (0,0) & (-1,1) & (1,-1) \\ (1,-1) & (0,0) & (-1,1) \\ (-1,1) & (1,-1) & (0,0) \end{pmatrix} \end{matrix}. \quad (3.15)$$

Notice that the game is symmetric, since $\pi_1(s,t) = \pi_2(t,s) \forall s,t \in S$. □

Now that we have defined what normal-form games are and what strategies are available to the players involved, a natural question to ask is how does one reason about such games. In non-interactive decision problems, the key notion is that of optimal action, i.e., an action that maximizes the individual's expected payoff for a given environment in which she operates. In interactive settings, the environment includes other individuals, all of whom are also hoping to maximize their payoffs. Thus the notion of optimal strategy for a given individual is not meaningful; the best strategy depends on the choices of others. Game theorists deal with this problem by identifying certain subsets of outcomes, called *solution concepts*. We describe a few such solution concepts.

- A strategy σ_i for player i is *strictly dominated* by σ'_i if

$$\pi_i(\sigma'_i; \hat{\sigma}_{-i}) > \pi_i(\sigma_i; \hat{\sigma}_{-i}) \quad \forall \hat{\sigma}_{-i} \in \hat{\Sigma}_{-i}, \quad (3.16)$$

i.e., whatever the other players do, player i is always better off using σ'_i rather than σ_i . A strategy σ_i for player i is *weakly dominated* by σ'_i if

$$\begin{aligned} \pi_i(\sigma'_i; \hat{\sigma}_{-i}) &\geq \pi_i(\sigma_i; \hat{\sigma}_{-i}) \quad \forall \hat{\sigma}_{-i} \in \hat{\Sigma}_{-i}, \text{ and} \\ \exists \hat{\sigma}'_{-i} \in \hat{\Sigma}_{-i} \text{ such that } \pi_i(\sigma'_i; \hat{\sigma}'_{-i}) &> \pi_i(\sigma_i; \hat{\sigma}'_{-i}). \end{aligned} \quad (3.17)$$

- A *Nash equilibrium* for an n -player game is a mixed-strategy profile

$\hat{\sigma}^* = (\sigma_1^*, \dots, \sigma_n^*)$ such that

$$\pi_i(\sigma_i^*; \hat{\sigma}_{-i}^*) \geq \pi_i(\sigma_i; \hat{\sigma}_{-i}^*) \quad \forall \sigma_i \in \Sigma_i. \quad (3.18)$$

In other words, given the strategy adopted by the players, none of them could do strictly better (i.e., increase their payoff) by adopting another strategy.

- A *strict Nash equilibrium* for an n -player game is a mixed-strategy profile

$\hat{\sigma}^* = (\sigma_1^*, \dots, \sigma_n^*)$ such that

$$\pi_i(\sigma_i^*; \hat{\sigma}_{-i}^*) > \pi_i(\sigma_i; \hat{\sigma}_{-i}^*) \quad \forall \sigma_i \neq \sigma_i^*. \quad (3.19)$$

It is clear that a strictly dominated strategy is never a best response to any strategy, whereas a weakly dominated strategy may be a best response to some strategy. This is why weakly dominated strategies may appear in Nash equilibria but strictly dominated strategies do not. We can solve some games by eliminating strictly dominated strategies as follows:

- We start with any n -player game, find all the strictly dominated strategies, and delete them.
- We then consider the reduced game in which there may be strategies that are now strictly dominated, despite not having been strictly dominated in the full game. We find these and delete them.

- We continue this process, repeatedly finding and removing strictly dominated strategies until none can be found.

The procedure described above works because of the fact that the set of Nash equilibria of the original game coincides with the set of Nash equilibria for the final reduced game, consisting only of strategies that survive iterated deletion of strictly dominated strategies. There is a problem with the iterated deletion of dominated strategies when it comes to deleting weakly dominated strategies, since the solution may depend on the order in which the strategies are eliminated.

- A strategy σ'_i for player i is a *best response* to $\hat{\sigma}_{-i}$ if

$$\sigma'_i \in \operatorname{argmax}_{\sigma_i \in \Sigma_i} \pi_i(\sigma_i; \hat{\sigma}_{-i}), \quad (3.20)$$

and a *unique best response* if σ'_i is unique.

- There is an alternative definition of a Nash equilibrium, which focusses on the strategies rather than the payoffs. The mixed-strategy profile $\hat{\sigma}^* = (\sigma_1^*, \dots, \sigma_n^*)$ is a *Nash equilibrium* if for all i , the mixed-strategy σ_i^* is the best response to σ_{-i}^* ; and a *strict Nash equilibrium* if it is a unique best response.
- For a two-player game in which each player has the same set Σ of mixed strategies, (σ^*, σ^*) is a *symmetric Nash equilibrium* if σ^* is a best response to itself; and a *strict symmetric Nash equilibrium* if it is a unique best response to itself.

We next state the fundamental theorem of classical game theory (also known as Nash's existence theorem) [Nas50], which states that every strategic-form game has at least one Nash equilibrium. For a proof of the theorem see [Mye97].

Theorem 3.3. (Nash's existence theorem) A finite normal-form game, i.e., a game with finite number of players and finite number of pure strategies for each player, has at least one Nash equilibrium involving pure or mixed strategies.

As Theorem 3.3 suggests, an n -player game with a finite number of pure strategies for each player, may have both pure-strategy and/or mixed-strategy equilibria. As a result, one should first check all pure outcomes (given by an n -tuple of pure strategies) to see which, if any, form equilibria. Then, to check for any mixed-strategy equilibria, we need to look for mixed strategies $\sigma_1, \dots, \sigma_n$, each of which is a best response to the others. For example, in a two-player two-strategy game, we look for mixing probabilities p_1 and p_2 that are best responses to each other. If there is a mixed-strategy equilibrium, we can determine Player 2's strategy (p_2) from the requirement that Player 1 randomizes. Player 1 will only randomize if her pure strategies have equal expected payoff (see Theorem 3.5 below). This equality of expected payoffs for Player 1 gives an equation that we can solve to determine p_2 . The same process gives an equation to solve for determining Player 1's strategy (p_1). If both of the obtained values, p_1 and p_2 , are strictly between 0 and 1, and are thus legitimate mixed strategies, then we have a mixed-strategy equilibrium.

It is only the equilibrium of strategies, and not the payoffs, which are of interest. It is possible to convert a difficult calculation into a simple one by means of a generalized affine transformation of payoffs in a game, without affecting its Nash equilibria, as stated by the following theorem.

Theorem 3.4. The set of Nash equilibria is unaffected if the payoffs $\pi_i(s_i; \hat{s}_{-i})$ are altered to $\pi'_i(s_i; \hat{s}_{-i})$ by a generalized affine transformation given by

$$\pi'_i(s_i; \hat{s}_{-i}) = \alpha_i \pi_i(s_i; \hat{s}_{-i}) + \beta(\hat{s}_{-i}) \quad \forall s_i \in S_i, \quad (3.21)$$

where $\alpha_i > 0$ and $\beta(\hat{s}_{-i}) \in \mathbb{R}$.

Theorem 3.5. (Equality of payoffs) Let $(\sigma_1^*, \dots, \sigma_n^*)$ be a Nash equilibrium, and let S_i^* be the support of σ_i^* . Then $\pi_i(s; \hat{\sigma}_{-i}^*) = \pi_i(\sigma_i^*; \hat{\sigma}_{-i}^*) \forall s \in S_i^*$.

Theorem 3.5 says that all strategies $s \in S_i^*$ give the same payoff as the randomizing strategy. Why does player i randomize then? The answer: if player i were to deviate from this strategy, then $\hat{\sigma}_{-i}^*$ would no longer be a best response from other players, and the equilibrium would disintegrate. This is why randomizing strategies are important for games, in a way that they were not for the single player decision problems.

It is interesting to classify outcomes in a game not just by their equilibrium properties, but also by whether they are “good for the society”. A choice of strategies is *Pareto-optimal*, named after the Italian economist Vilfredo Pareto, if there is no other choice of strategies in which all players receive payoffs at least as high, and at least one player receives a strictly higher payoff. A stronger condition is *social optimality*. A choice of strategies is a *social welfare maximizer* or *Hicks optimal*, named after economist John Hicks, if it maximizes the sum of the players’ payoffs. Outcomes that are Hicks optimal must also be Pareto-optimal: if such an outcome were not Pareto-optimal, there would be a different outcome in which all payoffs were at least as large, and one was larger—and this would be an outcome with a larger sum of payoffs. However, a Pareto-optimal outcome need not be Hicks optimal.

A Nash equilibrium is *Pareto-superior* to another Nash equilibrium if in the first equilibrium some players receive higher payoff and none get lower payoff compared to the second equilibrium. A Nash equilibrium is *payoff dominant* if it is Pareto-superior to all other Nash equilibria in the game, and *risk dominant* if players are more likely to choose the strategy if they are uncertain about the actions of the other players.

We next look at some examples to see how these solution concepts apply to various games. We start with a game with a unique pure-strategy Nash equilibrium.

Example 3.3. (Prisoner’s dilemma game) This game is motivated by the following story due to Albert Tucker, who was John Nash’s PhD advisor. Suppose two suspects have been apprehended by the police and are being interrogated in separate rooms. The police strongly suspect that these two individuals are responsible for a robbery, but there is not enough evidence to convict either of them of the robbery. However, they both resisted arrest and can be charged with that lesser crime, and be given a short-term jail sentence . Each of the suspects is told the following story. “If you confess, and your partner does not, then you will be released and your partner will be charged with the crime. Your confession will be sufficient to convict her of the robbery and she will be given a long-term jail sentence. If you both confess, then we do not need either of you to testify against the other, and you will both be convicted of the robbery, although in this case your sentence will be less because of your guilty plea. Finally, if neither of you confesses, then we cannot convict either of you of the robbery, so we will charge each of you with resisting arrest. Your partner is being offered the same deal. Do you want to confess?”

The story suggests a symmetric two-player game in which the players are the two suspects (call them Player 1 and Player 2). Each player has two strategies: not confess (C) or confess (D). The rank ordering of the payoffs to the two players is as follows:

$$\begin{aligned}\pi_1(D, C) &> \pi_1(C, C) > \pi_1(D, D) > \pi_1(C, D); \\ \pi_2(C, D) &> \pi_2(C, C) > \pi_2(D, D) > \pi_2(D, C).\end{aligned}\tag{3.22}$$

A particular assignment of values to the payoffs, while preserving their rank ordering and the essence of the story described above, gives us the following payoff matrix⁶ for the

⁶The general form of the payoff matrix for the prisoner’s dilemma game is written as

$$\pi = \begin{array}{c} C \quad D \\ \begin{array}{c} C \\ D \end{array} \left(\begin{array}{cc} (R, R) & (S, T) \\ (T, S) & (P, P) \end{array} \right),\end{array}$$

game:

$$\pi = \begin{array}{c} C \quad D \\ C \begin{pmatrix} (2,2) & (0,3) \\ (3,0) & (1,1) \end{pmatrix}. \end{array} \quad (3.23)$$

We can find the Nash equilibrium for the game by eliminating strictly dominated strategies. For Player 1, C is strictly dominated by D , and hence can be eliminated. Now if Player 1 plays D , then Player 2 is better off playing D , since $\pi_2(D,D) = 1$ which is greater than $\pi_2(D,C) = 0$. Therefore, (D,D) is a strict Nash equilibrium for the game, i.e., the two suspects must confess. To show that (D,D) is indeed a Nash equilibrium we must show that $\pi_1(D,D) \geq \pi_1(\sigma,D)$ for any (possibly mixed) strategy σ , and a similar condition for $\pi_2(D,D)$. Let $\sigma = pC + (1-p)D$, where $p \in [0, 1]$.

$$\begin{aligned} \pi_1(\sigma,D) &= p\pi(C,D) + (1-p)\pi(D,D) \\ &= p(0) + (1-p)(1) \\ &= 1-p \\ &= 0 \quad \forall \sigma \neq D. \\ &< \pi_1(D,D). \end{aligned} \quad (3.24)$$

Thus, (D,D) is a strict Nash equilibrium.

It is interesting to note that the two suspects would have been better off if they had both not confessed, in which case their payoffs would have been 2 each. However, even if the two suspects were personally committed to not confess (C)—hoping to achieve the outcome where both receive a payoff of 2—and even if one individual knew that the other was doing this, the other has the incentive to confess (D) so as to achieve a still higher

where T, R, P , and S , respectively called the *temptation*, *reward*, *punishment*, and *sucker* payoff, satisfy the condition $T > R > P > S$.

payoff of 3 for herself. Thus for the prisoner's dilemma game, the Hicks optimal outcome of (C, C) is at odds with the Nash equilibrium. The outcomes (D, D) , (D, C) , (C, D) are all Pareto-optimal. \square

The next two examples are of games with multiple (pure- and mixed-strategy) Nash equilibria.

Example 3.4. (Chicken game) In this game two cars are headed for each other at high speed. The loser is whoever chickens out first by swerving, and the winner is one who stays on track, i.e., does not swerve. There is a cost associated with swerving. If neither driver swerves, there is substantially higher cost of a collision.

The story suggests a symmetric two-player game in which the players are the car drivers (call them Player 1 and Player 2). Each player has two strategies: swerve (C) or not swerve (D). The rank ordering of the payoffs to the two players is as follows:

$$\begin{aligned}\pi_1(D, C) &> \pi_1(C, C) > \pi_1(C, D) > \pi_1(D, D); \\ \pi_2(C, D) &> \pi_2(C, C) > \pi_2(D, C) > \pi_2(D, D).\end{aligned}\tag{3.25}$$

A particular assignment of values to the payoffs, while preserving their rank ordering and the essence of the story described above, gives us the following payoff matrix for the game:

$$\pi = \begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} (2, 2) & (1, 3) \\ (3, 1) & (0, 0) \end{pmatrix}.\end{array}\tag{3.26}$$

If Player 2 chooses C , then the best response for Player 1 is to choose D because $\pi_1(D, C) = 3 > \pi_1(C, C) = 2$. On the other hand, if Player 2 chooses D , then the best response for Player 1 is to choose C because $\pi_1(C, D) = 1 > \pi_1(D, D) = 0$. A similar

argument also holds for Player 2. Therefore, (C, D) and (D, C) are both Nash equilibria. These equilibria agree with our intuition that if one player cooperates, i.e., swerves, the best thing for the other player to do is to not cooperate, i.e., not swerve.

The game also allows a mixed-strategy equilibrium, which we find as follows. Suppose Player 1 chooses C with probability p_1 , and Player 2 chooses C with probability p_2 . If Player 1 plays the pure strategy C against Player 2's mixed strategy, her payoff is $p_2 + 1$, and if she plays pure strategy D , her payoff is $3p_2$. Since Player 1 will randomize only if her pure strategies have equal expected payoff, $p_2 + 1 = 3p_2 \implies p_2 = \frac{1}{2}$. The situation is symmetric when we consider things from Player 2's point of view, i.e., we can show that $p_1 = \frac{1}{2}$. Thus, $((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$ is a mixed-strategy Nash equilibrium for the game. To prove that this is the case, we must show that $\pi_1(\sigma_1, \sigma_2) \geq \pi_1(\sigma'_1, \sigma_2)$ where $\sigma_1 = \sigma_2 = \frac{1}{2}C + \frac{1}{2}D$ and $\sigma'_1 = pC + (1-p)D$, where $p \in [0, 1]$, and a similar condition for Player 2.

$$\begin{aligned}\pi_1(\sigma_1, \sigma_2) &= \frac{1}{4}\pi(C, C) + \frac{1}{4}\pi(C, D) + \frac{1}{4}\pi(D, C) + \frac{1}{4}\pi(D, D) \\ &= \frac{1}{4}(2) + \frac{1}{4}(1) + \frac{1}{4}(3) + \frac{1}{4}(0) \\ &= \frac{3}{2}.\end{aligned}\tag{3.27}$$

$$\begin{aligned}\pi_1(\sigma'_1, \sigma_2) &= \frac{p}{2}\pi_1(C, C) + \frac{p}{2}\pi(C, D) + \frac{1-p}{2}\pi(D, C) + \frac{1-p}{2}\pi(D, D) \\ &= \frac{p}{2}(2) + \frac{p}{2}(1) + \frac{1-p}{2}(3) + \frac{1-p}{2}(0) \\ &= \frac{3}{2}.\end{aligned}\tag{3.28}$$

We can similarly show that $\pi_2(\sigma_1, \sigma_2) = \pi_1(\sigma_1, \sigma'_2) \forall \sigma'_2$. Therefore, (σ_1, σ_2) is a Nash equilibrium, but is not strict. □

Example 3.5. (Coordination game) Consider the following story from the writings of the famous Swiss philosopher Rousseau. Suppose that two people are out hunting; if they

work together, then they can catch a stag (which would be the highest-payoff outcome), but on their own each can catch a hare. The tricky part is that if one hunter tries to catch a stag on her own, she will get nothing, while the other one can still catch a hare.

The story suggests a symmetric two-player game, called the *stag-hunt* game, in which the players are the two hunters (call them Player 1 and Player 2). Each player has two strategies: hunt stag (C) or hunt hare (D). The rank ordering of the payoffs to the two players is as follows:

$$\begin{aligned}\pi_1(C, C) &> \pi_1(D, C) > \pi_1(D, D) > \pi_1(C, D); \\ \pi_2(C, C) &> \pi_2(C, D) > \pi_2(D, D) > \pi_2(D, C).\end{aligned}\tag{3.29}$$

A particular assignment of values to the payoffs, while preserving their rank ordering and the essence of the story described above, gives us the following payoff matrix for the game:

$$\pi = \begin{array}{c} \begin{array}{cc} C & D \\ C & \begin{pmatrix} (5, 5) & (0, 4) \\ (4, 0) & (2, 2) \end{pmatrix} \\ D & \end{array} \end{array}.\tag{3.30}$$

If Player 2 chooses C , then the best response for Player 1 is to also choose C because $\pi_1(C, C) = 5 > \pi_1(D, C) = 4$. On the other hand, if Player 2 chooses D , then the best response for Player 1 is to also choose D because $\pi_1(D, D) = 2 > \pi_1(C, D) = 0$. A similar argument also holds for Player 2. Therefore, (C, C) and (D, D) are both Nash equilibria. It must be noted that (C, C) is payoff dominant while (D, D) is risk dominant. Rational choice would seem to be the payoff dominant Nash equilibrium, but in evolutionary approaches, the players may end up in the risk dominant Nash equilibrium.

The game also allows a mixed-strategy equilibrium $((\frac{2}{3}, \frac{1}{3}), (\frac{2}{3}, \frac{1}{3}))$, identified using the procedure followed in Example 3.4. To prove that this is indeed a Nash equilibrium, we

must show that $\pi_1(\sigma_1, \sigma_2) \geq \pi_1(\sigma'_1, \sigma_2)$ where $\sigma_1 = \sigma_2 = \frac{2}{3}C + \frac{1}{3}D$ and $\sigma'_1 = pC + (1-p)D$, where $p \in [0, 1]$, and a similar condition for Player 2.

$$\begin{aligned}\pi_1(\sigma_1, \sigma_2) &= \frac{4}{9}\pi(C, C) + \frac{2}{9}\pi(C, D) + \frac{2}{9}\pi(D, C) + \frac{1}{9}\pi(D, D) \\ &= \frac{4}{9}(5) + \frac{2}{9}(0) + \frac{2}{9}(4) + \frac{1}{9}(2) \\ &= \frac{10}{3}.\end{aligned}\tag{3.31}$$

$$\begin{aligned}\pi_1(\sigma'_1, \sigma_2) &= \frac{2p}{3}\pi(C, C) + \frac{p}{3}\pi(C, D) + \frac{2(1-p)}{3}\pi(D, C) + \frac{1-p}{3}\pi(D, D) \\ &= \frac{2p}{3}(5) + \frac{p}{3}(0) + \frac{2(1-p)}{3}(4) + \frac{1-p}{3}(2) \\ &= \frac{10}{3}.\end{aligned}\tag{3.32}$$

We can similarly show that $\pi_2(\sigma_1, \sigma_2) = \pi_1(\sigma_1, \sigma'_2) \forall \sigma'_2$. Therefore, (σ_1, σ_2) is a Nash equilibrium, but is not strict.

The stag-hunt game is an example of a *coordination game* since the two players' shared goal is really to coordinate on the same strategy. The game is an *unbalanced coordination game* since the two coordinated alternatives are unequal in terms of their payoffs to the two individuals.

A coordination game in which the alternatives are equal is called a *balanced coordination game*. A motivating story for such a game is: suppose two individuals are each preparing slides for a joint project presentation. Each individual is unable to reach the other by phone, and needs to start working on the slides right away. Each can decide whether to prepare her half of the slides using Beamer or using Prezi. It will be much easier to merge the slides if both use the same software. In this two-player symmetric game the players are the two individuals preparing the presentation; let us call them players 1 and 2. Choosing Beamer (*B*) or choosing Prezi (*P*) form the two strategies. The

rank ordering of the payoffs to the two players is as follows:

$$\begin{aligned}\pi_1(P,P) &> \pi_1(B,B) > \pi_1(B,P) = \pi_1(P,B); \\ \pi_2(P,P) &> \pi_2(B,B) > \pi_2(P,B) = \pi_2(B,P).\end{aligned}\tag{3.33}$$

Notice that in this game, the players get the least payoff when they miscoordinate, unlike the stag-hunt game in which miscoordination is not the lowest-payoff outcome.

A variation of the presentation game in which the two individuals do not agree on which software they prefer, is called the *battle of the sexes* game, because of the following motivating story. A husband and wife want to see a movie together, and they need to choose between a romantic comedy and an action movie. They want to coordinate on their choice, but the (romance, romance) alternative gives a higher payoff to one of them while the (action, action) alternative gives a higher payoff to the other. Thus, we have a two-player game; let us call the husband Player 1 and the wife Player 2. The two strategies are R for watching the romance movie and A for watching the action movie. The rank ordering of the payoffs to the two players is as follows:

$$\begin{aligned}\pi_1(A,A) &> \pi_1(R,R) > \pi_1(A,R) = \pi_1(R,A); \\ \pi_2(R,R) &> \pi_2(A,A) > \pi_2(R,A) = \pi_2(A,R).\end{aligned}\tag{3.34}$$

□

The next two examples are of games that have no pure-strategy equilibria, but allow an equilibrium in mixed strategies.

Example 3.6. (Matching pennies game) Two players each place a penny on a table, either “heads up” (strategy H) or “tails up” (strategy T). If the pennies match, Player 1 wins the pennies; if the pennies differ, then Player 2 wins the pennies. The payoff structure of this

two-player asymmetric game is given by the matrix

$$\pi = \begin{array}{cc} & \begin{array}{cc} H & T \end{array} \\ \begin{array}{c} H \\ T \end{array} & \begin{pmatrix} (+1, -1) & (-1, +1) \\ (-1, +1) & (+1, -1) \end{pmatrix}. \end{array} \quad (3.35)$$

Clearly, this is a game in which the two players have completely opposing interests: one player only wins a penny when the other loses a penny. The game is therefore a zero-sum game.

We can easily check that there is no pure strategy pair that is a Nash equilibrium: (H, H) is not an equilibrium because Player 2 should switch to T ; (H, T) is not an equilibrium because Player 1 should switch to T ; (T, H) is not an equilibrium because Player 1 should switch to H ; and finally, (T, T) is not an equilibrium because Player 2 should switch to H .

Let us consider the mixed strategies $\sigma_1 = (p, 1 - p)$ for Player 1 and $\sigma_2 = (q, 1 - q)$ for Player 2. That is, Player 1 plays “heads” with probability p and Player 2 plays “heads” with probability q . The payoff to Player 1 is

$$\begin{aligned} \pi_1(\sigma_1, \sigma_2) &= pq - p(1 - q) - (1 - p)q + (1 - p)(1 - q) \\ &= 1 - 2q + 2p(2q - 1). \end{aligned} \quad (3.36)$$

Clearly, if $q < \frac{1}{2}$ then Player 1’s best response is to choose $p = 0$, i.e., “play tails”. On the other hand, if $q > \frac{1}{2}$ then Player 1’s best response is to choose $p = 1$, i.e., “play heads”. If $q = \frac{1}{2}$ then every mixed (and pure) strategy is a best response.

Now consider the payoff to Player 2.

$$\begin{aligned} \pi_2(\sigma_1, \sigma_2) &= -pq + p(1 - q) + (1 - p)q - (1 - p)(1 - q) \\ &= -1 + 2p + 2q(1 - 2p). \end{aligned} \quad (3.37)$$

Clearly, if $p < \frac{1}{2}$ then Player 2's best response is to choose $q = 1$, i.e., "play heads". On the other hand, if $p > \frac{1}{2}$ then Player 2's best response is to choose $q = 0$, i.e., "play tails". If $p = \frac{1}{2}$ then every mixed (and pure) strategy is a best response.

So the only pair of strategies for which each is best response to the other is $\sigma_1^* = \sigma_2^* = (\frac{1}{2}, \frac{1}{2})$. That is, $(\sigma_1^*, \sigma_2^*) = ((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$ is a Nash equilibrium and the expected payoff to each player is $\pi_1(\sigma_1^*, \sigma_2^*) = \pi_2(\sigma_1^*, \sigma_2^*) = 0$. □

Example 3.7. (Rock-paper-scissors game) Reconsider the rock-paper-scissors game from Example 3.2 with the payoff matrix

$$\pi = \begin{array}{c} \\ R \\ P \\ S \end{array} \begin{array}{ccc} R & P & S \\ \left(\begin{array}{ccc} (0,0) & (-1,1) & (1,-1) \\ (1,-1) & (0,0) & (-1,1) \\ (-1,1) & (1,-1) & (0,0) \end{array} \right) \end{array}. \quad (3.38)$$

We can go through a similar exercise as in Example 3.6 to show that no pure strategy profile (s, t) for $s, t \in \{R, P, S\}$ is a Nash equilibrium for this game. The game has a unique mixed Nash equilibrium given by $((\frac{1}{3}, \frac{1}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}))$, the justification for which can be found in [Nou07]. □

There are several possible interpretations of mixed-strategy equilibria in real-world situations [EK10]: when the players are genuinely playing a game, the players may be actively randomizing their actions; sometimes the mixed strategies are better viewed as proportions within a population; and finally, since a Nash equilibrium is often best thought of as an equilibrium in beliefs, if each player believes that her partner will play according to a particular Nash equilibrium, then she too will want to play according to it.

Now we turn our attention to some of the drawbacks of classical game theory. We assumed that in reasoning about behavior in strategic games, each player actually

succeeds in selecting the optimal strategies. In simple games played by experienced players, this seems like a reasonable assumption. But if the players are inexperienced and the games are complex, then the assumption is absolutely unwarranted. It is also unreasonable to assume that the individuals will always behave rationally. Recall that an individual is considered rational only if their preferences satisfy the completeness and transitivity conditions. It is very easy, for example, for an individual's preferences to violate the transitivity property. Consider an individual who is asked to choose between 255 shades of the color gray, ordered from the lightest to the darkest, and numbered from 1 to 255. She is quite likely to be indifferent between nearby shades since they are very hard to tell apart, but she might have a preference for shades that are farther apart. Thus, individuals cannot be expected to play games the way classical game theory predicts they should, and in fact they do not [GH01].

How can we then apply the ideas from classical game theory to non-human organisms? The theory has been applied, and quite successfully so, not just to higher primates [SPN04], but also to fish [Mil87], bats [Wil84], dung beetles [CFR09], bacteria [KRF02], viruses [TC99], and many other organisms. Clearly, rationality cannot be assumed in any of these cases. So the theory needs to be refined before it is applied to these contexts. Evolutionary game theory is such a refinement and we discuss it next.

3.3 Evolutionary Game Theory

In this section we explore evolutionary game theory, which shows that the basic ideas of game theory can be applied even to situations in which no individual is overtly reasoning or even making explicit decisions. Evolutionary game theory deals with settings in which individuals can exhibit different forms of behavior, and helps us analyze the ones that can persist in populations and the ones that can be driven out by others.

As the name suggests, evolutionary game theory has been applied most widely in the area of evolutionary biology, the domain in which the idea was first articulated by the British evolutionary biologist and geneticist John Maynard Smith and the American population geneticist George Robert Price [SP73]. Evolutionary biology is based on the idea that an organism's genes, i.e., its genotype, largely determines its observable characteristics, i.e., its phenotype, and hence its fitness in a given environment. Organisms that are most fit will tend to produce more offspring, causing the genes that provide greater fitness to increase in frequency in the population. Thus, fitter genes tend to win over time because they provide higher rates of reproduction.

In order to apply classical game theory in an evolutionary context, John Maynard Smith proposed the following changes:

- *Strategy* - In classical game theory, players have strategy sets from which they choose particular strategies. In evolutionary game theory, *populations* have strategy sets, which correspond to *phenotypic* variants, and *individuals* in the population have one of these different phenotypes as strategies, which they use in strategic interactions. Since the phenotype is determined by the individual's genotype, the strategy that an individual uses is genetically coded, and could have one of many different forms. For example, a behavioral strategy, the growth forms of a plant, the sex ratio of an organism and so on.
- *Payoffs* - In classical game theory, the payoffs are defined in terms of utility, which is an elusive concept. In evolutionary game theory, the notion of payoff is very clear. The payoff to an individual i using some strategy s_i against another individual j using some strategy s_j is defined to be the change in the Darwinian fitness of i as a result of the interaction, i.e., the payoff to i is the change in the expected number of offspring for i as a result of the interaction.

- *Equilibrium* - In classical game theory the Nash equilibrium is the most important solution concept. In evolutionary game theory, it is replaced by an *evolutionarily stable strategy* (ESS). A strategy is an ESS if a whole population using that strategy cannot be invaded by a small group of individuals using any mutant strategy, i.e., any variant strategy. The concept of an ESS, as we shall see, has interesting connections to the idea of Nash equilibrium, although it arises in quite a different way, and it is also linked to a dynamical theory of evolution.
- *Interactions* - In classical game theory, players play a single game against an opponent, however, in evolutionary game theory, the players are assumed to play many games against randomly chosen (by nature) opponents. That is, the individuals engage in many random pairwise interactions using strategies coded by their genome.

Applications of evolutionary game theory are not confined to evolutionary biology. It has many interesting applications in economics and social sciences as well. The reason for this is that the basic concepts of evolutionary game theory can also be viewed in cultural terms. For example, we can say that a *society* has a strategy set, i.e., the set of alternative cultural norms, and individuals “inherit” them or choose among them. Similarly, a cultural norm is evolutionarily stable if upon being adopted by all members of a society, no small group of individuals using an alternative cultural norm can invade. Thus evolutionary game theory, in this context, changes the problem from explaining the action of individuals to understanding the spread of behavioral norms in society.

In classical game theory the outcome of a pairwise game depends on the choices of rational and consciously reasoning individuals. The solution of this type of game—the Nash equilibrium—was based on the idea that each player uses a strategy that is a best response to the strategies chosen by the others. So neither would change what they were

doing. For a symmetric Nash equilibrium (σ^*, σ^*) , we can give an alternative interpretation of the Nash equilibrium by placing the game in a population context. In a population where everyone uses σ^* , the best thing for an individual to do is to follow the crowd. So if the population starts with everyone using σ^* , then it will remain that way, and the *population* is said to be in equilibrium. A natural question to ask is then: what happens if the population is close to, but not actually at, its equilibrium configuration?

Consider an infinite population of individuals that can use some set of pure strategies S . A *population profile* is a tuple \hat{x} that gives the probability $x(s)$ with which each strategy $s \in S$ is played in the population, i.e., the frequency of each strategy within the population. Note that a population profile need not correspond to a strategy adopted by any member of the population.

In this thesis we will mostly be dealing with *pairwise contest games* which describe a situation in which a given individual plays against an opponent that has been randomly selected (by Nature) from the population and the payoff depends on what both individuals do. In contrast a *game against the field* is one in which there is no specific “opponent” for a given individual—their payoff depends on what everyone in the population is doing. Pairwise contests are much more like games from classical game theory in that we can write the payoff $\pi(\sigma; \hat{x})$ associated with a strategy σ generating a population profile \hat{x} as

$$\pi(\sigma; \hat{x}) = \sum_{s \in S} \sum_{s' \in S} p(s)x(s')\pi(s, s'), \quad (3.39)$$

for suitably defined pairwise payoffs $\pi(s, s')$.

Starting with a population in which all the individuals adopt some strategy σ^* , suppose that a (genetic) mutation occurs and a small proportion ε of individuals use some strategy other than σ^* . The new population (i.e., after the appearance of the mutants) is called the *post-entry population*, and its profile is denoted by \hat{x}_ε .

As with any dynamical system, one interesting question is: what are the end points (if there are any) of evolution? One type of evolutionary end-point is an evolutionarily stable strategy (ESS). A mixed strategy σ^* is an ESS if there exists an $\bar{\epsilon}$ such that for every $0 < \epsilon < \bar{\epsilon}$ and every $\sigma \neq \sigma^*$

$$\pi(\sigma^*; \hat{x}_\epsilon) > \pi(\sigma; \hat{x}_\epsilon), \quad (3.40)$$

where \hat{x}_ϵ is the *post-entry population profile*. In other words, a strategy σ^* is an ESS if mutants that adopt any other strategy σ leave fewer offspring in the post-entry population, provided the proportion of mutants is sufficiently small.

We wish to find the conditions under which the population is stable. Let \hat{x}^* be the profile generated by a population of individuals who all adopt strategy σ^* . A necessary condition for evolutionary stability is

$$\sigma^* \in \operatorname{argmax}_{\sigma \in \Sigma} \pi(\sigma; \hat{x}^*). \quad (3.41)$$

So, at an equilibrium, the strategy adopted by individuals must be a best response to the population profile that it generates. Furthermore, we have the population equivalent of Theorem 3.5.

Theorem 3.6. (Bishop-Canning theorem) Let σ^* be a strategy that generates a population profile \hat{x}^* . Let S^* be the support of σ^* . If the population is stable, then $\pi(s; \hat{x}^*) = \pi(\sigma^*; \hat{x}^*) \forall s \in S^*$.

In the discussion of classical game theory we considered games that could be asymmetric (for example, the simplified game of poker). The players in these games could have different roles, with different strategy sets and different payoffs. However, in *symmetric games* there is no role asymmetry between the players, i.e., they have the same strategy sets and their payoffs are symmetric. Many games of interest in biological and social sciences are symmetric and we will restrict our attention to this class of games.

In a symmetric pairwise contest game, we write $\pi(\sigma_1, \sigma_2)$ for the payoff to the row player when she plays the (pure or mixed) strategy σ_1 and the opponent plays σ_2 . Due to the symmetry, the payoff of the column player is $\pi(\sigma_2, \sigma_1)$. In addition, in a symmetric game, at equilibrium each of the two players will play the same (pure or mixed) strategy, since there is no distinction between them that would allow them to play different strategies.

Example 3.8. (Beetle game) Consider a particular species of beetle and suppose that each beetle's fitness in a given environment is determined largely by the extent to which it can find food and use the nutrients from the food efficiently. Now suppose a small fraction of mutant beetles with large body size is introduced into the population, causing the beetles with the mutation to grow a significantly large body size. Thus, we now have two distinct kinds of beetles in the population: small ones and large ones. The beetles compete with each other for food. When beetles of the same size compete, they get equal share of the food. When a large beetle competes with a small beetle, the large beetle gets the majority of the food. In all cases, large beetles experience less of a fitness benefit from a given quantity of food, since some of it is directed into maintaining their expensive metabolism. Based on the two strategies (L and S) used, the payoffs to the beetles is described by the following payoff matrix:

$$\pi = \begin{array}{c} \begin{array}{cc} S & L \end{array} \\ \begin{array}{cc} S & (5,5) & (1,8) \\ L & (8,1) & (3,3) \end{array} \end{array}. \quad (3.42)$$

Note that the game has the same payoff structure as the prisoner's dilemma game, and thus (L, L) is a strict Nash equilibrium. We find the evolutionarily stable strategy as follows. Let a fraction ε of the population use strategy L and let $1 - \varepsilon$ use strategy S . The

expected payoffs π_S and π_L to a small and large beetle in the population is given by

$$\begin{aligned}\pi_S &= 5(1 - \varepsilon) + 1\varepsilon = 5 - 4\varepsilon, \text{ and} \\ \pi_L &= 8(1 - \varepsilon) + 3\varepsilon = 8 - 5\varepsilon.\end{aligned}\tag{3.43}$$

For small ε , the expected fitness of large beetles exceeds that of small beetles. Therefore, (S, S) is not an ESS.

Now suppose that a fraction ε uses S and $1 - \varepsilon$ uses L . The expected payoffs π_S and π_L to a small and large beetle in the population is now given by

$$\begin{aligned}\pi_S &= (1 - \varepsilon) + 5\varepsilon = 1 + 4\varepsilon, \text{ and} \\ \pi_L &= 3(1 - \varepsilon) + 8\varepsilon = 3 + 5\varepsilon.\end{aligned}\tag{3.44}$$

Again, for small ε , expected fitness of large beetles exceeds that of small beetles.

Therefore, (L, L) is an ESS. □

Example 3.9. (Hawk-dove game) The hawk-dove game is one of the first evolutionary games that was invented by John Maynard Smith to study conflict between animals [SP73]. The two basic types are called “hawks” and “doves”. However, they are not intended to represent two different species of animals; instead they represent two types of behavior (i.e., actions or pure strategies) that could be exhibited by animals of the same species. The terminology arises from human behavior where those who advocate preemptive military solutions to international problems are called “hawks” while those who prefer a more diplomatic approach are called “doves”⁷.

Individuals can use one of two possible pure strategies: H (be aggressive like a hawk); D (be non-aggressive like a dove). In general, an individual can use a randomized strategy which is to be aggressive with probability p , i.e., $\sigma = (p, 1 - p)$. A population consists of

⁷Doves are in reality very aggressive.

animals that are aggressive with probability x , i.e., $\hat{x} = (x, 1 - x)$, which can arise because in a monomorphic population, everyone uses the strategy $\sigma = (x, 1 - x)$.

At various times, individuals in a population may come into conflict over a resource of value v , such as food, or a breeding site. The outcome of a conflict depends on the types of the two individuals that meet. If a hawk and a dove meet, then the hawk gains the resource without a fight. If two doves meet then they share the resource. If two hawks meet then there is a fight and each individual has an equal chance of winning. The winner gets the resource and the loser pays a cost (e.g., injury) of c . The payoffs for a focal individual are then

$$\pi(\sigma; \hat{x}) = px \frac{v-c}{2} + p(1-x)v + (1-p)(1-x) \frac{v}{2}. \quad (3.45)$$

We assume $v < c$ (this is then “the hawk-dove game”). It is easy to see that there is no pure-strategy ESS. In a population of doves, $x = 0$, and

$$\begin{aligned} \pi(\sigma; \hat{x}) &= pv + p(1-x)v + (1-p) \frac{v}{2} \\ &= (1+p) \frac{v}{2}. \end{aligned} \quad (3.46)$$

so the best response to this population is to play hawk (i.e., individuals using the strategy $\sigma = (1, 0)$ will do best in this population). As a consequence, the proportion of more aggressive individuals will increase (i.e., x increases). In a population of hawks, $x = 1$, and

$$\pi(\sigma; \hat{x}) = p \frac{v-c}{2}. \quad (3.47)$$

so the best response to this population is to play dove (i.e., $p = 0$ —remember $v < c$).

Is there a mixed-strategy ESS, $\sigma^* = (p^*, 1 - p^*)$? For σ^* to be an ESS, it must be a best response to the population $\hat{x}^* = (p^*, 1 - p^*)$ that it generates. In the population \hat{x}^* , the

payoff to an arbitrary strategy $\sigma = (p, 1 - p)$ is

$$\begin{aligned}\pi(\sigma; \hat{x}^*) &= pp^* \frac{v-c}{2} + p(1-p^*)v + (1-p)(1-p^*) \frac{v}{2} \\ &= (1-p^*) \frac{v}{2} + p \frac{c}{2} \left(\frac{v}{c} - p^* \right).\end{aligned}\tag{3.48}$$

If $p^* < \frac{v}{c}$ then the best response is $p = 1$ (i.e., $p \neq p^*$). If $p^* > \frac{v}{c}$, then the best response is $p = 0$ (i.e., $p \neq p^*$). If $p^* = \frac{v}{c}$, then any choice of p (including p^*) gives the same payoff (i.e., $\pi(\sigma^*; \hat{x}^*) = \pi(\sigma; \hat{x}^*)$) and is a best response to \hat{x}^* . So we have

$$\sigma^* = \left(\frac{v}{c}, 1 - \frac{v}{c} \right).\tag{3.49}$$

To confirm that σ^* is an ESS we must show that, for every $\sigma \neq \sigma^*$, $\pi(\sigma^*; \hat{x}_\varepsilon) > \pi(\sigma; \hat{x}_\varepsilon)$, where the post-entry population profile is

$$\begin{aligned}\hat{x}_\varepsilon &= ((1-\varepsilon)p^* + \varepsilon p; ((1-\varepsilon)(1-p^*) + \varepsilon(1-p))) \\ &= (p^* + \varepsilon(p-p^*); 1-p^* + \varepsilon(p^*-p)).\end{aligned}\tag{3.50}$$

After a few lines of algebra and using the fact that $p^* = \frac{v}{c}$, we find

$$\begin{aligned}\pi(\sigma^*; \hat{x}_\varepsilon) - \pi(\sigma; \hat{x}_\varepsilon) &= \frac{\varepsilon c}{2} (p^* - p)^2 \\ &> 0 \quad \forall p \neq p^* \quad (\text{i.e., } \forall \sigma \neq \sigma^*),\end{aligned}\tag{3.51}$$

which proves that σ^* is an ESS. □

Theorem 3.7. If σ^* is an ESS in a pairwise contest, then $\forall \sigma \neq \sigma^*$ either

1. $\pi(\sigma^*, \sigma^*) > \pi(\sigma, \sigma^*)$, or
2. $\pi(\sigma^*, \sigma^*) = \pi(\sigma, \sigma^*)$ and $\pi(\sigma^*, \sigma) > \pi(\sigma, \sigma)$.

Conversely, if either (1) or (2) holds for each $\sigma \neq \sigma^*$ in a two-player game, then σ^* is an ESS in the corresponding population game.

The Nash equilibrium condition is $\pi(\sigma^*, \sigma^*) \geq \pi(\sigma, \sigma^*) \forall \sigma \neq \sigma^*$ so the condition $\pi(\sigma^*, \sigma) > \pi(\sigma, \sigma)$ in (2) is a supplementary requirement that eliminates some Nash equilibria from consideration. In other words, there may be a Nash equilibrium in the two-player game but no corresponding ESS in the population game. The supplementary condition is particularly relevant in the case of mixed-strategy Nash equilibria.

Theorem 3.7 gives us an alternative procedure for finding an ESS in a pairwise contest population game:

1. write down the associated two-player game;
2. find the symmetric Nash equilibria of the game;
3. test the Nash equilibria using conditions (1) and (2) from Theorem 3.7.

Any Nash equilibrium strategy σ^* that passes these tests is an ESS, leading to a population profile $\hat{x}^* = \hat{\sigma}^*$.

Example 3.10. Reconsider the hawk-dove game from Example 3.9. The associated two-player game has the following payoff matrix:

$$\pi = \begin{array}{cc} & \begin{array}{cc} H & D \end{array} \\ \begin{array}{c} H \\ D \end{array} & \left(\begin{array}{cc} (\frac{v-c}{2}, \frac{v-c}{2}) & (v, 0) \\ (0, v) & (\frac{v}{2}, \frac{v}{2}) \end{array} \right). \end{array} \quad (3.52)$$

It is easy to see that (for $v < c$) there are no symmetric pure-strategy Nash equilibria. To find a mixed-strategy Nash equilibrium, we use Theorem 3.5.

$$\begin{aligned} \pi_1(H, \sigma^*) &= \pi_1(D, \sigma^*) \\ \Leftrightarrow p^* \frac{v-c}{2} + (1-p^*)v &= (1-p^*) \frac{v}{c} \\ \Leftrightarrow p^* &= \frac{v}{c}. \end{aligned} \quad (3.53)$$

By the symmetry of the problem, we can deduce immediately that Player 2 also plays H with probability $p^* = \frac{1}{2}$. To show that $\sigma^* = (p^*, 1 - p^*)$ is an ESS, we must show that either condition (1) or condition (2) of Theorem 3.7 holds for every $\sigma \neq \sigma^*$. Because σ^* is a mixed strategy, Theorem 3.5 also tells us that $\pi(\sigma^*, \sigma^*) = \pi(\sigma, \sigma^*)$. So condition (1) does *not* hold, and the ESS condition becomes $\pi(\sigma^*, \sigma) > \pi(\sigma, \sigma)$. After a few lines of algebra, we find

$$\begin{aligned}\pi(\sigma^*, \sigma) - \pi(\sigma, \sigma) &= \frac{c}{2}(p^* - p)^2 \\ &> 0 \quad \forall p \neq p^*.\end{aligned}\tag{3.54}$$

which proves that σ^* is an ESS. □

A general symmetric two-player game with pure-strategy set $S = \{A, B\}$ has the following payoff matrix:

$$\pi = \begin{array}{cc} & \begin{array}{cc} A & B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{pmatrix} (a, a) & (b, c) \\ (c, b) & (d, d) \end{pmatrix} \end{array}.\tag{3.55}$$

By applying affine transformations, we can turn this into the equivalent game:

$$\pi = \begin{array}{cc} & \begin{array}{cc} A & B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{pmatrix} (a - c, a - c) & (0, 0) \\ (0, 0) & (d - b, d - b) \end{pmatrix} \end{array}.\tag{3.56}$$

It is easy to see that the ESS conditions given in Theorem 3.7 are unaffected by this transformation. Since we are considering generic games, we have $a \neq c$ and $b \neq d$. There are three possibilities to consider.

1. If $a - c > 0$, then $\pi(A, A) > \pi(B, A)$ and hence $\sigma = (1, 0)$ is an ESS, by condition (1) in Theorem 3.7.

2. If $d - b > 0$, then $\pi(B, B) > \pi(A, B)$ and hence $\sigma = (0, 1)$ is an ESS, by condition (1) in Theorem 3.7.

3. If $a - c < 0$ and $d - b < 0$, then there is a symmetric mixed strategy Nash equilibrium (σ^*, σ^*) with $\sigma^* = (p^*, 1 - p^*)$ and

$$p^* = \frac{d - b}{a - c + d - b}. \quad (3.57)$$

At this equilibrium, $\pi(\sigma^*, \sigma^*) = \pi(\sigma, \sigma^*)$ for any strategy σ , so we have to consider the inequality in condition (2) of Theorem 3.7. Now

$$\begin{aligned} \pi(\sigma^*, \sigma) - \pi(\sigma, \sigma) &= p(p^* - p)(a - c) + (1 - p)(p - p^*)(d - b) \\ &= (p^* - p)[p(a - c + d - b) - (d - b)] \\ &= -(a - c + d - b)(p^* - p)^2 \\ &> 0. \end{aligned} \quad (3.58)$$

So σ^* is an ESS.

Thus, all symmetric, two-strategy, pairwise games have an ESS. It is worth noting that condition (1) or (2) yields the *prisoner's dilemma class* of games; conditions (1) and (2) yield the *coordination class* of games; and finally, condition (3) yields the *chicken class* of games. The three games are thus considered to be canonical two-strategy games. Pairwise games with more than two strategies need not have an ESS, as following example illustrates.

Example 3.11. Reconsider the rock-paper-scissors game from Example 3.2. It has the following payoff matrix:

$$\pi = \begin{array}{c} \begin{array}{ccc} & R & P & S \\ R & (0, 0) & (-1, 1) & (1, -1) \\ P & (1, -1) & (0, 0) & (-1, 1) \\ S & (-1, 1) & (1, -1) & (0, 0) \end{array} \end{array}. \quad (3.59)$$

We want to show that the game has no ESS. First, it is clear that no pure strategy is an ESS, since $\pi(P, R) > \pi(R, R)$, $\pi(R, S) > \pi(S, S)$, and $\pi(S, P) > \pi(P, P)$. Let us consider a mixed strategy $\sigma = pR + qS + (1 - p - q)P$. If σ is an ESS then, it follows from Theorem 3.6 that $\pi(R, \sigma) = \pi(S, \sigma) = \pi(P, \sigma)$.

$$\begin{aligned} \pi(R, \sigma) = \pi(S, \sigma) &\implies p\pi(R, R) + q\pi(R, S) + (1 - p - q)\pi(R, P) \\ &= p\pi(S, R) + q\pi(S, S) + (1 - p - q)\pi(S, P) \\ &\implies -1 + p + 2q = 1 - 2p - q. \end{aligned} \tag{3.60}$$

Similarly,

$$\pi(S, \sigma) = \pi(P, \sigma) \implies 1 - 2p - q = p - q. \tag{3.61}$$

Solving Equations 3.60 and 3.61 we get $p = q = \frac{1}{3}$. So if σ is an ESS it must be $\sigma = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. Consider $\pi(\sigma, \sigma)$. It is equal to $\pi(R, \sigma)$ so we must look at the condition $\pi(\sigma, R) > \pi(R, R)$. However, $\pi(\sigma, R) = -\frac{1}{3} + \frac{1}{3} = 0$ and $\pi(R, R) = 0$. So $\pi(\sigma, R) \not> \pi(R, R)$ and σ is not an ESS. \square

Let us now look at some of the limitations of evolutionary game theory, which can (and have been) addressed in various ways.

1. Evolutionary game theory is a model of *phenotypic evolution*, i.e., the phenotype is the central concept as this is identified with strategy. Thus, we do not consider genes directly. However, since the phenotype is determined by the genotype, the change in the frequency of the different phenotypes implies a corresponding change in genotype frequencies. So we lose some information at the genotypic level. However, we also gain a lot. We can often make a reasonable conjecture for the strategies that an organism can use, and test this empirically, but we often have very little idea as to what genes are responsible for what.

2. Evolutionary game theory does not take into account the subtle features of sexual reproduction—it assumes that individuals inherit the strategy of their parents. This is exactly true in the case of asexual reproduction, but evolutionary game theory is used to model interactions among many sexual organisms.
3. Evolutionary game theory assumes the interactions to be random. One can consider more complex interaction schemes, such as assortative interactions or spatially structured interactions, which we do in Chapters 4, 5, and 6.
4. Evolutionary game theory also often assumes that the individuals interact in pairs, but we can generalize to groups of interacting individuals.

3.4 Replicator Dynamics

Although the concept of an evolutionarily stable strategy (ESS) implicitly assumes the existence of some kind of evolutionary dynamics, it provides an incomplete description. First, an ESS may not exist—the rock-paper-scissors game, for example, has no ESS—in which case the analysis tells us nothing about the evolution of the system described by the game. Second, the definition of an ESS deals only with monomorphic populations in which every individual uses the same strategy. But, if the ESS is a mixed strategy, then all the strategies in the support of the ESS obtain the same payoff as the ESS itself. It is unclear if a polymorphic population with the same population profile as that generated by the ESS can also be stable. Third, a game can have multiple ESSs—coordination games, for example, have two ESSs. It is unclear which ESS will be preferred. In order to address these issues we look at a specific type of evolutionary dynamics, called replicator dynamics.

A *replicator* is any entity that can make (approximately) accurate copies of itself. Example of replicators are: genomes in asexual organisms, genes in sexual organisms, a computer virus, a cultural norm, or a strategy in a game.

We define the fitness of a replicator to be the expected number of copies it makes of itself, i.e., its expected number of offspring. Consider a large population of n types of replicators $\mathcal{R}_1, \dots, \mathcal{R}_n$. We denote the fitness of replicator \mathcal{R}_i by f_i . Let x_i be the frequency of \mathcal{R}_i . The state of the population is specified by the population profile $\hat{x} = (x_1, \dots, x_n)$, which specifies the frequencies of the different types of replicators. Note that $x_i \geq 0$ and $\sum_{i=1}^n x_i = 1$. The population profile \hat{x} is a point in S_n , the n -simplex defined as $S_n = \{(x_1, \dots, x_n) | x_i \geq 0, \sum_{i=1}^n x_i = 1\}$. Since only $n - 1$ of the x_i 's are independent, S_n has dimension $n - 1$.

Example 3.12. The 2-simplex $S_2 = \{(x_1, x_2) | x_1, x_2 \geq 0, x_1 + x_2 = 1\}$ and the 3-simplex $S_3 = \{(x_1, x_2, x_3) | x_1, x_2, x_3 \geq 0, x_1 + x_2 + x_3 = 1\}$ are illustrated in Figure 4. □




Figure 4: 2-simplex S_2 and 3-simplex S_3 .

It is assumed that the replicators are programmed to use only pure strategies from a finite set $S = \{s_1, \dots, s_n\}$. Let n_i be the number of individuals using s_i . The total

population size is

$$N = \sum_{i=1}^n n_i, \quad (3.62)$$

and the proportion of individuals using s_i is

$$x_i = \frac{n_i}{N}. \quad (3.63)$$

Let β and δ be the background per capita birth and death rates in the population. The rate of change of the number of individuals using s_i is⁸

$$\dot{n}_i = (\beta - \delta + f_i)n_i, \quad (3.64)$$

and the rate of change of the total population size is given by

$$\begin{aligned} \dot{N} &= \sum_{i=1}^n \dot{n}_i \\ &= (\beta - \delta) \sum_{i=1}^n n_i + \sum_{i=1}^n n_i f_i \\ &= (\beta - \delta)N + N \sum_{i=1}^n x_i f_i \\ &= (\beta - \delta + \bar{f})N, \end{aligned} \quad (3.65)$$

where we have defined the *average fitness* of the population by

$$\bar{f} = \sum_{i=1}^n x_i f_i. \quad (3.66)$$

Thus the population grows or declines exponentially. This may not be very realistic, but we can improve the description by letting β and δ depend on N . So long as the fitness increments f_i depend only on the proportion x_i , and not on the actual numbers n_i , the game dynamics will be unchanged.

⁸We use a dot over a variable to denote its time derivative. For example, $\dot{x} = \frac{dx}{dt}$.

We are more interested in how the proportions of each type of replicators change over time. Now

$$\dot{n}_i = N\dot{x}_i + x_i\dot{N}. \quad (3.67)$$

So

$$\begin{aligned} N\dot{x}_i &= \dot{n}_i - x_i\dot{N} \\ &= (\beta - \delta + f_i)x_iN - x_i(\beta - \delta + \bar{f})N. \end{aligned} \quad (3.68)$$

Therefore,

$$\dot{x}_i = x_i(f_i - \bar{f}). \quad (3.69)$$

Equation 3.69 is called the *replicator equation*, originally introduced by Taylor and Jonker [TJ78], and followed by Christopher Zeeman, Peter Schuster, Josef Hofbauer, and Karl Sigmund. The equation has been enormously successful in modeling biological evolution. It says that the proportion of individuals using strategy s_i increases (decreases) if its fitness is bigger (smaller) than the average fitness in the population. Since a strategy in an evolutionary game is a replicator we can use the replicator equation to describe evolutionary game dynamics. Fixed points of the equation describe populations that are no longer evolving.

Assume that we have n types of replicators $\mathcal{R}_1, \dots, \mathcal{R}_n$ with frequencies x_1, \dots, x_n (at any given time) and fixed fitnesses f_1, \dots, f_n , respectively. Let \bar{f} be the average fitness of the population. We are assuming *frequency independent selection*, i.e., the fitnesses of the different replicators is fixed and does not depend on the frequencies of the replicators. The dynamics is described by the replicator equation (Equation 3.69). How does evolution work here?

Let us first consider the case $n = 2$, i.e., just two replicators \mathcal{R}_1 and \mathcal{R}_2 , with fitnesses f_1 and f_2 and frequencies x_1 and x_2 , respectively. Since $x_1 + x_2 = 1$, we can describe the state of the system by one of the frequencies, say $x = x_1$, and then $x_2 = 1 - x$. Also, $\bar{f} = xf_1 + (1 - x)f_2$. So, we have

$$\begin{aligned}
\dot{x} &= x(f_1 - \bar{f}) \\
&= x[f_1 - (xf_1 + (1 - x)f_2)] \\
&= x[(1 - x)f_1 - (1 - x)f_2] \\
&= x(1 - x)[f_1 - f_2] \\
&= rx(1 - x), \quad \text{where } r = f_1 - f_2.
\end{aligned} \tag{3.70}$$

This is just the logistic equation (see A.2) with $r = f_1 - f_2$. If $f_1 > f_2$ then $r > 0$ and the equilibrium point $x^* = 1$ is stable (and $x^* = 0$ is unstable), so the population evolves to all \mathcal{R}_1 . On the other hand, if $f_1 < f_2$ then $r < 0$ and the stability of the equilibrium points is reversed, and the system evolves to \mathcal{R}_2 . Thus, in this case, the system evolves to a state consisting of the replicator type with the highest fitness.

This is true in general. Assume we have n replicator types and one type, say i , has higher fitness than any of the others, i.e., $f_i > f_j, \forall j \neq i$. We can show that \mathcal{R}_i goes to fixation as follows:

$$\begin{aligned}
\frac{d}{dt} \left(\frac{x_j}{x_i} \right) &= \frac{x_i \frac{dx_j}{dt} - x_j \frac{dx_i}{dt}}{x_i^2} \\
&= \frac{x_i \dot{x}_j - x_j \dot{x}_i}{x_i^2} \\
&= \frac{x_i x_j [f_j - \bar{f}] - x_j x_i [f_i - \bar{f}]}{x_i^2} \\
&= \frac{x_i x_j [f_j - f_i]}{x_i^2}.
\end{aligned} \tag{3.71}$$

Since $f_i > f_j \forall j \neq i$, $\frac{d}{dt} \left(\frac{x_j}{x_i} \right) < 0$, which implies $x_j(t) \rightarrow 0$ as $t \rightarrow \infty$. Hence, $x_i(t) \rightarrow 1$ as $t \rightarrow \infty$ and \mathcal{R}_i goes to fixation.

Let us next consider *frequency dependent selection*, i.e., the fitnesses of the replicators depend on their frequencies at any given time. Consider a pairwise game with n pure strategies s_1, \dots, s_n , and a large population of individuals, where each individual uses one of these n strategies. Let π denote the payoff matrix for the game. The state of the population is $\hat{x} = \{x_1, \dots, x_n\}$, where x_i is the frequency of individuals using strategy s_i . The fitness f_i of strategy s_i (i.e., the expected number of offspring of any individual using strategy s_i) is

$$f_i = f_0 + \sum_{j=1}^n x_j \pi(s_i, s_j), \quad (3.72)$$

where f_0 is the fixed background fitness, same for all individuals in the population, and the average fitness of the population is

$$\bar{f} = \sum_{i=1}^n x_i f_i. \quad (3.73)$$

Now the dynamics of the system is given by the replicator equation

$$\dot{x}_i = x_i(f_i - \bar{f}). \quad (3.74)$$

Two-Strategy Pairwise Contests

Consider a population with profile \hat{x} playing a pairwise contest game with two pure strategies s_1 and s_2 . Let π be the payoff matrix for the game, and let $\pi(s_1, \hat{x})$ and $\pi(s_2, \hat{x})$ denote the fitnesses of the two strategies. Suppose $x \equiv x_1$, the frequency of individuals playing s_1 . Then the frequency x_2 of individuals playing s_2 is $x_2 = 1 - x$ and $\dot{x}_2 = -\dot{x}_1$. So we only need to consider a single differential equation

$$\dot{x} = x(\pi(s_1, \hat{x}) - \bar{\pi}(\hat{x})), \quad (3.75)$$

where $\bar{\pi}(\hat{x})$ is the average fitness of the population given by

$$\bar{\pi}(\hat{x}) = x\pi(s_1, \hat{x}) + (1 - x)\pi(s_2, \hat{x}) \quad (3.76)$$

Substituting Equation 3.76 into Equation 3.75 and simplifying the result, we get the replicator equation for strategy s_1 as

$$\dot{x} = x(1-x)(\pi(s_1, \hat{x}) - \pi(s_2, \hat{x})). \quad (3.77)$$

Example 3.13. Let us reconsider a general symmetric, two-strategy, pairwise game given by the following payoff matrix:

$$\pi = \begin{array}{c} \\ A & B \\ \left(\begin{array}{cc} (a, a) & (b, c) \\ (c, b) & (d, d) \end{array} \right). \end{array} \quad (3.78)$$

Let x_A, x_B be the frequencies of A, B , respectively, with $x_A + x_B = 1$. If we define $x \equiv x_A$ then $x_B = 1 - x$. In this case the replicator equation is:

$$\begin{aligned} \dot{x} &= x(1-x)[\pi(A, \hat{x}) - \pi(B, \hat{x})] \\ &= x(1-x)[(a-b-c+d)x + b-d]. \end{aligned} \quad (3.79)$$

This equation has equilibria at $x^* = 0, 1$, and a possible internal equilibrium $x^* \in (0, 1)$.

The allowed dynamics are as follows:

1. $a > c, d < b$: In this case A is globally stable.
2. $a < c, d > b$: In this case B is globally stable.
3. $a > c, d > b$: In this case A and B are bistable.
4. $a < c, d < b$: In this case the internal equilibrium is stable.

Figure 5 shows the phase line plots for each of the four cases. As we have seen before, (1) and (2) yield the *prisoner's dilemma class* of games; (3) yields the *coordination class* of games; and (4) yields the *chicken class* of games. \square




Figure 5: Phase line plots for a general symmetric, two-strategy, pairwise game.

Example 3.14. Consider a pairwise contest prisoner's dilemma game with payoff matrix

$$\pi = \begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} (3,3) & (0,5) \\ (5,0) & (1,1) \end{pmatrix}. \end{array} \quad (3.80)$$

Let x be the population of individuals using C , then

$$\pi(C; \hat{x}) = 3x + 0(1 - x) = 3x, \quad (3.81)$$

and

$$\pi(D; \hat{x}) = 5x + 1(1 - x) = 1 + 4x. \quad (3.82)$$

The rate of change of the proportion of individuals using C is

$$\begin{aligned} \dot{x} &= x(1 - x)(\pi(C; \hat{x}) - \pi(D; \hat{x})) \\ &= x(1 - x)(3x - (1 + 4x)) \\ &= -x(1 - x)(1 + x). \end{aligned} \quad (3.83)$$

The fixed points for the dynamical system are $x^* = 0$ and $x^* = 1$. We know that the unique Nash equilibrium for the prisoner's dilemma game is for everyone to defect (play D). This means $x^* = 0$ corresponds to a Nash equilibrium but $x^* = 1$ does not. We also see that $\dot{x} < 0$ for $x \in (0, 1)$, which implies that any population that is not at a fixed point of the dynamics will evolve towards the fixed point that corresponds to the Nash equilibrium and away from the other one. \square

Theorem 3.8. Let $S = \{s_1, s_2\}$ and let $\sigma^* = (p^*, 1 - p^*)$ be the strategy that uses s_1 with probability p^* and s_2 with probability $1 - p^*$. If (σ^*, σ^*) is a symmetric Nash equilibrium, then the population $\hat{x}^* = (x^*, 1 - x^*)$ with $x^* = p^*$ is a fixed point of the replicator dynamics $\dot{x} = x(1 - x)(\pi(s_1; \hat{x}) - \pi(s_2; \hat{x}))$.

Thus the Nash equilibria in two-player games and fixed points in the replicator dynamics are related. How about ESSs? Do all ESSs have an evolutionary end point and do all evolutionary end points have a corresponding ESS?

Example 3.15. Consider a pairwise contest with pure strategies A and B and with the payoff matrix

$$\pi = \begin{array}{cc} & \begin{array}{cc} A & B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{pmatrix} (3, 3) & (0, 0) \\ (0, 0) & (1, 1) \end{pmatrix} \end{array}. \quad (3.84)$$

The ESSs for the game are for everyone to play A or for everyone to play B . The mixed strategy $\sigma = (\frac{1}{4}, \frac{3}{4})$ is a Nash equilibrium but not an ESS. Let x be the fraction of individuals using A . Then the replicator dynamics is given by

$$\dot{x} = -x(1 - x)(1 - 4x) \quad (3.85)$$

with fixed points $x^* = 0$, $x^* = 1$, and $x^* = \frac{1}{4}$.

First, consider a population near $x^* = 0$. Let $x = x^* + \varepsilon$ where we must have $\varepsilon > 0$ so that $x > 0$. Then $\dot{x} = \dot{\varepsilon}$ since x^* is a constant. Thus we have

$$\dot{\varepsilon} = -\varepsilon(1 - \varepsilon)(1 - 4\varepsilon). \quad (3.86)$$

Since $\varepsilon \ll 1$, we can ignore terms proportional to ε^n where $n > 1$. Thus

$$\dot{\varepsilon} \approx -\varepsilon \quad (3.87)$$

which has solution

$$\varepsilon(t) = \varepsilon_0 e^{-t}. \quad (3.88)$$

This tells us that the dynamics reduces small deviations from the state $\hat{x} = (0, 1)$, i.e., $\varepsilon \rightarrow 0$ as $t \rightarrow \infty$. In other words, the fixed point $x^* = 0$ is asymptotically stable. We can similarly show that $x^* = 1$ is also asymptotically stable.

Now consider a population near $x^* = \frac{1}{4}$. Let $x = x^* + \varepsilon = \frac{1}{4} + \varepsilon$. There is no sign restriction on ε . Then we have

$$\dot{\varepsilon} \approx \frac{1}{16}\varepsilon \quad (3.89)$$

with solution

$$\varepsilon(t) = \varepsilon_0 e^{\frac{t}{16}}. \quad (3.90)$$

So $x^* = \frac{1}{4}$ is *not* asymptotically stable. So in this case we find that a strategy is an ESS if and only if the corresponding fixed point in the replicator dynamics is asymptotically stable. □

Theorem 3.9. For any two-strategy pairwise contest, a strategy is an ESS if and only if the corresponding fixed point in the replicator dynamics is asymptotically stable.

Let F be the set of fixed points and let A be the set of asymptotically stable fixed points in the replicator dynamics. Let N be the set of symmetric Nash equilibria and let E be the set of ESSs in the symmetric game corresponding to the replicator dynamics. Then, for any two-strategy pairwise contest game, the following relationships hold for a strategy σ^* and the corresponding population state \hat{x}^* :

1. $\sigma^* \Leftrightarrow \hat{x}^* \in A$;
2. $\hat{x}^* \in A \Rightarrow \sigma^* \in N$;
3. $\sigma^* \in N \Rightarrow \hat{x}^* \in F$.

More concisely, we have

$$E = A \subseteq N \subseteq F. \quad (3.91)$$

Games with More Than Two Strategies

If we increase the number of pure strategies to n , then we have n equations of the form

$$\dot{x}_i = f_i(x_1, \dots, x_n) \quad i = 1, \dots, n. \quad (3.92)$$

Using the constraint $\sum_{i=1}^n x_i = 1$, we can introduce a reduced state vector (x_1, \dots, x_{n-1}) and reduce the number of equations to $n - 1$.

$$\dot{x}_i = f_i(x_1, \dots, x_{n-1}) \quad i = 1, \dots, n - 1 \quad (3.93)$$

More compactly, in vector format, we have

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}). \quad (3.94)$$

An *invariant manifold* is a connected subset $M \subset S_n$ such that if $x(0) \in S_n$, then $x(t) \in S_n$ for all $t > 0$. To obtain a qualitative picture of the solutions of the dynamical

system, we obtain a linear approximation to the dynamical system, and consider the behavior of the solutions on (or close to) the invariant manifolds.

$$\dot{x}_i = \sum_{j=1}^{n-1} (x_j - x_j^*) \frac{\partial f_i}{\partial x_j}(\hat{x}^*) \quad (3.95)$$

If we define $\eta_i = x_i - x_j^*$, then

$$\dot{\eta} = \sum_{j=1}^{n-1} \eta \frac{\partial f_i}{\partial x_j}(\hat{x}^*), \quad (3.96)$$

which is a linear system $\dot{\eta} = A\eta$ with a fixed point at the origin. The matrix has constant components

$$A_{ij} = \frac{\partial f_i}{\partial x_j}(\hat{x}^*) \quad (3.97)$$

and its eigenvalues determine the behavior of the linearized system at the fixed point.

From Theorem A.2, the behavior of the full, non-linear system is the same, provided the fixed point is hyperbolic. Combining this information with the behavior of solutions on the other invariant manifolds is usually sufficient to determine a complete qualitative picture of the solution to the dynamical system. We illustrate this with an example.

Example 3.16. Consider the pairwise contest game with the payoff matrix

$$\pi = \begin{array}{c} \begin{array}{ccc} & A & B & C \\ A & (0,0) & (3,3) & (1,1) \\ B & (3,3) & (0,0) & (1,1) \\ C & (1,1) & (1,1) & (1,1) \end{array} \end{array}. \quad (3.98)$$

The replicator dynamics for the game is

$$\begin{aligned} \dot{x}_1 &= x_1(3x_2 + x_3 - \bar{\pi}(\hat{x})) \\ \dot{x}_2 &= x_2(3x_1 + x_3 - \bar{\pi}(\hat{x})) \\ \dot{x}_3 &= x_3(1 - \bar{\pi}(\hat{x})) \end{aligned} \quad (3.99)$$

with $\bar{\pi}(\hat{x}) = 6x_1x_2 + x_1x_3 + x_2x_3 + x_3$. Writing $x_1 = x, x_2 = y$ and $x_3 = 1 - x - y$, this system can be reduced to the two-variable dynamical system

$$\begin{aligned}\dot{x} &= x(1 - x + 2y - \bar{\pi}(x, y)) \\ \dot{y} &= y(1 + 2x - y - \bar{\pi}(x, y))\end{aligned}\tag{3.100}$$

with $\bar{\pi}(x, y) = 1 + 4xy - x^2 - y^2$.

The invariant manifolds are the fixed points $((0, 0), (0, 1), (1, 0), (\frac{1}{2}, \frac{1}{2}))$, the boundary lines $x = 0, y = 0$, and boundary line $x + y - 1 = 0$ since $\frac{d}{dt}(x + y) = \dot{x} + \dot{y} = (x + y - 1)(1 - \bar{\pi}(x, y)) = 0$. The line $x = y$ is also an invariant because $\dot{x} = \dot{y}$ on that line.

Consider the fixed point $(x^*, y^*) = (\frac{1}{2}, \frac{1}{2})$. Close to the point we have the linear approximation

$$\begin{pmatrix} \dot{\xi} \\ \dot{\eta} \end{pmatrix} = \begin{pmatrix} -1 & \frac{1}{2} \\ \frac{1}{2} & -1 \end{pmatrix} \begin{pmatrix} \xi \\ \eta \end{pmatrix}.\tag{3.101}$$

The eigenvalues of A are $\lambda_1 = -\frac{1}{2}$ and $\lambda_2 = -\frac{3}{2}$. Since both eigenvalues are real and negative, the fixed point is a stable node. The eigenvector corresponding to $\lambda_2 = -\frac{3}{2}$ is $\xi = -\eta$, which lies along the boundary line $x + y = 1$. The eigenvector corresponding to $\lambda = -\frac{1}{2}$ is $\xi = \eta$, which lies along $x = y$.

The fixed points $(x^*, y^*) = (1, 0)$ and $(x^*, y^*) = (0, 1)$ both have eigenvalues $\lambda_1 = 3$ and $\lambda_2 = 1$, so both are unstable nodes.

Close to the point $(x^*, y^*) = (0, 0)$ the linear approximation is

$$\begin{pmatrix} \dot{\xi} \\ \dot{\eta} \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \xi \\ \eta \end{pmatrix}.\tag{3.102}$$

which is not hyperbolic ($\lambda_1 = \lambda_2 = 0$). So the linearization tells us nothing about the stability properties of this fixed point.

On the line $y = 0$, $\dot{x} = x^2(x - 1)$, so $\dot{x} < 0$ for $0 < x < 1$. Similarly, on the line $x = 0$, we have $\dot{y} < 0$ for $0 < y < 1$. On the line $x = y$ we have $\dot{x} = x^2(1 - 2x)$, so x and y are both increasing for $0 < x, y < \frac{1}{2}$. On the line $x + y - 1 = 0$ we have

$$\begin{aligned}\dot{x} &= x(3 - 3x - \bar{\pi}(x, 1 - x)) \\ &= x(3 - 9x + 6x^2).\end{aligned}\tag{3.103}$$

Hence x is increasing (y is decreasing) for $0 < x < \frac{1}{2}$ and x is decreasing (y is increasing) for $\frac{1}{2} < x < 1$. Combining all this we have the qualitative picture of the dynamics shown in Figure 6. □




Figure 6: Behavior of the full system 3.100.

Theorem 3.10. If (σ^*, σ^*) is a symmetric Nash equilibrium, then the population state $\hat{x}^* = \hat{\sigma}^*$ is a fixed point of the replicator dynamics.

Theorem 3.11. If \hat{x}^* is an asymptotically stable fixed point of the replicator dynamics, then the symmetric strategy pair (σ^*, σ^*) with $\sigma^* = x^*$ is a Nash equilibrium.

Theorem 3.12. Every ESS corresponds to an asymptotically stable fixed point in the replicator dynamics. That is, if σ^* is an ESS, then the population with $\hat{x}^* = \hat{\sigma}^*$ is asymptotically stable.

For any pairwise contest game with more than two strategies, the following relationships hold for a strategy σ^* and the corresponding population state \hat{x}^* .

1. $\sigma^* \in E \implies \hat{x}^* \in A$;
2. $\hat{x}^* \in A \implies \sigma^* \in N$;
3. $\sigma^* \in N \implies \hat{x}^* \in F$.

Identifying a strategy with its corresponding population state, we can write these relations more concisely as

$$E \subseteq A \subseteq N \subseteq F. \quad (3.104)$$

Example 3.17. Reconsider the rock-paper-scissors game. Let x_1, x_2, x_3 be the proportions of R, S , and P players. Let the payoff matrix for the game be

$$\pi = \begin{array}{c} \\ R \\ P \\ S \end{array} \begin{array}{ccc} R & P & S \\ \left(\begin{array}{ccc} (0,0) & (-1,1) & (1,-1) \\ (1,-1) & (0,0) & (-1,1) \\ (-1,1) & (1,-1) & (0,0) \end{array} \right) \end{array}. \quad (3.105)$$

The replicator dynamics of the game is

$$\begin{aligned} \dot{x}_1 &= x_1(x_2 - x_3) \\ \dot{x}_2 &= x_2(x_3 - x_1) \\ \dot{x}_3 &= x_3(x_1 - x_2) \end{aligned} \quad (3.106)$$

with fixed points $((1, 0, 0), (0, 1, 0), (0, 0, 1), (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}))$.

It is easy to see by considering the boundaries that the first three points are not stable. For example, consider the invariant line $x_1 = 0$ where, for $0 < x_1, x_3 < 1$, we have $\dot{x}_2 > 0$ and $\dot{x}_3 < 0$. The results from the three invariant lines together imply that there is some kind of oscillatory behavior about the polymorphic fixed point $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$. If the fixed point is asymptotically stable, then trajectories will spiral into it, and if unstable, then trajectories will spiral out of it. The third possibility is that the solution trajectories form closed loops around the fixed point.

That this is in fact the case can be confirmed by observing that the time derivative of the relative entropy function along solution trajectories of the replicator dynamics is

$$\frac{d}{dt}V(x) = -\frac{1}{3}(x_3 - x_1) - \frac{1}{3}(x_1 - x_2) = 0. \quad (3.107)$$

The qualitative picture of the replicator dynamics for the rock-paper-scissors game is shown in Figure 7.

If the payoff is $\alpha < 0$ to both players when the result is a draw, then it can be shown that $\hat{x}^* = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is asymptotically stable in the replicator dynamics. On the other hand, if $\alpha > 0$, then the boundaries of the 3-simplex for the replicator dynamics are attractors. \square

Thus an evolutionary process can produce apparently rational (Nash equilibrium) behavior in a population composed of individuals not required to make consciously rational decisions. In populations where agents are assumed to have some critical faculties—such as human populations—the requirements of rationality are much less stringent than they are in classical game theory. Individuals are no longer required to be able to work through the (possibly infinite) sequence of reactions and counter-reactions to changes in behavior. They merely have to be able to evaluate the consequences of their actions, compare them to the results obtained by others two behaved differently and swap to a better (not




Figure 7: Behavior of the full system 3.106.

necessarily the best) strategy for the current situations. The population is stable when given that what everyone else is doing no individual would get a better result by adopting a different strategy.

3.5 Adaptive Dynamics

In this section we switch our attention to continuous-strategy games. These are games which consider strategies that form a mathematical continuum, i.e., a section of the real number line, such as, $[0, 1]$ or $(-\infty, \infty)$. This is in contrast with discrete or matrix games where the set of strategies is a finite, unordered list—for example, cooperate or defect. Most traits studied in evolutionary ecology are continuous (or nearly so), such as date of first flowering, time to maturity, an animal's body size, the bill dimensions of a bird, or the allocation of resources to roots, stems, and reproductive tissues on a plant. The

developments of matrix and continuous games have followed somewhat independent paths, with various researchers trying to relate the two fields [MB07].

Suppose the pure strategy (action) sets are a subset of the real line $[a, b]$. A pure strategy is then a choice $x \in [a, b]$ ⁹ and a mixed strategy is defined by giving a function $p(x)$ such that the probability that the choice lies between x and $x + dx$ is $p(x)dx$. The existence of Nash equilibria for games with continuous pure-strategy sets was proved independently by Debreu, Glicksburg, and Fan in 1952 (see [FT91, Mye97] for details).

Analyzing the evolution of a continuous trait x using the replicator equation is rather complicated, since the population profile in this case is a point in infinite-dimensional simplex. There is a much more direct way to attack the problem, and that is using the mathematical framework of adaptive dynamics.

Adaptive dynamics [GKM97, MB07] is a set of techniques developed during the 1990s for understanding the long-term consequences of small mutations in the traits expressing the phenotype. It links population dynamics to evolutionary dynamics and incorporates and the fundamental idea of frequency-dependent selection from evolutionary game theory. Adaptive dynamics is gaining ground as a versatile tool for evolutionary modeling.

Two fundamental ideas of adaptive dynamics are that the resident population can be assumed to be in a dynamical equilibrium when the new mutants appear, and that the eventual fate of such mutants can be inferred from their initial growth rate when rare in the environment consisting of the resident. The rate is known as the *invasion exponent* or *invasion fitness*. A population consisting of individuals with the same trait is called a *monomorphic population*. We assume the trait x to be a real number and the invading mutant trait will be denoted by y . The invasion fitness $f_x(y)$ ¹⁰, i.e., the expected growth

⁹Note that in the previous section on replicator dynamics, we used the letters x, y , etc., to denote frequency of strategies. Here the letters denote (continuous) trait values.

¹⁰Here again, the letter f has a different meaning than what it did in the previous section on replicator dynamics, where f_i meant the fitness of the replicator \mathcal{R}_i .

rate of an initially rare mutant in the environment set by the resident, from Equation 3.77 is given by

$$f_x(y) = \pi(y, x) - \pi(x, x), \quad (3.108)$$

where $\pi(x, y)$ is the payoff function that determines the payoff to the x -strategist against an opponent playing the strategy y . For each x , the invasion fitness can be thought of as the fitness landscape experienced by an initially rare mutant. The landscape changes with each successful invasion as in the case of evolutionary game theory, but in contrast with the classical view of evolution as an optimization process towards ever higher fitness.

The selection gradient is defined as the slope of the invasion fitness at $y = x$.

$$D(x) = \left. \frac{\partial f_x(y)}{\partial y} \right|_{y=x}. \quad (3.109)$$

The adaptive dynamics of the trait x is governed by $D(x)$. Note that in general, $\dot{x} = mD(x)$, where m depends on the population size and reflects the mutational process providing the raw material for evolutionary change. For constant population sizes, m is simply a parameter that scales time, and one can set $m = 1$ without loss of generality.

Traits or strategies x^* for which $D(x^*) = 0$ are called *singular strategies*. If there is no such solution, the trait x either always increases or decreases evolutionarily, depending on the sign of $D(x)$. If a singular trait x^* exists, it is *convergent stable*, and hence an attractor for the adaptive dynamics if

$$\left. \frac{dD}{dx} \right|_{x=x^*} < 0. \quad (3.110)$$

If the inequality is reversed, x^* is a *repeller*, i.e., the trait x evolves to even lower values if the initial trait of the population $x_0 < x^*$, and to even higher values if $x_0 > x^*$.

Convergence stability does not necessarily imply evolutionary stability. A convergent stable point x^* is either a minimum or maximum of the invasion fitness $f_x(y)$. If x^* is a

maximum, i.e., if

$$\left. \frac{\partial^2 f_{x^*}(y)}{\partial y^2} \right|_{y=x^*} < 0, \quad (3.111)$$

then x^* is *evolutionarily stable*, i.e., it cannot be invaded by any mutant. If, however, x^* is a minimum, i.e. if

$$\left. \frac{\partial^2 f_{x^*}(y)}{\partial y^2} \right|_{y=x^*} > 0, \quad (3.112)$$

then a population that is monomorphic for x^* can be invaded by mutants with trait value on either side of x^* . In this case, the population first converges evolutionarily towards x^* , but subsequently splits into two distinct and diverging phenotypic clusters. This phenomenon is called *evolutionary branching*, and the singular point is called an *evolutionary branching point* (EBP). Once the population has branched into two phenotypic clusters, the individuals in the population play a classical snowdrift game [HD04] so that either strategy could invade a population consisting of individuals having the other strategy. Furthermore, the average fitness of the individuals playing each of the two strategies is the same.

Adaptive dynamics is a formal framework for modeling evolutionary dynamics with special emphasis on the generation of diversity through branching of the evolutionary tree. It helps us better understand speciation by taking into account the fact that the fitness function itself is modified by the evolutionary process. A constant-fitness picture (each phenotype has a fixed fitness value, and the fittest type out-competes all others) is unable to explain the enormous diversity of life on earth: how could any type but the fittest survive? Although evolutionary branching is reminiscent of speciation, in the context of asexual population the species concept is not well defined. Applied to sexual population, the framework could describe evolution in allele space rather than the strategy space. Branching in the allele space can be interpreted as speciation only if the individuals in the separate do not interbreed.

We next look at an example of a continuous game and see how it is analyzed using adaptive dynamics.

Example 3.18. (War of attrition game) The war of attrition (WoA) game is used as a model of conflict. The model was originally formulated by John Maynard Smith [Smi74]. The two contestants in the game compete for a resource of value V by persisting while constantly accumulating a cost that is a function of the time t that the contest lasts. If the two individuals persist for the same time t , then the two receive half of the resource V minus the cost of persistence for time t . If the two individuals persist for times t_1 and t_2 such that $t_1 > t_2$, then individual persisting longer gets all of the resource V and only incurs the cost associated with the shorter persistence time t_2 while the other individual receives none of the resource but incurs the same cost.

The strategies of the two individuals, x and y , respectively, are the investments they make to obtain the resource (e.g., their persistence times). The payoff in such an interaction to the x strategist is given by

$$\pi(x, y) = V\Theta(x - y) - c \min(x, y), \quad (3.113)$$

where $\Theta(x)$ is the Heaviside step function

$$\Theta(x) = \begin{cases} 0 & \text{if } x < 0 \\ \frac{1}{2} & \text{if } x = 0 \\ 1 & \text{if } x > 0, \end{cases} \quad (3.114)$$

expressed using the half-maximum convention, and c is a constant of proportionality determining the cost of a given level of investment.

Note that this payoff function correctly determines the payoff structure of the game. If $x > y$, (i.e., the x strategist makes the larger investment), then this individual gets the

resource and incurs a cost cy , while the y strategist does not get the resource but incurs the same cost. If $x < y$, then the x strategist incurs a cost cx but does not get the resource, while the y strategist gets the resource and pays the same cost. Finally, if $x = y$ (i.e., the two individuals make the same investment), then on average each obtains the resource V with probability $\frac{1}{2}$, and both incur a cost cx .

In general, the cost need not be a linear function of the investment, and an immediate generalization of the WoA game is obtained by allowing the payoff of an x -strategist interacting with a y -strategist to be given by

$$\pi(x, y) = V\Theta(x - y) - C(\min(x, y)), \quad (3.115)$$

where C a function of the minimum of the two strategies, representing the cost associated with the investments. In the case when the cost function is the linear function $C(u) = cu$ we recover the classical WoA game with a mixed strategy Nash equilibrium given in terms of the distribution of persistence times chosen, as $p(t) = \frac{c}{V} \exp(-\frac{ct}{V})$ [Web07].

The minimum function $\min(x, y)$ can be expressed in terms of $\Theta(x)$ as

$$\min(x, y) = x\Theta(y - x) + y\Theta(x - y), \quad (3.116)$$

and thus we can write the payoff function in Equation 3.116 as

$$\pi(x, y) = V\Theta(x - y) - C(x\Theta(y - x) + y\Theta(x - y)). \quad (3.117)$$

We can replace $\Theta(x)$ in Equation 3.117 with its smooth equivalent $\theta(x)$ given by

$$\theta(x) = \frac{1}{1 + e^{-kx}}, \quad (3.118)$$

where k is the smoothing parameter. This parameter is also referred to as “inverse temperature” and models situations in which individuals are rational but payoffs have a small random component, or when gathering perfect information about decision

alternatives is costly. Note that $\theta(x)$ approximates $\Theta(x)$ as $k \rightarrow \infty$. Also, $\lim_{x \rightarrow -\infty} \theta(x) = 0$, $\lim_{x \rightarrow \infty} \theta(x) = 1$, $\theta(0) = \frac{1}{2}$, $\theta'(0) = \frac{k}{4}$, and $\theta''(0) = 0$.

We can thus model the WoA game as a continuous game with strategies varying continuously over $\mathbb{R}_+ = [0, \infty)$, and with payoff function given by

$$\pi(x, y) = V\theta(x - y) - C(x\theta(y - x) + y\theta(x - y)). \quad (3.119)$$

We can now analyze the dynamics of the game using the framework of adaptive dynamics. Consider a monomorphic population of x strategists, i.e., a population in which every individual persists for time x . The invasion fitness $f_x(y)$ of a mutant strategy y in this population of just x strategists is given by

$$\begin{aligned} f_x(y) &= \pi(y, x) - \pi(x, x) \\ &= V\theta(y - x) - C(y\theta(x - y) + x\theta(y - x)) - \left[\frac{V}{2} - C(x) \right]. \end{aligned} \quad (3.120)$$

The selection gradient $D(x)$ is given by

$$\begin{aligned} D(x) &= \left. \frac{\partial f_x(y)}{\partial y} \right|_{y=x} \\ &= \frac{Vk}{4} - \frac{C'(x)}{2}. \end{aligned} \quad (3.121)$$

To obtain the singular strategies x^* , we must solve $D(x^*) = 0$, i.e., $C'(x^*) = \frac{Vk}{2}$. If the cost function $C(x)$ is convex, i.e. $C''(x) > 0$, then $\left. \frac{dD}{dx} \right|_{x=x^*} = -\frac{C''(x^*)}{2} < 0$. Therefore, x^* is convergent stable.

$$\begin{aligned} \left. \frac{\partial^2 f_{x^*}}{\partial y^2} \right|_{y=x^*} &= \frac{\partial^2}{\partial y^2} \left\{ V\theta(y - x) - C(y\theta(x - y) + x\theta(y - x)) - \right. \\ &= \left. -\frac{C''(x^*)}{4} + \frac{kC'(x^*)}{2} \right\}. \end{aligned} \quad (3.122)$$

Therefore, x^* is an ESS if $C'(x^*) < \frac{C''(x^*)}{2k}$, or $C''(x^*) > 2kC'(x^*)$. Otherwise, x^* is an evolutionary branching point.

Let us first consider the case where the cost function is linear, i.e., $C(x) = cx$ where $c > 0$. Therefore, $C'(x) = c$, and $C''(x) = 0$. To obtain singular strategies x^* , we must solve $D(x^*) = 0$.

$$\begin{aligned} D(x^*) &= \frac{Vk}{4} - \frac{C'(x^*)}{2} \\ &= \frac{Vk}{4} - \frac{c}{2}. \end{aligned} \tag{3.123}$$

Since $D(x^*)$ does not depend on x^* , there are no singular strategies. The long-term trend of an initial strategy x can be obtained by considering the sign of $D(x)$. If $D(x) < 0$, i.e., if $c > \frac{Vk}{2}$, then x would evolve to lower and lower values. If on the other hand, $D(x) > 0$, i.e., $c < \frac{Vk}{2}$, then x would evolve to the higher and higher values. In the parlance of the WoA game, the individuals in the population will persist lesser and lesser if $c > \frac{Vk}{2}$, and longer and longer if $c < \frac{Vk}{2}$, where k is the smoothing parameter. For example, if $k = 4$ and $V = 0.25$, then the individuals will persist longer if $c < 0.5$, and less if $c > 0.5$.

Let us next consider a more interesting case in which the cost function $C(x)$ is quadratic, i.e., $C(x) = cx^2$ where $c > 0$. Therefore, $C'(x) = 2cx$, and $C''(x) = 2c$. To obtain singular strategies x^* , we must solve $D(x^*) = 0$.

$$\begin{aligned} D(x^*) &= \frac{Vk}{4} - \frac{C'(x^*)}{2} \\ &= \frac{Vk}{4} - cx^*. \end{aligned} \tag{3.124}$$

So, $x^* = \frac{Vk}{4c}$ is a singular strategy. Since $\frac{dD}{dx}|_{x=x^*} = -\frac{1}{2}C''(x^*) = -c < 0$, x^* is convergent stable.

$$\begin{aligned}\left.\frac{\partial^2 f_{x^*}}{\partial y^2}\right|_{y=x^*} &= -\frac{C''(x^*)}{4} + \frac{kC'(x^*)}{2} \\ &= \frac{1}{2}\left(\frac{k^2V}{2} - c\right).\end{aligned}\tag{3.125}$$

Therefore, if $\frac{k^2V}{2} > c$, i.e., $k^2V > 2c$, then x^* is an evolutionary branching point. If $\frac{k^2V}{2} < c$, i.e., $k^2V < 2c$, then x^* is an ESS. For example, suppose $k = 4$, and $c = 1$. If $V = 0.1$, then $x^* = 0.1$ is an ESS. If $V = 0.4$, then $x^* = 0.4$ is an evolutionary branching point.




Figure 8: Evolution of the distribution of persistence time x in a war of attrition game. Simulated on a complete network Γ_{WM} using parameters $n = 10000$, $x_0 = 0.5$, $\mu = 0.05$, $\sigma = 0.005$, $c = 1.0$, and $k = 4.0$. Darker shades of gray indicate higher concentration of individuals. (a) $V = 0.1$; the singular strategy (dashed line) $x^* = 0.1$ is an ESS. (b) $V = 0.4$; the population undergoes evolutionary branching at the singular strategy (dashed line) $x^* = 0.4$.

Figure 8 shows the results of simulating the war of attrition game on a complete network Γ_{WM} . The parameter values used in the simulations were: population size $n = 10000$, mutation rate $\mu = 0.05$, standard deviation for Gaussian mutations $\sigma = 0.005$, coefficient in the cost function $c = 1.0$, and smoothing parameter $k = 4.0$. The strategies x

were allowed to vary continuously over the interval $(0, 1)$, and the simulations were carried out for 60000 generations.

Figure 8(a) shows the time evolution of the distribution of persistence time x on a complete network, with the resource value V set to 0.1. The darker shades of gray indicate higher concentrations of individuals. The population, starting out with an initial monomorphic strategy of $x_0 = 0.5$, evolves over time to a stage in which all the individuals more or less play the strategy $x^* = 0.1$, and beyond this stage, the individuals do not deviate from this strategy by much. This is exactly what the analysis suggested; the singular strategy $x^* = 0.1$ is an ESS for the chosen parameters.

Figure 8(b) shows the time evolution of the distribution of persistence time x on a complete network, but with the resource value V set to 0.4. In this case, the population, starting out with an initial monomorphic strategy of $x_0 = 0.5$, evolves to a stage in which all the individuals play the strategy $x^* = 0.4$, and then the population undergoes evolutionary branching, again as suggested by the analysis. The population eventually evolves to a stage where a fraction of the population plays one strategy and the rest play another strategy. □

3.6 Agent-Based Models

In Section 3.4 we saw how a simple evolutionary process in which individuals who do well in a game are imitated by others who have had less success. This evolutionary process often converges to a Nash equilibrium. Thus, very little intelligence is required of the individuals who play the game. This dynamical process can be modeled on a computer in what is known as an *agent-based modeling* [Bon02].

Agent-based models (ABMs) are computer programs that use game theory to create artificial strategic actors called *agents*, set up the structure of their interactions and their

payoffs, and display and/or record the dynamics of the ensuing social order. An *evolutionary* ABM has, in addition, a *reproductive phase*, in which agents reproduce in proportion to their average success, old agents die off, and the new agents inherit the behavior of their parents, perhaps with some mutation. Thus, in an evolutionary ABM, there is a tendency for the more successful strategies to increase in frequency at the expense of the less successful strategies.

ABMs are important because most games are too complicated to admit closed-form analytical solutions. Moreover, sometimes the assumptions we make, such as, continuous time, infinite number of potential agents, etc., to get explicit algebraic solutions to games are sufficiently unrealistic that the ABM performs differently from the analytical solutions, and perhaps more like the real-life situations we are trying to model. Playing these games in real life can be very difficult, time consuming, and costly. ABMs offer us the best shot at mimicking how real-life agents might play such games. ABMs provide us a way of verifying the analytical results derived in cases where possible, and testing scenarios which do not permit mathematical analysis. Some of the benefits of agent-based modeling over other modeling techniques are:

- Agent-based modeling captures emergent phenomena. It is applicable when: individual behavior is nonlinear and can be characterized by thresholds, if-then rules, or nonlinear couplings; when individual behavior exhibits memory, path-dependence, hysteresis, non-markovian behavior, or temporal correlations, including learning, and adaption; and agent interactions are heterogeneous and can generate network effects. In such situations, aggregate differential equations will smooth out fluctuations, and thus not work.
- Agent-based modeling provides a natural description of the system. It can be used when behavior cannot be clearly defined, individual behavior is complex;

complexity of differential equations increases exponentially as the complexity of behavior increases, and becomes intractable. It can also be used when stochasticity applies to the agents' behavior; sources of randomness are applied to the right places as opposed to a noise term added more or less arbitrarily to an aggregate differential equation.

- ABMs are flexible. It is easy to add more agents to an ABM. The model provides a natural framework for tuning the complexity of agent behavior, degree of rationality, ability to learn and evolve, and rules for interactions. It is also possible to easily change levels of description and aggregation.

Some of the disadvantages of agent-based modeling include:

- Like most modeling techniques, agent-based modeling must also serve a specific purpose; a general-purpose model will not work. The model must be built at the right level of description.
- ABMs do not take into account soft factors, such as, irrational behavior, subjective choices, complex psychology, etc., which are difficult to quantify. Hence the nature of the output of ABMs vary from purely qualitative to quantitative results useable for decision making.
- ABMs look at the system not at the aggregate level but at the level of its constituent parts. Although the aggregate level could perhaps be described by just a few equations, the lower-level description involves describing the individual behavior of potentially many constituent units, thus placing heavy requirements on computational power.

It is very important to try and implement ABMs in two very different programming languages and ensure that the output of the two are the same. It is even better to have two

different people provide the two different implementations. This is particularly important when the model is too complicated to analyze mathematically. In such cases, being able to verify the results of the model using two different implementations of the ABM bolsters our confidence in our model. We have implemented our ABMs in Python and Java (see Appendix B).

Algorithm 3.1 A typical evolutionary ABM.

Input: A population network Γ of size n with strategy profile $\hat{x}^0 = \{x_1^0, x_2^0, \dots, x_n^0\}$ at generation 0; a payoff function π ; number of generations T ; and values for other model parameters.

Output: The strategy profile \hat{x}^T at time T .

```

1:  $L \leftarrow []$ 
2: for  $t$  from 1 to  $T$  do
3:    $L.append(\hat{x}^{t-1})$ 
4:   // Interaction rounds.
5:   for  $count$  from 1 to  $n$  do
6:     Two individuals  $i$  and  $j \in N(i)$  are selected (by nature) for a round of interaction,
       in which the two play a game, and  $i$  receives a payoff  $\pi(x_i^{t-1}, x_j^{t-1})$ .
7:   end for
8:   // Update rounds.
9:   for  $count$  from 1 to  $n$  do
10:    Two individuals  $i$  and  $j \in N(i)$  are selected (again by nature) for a round of
       update, in which  $i$  either updates its strategy  $x_i^t$  to that of  $j$  strategy  $x_j^{t-1}$  with a
       possible mutation, or retains its own strategy.
11:  end for
12: end for

```


Algorithm 3.1 provides the pseudo-code for a typical evolutionary ABM. We start with a population network of n individuals labeled $i = 1, \dots, n$ with a strategy profile $\hat{x}^0 = \{x_1^0, x_2^0, \dots, x_n^0\}$ at generation 0. Since we wish to allow the possibility of complex population structures, we identify the population with the set of vertices in a network Γ . The structure of Γ determines which individuals in the population can interact. Strictly speaking, two networks are required to specify the evolutionary dynamics: an interaction graph, Γ_I , specifies that two individuals in the population can interact by playing the game only if they are adjacent in Γ_I , and an updating graph, Γ_U , specifies that an individual in the population can update its strategy by comparing her state to the states only of those individuals adjacent to it in Γ_U . In this thesis, for simplicity, we shall assume that the interaction and updating networks are the same i.e., $\Gamma_I = \Gamma_U = \Gamma$. Given an individual $i \in \Gamma$, recall from Chapter 2 that the set of neighbors of i (i.e., the set of individuals adjacent to i in Γ) is denoted by $N(i)$.

The game is played for T generations, each involving n interactions and updates. During an interaction, say at generation t , an individual i plays a game (with a payoff function π) with a neighboring individual $j \in N(i)$ and obtains a payoff. The two individuals employ strategies x_i^{t-1} and x_j^{t-1} from the previous generation $t - 1$. During an update at generation t , an individual i either updates its strategy x^t in the generation t with strategy x_j^{t-1} of its neighbor $j \in N(i)$ in generation $t - 1$ with a possible (Gaussian) mutation, or retains her own strategy.

The details of how i and j are selected for update and how the update is actually carried out depends on the update rule \mathcal{U} . Let s denote the selection strength, and π_i denote the payoff of individual i obtained during the interaction round, and let f_i denote the individual's fitness. The relationship between the payoff and fitness is given by the

fitness-to-fecundity map

$$f_i = e^{s\pi_i}. \quad (3.126)$$

We consider the following update rules (see [SF07] for a detailed survey of these and other update rules):

- *Fermi (FE)* - The individuals i and $j \in N(i)$ are randomly picked and the probability $p_{i \leftarrow j}$ that i inherits j 's strategy is given by

$$p_{i \leftarrow j} = \frac{1}{1 + e^{-s(\pi_j - \pi_i)}}. \quad (3.127)$$

- *Replication (RE)* - The individuals i and $j \in N(i)$ are randomly picked and the probability $p_{i \leftarrow j}$ that i inherits j strategy is given by

$$p_{i \leftarrow j} = \begin{cases} 0 & \text{if } \pi_i \geq \pi_j \\ \frac{s(\pi_j - \pi_i)}{M - m} & \text{otherwise,} \end{cases} \quad (3.128)$$

where $M = \max_{k \in \Gamma} \pi_k$, and $m = \min_{l \in \Gamma} \pi_l$.

- *Imitation (IM)* - The individual i is randomly selected and the individual $j \in N(i) \cup i$ is picked with probability p given by

$$p = \frac{f_j}{\sum_{k \in N(i) \cup i} f_k}. \quad (3.129)$$

- *Birth-death (BD)* - The individual j is selected with probability p given by

$$p = \frac{f_j}{\sum_{k \in \Gamma} f_k}. \quad (3.130)$$

The individual $i \in N(j)$ is randomly selected. On a complete network Γ_{WM} , this update rule is also called the *Moran process*, named after Patrick Moran, an Australian statistician who made significant contributions to probability theory and its application to population and evolutionary genetics.

- *Death-birth (DB)* - The individual i is picked at random, and the individual $j \in N(i)$ is selected with probability p given by

$$p = \frac{f_j}{\sum_{k \in N(i)} f_k}. \quad (3.131)$$

The evolutionary dynamics on a complete network Γ_{WM} is generally independent of the update rule, but on other networks, there is usually a sensitive dependence of the dynamics on the update rule.

CHAPTER 4

DISCRETE GAMES

The only thing that will redeem mankind is cooperation.

- Bertrand Russell

4.1 Introduction

In this chapter we revisit the three symmetric, two-strategy, pairwise games we encountered in the previous chapter: prisoner's dilemma, chicken, and coordination games. We use these games as models for cooperative behavior by formulating them in terms of two strategies, namely cooperation and defection.

We consider a population of individuals playing the games through repeated pairwise interactions. In a well-mixed population, we assume the interactions to be assortative [Gra90], i.e., each individual interacts with other individuals of its own "type" with a certain probability called *degree of assortativity* or simply *assortativity*, and with random individuals from the population otherwise. In a population represented as a network, the structure of the network dictates who can interact with whom, and thus induces a degree of network assortativity in game interactions, which bears an inverse relationship to the average degree of the network ¹.

In each of the three games, the individuals playing the game have an incentive to be selfish and defect. As a result, everyone does worse than if they had all cooperated

¹This is because, on a network with a small average degree, an individual can only interact with small number of other individuals, whereas on a network with high average degree, an individual can interact with many other individuals. In the extreme case of a complete network, an individual can interact with every other individual in the network.

[Kol98]. We are interested in conditions that would lead to cooperative behavior among selfish individuals. As we shall see, assortative interactions (both in well-mixed populations and on networks) provide a mechanism for the promotion of cooperation and thus to Hicks (or socially) optimal outcomes. In addition, network properties such as clustering and homophily also have an effect on cooperation.

We propose agent-based models for emulating symmetric, two-strategy, pairwise games on a computer. We analyze these games for a well-mixed population using the replicator equation [TJ78], suitably modified to incorporate assortative interactions. We present results of simulating the games on model and empirical networks and provide a general discussion of the results. We conclude the chapter with suggestions for future inquiry.

4.2 Assortative Interactions

An *assortative interaction* [Gra90] is an interaction among individuals of the same type, which in the context of games would mean individuals playing the same strategy. Assortativity is closely related to the idea of kin selection [Ham63], which was suggested by William Donald Hamilton as a possible mechanism for the evolution of altruistic behavior. There are many ways in which interactions can be assortative. In a well-mixed population, we can quantify the extent to which an individual interacts with another individual of its own type using a parameter $r \in (0, 1)$ called the *degree of assortativity* or simply *assortativity*, which is interpreted as follows: with probability r , an individual interacts with another individual of its own type, and with probability $1 - r$, the individual interacts with a random individual from the population. We next derive the replicator equation for a symmetric, two-strategy game with assortative interactions.

Consider a well-mixed population of n individuals playing a symmetric, two-strategy (C and D), pairwise game with payoff matrix

$$\pi = \begin{array}{c} C \quad D \\ \begin{array}{cc} C & (a, a) \quad (b, c) \\ D & (c, b) \quad (d, d) \end{array} \end{array}, \quad (4.1)$$

where $a, b, c, d \in \mathbb{R}$. Let x denote the fraction of individuals playing the strategy C , so the fraction of the population playing the strategy D is $1 - x$. Let r denote the degree of assortativity. The fitnesses f_C^r and f_D^r of the strategies C and D , using Equation 3.72, are then given by

$$\begin{aligned} f_C^r &= r\pi(C, C) + (1 - r)[x\pi(C, C) + (1 - x)\pi(C, D)] \\ &= ra + (1 - r)[xa + (1 - x)b], \end{aligned} \quad (4.2)$$

and

$$\begin{aligned} f_D^r &= r\pi(D, D) + (1 - r)[x\pi(D, C) + (1 - x)\pi(D, D)] \\ &= rd + (1 - r)[xc + (1 - x)d]. \end{aligned} \quad (4.3)$$

The evolutionary dynamics of the population, using Equation 3.77, is given by

$$\begin{aligned} \dot{x} &= x(1 - x)(f_C^r - f_D^r) \\ &= x(1 - x)\{r(a - d) + (1 - r)[x(a - c) + (1 - x)(b - d)]\}, \end{aligned} \quad (4.4)$$

which is the replicator equation for symmetric, two-strategy, pairwise games with assortative interactions. Note that when $r = 0$, the equation reduces to

$$\dot{x} = x(1 - x)[x(a - b - c + d) + b - d], \quad (4.5)$$

which is the standard replicator equation for symmetric, two-strategy, pairwise games.

On a network, a combination of the structure of the network and the update rule enforce assortative interactions. An individual is more likely to interact with other individuals of its own type on a network with small average degree d than on a network with a large value for d . Thus, the network assortativity in game interactions bears an inverse relationship with the average degree of the network.

4.3 Agent-Based Model

An agent-based model [Bon02] provides a natural way of describing how individuals in a population might play a symmetric, two-strategy, pairwise game with payoff matrix given by Equation 4.1. One can simulate (B.4 describes how this is done) such a model on a computer in order to study the evolution of the frequencies of the C and D strategies.

We start with a population network Γ of size n with a strategy profile $\hat{x}^0 = \{x_1^0, x_2^0, \dots, x_n^0\}$ at generation 0, where $x_i^0 \in \{C, D\}$ for $1 \leq i \leq n$. Roughly, a fraction x_0 of individuals in the population play the strategy C and the remaining individuals play the strategy D . The update rule \mathcal{U} for the game is one of the five rules (FE, RE, IM, BD, or DB) we discussed in the previous chapter, the selection strength is s , and the degree of assortativity (applicable only if $\Gamma = \Gamma_{\text{WM}}$, i.e., if Γ is a complete network) is r .

The game is played for T generations, each involving n interactions and updates. During an interaction, say at generation t , an individual i plays the game with a neighboring individual $j \in N(i)$ and obtains a payoff π_i . On a complete network Γ_{WM} , the interaction is assortative, i.e., with probability r , the individual i plays the game with another individual j of its own type, and with probability $1 - r$, it plays the game with a randomly picked individual j . The two individuals i and j employ strategies x_i^{t-1} and x_j^{t-1} from the previous generation. During an update at generation t , an individual i updates its strategy x^t in the generation t with strategy x_j^{t-1} of its neighbor $j \in N(i)$. The updates we

consider are non-innovative, i.e., the strategies are copied/inherited without error. As a result, once the population fixates to a particular strategy, i.e., all the individuals in the population play the same strategy, the strategy profile of the population remains unchanged² in future generations.

Based on how the individuals i and j are selected for interactions and updates, we have two distinct algorithms for implementing the games on a computer. In the *sequential* approach, the n update rounds follow the n interaction rounds for a given generation. So the pair of individuals (i, j) considered for an interaction is generally not the same as the pair (i, j) of individuals considered for an update. Figure 9 shows a schematic of interactions and updates in the sequential approach. The box on the top depicts n rounds of interactions while the box below depicts n rounds of updates. A circled vertex represents the focal individual, i.e., the individual receiving a payoff during the interaction round, or the individual (possibly) inheriting a strategy during the update round.

Algorithm 4.2 provides the pseudo-code for implementing symmetric, two-strategy, pairwise games on a computer using sequential interactions and updates. In this case the update rule can be one of FE, RE, IM, BD, or DB. Each generation t consists of n interaction rounds followed by n update rounds (see Figure 9). In each interaction round, an individual i is randomly sampled from Γ . If Γ is a complete network Γ_{WM} , then i plays the game with an individual of its own strategy type³ with probability r and receives a payoff $\pi_i = \pi(x_i^{t-1}, x_i^{t-1})$. Otherwise, i plays the game with a randomly sampled individual $j \in N(i)$, and receives a payoff $\pi_i = \pi(x_i^{t-1}, x_j^{t-1})$. At the end of the n interactions, on average every individual in the population would have played the game

²The software implementations can take advantage of this fact and terminate the simulation upon fixation of strategies.

³We can either assume that such an individual always exists in the population, or that i actively seeks out such an individual if it exists, and forgoes playing the game at generation t otherwise. Both these assumptions result in the same dynamics up to rate of evolution of strategies. Our software implementation makes the former assumption.




Figure 9: Sequential interactions and updates per generation in a discrete game.

and thus received a payoff; if an individual does not interact in a generation, its payoff in that generation is zero.

Two individuals i and $j \in N(i)$ are selected based on the update rule \mathcal{U} (see Chapter 3) and i either ends up inheriting j 's strategy from generation $t - 1$ or retaining its own strategy.

At the start of each generation t , the strategy profile \hat{x}^{t-1} from the previous generation is saved in a list L of strategy profiles, and at the end of T generations the list L is returned as output.

Algorithm 4.2 Symmetric, two-strategy, pairwise games on complex networks, using sequential interactions and updates.

Input: A population network Γ of size n with strategy profile $\hat{x}^0 = \{x_1^0, x_2^0, \dots, x_n^0\}$ at generation 0, where $x_i^0 \in \{C, D\}$ for $1 \leq i \leq n$; a 2×2 payoff matrix π ; update rule \mathcal{U} (FE, RE, IM, BD, DB); selection strength s ; degree of assortativity r ; and number of generations T .

Output: A list $L = \{\hat{x}^0, \hat{x}^1, \dots, \hat{x}^{T-1}\}$.

```

1:  $L \leftarrow []$ 
2: for  $t$  from 1 to  $T$  do
3:    $L.append(\hat{x}^{t-1})$ 
4:   // Interactions.
5:   for  $count$  from 1 to  $n$  do
6:      $i \leftarrow$  random individual from  $\Gamma$ 
7:     if  $\Gamma = \Gamma_{WM}$  and  $random() < r$  then
8:        $\pi_i \leftarrow \pi(x_i^{t-1}, x_i^{t-1})$ 
9:     else
10:       $j \leftarrow$  random individual from  $N(i)$ 
11:       $\pi_i \leftarrow \pi(x_i^{t-1}, x_j^{t-1})$ 
12:    end if
13:  end for
14:  // Updates.
15:  for  $count$  from 1 to  $n$  do
16:    Two individuals  $i$  and  $j \in N(i)$  are selected based on the update rule  $\mathcal{U}$  (see Chapter 3) and  $i$  either inherits  $j$ 's strategy from generation  $t - 1$  or retains its own strategy
17:  end for

```

18: **end for**
 19: **return** L

In the *simultaneous* approach, at each generation, each of the n updates immediately follows two interactions, one in which an individual i interacts with another individual $k \in N(i)$ and obtains a payoff π_i and the other in which an individual $j \in N(i)$ ($j \neq k$) interacts with another individual $l \in N(j)$ ($l \neq i$) and obtains a payoff π_j . The same pair (i, j) is then considered for an update. Figure 10 shows a schematic of the interactions and updates in the simultaneous approach. A circled vertex again represents a focal individual, i.e., the individual receiving a payoff during the interaction round, or the individual (possibly) inheriting a strategy during the update round.

Algorithm 4.3 provides the pseudo-code for implementing symmetric, two-strategy, pairwise games on a computer using simultaneous interactions and updates (see Figure 10). In this case the update rule \mathcal{U} can either be FE or RE. Each generation t consists of n rounds, each of which involves an interaction immediately followed by an update. In each round, an individual i and an individual $j \in N(i)$ are randomly sampled from Γ . If Γ is a complete network Γ_{WM} , then i plays the game with an individual of its own strategy type with probability r and receives a payoff $\pi_i = \pi(x_i^{t-1}, x_i^{t-1})$. Otherwise, i plays the game with a randomly sampled individual $k \in N(i)$, and receives a payoff $\pi_i = \pi(x_i^{t-1}, x_k^{t-1})$. Similarly, j plays the game with an individual of its own strategy type with probability r and receives a payoff $\pi_j = \pi(x_j^{t-1}, x_j^{t-1})$. Otherwise, j plays the game with a randomly sampled individual $l \in N(i)$, and receives a payoff $\pi_j = \pi(x_j^{t-1}, x_l^{t-1})$. Note that on $\Gamma \neq \Gamma_{\text{WM}}$, i must be distinct from l and j must be distinct from k ; otherwise, the process continues with the next generation. The individual i updates its strategy x_i^t to the strategy x_j^{t-1} of a randomly sampled individual $j \in N(i)$, with probability $p_{i \leftarrow j}$ determined by the update rule, and with probability $1 - p_{i \leftarrow j}$, i retains its own strategy.




Figure 10: Simultaneous interactions and updates per generation in a discrete game.

At the start of each generation t , the strategy profile \hat{x}^{t-1} from the previous generation is saved in a list L of strategy profiles, and at the end of T generations the list L is returned as output.

Algorithm 4.3 Symmetric, two-strategy, pairwise games on complex networks, using simultaneous interactions and updates.

Input: A population network Γ of size n with strategy profile $\hat{x}^0 = \{x_1^0, x_2^0, \dots, x_n^0\}$ at generation 0, where $x_i^0 \in \{C, D\}$ for $1 \leq i \leq n$; a 2×2 payoff matrix π ; update rule \mathcal{U} (FE or RE); selection strength s ; degree of assortativity r ; and number of generations T .

Output: A list $L = \{\hat{x}^0, \hat{x}^1, \dots, \hat{x}^{T-1}\}$.

```

1:  $L \leftarrow \square$ 
2: for  $t$  from 1 to  $T$  do
3:    $L.append(\hat{x}^{t-1})$ 
4:   // Interactions.
5:   for  $count$  from 1 to  $n$  do
6:      $i \leftarrow$  random individual from  $\Gamma$ 
7:      $j \leftarrow$  random individual from  $N(i)$ 
8:     if  $\Gamma = \Gamma_{WM}$  and  $random() < r$  then
9:        $\pi_i \leftarrow \pi(x_i^{t-1}, x_i^{t-1})$ 
10:    else
11:       $k \leftarrow$  random individual from  $N(i)$ 
12:      if  $\Gamma \neq \Gamma_{WM}$  and  $k = j$  then
13:        continue
14:      end if
15:       $\pi_i \leftarrow \pi(x_i^{t-1}, x_k^{t-1})$ 
16:    end if
17:    if  $\Gamma = \Gamma_{WM}$  and  $random() < r$  then
18:       $\pi_j \leftarrow \pi(x_j^{t-1}, x_j^{t-1})$ 
19:    else
20:       $l \leftarrow$  random individual from  $N(j)$ 
21:      if  $\Gamma \neq \Gamma_{WM}$  and  $l = i$  then
22:        continue
23:      end if
24:       $\pi_j \leftarrow \pi(x_j^{t-1}, x_l^{t-1})$ 
25:    end if

```

26: The probability $p_{i \leftarrow j}$ that individual i inherits individual j 's strategy is calculated based on the update rule \mathcal{U} (see Chapter 3)

27: **if** $\text{random}() < p_{i \leftarrow j}$ **then**

28: $x_i^t \leftarrow x_j^{t-1}$

29: **else**

30: $x_i^t \leftarrow x_i^{t-1}$

31: **end if**

32: **end for**

33: **end for**

34: **return** L

4.4 Prisoner's Dilemma

An act of *cooperation* by an individual (donor) towards another individual (recipient) is one which benefits the recipient by an amount b , but incurs a cost c to the donor, where $b, c \in \mathbb{R}_+$ and $b > c$.

Two friends John and Bill are engaged in an act of cooperative behavior in which each of them can either cooperate or defect, i.e., not cooperate. If both John and Bill cooperate, then they both get a benefit of b but also incur a cost of c , and thus receive a payoff of $b - c$. If both John and Bill defect, then they neither benefit nor incur a cost, and thus receive a payoff of 0. If John cooperates and Bill defects, then Bill receives a benefit of b but incurs no cost while John receives no benefit but incurs a cost of c . Thus the payoff that John receives is $-c$ while the payoff that Bill receives is b . Finally, if Bill cooperates and John defects, then John receives a payoff of b while Bill receives $-c$.

If we denote the act of cooperation by C and that of defection by D , then we can define a symmetric, two-strategy, pairwise contest game with the payoff matrix:

$$\pi = \begin{array}{c} \\ \begin{array}{cc} & C & D \\ C & (b-c, b-c) & (-c, b) \\ D & (b, -c) & (0, 0) \end{array} \end{array}. \quad (4.6)$$

By Theorem 3.4, transforming the payoff matrix of a game using an affine transformation does not alter the equilibria of the game. So we can transform the above matrix by dividing each element by b to obtain the following payoff matrix:

$$\pi = \begin{array}{c} \\ \begin{array}{cc} & C & D \\ C & (1 - \frac{c}{b}, 1 - \frac{c}{b}) & (-\frac{c}{b}, 1) \\ D & (1, -\frac{c}{b}) & (0, 0) \end{array} \end{array}. \quad (4.7)$$

Finally, denoting the cost-to-benefit ratio $\frac{c}{b}$ by ρ , we can express payoff matrix for the game more compactly as

$$\pi = \begin{array}{c} C \\ D \end{array} \begin{array}{cc} C & D \\ \left(\begin{array}{cc} (1-\rho, 1-\rho) & (-\rho, 1) \\ (1, -\rho) & (0, 0) \end{array} \right), \end{array} \quad (4.8)$$

where $0 < \rho < 1$.

It is easy to see that a game with payoff matrix given by Equation 4.8 belongs to the prisoner's dilemma class of games we encountered in the previous chapter. Therefore, the set N of Nash equilibria for the game is $N = \{(D, D)\}$ and the set E of ESSs for the game is $E = \{(D, D)\}$. Comparing the payoff matrix for the game with the payoff matrix for a general symmetric, two-strategy, pairwise contest game given by Equation 4.1, we have $a = 1 - \rho$, $b = -\rho$, $c = 1$, and $d = 0$. Substituting for a, b, c , and d in Equation 4.5 gives us the evolutionary dynamics (without assortment) for the game as $\dot{x} = -x(1-x)\rho$, where x is the fraction of individuals playing the C strategy and $1-x$ is the fraction of individuals playing the D strategy. The set F of fixed points for the dynamics is $F = \{(0, 1), (1, 0)\}$ and the set A of asymptotically stable fixed points is $A = \{(0, 1)\}$. The four sets N, E, F , and A are related by $E = A \subseteq N \subseteq F$. Therefore, a population of individuals playing the prisoner's dilemma game with payoff matrix given by Equation 4.8, starting from an initial fraction $x_0 > 0$ of cooperators, will eventually end up in the all-defector ($x^* = 0$) equilibrium state. So a natural question to ask is: how can cooperation evolve in a population of individuals playing the prisoner's dilemma game?

Prisoner's dilemma has been extensively studied as a model for cooperative behavior [Now06b]. There are five mechanisms that have been proposed for the evolution of cooperation: kin selection, direct reciprocity, indirect reciprocity, network reciprocity, and group selection. Each mechanism suggests a simple rule—an inequality involving the

cost-to-benefit ratio ρ of the cooperative act—which specifies whether natural selection can lead to cooperation.

Kin Selection

Kin selection [Ham63] is the ingenious idea proposed by William Donald Hamilton that natural selection can favor cooperation if the donor and the recipient of a cooperative act are genetic relatives. Relatedness is defined as the probability of sharing a gene. For example, the probability that two brothers share the same gene by descent is $\frac{1}{2}$, while the same probability for cousins is $\frac{1}{8}$. More precisely, kin selection states that the coefficient of relatedness r must exceed the cost-to-benefit ratio ρ of the cooperative act, i.e., $r > \rho$. This simple rule is known as Hamilton's rule. The idea was informally articulated much earlier, supposedly in an English pub, by the irascible British geneticist and evolutionary biologist J.B.S. Haldane when he famously quipped “I will jump into the river to save two brothers or eight cousins”.

Direct Reciprocity

Cooperation is observed not only among relatives but also between unrelated individuals or even between members of different species. Such considerations led the American evolutionary sociobiologist Robert Trivers to propose direct reciprocity [Tri71] as another mechanism for the evolution of cooperation. He considered the iterated prisoner's dilemma game in which there are repeated encounters between the same two individuals. If one of them cooperates now, the other may cooperate later. Hence, it might pay off to cooperate.

In search for a good strategy for playing the iterated prisoner's dilemma, the American political scientist Robert Axelrod, in two computer tournaments [Axe06], discovered that

the “winning strategy” was the simplest of all, tit-for-tat (TFT). An individual playing this strategy always starts by cooperating and then does whatever the opponent has done in the previous round: cooperation for cooperation, defection for defection.

Unfortunately, the TFT strategy has a weakness. If there are erroneous moves, caused by “trembling hands” or “fuzzy minds”, then the performance of TFT declines. TFT cannot correct mistakes, because an accidental defection leads to a long sequence of retaliation. In such cases, TFT is replaced by generous-tit-for-tat (GTFT) [NS92]. An individual playing the GTFT strategy cooperates upon cooperation, but sometimes (with probability $1 - \rho$) also cooperates upon defection, where ρ is the cost-to-benefit ratio of the cooperative act.

TFT is subsequently replaced by the win-stay-lost-shift (WSLS) strategy. An individual playing WSLS repeats the previous move whenever it is doing well, but changes otherwise [NS93]. WSLS is more robust than either TFT or GTFT. TFT is an efficient catalyst of cooperation in a society where nearly everybody is a defector, but once cooperation is established WSLS is better able to maintain it.

Direct reciprocity can lead to the evolution of cooperation only if the probability w of another encounter between the same two individuals exceeds the cost-to-benefit ratio, i.e., $w > \rho$.

Indirect Reciprocity

Interactions among individuals are often asymmetric and fleeting. One individual is in a position to help another, but there is no possibility for a direct reciprocation. We often help strangers who are in need. We donate to charities that do not donate to us.

In the standard framework of indirect reciprocity, there are randomly chosen, pairwise encounters, where the same two individuals need not meet again. One individual acts as

donor and the other as recipient. The donor can decide whether or not to cooperate. The interaction is observed by a subset of the population who might inform others. Reputation allows evolution of cooperation by indirect reciprocity [NS98]. Theoretical and empirical studies of indirect reciprocity show that people who are more helpful are more likely to receive help [NS05].

Although simple forms of indirect reciprocity can be found in animals [NS05], only humans seem to engage in the full complexity of the game. Indirect reciprocity has substantial cognitive demands. Not only do the individuals have to remember their own interactions, but also monitor the ever-changing social network of the group.

Indirect reciprocity can promote cooperation if the probability q to know someone's reputation exceeds the cost-to-benefit ratio ρ of the cooperative act, i.e., $q > \rho$.

Network Reciprocity

The argument for natural selection of defection is based on a well-mixed population, in which everyone is equally likely to interact everyone else. This approximation is used by all standard approaches to evolutionary game dynamics [HS98, Smi82]. Real populations are not well-mixed. Spatial structures or social networks imply that some individuals interact more often than others. The approach of capturing this effect is evolutionary network theory [LHN05], and the resulting mechanism for the evolution of cooperation is called network reciprocity.

Games on networks are easy to study by computer simulation, but difficult to analyze mathematically because of the enormous number of possible configurations that can arise. Nonetheless, a simple rule $\frac{1}{d} > \rho$ determines if network reciprocity can favor cooperation, where d is the average degree of the network and ρ is the cost-to-benefit ratio of a cooperative act.

Group Selection

Selection acts not only on individuals but also on groups. A group of cooperators might be more successful than a group of defectors. In a simple model of group selection [TN06], a population is subdivided into groups. Cooperators help others in their own group.

Defectors do not help. Individuals reproduce in proportion to their payoff. Offspring are added to the same group. If a group reaches a certain size it can split into two. In this case, another group becomes extinct in order to constrain the total population size. Only individuals reproduce, but selection emerges on two levels. There is competition between groups because some groups grow faster and split more often. In particular, pure cooperator groups grow faster than pure defector groups, while in any mixed group defectors reproduce faster than cooperators. Therefore, selection on the lower intra-group level favors defectors, while selection on the higher inter-group level favors cooperators.

In the limit of weak selection and rare group splitting, group selection allows evolution of cooperation provided $\frac{1}{1+\frac{n}{m}} > \rho$, where n is the maximum group size, m the number of groups, and ρ is the cost-to-benefit ratio of a cooperative act.

In what follows, we consider assortative interactions in a well-mixed population, characterized by the parameter $r \in [0, 1]$, as a possible mechanism for the promoting cooperation. We carry out a complete analysis of the prisoner's dilemma game using the replicator equation modified to incorporate assortative interactions. Since the network assortativity in game interactions bears an inverse relationship to its average degree, knowing how the assortativity parameter r affects cooperation in a well-mixed population also informs us about how cooperation would be affected, albeit only qualitatively, when the game is played out on networks.

4.4.1 Analysis

Consider a population of individuals playing the prisoner's dilemma game with the payoff matrix given by Equation 4.8. Let $x(t)$ denote the fraction of cooperators at time t , and let $x(0) = x_0$. Let the degree of assortativity be r . Substituting for a, b, c , and d in Equation 4.4, we get the evolutionary dynamics for the population as

$$\dot{x} = x(1-x)(r-\rho). \quad (4.9)$$

The phase line plots for the above system, for different values of assortativity r , are shown in Figure 11. The system has two equilibria, $x^* = 0$ and $x^* = 1$, so if a population




Figure 11: Phase line plots for the system 4.9 for different values of assortativity r .

starts out with every individual cooperating or every individual defecting, then it remains that way. When $r = 0$ or when $r < \rho$, the equilibrium point $x^* = 0$ is stable, while the equilibrium point $x^* = 1$ is unstable. So a population, starting from an initial fraction $x_0 > 0$ of cooperators, will evolve to a state in which every individual defects. However, when $r > \rho$, the stability of the two equilibrium points is reversed, and the population will evolve to a state in which everyone cooperates.

Example 4.1. Consider a prisoner’s dilemma game with $\rho = 0.3$, i.e., with payoff matrix

$$\pi = \begin{array}{c} C \quad D \\ C \begin{pmatrix} (0.7, 0.7) & (-0.3, 1) \\ D \begin{pmatrix} (1, -0.3) & (0, 0) \end{pmatrix} \end{pmatrix}. \quad (4.10)$$

We simulate the game on a complete network Γ_{WM} with parameter values $n = 10000$, $s = 1$, $x_0 = 0.5$, $\mathcal{U} = \text{FE2}$, $T = 5000$, $\nu = 1$, and $\tau = 1$. Figures 12(a)-(c) show the evolution of the fraction x of cooperators for different values of assortativity r . As suggested by the analysis, when $r = 0$ and $r < \rho$ the population evolves to the all-defector state, and when $r > \rho$ the population evolves to the all-cooperator state. \square




Figure 12: Evolution of the fraction x of cooperators in a symmetric, two-strategy, pairwise prisoner’s dilemma game for different values of assortativity r . (a) $r = 0$. (b) $r = 0.2$. (c) $r = 0.4$. Simulation parameters: complete network Γ_{WM} , $\rho = 0.3$, $n = 10000$, $x_0 = 0.5$, $\mathcal{U} = \text{FE2}$, $s = 1$, $T = 5000$, $\nu = 1$, and $\tau = 1$.

Thus in a well-mixed population, assortative interactions provide a mechanism for promoting cooperative behavior. It is worth noting that the relation $r > \rho$, which provides a necessary condition for evolution of cooperation, has the same form as the equation for kin selection [Gra90, Ham63], which has long been suggested as a possible mechanism for the evolution of altruistic behavior.

On a network, as we remarked earlier, assortativity increases with decreasing values of the average degree d of the network. So when the prisoner's dilemma game is played on networks, we expect cooperation to emerge on networks with small values of d . We explore this and other network effects in detail in the following section.

4.4.2 Simulation Results

In this section we present the results of simulating the prisoner's dilemma game on networks having different structural properties. All the simulations were carried out using the following values for the parameters involved: payoff matrix π given by Equation 4.8, population size $n = 10000$, initial fraction of cooperators $x_0 = 0.5$, number of generations $T = 10000$, report frequency $\nu = 100$, number of trials $\tau = 1$, update rule $\mathcal{U} = \text{FE2}$, and selection strength $s = 1$. See B.4 for details on how to simulate our discrete game models. The results are presented using three different kinds of plots depicting the following:

1. Variation of the long-term value x_∞ of the fraction x of cooperators, averaged over the last 10% generations, with assortativity r on complete networks (average degree d on other networks) and the cost-to-benefit ratio ρ . The values of $r \in [0, 1]$ and $\rho \in (0, 1)$ were varied in steps of 0.01, and the value of d was varied from 20 to 4 in steps of -2. The intensity of the shade in the plot indicates the level of cooperation, with black corresponding to full cooperation ($x_\infty = 1$) and white corresponding to no cooperation ($x_\infty = 0$).
2. Variation of x_∞ with r (or d) for a fixed value of ρ . The dashed line in the case of complete networks indicates values predicted by analysis.
3. Variation of x_∞ with ρ for a fixed value of r (or d). The dashed line in the case of complete networks indicates values predicted by analysis. We denote the fractional

area under the x_∞ curve relative to the total area of the $[x_\infty, \rho]$ square by \mathcal{C} . We refer to \mathcal{C} as the *cooperation index* and use it as a measure of the extent to which a network promotes cooperation. For a well-mixed population, the value of $\mathcal{C}(r)$ for a given value of assortativity r can be calculated analytically as $\mathcal{C}(r) = \int_0^1 x_\infty(r, \rho) d\rho$, where $x_\infty(r, \rho) = 1$ if $r > \rho$ and 0 otherwise.

Complete Network

The first set of results concerns complete networks Γ_{WM} . Figures 13(a)(b) respectively show how the predicted and simulated values of the long-term fraction x_∞ of cooperators vary with assortativity r and cost-to-benefit ratio ρ . Figure 13(c) shows how x_∞ varies with r when $\rho = 0.25$. Figure 13(d) shows how x_∞ varies with ρ when $r = 0.25$.

Model Networks with Varying Average Degree

The next set of results concerns model networks with varying average degree d . Figures 14(a)(d)(g) show how the long-term fraction x_∞ of cooperators varies with average degree d and cost-to-benefit ratio ρ , on random regular Γ_{RR} , Erdős-Rényi Γ_{ER} , and Barabási-Albert Γ_{BA} networks respectively. Figures 14(b)(e)(h) show how x_∞ varies with d on Γ_{RR} , Γ_{ER} , and Γ_{BA} networks respectively, when $\rho = 0.1$. Figures 14(c)(f)(i) show how x_∞ varies with ρ on Γ_{RR} , Γ_{ER} , and Γ_{BA} networks respectively, when $d = 4$.

Model Networks with Clustering

The next set of results concerns power-law networks $\Gamma_{\text{BA}}^{\text{C}}$ with average degree $d = 4$, and varying clustering coefficient $C^{(2)}$. Figures 15(a)-(d) show how the long-term fraction x_∞




Figure 13: Variation of the long-term fraction x_∞ of cooperators with assortativity r and cost-to-benefit ratio ρ , on a complete network. (a) x_∞ (predicted) versus r and ρ . (b) x_∞ (simulated) versus r and ρ . (c) x_∞ versus r when $\rho = 0.25$. (d) x_∞ versus ρ when $r = 0.25$ ($\mathcal{C} = 0.252$). Simulation parameters: π given by Equation 4.8, $n = 10000$, $x_0 = 0.5$, $T = 10000$, $\nu = 100$, $\tau = 1$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

of cooperators varies with the cost-to-benefit ratio ρ on networks with clustering coefficient $C^{(2)} = 0, 0.2, 0.4$, and 0.6 .




Figure 14: Variation of the long-term fraction x_∞ of cooperators with average degree d and cost-to-benefit ratio ρ , on model networks. (a)-(c) Random regular network Γ_{RR} . (d)-(f) Erdős-Rényi network Γ_{ER} . (g)-(i) Barabási-Albert network Γ_{BA} . (a)(d)(g) x_∞ versus d and ρ . (b)(e)(h) x_∞ versus d when $\rho = 0.1$. (d)(f)(i) x_∞ versus ρ when $d = 4$ ($\mathcal{C} = 0.313, 0.283, 0.267$). Simulation parameters: π given by Equation 4.8, $n = 10000$, $x_0 = 0.5$, $T = 10000$, $\nu = 100$, $\tau = 1$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 15: Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on model power-law networks with clustering coefficient $C^{(2)}$. (a) $C^{(2)} = 0$ ($\mathcal{C} = 0.269$). (b) $C^{(2)} = 0.2$ ($\mathcal{C} = 0.264$). (c) $C^{(2)} = 0.4$ ($\mathcal{C} = 0.255$). (d) $C^{(2)} = 0.6$ ($\mathcal{C} = 0.248$). Simulation parameters: π given by Equation 4.8, $n = 10000$, $d = 4$, $x_0 = 0.5$, $T = 10000$, $\nu = 100$, $\tau = 1$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

Model Networks with Homophily

The next set of results concerns power-law networks Γ_{BA}^h with average degree $d = 4$, and varying homophily coefficient h . Figures 16(a)-(d) show how the long-term fraction x_∞ of

cooperators varies with the cost-to-benefit ratio ρ on networks with homophily coefficient $h = -0.1, 0, 0.1, \text{ and } 0.2$.




Figure 16: Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on model power-law networks with homophily coefficient h . (a) $h = -0.1$ ($\mathcal{C} = 0.293$). (b) $h = 0$ ($\mathcal{C} = 0.250$). (c) $h = 0.1$ ($\mathcal{C} = 0.239$). (d) $h = 0.2$ ($\mathcal{C} = 0.372$). Simulation parameters: π given by Equation 4.8, $n = 10000$, $d = 4$, $x_0 = 0.5$, $T = 10000$, $v = 100$, $\tau = 1$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

Empirical Networks

The last set of results concerns empirical networks. Figures 17(a)-(f) show how the long-term fraction x_∞ of cooperators varies with the cost-to-benefit ratio ρ on: the Γ^{internet} network (a); the $\Gamma^{\text{hepcoauth}}$ network (b); the $\Gamma^{\text{astrocoauth}}$ network (c); the Γ^{fb} network (d); the Γ^{p2p} network (e); and the Γ^{protein} network (f).




Figure 17: Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on empirical networks. See Table 2 for the basic properties of these networks. (a) the Γ^{internet} network ($\mathcal{C} = 0.276$). (b) the $\Gamma^{\text{hepcoauth}}$ network ($\mathcal{C} = 0.213$). (c) the $\Gamma^{\text{astrocoauth}}$ network ($\mathcal{C} = 0.058$). (d) the Γ^{fb} network ($\mathcal{C} = 0.027$). (e) the Γ^{p2p} network ($\mathcal{C} = 0.190$). (f) the Γ^{protein} network ($\mathcal{C} = 0.326$). Simulation parameters: π given by Equation 4.8, $x_0 = 0.5$, $T = 10000$, $\nu = 100$, $\tau = 1$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

4.5 Snowdrift Game

Two friends John and Bill are stuck in a car on their way home because the road is blocked by a snowdrift. John and Bill have a choice between cooperation and defection, where cooperation means get out of the car and shovel and defection means stay in the car and do not shovel. Let the benefit of getting home be b and the cost of shoveling in the cold be c , where $b > c$. If both John and Bill cooperate, then the amount of work is only half as much, so both of them get a benefit of b and incur a cost of $\frac{c}{2}$, thus receiving a payoff of $b - \frac{c}{2}$. If both John and Bill defect, then they both remain stuck in the snowdrift, and receive a payoff of 0. If John cooperates and Bill defects, then they both get a benefit of b but John is the only one paying a cost of c . So the payoff to Bill is b while the payoff to John is $b - c$. Finally if Bill cooperates and John defects, then John receives a payoff of b while Bill receives $b - c$.

If we denote the act of cooperation by C and the act of defection by D , then we can define a symmetric, two-strategy, pairwise contest game, called the *snowdrift game*, with the payoff matrix

$$\pi = \begin{array}{c} \\ \begin{array}{cc} C & D \\ C & \left(b - \frac{c}{2}, b - \frac{c}{2} \right) \quad (b - c, b) \\ D & (b, b - c) \quad (0, 0) \end{array} \end{array}. \quad (4.11)$$

By Theorem 3.4, transforming the payoff matrix of a game using an affine transformation does not alter the equilibria of the game. So we can transform the above matrix by dividing each element by b to obtain the following payoff matrix:

$$\pi = \begin{array}{c} \\ \begin{array}{cc} C & D \\ C & \left(1 - \frac{c}{2b}, 1 - \frac{c}{2b} \right) \quad \left(1 - \frac{c}{b}, 1 \right) \\ D & \left(1, 1 - \frac{c}{b} \right) \quad (0, 0) \end{array} \end{array}. \quad (4.12)$$

Finally, denoting the cost-to-benefit ratio $\frac{c}{b}$ by ρ , we can express the payoff matrix for the game more compactly as

$$\pi = \begin{array}{cc} & \begin{array}{c} C \\ D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} (1 - \frac{\rho}{2}, 1 - \frac{\rho}{2}) & (1 - \rho, 1) \\ (1, 1 - \rho) & (0, 0) \end{pmatrix}, \end{array} \quad (4.13)$$

where $0 < \rho < 1$.

It is easy to see that a game with payoff matrix given by Equation 4.13 belongs to the chicken class of games we encountered in the previous chapter. Therefore, the set N of Nash equilibria for the game is $N = \{(C, D), (D, C), (\frac{1-\rho}{1-\frac{\rho}{2}}, \frac{\frac{\rho}{2}}{1-\frac{\rho}{2}})\}$ and the set E of ESSs for the game is $E = \{(\frac{1-\rho}{1-\frac{\rho}{2}}, \frac{\frac{\rho}{2}}{1-\frac{\rho}{2}})\}$. Comparing the payoff matrix for the game with the payoff matrix for a general symmetric, two-strategy, pairwise contest game given by Equation 4.1, we have $a = 1 - \frac{\rho}{2}$, $b = 1 - \rho$, $c = 1$, and $d = 0$. Substituting for a, b, c , and d in Equation 4.5 gives us the evolutionary dynamics (without assortativity) for the game as $\dot{x} = x(1-x)[x(\frac{\rho}{2} - 1) + 1 - \rho]$, where x is the fraction of individuals playing the C strategy and $1-x$ is the fraction of individuals playing the D strategy. The set F of fixed points for the dynamics is $F = \{(0, 1), (1, 0), (\frac{1-\rho}{1-\frac{\rho}{2}}, \frac{\frac{\rho}{2}}{1-\frac{\rho}{2}})\}$ and the set A of asymptotically stable fixed points is $A = \{(\frac{1-\rho}{1-\frac{\rho}{2}}, \frac{\frac{\rho}{2}}{1-\frac{\rho}{2}})\}$. The four sets N, E, F , and A are related by $E = A \subseteq N \subseteq F$. Therefore, a population of individuals playing the snowdrift game with payoff matrix given by Equation 4.13, starting from an initial fraction $x_0 > 0$ of cooperators, will eventually end up in the internal equilibrium state ($x^* = \frac{1-\rho}{1-\frac{\rho}{2}}$) in which cooperators and defectors coexist. So a natural question to ask is: how can the level of cooperation, i.e., the fraction of cooperators at equilibrium, be increased through an external mechanism?

Games in the chicken class, such as the hawk-dove game, have been studied extensively in the fields of political science, economics, and biology as models of conflict

between two players [Rus59, SP73, SF07]. The snowdrift game, which also belongs to the chicken class, has been studied as a model for cooperation by [HD04], in which the authors consider the effect of spatial structure on the snowdrift game and make the interesting observation that spatial structure can inhibit cooperation.

In what follows, we consider assortative interactions in a well-mixed population, characterized by the parameter $r \in [0, 1]$, as a possible mechanism for the promoting cooperation. We carry out a complete analysis of the snowdrift game using the replicator equation modified to include assortative interactions. Since the network assortativity in game interactions bears an inverse relationship to its average degree, knowing how the assortativity parameter r affects cooperation in a well-mixed population also informs us about how cooperation would be affected, albeit only qualitatively, when the game is played out on networks.

4.5.1 Analysis

Consider a population of individuals playing the snowdrift game with the payoff matrix given by Equation 4.13. Let $x(t)$ denote the fraction of cooperators at time t , and let $x(0) = x_0$. Let the degree of assortativity be r . Substituting for a, b, c , and d in Equation 4.4, we get the evolutionary dynamics for the population as

$$\dot{x} = x(1-x) \left[x(1-r) \left(\frac{\rho}{2} - 1 \right) + \frac{r\rho}{2} - \rho \right]. \quad (4.14)$$

The phase line plots for the above system, for different values of assortativity r , are shown in Figure 18. The system has three equilibria: $x^* = 0$ and $x^* = 1$ at the boundaries, and an internal equilibrium

$$x^* = \frac{\frac{r\rho}{2} + 1 - \rho}{(1-r) \left(1 - \frac{\rho}{2} \right)}, \quad (4.15)$$




Figure 18: Phase line plots for the system 4.14 for different values of assortativity r .

where $0 < \rho < 1$ and $x^* \in [0, 1] \implies 2r < \rho < \frac{2}{2-r}$. If a population starts out with a fraction x^* of individuals cooperating, then it remains that way. It is easy to verify that the boundary equilibria are unstable, while the internal equilibrium is stable. So if a population starts out with a fraction $x_0 \neq x^*$ of the individuals cooperating, then it will evolve to a state of coexistence determined by internal equilibrium, in which fraction x^* of the population are cooperators and the rest are defectors.

Since $\frac{\partial x^*}{\partial r} > 0$, increasing the assortativity r has the effect of increasing the fraction of cooperators at equilibrium. And since $\frac{\partial x^*}{\partial \rho} < 0$, increasing the cost-to-benefit ratio ρ has the effect of decreasing the fraction of cooperators at equilibrium.

Example 4.2. Consider the snowdrift game with $\rho = 0.5$, i.e., with payoff matrix

$$\pi = \begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} (0.75, 0.75) & (0.5, 1) \\ (1, 0.5) & (0, 0) \end{pmatrix}. \end{array} \quad (4.16)$$

We simulate the game on a complete network Γ_{WM} with parameter values $n = 10000$, $s = 1$, $x_0 = 0.5$, $\mathcal{U} = \text{FE2}$, $T = 5000$, $\nu = 1$, and $\tau = 1$. Figures 19(a)-(c) show the evolution of the fraction x of cooperators for different values of assortativity r . The fraction x of cooperators at equilibrium not only agrees with the predicted values, but also increases with increasing r , exactly as suggested by the analysis.




Figure 19: Evolution of the fraction x of cooperators in a symmetric, two-strategy, pairwise snowdrift game, for different values of assortativity r . (a) $r = 0$ ($x^* \approx 0.67$). (b) $r = 0.1$ ($x^* \approx 0.78$). (c) $r = 0.2$ ($x^* \approx 0.92$). Simulation parameters: complete network Γ_{WM} , $\rho = 0.5$, $n = 10000$, $x_0 = 0.5$, $\mathcal{U} = \text{FE2}$, $s = 1$, $T = 5000$, $\nu = 1$, and $\tau = 1$.

□

Thus in a well-mixed population, assortative interactions provide a mechanism for increasing the levels of cooperative behavior. Since network assortativity in game interactions bears an inverse relationship with the average degree of the network, when the snowdrift game is played on a network, we expect levels of cooperation to increase on networks with small values of d . We explore this and other network effects in detail in the following section.

4.5.2 Simulation Results

In this section we present the results of simulating the snowdrift game on networks having different structural properties. All the simulations were carried out using the following values for the parameters involved: payoff matrix π given by Equation 4.13, population size $n = 10000$, initial fraction of cooperators $x_0 = 0.5$, number of generations $T = 10000$, report frequency $\nu = 100$, number of trials $\tau = 1$, update rule $\mathcal{U} = \text{FE2}$, and selection

strength $s = 1$. See B.4 for details on how to simulate our discrete game models. The results are presented using three different kinds of plots depicting the following:

1. Variation of the long-term value x_∞ of the fraction x of cooperators, averaged over the last 10% generations, with assortativity r on complete networks (average degree d on other networks) and the cost-to-benefit ratio ρ . The values of $r \in [0, 0.5]$ and $\rho \in (0, 1)$ were varied in steps of 0.01, and the value of d was varied from 20 to 4 in steps of -2. The intensity of the shade in the plot indicates the level of cooperation, with black corresponding to full cooperation ($x_\infty = 1$) and white corresponding to no cooperation ($x_\infty = 0$).
2. Variation of x_∞ with r (or d) for a fixed value of ρ . The dashed line in the case of complete networks indicates values predicted by analysis.
3. Variation of x_∞ with ρ for a fixed value of r (or d). The dashed line in the case of complete networks indicates values predicted by analysis. We denote the fractional area under the x_∞ curve relative to the total area of the $[x_\infty, \rho]$ square by \mathcal{C} . We refer to \mathcal{C} as the *cooperation index* and use it as a measure of the extent to which a network promotes cooperation. For a well-mixed population, the value of $\mathcal{C}(r)$ for a given value of assortativity r can be calculated analytically as $\mathcal{C}(r) = \int_0^1 x_\infty(r, \rho) d\rho$, where $x_\infty(r, \rho)$ is given by Equation 4.15.

Complete Network

The first set of results concerns complete networks Γ_{WM} . Figures 20(a)(b) respectively show how the predicted and simulated values of the long-term fraction x_∞ of cooperators vary with assortativity r and cost-to-benefit ratio ρ . Figure 20(c) shows how x_∞ varies with r when $\rho = 0.75$. Figure 20(d) shows how x_∞ varies with ρ when $r = 0.25$.




Figure 20: Variation of the long-term fraction x_∞ of cooperators with assortativity r and cost-to-benefit ratio ρ , on a complete network. (a) x_∞ (predicted) versus r and ρ . (b) x_∞ (simulated) versus r and ρ . (c) x_∞ versus r when $\rho = 0.75$. (d) x_∞ versus ρ when $r = 0.25$ ($\mathcal{C} = 0.701$). Simulation parameters: π given by Equation 4.13, $n = 10000$, $x_0 = 0.5$, $T = 10000$, $\nu = 100$, $\tau = 1$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

Model Networks with Varying Average Degree

The next set of results concerns model networks with varying average degree d . Figures 21(a)(d)(g) show how the long-term fraction x_∞ of cooperators varies with average degree d and cost-to-benefit ratio ρ , on random regular Γ_{RR} , Erdős-Rényi Γ_{ER} , and

Barabási-Albert Γ_{BA} networks respectively. Figures 21(b)(e)(h) show how x_∞ varies with d on Γ_{RR} , Γ_{ER} , and Γ_{BA} networks respectively, when $\rho = 0.75$. Figures 21(c)(f)(i) show how x_∞ varies with ρ on Γ_{RR} , Γ_{ER} , and Γ_{BA} networks respectively, when $d = 4$.

Model Networks with Clustering

The next set of results concerns power-law networks $\Gamma_{\text{BA}}^{\text{C}}$ with average degree $d = 4$, and varying clustering coefficient $C^{(2)}$. Figures 22(a)-(d) show how the long-term fraction x_∞ of cooperators varies with the cost-to-benefit ratio ρ on networks with clustering coefficient $C^{(2)} = 0, 0.2, 0.4$, and 0.6 .

Model Networks with Homophily

The next set of results concerns power-law networks $\Gamma_{\text{BA}}^{\text{h}}$ with average degree $d = 4$, and varying homophily coefficient h . Figures 23(a)-(d) show how the long-term fraction x_∞ of cooperators varies with the cost-to-benefit ratio ρ on networks with homophily coefficient $h = -0.1, 0, 0.1$, and 0.2 .

Empirical Networks

The last set of results concerns empirical networks. Figures 24(a)-(f) show how the long-term fraction x_∞ of cooperators varies with the cost-to-benefit ratio ρ on: the Γ^{internet} network (a); the $\Gamma^{\text{hepcoauth}}$ network (b); the $\Gamma^{\text{astrocoauth}}$ network (c); the Γ^{fb} network (d); the Γ^{p2p} network (e); and the Γ^{protein} network (f).




Figure 21: Variation of the long-term fraction x_∞ of cooperators with average degree d and cost-to-benefit ratio ρ , on model networks. (a)-(c) Random regular network. Γ_{RR} . (d)-(f) Erdős-Rényi network Γ_{ER} . Barabási-Albert network Γ_{BA} . (a)(d)(g) x_∞ versus d and ρ . (b)(e)(h) x_∞ versus d when $\rho = 0.75$. (d)(f)(i) x_∞ versus ρ when $d = 4$ ($\mathcal{C} = 0.886, 0.814, 0.851$). Simulation parameters: π given by Equation 4.13, $n = 10000$, $x_0 = 0.5$, $T = 10000$, $v = 100$, $\tau = 1$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 22: Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on model power-law networks with clustering coefficient $C^{(2)}$. (a) $C^{(2)} = 0$ ($\mathcal{C} = 0.850$). (b) $C^{(2)} = 0.2$ ($\mathcal{C} = 0.852$). (c) $C^{(2)} = 0.4$ ($\mathcal{C} = 0.856$). (d) $C^{(2)} = 0.6$ ($\mathcal{C} = 0.859$). Simulation parameters: π given by Equation 4.13, $n = 10000$, $d = 4$, $x_0 = 0.5$, $T = 10000$, $\nu = 100$, $\tau = 1$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 23: Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on model power-law networks with homophily coefficient h . (a) $h = -0.1$ ($\mathcal{C} = 0.868$). (b) $h = 0$ ($\mathcal{C} = 0.836$). (c) $h = 0.1$ ($\mathcal{C} = 0.825$). (d) $h = 0.2$ ($\mathcal{C} = 0.871$). Simulation parameters: π given by Equation 4.13, $n = 10000$, $d = 4$, $x_0 = 0.5$, $T = 10000$, $\nu = 100$, $\tau = 1$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 24: Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on empirical networks. See Table 2 for the basic properties of these networks. (a) the Γ^{internet} network ($\mathcal{C} = 0.687$). (b) the $\Gamma^{\text{hepcoauth}}$ network ($\mathcal{C} = 0.734$). (c) the $\Gamma^{\text{astrocoauth}}$ network ($\mathcal{C} = 0.644$). (d) the Γ^{fb} network ($\mathcal{C} = 0.625$). (e) the Γ^{p2p} network ($\mathcal{C} = 0.658$). (f) the Γ^{protein} network ($\mathcal{C} = 0.810$). Simulation parameters: π given by Equation 4.13, $x_0 = 0.5$, $T = 10000$, $\nu = 100$, $\tau = 1$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

4.6 Sculling Game

Two friends John and Bill are rowing to get to their destination. Reaching their destination at all has a value of $\frac{b}{2}$ to each of them. If they reach their destination before a certain time—which will happen only if they both row—then they can attend a party they would like to go to which has an additional value $\frac{3b}{2}$ to each of them. It is assumed that if they do not reach their destination at all, then the benefit to both is 0. Furthermore, the effort of rowing carries a cost c to the rower. Thus, if both row they reach their destination on time and can go to the party and each gets a benefit $2b$. They each pay a cost c for rowing, giving a payoff for each of $2b - c$. If John rows and Bill does not then they will eventually reach their destination, but late for the party, giving them each a benefit of $\frac{b}{2}$. John will pay a cost c for rowing and so get a payoff of $\frac{b}{2} - c$. Bill will pay no cost and thus get a payoff of $\frac{b}{2}$. If neither rows, they drift in the river and do not reach their destination. Neither gets any benefit nor pays any cost for rowing, and thus each gets a payoff of 0.

If we treat rowing as the cooperative strategy and denote it by C , and treat not rowing as the defective strategy and denote it by D , then we can define a symmetric, two-strategy (C and D), pairwise contest game, called the *sculling game*⁴, with the payoff matrix

$$\pi = \begin{array}{cc} & \begin{array}{cc} C & D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \left(\begin{array}{cc} (2b - c, 2b - c) & (\frac{b}{2} - c, \frac{b}{2}) \\ (\frac{b}{2}, \frac{b}{2} - c) & (0, 0) \end{array} \right). \end{array} \quad (4.17)$$

By Theorem 3.4, transforming the payoff matrix of a game using an affine transformation does not alter the equilibria of the game. So we can transform the above

⁴Sculling generally refers to a method of using oars to propel watercraft in which the oar or oars touch the water on both the port and starboard sides of the craft, or over the stern. By extension, the oars themselves are also often referred to as sculls when used in this manner, and the boat itself may be referred to as a scull. [Source: Wikipedia]

matrix by dividing each element by b to obtain the following payoff matrix:

$$\pi = \begin{array}{cc} & \begin{array}{c} C \\ D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} (2 - \frac{c}{b}, 2 - \frac{c}{b}) & (\frac{1}{2} - \frac{c}{b}, \frac{1}{2}) \\ (\frac{1}{2}, \frac{1}{2} - \frac{c}{b}) & (0, 0) \end{pmatrix} \end{array}. \quad (4.18)$$

Finally, denoting the cost-to-benefit ratio $\frac{c}{b}$ by ρ , we can express payoff matrix for the game more compactly as

$$\pi = \begin{array}{cc} & \begin{array}{c} C \\ D \end{array} \\ \begin{array}{c} C \\ D \end{array} & \begin{pmatrix} (2 - \rho, 2 - \rho) & (\frac{1}{2} - \rho, \frac{1}{2}) \\ (\frac{1}{2}, \frac{1}{2} - \rho) & (0, 0) \end{pmatrix} \end{array}, \quad (4.19)$$

where $\frac{1}{2} < \rho < \frac{3}{2}$.

It is easy to see that a game with payoff matrix given by Equation 4.19 belongs to the coordination class of games we encountered in the previous chapter. Therefore, the set N of Nash equilibria for the game is $N = \{(C, C), (D, D), (\rho - \frac{1}{2}, \frac{3}{2} - \rho)\}$ and the set E of ESSs for the game is also $E = \{(C, C), (D, D), (\rho - \frac{1}{2}, \frac{3}{2} - \rho)\}$. Comparing the payoff matrix for the game with the payoff matrix for a general symmetric, two-strategy, pairwise contest game given by Equation 4.1, we have $a = 2 - \rho$, $b = \frac{1}{2} - \rho$, $c = \frac{1}{2}$, and $d = 0$. Substituting for a, b, c , and d in Equation 4.5 gives us the evolutionary dynamics (without assortativity) for the game as $\dot{x} = x(1-x)[x + \frac{1}{2} - \rho]$, where x is the fraction of individuals playing the C strategy and $1 - x$ is the fraction of individuals playing the D strategy. The set F of fixed points for the dynamics is $F = \{(0, 1), (1, 0), (\rho - \frac{1}{2}, \frac{3}{2} - \rho)\}$ and the set A of asymptotically stable fixed points is $A = \{(0, 1), (1, 0)\}$. The four sets N, E, F , and A are related by $E = A \subseteq N \subseteq F$. Since the internal equilibrium state ($x^* = \rho - \frac{1}{2}$) is unstable, a population of individuals playing the sculling game with payoff matrix given by Equation 4.19, starting from an initial fraction $x_0 > 0$ of cooperators, will eventually end up in the all-defector ($x = 0$) state if $x_0 < x^*$, and in the all-cooperator ($x = 1$) state if

$x_0 > x^*$. So a natural question to ask is: can the basin of attraction around the all-cooperator ($x = 1$) state be widened through an external mechanism such that the Hicks optimal outcome in which every individual cooperates will be favored?

Games in the coordination class, such as the stag-hunt game, have been studied extensively in social science, including economics, to model situations in which all parties can realize mutual gains, but only by making mutually consistent decisions [Sky03, SF07]. Coordination games, as far as we know, have not been used as models of cooperative behavior, and our formulation of the sculling game as a model for cooperation is new.

In what follows, we consider assortative interactions in a well-mixed population, characterized by the parameter $r \in [0, 1]$, as a possible mechanism for the promoting cooperation. We carry out a complete analysis of the sculling game using the replicator equation modified to include assortative interactions. Since the network assortativity bears an inverse relationship to its average degree, knowing how the assortativity parameter r affects cooperation in a well-mixed population also informs us about how cooperation would be affected, albeit only qualitatively, when the game is played out on networks.

4.6.1 Analysis

Consider a population of individuals playing the sculling game with the payoff matrix given by Equation 4.19. Let $x(t)$ denote the fraction of cooperators at time t , and let $x(0) = x_0$. Let the degree of assortativity be r . Comparing the payoff matrix for the game with the payoff matrix for a general symmetric, two-strategy, pairwise contest game given by Equation 4.1, we have $a = 2 - \rho$, $b = \frac{1}{2} - \rho$, $c = \frac{1}{2}$, and $d = 0$. Substituting for a, b, c , and d in Equation 4.4, we get the evolutionary dynamics for the population as

$$\dot{x} = x(1-x) \left[x(1-r) + \frac{3r}{2} - \rho + \frac{1}{2} \right]. \quad (4.20)$$

The phase line plots for the above system, for different values of assortativity r , are shown in Figure 25. The system has three equilibria: $x^* = 0$ and $x^* = 1$ at the




Figure 25: Phase line plots for the system 4.20 for different values of assortativity r .

boundaries, and the internal equilibrium

$$x^* = \frac{-\frac{3r}{2} + \rho - \frac{1}{2}}{1 - r}, \quad (4.21)$$

where $\frac{1}{2} < \rho < \frac{3}{2}$ and $x^* \in [0, 1] \implies \frac{1}{2} + \frac{3r}{2} < \rho < \frac{3}{2} + r$. If a population starts out with a fraction x^* of individuals cooperating, then it remains that way. It is easy to verify that the boundary equilibria are stable, while the internal equilibrium is unstable. So if a population starts out with a fraction $x_0 \neq x^*$ of the individuals cooperating, then it will evolve to the all-defector state if $x_0 < x^*$ and to the all-cooperator state if $x_0 > x^*$.

Since $\frac{\partial x^*}{\partial r} < 0$, increasing the assortativity r has the effect of increasing the basin of attraction around $x^* = 1$, thus decreasing the value of the initial fraction x_0 of cooperators that is needed for the population to evolve to the all-cooperator state. And since $\frac{\partial x^*}{\partial \rho} > 0$, increasing the cost-to-benefit ratio ρ has the effect of increasing the basin of attraction around $x^* = 0$, thus increasing the value of the initial fraction x_0 of cooperators that is need for the population to evolve to the all-cooperator state.

Example 4.3. Consider the sculling game with $\rho = 1$, i.e., with payoff matrix

$$\pi = \begin{array}{c} C \\ D \end{array} \begin{array}{cc} C & D \\ \begin{pmatrix} 1, 1 \\ 0.5, -0.5 \end{pmatrix} & \begin{pmatrix} -0.5, 0.5 \\ 0, 0 \end{pmatrix} \end{array}, \quad (4.22)$$

We simulate the game on a complete network Γ_{WM} with parameter values $n = 10000$, $s = 1$, $x_0 = 0.3$, $\mathcal{U} = \text{FE2}$, $T = 5000$, $\nu = 1$, and $\tau = 1$. Figures 26(a)-(c) show the evolution of the fraction x of cooperators for different values of assortativity r . When $r = 0$ and $r = 0.1$, the initial fraction x_0 of cooperators is smaller than the value of the internal equilibrium x^* , i.e., $x_0 < x^*$, so the population evolves to the all-defector state. However, when $r = 0.2$, $x_0 > x^*$, so the population evolves to the all-cooperator state.




Figure 26: Evolution of the fraction x of cooperators in a symmetric, two-strategy, pairwise sculling game, for different values of assortativity r . (a) $r = 0$. (b) $r = 0.1$. (c) $r = 0.2$. Simulation parameters: complete network Γ_{WM} , $\rho = 1$, $n = 10000$, $x_0 = 0.3$, $\mathcal{U} = \text{FE2}$, $s = 1$, $T = 5000$, $\nu = 1$, and $\tau = 1$.

□

Thus in a well-mixed population, assortative interactions provide a mechanism that favors the all-cooperator state by increasing the basin of attraction around the $x^* = 1$ equilibrium. Since network assortativity bears an inverse relationship with the average

degree of the network, when the sculling game is played on a network, we expect all-cooperator outcomes on networks with small values of d . We explore this and other network effects in detail in the following section.

4.6.2 Simulation Results

In this section we present the results of simulating the sculling game on networks having different structural properties. All the simulations, unless specified otherwise, were carried out using the following values for the parameters involved: payoff matrix π given by Equation 4.19, population size $n = 10000$, initial fraction of cooperators $x_0 = 0.5$, number of generations $T = 10000$, report frequency $\nu = 100$, number of trials $\tau = 5$, update rule $\mathcal{U} = \text{FE2}$, and selection strength $s = 1$. See B.4 for details on how to simulate our discrete game models. The results are presented using three different kinds of plots depicting the following:

1. Variation of the long-term value x_∞ of the fraction x of cooperators, averaged over the last 10% generations, with assortativity r on complete networks (average degree d on other networks) and the cost-to-benefit ratio ρ . The values of $r \in [0, 0.66]$ and $\rho \in (0.5, 1.5)$ were varied in steps of 0.01, and the value of d was varied from 20 to 4 in steps of -2. The intensity of the shade in the plot indicates the level of cooperation, with black corresponding to full cooperation ($x_\infty = 1$) and white corresponding to no cooperation ($x_\infty = 0$).
2. Variation of x_∞ with r (or d) for a fixed value of ρ . The dashed line in the case of complete networks indicates values predicted by analysis.
3. Variation of x_∞ with ρ for a fixed value of r (or d). The dashed line in the case of complete networks indicates values predicted by analysis. We denote the fractional

area under the x_∞ curve relative to the total area of the $[x_\infty, \rho]$ square by \mathcal{C} . We refer to \mathcal{C} as the *cooperation index* and use it as a measure of the extent to which a network promotes cooperation. For a well-mixed population, the value of $\mathcal{C}(r)$ for a given value of assortativity r can be calculated analytically as

$\mathcal{C}(r) = \int_{0.5}^{1.5} x_\infty(r, \rho) d\rho$, where $x_\infty(r, \rho) = 1$ if $x_0 > x^*$ given by Equation 4.21 and 0 otherwise.

Complete Network

The first set of results concerns complete networks Γ_{WM} . Figures 27(a)(b) respectively show how the predicted and simulated values of the long-term fraction x_∞ of cooperators vary with assortativity r and cost-to-benefit ratio ρ . Figure 27(c) shows how x_∞ varies with r when $\rho = 1.2$. Figure 27(d) shows how x_∞ varies with ρ when $r = 0.25$.

Model Networks with Varying Average Degree

The next set of results concerns model networks with varying average degree d . Figures 28(a)(d)(g) show how the long-term fraction x_∞ of cooperators varies with average degree d and cost-to-benefit ratio ρ , on random regular Γ_{RR} , Erdős-Rényi Γ_{ER} , and Barabási-Albert Γ_{BA} networks respectively. Figures 28(b)(e)(h) show how x_∞ varies with d on Γ_{RR} , Γ_{ER} , and Γ_{BA} networks respectively, when $\rho = 1.1$. Figures 28(c)(f)(i) show how x_∞ varies with ρ on Γ_{RR} , Γ_{ER} , and Γ_{BA} networks respectively, when $d = 4$.

Model Networks with Clustering

The next set of results concerns power-law networks $\Gamma_{\text{BA}}^{\text{C}}$ with average degree $d = 4$, and varying clustering coefficient $C^{(2)}$. Figures 29(a)-(d) show how the long-term fraction x_∞




Figure 27: Variation of the long-term fraction x_∞ of cooperators with assortativity r and cost-to-benefit ratio ρ , on a complete network. (a) x_∞ (predicted) versus r and ρ . (b) x_∞ (simulated) versus r and ρ . (c) x_∞ versus r when $\rho = 1.2$. (d) x_∞ versus ρ when $r = 0.25$ ($\mathcal{C} = 0.634$). Simulation parameters: π given by Equation 4.19, $n = 10000$, $x_0 = 0.5$, $T = 10000$, $\nu = 100$, $\tau = 5$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

of cooperators varies with the cost-to-benefit ratio ρ on networks with clustering coefficient $C^{(2)} = 0, 0.2, 0.4$, and 0.6 .




Figure 28: Variation of the long-term fraction x_∞ of cooperators with average degree d and cost-to-benefit ratio ρ , on model networks. (a)-(c) Random regular network Γ_{RR} . (d)-(f) Erdős-Rényi network Γ_{ER} . (g)-(i) Barabási-Albert network Γ_{BA} . (a)(d)(g) x_∞ versus d and ρ . (b)(e)(h) x_∞ versus d when $\rho = 1.1$. (d)(f)(i) x_∞ versus ρ when $d = 4$ ($\mathcal{C} = 0.819, 0.740, 0.771$). Simulation parameters: π given by Equation 4.19, $n = 10000$, $x_0 = 0.5$, $T = 10000$, $v = 100$, $\tau = 5$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 29: Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on model power-law networks with clustering coefficient $C^{(2)}$. (a) $C^{(2)} = 0$ ($\mathcal{C} = 0.769$). (b) $C^{(2)} = 0.2$ ($\mathcal{C} = 0.763$). (c) $C^{(2)} = 0.4$ ($\mathcal{C} = 0.767$). (d) $C^{(2)} = 0.6$ ($\mathcal{C} = 0.759$). Simulation parameters: π given by Equation 4.19, $n = 10000$, $d = 4$, $x_0 = 0.5$, $T = 10000$, $\nu = 100$, $\tau = 5$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

Model Networks with Homophily

The next set of results concerns power-law networks Γ_{BA}^h with average degree $d = 4$, and varying homophily coefficient h . Figures 30(a)-(d) show how the long-term fraction x_∞ of

cooperators varies with the cost-to-benefit ratio ρ on networks with homophily coefficient $h = -0.1, 0, 0.1, \text{ and } 0.2$.




Figure 30: Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on model power-law networks with homophily coefficient h . (a) $h = -0.1$ ($\mathcal{C} = 0.800$). (b) $h = 0$ ($\mathcal{C} = 0.742$). (c) $h = 0.1$ ($\mathcal{C} = 0.711$). (d) $h = 0.2$ ($\mathcal{C} = 0.807$). Simulation parameters: π given by Equation 4.19, $n = 10000$, $d = 4$, $x_0 = 0.5$, $T = 10000$, $v = 100$, $\tau = 5$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

Empirical Networks

The last set of results concerns empirical networks. Figures 31(a)-(f) show how the long-term fraction x_∞ of cooperators varies with the cost-to-benefit ratio ρ on: the Γ^{internet} network (a); the $\Gamma^{\text{hepcoauth}}$ network (b); the $\Gamma^{\text{astrocoauth}}$ network (c); the Γ^{fb} network (d); the Γ^{p2p} network (e); and the Γ^{protein} network (f).

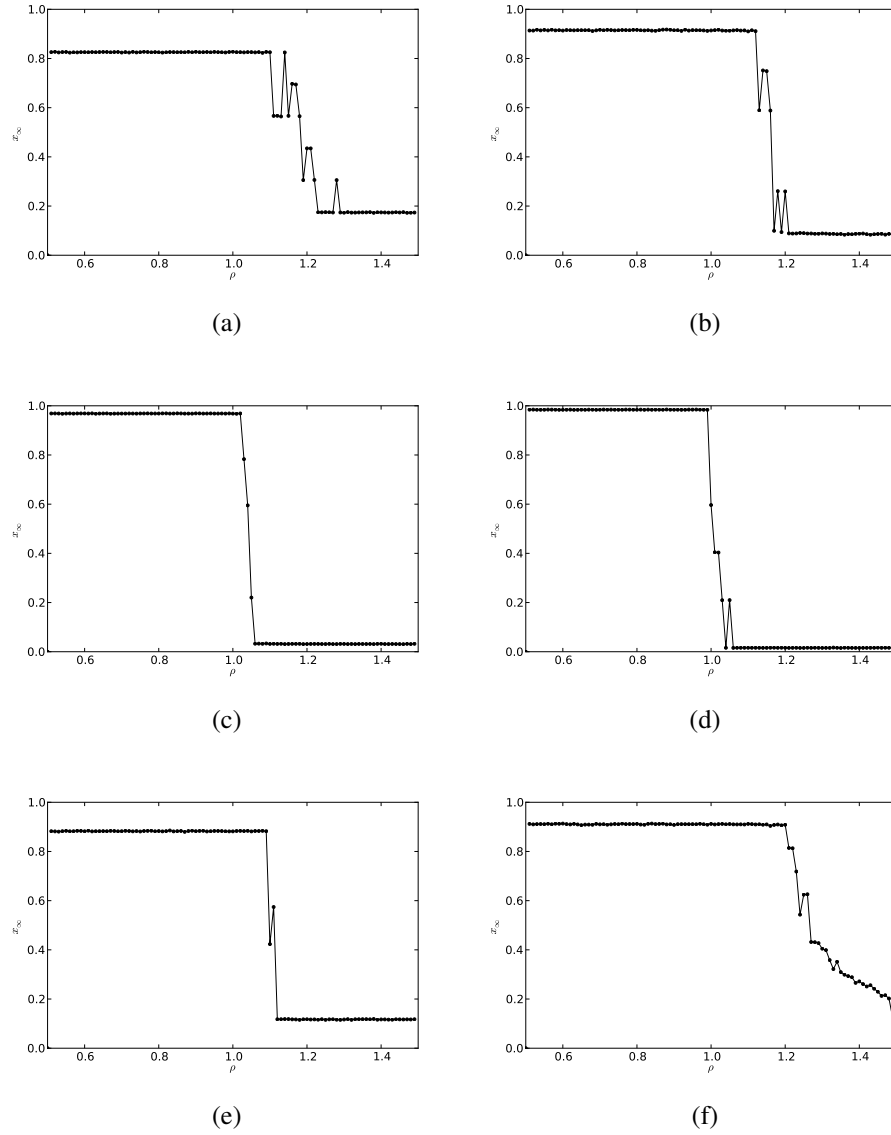


Figure 31: Variation of the long-term fraction x_∞ of cooperators with cost-to-benefit ratio ρ , on empirical networks. See Table 2 for the basic properties of these networks. (a) the Γ^{internet} network ($\mathcal{C} = 0.609$). (b) the $\Gamma^{\text{hepcoauth}}$ network ($\mathcal{C} = 0.626$). (c) the $\Gamma^{\text{astrocoauth}}$ network ($\mathcal{C} = 0.534$). (d) the Γ^{fb} network ($\mathcal{C} = 0.508$). (e) the Γ^{p2p} network ($\mathcal{C} = 0.575$). (f) the Γ^{protein} network ($\mathcal{C} = 0.747$). Simulation parameters: π given by Equation 4.19, $x_0 = 0.5$, $T = 10000$, $\nu = 100$, $\tau = 5$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

4.7 Discussion

For each of the three (prisoner's dilemma, snowdrift, and sculling) games considered, the simulation results (see Figures 13(a)-(d), 20(a)-(d), 27(a)-(d)) on complete networks Γ_{WM} agree very well with the values predicted by analysis on well-mixed populations. In case of the prisoner's dilemma game, a population, starting from an initial fraction $x_0 > 0$ of cooperators, eventually evolves to the all-cooperator ($x^* = 1$) equilibrium state when the inequality $r > \rho$ is satisfied, where r is the degree of assortative interactions and ρ is the cost-to-benefit ratio. When $r < \rho$, population evolves to the all-defector ($x^* = 0$) equilibrium state. In case of the snowdrift game, as predicted by the analysis, an increase in r results in an increase in the internal equilibrium state x^* , and an increase in ρ results in a decrease in x^* . In case of the sculling game, again as predicted by analysis, as r increases, the basin of attraction around the all-cooperator ($x^* = 1$) equilibrium state widens and beyond a certain stage the initial fraction $x_0 > 0$ of cooperators lies within the basin, and the population evolves to the all-cooperator ($x^* = 1$) equilibrium state. On the other hand, as ρ increases, the basin of attraction around the all-defector ($x^* = 0$) equilibrium state widens and beyond a certain stage the initial fraction $x_0 > 0$ of cooperators lies within the basin, and the population evolves to the all-defector ($x^* = 0$) equilibrium state.

The simulation results (see Figures 14(a)-(i), 21(a)-(i), 28(a)-(i)) on model (random regular, Erdős-Rényi, Barabási-Albert) networks with varying average degree d suggest that the long-term fraction x_∞ of cooperators increases as d decreases. This makes sense because as the average degree of a network decreases, the network assortativity in game interactions increases, and from the results on complete networks we know that cooperation increases with increasing assortativity. For d as small as 20, the networks yield results similar to complete networks with zero assortativity. This again makes sense

because as the average degree of a network increases, it more and more approximates a complete network. The figures also suggest that x_∞ decreases with increasing cost-to-benefit ratio ρ , similar to what we witnessed on complete networks. Furthermore, except for small values of d , the degree distribution of the networks has very little effect on the value of x_∞ ; when $d = 4$, the random regular network has the highest values for the cooperation index \mathcal{C} for all three games.

In case of the prisoner's dilemma game on model networks with average degree d , our results agree well with what is suggested by network reciprocity [Now06b], i.e., the networks promote cooperation when $\frac{1}{d} > \rho$. In case of the snowdrift game, our results on random regular, Erdős-Rényi, and Barabási-Albert networks suggest that network structure promote cooperation. This is in contradiction with the observation made by Doebeli and Hauert in [HD04], in which the authors note that spatial structure inhibits cooperation. The resolution of these observations could be that the argument of Doebeli and Hauert assumes that the network is a lattice, and it probably only holds (at most) for a planar network.

A comparison of the cooperation indices \mathcal{C} from the simulation results (see Figures 15(a)-(d), 22(a)-(d), 29(a)-(d)) on model networks with clustering coefficient $C^{(2)}$ suggests that in case of prisoner's dilemma and sculling games, an increase in clustering inhibits cooperation, while in case of snowdrift game, an increase in clustering promotes cooperation.

A comparison of the cooperation indices \mathcal{C} from the simulation results (see Figures 16(a)-(d), 23(a)-(d), 30(a)-(d)) on model networks with homophily coefficient h suggests that in case of all three games, networks that are neither assortative nor disassortative, i.e., $h = 0$, inhibit cooperation, while assortative (positive or negative) networks promote cooperation.

A comparison of the cooperation indices \mathcal{C} from the simulation results (see Figures 17(a)-(f), 24(a)-(f), 31(a)-(f)) on empirical networks further support the fact that networks with small values of average degree d promote cooperation, while networks with high values of d inhibit cooperation. For example, the Γ^{protein} network, which has the smallest average degree ($d \approx 3.56$), has the highest values for \mathcal{C} , while the Γ^{fb} network, which has the highest average degree ($d \approx 46.51$), has the smallest values for \mathcal{C} . The six empirical networks have very different sizes, and yet the trend of x_∞ with respect to their average degree d is comparable to the trend we observed in the case of model networks, which were all of size 10000. This suggests that the results of our models, as long as all other parameters are fixed, is independent of the population size.

Looking at the simulation results on Erdős-Rényi networks Γ_{ER} having small average degree d and on the six empirical networks, we notice that the two strategies C and D never fixated when the prisoner's dilemma, snowdrift, and sculling games were played on these networks. In other words, the long-term fraction x_∞ of cooperators on these networks was strictly between 0 and 1. The reason for this is that the networks consist of vertices with degree less than two, and the update rule $\mathcal{U} = \text{FE2}$ that we used in our simulations excludes such vertices from the interaction and update rounds.

While the game dynamics on complete networks is independent of the update rule \mathcal{U} , it generally does depend on \mathcal{U} when the games are played on other networks. To illustrate this, we simulated the prisoner's dilemma, snowdrift, and sculling games on a complete network Γ_{WM} and a Barabási-Albert Γ_{BA} network, using five different update rules $\mathcal{U} = \text{FE1}, \text{FE2}, \text{RE1}, \text{RE2},$ and BD . We were interested in how the variation of the long-term fraction x_∞ of cooperators with the cost-to-benefit ratio ρ is affected by the choice of update rule.

	FE1	FE2	RE1	RE2	BD
Prisoner's dilemma game	0.250	0.252	0.241	0.248	0.240
Snowdrift game	0.649	0.701	0.549	0.671	0.501
Sculling game	0.708	0.634	0.835	0.628	0.899

Table 3: Effect of update rule on discrete-game dynamics on a complete network. The value for each (game, update rule) pair is the cooperation index \mathcal{C} . Simulation parameters: π given by Equations 4.8, 4.13, and 4.19 respectively; $n = 10000$; $r = 0.25$; $x_0 = 0.5$; $r = 0.25$; $T = 10000$; $v = 100$; $\tau = 1$; $s = 1$; and ρ was varied within the allowed range in steps of 0.01.

Table 3 shows the cooperation indices \mathcal{C} for each of the (game, update rule) pairs when the games are played on a Γ_{WM} network, and Table 4 shows the same when the games are played on a Γ_{BA} network. On a Γ_{WM} network, the game dynamics is relatively independent of the choice of update rule. However, on a Γ_{BA} network, the dynamics of the prisoner's dilemma game is strongly affected by the update rule that is used, where only the FE2, RE2, and BD rules seem to promote any cooperation.

4.8 Future Directions

In this section we propose two new lines of inquiry, along with an outline of a possible approach for each.

1. The symmetric, two-strategy, pairwise games we considered in this chapter were all analyzed for a well-mixed population, using the deterministic framework of replicator dynamics, in which the abundance of individuals in a population is a continuous variable. It is assumed in this framework that the population size is

	FE1	FE2	RE1	RE2	BD
Prisoner's dilemma game	0.022	0.267	0.030	0.278	0.176
Snowdrift game	0.630	0.851	0.509	0.751	0.616
Sculling game	0.551	0.771	0.623	0.844	0.723

Table 4: Effect of update rule on discrete-game dynamics on a Barabási-Albert network. The value for each (game, update rule) pair is the cooperation index \mathcal{C} . Simulation parameters: π given by Equations 4.8, 4.13, and 4.19 respectively; $n = 10000$; $d = 4$; $x_0 = 0.5$; $T = 10000$; $v = 100$; $\tau = 1$; $s = 1$; and ρ was varied within the allowed range in steps of 0.01.

infinite. Simulation results on large complete networks, which model an infinitely large well-mixed population, agree well with the analysis. Such agreement, however, is not to be expected when the size of the network is small. We would like to know the effects of small population sizes on the evolutionary dynamics of symmetric, two-strategy, pairwise games with assortative interactions. In this case, the dynamics is no longer given by differential equations, but require a stochastic formulation [NST04] as highlighted below.

Consider a population of fixed size n , with two types of individuals: cooperators C and defectors D . At each time step t , a random individual is chosen for reproduction and another random individual is chosen for elimination, using sampling with replacement. The only stochastic variable is the number of C individuals, which is denoted by i .

The Moran process is defined on the state space $i = 0, \dots, n$. The $(N + 1) \times (N + 1)$ transition matrix $M = [M_{ij}]$ is a tridiagonal, stochastic matrix that

determines the probability $p_{i,j}$ of moving from any state i to any other state j . There are four possibilities for picking the individuals in each time step:

(C,C) , (D,D) , (C,D) , and (D,D) , which define the probabilities $p_{i,i}$, $p_{i,i-1}$, and $p_{i,i+1}$, with $p_{i,i} = 1 - p_{i,i-1} - p_{i,i+1}$. We consider frequency-dependent selection, so these probabilities are expressed in terms of the elements in the payoff matrix π for a game, the degree of assortative interactions r , and the fitnesses of the strategies C and D .

The states $i = 0$ and $i = n$ are “absorbing states”. Once the process has reached such a state, it will stay there forever. The state $i = 1, \dots, n - 1$ are called transient states. Since the stochastic process has two absorbing states, we can ask: starting in state i , what is the probability of reaching state n ? In other words, given that we start with i number of C individuals, what is the probability that eventually the whole population will consist of only C individuals?

We denote by x_i the probability of reaching state n starting from state i . Clearly, $1 - x_i$ denotes the probability of reaching state 0 starting from state i . In vector notation, the stochastic dynamics is given by

$$\mathbf{x}_{t+1} = M\mathbf{x}_t, \tag{4.23}$$

where $\mathbf{x}_t = \{x_0, x_1, \dots, x_n\}$ denotes the absorption probabilities at time t . The absorption probabilities are given by the right-hand eigenvectors associated with the largest eigenvalue of M , which is 1 since it is a stochastic matrix.

Let ρ_C denote the probability that a single C individual will take over the whole population. We call this the fixation probability of C . Similarly, we can define the fixation probability for D as the probability that a single D individual will take over the entire population. If $\frac{\rho_D}{\rho_C} > 1$, then it is more likely that a single D individual will

take over and become fixed, and if the inequality is reversed, it is more likely that a single C individual will become fixed.

We can also consider mutations of the resident individuals, in which each resident mutates to the other type with probability μ , and stays the same type with probability $1 - \mu$. We would be interested in the relative stability of a population of cooperators when it is invaded by a mutant defector. We would also be interested in whether a mutant cooperator can take over a population of defectors.

The stochastic dynamics described above is easily implemented in software, which will help us confirm the results of the stochastic analysis.

2. As noted in Section 4.7, the dynamics of games on complete networks, which model well-mixed populations, is unaffected by the choice of update rules. This is, however, not the case on other networks, in which case the dynamics sensitively depends on the update rule that is used. Furthermore, the structure of the network also affects the dynamics of the game in significant ways. So any attempt at analyzing the evolutionary dynamics of games on networks must take into account both the structure of the networks and also the details of the update rule that is employed [ON06b].

In our second line of inquiry, we would like to analyze symmetric, two-strategy, pairwise games on networks with average degree d . We could follow the program highlighted in [ON06b] and derive a system of ordinary differential equations that describe how the expected frequency x_i of each strategy in a game on a network changes over time. The program uses pair approximation on regular graphs of degree d [Ell01, MOS92, OHL06, ON06a], and works as follows.

We introduce the $n \times n$ matrix $B = [b_{ij}]$ that incorporates the details of the update rule \mathcal{U} under consideration. We further introduce the quantities

$$g_i = \sum_{j=1}^n x_j b_{ij}, \quad (4.24)$$

which characterize the local competition among the strategies, such that $\sum_{i=1}^n x_i g_i = 0$.

In the limit of weak selection, we derive the replicator equation for the expected frequency $x_i(t)$ of strategy i on an infinitely large graph of degree d at time t as

$$\dot{x}_i = x_i(f_i + g_i - \phi) \quad i = 1, \dots, n \quad (4.25)$$

where f_i denotes the fitness of strategy i , and ϕ is the average fitness of the population given by $\phi = \sum_{i=1}^n x_i(f_i + g_i) = \sum_{i=1}^n x_i f_i$, since $\sum_{i=1}^n x_i g_i = 0$.

Equation 4.25, by a transformation of the payoff matrix $\pi = [a_{ij}] \rightarrow [a_{ij} + b_{ij}]$ for the game, can also be written as

$$\dot{x}_i = x_i \left[\sum_{j=1}^n x_j (a_{ij} + b_{ij}) - \phi \right], \quad (4.26)$$

which is called the replicator equation on networks [ON06b]. It is well-known that pair approximations give good results for random regular networks as n increases and the probability of short loops becomes negligible. Based on the results of simulating the games on different networks, we expect the results to also hold well for the Erdős-Rényi and Barabási-Albert networks with average degree d .

CHAPTER 5

CONTINUOUS GAMES

Seeing this gradation and diversity of structure in one small, intimately related group of birds, one might really fancy that from an original paucity of birds in this archipelago, one species had been taken and modified for different ends.

- Charles Darwin

5.1 Introduction

Cooperative behavior is not always discrete in nature. For example, when vampire bats share blood with one another [Wil90], or when individuals share files over a peer-to-peer network [RF02, SR04], they exhibit varying levels of cooperative behavior. Besides not being able to take degrees of cooperative behavior into account, discrete game models suffer from another problem in that they consider the invasiveness and stability of fully developed, highly cooperative interactions, despite the fact that the gradual evolution of cooperation from an initially selfish state represents a more plausible evolutionary scenario. A framework which allows different degrees of cooperation is thus more natural. Once variable levels of cooperation are considered it becomes possible to study the crucial issue of how cooperation can evolve gradually from an initial state consisting of non-cooperative entities.

Situations in which individuals exhibit varying degrees of cooperative behavior can be described using continuous games. The level of cooperation—called an *investment*—of an individual is that individual's strategy in the game. An investment is represented as a continuous variable; the cost and benefit associated with an investment are represented as

continuous functions of the investment; and the payoff to an individual is simply the difference between the benefit and cost, and thus is also a continuous function.

We consider three continuous games: the continuous prisoner's dilemma game in which a cooperative investment made by an individual (donor) towards another individual (recipient) benefits the recipient but costs the donor; the continuous snowdrift game in which the cooperative investment benefits both the donor and recipient but costs the donor; and the continuous tragedy of the commons game in which the cooperative investment¹ benefits the donor but costs both the donor and recipient.

We consider, as we did in the previous chapter on discrete games, a population of individuals playing the continuous games through repeated pairwise interactions. In a well-mixed population, we again assume the interactions to be assortative [Gra90], i.e., each individual interacts with other individuals of their own "type" with a certain probability, and with random individuals from the population otherwise.

In each of the three games, the individuals playing the game have an incentive to be selfish and not make any cooperative investments. As a result, everyone does worse than if they had all cooperated. We are interested in conditions that would lead to higher levels of cooperative behavior among selfish individuals [Kol98]. Assortative interactions (both in well-mixed populations and on networks), as we shall see, provide a mechanism for the promotion and maintenance of cooperative investments, and thus lead to Hicks (or socially) optimal outcomes. In addition network properties such as clustering and homophily also affect cooperative investments.

We propose agent-based models for emulating on a computer how individuals in a population might play the three continuous games. We analyze the games for a well-mixed population for general cost and benefit functions, using the framework of

¹In this context investment means the level of consumption of a limited common-pool resource, and cooperative investments mean lower levels of consumption.

adaptive dynamics [GKM97, MB07], suitably modified to incorporate assortative interactions. We also consider a few specific cost and benefit functions for which the analysis is tractable. We present results of simulating the games on model and empirical networks and provide a general discussion of the results. We conclude the chapter with suggestions for future inquiry.

5.2 Assortative Interactions

Here again, as in the case of discrete games, we consider assortative interactions [Gra90]. In a well-mixed population we quantify the extent to which an individual interacts with another individual of its own type (strategy in the context of games) using a parameter $r \in (0, 1)$ called the *degree of assortativity* or simply *assortativity*, which is interpreted the same way as it was in the case of discrete games. We proceed to develop the framework of adaptive dynamics with assortative interactions.

Consider a well-mixed population of individuals playing a continuous game. Let the interactions be pairwise with $\pi(x, y)$ denoting the payoff to the x -strategist when it plays the game with a y -strategist. Let r denote the degree of assortativity. Let the population be monomorphic at the outset, i.e., all the individuals play the same (resident) strategy, say x . Let us now consider a small fraction ε of mutant individuals in the population playing strategy y . The fitnesses f_x^r and f_y^r of the x and y strategies, using Equation 3.72, are given by

$$f_x^r = r\pi(x, x) + (1 - r)[(1 - \varepsilon)\pi(x, x) + \varepsilon\pi(x, y)], \quad (5.1)$$

and

$$f_y^r = r\pi(y, y) + (1 - r)[(1 - \varepsilon)\pi(y, x) + \varepsilon\pi(y, y)]. \quad (5.2)$$

Since there are just two strategies—the resident x and the mutant y , from Equation 3.72, the invasion fitness $f_x^r(y)$ of the mutant strategy y , i.e., the fitness of y relative to the average fitness of the population is given by²

$$\begin{aligned} f_x^r(y) &= \lim_{\varepsilon \rightarrow 0} f_y^r - f_x^r \\ &= r\pi(y, y) + (1 - r)\pi(y, x) - \pi(x, x). \end{aligned} \quad (5.3)$$

Note that when $r = 0$, the above equation reduces to

$$f_x(y) = \pi(y, x) - \pi(x, x), \quad (5.4)$$

which is the standard equation for the invasion fitness of a mutant y -strategist in a population of predominantly x -strategists.

5.3 Agent-Based Model

An agent-based model [Bon02] provides a natural way of describing how individuals in a population might play a pairwise continuous game with payoff function π . One can simulate (see section B.4) such a model on a computer in order to study the evolution of the strategies.

We start with a population network Γ of size n with a strategy profile $\hat{x}^0 = \{x_1^0, x_2^0, \dots, x_n^0\}$ at generation 0, where $x_1^0 = x_2^0 = \dots = x_n^0 = x_0 \in (0, x_m)$ and $x_m \in \mathbb{R}_{>0}$. Thus at generation $t = 0$, we have a monomorphic population with every individual playing the strategy x_0 . The payoff function for the pairwise, continuous game is $\pi(x, y)$, the update rule \mathcal{U} for the game is one of the five rules (FE, RE, IM, BD, or DB) we discussed in the Chapter 3, the selection strength is s , and the degree of assortativity (applicable only if $\Gamma = \Gamma_{\text{WM}}$, i.e., if Γ is a complete network) is r , the rate μ of Gaussian mutations of strategies, and the standard deviation σ for the Gaussian mutations.

²Note that the f 's on the two sides of Equation 5.3 are different.

The game is played for T generations, each involving n interactions and updates. During an interaction, say at generation t , an individual i plays the game with a neighboring individual $j \in N(i)$ and obtains a payoff π_i . On a complete network Γ_{WM} , the interaction is assortative, i.e., with probability r , the individual i plays the game with another individual j of its own type, and with probability $1 - r$, it plays the game with a randomly picked individual j . The two individuals employ strategies x_i^{t-1} and x_j^{t-1} from the previous generation. During an update at generation t , an individual i updates its strategy x^t in the generation t with strategy x_j^{t-1} of its neighbor $j \in N(i)$, with a possible mutation of the strategy, picked from a Gaussian distribution with mean x_j^{t-1} and standard deviation σ . The updates we consider are thus innovative.

Just as in the case of the discrete games, we have two distinct algorithms for implementing the games on a computer. In the *sequential* approach, the n update rounds follow the n interaction rounds for a given generation. So the pair of individuals (i, j) considered for an interaction is generally not the same as the pair (i, j) of individuals considered for an update. Figure 32 shows a schematic of interactions and updates in the sequential approach. The box on the top depicts n rounds of interactions while the box below depicts n rounds of updates. A circled vertex represents the focal individual, i.e., the individual receiving a payoff during the interaction round, or the individual inheriting a strategy during the update round.

Algorithm 5.4 provides the pseudo-code for implementing pairwise continuous games on a computer using sequential interactions and updates. In this case the update rule can be one of FE, RE, IM, BD, or DB. Each generation t consists of n interaction rounds followed by n update rounds (see Figure 9). In each interaction round, an individual i is randomly sampled from Γ . If Γ is a complete network Γ_{WM} , then i plays the game with an




Figure 32: Sequential interactions and updates in continuous games.

individual of their own strategy type³ with probability r and receives a payoff $\pi_i = \pi(x_i^{t-1}, x_i^{t-1})$. Otherwise, i plays the game with a randomly sampled individual $j \in N(i)$, and receives a payoff $\pi_i = \pi(x_i^{t-1}, x_j^{t-1})$. At the end of the n interactions, roughly every individual in the population would have played the game and thus received a payoff.

Two individuals i and $j \in N(i)$ are selected based on the update rule \mathcal{U} (see Chapter 3) and i either ends up inheriting j 's (possibly mutated) strategy from generation $t - 1$ or retaining their own strategy.

³We can either assume that such an individual always exists in the population, or that i actively seeks out such an individual if it exists, and forgoes playing the game at generation t otherwise. Both these assumptions result in the same dynamics up to rate of evolution of strategies. Our software implementation makes the former assumption.

At the start of each generation t , the strategy profile \hat{x}^{t-1} from the previous generation is saved in a list L of strategy profiles, and at the end of T generations the list L is returned as output.

Algorithm 5.4 Continuous games on complex networks, using sequential interactions and updates.

Input: A population network Γ of size n with strategy profile $\hat{x}^0 = \{x_1^0, x_2^0, \dots, x_n^0\}$ at generation 0, where $x_1^0 = x_2^0 = \dots = x_n^0 = x_0 \in (0, x_m)$ and $x_m \in \mathbb{R}_{>0}$; a continuous payoff function $\pi(x, y)$; update rule \mathcal{U} (FE, RE, IM, BD, DB); selection strength s ; probability r of assortative interactions; rate μ and standard deviation σ for Gaussian mutations of strategies; and number of generations T .

Output: A list $L = \{\hat{x}^0, \hat{x}^1, \dots, \hat{x}^{T-1}\}$.

```

1:  $L \leftarrow []$ 
2: for  $t$  from 1 to  $T$  do
3:    $L.append(\hat{x}^{t-1})$ 
4:   // Interaction rounds.
5:   for  $count$  from 1 to  $n$  do
6:      $i \leftarrow$  random individual from  $\Gamma$ 
7:     if  $\Gamma = \Gamma_{WM}$  and  $random() < r$  then
8:        $\pi_i \leftarrow \pi(x_i^{t-1}, x_i^{t-1})$ 
9:     else
10:       $j \leftarrow$  random individual from  $N(i)$ 
11:       $\pi_i \leftarrow \pi(x_i^{t-1}, x_j^{t-1})$ 
12:    end if
13:  end for
14: end for
15: // Update rounds.

```

```

16: for count from 1 to n do
17:   Two individuals  $i$  and  $j \in N(i)$  are selected based on the update rule  $\mathcal{U}$  (see
      Chapter 3) and  $i$  either inherits  $j$ 's (possibly mutated) strategy from generation  $t - 1$ 
      or retains their own strategy
18: end for
19: return  $L$ 

```

In the *simultaneous* approach, at each generation, each of the n updates immediately follows two interactions, one in which an individual i interacts with another individual $k \in N(i)$ and obtains a payoff π_i and the other in which an individual $j \in N(i)$ ($j \neq k$) interacts with another individual $l \in N(j)$ ($l \neq i$) and obtains a payoff π_j . The same pair (i, j) is then considered for an update. Figure 33 shows a schematic of the interactions and updates in the simultaneous approach. A circled vertex again represents a focal individual, i.e., the individual receiving a payoff during the interaction round, or the individual inheriting a strategy during the update round.

Algorithm 5.5 provides the pseudo-code for implementing pairwise continuous games on a computer using simultaneous interactions and updates (see Figure 10). In this case the update rule \mathcal{U} can either be FE or RE. Each generation t consists of n rounds, each of which involves an interaction immediately followed by an update. In each round, an individual i and an individual $j \in N(i)$ are randomly sampled from Γ . If Γ is a complete network Γ_{WM} , then i plays the game with an individual of its own strategy type with probability r and receives a payoff $\pi_i = \pi(x_i^{t-1}, x_i^{t-1})$. Otherwise, i plays the game with a randomly sampled individual $k \in N(i)$, and receives a payoff $\pi_i = \pi(x_i^{t-1}, x_k^{t-1})$. Similarly, j plays the game with an individual of its own strategy type with probability r and receives a payoff $\pi_j = \pi(x_j^{t-1}, x_j^{t-1})$. Otherwise, j plays the game with a randomly sampled individual $l \in N(j)$, and receives a payoff $\pi_j = \pi(x_j^{t-1}, x_l^{t-1})$. Note that on




Figure 33: Simultaneous interactions and updates in continuous games.

$\Gamma \neq \Gamma_{\text{WM}}$, i must be distinct from l and j must be distinct from k . Otherwise, the process continues with the next generation. The individual i updates its strategy x_i^t to the (possibly mutated) strategy x_j^{t-1} of a randomly sampled individual $j \in N(i)$, with probability $p_{i \leftarrow j}$ determined by the update rule, and with probability $1 - p_{i \leftarrow j}$, i retains its own strategy.

At the start of each generation t , the strategy profile \hat{x}^{t-1} from the previous generation is saved in a list L of strategy profiles, and at the end of T generations the list L is returned as output.

Algorithm 5.5 Continuous games on complex networks, using simultaneous interactions and updates.

Input: A population network Γ of size n with strategy profile $\hat{x}^0 = \{x_1^0, x_2^0, \dots, x_n^0\}$ at generation 0, where $x_1^0 = x_2^0 = \dots = x_n^0 = x_0 \in (0, x_m)$ and $x_m \in \mathbb{R}_{>0}$; a continuous payoff function $\pi(x, y)$; update rule \mathcal{U} (FE or RE); selection strength s ; probability r of assortative interactions; rate μ and standard deviation σ for Gaussian mutations of strategies; and number of generations T .

Output: A list $L = \{\hat{x}^0, \hat{x}^1, \dots, \hat{x}^{T-1}\}$.

```

1:  $L \leftarrow []$ 
2: for  $t$  from 1 to  $T$  do
3:    $L.append(\hat{x}^{t-1})$ 
4:   // Interaction and update rounds.
5:   for  $count$  from 1 to  $n$  do
6:      $i \leftarrow$  random individual from  $\Gamma$ 
7:      $j \leftarrow$  random individual from  $N(i)$ 
8:     if  $\Gamma = \Gamma_{WM}$  and  $random() < r$  then
9:        $\pi_i \leftarrow \pi(x_i^{t-1}, x_i^{t-1})$ 
10:    else
11:       $k \leftarrow$  random individual from  $N(i)$ 
12:      if  $\Gamma \neq \Gamma_{WM}$  and  $k = j$  then
13:        continue
14:      end if
15:       $\pi_i \leftarrow \pi(x_i^{t-1}, x_k^{t-1})$ 
16:    end if
17:    if  $\Gamma = \Gamma_{WM}$  and  $random() < r$  then
18:       $\pi_j \leftarrow \pi(x_j^{t-1}, x_j^{t-1})$ 
19:    else
20:       $l \leftarrow$  random individual from  $N(j)$ 

```

```

21:     if  $\Gamma \neq \Gamma_{\text{WM}}$  and  $l = i$  then
22:         continue
23:     end if
24:      $\pi_j \leftarrow \pi(x_j^{t-1}, x_l^{t-1})$ 
25: end if
26: The probability  $p_{i \leftarrow j}$  that individual  $i$  inherits individual  $j$ 's strategy is calculated
    based on the update rule  $\mathcal{U}$  (see Chapter 3)
27: if  $\text{random}() < p_{i \leftarrow j}$  then
28:      $x_i^t \leftarrow x_j^{t-1}$ 
29:     if  $\text{random}() < \mu$  then
30:          $x_i^t \leftarrow \max(0, \min(\text{gaussian}(x_i^t, \sigma), x_m))$ 
31:     end if
32: else
33:      $x_i^t \leftarrow x_i^{t-1}$ 
34: end if
35: end for
36: end for
37: return  $L$ 

```

5.4 Prisoner's Dilemma

Two friends John and Bill are software developers. In addition to their day jobs, they also invest time into open-source software development. The “free”⁴ software that John develops is extensively used by Bill but not by John himself, and the one that Bill develops is heavily used by John but not by Bill himself. Such interactions among pairs of individuals in which the cooperative investment made by an individual (donor) benefits the other individual (recipient) but is costly to the donor, can be described using a pairwise, continuous prisoner's dilemma (CPD) game [IKD04, KDK99, KD02].

The CPD game consists of two individuals, each making an investment $x \in \mathbb{R}_{\geq 0}$. The investment x has the following effects: the payoff of the investor (donor) is reduced by $C(x)$, where C is a function that specifies the cost of making the investment; and the payoff of the beneficiary (recipient) is increased by $B(x)$, where B is a function that specifies the benefit resulting from the investment. Therefore, if two interacting individuals make investments x and y , then the payoff $\pi(x, y)$ and $\pi(y, x)$ to the individuals are respectively given by

$$\begin{aligned}\pi(x, y) &= B(y) - C(x), \text{ and} \\ \pi(y, x) &= B(x) - C(y).\end{aligned}\tag{5.5}$$

We assume that the cost and benefit functions are such that for investments between 0 and an upper limit x_m , $B(x) > C(x)$. This assumption is a necessary condition for cooperation to evolve; otherwise, if every individual invests x , the payoff of every individual is lower if $x > 0$ than if $x = 0$.

For any positive investments $x_1 < x_2 < x_m$, where x_m is the investment that maximizes $B(x) - C(x)$, we have $\pi(x_1, x_2) > \pi(x_2, x_2) > \pi(x_1, x_1) > \pi(x_2, x_1)$. Thus, restricting the

⁴Free as in beer and speech.

investments to only two values x_1 (defect) and x_2 (cooperate), results in the standard prisoner's dilemma game we encountered in the previous chapter. Therefore, the CPD game can be viewed as a generalization of the standard prisoner's dilemma game in which any level of investment can be made.

In a well-mixed population, the zero-investment strategy is the winning one. Starting at any level, investments will gradually evolve to zero, since the defectors will benefit from the investment of the cooperators without bearing the costs. The zero-investment strategy is also evolutionarily stable, i.e., an ESS. So a natural question to ask is: is there a way to promote and maintain non-zero investments levels in a population of individuals playing the CPD game?

In [KDK99] the authors show that when the CPD game is played on two-dimensional spatial lattices, investments evolve readily from very low levels to significant levels corresponding to cooperation. This occurs under very general assumptions as long as the spatial lattices are large enough for the interactions between individuals to be local enough compared with the spatial extent of the whole population. In [KD02] the authors consider the CPD game in a well-mixed population and demonstrate both analytically and by simulation, that payoff-based strategies, which embody the intuitively appealing idea that individuals invest more in cooperative interactions when they profit from these interactions, provide a natural explanation for the gradual evolution of cooperation from an initially noncooperative state and for the maintenance of cooperation thereafter. In [IKD04] the authors examine the effect of increasing the neighborhood size when the CPD game is played on a two-dimensional lattice, and find that the mean-field limit of no cooperation is reached for a critical neighborhood size of about five neighbors on each side in a Moore neighborhood⁵, which does not depend on the size of the lattice. The

⁵The Moore neighborhood comprises the eight individuals surrounding a central individual on a two-dimensional square lattice.

authors also find the related result that in a network of players, the critical average degree of vertices for which defection is the final state does not depend on network size, but only on the network topology. This critical average degree is considerably (about ten times) higher for clustered (social) networks, than for random networks. This result strengthens the argument made in [KDK99] that clustering is the mechanism which makes the development and maintenance of the cooperation possible. In the lattice topology, the authors further observed that when the neighborhood sizes for “interaction” and “update” differ by more than 0.5, cooperation is not sustainable, even for neighborhood sizes that are below the mean-field limit of defection. Finally, the authors note that the series of interaction and update neighborhoods converge, and a final cooperative state with considerable levels of average investment is achieved.

In what follows, we consider assortative interactions in a well-mixed population, characterized by the parameter $r \in [0, 1]$, as a possible mechanism for promoting and maintaining cooperation. We carry out a complete analysis of the CPD game using the framework of adaptive dynamics modified to include assortative interactions. Since the network-induced assortativity in game interactions bears an inverse relationship to its average degree, knowing how the assortativity parameter r affects cooperation in a well-mixed population also informs us about how cooperation would be affected, albeit only qualitatively, when the game is played out on networks.

5.4.1 Analysis

Consider a population of individuals playing the CPD game with payoff function as given by Equation 5.5. Suppose that each individual makes an investment x , i.e., the population is monomorphic in strategy x , called the resident strategy. Let $r \in [0, 1]$ be the degree of assortative interactions among the individuals. Now suppose a rare mutant strategy y

appears in the population. We are interested in knowing the eventual fate of the mutant strategy y . From Equations 5.3 and 5.5, we can write the invasion fitness $f'_x(y)$ of the mutant strategy as

$$f'_x(y) = r[B(y) - B(x)] - C(y) + C(x). \quad (5.6)$$

From Equation 3.109, the selection gradient $D(x)$ is given by

$$\begin{aligned} D(x) &= \left. \frac{\partial f_x(y)}{\partial y} \right|_{y=x} \\ &= rB'(x) - C'(x). \end{aligned} \quad (5.7)$$

In order to obtain the singular strategies x^* for the dynamics, we solve $D(x^*) = 0$, i.e.,

$$rB'(x^*) - C'(x^*) = 0. \quad (5.8)$$

From Equations 3.110 and 3.111, the singular strategy x^* is convergent stable if

$\left. \frac{dD}{dx} \right|_{x=x^*} < 0$ and evolutionarily stable if $\left. \frac{\partial^2 f_{x^*}(y)}{\partial y^2} \right|_{y=x^*} < 0$. But since for the CPD game, $\left. \frac{dD}{dx} \right|_{x=x^*} = \left. \frac{\partial^2 f_{x^*}(y)}{\partial y^2} \right|_{y=x^*} = rB''(x^*) - C''(x^*)$, the singular strategy x^* is convergent stable if and only if it is evolutionarily stable, i.e., when

$$rB''(x^*) - C''(x^*) < 0. \quad (5.9)$$

We next consider specific forms of cost and benefit functions $C(x)$ and $B(x)$ which are realistic and for which the analysis is mathematically tractable.

Linear Cost and Benefit Functions

Suppose the cost and benefit functions are linear functions of the investment x , i.e., suppose $C(x) = cx$ and $B(x) = bx$, where $b > c$. Figure 34 shows the two functions, along with the payoff function $\pi(x, x)$ when both of the interacting individuals are investing an

amount x . Linear benefit functions arise as approximations to more general concave functions when the investments levels are low. Linear cost functions are also interesting because they arise as approximations to more general cost functions.




Figure 34: Linear cost, linear benefit, and payoff functions for the CPD game. $C(x) = 0.4x$, $B(x) = x$. Hicks optimal strategy $x_H = 1$.

From Equation 5.7, we have $D(x) = rb - c$. Since $D(x)$ is constant, there are no singular strategies for the game, and the evolutionary fate of a mutant strategy is determined solely by the sign of $D(x)$. An initially monomorphic population in which every individual is investing an amount $0 < x_0 < x_m$, will evolve to a state in which everyone cooperates by investing an amount x_m if $D(x) > 0$, i.e., $r > \rho$, where $\rho = \frac{c}{b}$ ⁶. If on the other hand the inequality is reversed, then the population will evolve to a state in which everyone defects by making zero investment. Note that the condition that promotes cooperative investments is simply Hamilton's rule [Ham63]. We illustrate this with an example.

⁶We will refer to ρ as the cost-to-benefit ratio, although it is really the ratio of the derivative of the cost function over the derivative of the benefit function at $x = 0$, i.e., $\rho = \frac{C'(x)}{B'(x)} \Big|_{x=0}$.

Example 5.1. Consider a CPD game with linear cost function $C(x) = 0.3x$ and linear benefit function $B(x) = x$, i.e., $\rho = \frac{c}{b} = 0.3$. Suppose every individual in the population initially invests an amount $x_0 = 0.2$. We simulate the evolutionary dynamics of the game on a complete network Γ_{WM} with parameter values $n = 10000$, $s = 1$, $x_0 = 0.2$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $\nu = 100$, and $\mathcal{U} = \text{FE2}$. Figures 35(a)-(c) show the evolution of the distribution of strategies x for different values of assortativity r . As suggested by the analysis, when $r = 0$ and $r < \rho$ the population evolves to a state in which every individual makes zero investment, and when $r > \rho$ the population evolves to the state in which every individual makes the maximum investment $x_m = 1$.




Figure 35: Evolution of the distribution of strategies x in a CPD game with linear cost and benefit functions. (a) $r = 0$. (b) $r = 0.2$. (c) $r = 0.4$. Simulation parameters: complete network Γ_{WM} , $n = 10000$, $C(x) = 0.3x$ and $B(x) = x$, $s = 1$, $x_0 = 0.2$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $\nu = 100$, and $\mathcal{U} = \text{FE2}$.

□

Quadratic Cost and Benefit Functions

Now suppose the cost and benefit functions are quadratic functions of the investment x , i.e, suppose $C(x) = c_1x^2$ and $B(x) = -b_2x^2 + b_1x$, where $c_1, b_1, b_2, > 0$. Figure 36 shows the two functions, along with the payoff function $\pi(x, x)$ when both of the interacting

individuals are investing an amount x . There is good evidence that in many situations the benefit function exhibits diminishing returns for sufficiently large levels of investment [Alt79, SR83, Wei81] and the cost is well described by a quadratic function [SM76].




Figure 36: Quadratic cost, quadratic benefit, and payoff functions for the CPD game. $C(x) = x^2$, $B(x) = -x^2 + 2x$. Hicks optimal strategy $x_H = 0.5$.

Substituting for $B'(x^*)$ and $C'(x^*)$ in Equation 5.8, we obtain the singular strategy x^* for the game as

$$x^* = \frac{rb_1}{2rb_2 + 2c_2}. \quad (5.10)$$

If we let $b_1 = 2b_2$, we have $x_m = 1$. Therefore,

$$x^* = \frac{1}{1 + \frac{2c_1}{rb_1}}, \quad (5.11)$$

and $0 < x^* < 1$. Now since $rB''(x^*) - C''(x^*) = -2b_2r - 2c_1 < 0$, Equation 5.9 suggests that the singular strategy x^* is convergent stable and hence also evolutionarily stable. As a result, an initially monomorphic population in which every individual is investing an

amount $0 < x_0 < x_m$, will evolve to a state in which everyone cooperates by investing an amount x^* given by Equation 5.11. We illustrate this with an example.

Example 5.2. Consider a CPD game with quadratic cost function $C(x) = 0.3x$ and quadratic benefit function $B(x) = -x^2 + 2x$. Suppose every individual in the population initially invests an amount $x_0 = 0.1$. We simulate the game on a complete network Γ_{WM} with parameter values $n = 10000$, $s = 1$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $\nu = 100$, and $\mathcal{U} = \text{FE2}$. Figures 37(a)(b) show the evolution of the distribution of strategies x for assortativity $r = 0$ and $r = 0.5$. As suggested by the analysis, when $r = 0$ the population evolves to a state in which every individual is making zero investment, and when $r = 0.5$ the population evolves to the state in which every individual is making an investment given by the singular strategy $x^* = 0.33$, indicated by the dashed line. Note that $\pi(x^*, x^*) < \pi(x_H, x_H)$, where $x_H = 0.5$ is the Hicks optimal strategy.




Figure 37: Evolution of the distribution of strategies x in a CPD game with quadratic cost and benefit functions. (a) $r = 0$. (b) $r = 0.5$. Simulation parameters: complete network Γ_{WM} , $n = 10000$, $C(x) = x^2$ and $B(x) = -x^2 + 2x$, $s = 1$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $\nu = 100$, and $\mathcal{U} = \text{FE2}$.

□

Thus in a well-mixed population, assortative interactions provide a mechanism for promoting and maintaining cooperative investments. On a network, assortativity in game interactions increases with decreasing values of the average degree d of the network. So when the CPD game is played on networks, we expect cooperative investments to evolve on networks with small values of d . We will explore this and other network effects in detail in the following section.

5.4.2 Simulation Results

In this section we present the results of simulating the CPD game on networks having different structural properties. We first consider a linear cost function $C(x) = cx$ and a linear benefit function $B(x) = bx$, such that $b, c \in \mathbb{R}_{\geq 0}$ and $b > c$. We denote the ratio $\frac{c}{b}$ by ρ and refer to it as the *cost-to-benefit ratio*. All the simulations were carried out using the following values for the parameters involved: payoff function π given by Equation 5.5, population size $n = 10000$, initial strategy $x_0 = 0.1$, maximum strategy $x_m = 1$, mutation probability $\mu = 0.01$, standard deviation for mutations $\sigma = 0.005$, number of generations $T = 200000$, report frequency $\nu = 100$, update rule $\mathcal{U} = \text{FE2}$, and selection strength $s = 1$. See B.5 for details on how to simulate our continuous game models on complex networks. We present the results using three different kinds of plots depicting the following:

1. Variation of the long-term value \bar{x}_∞ of the mean strategy \bar{x} , averaged over the last 10% generations, with assortativity r on complete networks (average degree d on other networks) and the cost-to-benefit ratio ρ . The values of $r \in [0, 1]$ and $\rho \in (0, 1)$ were varied in steps of 0.02, and the value of d was varied from 20 to 4 in steps of -2. The intensity of the shade in the plot indicates the level of cooperation, with black corresponding to maximum cooperation ($\bar{x}_\infty = x_m$) and white corresponding to no cooperation ($\bar{x}_\infty = 0$).
2. Variation of \bar{x}_∞ with r (or d) for a fixed value of ρ . The dashed line in the case of complete networks indicates values predicted by analysis.
3. Variation of \bar{x}_∞ with ρ for a fixed value of r (or d). The dashed line in the case of complete networks indicates values predicted by analysis. We denote the fractional area under the \bar{x}_∞ curve relative to the total area of the $[\bar{x}_\infty, \rho]$ square by \mathcal{C} . We refer

to \mathcal{C} as the *cooperation index* and use it as a measure of the extent to which a network promotes cooperation. For a well-mixed population, the value of $\mathcal{C}(r)$ for a given value of assortativity r can be calculated analytically as $\mathcal{C}(r) = \int_0^1 \bar{x}_\infty(r, \rho) d\rho$, where $\bar{x}_\infty(r, \rho) = x_m$ if $r > \rho$ and 0 otherwise.

The results for the CPD game played with linear cost and benefit functions thus have the same format as the ones for the discrete prisoner's dilemma game we encountered in the previous chapter. The difference is that with the discrete prisoner's dilemma game, we were interested how the fraction x of cooperators changed with time, starting from an initial fraction x_0 of cooperators—we considered frequency-dependent selection without mutations. In case of the CPD game we are interested in how the levels of cooperative investments x made by individuals changes with time, starting from a (monomorphic) state in which every individual in the population invests an arbitrarily small amount x_0 ; in this case mutation is the key ingredient that effects change.

Complete Network

The first set of results concerns complete networks Γ_{WM} . Figures 38(a)(b) respectively show how the predicted and simulated values of the long-term mean strategy \bar{x}_∞ vary with assortativity r and cost-to-benefit ratio ρ . Figure 38(c) shows how \bar{x}_∞ varies with r when $\rho = 0.26$. Figure 38(d) shows how \bar{x}_∞ varies with ρ when $r = 0.26$.

Model Networks with Varying Average Degree

The next set of results concerns model networks with varying average degree d . Figures 39(a)(d)(g) show how the long-term mean strategy \bar{x}_∞ varies with average degree d and cost-to-benefit ratio ρ , on random regular Γ_{RR} , Erdős-Rényi Γ_{ER} , and Barabási-Albert Γ_{BA} networks respectively. Figures 39(b)(e)(h) show how \bar{x}_∞ varies with d on Γ_{RR} , Γ_{ER} ,




Figure 38: Variation of the long-term mean strategy \bar{x}_∞ with assortativity r and cost-to-benefit ratio ρ , on a complete network. (a) \bar{x}_∞ (predicted) versus r and ρ . (b) x_∞ (simulated) versus r and ρ . (c) x_∞ versus r when $\rho = 0.26$. (d) \bar{x}_∞ versus ρ when $r = 0.26$ ($\mathcal{C} = 0.241$). Simulation parameters: $C(x) = cx$ and $B(x) = bx$ with $b > c$, π given by Equation 5.5, $n = 10000$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $v = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

and Γ_{BA} networks respectively, when $\rho = 0.2$. Figures 39(c)(f)(i) show how \bar{x}_∞ varies with ρ on Γ_{RR} , Γ_{ER} , and Γ_{BA} networks respectively, when $d = 4$.




Figure 39: Variation of the long-term mean strategy \bar{x}_∞ with average degree d and cost-to-benefit ratio ρ , on model networks. (a)-(c) Random regular network Γ_{RR} . (d)-(f) Erdős-Rényi network Γ_{ER} . (g)-(i) Barabási-Albert network Γ_{BA} . (a)(d)(g) \bar{x}_∞ versus d and ρ . (b)(e)(h) \bar{x}_∞ versus d when $\rho = 0.2$. (d)(f)(i) \bar{x}_∞ versus ρ when $d = 4$ ($\mathcal{C} = 0.272, 0.219, 0.212$). Simulation parameters: $C(x) = cx$ and $B(x) = bx$ with $b > c$, π given by Equation 5.5, $n = 10000$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $v = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

Model Networks with Clustering

The next set of results concerns power-law networks Γ_{BA}^C with average degree $d = 4$, and varying clustering coefficient $C^{(2)}$. Figures 40(a)-(d) show how the long-term mean strategy \bar{x}_∞ varies with the cost-to-benefit ratio ρ on networks with clustering coefficient $C^{(2)} = 0, 0.2, 0.4$, and 0.6 .

Model Networks with Homophily

The next set of results concerns power-law networks Γ_{BA}^h with average degree $d = 4$, and varying homophily coefficient h . Figures 41(a)-(d) show how the long-term mean strategy \bar{x}_∞ varies with the cost-to-benefit ratio ρ on networks with homophily coefficient $h = -0.1, 0, 0.1$, and 0.2 .

Empirical Networks

The last set of results concerns empirical networks. Figures 42(a)-(f) show how the long-term mean strategy \bar{x}_∞ varies with the cost-to-benefit ratio ρ on: the Γ^{internet} network (a); the $\Gamma^{\text{hepcoauth}}$ network (b); the $\Gamma^{\text{astrocoauth}}$ network (c); the Γ^{fb} network (d); the Γ^{p2p} network (e); and the Γ^{protein} network (f).




Figure 40: Variation of the long-term mean strategy \bar{x}_∞ with cost-to-benefit ratio ρ , on model power-law networks with clustering coefficient $C^{(2)}$. (a) $C^{(2)} = 0$ ($\mathcal{C} = 0.215$). (b) $C^{(2)} = 0.2$ ($\mathcal{C} = 0.199$). (c) $C^{(2)} = 0.4$ ($\mathcal{C} = 0.195$). (d) $C^{(2)} = 0.6$ ($\mathcal{C} = 0.176$). Simulation parameters: $C(x) = cx$ and $B(x) = bx$ with $b > c$, π given by Equation 5.5, $n = 10000$, $d = 4$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 41: Variation of the long-term mean strategy \bar{x}_∞ with cost-to-benefit ratio ρ , on model power-law networks with homophily coefficient h . (a) $h = -0.1$ ($\mathcal{C} = 0.244$). (b) $h = 0$ ($\mathcal{C} = 0.193$). (c) $h = 0.1$ ($\mathcal{C} = 0.165$). (d) $h = 0.2$ ($\mathcal{C} = 0.122$). Simulation parameters: $C(x) = cx$ and $B(x) = bx$ with $b > c$, π given by Equation 5.5, $n = 10000$, $d = 4$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 42: Variation of the long-term mean strategy \bar{x}_∞ with cost-to-benefit ratio ρ , on empirical networks. See Table 2 for the basic properties of these networks. (a) the Γ^{internet} network ($\mathcal{C} = 0.111$). (b) the $\Gamma^{\text{hepcoauth}}$ network ($\mathcal{C} = 0.122$). (c) the $\Gamma^{\text{astrocoauth}}$ network ($\mathcal{C} = 0.029$). (d) the Γ^{fb} network ($\mathcal{C} = 0.012$). (e) the $\Gamma^{\text{p}2\text{p}}$ network ($\mathcal{C} = 0.103$). (f) the Γ^{protein} network ($\mathcal{C} = 0.191$). Simulation parameters: $C(x) = cx$ and $B(x) = bx$ with $b > c$, π given by Equation 5.5, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

We next consider a quadratic cost function $C(x) = x^2$ and a quadratic benefit function $B(x) = -x^2 + 2x$. All the simulations, unless specified otherwise, were carried out using the following values for the parameters involved: payoff function π given by Equation 5.5, population size $n = 10000$, initial strategy $x_0 = 0.1$, maximum strategy $x_m = 1$, mutation probability $\mu = 0.01$, standard deviation for mutations $\sigma = 0.005$, number of generations $T = 200000$, report frequency $\nu = 100$, update rule $\mathcal{U} = \text{FE2}$, and selection strength $s = 1$. We present the results using two different kinds of plots depicting the following:

1. Variation of the distribution of the long-term values x_∞ of strategies x , averaged over the last 10% generations, with assortativity r on complete networks (average degree d on other networks). The value of $r \in [0, 1]$ was varied in steps of 0.01, and the value of d was varied from 20 to 4 in steps of -2. The intensity of the shade in the plot indicates the number of individuals playing strategies within a particular range, with black (white) corresponding to every (no) individual being in that range. The dashed line in the case of a complete network corresponds to the predicted values of the singular strategy x^* .
2. Evolution of the distribution of strategies x . We indicate the long-term value \bar{x}_∞ of the mean strategy \bar{x} and use this value to compare the game dynamics across networks.

Complete Network and Model Networks with Varying Average Degree

The first set of results concerns complete networks Γ_{WM} , and model networks with varying average degree d . Figure 43(a) shows how the distribution of the long-term values \bar{x}_∞ of strategies x varies with assortativity r on a complete network. Figures 43(b)-(d) show how the distribution of x_∞ varies with average degree d , on random regular Γ_{RR} , Erdős-Rényi Γ_{ER} , and Barabási-Albert Γ_{BA} networks respectively.




Figure 43: Variation of the distribution of the long-term values x_∞ of strategies with assortativity r (average degree d) on complete (model) networks. (a) Complete network Γ_{WM} . (b) Random regular network Γ_{RR} . (c) Erdős-Rényi network Γ_{ER} . (d) Barabási-Albert network Γ_{BA} . Simulation parameters: $C(x) = x^2$ and $B(x) = -x^2 + 2x$, π given by Equation 5.5, $n = 10000$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

Model Networks with Clustering

The next set of results concerns power-law networks Γ_{BA}^C with average degree $d = 4$, and varying clustering coefficient $C^{(2)}$. Figures 44(a)-(d) show the evolution of the distribution of strategies x on networks with clustering coefficient $C^{(2)} = 0, 0.2, 0.4$, and 0.6 .




Figure 44: Evolution of the distribution of strategies x on model power-law networks with clustering coefficient $C^{(2)}$. (a) $C^{(2)} = 0$ ($\bar{x}_\infty = 0.197$). (b) $C^{(2)} = 0.2$ ($\bar{x}_\infty = 0.220$). (c) $C^{(2)} = 0.4$ ($\bar{x}_\infty = 0.197$). (d) $C^{(2)} = 0.6$ ($\bar{x}_\infty = 0.214$). Simulation parameters: $C(x) = x^2$ and $B(x) = -x^2 + 2x$, π given by Equation 5.5, $n = 10000$, $d = 4$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

Model Networks with Homophily

The next set of results concerns power-law networks Γ_{BA}^h with average degree $d = 4$, and varying homophily coefficient h . Figures 45(a)-(d) show the evolution of the distribution of strategies x on networks with homophily coefficient $h = -0.1, 0, 0.1$, and 0.2 .




Figure 45: Evolution of the distribution of strategies x on model power-law networks with homophily coefficient h . (a) $h = -0.1$ ($\bar{x}_\infty = 0.207$). (b) $h = 0$ ($\bar{x}_\infty = 0.205$). (c) $h = 0.1$ ($\bar{x}_\infty = 0.196$). (d) $h = 0.2$ ($\bar{x}_\infty = 0.155$). Simulation parameters: $C(x) = x^2$ and $B(x) = -x^2 + 2x$, π given by Equation 5.5, $n = 10000$, $d = 4$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

Empirical Networks

The last set of results concerns empirical networks. Figures 46(a)-(f) show the evolution of the distribution of strategies x on: the Γ^{internet} network (a); the $\Gamma^{\text{hepcoauth}}$ network (b);

the $\Gamma^{\text{astrocoauth}}$ network (c); the Γ^{fb} network (d); the Γ^{p2p} network (e); and the Γ^{protein} network (f).




Figure 46: Evolution of the distribution of strategies x on empirical networks. Ssee Table 2 for the basic properties of these networks. (a) the Γ^{internet} network ($\bar{x}_\infty = 0.138$). (b) the $\Gamma^{\text{hepcoauth}}$ network ($\bar{x}_\infty = 0.114$). (c) the $\Gamma^{\text{astrocoauth}}$ network ($\bar{x}_\infty = 0.045$). (d) the Γ^{fb} network ($\bar{x}_\infty = 0.023$). (e) the $\Gamma^{\text{P}^2\text{p}}$ network ($\bar{x}_\infty = 0.082$). (f) the Γ^{protein} network ($\bar{x}_\infty = 0.173$). Simulation parameters: $C(x) = x^2$ and $B(x) = -x^2 + 2x$, π given by Equation 5.5, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

5.5 Snowdrift Game

Two friends John and Bill are working on a joint software project. The effort that each invests into the project benefits both of them, since they both would like to see the project come to fruition. However, the investment only costs the investor. Such interactions among pairs of individuals in which the cooperative investment made by an individual (donor) benefits both the donor and the other individual (recipient), but involves a cost only to the donor, can be described using a pairwise, continuous snowdrift (CSD) game [DHK04].

The CSD game consists of two individuals, each making an investment $x \in \mathbb{R}_{\geq 0}$. The investment has the following effects: the payoff of the investor (donor) is reduced by $C(x)$, where C is a function that specifies the cost of making the investment; and the payoff of the beneficiary (recipient) is increased by $B(2x)$, where B is a function that specifies the benefit resulting from the investment. Therefore, if two interacting individuals make investments x and y , then the payoff $\pi(x, y)$ and $\pi(y, x)$ to the individuals are respectively given by

$$\begin{aligned}\pi(x, y) &= B(x + y) - C(x), \text{ and} \\ \pi(y, x) &= B(x + y) - C(y).\end{aligned}\tag{5.12}$$

We assume that the cost and benefit functions are smooth, strictly increasing, and satisfy $B(0) = C(0) = 0$.

In [DHK04] the authors study the CSD game in the context of a well-mixed population. They provide a complete classification of the adaptive dynamics for quadratic cost and benefit functions. Depending on the coefficients of these functions, the end state of the population can be one of: evolutionary branching; convergent stable ESS;

evolutionary repeller leading to bi-stable evolutionary dynamics; and in the absence of a singular strategy, either uniform selection for defectors or cooperators.

In what follows, we consider assortative interactions in a well-mixed population, characterized by the parameter $r \in [0, 1]$, as a possible mechanism for promoting and maintaining cooperation. We carry out a complete analysis of the CSD game using the framework of adaptive dynamics modified to include assortative interactions. Since the network-induced assortativity in game interactions bears an inverse relationship to its average degree, knowing how the assortativity parameter r affects cooperation in a well-mixed population also informs us about how cooperation would be affected, albeit only qualitatively, when the game is played out on networks.

5.5.1 Analysis

Consider a population of individuals playing the CSD game with the payoff function given by Equation 5.12. Suppose that each individual makes an investment x , i.e., the population is monomorphic in strategy x , called the resident strategy. Let $r \in [0, 1]$ be the degree of assortative interactions among the individuals. Now suppose a rare mutant strategy y appears in the population. We are interested in knowing the eventual fate of the mutant strategy y . From Equations 5.3 and 5.5, we can write the invasion fitness $f_x^r(y)$ of the mutant strategy as

$$f_x^r(y) = (1 - r)B(x + y) + rB(2y) - C(y) - B(2x) + c(x). \quad (5.13)$$

From Equation 3.109, the selection gradient $D(x)$ is given by

$$\begin{aligned} D(x) &= \left. \frac{\partial f_x^r(y)}{\partial y} \right|_{y=x} \\ &= (1 + r)B'(2x) - C'(x). \end{aligned} \quad (5.14)$$

In order to obtain the singular strategies x^* for the dynamics, we solve $D(x^*) = 0$, i.e.,

$$(1+r)B'(2x^*) - C'(x^*) = 0. \quad (5.15)$$

From Equation 3.110, the singular strategy x^* is convergent stable if $\left. \frac{dD}{dx} \right|_{x=x^*} < 0$, i.e.,

$$2(1+r)B''(2x^*) - C''(x^*) < 0, \quad (5.16)$$

and a repeller if the inequality is reversed.

From Equation 3.111, the singular strategy x^* is evolutionarily stable if $\left. \frac{\partial^2 f_{x^*}(y)}{\partial y^2} \right|_{y=x^*} < 0$, i.e.,

$$B''(2x^*) + 3rB''(2x^*) - C''(x^*) < 0, \quad (5.17)$$

and an evolutionary branching point (EBP) if the inequality is reversed.

We next consider specific forms of cost and benefit functions $C(x)$ and $B(x)$ which are realistic and for which the analysis is mathematically tractable.

Quadratic Cost and Benefit Functions

Suppose the cost and benefit functions are quadratic functions of the investment x , i.e., suppose $C(x) = -c_2x^2 + c_1x$ and $B(x) = -b_2x^2 + b_1x$, where $c_1, c_2, b_1, b_2, > 0$. Figures 47(a)(b) show the two functions, along with the payoff function $\pi(x, x)$ when both of the interacting individuals are investing an amount x . Saturating benefits are clearly realistic [Alt79, SR83, Wei81], whereas costs could often be expected to accelerate. However, diminishing additional costs of larger cooperative investments are reasonable whenever the initiation of cooperative acts, is more costly than subsequent increases in cooperative investments.

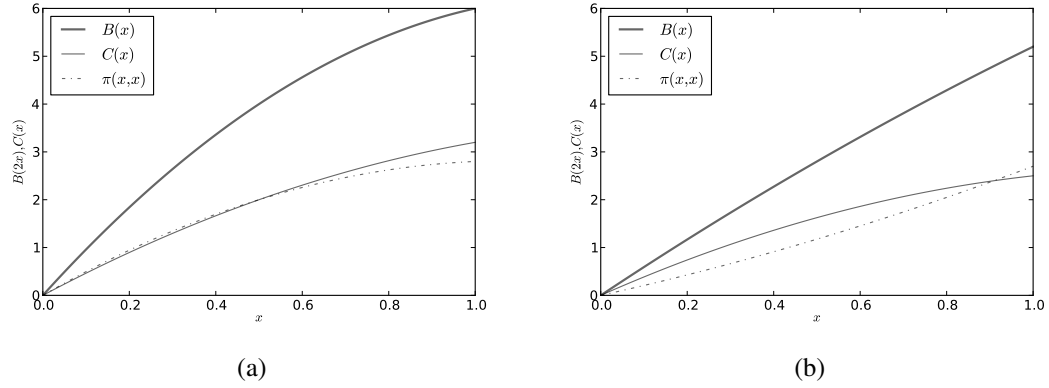


Figure 47: Quadratic cost, quadratic benefit, and payoff functions for the CSD game. (a) $C(x) = -1.6x^2 + 4.8x$, $B(x) = -x^2 + 5x$, Hicks optimal strategy $x_H = 1$. (b) $C(x) = -1.5x^2 + 4x$, $B(x) = -0.2x^2 + 3x$, Hicks optimal strategy $x_H = 1$.

Substituting for $B'(x^*)$ and $C'(x^*)$ in Equation 5.15, we obtain the singular strategy x^* for the game as

$$x^* = \frac{(1+r)b_1 - c_1}{4(1+r)b_2 - 2c_2}. \quad (5.18)$$

From Equation 5.16, we see that the singular strategy x^* is convergent stable if

$$b_2 > \frac{c_2}{2(1+r)}, \quad (5.19)$$

and a repeller if the inequality is reversed.

From Equation 5.17, we see that the singular strategy x^* is an ESS if

$$b_2 > \frac{c_2}{1+3r}, \quad (5.20)$$

and an evolutionary branching point (EBP) if the inequality is reversed.

Thus the state to which an initially monomorphic population in which every individual is investing an amount $0 < x_0 < x_m$ will evolve crucially depends on the coefficients

c_1, c_2, b_1 , and b_2 of the cost and benefit functions. If Equation 5.19 is satisfied, then the population evolves to a state in which everyone is investing x^* , and the fate of the population thereafter depends on whether the Equation 5.20 is satisfied or not. If it is, then x^* is evolutionarily stable, so the population remains in that state. Otherwise, x^* is an EBP and the population splits into two phenotypic clusters which evolve away from each other. If Equation 5.19 is not satisfied, then x^* is a repeller and the fate of the population depends on how the initial (resident) strategy x_0 compares with x^* . If $x_0 < x^*$, then the population evolves to the zero-investment state. On the other hand, if $x_0 > x^*$, then the population evolves to the maximum-investment state. We illustrate these possibilities with an example.

Example 5.3. Consider a CSD game with quadratic cost function $C(x) = -1.6x^2 + 4.8x$ and quadratic benefit function $B(x) = -x^2 + 5x$. Suppose every individual in the population initially invests an amount $x_0 = 0.1$. We simulate the game on a complete network Γ_{WM} with parameter values $n = 10000$, $s = 1$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $v = 100$, and $\mathcal{U} = \text{FE2}$. If $r = 0$ the singular strategy $x^* = 0.25$ is an EBP, so the population evolves toward x^* and then branches out into two phenotypic clusters (Figure 48(a)). When $r = 0.3$, the singular strategy $x^* = 0.85$ is an ESS, so the population evolves toward x^* and remains there (Figure 48(b)). Note that $\pi(x^*, x^*) < \pi(x_H, x_H)$, where $x_H = 1$ is the Hicks optimal strategy.

Now consider the cost and benefit functions to be $C(x) = -1.5x^2 + 4x$ and $B(x) = -0.2x^2 + 3x$, and let $r = 0$. The singular strategy $x^* = 0.45$ is a repeller in this case, so if a population starts with the strategy $x_0 = 0.3$, it evolves to the zero-investment state (Figure 48(c)). On the other hand, if $x_0 = 0.6$, then the population evolves to the maximum-investment state (Figure 48(d)).

□




Figure 48: Evolution of the distribution of strategies x in a CSD game with quadratic cost and quadratic benefit functions. $C(x) = -c_2x^2 + c_1x^2$ and $B(x) = -b_2x^2 + b_1x$. (a) $c_1 = 4.8, c_2 = 1.6, b_1 = 5, b_2 = 1, r = 0, x_0 = 0.1$ ($x^* = 0.25$ is an EBP). (b) $c_1 = 4.8, c_2 = 1.6, b_1 = 5, b_2 = 1, r = 0.3, x_0 = 0.1$ ($x^* = 0.85$ is an ESS). (c) $c_1 = 4, c_2 = 1.5, b_1 = 3, b_2 = 0.2, r = 0, x_0 = 0.3$ ($x^* = 0.45$ is a repeller). (d) $c_1 = 4, c_2 = 1.5, b_1 = 3, b_2 = 0.2, r = 0, x_0 = 0.6$ ($x^* = 0.45$ is a repeller). Simulation parameters: complete network Γ_{WM} , $n = 10000$, $s = 1, x_m = 1, \mu = 0.01, \sigma = 0.005, T = 300000, \nu = 100$, and $\mathcal{U} = \text{FE2}$.

Thus in a well-mixed population, by changing the degree of assortative interactions, we can vary the nature of the singular strategy, each resulting in a different end state for the evolutionary dynamics. The singular strategy can be one of: an EBP in which case

high and low investors coexist—this scenario is referred to as “tragedy of the commune” [DHK04]; an ESS in which case everyone invests the same amount given by the singular strategy; or an evolutionary repeller in which case the individuals either invest nothing or invest the maximum amount depending on how the initial strategy compares with the singular strategy. On a network, assortativity in game interactions increases with decreasing values of the average degree d of the network. So when the CSD game is played on networks, we expect to see similar kinds of behavior by changing d . We will explore this and other network effects in the following section.

5.5.2 Simulation Results

In this section we present the results of simulating the CSD game on networks having different structural properties. We first consider a quadratic cost function

$C(x) = -1.6x^2 + 4.8x$ and a quadratic benefit function $B(x) = -x^2 + 5x$. All the

simulations were carried out using the following values for the parameters involved:

payoff function π given by Equation 5.12, population size $n = 10000$, initial strategy

$x_0 = 0.1$, maximum strategy $x_m = 1$, mutation probability $\mu = 0.01$, standard deviation for

mutations $\sigma = 0.005$, number of generations $T = 300000$, report frequency $\nu = 100$,

update rule $\mathcal{U} = \text{FE2}$, and selection strength $s = 1$. See B.5 for details on how to simulate

our continuous game models on complex networks. We present the results using two

different kinds of plots depicting the following:

1. Variation of the distribution of the long-term values x_∞ of strategies x , averaged over the last 10% generations, with assortativity r on complete networks (average degree d on other networks). The value of $r \in [0, 0.6]$ was varied in steps of 0.01, and the value of d was varied from 20 to 4 in steps of -2. As r increases (d decreases), the singular strategy x^* transitions from an EBP to an ESS. The intensity of the shade in

the plot indicates the number of individuals playing strategies within a particular range, with black (white) corresponding to every (no) individual being in that range. The dashed line in the case of a complete network corresponds to the predicted values of the singular strategy x^* .

2. Evolution of the distribution of strategies x . If in the long term the individuals in the population play a single strategy, we characterize the state of the population by \bar{x}_∞ . If on the other hand the population undergoes evolutionary branching, then we characterize the state of the population using x_{BP} , the strategy x at which the population undergoes branching. We use these values to compare the game dynamics across networks.

Complete Network and Model Networks with Varying Average Degree

The first set of results concerns complete networks Γ_{WM} , and model networks with varying average degree d . Figures 49(a) shows how the distribution of the long-term values \bar{x}_∞ of strategies x varies with assortativity r on a complete network. Figures 49(b)-(d) show how the distribution of x_∞ varies with average degree d , on random regular Γ_{RR} , Erdős-Rényi Γ_{ER} , and Barabási-Albert Γ_{BA} networks respectively.

Model Networks with Clustering

The next set of results concerns power-law networks Γ_{BA}^C with varying clustering coefficient $C^{(2)}$. Figures 50(a)-(d) show the evolution of the distribution of strategies x on networks with average degree $d = 4$ and clustering coefficient $C^{(2)} = 0, 0.2, 0.4$, and 0.6 . Figures 51(a)-(d) show the evolution of the distribution of strategies x on networks with average degree $d = 10$ and clustering coefficient $C^{(2)} = 0.1, 0.2, 0.3$, and 0.4 .




Figure 49: Variation of the distribution of the long-term values x_∞ of strategies with assortativity r (average degree d) on complete (model) networks. (a) Complete network Γ_{WM} . (b) Random regular network Γ_{RR} . (c) Erdős-Rényi network Γ_{ER} . (d) Barabási-Albert network Γ_{BA} . Simulation parameters: $C(x) = -1.6x^2 + 4.8x$ and $B(x) = -x^2 + 5x$, π given by Equation 5.12, $n = 10000$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

Model Networks with Homophily

The next set of results concerns power-law networks Γ_{BA}^h with varying homophily coefficient h . Figures 52(a)-(d) show the evolution of the distribution of strategies x on




Figure 50: Evolution of the distribution of strategies x on model power-law networks with average degree $d = 4$ and clustering coefficient $C^{(2)}$. (a) $C^{(2)} = 0$ ($\bar{x}_\infty = 0.821$). (b) $C^{(2)} = 0.2$ ($\bar{x}_\infty = 0.797$). (c) $C^{(2)} = 0.4$ ($\bar{x}_\infty = 0.796$). (d) $C^{(2)} = 0.6$ ($\bar{x}_\infty = 0.78$). Simulation parameters: $C(x) = -1.6x^2 + 4.8x$ and $B(x) = -x^2 + 5x$, π given by Equation 5.12, $n = 10000$, $d = 4$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

networks with average degree $d = 4$ and homophily coefficient $h = -0.1, 0, 0.1$, and 0.2 . Figures 53(a)-(d) show the evolution of the distribution of strategies x on networks with average degree $d = 10$ and homophily coefficient $h = -0.1, 0, 0.1$, and 0.2 .




Figure 51: Evolution of the distribution of strategies x on model power-law networks with average degree $d = 10$ and clustering coefficient $C^{(2)}$. (a) $C^{(2)} = 0.1$ ($x_{BP} = 0.574$). (b) $C^{(2)} = 0.2$ ($x_{BP} = 0.516$). (c) $C^{(2)} = 0.3$ ($x_{BP} = 0.543$). (d) $C^{(2)} = 0.4$ ($\bar{x}_\infty = 0.551$). Simulation parameters: $C(x) = -1.6x^2 + 4.8x$ and $B(x) = -x^2 + 5x$, π given by Equation 5.12, $n = 10000$, $d = 10$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

Empirical Networks

The last set of results concerns empirical networks. Figures 54(a)-(f) show the evolution of the distribution of strategies x on: the Γ^{internet} network (a); the $\Gamma^{\text{hepcoauth}}$ network (b);




Figure 52: Evolution of the distribution of strategies x on model power-law networks with average degree $d = 4$ and homophily coefficient h . (a) $h = -0.1$ ($\bar{x}_\infty = 0.82$). (b) $h = 0$ ($\bar{x}_\infty = 0.8$). (c) $h = 0.1$ ($\bar{x}_\infty = 0.751$). (d) $h = 0.2$ ($x_{BP} = 0.7$). Simulation parameters: $C(x) = -1.6x^2 + 4.8x$ and $B(x) = -x^2 + 5x$, π given by Equation 5.12, $n = 10000$, $d = 4$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

the $\Gamma^{\text{astrocoauth}}$ network (c); the Γ^{fb} network (d); the Γ^{p2p} network (e); and the Γ^{protein} network (f).

We next consider a quadratic cost function $C(x) = -1.5x^2 + 4x$ and a quadratic benefit function $B(x) = -0.2x^2 + 3x$. All the simulations were carried out using the following




Figure 53: Evolution of the distribution of strategies x on model power-law networks with average degree $d = 10$ and homophily coefficient h . (a) $h = -0.1$ ($x_{BP} = 0.491$). (b) $h = 0$ ($x_{BP} = 0.545$). (c) $h = 0.1$ ($x_{BP} = 0.546$). (d) $h = 0.2$ ($x_{BP} = 0.54$). Simulation parameters: $C(x) = -1.6x^2 + 4.8x$ and $B(x) = -x^2 + 5x$, π given by Equation 5.12, $d = 10$, $n = 10000$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

values for the parameters involved: payoff function π given by Equation 5.12, population size $n = 10000$, initial strategy $x_0 = 0.3$, maximum strategy $x_m = 1$, mutation probability $\mu = 0.01$, standard deviation for mutations $\sigma = 0.005$, number of generations $T = 300000$, report frequency $\nu = 100$, update rule $\mathcal{U} = \text{FE2}$, and selection strength $s = 1$. We present the results using two different kinds of plots depicting the following:

1. Variation of the distribution of the long-term values x_∞ of strategies x , averaged over the last 10% generations, with assortativity r on complete networks (average degree d on other networks). The value of $r \in [0, 0.3]$ was varied in steps of 0.01, and the value of d was varied from 20 to 4 in steps of -2. For the entire range of r and d , the singular strategy x^* is a repeller. When $r = 0$ ($d = 20$), the $x_0 < x^*$, so the population evolves to the state in which everyone makes zero investment. But as r increases (d decreases), a stage is reached beyond which $x_0 > x^*$, and the population evolves to a state in which everyone makes the maximum investment $x_m = 1$. The intensity of the shade in the plot indicates the number of individuals playing strategies within a particular range, with black (white) corresponding to every (no) individual being in that range. The dashed line in the case of a complete network corresponds to the predicted values of the singular strategy x^* .
2. Evolution of the distribution of strategies x . We indicate the long-term value \bar{x}_∞ of the mean strategy \bar{x} and use this value to compare the game dynamics across networks.

Complete Network and Model Networks with Varying Average Degree

The first set of results concerns complete networks Γ_{WM} , and model networks with varying average degree d . Figures 55(a) shows how the distribution of the long-term values \bar{x}_∞ of strategies x varies with assortativity r on a complete network. Figures 55(b)-(d) show how the distribution of x_∞ varies with average degree d , on random regular Γ_{RR} , Erdős-Rényi Γ_{ER} , and Barabási-Albert Γ_{BA} networks respectively.

Model Networks with Clustering

The next set of results concerns power-law networks Γ_{BA}^C with varying clustering coefficient $C^{(2)}$. Figures 56(a)-(d) show the evolution of the distribution of strategies x on networks with average degree $d = 4$ and clustering coefficient $C^{(2)} = 0, 0.2, 0.4$, and 0.6 . Figures 57(a)-(d) show the evolution of the distribution of strategies x on networks with average degree $d = 10$ and clustering coefficient $C^{(2)} = 0.1, 0.2, 0.3$, and 0.4 .

Model Networks with Homophily

The next set of results concerns power-law networks Γ_{BA}^h with varying homophily coefficient h . Figures 58(a)-(d) show the evolution of the distribution of strategies x on networks with average degree $d = 4$ and homophily coefficient $h = -0.1, 0, 0.1$, and 0.2 . Figures 59(a)-(d) show the evolution of the distribution of strategies x on networks with average degree $d = 10$ and homophily coefficient $h = -0.1, 0, 0.1$, and 0.2 .

Empirical Networks

The last set of results concerns empirical networks. Figures 60(a)-(f) show the evolution of the distribution of strategies x on: the Γ^{internet} network (a); the $\Gamma^{\text{hepcoauth}}$ network (b); the $\Gamma^{\text{astrocoauth}}$ network (c); the Γ^{fb} network (d); the Γ^{p2p} network (e); and the Γ^{protein} network (f).




Figure 54: Evolution of the distribution of strategies x on empirical networks. See Table 2 for the basic properties of these networks. (a) the Γ^{internet} network ($\bar{x}_\infty = 0.734$). (b) the $\Gamma^{\text{hepcoauth}}$ network ($\bar{x}_\infty = 0.744$). (c) the $\Gamma^{\text{astrocoauth}}$ network ($x_{BP} = 0.4$). (d) the Γ^{fb} network ($x_{BP} = 0.3$). (e) the Γ^{p2p} network ($x_{BP} = 0.6$). (f) the Γ^{protein} network ($x_{BP} = 0.75$). Simulation parameters: $C(x) = -1.6x^2 + 4.8x$ and $B(x) = -x^2 + 5x$, π given by Equation 5.12, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 55: Variation of the distribution of the long-term values x_∞ of strategies with assortativity r (average degree d) on complete (mode) networks. (a) Complete network Γ_{WM} . (b) Random regular network Γ_{RR} . (c) Erdős-Rényi network Γ_{ER} . (d) Barabási-Albert network Γ_{BA} . Simulation parameters: $C(x) = -1.5x^2 + 4x$ and $B(x) = -0.2x^2 + 3x$, π given by Equation 5.12, $n = 10000$, $x_0 = 0.3$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 56: Evolution of the distribution of strategies x on model power-law networks with average degree $d = 4$ and clustering coefficient $C^{(2)}$. (a) $C^{(2)} = 0$ ($\bar{x}_\infty = 1$). (b) $C^{(2)} = 0.2$ ($\bar{x}_\infty = 1$). (c) $C^{(2)} = 0.4$ ($\bar{x}_\infty = 1$). (d) $C^{(2)} = 0.6$ ($\bar{x}_\infty = 1$). Simulation parameters: $C(x) = -1.5x^2 + 4x$ and $B(x) = -0.2x^2 + 3x$, π given by Equation 5.12, $n = 10000$, $d = 4$, $x_0 = 0.3$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 57: Evolution of the distribution of strategies x on model power-law networks with average degree $d = 10$ and clustering coefficient $C^{(2)}$. (a) $C^{(2)} = 0.1$ ($\bar{x}_\infty = 0$). (b) $C^{(2)} = 0.2$ ($\bar{x}_\infty = 0$). (c) $C^{(2)} = 0.3$ ($\bar{x}_\infty = 0$). (d) $C^{(2)} = 0.4$ ($\bar{x}_\infty = 0$). Simulation parameters: $C(x) = -1.5x^2 + 4x$ and $B(x) = -0.2x^2 + 3x$, π given by Equation 5.12, $n = 10000$, $d = 10$, $x_0 = 0.3$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 58: Evolution of the distribution of strategies x on model power-law networks with average degree $d = 4$ and homophily coefficient h . (a) $h = -0.1$ ($\bar{x}_\infty = 1$). (b) $h = 0$ ($\bar{x}_\infty = 1$). (c) $h = 0.1$ ($\bar{x}_\infty = 1$). (d) $h = 0.2$ ($\bar{x}_\infty = 1$). Simulation parameters: $C(x) = -1.5x^2 + 4x$ and $B(x) = -0.2x^2 + 3x$, π given by Equation 5.12, $n = 10000$, $d = 4$, $x_0 = 0.3$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 59: Evolution of the distribution of strategies x on model power-law networks with average degree $d = 10$ and homophily coefficient h . (a) $h = -0.1$ ($\bar{x}_\infty = 0$). (b) $h = 0$ ($\bar{x}_\infty = 0$). (c) $h = 0.1$ ($\bar{x}_\infty = 0$). (d) $h = 0.2$ ($\bar{x}_\infty = 0$). Simulation parameters: $C(x) = -1.5x^2 + 4x$ and $B(x) = -0.2x^2 + 3x$, π given by Equation 5.12, $n = 10000$, $d = 10$, $x_0 = 0.3$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 60: Evolution of the distribution of strategies x on empirical networks. See Table 2 for the basic properties of these networks. (a) the Γ^{internet} network ($\bar{x}_\infty = 1$). (b) the $\Gamma^{\text{hepcoauth}}$ network ($\bar{x}_\infty = 1$). (c) the $\Gamma^{\text{astrocoauth}}$ network ($\bar{x}_\infty = 0$). (d) the Γ^{fb} network ($\bar{x}_\infty = 0$). (e) the $\Gamma^{\text{P}2\text{P}}$ network ($\bar{x}_\infty = 0$). (f) the Γ^{protein} network ($\bar{x}_\infty = 1$). Simulation parameters: $C(x) = -1.5x^2 + 4x$ and $B(x) = -0.2x^2 + 3x$, π given by Equation 5.12, $x_0 = 0.3$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

5.6 Tragedy of the Commons

Two friends John and Bill share an internet connection at home. The total bandwidth available to them is of course limited. Each of them benefits from consuming the shared bandwidth, but the incurring costs (slowdown in connectivity) are diluted and shared among both. Such situations involving pairs of individuals sharing a common-pool resource in which one individual's consumption of the resource benefits that individual but costs both the individuals involved, can be described using a pairwise, continuous tragedy of the commons (CTOC) game [KDH10].

The CTOC game consists of two individuals, each making an investment⁷ $x \in \mathbb{R}_{\geq 0}$. The investment has the following effects: the payoff of the investor (donor) is reduced by $C(2x)$, where C is a function that specifies the cost of making the investment; and the payoff of the beneficiary (recipient) is increased by $B(x)$, where B is a function that specifies the benefit resulting from the investment. Therefore, if two interacting individuals make investments x and y , then the payoff $\pi(x, y)$ and $\pi(y, x)$ to the individuals are respectively given by

$$\begin{aligned}\pi(x, y) &= B(x) - C(x + y), \text{ and} \\ \pi(y, x) &= B(y) - C(x + y).\end{aligned}\tag{5.21}$$

We assume that at least for small x , $B(x) > C(x)$. This simply reflects the fact that the public resource is valuable and thus creates an incentive for individuals to consume and take advantage of the resource.

In the classic tragedy of the commons [Har68], each individual benefits from consuming the common-pool resource, but the incurring costs are diluted and shared

⁷Recall that investment in this context means the level of consumption of a limited common-pool resource, and cooperative investments mean lower levels of consumption.

among all interacting individuals. As a result, the common resource is prone to exploitation and self-interest drives over-consumption of the resource to the detriment of all. In [KDH10] the authors consider the CTOC game in a well-mixed setting and demonstrate both analytically and by simulation that the evolutionary outcomes based on a continuous range of consumption levels of common-pool resources are often strikingly different from the classic tragedy of the commons. The authors reveal a second tragedy: not only is the common resource overexploited, but selection may result in states in which high and low consumers stably coexist.

In what follows, we consider assortative interactions in a well-mixed population, characterized by the parameter $r \in [0, 1]$, as a possible mechanism for promoting and maintaining cooperation. We carry out a complete analysis of the CTOC game using the framework of adaptive dynamics modified to include assortative interactions. Since the network-induced assortativity in game interactions bears an inverse relationship to its average degree, knowing how the assortativity parameter r affects cooperation in a well-mixed population also informs us about how cooperation would be affected, albeit only qualitatively, when the game is played out on networks.

5.6.1 Analysis

Consider a population of individuals playing the CTOC game with the payoff function given by Equation 5.21. Suppose that each individual makes an investment x , i.e., the population is monomorphic in strategy x , called the resident strategy. Let $r \in [0, 1]$ be the degree of assortative interactions among the individuals. Now suppose a rare mutant strategy y appears in the population. We are interested in knowing the eventual fate of the mutant strategy y . From Equations 5.3 and 5.5, we can write the invasion fitness $f_x^r(y)$ of

the mutant strategy as

$$f_x^r(y) = B(y) - rC(2y) - (1-r)C(x+y) - B(x) + C(2x). \quad (5.22)$$

From Equation 3.109, the selection gradient $D(x)$ is given by

$$\begin{aligned} D(x) &= \left. \frac{\partial f_x(y)}{\partial y} \right|_{y=x} \\ &= B'(x) - (1+r)C'(2x). \end{aligned} \quad (5.23)$$

In order to obtain the singular strategies x^* for the dynamics, we solve $D(x^*) = 0$, i.e.,

$$B'(x^*) - (1+r)C'(2x^*) = 0. \quad (5.24)$$

From Equation 3.110, the singular strategy x^* is convergent stable if $\left. \frac{dD}{dx} \right|_{x=x^*} < 0$, i.e.,

$$B''(x^*) - 2(1+r)C''(2x^*) < 0, \quad (5.25)$$

and a repeller if the inequality is reversed.

From Equation 3.111, the singular strategy x^* is evolutionarily stable if

$$\begin{aligned} \left. \frac{\partial^2 f_{x^*}(y)}{\partial y^2} \right|_{y=x^*} &< 0, \text{ i.e.,} \\ B''(x^*) - 3rC'''(2x^*) - C''(2x^*) &< 0, \end{aligned} \quad (5.26)$$

and an evolutionary branching point (EBP) if the inequality is reversed.

We next consider specific forms of cost and benefit functions $C(x)$ and $B(x)$ which are realistic and for which the analysis is mathematically tractable.

Quadratic Cost and Cubic Benefit Functions

Suppose the cost function is a quadratic function of the investment x and the benefit function is a cubic function of the investment x , i.e, suppose $C(x) = c_1x^2$ and

$B(x) = -b_3x^3 + b_2x^2 + b_1x$, where $c_1, b_1, b_2, b_3 > 0$. Figure 61 shows the two functions, along with the payoff function $\pi(x, x)$ when both of the interacting individuals are investing an amount x . Accelerating costs represent a realistic assumption and are often observed in nature [SM76]. Benefits also often accelerate initially and then saturate, as in the case of Sigmoid benefit functions [Smi82].




Figure 61: Quadratic cost, cubic benefit, and payoff functions for the CTOC game. $C(x) = x^2$, $B(x) = -0.0834x^3 + 2x^2 + x$. Hicks optimal strategy $x_H = 0.25$.

To simplify the analysis, we take $b_2 = 2b_1$ and $c_1 = b_1$. Substituting for $B'(x^*)$ and $C'(x^*)$ in Equation 5.24, we obtain the singular strategy x^* for the game as

$$x^* = \frac{\sqrt{16r^2b_1^2 + 12b_1b_3} - 4rb_1}{6b_3}. \quad (5.27)$$

From Equation 5.25, since $B''(x^*) - 3rC''(2x^*) - C''(2x^*) = -6b_3x^* - 4rb_1 < 0$ we see that the singular strategy x^* is convergent stable.

From Equation 5.17, we see that the singular strategy x^* is an ESS if

$$b_1 < \frac{\sqrt{16r^2b_1^2 + 12b_1b_3}}{2(1-r)}, \quad (5.28)$$

and an evolutionary branching point (EBP) if the inequality is reversed.

Thus the state to which an initially monomorphic population in which every individual is investing an amount $0 < x_0 < x_m$ will evolve crucially depends on the coefficients c_1, b_1 , and b_3 of the cost and benefit functions. The population first evolves to a state in which everyone is investing x^* , and the fate of the population thereafter depends on whether the Equation 5.28 is satisfied or not. If it is, then x^* is evolutionarily stable, so the population remains in that state. Otherwise, x^* is an EBP and the population splits into two phenotypic clusters which evolve away from each other. We illustrate these possibilities with an example.

Example 5.4. Consider a CTOC game with quadratic cost function $C(x) = x^2$ and cubic benefit function $B(x) = -0.0834x^3 + 2x^2 + x$. Suppose every individual in the population initially invests an amount $x_0 = 0.1$. We simulate the game on a complete network Γ_{WM} with parameter values $n = 10000$, $s = 1$, $x_0 = 0.1$, $x_m = 3$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $v = 100$, and $\mathcal{U} = \text{FE2}$. If $r = 0$ the singular strategy $x^* = 2$ is an EBP, so the population evolves toward x^* and then branches out into two phenotypic clusters (Figure 62(a)). When $r = 0.4$, the singular strategy $x^* = 0.57$ is an ESS, so the population evolves toward x^* and remains there (Figure 62(b)). Note that $\pi(x^*, x^*) < \pi(x_H, x_H)$, where $x_H = 0.25$ is the Hicks optimal strategy.

□

Thus in a well-mixed population, by changing the degree of assortative interactions, we can vary the nature of the singular strategy, each resulting in a different end state for the evolutionary dynamics. The singular strategy can be: an EBP in which case high and low investors, i.e., individuals who overexploit the common-pool resource and who consume less, can coexist—this scenario poses a second tragedy, and is referred to as “tragedy of the commune” [KDH10]; or an ESS in which case everyone invests the same




Figure 62: Evolution of the distribution of strategies x in a CTOC game with quadratic cost and cubic benefit functions. (a) $r = 0$ ($x^* = 2$ is an EBP); (b) $r = 0.4$ ($x^* = 0.57$ is an ESS). Simulation parameters: complete network Γ_{WM} , $n = 10000$, $C(x) = x^2$ and $B(x) = -0.0834x^3 + 2x^2 + x$, $s = 1$, $x_0 = 0.1$, $x_m = 3$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, and $\mathcal{U} = \text{FE2}$.

amount given by the singular strategy. On a network, assortativity in game interactions increases with decreasing values of the average degree d of the network. So when the CTOC game is played on networks, we expect to see similar kinds of behavior by changing d . We will explore this and other network effects in the following section.

5.6.2 Simulation Results

In this section we present the results of simulating the CTOC game on networks having different structural properties. We consider a quadratic cost function $C(x) = x^2$ and a cubic benefit function $B(x) = -0.0834x^3 + 2x^2 + x$. All the simulations were carried out using the following values for the parameters involved: payoff function π given by Equation 5.21, population size $n = 10000$, initial strategy $x_0 = 0.1$, maximum strategy $x_m = 3$, mutation probability $\mu = 0.01$, standard deviation for mutations $\sigma = 0.005$, number of generations $T = 300000$ unless specified otherwise in the figure legends, report frequency $\nu = 100$, update rule $\mathcal{U} = \text{FE2}$, and selection strength $s = 1$. See B.5 for details on how to simulate our continuous game models on complex networks. We present the results using two different kinds of plots depicting the following:

1. Variation of the distribution of the long-term values x_∞ of strategies x , averaged over the last 10% generations, with assortativity r on complete networks (average degree d on other networks). The value of $r \in [0, 1]$ was varied in steps of 0.01, and the value of d was varied from 20 to 4 in steps of -2. As r increases (d decreases), the singular strategy x^* transitions from an EBP to an ESS. The intensity of the shade in the plot indicates the number of individuals playing strategies within a particular range, with black (white) corresponding to every (no) individual being in that range. The dashed line in the case of a complete network corresponds to the predicted values of the singular strategy x^* .
2. Evolution of the distribution of strategies x . If in the long term the individuals in the population play a single strategy, we characterize the state of the population by \bar{x}_∞ . If on the other hand the population undergoes evolutionary branching, then we characterize the state of the population using x_{BP} , the strategy x at which the

population undergoes branching. We use these values to compare the game dynamics across networks.

Complete Network and Model Networks with Varying Average Degree

The first set of results concerns complete networks Γ_{WM} , and model networks with varying average degree d . Figures 63(a) shows how the distribution of the long-term values \bar{x}_∞ of strategies x varies with assortativity r on a complete network. Figures 63(b)-(d) show how the distribution of x_∞ varies with average degree d , on random regular Γ_{RR} , Erdős-Rényi Γ_{ER} , and Barabási-Albert Γ_{BA} networks respectively.

Model Networks with Clustering

The next set of results concerns power-law networks $\Gamma_{\text{BA}}^{\text{C}}$ with varying clustering coefficient $C^{(2)}$. Figures 64(a)-(d) show the evolution of the distribution of strategies x on networks with average degree $d = 4$ and clustering coefficient $C^{(2)} = 0, 0.2, 0.4$, and 0.6 . Figures 65(a)-(d) show the evolution of the distribution of strategies x on networks with average degree $d = 10$ and clustering coefficient $C^{(2)} = 0.1, 0.2, 0.3$, and 0.4 .

Model Networks with Homophily

The next set of results concerns power-law networks $\Gamma_{\text{BA}}^{\text{h}}$ with varying homophily coefficient h . Figures 66(a)-(d) show the evolution of the distribution of strategies x on networks with average degree $d = 4$ and homophily coefficient $h = -0.1, 0, 0.1$, and 0.2 . Figures 67(a)-(d) show the evolution of the distribution of strategies x on networks with average degree $d = 10$ and homophily coefficient $h = -0.1, 0, 0.1$, and 0.2 .




Figure 63: Variation of the distribution of the long-term values x_∞ of strategies with assortativity r (average degree d) on complete (model) networks. (a) Complete network Γ_{WM} . (b) Random regular network Γ_{RR} . (c) Erdős-Rényi network Γ_{ER} . (d) Barabási-Albert network Γ_{BA} . Simulation parameters: $C(x) = x^2$ and $B(x) = -0.0834x^3 + 2x^2 + x$, π given by Equation 5.21, $n = 10000$, $x_0 = 0.1$, $x_m = 3$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

Empirical Networks

The last set of results concerns empirical networks. Figures 68(a)-(f) show the evolution of the distribution of strategies x on: the Γ^{internet} network (a); the $\Gamma^{\text{hepcoauth}}$ network (b);




Figure 64: Evolution of the distribution of strategies x on model power-law networks with average degree $d = 4$ and clustering coefficient $C^{(2)}$. (a) $C^{(2)} = 0$ ($\bar{x}_\infty = 0.815$). (b) $C^{(2)} = 0.2$ ($\bar{x}_\infty = 0.788$). (c) $C^{(2)} = 0.4$ ($\bar{x}_\infty = 0.841$). (d) $C^{(2)} = 0.6$ ($\bar{x}_\infty = 0.832$). Simulation parameters: $C(x) = x^2$ and $B(x) = -0.0834x^3 + 2x^2 + x$, π given by Equation 5.21, $n = 10000$, $d = 4$, $x_0 = 0.1$, $x_m = 3$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

the $\Gamma^{\text{astrocoauth}}$ network (c); the Γ^{fb} network (d); the Γ^{p2p} network (e); and the Γ^{protein} network (f).




Figure 65: Evolution of the distribution of strategies x on model power-law networks with average degree $d = 10$ and clustering coefficient $C^{(2)}$. (a) $C^{(2)} = 0.1$ ($x_{BP} = 1.425$). (b) $C^{(2)} = 0.2$ ($x_{BP} = 1.389$). (c) $C^{(2)} = 0.3$ ($x_{BP} = 1.401$). (d) $C^{(2)} = 0.4$ ($x_{BP} = 1.419$). Simulation parameters: $C(x) = x^2$ and $B(x) = -0.0834x^3 + 2x^2 + x$, π given by Equation 5.21, $n = 10000$, $d = 10$, $x_0 = 0.1$, $x_m = 3$, $\mu = 0.01$, $\sigma = 0.005$, $T = 500000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 66: Evolution of the distribution of strategies x on model power-law networks with average degree $d = 4$ and homophily coefficient h . (a) $h = -0.1$ ($\bar{x}_\infty = 0.769$). (b) $h = 0$ ($x_{BP} = 0.785$). (c) $h = 0.1$ ($x_{BP} = 0.909$). (d) $h = 0.2$ ($x_{BP} = 0.176$). Simulation parameters: $C(x) = x^2$ and $B(x) = -0.0834x^3 + 2x^2 + x$, π given by Equation 5.21, $n = 10000$, $d = 4$, $x_0 = 0.1$, $x_m = 3$, $\mu = 0.01$, $\sigma = 0.005$, $T = 500000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.




Figure 67: Evolution of the distribution of strategies x on model power-law networks with average degree $d = 10$ and homophily coefficient h . (a) $h = -0.1$ ($x_{BP} = 1.373$). (b) $h = 0$ ($x_{BP} = 1.380$). (c) $h = 0.1$ ($x_{BP} = 1.391$). (d) $h = 0.2$ ($x_{BP} = 1.4$). Simulation parameters: $C(x) = x^2$ and $B(x) = -0.0834x^3 + 2x^2 + x$, π given by Equation 5.21, $n = 10000$, $d = 10$, $x_0 = 0.1$, $x_m = 3$, $\mu = 0.01$, $\sigma = 0.005$, $T = 500000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

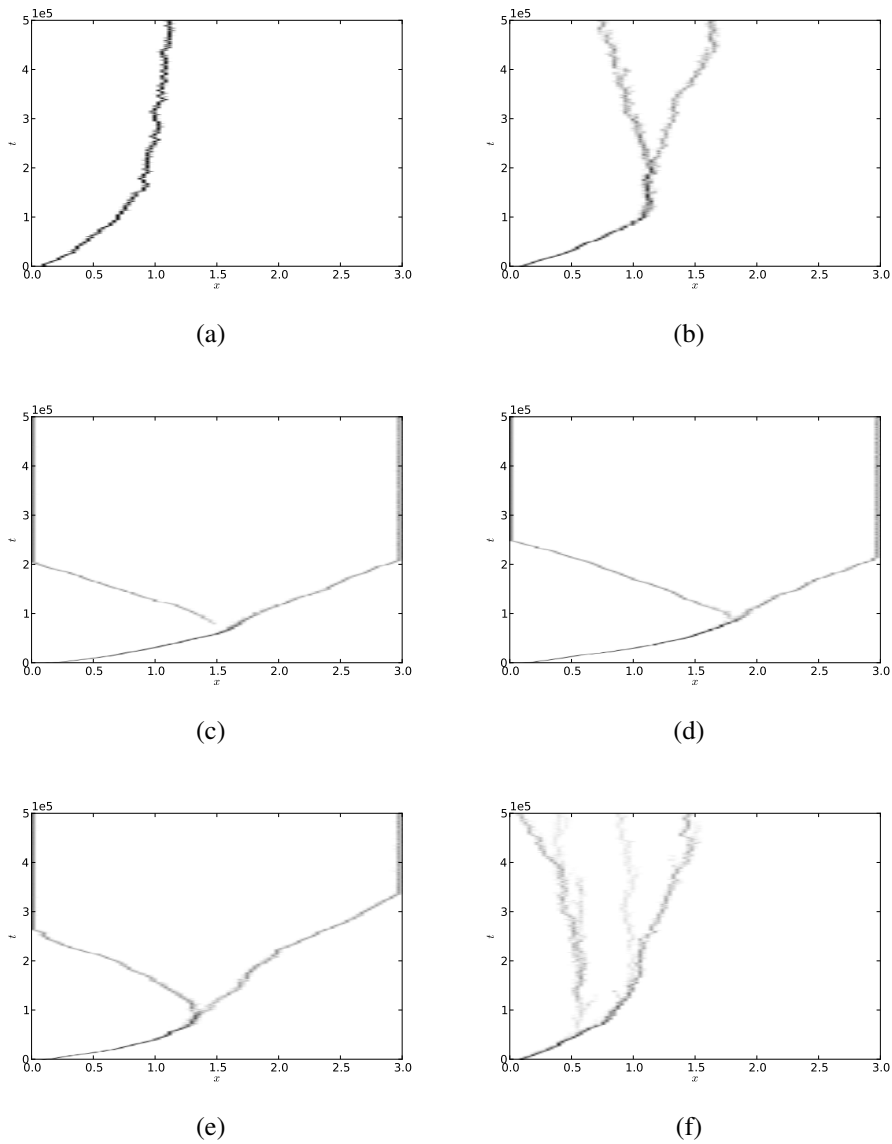


Figure 68: Evolution of the distribution of strategies x on empirical networks. See Table 2 for the basic properties of these networks. (a) the Γ^{internet} network ($\bar{x}_\infty = 0.773$). (b) the $\Gamma^{\text{hepcoauth}}$ network ($x_{BP} = 1.125$). (c) the $\Gamma^{\text{astrocoauth}}$ network ($x_{BP} = 1.5$). (d) the Γ^{fb} network ($x_{BP} = 1.8$). (e) the Γ^{p2p} network ($x_{BP} = 1.3$). (f) the Γ^{protein} network ($x_{BP} = 0.6$). Simulation parameters: $C(x) = x^2$ and $B(x) = -0.0834x^3 + 2x^2 + x$, π given by Equation 5.21, $x_0 = 0.1$, $x_m = 3$, $\mu = 0.01$, $\sigma = 0.005$, $T = 500000$, $\nu = 100$, $\mathcal{U} = \text{FE2}$, and $s = 1$.

5.7 Discussion

For the CPD game with linear cost $C(x) = cx$ and benefit $B(x) = bx$ functions, the simulation results (see Figures 38(a)-(d)) on complete networks Γ_{WM} agree very well with the values predicted by analysis on well-mixed populations. A population, starting from a state in which everyone invests a small amount $x_0 > 0$, eventually evolves to the state in which everyone invests the maximum amount $x_m = 1$ when the inequality $r > \rho$ is satisfied, where r is the degree of assortative interactions and ρ is the cost-to-benefit ratio $\frac{c}{b}$. When $r < \rho$, population evolves to the zero-investment state.

The simulation results (see Figures 39(a)-(i)) on model (random regular, Erdős-Rényi, Barabási-Albert) networks with varying average degree d suggest that the average long-term cooperative investment \bar{x}_∞ increases as d decreases. This makes sense because as the average degree of a network decreases, the network assortativity increases, and from the results on complete networks we know that cooperative investments increase with increasing assortativity. For d around 20, the networks yield results similar to complete networks with zero assortativity. This again makes sense because as the average degree of a network increases, it more and more approximates a complete network. The figures also suggest that \bar{x}_∞ decreases with increasing cost-to-benefit ratio ρ , similar to what we witnessed on complete networks. Furthermore, except for small values of d , the degree distribution of the networks has very little effect on the value of x_∞ ; when $d = 4$, the random regular network has the highest value for the cooperation index \mathcal{C} .

A comparison of the cooperation indices \mathcal{C} from the simulation results (see Figures 40(a)-(d)) on model networks with clustering coefficient $C^{(2)}$ suggests that an increase in clustering inhibits cooperative investments. The results (see Figures 41(a)-(d)) on model networks with homophily coefficient h suggests that an increase in homophily (from negative to positive) also inhibits cooperative investments.

A comparison of the cooperation indices \mathcal{C} from the simulation results (see Figures 42(a)-(f)) on empirical networks further support the fact that networks with small values of average degree d promote cooperation, while networks with high values of d inhibit cooperation. For example, the Γ^{protein} network, which has the smallest average degree ($d \approx 3.56$), has the highest values for \mathcal{C} , while the Γ^{fb} network, which has the highest average degree ($d \approx 46.51$), has the smallest values for \mathcal{C} . The six empirical networks have very different sizes, and yet the trend of \bar{x}_∞ with respect to their average degree d is comparable to the trend we observed in the case of model networks, which were all of size 10000. This suggests that the results of our models, as long as all other parameters are fixed, is independent of the population size.

For the CPD game with quadratic cost $C(x) = x^2$ and benefit $B(x) = -x^2 + 2x$ functions, the simulation results (see Figure 43(a)) on a complete network Γ_{WM} agrees very well with the values predicted by analysis on well-mixed populations. A population, starting from a state in which everyone invests a small amount $x_0 > 0$, eventually evolves to the state in which everyone invests an amount given by the singular strategy x^* , which increases with an increase in assortativity r . The simulation results (see Figures 43(b)-(d)) on model (random regular, Erdős-Rényi, Barabási-Albert) networks with varying average degree d suggest that the average long-term cooperative investment \bar{x}_∞ increases as d decreases. This is what we expect because network-induced assortativity in game interactions increases as average degree decreases. Furthermore, except for small values of d , the degree distribution of the networks has very little effect on the value of \bar{x}_∞ ; when $d = 4$, the random regular network has the highest value.

A comparison of \bar{x}_∞ values from the simulation results (see Figures 44(a)-(d)) on model networks with clustering coefficient $C^{(2)}$ is inconclusive. The results (see Figures 45(a)-(d)) on model networks with homophily coefficient h suggests that an increase in homophily (from negative to positive) inhibits cooperative investments.

A comparison of \bar{x}_∞ values from the simulation results (see Figures 46(a)-(f)) on empirical networks further support the fact that networks with small values of average degree d promote cooperation, while networks with high values of d inhibit cooperation. For example, the Γ^{protein} network, which has the smallest average degree ($d \approx 3.56$), has the highest values for \bar{x}_∞ , while the Γ^{fb} network, which has the highest average degree ($d \approx 46.51$), has the smallest values for \bar{x}_∞ . The trend of x_∞ with respect to the average degree d of the networks is comparable to the trend observed in the case of model networks of size 10000. This suggests that the results are independent of population size.

In summary, on complete networks, the cooperative investments that individuals make approaches the Hicks optimal strategy x_H , which for the parameters chosen is 0.5, as the assortativity r increases. On other networks, the x_H is approached as the average degree d of the networks decreases. Clustering coefficient $C^{(2)}$ and homophily coefficient h of networks also have an effect, albeit weaker than r or d , in driving the cooperative investments towards or away from x_H .

For the CSD game with quadratic cost $C(x) = -1.6x^2 + 4.8x$ and benefit $B(x) = -x^2 + 5x$ functions, the simulation results (see Figure 49(a)) on a complete network Γ_{WM} agrees very well with the values predicted by analysis on well-mixed populations. As the assortativity r increases, the singular strategy x^* transitions at $r = 0.2$ from being an EBP to being an ESS. So the population, starting from a state in which everyone invests a small amount $x_0 > 0$, undergoes evolutionary branching for values of $r < 0.2$, and for $r > 0.2$ evolves to the state in which everyone invests an amount given by the singular strategy x^* . The simulation results (see Figures 50(b)-(d)) on model (random regular, Erdős-Rényi, Barabási-Albert) networks with varying average degree d show that the population undergoes evolutionary branching when $d > 4$ and at $d = 4$, everyone in the population invests an amount \bar{x}_∞ . This trend makes qualitative sense because network-induced assortativity in game interactions increases as average degree decreases.

Furthermore, except for small values of d , the degree distribution of the networks has very little effect on the dynamics. For $d = 4$ at which all three networks exhibit an ESS, the random regular network has the highest value for \bar{x}_∞ .

A comparison of \bar{x}_∞ values from the simulation results (see Figures 51(a)-(d)) on model networks with average degree $d = 4$ and with clustering coefficient $C^{(2)}$ suggests that an increase in clustering inhibits cooperative investments, but do not change the nature of the singular strategy, which is an ESS throughout. The results (see Figures 52(a)-(d)) on model networks with average degree $d = 10$ and with clustering coefficient $C^{(2)}$ suggests that an increase in clustering can change the nature of a singular strategy from being an EBP to being an ESS.

The results (see Figures 53(a)-(d)) on model networks with average degree $d = 4$ and with homophily coefficient h suggests that an increase in homophily (from negative to positive) can change the nature of the singular strategy from being an EBP to being an ESS. The results (see Figures 54(a)-(d)) on model networks with average degree $d = 10$ and with homophily coefficient h suggests that an increase in homophily (from negative to positive) does not change the nature of a singular strategy, which is an EBP throughout.

A comparison of \bar{x}_∞ values from the simulation results (see Figures 55(a)-(f)) on empirical networks further support the fact that networks with small values of average degree d exhibit an ESS in general, while networks with high values of d exhibit evolutionary branching. For example, the Γ^{internet} network, which has the second smallest average degree ($d \approx 4.219$), exhibits an ESS, while the Γ^{fb} network, which has the highest average degree ($d \approx 46.51$), exhibits evolutionary branching. It is interesting to note that the Γ^{protein} network, which has the smallest value of d exhibits evolutionary branching. This is probably because of secondary effects due to clustering and homophily. The trend of x_∞ with respect to the average degree d of the networks, which have very different

sizes, is in general comparable to the trend observed in the case of model networks of size 10000. This suggests that the results are independent of population size.

In summary, on complete networks, the cooperative investments that individuals make approaches the Hicks optimal strategy x_H , which for the parameters chosen is 1, as the assortativity r increases. For small values of r , there are some individuals that make high investments, while others invest very little. On other networks, the value of x_H is approached as the average degree d of the networks decreases. Here again, at high values of d , the population exhibits evolutionary branching, with individuals that make high investments coexisting with ones that make low investments. Clustering coefficient $C^{(2)}$ and homophily coefficient h of networks also have some effect on the dynamics, albeit weaker than r or d .

For the CSD game with quadratic cost $C(x) = -1.5x^2 + 4x$ and benefit $B(x) = -0.2x^2 + 3x$ functions, the simulation results (see Figure 55(a)) on a complete network Γ_{WM} agrees very well with the values predicted by analysis on well-mixed populations. For the entire range of r values considered, the singular strategy x^* is a repeller, and this value decreases as r increases. So there is a threshold value of r below which the initial strategy x_0 satisfies $x_0 < x^*$ and the population in this case evolves to the zero-investment state. Above the threshold, $x_0 > x^*$ and the population evolves to the maximum-investment state. The simulation results (see Figures 55(b)-(d)) on model (random regular, Erdős-Rényi, Barabási-Albert) networks with varying average degree d show exactly what we expect from knowing that network-induced assortativity in game interactions increases with decreasing average degree. There is a threshold value of d below which the population evolves to the zero-investment state, and above which it evolves to the maximum investment state. Furthermore, the degree distribution of the networks has no effect on the dynamics.

A comparison of \bar{x}_∞ values from the simulation results (see Figures 56-59) on model networks with average degree $d = 4, 10$ and with clustering coefficient $C^{(2)}$, and on model networks with average degree $d = 4, 10$ and with homophily coefficient h suggests that when the singular strategy is a repeller, the dynamics is unaffected by clustering or homophily.

A comparison of \bar{x}_∞ values from the simulation results (see Figures 60(a)-(f)) on empirical networks show that if the initial strategy x_0 has the right value, then on networks with smaller average degree d , the singular strategy x^* is more likely to satisfy $x_0 > x^*$, and the population in this case evolves to the maximum-investment state. On the other hand, for networks with higher values of d , the value of x^* is likely to be higher than the value of x_0 and the population in this case evolves to the zero-investment state. For example, with $x_0 = 0.3$, the Γ^{protein} network, which has the smallest average degree ($d \approx 3.56$), the population evolves to the maximum-investment state, while on the Γ^{fb} network, which has the highest average degree ($d \approx 46.51$), the population evolves to the zero-investment state. The trend of x_∞ with respect to the average degree d of the networks, which have very different sizes, is in general comparable to the trend observed in the case of model networks of size 10000. This suggests that the results are independent of population size.

In summary, on complete networks, the individuals switch from making zero investments to making maximum (Hicks optimal) investments beyond a certain threshold value of assortativity. On other networks, a similar trend is observed with decreasing average degree d . Clustering coefficient $C^{(2)}$ and homophily coefficient h of networks seem to have no effect when the singular strategy is a repeller.

For the CTOC game with quadratic cost $C(x) = x^2$ and cubic benefit $B(x) = -0.0834x^3 + 2x^2 + x$ functions, the simulation results (see Figure 63(a)) on a complete network Γ_{WM} agrees very well with the values predicted by analysis on

well-mixed populations. As the assortativity r increases, the singular strategy x^* decreases and transitions at $r = 0.3$ from being an EBP to being an ESS. So the population, starting from a state in which everyone invests a small amount $x_0 > 0$, undergoes evolutionary branching for values of $r < 0.3$, and for $r > 0.3$ evolves to the state in which everyone invests an amount given by the singular strategy x^* . The simulation results (see Figures 63(b)-(d)) on model (random regular, Erdős-Rényi, Barabási-Albert) networks with varying average degree d show that except for Barabási-Albert networks, the population undergoes evolutionary branching for all values of d . On Barabási-Albert networks, when $d = 4$, the population exhibits an ESS. Though network-induced assortativity in game interactions increases with decreasing average degree, the level of assortativity isn't sufficient on networks other than Barabási-Albert networks to change the nature of the singular strategy from an EBP to an ESS.

A comparison of \bar{x}_∞ values from the simulation results (see Figures 64(a)-(d)) on model networks with average degree $d = 4$ and with clustering coefficient $C^{(2)}$ is inconclusive. The results (see Figures 65(a)-(d)) on model networks with average degree $d = 10$ and with clustering coefficient $C^{(2)}$ suggests that an increase in clustering has no effect on the nature of the singular strategy; it is an EBP throughout.

The results (see Figures 66(a)-(d)) on model networks with average degree $d = 4$ and with homophily coefficient h suggests that an increase in homophily (from negative to positive) changes the nature of the singular strategy from being an ESS to being an EBP. The results (see Figures 67(a)-(d)) on model networks with average degree $d = 10$ and with homophily coefficient h suggests that an increase in homophily (from negative to positive) does not change the nature of a singular strategy that is an EBP.

A comparison of \bar{x}_∞ values from the simulation results (see Figures 68(a)-(f)) on empirical networks further support the fact that networks with small values of average

degree d exhibit an ESS in general, while networks with high values of d exhibit evolutionary branching. For example, the Γ^{internet} network, which has the second smallest average degree ($d \approx 4.219$), exhibits an ESS, while the Γ^{fb} network, which has the highest average degree ($d \approx 46.51$), exhibits evolutionary branching. Here again, it is interesting to note that the Γ^{protein} network, which has the smallest value of d exhibits evolutionary branching. This is probably because of secondary effects due to clustering and homophily. The trend of x_∞ with respect to the average degree d of the networks, which have very different sizes, is in general comparable to the trend observed in the case of model networks of size 10000. This suggests that the results are independent of population size.

In summary, on complete networks, the cooperative investments that individuals make approaches the Hicks optimal strategy x_H , which for the parameters chosen is 0.25, as the assortativity r increases. For small values of r , there are some individuals that make low cooperative investments (i.e., overexploit the common-pool resource), while others make very high cooperative investments (i.e., consume less of the common-pool resource). On Barabási-Albert networks, x_H is approached as the average degree d of the networks decreases. On networks with high values of d , the population exhibits evolutionary branching, with some individuals overexploiting the common-pool resource and others consuming less of it. Clustering coefficient $C^{(2)}$ and homophily coefficient h of networks also have an effect on the dynamics, albeit weaker than r or d .

When the CPD, CSD, and CTOC games are played on Erdős-Rényi networks Γ_{ER} having small average degree d and on the six empirical networks, there is a fraction of individuals that never evolve away from the initial monomorphic strategy x_0 . The reason for this is that the networks consist of vertices with degree less than two, and the update rule $\mathcal{U} = \text{FE2}$ that we used in our simulations excludes such vertices from the interaction and update rounds. The strategies of the individuals corresponding to these vertices have been suppressed in the results.

	FE1	FE2	RE1	RE2	BD
CPD game	0.331	0.336	0.335	0.334	0.334
CSD game	0.584	0.549	0.472	0.523	0.582
CTOC game	1.294	1.345	1.257	1.288	1.201

Table 5: Effect of update rule on continuous-game dynamics on a complete network. The value for each (game, update rule) pair indicates the long-term value \bar{x}_∞ of mean strategy \bar{x} for the CPD game, and strategy x_{BP} at which the population undergoes evolutionary branching for the CSD and CTOC games. Simulation parameters: (CPD game) $r = 0.5$, $C(x) = x^2$ and $B(x) = -x^2 + 2x$, π given by Equation 5.5, $n = 10000$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $\nu = 100$, and $s = 1$; (CSD game) $r = 0.3$, $C(x) = -1.6x^2 + 4.8x$ and $B(x) = -x^2 + 5x$, π given by Equation 5.12, $n = 10000$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, and $s = 1$; (CTOC game) $r = 0.1$, $C(x) = x^2$ and $B(x) = -0.0834x^3 + 2x^2 + x$, π given by Equation 5.21, $n = 10000$, $x_0 = 0.1$, $x_m = 3$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, and $s = 1$.

While the continuous game dynamics on complete networks is independent of the update rule \mathcal{U} , it does depend on \mathcal{U} when the games are played on other networks. To illustrate this, we simulated the CPD, CSD, and CTOC games on a complete network Γ_{WM} and a Barabási-Albert Γ_{BA} network, using five different update rules $\mathcal{U} = \text{FE1, FE2, RE1, RE2, and BD}$. Table 5 shows the long-term value \bar{x}_∞ of mean strategy \bar{x} for the CPD game, and the evolutionary branching point x_{BP} for the CSD and CTOC games for each of the (game, update rule) pairs when the games are played on a Γ_{WM} network. Table 4 shows the same when the games are played on a Γ_{BA} network. On a Γ_{WM} network, the game dynamics is relatively independent of the choice of update rule. However, on a Γ_{BA} network, the dynamics of all three games is significantly affected by the update rule.

	FE1	FE2	RE1	RE2	BD
CPD game	0.013	0.192	0.009	0.240	0.101
CSD game	0.186	0.576	0.118	0.481	0.400
CTOC game	2.011	1.688	1.906	1.500	1.625

Table 6: Effect of update rule on continuous-game dynamics on a Barabási-Albert network. The value for each (game, update rule) pair indicates the long-term value \bar{x}_∞ of mean strategy \bar{x} for the CPD game, and strategy x_{BP} at which the population undergoes evolutionary branching for the CSD and CTOC games. Simulation parameters: (CPD game) $d = 4$, $C(x) = x^2$ and $B(x) = -x^2 + 2x$, π given by Equation 5.5, $n = 10000$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 200000$, $\nu = 100$, and $s = 1$; (CSD game) $d = 10$, $C(x) = -1.6x^2 + 4.8x$ and $B(x) = -x^2 + 5x$, π given by Equation 5.12, $n = 10000$, $x_0 = 0.1$, $x_m = 1$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, and $s = 1$; (CTOC game) $d = 20$, $C(x) = x^2$ and $B(x) = -0.0834x^3 + 2x^2 + x$, π given by Equation 5.21, $n = 10000$, $x_0 = 0.1$, $x_m = 3$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $\nu = 100$, and $s = 1$.

5.8 Future Directions

In this section we suggest two new lines of inquiry, along with an outline of how we might want to approach them.

1. We analyzed our continuous game models on well-mixed populations, using the framework of adaptive dynamics [GKM97, MB07], suitably modified to take into account assortative interactions. Since the assortativity on a network bears an inverse relationship to its average degree d , knowing how r , the degree of assortative interactions in well-mixed populations, affects the game dynamics gives us a qualitative sense of how the dynamics would behave on a network.

It would be nice to get a quantitative idea of how the game dynamics on a network is affected by its structural properties. To do this, we need to develop a framework of adaptive dynamics for networks, and this is how we would approach the problem. Following the procedure highlighted in Section 4.8 of Chapter 4 we can derive the replicator equation for a regular network Γ with average degree d , taking into account the update rule \mathcal{U} for the evolutionary dynamics; recall that the evolutionary dynamics on networks depends on the update rule. Once we have the replicator equation, we can obtain the invasion fitness $f_y(x)$ of a rare mutant strategy y , which will then include the details of the network Γ and the update rule \mathcal{U} . With the expression for the invasion fitness of the mutant strategy in hand, we can carry out the program of adaptive dynamics as usual to figure out the evolutionary fate of the mutant strategy. Though the analysis is tractable only for simple networks such as Bethe lattices or Cayley trees [Bet35], it should serve as a good approximation for an arbitrary network with average degree d provided the probability of short loops in the network becomes negligible as the network size n increases.

2. The three continuous games we explored in this chapter only considered pairwise interactions, i.e., interactions among pairs of individuals in a population. One can easily think of situations in which groups of individuals engage in a cooperative act among each other (CPD game), make cooperative investments toward a common good (CSD game), or collectively consume a limited common-pool resource (CTOC game). Such situations involving groups of individuals can be described by generalizing the CPD, CSD, and CTOC games to groups of $N > 2$ individuals making investments x_1, \dots, x_N . In case of a well-mixed population, we might interpret the degree r of assortative interactions as follows: with probability r an individual interacts with $N - 1$ other individuals of its own type, and with probability $1 - r$ it interacts with $N - 1$ individuals randomly picked from the population. The invasion fitness $f_y^r(x)$ of a rare mutant strategy y is thus given by

$$f_y^r(x) = r\pi(y; \underbrace{y, \dots, y}_{N-1 \text{ terms}}) + (1 - r)\pi(y; \underbrace{x, \dots, x}_{N-1 \text{ terms}}) - \pi(x; \underbrace{x, \dots, x}_{N-1 \text{ terms}}). \quad (5.29)$$

Note that when $N = 2$, the equation reduces to the invasion fitness given by Equation 5.3.

The payoff $\pi(x_i; \underline{x}_i)$ to the i th individual playing strategy x_i , for the CPD, CSD, and CTOC games, is respectively given by

$$\pi(x_i; \underline{x}_i) = B(x_1 + \dots + x_{i-1} + x_{i+1} + \dots + x_N) - C(x_i), \quad (5.30)$$

$$\pi(x_i; \underline{x}_i) = B(x_1 + \dots + x_N) - C(x_i), \text{ and}$$

$$\pi(x_i; \underline{x}_i) = B(x_i) - C(x_1 + \dots + x_N),$$

where $C(x)$ and $B(x)$ are the cost and benefit functions for the games, and $\underline{x}_i = x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_N$ denotes the strategies of the individuals not including the i th individual. Note the when $N = 2$, the payoff functions reduce to Equations 5.21, 5.12, and 5.21 respectively for the pairwise CPD, CSD, and CTOC games.

It is worth noting that the N -player CPD game is a generalization of the N -player public goods game [Bin92, FG00], used extensively in economics to study cooperative behavior, and in which the payoff $\pi(x_i; \underline{x}_i)$ to the i th individual contributing an amount x_i toward a public good, is given by

$$\pi(x_i; \underline{x}_i) = \frac{k}{N} \left(x_1 + \cdots + x_{i-1} + x_{i+1} + \cdots + x_N \right) - \left(1 - \frac{k}{N} \right) x_i, \quad (5.31)$$

where $k < N$ is the factor by which the contributions are multiplied before the public good payoff is evenly divided among the players, and $\underline{x}_i = x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_N$ denotes the contributions made by the individuals not including the i th individual.

With the invasion fitness of the mutant strategy y and payoff functions defined as above, the analysis of the games involving groups of $N > 2$ interacting individuals can be carried out as usual using the framework of adaptive dynamics.

When the games are played on a network Γ , the notion of a group is naturally imposed by the network structure. The group to which an individual belongs is determined by the individual's neighborhood. For example, the payoff $\pi(x_i; \underline{x}_i)$ of the i th individual playing strategy x_i , for the CPD, CSD, and CTOC games, are respectively given by

$$\pi(x_i; \underline{x}_i) = B(x_1 + \cdots + x_{k_i}) - C(x_i), \quad (5.32)$$

$$\pi(x_i; \underline{x}_i) = B(x_1 + \cdots + x_{k_i} + x_i) - C(x_i), \text{ and}$$

$$\pi(x_i; \underline{x}_i) = B(x_i) - C(x_1 + \cdots + x_{k_i} + x_i),$$

where $C(x)$ and $B(x)$ are the cost and benefit functions for the games, and $\underline{x}_i = x_1, \dots, x_{k_i}$ denote the strategies of the k_i neighbors of i .

CHAPTER 6

DISEASE DYNAMICS

Prevention is better than cure.

- Desiderius Erasmus

6.1 Introduction

In this chapter we turn our attention to yet another kind of dynamical process on complex networks, namely the spreading of infectious diseases in a population. Mathematical modeling of epidemics is an active area of research spanning multiple disciplines. Epidemiologists, computer scientists, and social scientists share a common interest in studying spreading phenomena and use similar models to describe the diffusion of viruses, knowledge, and innovation [KSC97, Kle07, LM01].

In this chapter we first look at compartmental models from epidemiology—the study of the patterns, causes, and effects of health and disease conditions in populations—in which the within-host dynamics of a disease is represented as changes between a few disease states or compartments, such as *susceptible* (S), *infected* (I), *recovered* (R), and so forth [AM92, Bai75, DH00]. The individual-level disease parameters such as the transmission rate of a disease are fixed in such models, and one is primarily interested in the change of the number of infected individuals in the population as a function of time. The main aim in these models is therefore to understand the properties of epidemics at equilibrium, the existence of a non-zero density of infected individuals, the presence or

absence of a global outbreak, etc. These models assume a fully-mixed population in which every individual has an equal chance of coming into contact with every other individual.

In reality, most diseases spread over networks of contacts between individuals. For example, contagious diseases and parasites are communicated when people come into close contact, and HIV and other sexually transmitted diseases (STDs) are communicated when people have sex. Such diseases are better understood using epidemiological models that take into account the connection patterns among the individuals in a population. The introduction of network structure into the models leads to important new results concerning the basic properties of disease spreading processes [AM92, BPV03, MA84, ML01, PV01b].

We propose a stochastic, discrete-time, evolutionary variant of the susceptible-infected-recovered (SIR) compartmental model to study disease dynamics on complex networks. We focus on the evolutionary dynamics of disease-causing parasites—anything that lives and multiplies inside another organism called the “host” and usually causes harm. We assume that the host does not evolve on the time scale under consideration, which is a reasonable assumption, since parasites generally evolve much faster than their hosts. The model is individual based and focuses on how the disease-causing pathogen evolves in virulence [Ewa93, Now06a] in response to the selection pressures acting on it. In addition to studying the evolution of virulence, we use our model to evaluate the effectiveness of different vaccination strategies in combating a disease [AM85, DB02, PV02b]. We analyze our model for a complete network and describe how the model can be implemented on a computer for in-silico exploration of disease dynamics. We present results from agent-based simulations on model and empirical networks, and provide a general discussion of the results. We conclude the chapter with suggestions for future inquiry.

6.2 Models of Disease Dynamics

The biological details of what goes on when a host individual contracts an infection is rather complicated. The pathogen responsible for the infection typically multiplies within the body while the immune system tries to combat it. If we want to understand fully how diseases spread through populations we need to take all of these details into account, which is practically infeasible. Fortunately, there are more tractable approaches based on simplified models which provide a good insight on disease dynamics in many cases and it is on such models that we focus in this section.

Mathematical modeling of epidemics was first undertaken by Anderson McKendrick, a doctor and amateur mathematician who made seminal contributions to the field in the early twentieth century. The theories that he and others developed form the cornerstone of classical mathematical epidemiology. The books by Bailey [Bai75], Anderson and May [AM92], and the review article by Hethcote [Het00] are excellent resources on the subject.

Classical epidemiology [AM92, Bai75, DH00] avoids discussing contact network—networks through which diseases typically spread—by making use of a *fully mixed* or *mass-action* approximation, in which it is assumed that every individual has an equal chance, per unit time, of coming into contact with every other, i.e., people meet completely at random. In the real world, people have contact only with a small fraction of the population of the world, and that fraction is not chosen at random, which is where networks come into the picture. A familiarity with the classical models is still useful, since the model we propose to study disease dynamics is based on these classical models.

Before we discuss the models, we must describe an important metric that is used in the study of epidemics, namely the *basic reproduction number* R_0 [Bai75], which is defined as follows. Consider the spread of a disease when it is just starting out, i.e., when there are only a few cases of the disease and the rest of the population is susceptible—a *naive*

Disease	Transmission	R_0
Measles	Airborne	12 - 18
Pertussis	Airborne droplet	12 - 17
Diphtheria	Saliva	6 - 7
Smallpox	Social contact	5 - 7
Polio	Fecal-oral route	5 - 7
Rubella	Airborne droplet	5 - 7
Mumps	Airborne droplet	4 - 7
HIV/AIDS	Sexual contact	2 - 5
SARS	Airborne droplet	2 - 5 [WT04]
Influenza (1918 pandemic strain)	Airborne droplet	2 - 3 [MRL04]

Table 7: Values of R_0 for well-known infectious diseases. Unless otherwise noted, the R_0 values are from [CW].

population in epidemiology jargon—and consider a susceptible individual who catches the disease in this early stage of the outbreak. The basic reproduction number is defined to be the average number of additional individuals that such an individual passes the disease onto before she recovers. In general, it is possible to show that the extinction probability of an epidemic starting with n infected individuals is equal to R_0^{-n} [Bai75]. For example, in the case of a single infected individual, even for values of R_0 as high as 2 the outbreak probability is just 50%. Table 7 lists the value of R_0 for various known diseases.

The point $R_0 = 1$ which separates the growing and shrinking behaviors of a disease marks the *epidemic threshold* between regimes in which the disease either multiplies or dies out. The epidemic threshold is a central concept in epidemiology since it provides a

reference frame to understand the evolution of contagion processes and to focus on policies aimed at an effective reduction of the basic reproductive number in order to stop epidemic outbreaks. Generally, the larger the value of R_0 , the harder it is to control the epidemic. Assuming a well-mixed population, the proportion v of the population that needs to be vaccinated to eradicate the disease is given by

$$v \geq v^* = 1 - \frac{1}{R_0}. \quad (6.1)$$

If $v < v^*$, the disease will remain endemic, although at lower levels than in the case of absence of vaccinations. For example, to prevent a disease with $R_0 = 5$ from spreading, at least 80% of the population must be vaccinated.

6.2.1 SI Model

The simplest model in epidemiology is the susceptible-infected (SI) model in which the disease dynamics is represented as changes between two disease states, namely *susceptible* (S) and *infected* (I) [AM92, Bai75, DH00]. An individual in the susceptible state is someone who does not have the disease yet but could catch it if she comes into contact with someone who does. An individual in the infected state is someone who has the disease and can potentially pass it on if she comes into contact with a susceptible individual. Figure 69 shows the allowed transitions between states in the SI model.




Figure 69: Allowed transitions between states in the SI model.

Consider a disease spreading through a population of n individuals. Let $S(t)$ be the number of individuals who are susceptible at time t and let $X(t)$ be the number of infected

individuals. Since the disease-spreading process is a random one, these numbers are not uniquely determined. To circumvent this problem, S and X are interpreted as the average or expected number of susceptible and infected individuals.

Suppose that people meet and make contacts sufficiently to result in the spread of disease entirely at random with a per-individual transmission rate β , meaning that each infected individual, on average, transmits the disease to β randomly chosen individuals per unit time. The disease is transmitted when an infected individual comes into contact with a susceptible one. If the total population consists of n individuals, then the average probability of an individual being susceptible is $\frac{S}{n}$, and hence an infected person transmits the disease to $\frac{\beta S}{n}$ susceptible individuals per unit time. Since there are on average X infected individuals in the population, the overall average rate of new infections is $\frac{\beta SX}{n}$ and we can write a differential equation¹ describing the rate of change of X as

$$\dot{X} = \beta \frac{SX}{n}. \quad (6.2)$$

At the same time the number of susceptible individuals goes down at the same rate,

$$\dot{S} = -\beta \frac{SX}{n}. \quad (6.3)$$

It is often convenient to define variables representing the fractions of susceptibles and infected individuals as:

$$s = \frac{S}{n}, \quad x = \frac{X}{n}, \quad (6.4)$$

in terms of which 6.2 and 6.3 can be written as

$$\begin{aligned} \dot{s} &= -\beta sx \\ \dot{x} &= \beta sx. \end{aligned} \quad (6.5)$$

¹We use the notation \dot{X} to denote the derivative of X with respect to time.

Since $S + X = n$, or equivalently $s + x = 1$, we do not need both of these equations.

Eliminating s from the equations by writing $s = 1 - x$, we have

$$\dot{x} = \beta(1 - x)x, \tag{6.6}$$

which is just the *logistic growth equation* (see A.3) with the solution

$$x(t) = \frac{x_0 e^{\beta t}}{1 - x_0 + x_0 e^{\beta t}}, \tag{6.7}$$

where $x_0 = x(0)$. This equation produces an S-shaped “logistic growth curve” for the fraction of infected individuals, as the following example illustrates.

Example 6.1. Consider the SI model of disease spread through a population in which 1% of the population is infected at the beginning and the rest of the individuals are in the susceptible state, i.e., $x(0) = 0.01$ and $s(0) = 0.99$. Let $\beta = 0.8$. Figure 70 shows the trajectories for the system, i.e., the evolution of the fractions s, x of the population in the two disease states. The x curve increases exponentially for a short time, corresponding to




Figure 70: The trajectories for the SI model. $s(0) = 0.99, x(0) = 0.01$, and $\beta = 0.8$.

the initial phase in which most of the population is susceptible. The curve then saturates

as the number of susceptibles declines and the disease has a harder and harder time finding new victims. □

For the SI model, R_0 is formally infinite since an infected individual can infect an arbitrary number of others. For a population of finite size, however, the empirical value of R_0 is finite.

6.2.2 SIR Model

One way in which the simple SI model can be extended is by introducing recovery from disease. In the SI model, individuals, once infected, are infected forever. For many real diseases, however, people recover from infections after a certain time because their immune system fights off the disease-causing agent. Furthermore, people often retain their immunity to the disease after such a recovery so that they cannot catch it again. To represent this behavior in our model we introduce a third disease state, usually denoted R for recovered. The corresponding three-state model is called the *susceptible-infected-recovered* or SIR model [AM92, Bai75, DH00].

The R state in the SIR model can also represent death, and thus accommodate diseases with mixed outcomes where people sometimes recover and sometimes die—from a mathematical standpoint we do not care whether the individuals in the R state are recovered or dead, since either way they are effectively removed from the pool of potential hosts for the disease. For this reason the model is also called *susceptible-infected-removed* model.

The dynamics of the fully-mixed SIR model has two stages. In the first stage, susceptible individuals become infected when they come in contact with infected individuals; the transmission rate of infected individuals is β as in the SI model. In the second stage, infected individuals recover, with γ denoting the probability of an infected

individual recovering per unit time. Figure 71 shows the allowed transitions between states in the SIR model.



Figure 71: Allowed transitions between states in the SIR model.

The differential equations for the SIR model in terms of the fractions s, x, r of the individuals in the three states are

$$\dot{s} = -\beta sx, \quad (6.8)$$

$$\dot{x} = \beta sx - \gamma x, \quad (6.9)$$

$$\dot{r} = \gamma x, \quad (6.10)$$

and in addition the three variables satisfy

$$s + x + r = 1. \quad (6.11)$$

To solve these equations we eliminate x between Equation 6.8 and Equation 6.10, giving

$$\frac{1}{s}\dot{s} = -\frac{\beta}{\gamma}\dot{r}, \quad (6.12)$$

and then integrate both sides with respect to t to get:

$$s = s_0 e^{-\frac{\beta r}{\gamma}}, \quad (6.13)$$

where $s_0 = s(0)$ and we have chosen the constant of integration so that there are no individuals in the recovered state at $t = 0$.

Putting $x = 1 - s - r$ in Equation 6.10 and using Equation 6.13, we get

$$\dot{r} = \gamma(1 - r - s_0 e^{-\beta r/\gamma}). \quad (6.14)$$

If we can solve this equation for r then we can find s from Equation 6.13 and x from Equation 6.11. Unfortunately, we cannot find a closed-form solution for the equation, and have to resort to numerical means instead.

Example 6.2. Consider the SIR model of disease spread through a population in which 1% of the population is infected at the beginning and the rest of population is susceptible to the disease. Let $\beta = 0.8$ and $\gamma = 0.2$. We can solve the differential equations for the SIR model numerically, and the trajectories are shown in Figure 72. The fraction of




Figure 72: The trajectories for the SIR model. $s(0) = 0.99, x(0) = 0.01, r(0) = 0, \beta = 0.8,$ and $\gamma = 0.2$.

susceptibles in the population decreases monotonically as they are infected and the fraction of recovered individuals increases monotonically. The fraction infected, however, goes up at first as people get infected, then down again as they recover, and eventually goes to zero as $t \rightarrow \infty$. Note that the number of susceptibles does not go to zero, a close inspection of which shows that the curve for $s(t)$ ends a little above the axis. This is because when $x \rightarrow 0$ there are no infected individuals left to infect the remaining susceptibles. Any individuals who survive to late enough times without being infected

will probably never get the disease at all. Similarly the fraction of recovered individuals does not quite reach one as $t \rightarrow \infty$. □

We can calculate R_0 for the SIR model as follows. If an individual remains infectious for a time τ then the expected number of others she will have contact with during that time is $\beta\tau$. The definition for R_0 is for a naive population in which all of the individuals one has contact with are susceptible, and hence $\beta\tau$ is also the total number of individuals the infected individual will infect.

Given the value of γ we can calculate the length of time τ that an infected individual is likely to remain infected before they recover. The probability of recovering in any time interval $\delta\tau$ is $\gamma\delta\tau$, and the probability of not doing so is $1 - \gamma\delta\tau$. Thus the probability that the individual is still infected after a total time τ is

$$\lim_{\delta\tau \rightarrow 0} (1 - \gamma\delta\tau)^{\tau/\delta\tau} = e^{-\gamma\tau}, \quad (6.15)$$

and the probability $p(\tau)d\tau$ that the individual remains infected this long and then recovers in the interval between τ and $\tau + \delta\tau$ is the quantity times $\gamma d\tau$:

$$p(\tau)d\tau = \gamma e^{-\gamma\tau} d\tau, \quad (6.16)$$

which is a standard exponential distribution. We average $\beta\tau$ over the distribution of τ to get the average number R_0 as

$$R_0 = \beta\gamma \int_0^{\infty} \tau e^{-\gamma\tau} d\tau = \frac{\beta}{\gamma}. \quad (6.17)$$

The epidemic threshold for the SIR model is at $R_0 = 1$, which corresponds to the point $\beta = \gamma$. Note that if $\gamma = 0$, then R_0 is infinite, which makes sense because the model in this case reduces to the SI model for which R_0 is indeed infinite.

6.2.3 SIR Model with Births and Deaths

We next consider a variant of the SIR model in which we introduce births and deaths. In this model, an individual in any disease state can die of natural causes. In addition, an infected individual can die because of the virulence induced by the disease [AM92, Bai75, DH00]. When an individual dies they are immediately replaced by a new-born, susceptible individual. We let u denote the probability per unit time of an individual dying of natural causes, and c denote the disease-induced death probability per individual per unit time. The allowed transitions for this model is shown in Figure 73.




Figure 73: Allowed transitions between states in the SIR model with births and deaths.

The differential equations for the model in terms of the fractions s, x, r of the individuals in the three states are

$$\dot{s} = -\beta sx + (u+c)x + ur, \quad (6.18)$$

$$\dot{x} = \beta sx - (\gamma + u + c)x, \quad (6.19)$$

$$\dot{r} = \gamma x - ur, \quad (6.20)$$

and the three variables as usual satisfy the relation $s + x + r = 1$. These equations can be solved numerically to obtain the evolution of the fractions of susceptible, infected, and recovered individuals.

Example 6.3. Consider the SIR model with births and deaths of disease spread through a population in which 1% of the population is infected at the beginning and the rest of population is susceptible to the disease. Let $\beta = 0.8, \gamma = 0.2, u = 0.01$, and $c = 0.11$. We can evaluate the differential equations for the model numerically, and the trajectories are shown in Figure 74. Notice how, unlike the standard SIR model, the curves oscillate




Figure 74: The trajectories for the SIR model with births and deaths. $s(0) = 0.99, x(0) = 0.01, r(0) = 0, \beta = 0.8, c = 0.11, \gamma = 0.2$, and $u = 0.01$.

before reaching the equilibrium values. Moreover, the fractions of susceptible and recovered individuals at equilibrium are respectively higher and lower than those in the standard SIR model. □

We can calculate R_0 for the model as follows. The average lifetime τ of an infected individual is $\frac{1}{\gamma+u+c}$, since the probability of an individual leaving the infected state per

unit time is $\gamma + u + c$. The rate at which one infected individual produces new infections is βs . The product of these two quantities is the average number of new infectious caused by a single infected individual in her lifetime if there are s uninfected individuals. The equilibrium abundance of uninfected individuals prior to the arrival of the infection consists of the entire population. Therefore,

$$R_0 = \frac{\beta}{\gamma + u + c}. \quad (6.21)$$

Note that when $u = c = 0$, we recover the R_0 for the standard SIR model without births and deaths.

We can understand parasite evolution in this model by studying the epidemiological dynamics of at least two parasite strains S_1 and S_2 competing for the same host [ALM94, Ewa93, LM94, Now06a]. Let β_1 and β_2 be the transmission rates of the two strains, let c_1 and c_2 be their virulences, and let x_1 and x_2 be the fractions of individuals infected by the two strains. The differential equations for the model in terms of the fractions s, x_1, x_2, r of individuals in the three states are

$$\dot{s} = (u + c_1 - \beta_1 s)x_1 + (u + c_2 - \beta_2 s)x_2 + ur, \quad (6.22)$$

$$\dot{x}_1 = (\beta_1 s - \gamma - u - c_1)x_1, \quad (6.23)$$

$$\dot{x}_2 = (\beta_2 s - \gamma - u - c_2)x_2, \quad (6.24)$$

$$\dot{r} = \gamma x_1 + \gamma x_2 - ur, \quad (6.25)$$

The basic reproductive numbers R_0^1 and R_0^2 for the two strains, using Equation 6.21, are given by

$$R_0^1 = \frac{\beta_1}{\gamma + u + c_1}, \quad (6.26)$$

and

$$R_0^2 = \frac{\beta_2}{\gamma + u + c_2}. \quad (6.27)$$

At equilibrium, the time derivatives $\dot{s}, \dot{x}_1, \dot{x}_2$, and \dot{r} must be zero. Furthermore, stable coexistence between S_1 and S_2 requires that both x_1 and x_2 be positive at equilibrium. From $\dot{x}_1 = 0$ and $x_1 > 0$, we obtain $s = \frac{\gamma+u+c_1}{\beta_1}$. But from $\dot{x}_2 = 0$ and $x_2 > 0$, we obtain $s = \frac{\gamma+u+c_2}{\beta_2}$. Both conditions can hold simultaneously only if $R_0^1 = R_0^2$. Generally, however, we expect $R_0^1 \neq R_0^2$, in which case coexistence is not possible. The evolutionary fate of the two strains is described by one of the following cases:

- If $R_0^1 < 1$ and $R_0^2 < 1$, then the only stable equilibrium is the uninfected population, since $\dot{x}_1 < 0$ and $\dot{x}_2 < 0$.
- If $R_0^1 > 1 > R_0^2$, then S_2 becomes extinct and S_1 survives, since $\dot{x}_2 < 0$ and $\dot{x}_1 > 0$.
- If $R_0^1 < 1 < R_0^2$, then S_1 becomes extinct and S_2 survives, since $\dot{x}_1 < 0$ and $\dot{x}_2 > 0$.
- If $R_0^1 > 1$ and $R_0^2 > 1$, then the strain with the higher basic reproductive number will outcompete the strain with the lower basic reproductive number. For example, if $R_0^2 > R_0^1$, then all infected individuals will eventually carry S_2 , while S_1 will become extinct.

Evolution thus tends to maximize R_0 . If there is no constraint between the transmission rate β and virulence c , then the evolutionary dynamics will increase β and reduce c . This represents the conventional wisdom that infectious diseases will evolve to become less virulent. In general, however, we expect an association between β and c ; usually the harm done to hosts (c) is associated with the production of transmission stages (β). For certain functional relations between β and c there is an evolutionarily stable degree of virulence c^* , corresponding to the maximum value of R_0 . Other situations allow evolution toward the extreme values of very high or low virulences. The detailed dynamics depend on the shape of β as a function of c . For example, if $\beta(c) = \beta_0 c^\alpha$, where $\beta_0, \alpha \in (0, 1)$, the maximum R_0 is achieved at an intermediate optimum level of

virulence given by $c^* = \frac{\alpha(\gamma+u)}{1-\alpha}$. We will restrict our attention to this class of saturating functions for $\beta(c)$, since it has a nice biological interpretation based on the law of diminishing returns [LHN95]. Figure 75 shows how R_0 varies with c for $\gamma = 0.2$, $u = 0.01$, and $\alpha = 0.5$. The dashed line and the inset in the plot show the value of c^* , i.e., the value of virulence that maximizes R_0 .




Figure 75: The R_0 curve for the SIR model with births and deaths. $\beta(c) = \sqrt{c}$, $\gamma = 0.2$, and $u = 0.01$.

6.2.4 Our Model

We are now ready to present our own model for studying the evolutionary dynamics of diseases. The model is based on the SIR model with births and deaths that we encountered in the previous section, and as a result shares the same epidemiological parameters with that model. However, the model is different from the equation-based models in the following fundamental ways:

- Unlike the equation-based approaches, it models disease dynamics at the level of individuals in the population.
- It does not assume a well-mixed population. It instead represents a population as a network in which the vertices correspond to the individuals in the population and the edges correspond to the connection patterns among the individuals. Our model can thus be used to study diseases that spread over contact networks.
- It is a discrete-time, stochastic model, and hence the parameters that were interpreted as rates in the deterministic, equation-based models have to be interpreted as probabilities. In the limit where the time interval between generations tends to zero, the model well approximates the continuous-time model.
- Unlike the equation-based approaches which fix the epidemiological parameters and simply focus on the time evolution of the number of individuals in the different disease states, the main purpose of our model is to understand the evolution of the virulence of a disease.
- Our model employs ideas from evolutionary game theory [Smi74] to study the evolutionary dynamics of a disease. A population is a collection of individuals, each in one of the disease states under consideration. We assume that each infected host individual carries a single strain of the disease, and collectively refer to the pathogens within a single host as a player. The strategy of a player is its transmission probability. The payoff or fitness of a player is its virulence, which is determined using the payoff function that relates the transmission probability to the virulence. Finally, the long-term value of the mean virulence in the population corresponds to an evolutionarily stable strategy (ESS).

We now proceed to describe our model. Let Γ denote a contact network of n vertices through which a disease spreads. The vertices, labeled $i = 1, \dots, n$, represent the individuals in a population and the edges represent the connection patterns among the individuals. Since our model is based on the SIR model, each vertex is either a susceptible (S), infected (I), or a recovered (R) vertex. An individual can in addition be in a vaccinated (V) state if the individual has been vaccinated for the disease².

Recall that the neighbors of a vertex $i \in \Gamma$ is denoted by $N(i)$. We denote the set of neighbors of i that are in state X by $N_X(i)$, and the number of such neighbors by $n_X(i)$, i.e., $n_X(i) = |N_X(i)|$. For example, $N_I(i)$ denotes the set of infected neighbors of i and $n_I(i)$ denotes the number of infected neighbors of i .

The basic parameters in the model are: the infection probability β_i which may vary for different individuals in state I ; the disease-induced death probability c_i which may also vary for different individuals in state I ; the natural death probability u ; and the recovery probability γ .

The model is a stochastic process on Γ , carried out for T generations. At each generation t of the process, a vertex i in Γ is picked at random and one of the following occurs depending on the state of i :

If i is a S vertex, it can die with probability u and be reborn as a S vertex again. If this happens, the process starts over again with another vertex in Γ .

If the vertex does not die, it can become infected by any infected neighbor it may have. The probability p_i that the S vertex i becomes infected is simply the complement of the probability that the vertex is not infected by any of its infected neighbors, and is thus is

²A vaccinated individual cannot propagate the disease to her neighbors. Moreover, when vaccinated individuals die, they are replaced by new born individuals who inherit their immunity to the disease, and hence also belong to the vaccinated state. Thus, when an individual is vaccinated, we are effectively deleting the corresponding vertex from the contact network Γ .

given by

$$p_i = 1 - \prod_{j \in N_I(i)} (1 - \beta_j). \quad (6.28)$$

If i does not get infected, the process starts over again with another vertex in Γ .

If i does get infected, it assumes the transmission probability of one of the infected neighbors $j \in N_I(i)$ in proportion to the values of β_j , i.e., the probability $p_{i \leftarrow j}$ that i will assume j 's transmission probability is given by

$$p_{i \leftarrow j} = \frac{\beta_j}{\sum_{l=1}^{n_I(i)} \beta_l}. \quad (6.29)$$

The β value is selected using the roulette-wheel selection method. This involves picking a uniform random number $r \in [0, L)$, where $L = \sum_{l=1}^{n_I(i)} \beta_l$, and choosing the β value corresponding to the interval in which r lands (β_k in Figure 76). This procedure determines (stochastically) which neighbor j of i infects i .




Figure 76: Roulette-wheel selection of transmission probability for the S -vertex i from its neighborhood.

If a mutation of the transmission probability does occur (with probability μ) then the new β value of i is taken to be a random number picked from a the normal distribution $N(\beta_j, \sigma)$, i.e., from a normal distribution whose mean is the transmission probability of j , and whose standard deviation is σ .

We are assuming that the values of β_i and c_i for any given vertex i are coupled through the relation

$$\beta_i = \beta_0 c_i^\alpha, \quad (6.30)$$

where $\beta_0, \alpha \in (0, 1)$.

Thus if the β value assigned to i has mutated then the corresponding new c value for i will be determined from the new β value using Equation 6.30.

If no mutation occurs (with probability $1 - \mu$) then the β value of i is simply replaced by that of j , i.e., by β_j . Since β and c are coupled, we also replace the c value of i by the c value of j .

If the vertex i is in state I it can die either of natural causes or due to the disease, i.e., with probability $1 - (1 - d)(1 - c_i)$. If the vertex dies, it is reborn as a S vertex, and the process starts over again with another vertex in Γ . If the vertex does not die it can recover with probability γ and become a R vertex.

If the vertex i is in the R state it can die with probability u and be reborn a S vertex.

If the vertex i is in the V state, the process starts over again with a new vertex in Γ .

The process is set in motion by infecting a single individual in a population Γ of n susceptible individuals. This seed individual is picked as the random neighbor of a randomly picked individual. On scale-free networks—social networks through which human diseases spread and technological and information networks through which computer viruses spread are typically scale-free—the seed individual is likely to be a highly connected individual [CHB03, MKC04]. The seed individual is assigned a transmission probability β , and hence the virulence $c = \left(\frac{\beta}{\beta_0}\right)^{\frac{1}{\alpha}}$. The model is run for T generations. The mean virulence \bar{c} of the population at any generation is calculated as

$\bar{c} = \frac{\sum_{i \in N_I(i)} c_i}{N_I(i)}$. We are interested in the long-term value of the mean virulence, which we denote by \bar{c}_∞ .

Algorithm 6.6 suggests a recipe for implementing the model as a computer program. We provide Python- and Java-based implementations of the model, the details of which can be found in B.6.

Algorithm 6.6 Disease dynamics on complex networks

Input: A population network Γ of size n with a single individual infected individual i with infection probability β , and the rest in the susceptible (S) state, i.e., at generation 0, the population has the state profile $\mathbf{s}^0 = \{s_1^0, s_2^0, \dots, s_N^0\}$ with $s_1^0 = s_2^0 = \dots = s_{i-1}^0 = s_{i+1}^0 = \dots = s_N^0 = R$ and $s_i^0 = I$ and the virulence profile $\mathbf{c}^0 = \{c_1^0, c_2^0, \dots, c_N^0\}$ with $c_1^0 = c_2^0 = \dots = c_{i-1}^0 = c_{i+1}^0 = \dots = c_N^0 = 0$ and $c_i = \left(\frac{\beta}{\beta_0}\right)^{\frac{1}{\alpha}}$; the natural death probability u ; recovery probability γ ; constants β_0 and α ; rate μ and standard deviation σ for Gaussian mutations of transmission probability; and number of generations T .

Output: The lists $L_1 = \{\mathbf{s}^0, \mathbf{s}^1, \dots, \mathbf{s}^{T-1}\}$, and $L_2 = \{\mathbf{c}^0, \mathbf{c}^1, \dots, \mathbf{c}^{T-1}\}$.

```

1:  $L_1 \leftarrow [], L_2 \leftarrow []$ 
2: for  $t$  from 1 to  $T$  do
3:    $L_1.append(\mathbf{s}^{t-1}), L_2.append(\mathbf{c}^{t-1})$ 
4:   for  $count$  from 1 to  $N$  do
5:      $i \leftarrow$  random individual from  $\Gamma$ 
6:     if  $s_i^t = S$  then
7:       if  $random() < u$  then
8:         continue
9:       end if
10:       $p \leftarrow \prod_{j \in N_I(i)} (1 - \beta_j^{t-1})$ 

```

```

11:     if  $random() < 1 - p$  then
12:          $j \leftarrow$  a random individual from  $N_I(i)$  picked in proportion to infection
           probability (Equation 6.29)
13:          $s_i^t \leftarrow I, \beta_i^t \leftarrow \beta_j^{t-1}$ 
14:         if  $random() < \mu$  then
15:              $\beta_i^t \leftarrow \max(0, \min(gaussian(\beta_i^t, \sigma), 1))$ 
16:         end if
17:          $c_i^t \leftarrow \left(\frac{\beta_i^t}{\beta_0}\right)^{\frac{1}{\alpha}}$ 
18:         end if
19:     else if  $s_i^{t-1} = I$  then
20:         if  $random() < 1 - (1 - u)(1 - c_i^{t-1})$  then
21:              $s_i^t \leftarrow S, \beta_i^t \leftarrow 0, c_i^t \leftarrow 0$ 
22:         else if  $random() < \gamma$  then
23:              $s_i^t \leftarrow R, \beta_i^t \leftarrow 0, c_i^t \leftarrow 0$ 
24:         end if
25:     else if  $s_i^{t-1} = R$  then
26:         if  $random() < u$  then
27:              $s_i^t \leftarrow S, \beta_i^t \leftarrow 0, c_i^t \leftarrow 0$ 
28:         end if
29:     else if  $s_i^{t-1} = V$  then
30:         continue
31:     end if
32: end for
33: end for
34: return  $L_1, L_2$ 

```

There are few points worth noting about the model. First, the population size remains constant over time in spite of births and deaths. This is because every dead individual is immediately replaced by a new born individual. Second, the structure of the contact network Γ also remains constant over time³. Third, the mutations of the transmission probabilities provide an essential ingredient for the evolution of virulence of a disease. Without mutations, i.e., if $\mu = 0$, the model reduces to the standard SIR model with births and deaths with fixed values of β and c .

We can derive R_0 for a population represented by a complete network Γ_{WM} as follows. The mean infection period τ for an infected individual is the average time it takes for an infected individual to leave the I state to which the individual belongs and either join the R state or the S state. Let $p_{I \rightarrow R,S}$ denote the probability that the individual leaves the I state, $p_{I \rightarrow R}$ denote the probability that the individual enters the R state, and $p_{I \rightarrow S}$ denote the probability that the individual enters the S state. The infected individual can join the R state by recovering with a probability γ , so $p_{I \rightarrow R} = \gamma$. The infected individual can join the S state by dying and being reborn; the individual can die of natural causes with probability u or due the infection with probability c , so $p_{I \rightarrow S} = 1 - (1 - u)(1 - c)$. Therefore, $p_{I \rightarrow R,S}$ is obtained as follows:

$$\begin{aligned}
p_{I \rightarrow R,S} &= p_{I \rightarrow R} + p_{I \rightarrow S} - p_{I \rightarrow R} p_{I \rightarrow S} \\
&= \gamma + [1 - (1 - u)(1 - c)] - \gamma[1 - (1 - u)(1 - c)] \\
&= \gamma + 1 - (1 - u)(1 - c) - \gamma + \gamma(1 - u)(1 - c) \\
&= 1 - (1 - u)(1 - c) + \gamma(1 - u)(1 - c) \\
&= 1 - [(1 - u)(1 - c) - \gamma(1 - u)(1 - c)] \\
&= 1 - (1 - \gamma)(1 - u)(1 - c).
\end{aligned} \tag{6.31}$$

³In Section 6.4 we explore the possibility where the network is slightly perturbed in each generation.

So the mean infection period is given by $\tau = \frac{1}{P_{I \rightarrow R,S}} = \frac{1}{1 - (1 - \gamma)(1 - u)(1 - c)}$. Since the network is complete, the infected individual is in contact with every other individual in the population, and since the population is naive, all of these individuals are in the susceptible state, i.e., the fraction s of susceptible individuals is 1. Therefore, the number of new infections caused by the infected individual per unit time is β . Thus

$$\begin{aligned}
 R_0 &= \text{number of infections per unit time} \times \text{mean infection period} \\
 &= \beta \tau \\
 &= \frac{\beta}{1 - (1 - \gamma)(1 - u)(1 - c)} \\
 &= \frac{\beta_0 c^\alpha}{1 - (1 - \gamma)(1 - u)(1 - c)}. \tag{6.32}
 \end{aligned}$$

For small values of γ, d and c , we have the following approximation:

$$1 - (1 - \gamma)(1 - u)(1 - c) \approx \gamma + u + c. \tag{6.33}$$

Therefore, the expression for R_0 can be approximated by

$$R_0 = \frac{\beta_0 c^\alpha}{\gamma + u + c}, \tag{6.34}$$

which is what we derived for R_0 for the standard SIR model with births and deaths. Thus, for small values of the probabilities γ, u , and c , we expect the discrete-time model to closely approximate the continuous-time model. Moreover, we expect, as in the case of the continuous-time model, the virulence of disease to evolve so as to maximize R_0 , and the following example demonstrates that this is indeed the case. Calculating the expression for c^* that maximizes R_0 is a trivial exercise in calculus, and is given by

$$c^* = \frac{\alpha(\gamma + u)}{1 - \alpha}. \tag{6.35}$$

Example 6.4. Consider a disease spreading over a complete network Γ_{WM} having 10000 vertices. Let $\beta = 0.075, \gamma = 0.08, u = 0.01, \beta(c) = \sqrt{c}, \mu = 0.0001$, and $\sigma = 0.05$. Figure 77 shows how R_0 varies with c . The value c^* of virulence that maximizes R_0 is $c^* = 0.09$.




Figure 77: The R_0 curve for our model. $\beta(c) = \sqrt{c}$, $\gamma = 0.2$, and $u = 0.01$.

We simulate the disease dynamics on a computer for $T = 300000$ generations using the software implementation of our model. The report frequency $\nu = 100$. Figure 78(a) shows the evolution of the fractions s, x, r of the population in the different disease states. Figure 78(b) shows the evolution of the mean virulence \bar{c} . The long-term value \bar{c}_∞ of the mean virulence, indicated by the inset in the figure, is $\bar{c}_\infty = 0.097$, which agrees very well with c^* , the value of virulence that maximizes R_0 . □




Figure 78: Disease dynamics on a complete network. Simulation parameters: $n = 10000$, $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 300000$, and $\nu = 100$. (a) Evolution of the fractions s, x, r of the population in the three disease states. (b) Evolution of the mean virulence \bar{c} .

6.3 Simulation Results

In this section we present the results of simulating our disease dynamics model on networks having different structural properties. Our inquiry is two-fold: we want to know how the structural properties of a network shape the long-term value \bar{c}_∞ of the mean virulence of a disease spreading over the network; and what is the most effective way of carrying out vaccination programs so that the disease is eliminated with the least number of vaccinations administered.

All the simulations, unless specified otherwise, were carried out using the following values for the parameters involved: population size $n = 100000$, transmission probability $\beta = 0.075$, recovery probability $\gamma = 0.08$, natural death probability $u = 0.01$, $\beta(c) = \sqrt{c}$, mutation probability $\mu = 0.0001$, standard deviation for mutations $\sigma = 0.05$, number of generations $T = 300000$, and report frequency $\nu = 100$. See B.6 for details on how to simulate our model of disease dynamics.

For the parameters chosen, the value c^* of virulence that maximizes R_0 is $c^* = 0.09$, and this value is marked by a dashed line in the plots. In the plots that show the evolution of mean virulence \bar{c} , the inset in the plots indicates the value \bar{c}_∞ , the average value of the mean virulence \bar{c} over the last 10% of the generations.

6.3.1 Evolution of Virulence

Complete Network

The results shown in each of the plots in this section is an average over ten runs. The first set of results concerns complete networks Γ_{WM} . Figure 79 shows the evolution of mean virulence \bar{c} on a complete network.




Figure 79: Evolution of mean virulence \bar{c} on a complete network. Simulation parameters: $n = 10000$, $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 300000$, and $\nu = 100$.

Model Networks with Varying Average Degree

The next set of results concerns model networks with varying average degree d . Figures 80(a)-(d) show the evolution of mean virulence \bar{c} on random regular networks Γ_{RR} with average degree $d = 4, 6, 8$, and 10 . Figures 81(a)-(d) show the evolution of mean virulence \bar{c} on Erdős-Rényi networks Γ_{ER} with average degree $d = 4, 6, 8$, and 10 . Figures 82(a)-(d) show the evolution of mean virulence \bar{c} on Barabási-Albert networks Γ_{BA} with average degree $d = 4, 6, 8$, and 10 . Figure 83 shows how the long-term value \bar{c}_∞ of mean virulence varies with average degree d , for Γ_{RR} , Γ_{ER} , and Γ_{BA} networks. The solid line indicates the value of \bar{c}_∞ on a complete network Γ_{WM} .




Figure 80: Evolution of mean virulence \bar{c} on random regular networks with varying average degree d . (a) $d = 4$. (b) $d = 6$. (c) $d = 8$. (d) $d = 10$. Simulation parameters: $n = 10000$, $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 300000$, and $\nu = 100$.

Model Networks with Clustering

The next set of results concerns power-law networks Γ_{BA}^C with average degree $d = 10$, and varying clustering coefficient $C^{(2)}$. Figures 84(a)-(d) show the evolution of mean virulence \bar{c} on networks with clustering coefficient $C^{(2)} = 0.1, 0.2, 0.3$, and 0.4 . Figure 85 shows how the long-term value \bar{c}_∞ of mean virulence varies with clustering coefficient $C^{(2)}$.




Figure 81: Evolution of mean virulence \bar{c} on Erdős-Rényi networks with varying average degree d . (a) $d = 4$. (b) $d = 6$. (c) $d = 8$. (d) $d = 10$. Simulation parameters: $n = 10000$, $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 300000$, and $\nu = 100$.

Model Networks with Homophily

The next set of results concerns power-law networks Γ_{BA}^h with average degree $d = 10$, and varying homophily coefficient h . Figures 86(a)-(d) show the evolution of mean virulence \bar{c} on networks with homophily coefficient $h = -0.1, 0, 0.1$, and 0.2 . Figure 87 shows how the long-term value \bar{c}_∞ of mean virulence varies with homophily coefficient h .




Figure 82: Evolution of mean virulence \bar{c} on Barabási-Albert networks with varying average degree d . (a) $d = 4$. (b) $d = 6$. (c) $d = 8$. (d) $d = 10$. Simulation parameters: $n = 10000$, $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 300000$, and $\nu = 100$.

Empirical Networks

The last set of results concerns empirical networks. Figures 88(a)-(f) show the evolution of mean virulence \bar{c} on: the Γ^{internet} network (a); the $\Gamma^{\text{hepcoauth}}$ network (b); the $\Gamma^{\text{astrocoauth}}$ network (c); the Γ^{fb} network (d); the Γ^{p2p} network (e); and the Γ^{protein} network (f). Figure




Figure 83: Long-term value \bar{c}_∞ of mean virulence versus average degree d , for random regular, Erdős-Rényi, and Barabási-Albert networks. The solid line indicates the value of \bar{c}_∞ on a complete network Γ_{WM} . Simulation parameters: $n = 10000, \beta = 0.075, \gamma = 0.08, u = 0.01, \beta(c) = \sqrt{c}, \mu = 0.0001, \sigma = 0.05, T = 300000$, and $v = 100$.

89 shows the long-term values \bar{c}_∞ of mean virulence for each of the six empirical networks.




Figure 84: Evolution of mean virulence \bar{c} on power-law networks with clustering coefficient $C^{(2)}$. (a) $C^{(2)} = 0.1$. (b) $C^{(2)} = 0.2$. (c) $C^{(2)} = 0.3$. (d) $C^{(2)} = 0.4$. Simulation parameters: $n = 10000$, $d = 10$, $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 300000$, and $\nu = 100$.




Figure 85: Long-term value \bar{c}_∞ of mean virulence versus clustering coefficient $C^{(2)}$, on power-law networks. Simulation parameters: $n = 10000$, $d = 10$, $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 300000$, and $\nu = 100$.




Figure 86: Evolution of mean virulence \bar{c} on power-law networks with homophily coefficient h . (a) $h = -0.1$. (b) $h = 0$. (c) $h = 0.1$. (d) $h = 0.2$. Simulation parameters: $n = 10000$, $d = 10$, $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 300000$, and $\nu = 100$.




Figure 87: Long-term value \bar{c}_∞ of mean virulence versus homophily coefficient h , on power-law networks. Simulation parameters: $n = 10000$, $d = 10$, $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 300000$, and $\nu = 100$.

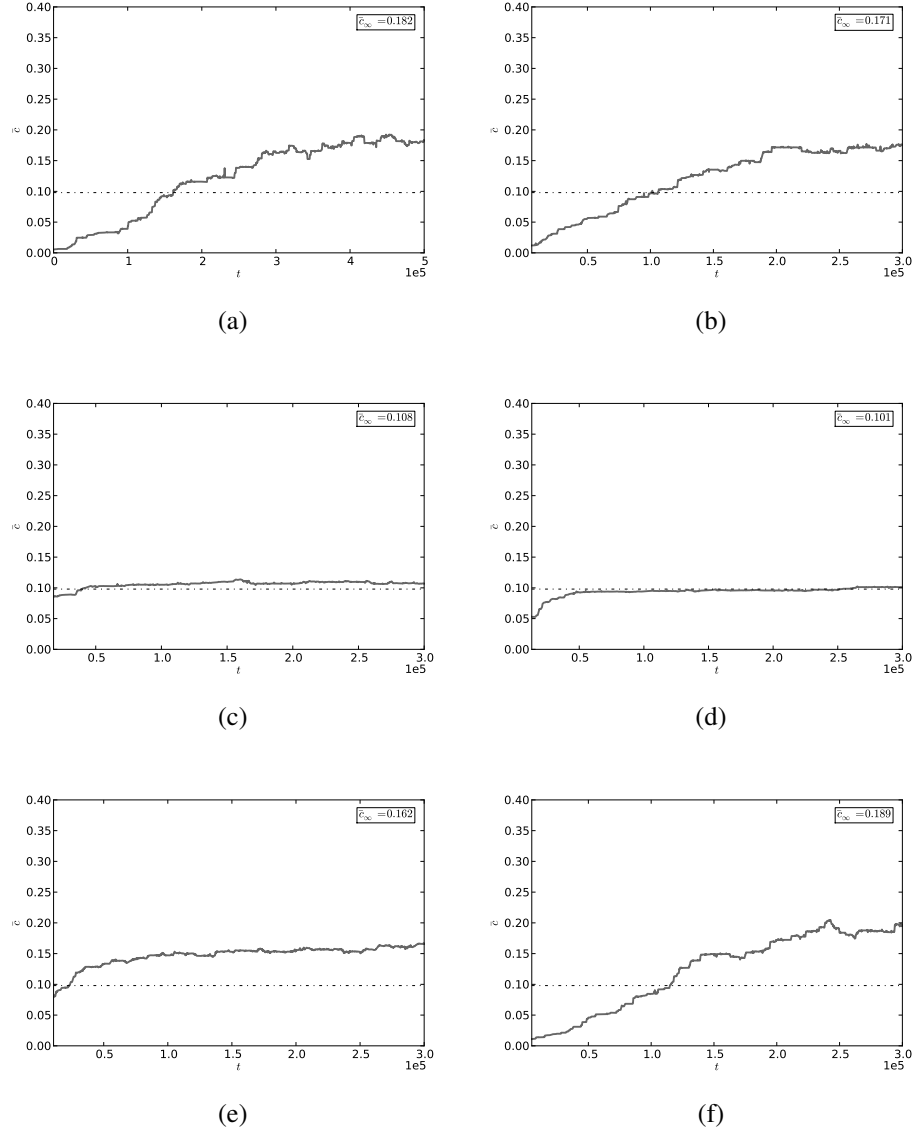


Figure 88: Evolution of mean virulence \bar{c} on empirical networks. See Table 2 for the basic properties of these networks. (a) the Γ^{internet} network. (b) the $\Gamma^{\text{hepcoauth}}$ network. (c) $\Gamma^{\text{astrocoauth}}$ network. (d) the Γ^{fb} network. (e) the Γ^{p2p} network. (f) the Γ^{protein} network. Simulation parameters: $\beta = 0.075, \gamma = 0.08, u = 0.01, \beta(c) = \sqrt{c}, \mu = 0.0001, \sigma = 0.05, T = 300000$, and $\nu = 100$.




Figure 89: Long-term value \bar{c}_∞ of mean virulence for empirical networks. See Table 2 for the basic properties of these networks. Simulation parameters: $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 300000$, and $\nu = 100$.

6.3.2 Vaccination Strategies

We next pursue our second line of inquiry, namely the effectiveness of different vaccination strategies in the eradication of a disease [AM85, DB02, PV02b]. In the event of an epidemic outbreak, health officials often have to work with a limited supply of vaccinations, and have to respond as quickly as possible. So the efficiency with which vaccinations are administered is of paramount importance.

We explore different vaccination strategies in the context of our model as follows. We infect a single individual in a naive population with a disease and let it spread, giving it enough time (say T_1 generations) so that the mean virulence \bar{c} attains its equilibrium value \bar{c}_∞ . We next vaccinate a fraction ν of individuals, i.e., change their disease state to V . We then let the simulation continue for another T_2 generations. We are interested in the smallest value of ν (we call it ν^*) that would wipe out the disease, i.e., drive \bar{c}_∞ to zero.

Human diseases spread over social networks [EGK04, LEA03] and computer viruses spread over technological and information networks [KSC97, Kle07]. These networks typically have a power-law degree distribution. We therefore restrict our study of vaccination strategies to Barabási-Albert networks and to empirical networks which mimic real-world contact networks. In particular we consider a Barabási-Albert network with average degree $d = 10$, and the social network $\Gamma^{\text{astrocoauth}}$ of coauthorships among astrophysicists for our study.

We consider three different strategies for vaccinating individuals in a population: in the *degree* (\mathcal{V}_D) strategy, the individuals are vaccinated in the reverse order of their degree centrality, i.e., their degree [AM92]. We use a *simultaneous approach*, meaning the degree centrality of the individuals is calculated once in advance before the vaccinations are carried out; in the *referral* (\mathcal{V}_R) strategy [CHB03, MKC04], each individual that is picked for vaccination is a random neighbor of a randomly picked individual—in power-law

networks individuals picked using the referral method are typically highly-connected individuals; finally, in the *random* (\mathcal{V}_N) strategy, each individual is picked at random [AM92].

The results shown in each of the plots in this section is based on a single run. Figures 90(a)-(d) show the evolution of mean virulence \bar{c} on a Barabási-Albert network Γ_{BA} with ν of the population vaccinated after $T_1 = 100000$ generations. In (a) no vaccinations were administered; in (b) 20% of the population was vaccinated using the \mathcal{V}_D strategy; in (c) 25% of the population was vaccinated using the \mathcal{V}_R strategy; and in (d) 50% of the population vaccinated using the \mathcal{V}_N strategy. Figure 91 shows how the long-term value \bar{c}_∞ of mean virulence varies with the fraction ν of individuals vaccinated, for each of the three vaccination strategies.

Figures 92(a)-(d) show the evolution of mean virulence \bar{c} on the $\Gamma^{\text{astrocoauth}}$ network with fraction ν of the population vaccinated after $T_1 = 100000$ generations. In (a) no vaccinations were administered; in (b) 20% of the population was vaccinated using the \mathcal{V}_D strategy; in (c) 50% of the population was vaccinated using the \mathcal{V}_R strategy; and in (d) 55% of the population was vaccinated using the \mathcal{V}_D strategy. Figure 93 shows how the long-term value \bar{c}_∞ of mean virulence varies with the fraction ν of individuals vaccinated, for the three vaccination strategies.




Figure 90: Evolution of mean virulence \bar{c} on a Barabási-Albert network with vaccinations. Fraction ν of the population was vaccinated after $T_1 = 100000$ generations. (a) No vaccinations were administered. (b) 20% of the population was vaccinated using the \mathcal{V}_D strategy. (c) 25% of the population was vaccinated using the \mathcal{V}_R strategy. (d) 50% of the population was vaccinated using the \mathcal{V}_N strategy. Simulation parameters: $n = 10000$, $d = 10$, $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 200000$, and $\nu = 100$.




Figure 91: Long-term value of mean virulence \bar{c}_∞ versus the fraction v of individuals vaccinated on a Barabási-Albert network. Simulation parameters: $n = 10000, d = 10, \beta = 0.075, \gamma = 0.08, u = 0.01, \beta(c) = \sqrt{c}, \mu = 0.0001, \sigma = 0.05, T = 200000$, and $v = 100$.




Figure 92: Evolution of mean virulence \bar{c} on an empirical network with vaccinations. Fraction v of the population was vaccinated after $T_1 = 100000$ generations. (a) No vaccinations were administered. (b) 20% of the population was vaccinated using the \mathcal{V}_D strategy. (c) 50% of the population was vaccinated using the \mathcal{V}_R strategy. (d) 55% of the population was vaccinated using the \mathcal{V}_N strategy. Simulation parameters: $n = 17903$, $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 200000$, and $v = 100$.




Figure 93: Long-term value \bar{c}_∞ of mean virulence versus the fraction v of individuals vaccinated on an empirical network. Simulation parameters: $n = 17903$, $\beta = 0.075$, $\gamma = 0.08$, $u = 0.01$, $\beta(c) = \sqrt{c}$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 200000$, and $v = 100$.

6.4 Discussion

On a complete network Γ_{WM} , the long-term value \bar{c}_∞ of mean virulence obtained from simulations agrees (see Figure 79) very well with the predicted value c^* of virulence that maximizes R_0 (Equation 6.32).

On random regular Γ_{RR} , Erdős-Rényi Γ_{ER} , and Barabási-Albert Γ_{BA} networks, the long-term value \bar{c}_∞ of mean virulence is higher (see Figures 80-83) for small values of average degree d , and approaches the value on a complete network Γ_{WM} as d increases. For example, on Γ_{BA} networks with average degree $d = 4$, the value of \bar{c}_∞ is approximately 3.25 times the value on a Γ_{WM} network, while on Γ_{BA} networks with average degree $d = 20$, it is only about 1.2 times the value on Γ_{WM} . This result makes sense because with the increase in average degree, a network tends more and more towards a complete network. What is surprising is that the value of d at which networks behave like a complete network is rather small, about 20. Complete networks thus offer us a baseline for the long-term value of virulence against which we can compare the value on other networks, since if for a particular disease, we have knowledge of the parameters such as γ, u and α , we can estimate a lower bound for \bar{c}_∞ .

The fact that the virulence level on power-law networks is higher at low values of average degree has important practical consequences since both human diseases and computer viruses typically spread over contact networks having power-law degree distributions, and small average degree.

It is well documented that heterogeneous networks such as power-law networks favor the spreading of a disease [AM92, ML01, PV01b, PV02a]. Our model suggests that as long as the mixing among individuals is non-random and the average degree of the contact network is small, the levels of virulence of a disease will be high.

Except for very small values of the average degree d , the effect of degree distribution on the value of \bar{c}_∞ is insignificant.

On power-law networks Γ_{BA}^C with fixed average degree d , the value of \bar{c}_∞ increases (see Figures 84 and 85) with increasing clustering coefficient $C^{(2)}$. There is roughly a 17% increase in the value of \bar{c}_∞ on a network with $C^{(2)} = 0.4$ than on a network with no clustering. This has important consequences for human diseases, which spread over social networks that exhibit high clustering.

On a power-law network Γ_{BA}^h with fixed average degree d , the value of \bar{c}_∞ decreases (see Figures 86 and 87) with increasing homophily coefficient h . There is roughly a 15% decrease in the value of \bar{c}_∞ from a network with $h = -0.1$ to a network with $h = 0.2$. This suggests that a human disease spreading over a (typically homophilic) social network will be less virulent than a computer virus with same disease parameters spreading over a (typically heterophilic) technological or information network.

The simulation results on empirical networks (see figures 88-89) further support our claim that the value of \bar{c}_∞ increases with average degree d . On Γ^{protein} , which has the smallest average degree ($d \approx 3.56$), the value of \bar{c}_∞ is the highest. On Γ^{fb} , which has the highest average degree ($d \approx 46.51$), the value of \bar{c}_∞ is the lowest, and comparable to the value on a complete network. These empirical networks have very different sizes, and yet the trend of \bar{c}_∞ with respect to their average degree d is comparable to the trend we observed in the case of model networks, which were all of size 10000. This suggests that the results of the model, as long as all other parameters are fixed, is independent of the population size.

In our model it was assumed that over the time scale over which the disease spreads through a population, the contact network representing the population remains fixed in terms of its structural properties. One could argue that this assumption is somewhat

unrealistic, since people’s contacts often change over time [KKK02, Moo02, WM92]. We explored this possibility in our model by carrying out a specified number s of *double-edge swap* operations on the network per generation. Each such operation removes two randomly chosen edges (u, v) and (x, y) and creates the new edges (u, x) and (v, y) . These operations change the structure of the network without altering its average degree. Figure 94 compares the value of \bar{c}_∞ against the value of s on a Barabási-Albert Γ_{BA} network with average degree $d = 10$. It indicates that for small values of s (about 100 swaps per generation), the model behaves similarly to the one without the edge swaps, thus suggesting that the model is robust against small fluctuations in the structure of the underlying network. Of course, as s increases the underlying network changes so much that it begins to approximate a complete network, and that is what we observe in the figure.




Figure 94: Long-term value \bar{c}_∞ of mean virulence versus the number s of double-edge swaps per generation on power-law networks. Simulation parameters: $n = 10000, d = 10, \beta = 0.075, \gamma = 0.08, u = 0.01, \beta(c) = \sqrt{c}, \mu = 0.0001, \sigma = 0.05, T = 300000$, and $\nu = 100$.

Turning our attention to the results on vaccination strategies, for all three vaccination strategies (\mathcal{V}_D , \mathcal{V}_R , and \mathcal{V}_N) the value of \bar{c}_∞ increases (Figures 90-93) with increasing ν before it starts to decrease and eventually drops to zero. For example, on the $\Gamma^{\text{astrocoauth}}$ network, the value of \bar{c}_∞ is about three times higher when 20% of the population is vaccinated using the \mathcal{V}_D strategy than its value when no vaccinations were administered. This makes sense because as vaccinations are carried out, the corresponding vertices are effectively removed from the network, as a result of which the average degree of the network decreases. From our earlier result, we know that for networks with low average degree the value of \bar{c}_∞ is higher. This observation, unfortunately, has serious implications for real-world diseases, suggesting that if enough vaccinations are not administered, then problem could be made worse by increasing the level of virulence of the disease rather than decreasing it.

Among the three strategies, \mathcal{V}_D is the most effective, \mathcal{V}_R is the second best, and \mathcal{V}_N is the least effective. Adopting the \mathcal{V}_D strategy requires knowledge of the entire network, but since maps of contact networks over which human diseases and computer viruses spread are becoming more and more readily available [EGK04, LEA03, KSC97, Kle07], implementing the \mathcal{V}_D strategy is not a far-fetched proposition. Thus, even though power-law networks allow for higher levels of virulence, the heterogeneities in their structure make them vulnerable to targeted attacks, and this can be put to our advantage in trying to combat a disease. Human diseases and computer viruses typically spread over power-law networks, so vaccination deployment programs can take advantage of the vulnerabilities of these networks and favor the \mathcal{V}_D and \mathcal{V}_R strategies over the naive \mathcal{V}_N strategy.

6.5 Future Directions

In this section we propose three new lines of inquiry for our model of disease dynamics, along with an outline of ways to approach them.

1. We did not establish analytically that on a complete network Γ_{WM} , a disease-causing parasite in our model evolves in virulence so as to maximize R_0 , given by Equation 6.32. We argued instead that this should be the case for small values of the (transmission, recovery, and death) probabilities involved, since in that case our model approximates the continuous-time model for which we did establish the result analytically. We supported our argument with simulation results.

A future project would be to establish the result explicitly for our model. This could be done by treating the disease dynamics as a Markov process, with the transition matrix M for the process specifying how individuals in a population transition with respect to their disease states from time t to time $t + 1$. This matrix is easily written down in terms of the disease parameters, i.e., the transmission, recovery, and (natural and disease-induced) death probabilities. We could consider two strains S_1 and S_2 of the parasite, having transmission rates β_1 and β_2 , and disease-induced death rates c_1 and c_2 respectively. Comparing the steady-state vectors—a three-dimensional vector whose components specify the fractions s, x, r of the individuals in the population in the three disease states—of the transition matrices M_1 and M_2 for the two strains of the parasite, we could show that the virulence does indeed evolve so as to maximize R_0 .

2. For a given set of simulation parameters, the results obtained from our model on a network $\Gamma \neq \Gamma_{WM}$ agree with the results predicted for a complete network Γ_{WM} as long as the average degree d of Γ is sufficiently large ($d > 20$). The agreement

breaks down for small values of d . In the small- d regime, where the network effects are the strongest, the analysis must take network structure into account.

Various approaches have been put forth for extending the classical epidemiological models to networks. These include the study of late-time and time-dependent properties of epidemics on networks, and degree-based approximation [New10].

In the study of the late-time properties of epidemics on networks, one is interested in the long-term behavior of the disease on networks. For example, in the SI model as $t \rightarrow \infty$ every individual who can be infected by the disease is infected. Thus in the limit of long times the disease will spread from every initial carrier to infect all reachable vertices, i.e., all vertices in the component of which the carrier belongs. Most networks have one large component that contains a significant fraction S of all vertices in the network, and a selection of smaller components. If the initial carrier of the disease is chosen uniformly at random from the network, the probability that it will fall in the large component, and hence cause a large outbreak is S , and the size of the outbreak as a fraction of the network will also be S . Conversely, with probability $1 - S$ the initial carrier will fall in one of the small components and the outbreak will be small.

In the SIR model, it is in general not true that the susceptible neighbor of an infected individual will always get infected in the end. The probability that such neighbors may never catch the disease is $e^{-\beta\tau}$, where β is the transmission rate and τ is the amount of time for which the infected individual remains infected. Thus the probability that the disease is transmitted is $\Phi = 1 - e^{-\beta\tau}$. Now, one could take the network and “color in” or “occupy” each edge with probability Φ , or not with probability $1 - \Phi$. The disease dynamics can then be treated as an ordinary bond

percolation process [SA94]. The occupied edges represent contacts sufficient to spread the disease, but not necessarily actual disease transmission: if the disease does not reach either end of an occupied edge then the disease will not be transmitted along that edge, so edge occupation only represents the potential for transmission if the disease reaches an edge. Now consider the spread of a disease that starts at a randomly chosen vertex. The set of vertices to which the disease will ultimately spread is precisely the set connected to the initial vertex by any path of occupied edges. The end result is that the disease infects all members of the bond percolation cluster to which the initial carrier belongs. The percolation transition for bond percolation on the network corresponds precisely to the epidemic threshold for a disease on the same network where the edge occupation probability Φ is given in terms of the transmission rate β and recovery time τ , and the sizes of outbreaks are given by the sizes of the bond percolation clusters.

In the study of the time-dependent properties of epidemic models, one is interested in the time evolution of the probabilities for vertices to be in specific disease states. One can imagine having repeated outbreaks of the same disease on the same network, starting from the same initial conditions, and calculating for example the average probabilities $s_i(t)$ and $x_i(t)$ that vertex i is susceptible or infected at time t . Given the adjacency matrix A of a network one can write down differential equations for the evolution of such quantities as a straightforward generalization of the equations for the corresponding classical models for well-mixed populations. Unfortunately, these equations are not solvable in closed form for a general matrix A . Moreover, even the results obtained from numerical integration of such equations do not agree with the results obtained by simulating the disease with the same set of disease parameters. The reason for this disagreement is that the equations contain the product of average quantities such as

s_i and x_i , and it is implicitly assumed that the product of averages is equal to the average of products. In the fully-mixed model this is true (for large n) because of the mixing itself, but is not true for arbitrary networks, because the probabilities are not independent. In general, these quantities will be correlated between neighboring vertices. These correlations are incorporated into the calculations, at least approximately, by using the so-called pair approximation or moment closure method [Ell01].

In the degree-based approximation method, one focuses on outbreaks taking place in the giant component of the network. The approximation assumes that all vertices of the same degree have the same probability of infection at any given time. Suppose $s_k(t)$ and $x_k(t)$ are the probabilities that a vertex with degree k is susceptible or infected, respectively, at time t . Now consider a susceptible vertex A . To become infected, A has to contract the infection from one of its infected neighbors. The probability that a particular neighbor B is infected depends on B 's degree. But by assumption vertex A is not infected and so B cannot have caught the disease from A . If B is infected it must have caught the disease from one of its remaining neighbors. B 's probability of infection is thus x_k , but where k indicates the excess degree—the number of edges it has other than the edge we followed from A to reach it—and not the total degree. The advantage of this approach is that the probability of B being infected depends only on B 's excess degree and not on A 's degree, and the equations obtained are analytically tractable.

All of the the network-based approaches we discussed above fix the disease parameters. In our model we consider the evolution of virulence of a disease. So the most direct way of extending our model to networks is as follows. Consider two strains S_1 and S_2 of the disease-causing parasite. Following the program of the study of time-dependent properties of epidemics on networks, write down the differential

equations for the probabilities s_i, x_i^1, x_i^2, r_i that a vertex i is susceptible, infected (by S_1 or S_2), or recovered, in terms of the adjacency matrix A for the network Γ under consideration. Find the condition under which one strain can invade the other. This condition might hint at an expression for the basic reproductive number R_0 for the network Γ in terms of: the recovery probability γ , the natural death probability u , the disease-induced death probability c , and the average degree d of the network. We would thus have an analytical expression for the value c^* that maximizes R_0 on networks. Except for simple networks without cycles, such as Bethe lattices or Caylee trees [Bet35], we do not expect to be able to carry out the analysis in general for arbitrary networks, although we could use pair approximations [Ell01] to refine the results to obtain greater levels of accuracy. From simulating our model on networks, we see that the average degree of a network, more than any other property of the network, determines to a first approximation, the long-term value \bar{c}_∞ of the mean virulence. So the expression for c^* obtained for a Caylee tree of average degree d , albeit simplistic, should be a reasonably good approximation for that on an arbitrary network with average d .

We could alternatively follow the degree-based approximation program discussed above and write down the differential equations for $s_k(t), x_k^1(t), x_k^2(t)$, and $r_k(t)$, which represent the probabilities that individuals with degree k are susceptible, infected (by S_1 or S_2), or recovered at time t . Carrying out an invasion analysis should give us the condition under which one strain can invade the other, and thus provide us, possibly, with an expression for R_0 on networks.

3. In exploring different vaccination strategies, we only considered a single centrality-measure-based vaccination strategy, namely the strategy \mathcal{V}_D based on the degree centrality of the vertices in a network. Recall that the concept of a centrality

measure attempts to identify which vertices in a network are the most important or central [WF94]. A number of measures of centrality, besides degree, have been proposed for networks: eigenvector centrality, closeness centrality, and betweenness centrality. *Eigenvector centrality* is a widely employed centrality measure [Bon87]. It can be viewed in a sense as a refinement of degree centrality. Whereas degree centrality ranks a vertex as being important if it is connected to many other vertices, eigenvector centrality is based on the more subtle notion that a vertex should be viewed as important if it is linked to other vertices which are themselves important. *Closeness centrality* of a vertex is defined as the reciprocal of the mean geodesic between the vertex and all other vertices in the network⁴. The quantity takes high values for vertices that are only a short geodesic distance from many other vertices in the network, and is a natural measure of centrality which is widely used in network studies. *Betweenness centrality* measures how many short paths between vertices in a network pass through a given vertex [Fre77]. It is worth highlighting that while degree centrality of a vertex is a local measure, the other centrality measures are non-local and take into account the global structure of the network. As noted earlier, obtaining a representation of contact networks over which diseases spread is not at all a far-fetched proposition these days [EGK04, LEA03, KSC97, Kle07]. Once we have the contact network for a disease, these centrality measures are easily calculated using software tools such as NetworkX (<http://networkx.lanl.gov/>) and iGraph (<http://igraph.sourceforge.net/>).

As a future project, we could extend our study of vaccination strategies by considering strategies based on the non-local centrality measures described above.

⁴If a network has more than one component then the closeness centrality of a vertex is calculated as the reciprocal of the mean geodesic distance from the vertex to all other vertices in the same component.

As with \mathcal{V}_D , in the strategies \mathcal{V}_E , \mathcal{V}_C , and \mathcal{V}_B based on eigenvector, closeness, and betweenness centrality measures, we would vaccinate a fraction v of individuals in reverse order of the respective centrality measure and see how that affects the long-term value \bar{c}_∞ of the mean virulence. In addition, we could consider *simultaneous* and *sequential* approaches to evaluating the centrality measures of individuals. In the former, the values are calculated just once, before the vaccinations are carried out. In the latter, the centrality measures are recalculated after each vaccination; since the vaccination of an individual is tantamount to removing the corresponding vertex from the contact network, centrality measures of the individuals left in the network will be different each time an individual is vaccinated.

APPENDIX A

DYNAMICAL SYSTEMS

Prediction is very difficult, especially about the future.

- Niels Bohr

A.1 Introduction

In this appendix we provide a brief introduction to *dynamical systems*, a subject that deals with systems that evolve with time. We only consider continuous-time dynamical systems, and restrict our attention to those that do not explicitly depend on time, i.e., autonomous systems. We start out simple by looking at one-dimensional dynamical systems, which are relatively easy to understand. We then proceed to look at dynamical systems in two or more dimensions. For a comprehensive treatment of these topics, see [Str94, Ban05], and for proofs of theorems that we only state but do not prove, see [BC87, LS88].

A.2 One-Dimensional Dynamical Systems

A *one-dimensional dynamical system* or a *first-order system* is represented by the differential equation

$$\dot{x} = f(x), \tag{A.1}$$

where $x(t)$ is a real-valued function of time t , and $f(x)$ is a smooth, real-valued function of x .

We first look at an example that illustrates a qualitative (geometric) method for analyzing one-dimensional dynamical systems by interpreting the corresponding differential equations as a vector fields.

Example A.1. Consider the nonlinear differential equation

$$\dot{x} = \sin x, \tag{A.2}$$


which has an exact solution given by

$$x(t) = 2 \cot^{-1} \left[\cot \left(\frac{x_0}{2} \right) e^{-t} \right], \tag{A.3}$$


where $x_0 = x(0)$. Though exact, the solution is a nightmare to interpret. In contrast, a graphical analysis (also known as phase line analysis) of Equation A.2 is clear and simple (see Figure A.95(a)). We think of t as time, x as the position of an imaginary particle moving along the real line, and \dot{x} as the velocity of that particle. Equation A.2 then represents a *vector field* on the line. It specifies the velocity vector \dot{x} at each x . To sketch the vector field, it is convenient to plot \dot{x} versus x , and draw arrows on the x -axis to indicate the corresponding velocity vector at each x . The arrows point right when $\dot{x} > 0$ and left when $\dot{x} < 0$. At points where $\dot{x} = 0$, there is no flow, and such points are called *fixed points*, and are denoted by x^* . There are two kinds of fixed points: *stable* fixed points (also called *attractors* or *sinks*) represented by filled circles, and *unstable* fixed points (also called *repellers* or *sources*) represented by open circles.

Equipped with a *phase line* plot, such as Figure A.95(a), we can easily understand the solution to Equation A.2. We start our imaginary particle at x_0 and watch it carried along by the flow. The qualitative form of solutions for any initial condition, called the *phase portrait*, is shown in Figure A.95(b). □

The method described in Example A.1 can be extended to any one-dimensional system $\dot{x} = f(x)$. We just need to draw the graph of $f(x)$ as shown in Figure A.96, and then use it



(a)



(b)

Figure A.95: Phase line plot and phase portrait for the system A.2. (a) Phase line plot. (b) Phase portrait.

to sketch the vector field on the real line. The flow is to the right when $f(x) > 0$ and to the left when $f(x) < 0$. To find the solutions to $\dot{x} = f(x)$ starting from an arbitrary initial condition x_0 , we place an imaginary particle (called *phase point*) at x_0 and watch it carried along by the flow. As time goes by, the phase point moves along the x -axis according to

some function $x(t)$, called the *trajectory* based at x_0 , and it represents the solution of the differential equation starting from the initial condition x_0 .




Figure A.96: Phase line plot for $\dot{x} = f(x)$.

In terms of the original differential equation, fixed points represent *equilibrium* solutions (sometimes called steady, constant, or rest solutions). An equilibrium point is *stable* if all sufficiently small perturbations away from it damp out in time, and *unstable* if the perturbations grow in time.

Linear Stability Analysis

One often desires a quantitative measure of the stability of fixed points of a dynamical system. Such a measure can be obtained by *linearizing* the system about its fixed points.

Let x^* be a fixed point, and let $\eta(t) = x(t) - x^*$ be a small perturbation away from x^* . To see whether the perturbation grows or decays, we derive a differential equation for η .

$$\dot{\eta} = \frac{d}{dt}(x - x^*) = \dot{x}, \tag{A.4}$$

since x^* is constant. Thus, $\dot{\eta} = \dot{x} = f(x) = f(x^* + \eta)$. Using Taylor series expansion we obtain

$$f(x^* + \eta) = f(x^*) + \eta f'(x^*) + O(\eta^2), \quad (\text{A.5})$$

where $O(\eta^2)$ denotes quadratically small terms in η . Since $f(x^*) = 0$ we have

$$\dot{\eta} = \eta f'(x^*) + O(\eta^2). \quad (\text{A.6})$$

Now if $f'(x) \neq 0$, then $O(\eta^2)$ terms are negligible and we may write

$$\dot{\eta} = \eta f'(x^*). \quad (\text{A.7})$$

This linear equation in η is called the *linearization about x^** . It shows that the perturbation grows exponentially if $f'(x^*) > 0$ and decays if $f'(x^*) < 0$. If $f'(x^*) = 0$, then the $O(\eta^2)$ terms are not negligible and a nonlinear analysis is needed to determine stability. The reciprocal $\frac{1}{|f'(x^*)|}$ is called *characteristic time scale* and it determines the time required for $x(t)$ to vary significantly in the neighborhood of x^* .

Example A.2. Let us reconsider the dynamical system $\dot{x} = \sin x$ from Example A.1. The fixed points of the system are given by $x^* = k\pi$ where k is an integer. Also,

$$f'(x^*) = \cos k\pi = \begin{cases} 1 & \text{if } k \text{ is even,} \\ -1 & \text{if } k \text{ is odd.} \end{cases} \quad (\text{A.8})$$

Therefore, x^* is unstable if k is even and stable if k is odd, which agrees with the qualitative analysis of the system presented in Example A.1. □

Example A.3. A simple model¹ for population growth of organisms is given by the *logistic equation*

$$\dot{N} = rN \left(1 - \frac{N}{K}\right), \quad (\text{A.9})$$

¹The corresponding discrete-time model gives rise to complex dynamics, including chaos.

where $N(t)$ is the population at time t , K is the *carrying capacity*, and $r > 0$ is the growth rate. This model was first suggested by the Belgian mathematician Pierre François Verhulst in 1838 to describe the growth of human population. The model has an exact solution given by

$$N(t) = \frac{KN_0 e^{rt}}{K + N_0(e^{rt} - 1)}, \quad (\text{A.10})$$

where $N_0 = N(0)$. Figure A.97(a) shows a plot of \dot{N} versus N . Note that we plot only $N \geq 0$ since it does not make much sense talking about negative populations. Fixed points of Equation A.9 are $N^* = 0$ and $N^* = K$. By looking at the flow, we can see that $N^* = 0$ is unstable and $N^* = K$ is stable. In biological terms, $N = 0$ is an unstable equilibrium; a small population will grow exponentially fast and run away from $N = 0$. On the other hand, if N is perturbed slightly from K , the perturbation will decay monotonically and $N(t) \rightarrow K$ as $t \rightarrow \infty$. Also, for $N_0 < \frac{K}{2}$, the phase point moves faster and faster until it crosses $N = \frac{K}{2}$. Then the phase point slows down and eventually crawls towards $N = K$. Figure A.97(b) shows the trajectories of the system.


Quantitatively, $f'(N) = r - 2r\frac{N}{K}$. So, $f'(0) = r > 0$ and $f'(K) = -r < 0$. Thus, $N^* = 0$ is unstable and $N^* = K$ is stable. The characteristic time scale is $\frac{1}{|f'(N^*)|} = \frac{1}{r}$ in either case. □

The following theorem describes the conditions that guarantee the existence and uniqueness of solution for a one-dimensional dynamical system.


Theorem A.1. (Existence and uniqueness theorem) Consider the initial value problem

$$\dot{x} = f(x), \quad x(0) = x_0. \quad (\text{A.11})$$

Suppose that $f(x)$ and $f'(x)$ are continuous on an open interval \mathbb{R} of the x -axis, and suppose $x_0 \in \mathbb{R}$. Then the initial value problem has a solution $x(t)$ on some time interval $(-\tau, \tau)$ about $t = 0$, and the solution is unique.



(a)



(b)

Figure A.97: Phase line plot and phase portrait for the system A.9. (a) Phase line plot. (b) Phase portrait.

A.3 Linear Systems

In one-dimensional phase spaces the flow is confined—all trajectories are forced to move monotonically or remain constant. In higher-dimensional phase spaces, trajectories have

much more freedom, and so a wider range of dynamical behavior is possible. We first consider the simplest class of higher-dimensional systems, namely linear systems in two dimensions. These are interesting in their own right, and also play an important role in the classification of fixed points of nonlinear systems in two dimensions. A *two-dimensional linear system* has the form

$$\begin{aligned}\dot{x} &= f(x, y) \\ \dot{y} &= g(x, y)\end{aligned}\tag{A.12}$$

where f and g are linear functions of x and y . Since f and g are linear, we can express the system as

$$\begin{aligned}\dot{x} &= ax + by \\ \dot{y} &= cx + dy\end{aligned}\tag{A.13}$$

where $a, b, c, d \in \mathbb{R}$ are constants, or more compactly in matrix form as

$$\dot{\mathbf{x}} = A\mathbf{x},\tag{A.14}$$

where

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \text{ and } \mathbf{x} = \begin{pmatrix} x \\ y \end{pmatrix}.\tag{A.15}$$

The system is *linear* in the sense that if \mathbf{x}_1 and \mathbf{x}_2 are two solutions, then so is any linear combination $c_1\mathbf{x}_1 + c_2\mathbf{x}_2$. The zero vector $\mathbf{0}$ is always a fixed point for any choice of A .

The solutions of $\dot{\mathbf{x}} = A\mathbf{x}$ can be visualized as trajectories moving on the (x, y) plane, called the *phase plane*. The fixed point \mathbf{x}^* is an *attracting* fixed point if all trajectories that start near \mathbf{x}^* approach it as $t \rightarrow \infty$, i.e., $\mathbf{x} \rightarrow \mathbf{x}^*$ as $t \rightarrow \infty$. Given an attracting fixed point \mathbf{x}^* , the set of initial conditions \mathbf{x}_0 such that $\mathbf{x}(t) \rightarrow \mathbf{x}^*$ as $t \rightarrow \infty$ is called the *basin of*

attraction of \mathbf{x}^* . If \mathbf{x}^* attracts all trajectories in the phase plane, it is said to be *globally attracting*. A fixed point \mathbf{x}^* is *Lyapunov stable* if all trajectories that start sufficiently close to \mathbf{x}^* remain close to it at all times. When a fixed point is Lyapunov stable but not attracting, it is called *neutrally stable*. If a fixed point is both Lyapunov stable and attracting, it is called *stable*, or *asymptotically stable*. Finally, \mathbf{x}^* is *unstable* if it is neither attracting nor Lyapunov stable.

The *straight-line trajectories* are those that start on a line and stay on the line forever, exhibiting exponential growth or decay along the line. We would like to find the trajectories of the form

$$\mathbf{x}(t) = e^{\lambda t} \mathbf{v}, \quad (\text{A.16})$$

where $\mathbf{v} \neq \mathbf{0}$ is some fixed vector, and λ is a growth rate. If such solutions exist, they correspond to exponential motion along the line spanned by vector \mathbf{v} . To find the conditions on \mathbf{v} and λ , we substitute $\mathbf{x}(t) = e^{\lambda t} \mathbf{v}$ into $\dot{\mathbf{x}} = A\mathbf{x}$, and obtain $\lambda e^{\lambda t} \mathbf{v} = e^{\lambda t} A\mathbf{v}$ or

$$A\mathbf{v} = \lambda \mathbf{v}, \quad (\text{A.17})$$

which says that the desired straight line solution exists if \mathbf{v} is an *eigenvector* of A with the corresponding *eigenvalue* λ .

For a 2×2 matrix

$$A = \begin{pmatrix} a & b \\ c & d \end{pmatrix}, \quad (\text{A.18})$$

the eigenvalues λ_1 and λ_2 can be calculated as

$$\lambda_1 = \frac{\tau + \sqrt{\tau^2 - 4\Delta}}{2}, \quad \lambda_2 = \frac{\tau - \sqrt{\tau^2 - 4\Delta}}{2} \quad (\text{A.19})$$

where $\tau = \text{trace}(A) = a + d$, and $\Delta = \det(A) = ad - bc$.

When $\lambda_1 \neq \lambda_2^2$, any initial condition \mathbf{x}_0 can be written as a linear combination of eigenvectors, say $\mathbf{x}_0 = c_1 \mathbf{v}_1 + c_2 \mathbf{v}_2$. This allows us to write down the general solution for $\mathbf{x}(t)$ as

$$\mathbf{x}(t) = c_1 e^{\lambda_1 t} \mathbf{v}_1 + c_2 e^{\lambda_2 t} \mathbf{v}_2 = \sum_{i=1}^2 c_i e^{\lambda_i t} \mathbf{v}_i. \quad (\text{A.20})$$

For an n -dimensional linear dynamical system $\dot{\mathbf{x}} = A\mathbf{x}$, where A is an $n \times n$ matrix, the general solution is given by

$$\mathbf{x}(t) = \sum_{i=1}^n c_i e^{\lambda_i t} \mathbf{v}_i, \quad (\text{A.21})$$

where the \mathbf{v}_i 's are the eigenvectors and the λ_i 's are the corresponding eigenvalues of the matrix A .

In general, the eigenvalues and eigenvectors of the matrix can be complex. However, the constants c_i will also be complex in such a way that the final expression for $\mathbf{x}(t)$ is always real.

We are not concerned with the exact solution of a linear dynamical system because it will, in general, only be an approximation to the full nonlinear system that we are really interested in. What is of interest is the qualitative behavior of the system near the fixed point. The solution is the sum of terms involving $e^{\lambda t}$ where each complex eigenvalue λ can be written as $\lambda = \alpha + i\omega$. Thus, $e^{\lambda t} = e^{(\alpha+i\omega)t} = e^{\alpha t}(\cos \omega t + i \sin \omega t)$. Table A.8³ provides the classification of the fixed points of a linear dynamical system $\dot{\mathbf{x}} = A\mathbf{x}$ according to the real and imaginary parts of the eigenvalues of A .

²For simplicity, we ignore the “degenerate cases” in which there is a single repeated eigenvalue or one or more of the eigenvalues is zero.

³For a saddle point the trajectories move towards the fixed point in one direction called the *stable manifold* and away from it in another direction called the *unstable manifold*.

All α	ω	Nature of Fixed Point
< 0	$= 0$	Stable node
> 0	$= 0$	Unstable node
$\neq 0$	$= 0$	Hyperbolic or saddle point
< 0	$\neq 0$	Stable spiral
> 0	$\neq 0$	Unstable spiral
$= 0$	$\neq 0$	Center

Table A.8: Classification of fixed points of $\dot{\mathbf{x}} = A\mathbf{x}$ according to the real and imaginary parts of the eigenvalues of A . “All $\alpha \neq 0$ ” means some are positive and some negative.

A.4 Higher Dimensional Dynamical Systems

An n -dimensional nonlinear dynamical system given by

$$\begin{aligned}
 \dot{x}_1 &= f_1(x_1, \dots, x_n) \\
 \dot{x}_2 &= f_2(x_1, \dots, x_n) \\
 &\vdots \\
 \dot{x}_n &= f_n(x_1, \dots, x_n),
 \end{aligned}
 \tag{A.22}$$

where f_1, \dots, f_n are nonlinear functions. More compactly

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}),
 \tag{A.23}$$

where $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_n(\mathbf{x}))$. Here \mathbf{x} represents a point in n -dimensional phase space, and $\dot{\mathbf{x}}$ is the velocity vector at that point. The fixed points \mathbf{x}^* of the system satisfy $\mathbf{f}(\mathbf{x}^*) = \mathbf{0}$ and correspond to the equilibrium points of the system.

Linear Stability Analysis

Here we extend the *linearization* technique developed in Section A.2 for one-dimensional systems, so that we can approximate the phase portrait near a fixed point by that of a corresponding linear system.

Consider the nonlinear system given by Equation A.22, and suppose that $\mathbf{x}^* = (x_1^*, \dots, x_n^*)$ is a fixed point. Let

$$\begin{aligned}\eta_1 &= x_1 - x_1^* \\ \eta_2 &= x_2 - x_2^* \\ &\vdots \\ \eta_n &= x_n - x_n^*\end{aligned}\tag{A.24}$$

denote the components of a small perturbation from the fixed point. To see whether the perturbation grows or decays with time, we need to derive the differential equations for η_1, \dots, η_n .

$$\begin{aligned}\dot{\eta}_1 &= \dot{x}_1 \\ &= f_1(x_1^* + \eta_1, \dots, x_n^* + \eta_n) \\ &= f_1(x_1^*, \dots, x_n^*) + \eta_1 \frac{\partial f_1}{\partial x_1} + \dots + \eta_n \frac{\partial f_1}{\partial x_n} + O(\dots) \\ &\quad \text{(using Taylor series expansion)} \\ &= \eta_1 \frac{\partial f_1}{\partial x_1} + \dots + \eta_n \frac{\partial f_1}{\partial x_n} + O(\dots) \\ &\quad \text{(since } f_1(x_1^*, \dots, x_n^*) = 0\text{)}\end{aligned}\tag{A.25}$$

where that the partials are evaluated at (x_1^*, \dots, x_n^*) . Similarly,

$$\begin{aligned}\dot{\eta}_2 &= \eta_2 \frac{\partial f_2}{\partial x_1} + \dots + \eta_n \frac{\partial f_2}{\partial x_n} + O(\dots) \\ &\vdots \\ \dot{\eta}_n &= \eta_n \frac{\partial f_n}{\partial x_1} + \dots + \eta_n \frac{\partial f_n}{\partial x_n} + O(\dots).\end{aligned}\tag{A.26}$$

Hence the perturbation (η_1, \dots, η_n) evolves according to

$$\begin{pmatrix} \dot{\eta}_1 \\ \vdots \\ \dot{\eta}_n \end{pmatrix} = \begin{pmatrix} \frac{\partial f_1}{\partial x_1} & \dots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_n}{\partial x_1} & \dots & \frac{\partial f_n}{\partial x_n} \end{pmatrix} \begin{pmatrix} \eta_1 \\ \vdots \\ \eta_n \end{pmatrix} + O(\dots).\tag{A.27}$$

The $n \times n$ matrix in Equation A.27 is called the *Jacobian matrix* at the fixed point (x_1^*, \dots, x_n^*) , and will be denoted by $J(x_1^*, \dots, x_n^*)$. Since the higher order terms $O(\dots)$ in Equation A.27 are tiny, we can neglect them and obtain the linearized system

$$\begin{pmatrix} \dot{\eta}_1 \\ \vdots \\ \dot{\eta}_n \end{pmatrix} = \begin{pmatrix} \frac{\partial f_1}{\partial x_1} & \dots & \frac{\partial f_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_n}{\partial x_1} & \dots & \frac{\partial f_n}{\partial x_n} \end{pmatrix} \begin{pmatrix} \eta_1 \\ \vdots \\ \eta_n \end{pmatrix}\tag{A.28}$$

whose dynamics can be analyzed using methods from Section A.3. It is safe to neglect the higher order terms as long as the fixed points for the dynamical system are not any of the borderline cases (centers, degenerate nodes, stars, or non-isolated fixed points).

When the fixed points of the system are hyperbolic, they are sturdy, i.e., their stability type is unaffected by small nonlinear terms. According to the Hartman-Grobman theorem (stated below), the local phase portrait near a hyperbolic fixed point is “topologically equivalent”⁴ to the phase portrait of the linearization about it. In other words, the theorem

⁴There is a homeomorphism—a continuous map with a continuous inverse—from one local phase portrait onto the other such that the trajectories of the first map onto trajectories of the second and the sense of time is preserved.

justifies the use of linearization approach to discover the properties of fixed points in a dynamical system in most cases. The exceptions are situations where one or more eigenvalues are purely imaginary. In such cases, the stability (or otherwise) of the fixed point must be determined by other means.

Theorem A.2. (Hartman-Grobman theorem) If x^* is a hyperbolic fixed point for the system of autonomous differential equations (A.22), then there is an open set U containing x^* and a homeomorphism h with domain U , such that the solutions of the differential equations (A.22) are mapped by h to solutions of the linearized system (A.27) in U .

The following theorem provides the conditions for the existence and uniqueness of solution for an n -dimensional nonlinear dynamical system.

Theorem A.3. (Existence and uniqueness theorem) Consider the initial value problem

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}), \quad \mathbf{x}(0) = \mathbf{x}_0. \quad (\text{A.29})$$

Suppose that $\mathbf{f}(\mathbf{x})$ is continuous and that all of its partial derivatives $\frac{\partial f_i}{\partial x_j}$ for $i, j = 1, \dots, n$, are continuous for \mathbf{x} in some open connected set $D \subset \mathbb{R}^n$. Then for $\mathbf{x}_0 \in D$, the initial value problem has a solution $\mathbf{x}(t)$ on some time interval $(-\tau, \tau)$ about $t = 0$, and the solution is unique.

To obtain the qualitative understanding of a nonlinear dynamical system with hyperbolic fixed points, we must follow a four-step procedure:

1. Find the fixed points of the nonlinear system.
2. Derive the linear approximation of the system close to each fixed point.
3. Determine the properties of the fixed points in the linearized system (see Table A.8) based on the real and imaginary components α and ω of the eigenvalues λ of the linearized system.

4. Combine this information to produce the sketch of the full, nonlinear system.

We illustrate this with an example.

Example A.4. Consider the two-dimensional dynamical system

$$\begin{aligned}\dot{x} &= x(1-x)(1-2y) \\ \dot{y} &= y(1-y)(1-2x)\end{aligned}\tag{A.30}$$

defined on the unit square (i.e., $0 \leq x, y \leq 1$). The fixed points (x^*, y^*) of this system are the set of points $\{(0, 0), (0, 1), (1, 0), (1, 1), (\frac{1}{2}, \frac{1}{2})\}$. We next linearize the system close to each of the fixed points and determine their nature.

- $(x^*, y^*) = (0, 0)$. The Jacobian matrix at this fixed point is

$$J(0, 0) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},\tag{A.31}$$

with eigenvalues $\lambda_1 = \lambda_2 = 1$. The fixed point is thus an unstable node.

- $(x^*, y^*) = (0, 1)$. The Jacobian matrix at this fixed point is

$$J(0, 1) = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix},\tag{A.32}$$

with eigenvalues $\lambda_1 = \lambda_2 = -1$. The fixed point is thus a stable node.

- $(x^*, y^*) = (1, 0)$. The Jacobian matrix at this fixed point is

$$J(1, 0) = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix},\tag{A.33}$$

with eigenvalues $\lambda_1 = \lambda_2 = -1$. The fixed point is thus a stable node.

- $(x^*, y^*) = (1, 1)$. The Jacobian matrix at this fixed point is

$$J(1, 1) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad (\text{A.34})$$

with eigenvalues $\lambda_1 = \lambda_2 = 1$. The fixed point is an unstable node.

- $(x^*, y^*) = (\frac{1}{2}, \frac{1}{2})$. The Jacobian matrix at this fixed point is

$$J\left(\frac{1}{2}, \frac{1}{2}\right) = \begin{pmatrix} 0 & -\frac{1}{2} \\ -\frac{1}{2} & 0 \end{pmatrix}, \quad (\text{A.35})$$

with eigenvalues $\lambda_1 = \frac{1}{2}$ and $\lambda_2 = -\frac{1}{2}$, and corresponding eigenvectors

$$\mathbf{v}_1 = \begin{pmatrix} 1 \\ -1 \end{pmatrix} \quad \text{and} \quad \mathbf{v}_2 = \begin{pmatrix} 1 \\ 1 \end{pmatrix}. \quad (\text{A.36})$$

The fixed point is thus hyperbolic with stable manifold $x = y$ and unstable manifold $x = -y$.

The behavior of the (linearized) system near each of the fixed points is shown in Figure A.98(a). Using Theorem A.2, we obtain the qualitative behavior of the full system, which is shown in Figure A.98(b). □

As we have already remarked, the stability (or otherwise) of a non-hyperbolic fixed point cannot be determined by linearization. The linearized system about a non-hyperbolic fixed point has eigenvalues with real parts that are zero, and the stability of the system is determined by higher order terms in the Taylor series expansion about the fixed point. In such cases we can make use of the *direct Lyapunov method*.

Theorem A.4. (Lyapunov stability theorem) Let $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$ be a dynamical system with a fixed point at \mathbf{x}^* . If we can find a scalar function $V(\mathbf{x})$, defined for the allowed states of the system close to \mathbf{x}^* , such that

1. $V(\mathbf{x}^*) = 0$
2. $V(\mathbf{x}) > 0$ for $\mathbf{x} \neq \mathbf{x}^*$
3. $\frac{dV}{dt} < 0$ for $\mathbf{x} \neq \mathbf{x}^*$

then the fixed point \mathbf{x}^* is asymptotically stable.





Figure A.98: Behavior of the system A.30 near fixed points and as a whole. (a) Behavior near fixed points. (b) Behavior as a whole.

APPENDIX B

SIMULATION SOFTWARE

In the computer field, the moment of truth is a running program; all else is prophecy.

- Herbert Simon

B.1 Introduction

In this appendix we discuss the software we have implemented and used for simulating the agent-based models described in Sections 4.3, 5.3, and 6.2.4. We provide information on where to obtain the software from, what is contained in the distribution, and how to configure and run the software. The code is made available under the GNU General Public License.

B.2 Obtaining the Software

We provide two different implementations of our agent-based models: one in Python and the other in Java. In order to run our Python code you must have the following software installed on your computer:

- Python (v2.7.3),
- NumPy (v1.6.2),
- SciPy (v0.11.0),
- matplotlib (v1.2), and

- NetworkX (v1.7);

and to run our Java code you must have the following:

- Java (version 7 update 9),
- google-gson (v2.2.2)¹, and
- JUNG2 (v2.0.1)².

The numbers in parentheses indicate the versions of the software that our code has been tested against.

There are a few differences between the Python- and Java-based implementations, some stemming from the differences in the implementation languages—Python and Java in our case—and others from specific design choices. For example, the Java-based implementation does not provide plotting capabilities; instead, we supply Python scripts to generate plots from the results. The implementations also differ in terms of their performance. The Java-based implementation runs much faster than the Python-based implementation since Java is a compiled language and also because the Java Virtual Machine (JVM) implements Just-In-Time (JIT) compilation for the native architecture³.

The Python- and Java-based implementations are packaged into a single archive `edcn.tar.gz` (`edcn` is short for evolutionary dynamics on complex networks), which can be downloaded from the following website:

<http://www.cs.umb.edu/~swamir/research.html>

¹`gson-xxx.jar` must be included in the CLASSPATH environment variable.

²`collections-generic-xxx.jar`, `jung-algorithms-xxx.jar`, `jung-io-xxx.jar`, `jung-api-xxx.jar`, and `jung-graph-impl-xxx.jar` must be included in the CLASSPATH environment variable.

³Performance of our Python code can be improved by running it using PyPy, a fast alternative implementation of the Python language.

Extracting the archive into a folder (say `/home/darwin`) on your file system produces a folder `edcn` under `/home/darwin`. Throughout this appendix we use the alias `$edcn` to refer to the folder `/home/darwin/edcn`. The folders⁴ `$edcn/discrete/python`, `$edcn/continuous/python`, and `$edcn/disease/python` contain Python code for simulating symmetric, two-strategy, pairwise games (Section 4.3), continuous games (Section 5.3), and disease dynamics (Section 6.2.4) respectively. The corresponding Java versions are under the folders⁵ `$edcn/discrete/java`, `$edcn/continuous/java`, and `$edcn/disease/java`. The following sections describe the Python- and Java-based implementations in detail.

B.3 Common Utilities

Python Code

The folder `$edcn/util/python` contains common utilities used by rest of our Python code. The module `network.py` provides a high-level abstraction (`Vertex`) for representing vertices in a network, and several abstractions (`Network` and its sub-classes) for representing different kinds of networks. This module is used by the simulation code to build complex networks. The module supports several network topologies, and the ones relevant to this thesis are listed in Table B.9, along with the corresponding parameters.

Table B.9: Some of the network topologies supported by the `network.py` module, along with the corresponding parameters.

Network Topology	Parameters
------------------	------------

⁴These folders must be included in the `PYTHONPATH` environment variable.

⁵These folders must be included in the `CLASSPATH` environment variable.

<p>Complete: builds a complete network Γ_{WM} in which every vertex is connected to every other vertex in the network (Section 2.4.1)</p>	
<p>Erdos_Renyi: builds a network Γ_{ER} with a binomial/Poisson degree distribution (Section 2.4.2)</p>	<p>vertices: list of <code>Vertex</code> objects representing the vertices in the network p: probability that two randomly picked vertices have an edge between them seed: seed for the random number generator used to build the network</p>
<p>Barabasi_Albert: builds a network Γ_{BA} with a power-law degree distribution (Section 2.4.4)</p>	<p>vertices: list of <code>Vertex</code> objects representing the vertices in the network m: number of attachments made from the new vertex to the existing vertices during the network construction process seed: seed for the random number generator used to build the network</p>

<p>Powerlaw_Homophilic: builds a network Γ_{BA}^h with a power-law degree distribution (Section 2.4.6) and with homophily</p>	<p>vertices: list of <code>Vertex</code> objects representing the vertices in the network</p> <p>m: number of attachments made from the new vertex to the existing vertices during the network construction process</p> <p>h: desired value for homophily</p> <p>maxiter: maximum number of re-wiring attempts</p> <p>seed: seed for the random number generator used to build the network</p>
<p>Powerlaw_Clustered: builds a network Γ_{BA}^C with a power-law degree distribution (Section 2.4.5) and with clustering</p>	<p>vertices: list of <code>Vertex</code> objects representing the vertices in the network</p> <p>m: number of attachments made from the new vertex to the existing vertices during the network construction process</p> <p>p: probability of adding a triangle after adding a random edge</p> <p>seed: seed for the random number generator used to build the network</p>
<p>Random_Regular: builds a random, generally non-planar network Γ_{RR} (Section 2.4.3) in which every vertex has the specified degree</p>	<p>vertices: list of <code>Vertex</code> objects representing the vertices in the network</p> <p>degree: degree of each vertex</p> <p>seed: seed for the random number generator used to build the network</p>

GraphML_File: builds a network from a GraphML file	name: name of the GraphML file
----------------------------------------------------	--------------------------------

The script `gen_network.py` can be used to generate any of the networks listed in Table B.9 and save it as a GraphML file. The script supports the following command-line syntax:

```
python gen_network.py <topology> <size> <params> <outfile>
```

where `topology` is the desired network topology, `size` is the network size, `params` is a list of parameters for the network, and `outfile` is the name of the output GraphML file. For example, the following command produces a 10000-vertex Barabási-Albert network with average degree 4 using a seed of 42, and saves the network in `network.graphml`.

```
> python gen_network.py "Barabasi_Albert" 10000 \
"{\"m\":2,\"seed\":42}" "network.graphml"
```

The script `network_stats.py` can be used to print to STDOUT the values of some basic properties of the input network in GraphML format. The script supports the following command-line syntax:

```
> python network_stats.py <infile>
```

For example, the following command prints the basic properties of the network generated above:

```
> python network_stats.py network.graphml
number of vertices = 10000
number of edges = 19996
average degree = 3.999200
number of components = 1
```

```
fractional size of largest component = 1.000000
diameter = 5.097388
average clustering coefficient = 0.003618
homophily = -0.051038
```

Java Code

The folder `$edcn/util/java` contains common utilities used by rest of our Java code. The module `Network.java` provides a high-level abstraction (`Network`) for a network, and in particular, abstractions (`CompleteNetwork` and `GraphMLNetwork`) for dealing with complete networks and networks built from GraphML files. The module `Random.java` supports functions for random number generation, and the module `Stats.java` supports basic statistical computations.

B.4 Simulating Discrete Games

Python Code

The folder `$edcn/discrete/python` contains Python code for simulating symmetric, two-strategy, pairwise games (Section 4.3). The script `experiment.py` provides a template for defining the simulation parameters (see Table B.10). Once the parameters for a game are defined, the game can be simulated using the following command:

```
> python experiment.py [--realtime]
```

where the optional `--realtime` argument runs the simulation in real-time mode instead of the default off-line mode.

Table B.10: Parameters for simulating symmetric, two-strategy, pairwise games using our Python code.

Parameter	Description
payoff	2×2 payoff matrix π for the game, specified as $[[a, b], [c, d]]$
selection_strength	selection strength s
population	population size n
init_cooperators	initial fraction x_0 of cooperators
generations	number of generations T
report_freq	frequency ν (in generations) at which the results are saved/displayed
assortativity	probability r with which an individual interacts with another individual playing the same strategy as its own; only relevant for complete networks
update_rule	update rule \mathcal{U} for the evolutionary dynamics; must be one of FE1, FE2, RE1, RE2, IM, BD, or DB; FE2 and RE2 work with simultaneous interactions and updates (Algorithm 4.3) and the rest work with sequential interactions and updates (Algorithm 4.2)
seed	seed for the random number generator
network_topology	network topology Γ (see Table B.9)
network_params	network parameters (see Table B.9)

The script `discrete_omain.py` runs the simulation in off-line mode. The results of the simulation are saved in `results.pkl`, and a plot showing the evolution of the fraction x of cooperators is saved in `plot.pdf`. The inset in the plot shows the value of x averaged

over the last 10% of the generations, denoted by x_∞ . The script can also be used to extend a previously run simulation to include more generations. This can be done by updating the generations parameter in the `experiment.py` file for the simulation to a new value and running the modified file using the above command.

The script `discrete_rmain.py` runs the simulation in real-time mode. The results of the simulation are not saved in this case, but are rendered in real time as a plot on the screen showing the evolution of x . The script can also be used to re-play an off-line simulation.

Example B.1. Consider the following file `dpd.py`⁶ that defines the parameters for simulating a discrete prisoner's dilemma game with $\rho = 0.3$, i.e., with the payoff matrix

$$\pi = \begin{array}{c} C \quad D \\ \begin{array}{cc} C & (0.7, 0.7) & (-0.3, 1) \\ D & (1, -0.3) & (0, 0) \end{array} \end{array}, \quad (\text{B.1})$$

on a complete network Γ_{WM} . The parameters are $n = 10000$, $s = 1$, $x_0 = 0.5$, $r = 0.4$, $T = 5000$, $\nu = 1$, $\tau = 1$, and $\mathcal{U} = \text{FE2}$:

```
# dpd.py
import discrete, discrete_omain, discrete_rmain, sys, time
def main(args):
    params = {
        "payoff" : [[0.7, -0.3], [1.0, 0.0]],
        "selection_strength" : 1.0,
        "population" : 10000,
        "init_cooperators" : 0.5,
        "generations" : 5000,
        "report_freq" : 1,
        "assortativity" : 0.4,
```

⁶This file was created from the template `$edcn/discrete/python/experiment.py`.

```

    "update_rule" : "FE2",
    "seed" : int(time.time()),
    "network_topology" : "Complete",
    "network_params" : {}
}
if len(args) == 1 and args[0] == "--realtime":
    discrete_rmain.run(params)
else:
    discrete_omain.run(params)
if __name__ == "__main__":
    main(sys.argv[1:])

```

We simulate the game using the following command:

```
> python dpd.py
```

which produces `results.pkl` containing the results of the simulation, and `plot.pdf` showing the evolution (see Figure B.99) of x . □

Java Code

The folder `$edcn/discrete/java` contains Java code for simulating symmetric, two-strategy, pairwise games (Section 4.3). `params.json` provides a template for defining the simulation parameters (see Table B.11), and `Params.java` facilitates reading them. Once the parameters for a game are defined, the game can be simulated using the following command:

```
> java Discrete params.json > results.txt
```

which saves the results (averaged over all trials) of the simulation in `results.txt`.




Figure B.99: Evolution of the fraction x of cooperators in a discrete prisoner’s dilemma game. Simulated on a complete network Γ_{WM} using parameters $n = 10000$, $\rho = 0.3$, $s = 1$, $x_0 = 0.5$, $r = 0.4$, $T = 5000$, $\nu = 1$, and $\mathcal{U} = \text{FE2}$.

Table B.11: Parameters for simulating symmetric, two-strategy, pairwise games using our Java code.

Parameter	Description
payoff	2×2 payoff matrix π for the game, specified as $[[a, b], [c, d]]$
selection_strength	selection strength s
population	population size n
init_cooperators	initial fraction x_0 of cooperators
generations	number of generations T per trial
report_freq	frequency ν (in generations) at which the results are saved
trials	number of trials τ

assortativity	probability r with which an individual interacts with another individual playing the same strategy as its own; only relevant for complete networks
update_rule	update rule \mathcal{U} for the evolutionary dynamics; must be one of FE1, FE2, RE1, RE2, IM, BD, or DB; FE2 and RE2 work with simultaneous interactions and updates (Algorithm 4.3) and the rest work with sequential interactions and updates (Algorithm 4.2)
network_topology	network topology Γ ; must be Complete or GraphML_File
graphml_file	null or name of the GraphML file

The Python script `plots.py` can be used to generate `plot.pdf` showing the evolution of the fraction x of cooperators from `results.txt`. The command-line syntax for the script is as follows:

```
python plots.py <results file> <start time> <end time>
```

where `<results file>` is the name of the file storing the simulation results produced by `Discrete`, and `<start time>` and `<end time>` specify the time frame (in generations) for the plot. The inset in the plot shows the value of x averaged over the last 10% of the generations, denoted by x_∞ .

Example B.2. Consider the following file `dpd.json`⁷ that defines the parameters for simulating the discrete prisoner’s dilemma game from Example B.1:

```
// dpd.json
{
  payoff : [[0.7, -0.3], [1.0, 0.0]],
```

⁷This file was created from the template `$edcn/discrete/java/params.json`.

```
selection_strength : 1.0,  
population : 10000,  
init_cooperators : 0.5,  
generations : 5000,  
report_freq : 1,  
trials : 1,  
assortativity : 0.4,  
update_rule : "FE2",  
network_topology : "Complete",  
graphml_file : null  
}
```

We simulate the game using the following command:

```
> java Discrete dpd.json > results.txt
```

which produces `results.txt` containing the results of the simulation. The following command:

```
> python plots.py results.txt 0 600
```

produces `plot.pdf` showing the evolution (see Figure B.100) of x from generation $t = 0$ to generation $t = 600$. □

B.5 Simulating Continuous Games

Python Code

The folder `$edcn/continuous/python` contains Python code for simulating continuous games (Section 5.3). The script `experiment.py` provides a template for defining the




Figure B.100: Evolution of the fraction x of cooperators in a discrete prisoner’s dilemma game. Simulated on a complete network Γ_{WM} using parameters $n = 10000$, $\rho = 0.3$, $s = 1$, $x_0 = 0.5$, $r = 0.4$, $T = 5000$, $\nu = 1$, and $\mathcal{U} = \text{FE2}$.

simulation parameters (see Table B.12). Once the parameters for a game are defined, the game can be simulated using the following command:

```
> python experiment.py [--realtime]
```

where the optional `--realtime` argument runs the simulation in real-time mode instead of the default off-line mode.

Table B.12: Parameters for simulating continuous games using our Python code.

Parameter	Description
payoff	payoff function π for the continuous game
selection_strength	selection strength s
population	population size n
generations	number of generations T

report_freq	frequency ν (in generations) at which results are saved/displayed
init_trait	initial strategy x_0
max_trait	maximum strategy x_m
assortativity	probability r with which an individual interacts with another individual playing the same strategy as its own; only relevant for complete networks
group	number g of neighbors each individual interacts with during each round of interactions; if None, each individual will interact with its entire neighborhood; the average value of the neighbors' strategies is used in calculating the payoff; note that this setting is not applicable for the FE2 and RE2 update rules, and must be set to 1 if r is positive
update_rule	update rule \mathcal{U} for the evolutionary dynamics; must be one of FE1, FE2, RE1, RE2, IM, BD, or DB; FE2 and RE2 work with simultaneous interactions and updates (Algorithm 5.5) and the rest work with sequential interactions and updates (Algorithm 5.4)
mutation	rate μ for Gaussian mutation of strategies
stddev	standard deviation σ for the mutations
seed	seed for the random number generator
network_topology	network topology Γ (see Table B.9)
network_params	network parameters (see Table B.9)

The script `continuous_omain.py` runs the simulation in off-line mode. The results of the simulation are saved in `results.pkl`, a plot showing the evolution of mean strategy \bar{x} is saved in `plot1.pdf`, and a plot showing the evolution of the distribution of strategies is saved in `plot2.pdf`. The inset in `plot1.pdf` shows the average value of the mean strategy \bar{x} over the last 10% of the generations, denoted by \bar{x}_∞ . The script can also be used to extend a previously run simulation to include more generations. This can be done by updating the `generations` parameter in the `experiment.py` file for the simulation to a new value and running the modified file using the above command.

The script `continuous_rmain.py` runs the simulation in real-time mode. The results of the simulation are not saved in this case, but are rendered in real time as plots on the screen, one showing the evolution of the strategy distribution, and the other showing the strategy histogram as it varies with the time. The script can also be used to re-play an off-line simulation.

Example B.3. Consider the following file `ctoc.py`⁸ that defines the parameters for simulating the continuous tragedy of the commons game with payoff function $\pi(x, y) = B(x) - C(x + y)$, where $B(x) = -\frac{1}{12}x^3 + 2x^2 + x$ and $C(x) = x^2$, on a complete network Γ_{WM} . The parameters are $n = 10000$, $s = 1$, $x_0 = 1.5$, $x_m = 4$, $r = 0$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $v = 100$, and $\mathcal{U} = \text{FE2}$:

```
# ctoc.py
import continuous, continuous_omain, continuous_rmain, sys, time
def main(args):
    def B(x):
        return -0.0834 * x ** 3 + 2.0 * x ** 2 + x
    def C(x):
        return x ** 2
    payoff = lambda x, y: B(x) - C(x + y)
```

⁸This file was created from the template `$edcn/continuous/python/experiment.py`.

```

params = {
    "selection_strength" : 1.0,
    "population" : 10000,
    "generations" : 300000,
    "report_freq" : 100,
    "init_trait" : 1.5,
    "max_trait" : 4.0,
    "assortativity" : 0.0,
    "group" : None,
    "update_rule" : "FE2",
    "mutation" : 0.01,
    "stddev" : 0.005,
    "seed" : int(time.time()),
    "network_topology" : "Complete",
    "network_params" : {}
}

if len(args) == 1 and args[0] == "--realtime":
    continuous_rmain.run(params, payoff)
else:
    continuous_omain.run(params, payoff)

if __name__ == "__main__":
    main(sys.argv[1:])

```

We simulate the game using the following command:

```
> python ctoc.py
```

which produces `results.pkl` containing the results of the simulation, `plot1.pdf` showing the evolution (see Figure B.101(a)) of \bar{x} , and `plot2.pdf` showing the evolution (see Figure B.101(b)) of the distribution of strategies. □




Figure B.101: Evolution of the mean strategy \bar{x} and the distribution of strategies x in a continuous tragedy of the commons game. Simulated on a complete network Γ_{WM} using parameters $\pi(x, y) = B(x) - C(x + y)$, where $B(x) = -\frac{1}{12}x^3 + 2x^2 + x$ and $C(x) = x^2$, $n = 10000$, $s = 1$, $x_0 = 1.5$, $x_m = 4$, $r = 0$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $v = 100$, and $\mathcal{U} = \text{FE2}$. (a) Evolution of \bar{x} . (b) Evolution of the distribution of x .

Java Code

The folder `$edcn/continuous/java` contains Java code for simulating continuous games (Section 5.3). `params.json` provides a template for defining the simulation parameters (see Table B.13), and `Params.java` facilitates reading them. Once the parameters for a game are defined, the game can be simulated using the following command:

```
> java Continuous params.json > results.txt
```

which saves the results of the simulation in `results.txt`.

Table B.13: Parameters for simulating continuous games using our Java code.

Parameter	Description
-----------	-------------

payoff	name of the payoff function π (Payoff1 or Payoff2 for prisoner's dilemma, Payoff3 for snowdrift game, and Payoff4 for tragedy of the commons) for the game
payoff_params	parameters (see Payoff*.java for details) for the payoff function specified as $[a, b, \dots]$
selection_strength	selection strength s
population	population size n
generations	number of generations T
report_freq	frequency v (in generations) at which results are saved/displayed
init_trait	initial strategy x_0
max_trait	maximum strategy x_m
assortativity	probability r with which an individual interacts with another individual playing the same strategy as its own; only relevant for complete networks
group	number g of neighbors each individual interacts with during each round of interactions; if 0, each individual will interact with its entire neighborhood; the average value of the neighbors' strategies is used in calculating the payoff; note that this setting is not applicable for the FE2 and RE2 update rules, and must be set to 1 if r is positive

update_rule	update rule \mathcal{U} for the evolutionary dynamics; must be one of FE1, FE2, RE1, RE2, IM, BD, or DB; FE2 and RE2 work with simultaneous interactions and updates (Algorithm 5.5) and the rest work with sequential interactions and updates (Algorithm 5.4)
mutation	rate μ for Gaussian mutation of strategies
stddev	standard deviation σ for the mutations
network_topology	network topology Γ ; must be Complete or GraphML_File
graphml_file	null or name of the GraphML file

The modules `Payoff1.java` and `Payoff2.java` define the payoff functions for the prisoner’s dilemma game, `Payoff3.java` defines the payoff function for the snowdrift game, and `Payoff4.java` defines the payoff function for the tragedy of the commons game.

The Python script `plots.py` can be used to generate `plot1.pdf` showing the evolution of mean strategy \bar{x} , and `plot2.pdf` showing the evolution of the distribution of strategies. The command-line syntax for the script is as follows:

```
python plots.py <results file> <start time> <end time> [<xstar>]
```

where `<results file>` is the name of the file storing the simulation results produced by `Continuous`, `<start time>` and `<end time>` specify the time frame (in generations) for the plots, and the optional `<xstar>` argument specifies the singular strategy—if one exists—for the game, indicated in the plots by a dashed line. The inset in `plot1.pdf` shows the average value of the mean strategy \bar{x} over the last 10% of the generations, denoted by \bar{x}_∞ .

Example B.4. Consider the following file `ctoc.json`⁹ that defines the parameters for simulating the continuous tragedy of the commons game from Example B.3:

```
// ctoc.json
{
  payoff : "Payoff4",
  payoff_params : [1.0, 2.0, 0.0834, 1.0],
  selection_strength : 1.0,
  population : 10000,
  generations : 300000,
  report_freq : 100,
  init_trait : 1.5,
  max_trait : 4.0,
  assortativity : 0.0,
  group : 0,
  update_rule : "FE2",
  mutation : 0.01,
  stddev : 0.005,
  network_topology : "Complete",
  graphml_file : null
}
```

We simulate the game using the following command:

```
> java Continuous ctoc.json > results.txt
```

which produces `results.txt` containing the results of the simulation. The following command:

```
> python plots.py results.txt 0 200000 2.0
```

⁹This file was created from the template `$edcn/continuous/java/params.json`.

produces `plot1.pdf` showing the evolution (see Figure B.102(a)) of \bar{x} , and `plot2.pdf` showing the evolution (see Figure B.102(b)) of the distribution of strategies, each from generation $t = 0$ to generation $t = 200000$. □




Figure B.102: Evolution of the mean strategy \bar{x} and the distribution of strategies x in a continuous tragedy of the commons game. Simulated on a complete network Γ_{WM} using parameters $\pi(x, y) = B(x) - C(x + y)$, where $B(x) = -\frac{1}{12}x^3 + 2x^2 + x$ and $C(x) = x^2$, $n = 10000$, $s = 1$, $x_0 = 1.5$, $x_m = 4$, $r = 0$, $\mu = 0.01$, $\sigma = 0.005$, $T = 300000$, $v = 100$, and $\mathcal{U} = \text{FE2}$. (a) Evolution of \bar{x} . (b) Evolution of the distribution of x .

B.6 Simulating Disease Dynamics

Python Code

The folder `$edcn/disease/python` contains Python code for simulating disease dynamics (Section 6.2.4). The script `experiment.py` provides a template for defining the simulation parameters (see Table B.14). Once the parameters for a disease are defined, the disease dynamics can be simulated using the following command:

```
> python experiment.py [--realtime]
```

where the optional `--realtime` argument runs the simulation in real-time mode instead of the default off-line mode.

Table B.14: Parameters for simulating disease dynamics using our Python code.

Parameter	Description
population	population size n
generations	number of generations T
report_freq	frequency ν (in generations) at which results are saved/displayed
infection	infection probability β of the single infected individual at generation $t = 0$
b, a	constants β_0 and α in the function $\beta_i = \beta_0 c_i^\alpha$ relating infection probability β_i of individual i to its virulence c_i
recovery	recovery probability γ
death	probability u of dying due to natural causes
mutation	rate μ for the Gaussian mutation of infection probability
stddev	standard deviation σ for the Gaussian mutations
seed	seed for the random number generator
network_topology	network topology Γ (see Table B.9)
network_params	network parameters (see Table B.9)

The script `disease_omain.py` runs the simulation in off-line mode. The results of the simulation are saved in `results.pkl`, a plot showing the evolution of the fraction of susceptible s , infected x , and recovered r individuals, is saved in `plot1.pdf`, and a plot showing the evolution of the mean virulence \bar{c} and its average value \bar{c}_∞ over the last 10%

of the generations is saved in `plot2.pdf`. The script can also be used to extend a previously run simulation to include more generations. This can be done by updating the `generations` parameter in the `experiment.py` file for the simulation to a new value and running the modified file using the above command.

The script `disease_rmain.py` runs the simulation in real-time mode. The results of the simulation are not saved in this case, but are rendered in real time as plots on the screen, one showing the evolution of s , x , and r , and the other showing the evolution of the mean virulence \bar{c} . The script can also be used to re-play an off-line simulation.

The script `r0_plot.py` produces `r0.pdf` showing how the basic reproductive ratio R_0 varies with the virulence c . The command-line syntax for the script is:

```
python r0_plot.py <a> <b> <recovery> <death>
```

The module `disease.py` can be used for studying the effects of vaccination on the spread of a disease. This is done by simulating the disease dynamics for T_1 generations, vaccinating a fraction of the population using one of three strategies (degree, referral, random), and then continuing the simulation for T_2 generations. Vaccinations are carried out on the command line by running one of the following functions defined in `disease.py`:

- `state_stats(infile)` lists the number of individuals in the susceptible, infected, recovered, and vaccinated states in `infile`.
- `degree_vaccination(infile, outfile, v)` vaccinates the top v percent of the highly-connected individuals in `infile`, and saves the changes in `outfile`.
- `referral_vaccination(infile, outfile, v)` vaccinates v percent of the population in `infile`—each vaccinated individual is a randomly picked neighbor of a randomly picked individual—and saves the changes in `outfile`.

- `random_vaccination(infile, outfile, v)` vaccinates v percent of the population in `infile` uniformly at random, and saves the changes in `outfile`.

The `infile` and `outfile` arguments to each of the above functions denote input and output (`results.pkl`) files that store simulation results.

Example B.5. Consider the following file `dd.py`¹⁰ that defines the parameters for simulating disease dynamics on a complete network Γ_{WM} . The parameters are $n = 10000$, $\beta_0 = 1$, $\alpha = 0.5$, $\beta = 0.075$, $\gamma = 0.08$, $d = 0.01$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 100000$, and $v = 100$:

```
# dd.py
import disease, disease_omain, disease_rmain, sys, time
def main(args):
    params = {
        "population" : 10000,
        "generations" : 100000,
        "report_freq" : 100,
        "infection" : 0.075,
        "b" : 1.0,
        "a" : 0.5,
        "recovery" : 0.08,
        "death" : 0.01,
        "mutation" : 0.0001,
        "stddev" : 0.05,
        "seed" : int(time.time()),
        "network_topology" : "Complete",
        "network_params" : {}
    }
    if len(args) == 1 and args[0] == "--realtime":
        disease_rmain.run(params)
```

¹⁰This file was created from the template `$edcn/disease/python/experiment.py`.

```

else:
    disease_omain.run(params)
if __name__ == "__main__":
    main(sys.argv[1:])

```

We simulate the disease dynamics using the following command:

```
> python dd.py
```

which produces results.pkl containing the results of the simulation, plot1.pdf showing the evolution (see Figure B.103(a)) of s , x , and r , and plot2.pdf showing the evolution (see Figure B.103(b)) of \bar{c} . □




Figure B.103: Evolution of the fraction of susceptible s , infected x , and recovered r individuals, and evolution of the mean virulence \bar{c} . Simulated on a complete Γ_{WM} using parameters $n = 10000$, $\beta_0 = 1$, $\alpha = 0.5$, $\beta = 0.075$, $\gamma = 0.08$, $d = 0.01$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 100000$, and $\nu = 100$. (a) Evolution of s , x , and r . (b) Evolution of \bar{c} .

Java Code

The folder `$edcn/disease/java` contains Java code for simulating disease dynamics (Section 6.2.4). `params.json` provides a template for defining the simulation parameters (see Table B.15), and `Params.java` facilitates reading them. Once the parameters for a disease are defined, the disease dynamics can be simulated using the following command:

```
> java Disease params.json > results.txt
```

which saves the results of the simulation in `results.txt`.

Table B.15: Parameters for simulating disease dynamics using our Java code.

Parameter	Description
population	population size n
generations	number of generations
report_freq	frequency (in generations) at which results are saved/displayed
infection	infection probability β of the single infected individual at generation $t = 0$
b, a	constants β_0 and α in the function $\beta_i = \beta_0 c_i^\alpha$ relating infection probability β_i of individual i to their virulence c_i
recovery	recovery probability γ
death	probability u of dying due to natural causes
mutation	rate μ for Gaussian mutation of infection probability
stddev	standard deviation σ for the mutations
seed	seed for the random number generator
network_topology	network topology Γ ; must be Complete or GraphML_File
graphml_file	null or name of the GraphML file

The Python script `plots.py` can be used to generate `plot1.pdf` showing the evolution of the fraction of susceptible s , infected x , and recovered r individuals, and `plot2.pdf` showing the evolution of the mean virulence \bar{c} and its average value \bar{c}_∞ over the last 10% of the generations. The command-line syntax for the script is as follows:

```
python plots.py <results file> <start time> <end time>
```

where `<results file>` is the name of the file storing the simulation results produced by `Disease`, and `<start time>` and `<end time>` specify the time frame (in generations) for the plots.

Example B.6. Consider the following file `dd.json`¹¹ that defines the parameters for simulating disease dynamics from Example B.5:

```
// dd.json
{
  population : 10000,
  generations : 100000,
  report_freq : 100,
  infection : 0.075,
  b : 1.0,
  a : 0.5,
  recovery : 0.08,
  death : 0.01,
  mutation : 0.0001,
  stddev : 0.05,
  network_topology : "Complete",
  graphml_file : null
}
```

We simulate the disease dynamics using the following command:

¹¹This file was created from the template `$edcn/disease/java/params.json`.

```
> java Disease dd.json > results.txt
```

which produces `results.txt` containing the results of the simulation. The following command:

```
> python plots.py results.txt 0 40000
```

produces `plot1.pdf` showing the evolution (see Figure B.104(a)) of s , x , and r , and `plot2.pdf` showing the evolution (see Figure B.104(b)) of \bar{c} , each from generation $t = 0$ to generation $t = 40000$. □




Figure B.104: Evolution of the fraction of susceptible s , infected x , and recovered r individuals, and evolution of the mean virulence \bar{c} . Simulated on a complete Γ_{WM} using parameters $n = 10000$, $\beta_0 = 1$, $\alpha = 0.5$, $\beta = 0.075$, $\gamma = 0.08$, $d = 0.01$, $\mu = 0.0001$, $\sigma = 0.05$, $T = 100000$, and $\nu = 100$. (a) Evolution of s , x , and r . (b) Evolution of \bar{c} .

REFERENCE LIST

- [AB02] R. Albert and A.L. Barabási. “Statistical mechanics of complex networks.” *Reviews of modern physics*, **74**(1):47–97, 2002.
- [AG05] L.A. Adamic and N. Glance. “The political blogosphere and the 2004 US election: divided they blog.” In *Proceedings of the 3rd international workshop on Link discovery*, pp. 36–43. ACM, 2005.
- [ALM94] R. Antia, B.R. Levin, and R.M. May. “Within-host population dynamics and the evolution and maintenance of microparasite virulence.” *American Naturalist*, pp. 457–472, 1994.
- [Alt79] S.A. Altmann. “Altruistic behaviour: the fallacy of kin deployment.” *Animal Behaviour*, **27**:958–959, 1979.
- [AM79] R.M. Anderson and R.M. May. “Population biology of infectious diseases: Part I.” *Nature*, **280**(5721):361–367, 1979.
- [AM85] R.M. Anderson and R.M. May. “Vaccination and herd immunity to infectious diseases.” *Nature*, **318**(6044):323–329, 1985.
- [AM92] R.M. Anderson and R.M. May. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press, 1992.
- [ASB00] L.A.N. Amaral, A. Scala, M. Barthélémy, and H.E. Stanley. “Classes of small-world networks.” *Proceedings of the National Academy of Sciences*, **97**(21):11149–11152, 2000.
- [Axe06] R.M. Axelrod. *The Evolution of Cooperation*. Basic Books, 2006.
- [BA99] A.L. Barabási and R. Albert. “Emergence of scaling in random networks.” *Science*, **286**(5439):509–512, 1999.
- [Bai75] N.T.J. Bailey. *The Mathematical Theory of Infectious Diseases and its Applications*. Griffin, 1975.
- [BAJ99] A.L. Barabási, R. Albert, and H. Jeong. “Diameter of the world-wide web.” *Nature*, **401**(9):130–131, 1999.
- [BAJ00] A.L. Barabási, R. Albert, and H. Jeong. “Scale-free characteristics of random networks: the topology of the world-wide web.” *Physica A: Statistical Mechanics and its Applications*, **281**(1):69–77, 2000.

- [Ban05] S. Banerjee. *Dynamics for Engineers*. Wiley, 2005.
- [BBP04] M. Barthélemy, A. Barrat, R. Pastor-Satorras, and A. Vespignani. “Velocity and hierarchical spread of epidemic outbreaks in scale-free networks.” *Physical Review Letters*, **92**(17):178701, 2004.
- [BBP05] M. Barthelemy, A. Barrat, R. Pastor-Satorras, and A. Vespignani. “Dynamical patterns of epidemic outbreaks in complex heterogeneous networks.” *Journal of theoretical biology*, **235**(2):275–288, 2005.
- [BBV08] A. Barrat, M. Barthélemy, and A. Vespignani. *Dynamical Processes on Complex Networks*. Cambridge University Press, 2008.
- [BC87] R.L. Borelli and C.S. Coleman. *Differential Equations: A Modeling Approach*. Prentice-Hall, 1987.
- [BE12] D.M. Boyd and N.B. Ellison. “Social network sites.” *Online Communication and Collaboration: A Reader*, 2012.
- [Bec] B.J. Becker. “Infectious and Epidemic Disease in History.” <https://eee.uci.edu/clients/bjbecker/PlaguesandPeople/lecture3.html>.
- [Bet35] H.A. Bethe. “Statistical theory of superlattices.” *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, **150**(871):552–575, 1935.
- [Bin92] K.G. Binmore. *Fun and Games: A Text on Game Theory*. McGraw-Hill, 1992.
- [BKM00] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. “Graph structure in the web.” *Computer Networks*, **33**(1):309–320, 2000.
- [Bol79] B. Bollobás. *A Probabilistic Proof of an Asymptotic Formula for the Number of Labelled Regular Graphs*. Matematisk Institut, Aarhus Universitet, 1979.
- [Bon87] P. Bonacich. “Power and centrality: A family of measures.” *American Journal of Sociology*, **92**:1170–1182, 1987.
- [Bon02] E. Bonabeau. “Agent-based modeling: Methods and techniques for simulating human systems.” *Proceedings of the National Academy of Sciences*, **99**(3):7280–7287, 2002.
- [BP83] H.J. Bremermann and J. Pickering. “A game-theoretical model of parasite virulence.” *Journal of Theoretical Biology*, **100**(3):411–426, 1983.

- [BPV03] M. Boguná, R. Pastor-Satorras, and A. Vespignani. “Absence of epidemic threshold in scale-free networks with degree correlations.” *Physical Review Letters*, **90**(2):28701, 2003.
- [CBB09] R. Carvalho, L. Buzna, F. Bono, E. Gutiérrez, W. Just, and D. Arrowsmith. “Robustness of trans-European gas networks.” *Physical review E*, **80**(1):016106, 2009.
- [CCG02] Q. Chen, H. Chang, R. Govindan, and S. Jamin. “The origin of power laws in Internet topologies revisited.” In *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings*, volume 2, pp. 608–617. IEEE, 2002.
- [CCK06] P. Cano, O. Celma, M. Koppenberger, and M.B. Javier. “Topology of music recommendation networks.” *Chaos: An Interdisciplinary Journal of Nonlinear Science*, **16**(1):013107, 2006.
- [CD06] G. Clarkson and D. Dekorte. “The problem of patent thickets in convergent technologies.” *Annals of the New York Academy of Sciences*, **1093**(1):180–200, 2006.
- [CEB00] R. Cohen, K. Erez, D. Ben-Avraham, and S. Havlin. “Resilience of the Internet to random breakdowns.” *Physical Review Letters*, **85**(21):4626–4628, 2000.
- [CFR09] M. Crowe, M. Fitzgerald, D.L. Remington, G.D. Ruxton, and J. Rychtář. “Game theoretic model of brood parasitism in a dung beetle *Onthophagus taurus*.” *Evolutionary Ecology*, **23**(5):765–776, 2009.
- [CHB03] R. Cohen, S. Havlin, and D. Ben-Avraham. “Efficient immunization strategies for computer networks and populations.” *Physical review letters*, **91**(24):247901, 2003.
- [CHK01] D.S. Callaway, J.E. Hopcroft, J.M. Kleinberg, M.E.J. Newman, and S.H. Strogatz. “Are randomly grown graphs really random?” *Physical Review E*, **64**(4):041902, 2001.
- [CNS00] D.S. Callaway, MEJ Newman, S.H. Strogatz, and D.J. Watts. “Network robustness and fragility: Percolation on random graphs.” *Physical Review Letters*, **85**:5468–5471, 2000.
- [Cre03] R. Cressman. *Evolutionary dynamics and extensive form games*. MIT Press, 2003.

- [Cro10] C.D. Crowder. “List of Ten of the Worst Computer Viruses.”, 2010.
[http://voices.yahoo.com/
list-ten-worst-computer-viruses-5526277.html?cat=15.](http://voices.yahoo.com/list-ten-worst-computer-viruses-5526277.html?cat=15)
- [CSN09] A. Clauset, C.R. Shalizi, and M.E.J. Newman. “Power-law distributions in empirical data.” *SIAM review*, **51**(4):661–703, 2009.
- [CW] The CDC and the World Health Organization. “History and Epidemiology of Global Smallpox Eradication.” [http://www.bt.cdc.gov/agent/
smallpox/training/overview/pdf/eradicationhistory.pdf](http://www.bt.cdc.gov/agent/smallpox/training/overview/pdf/eradicationhistory.pdf). Slide 16-17.
- [Dar69] C. Darwin. *On the Origin of Species by Means of Natural Selection*. Dover, 1869.
- [DB02] Z. Dezsó and A.L. Barabási. “Halting viruses in scale-free networks.” *Physical Review E*, **65**(5):055103, 2002.
- [DC96] A.P. Dobson and E.R. Carper. “Infectious diseases and human population history.” *Bioscience*, **46**(2):115–126, 1996.
- [DH00] O. Diekmann and J.A.P. Heesterbeek. *Mathematical Epidemiology of Infectious Diseases: Model Building, Analysis and Interpretation*, volume 5. Wiley, 2000.
- [DHK04] M. Doebeli, C. Hauert, and T. Killingback. “The evolutionary origin of cooperators and defectors.” *Science*, **306**(5697):859, 2004.
- [DP] Centers for Disease Control and Prevention. “Tuberculosis.” <http://www.cdc.gov/tb/?404>; [http://www.cdc.gov:
80/TB/pubs/mdrtb/default.htm](http://www.cdc.gov:80/TB/pubs/mdrtb/default.htm).
- [Dug97] L.A. Dugatkin. *Cooperation Among Animals: An Evolutionary Perspective*. Oxford University Press, 1997.
- [EGK04] S. Eubank, H. Guclu, V.S.A. Kumar, M.V. Marathe, A. Srinivasan, Z. Toroczkai, and N. Wang. “Modelling disease outbreaks in realistic urban social networks.” *Nature*, **429**(6988):180–184, 2004.
- [EK10] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, 2010.
- [Ell01] S.P. Ellner. “Pair approximation for lattice models with multiple interaction scales.” *Journal of Theoretical Biology*, **210**(4):435–447, 2001.

- [EMB02] H. Ebel, L. Mielsch, and S. Bornholdt. “Scale-free topology of e-mail networks.” *Physical Review E*, **66**, 2002.
- [ER59] P. Erdős and A. Rényi. “On random graphs.” *Publicationes Mathematicae*, **6**:290–297, 1959.
- [ER60] P. Erdős and A. Rényi. “On the evolution of random graphs.” In *Publications of the mathematical Institute of the Hungarian Academy of Sciences*, volume 5, pp. 17–61, 1960.
- [ER61] P. Erdős and A. Rényi. “On the strength of connectedness of a random graph.” *Acta Mathematica Hungarica*, **12**(1):261–267, 1961.
- [Ewa93] P.W. Ewald. “The evolution of virulence.” *Scientific American*, **268**(4):86–93, 1993.
- [FFF99] M. Faloutsos, P. Faloutsos, and C. Faloutsos. “On power-law relationships of the Internet topology.” In *ACM SIGCOMM Computer Communication Review*, volume 29, pp. 251–262. ACM, 1999.
- [FG00] E. Fehr and S. Gächter. “Cooperation and punishment in public goods experiments.” *American Economic Review*, **90**(10):980–994, 2000.
- [FGD06] J.C. Flack, M. Girvan, F.B.M. De Waal, and D.C. Krakauer. “Policing stabilizes construction of social niches in primates.” *Nature*, **439**(7075):426–429, 2006.
- [FJ08] J.H. Fowler and S. Jeon. “The authority of Supreme Court precedent.” *Social Networks*, **30**(1):16–30, 2008.
- [FJS07] J.H. Fowler, T.R. Johnson, J.F. Spriggs II, S. Jeon, and P.J. Wahlbeck. “Network analysis and the law: Measuring the legal importance of Supreme Court precedents.” *Political Analysis*, **15**(3):324–46, 2007.
- [Fra98] S.A. Frank. *Foundations of Social Evolution*. Princeton University Press, 1998.
- [Fre77] L.C. Freeman. “A set of measures of centrality based on betweenness.” *Sociometry*, **40**:35–41, 1977.
- [FT91] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
- [GH01] J.K. Goeree and C.A. Holt. “Ten little treasures of game theory and ten intuitive contradictions.” *The American Economic Review*, **91**(5):1402–1422, 2001.

- [Gib92] R. Gibbons. *Game Theory for Applied Economists*. Princeton University Press, 1992.
- [Gin09] H. Gintis. *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction*. Princeton University Press, 2009.
- [GKM97] S.A.H. Geritz, E. Kisdi, G. Mesze´ N.A, and J.A.J. Metz. “Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree.” *Evolutionary ecology*, **12**(1):35–57, 1997.
- [Gra73] T.R. Grant. “Dominance and association among members of a captive and a free-ranging group of grey kangaroos (*Macropus giganteus*).” *Animal Behaviour*, **21**(3):449–456, 1973.
- [Gra90] A. Grafen. “Do animals really recognize kin?” *Animal Behaviour*, **39**(1):42–54, 1990.
- [Gro02] J.W. Grossman. “The evolution of the mathematical research collaboration graph.” *Congressus Numerantium*, pp. 201–212, 2002.
- [Gru08] J. Grujić. “Movies recommendation networks as bipartite graphs.” *Computational Science–ICCS 2008*, pp. 576–583, 2008.
- [Ham63] W.D. Hamilton. “The evolution of altruistic behavior.” *The American Naturalist*, **97**(896):354–356, 1963.
- [Har68] Garrett Hardin. “The tragedy of the commons.” *Science*, **162**:1243–1248, 1968.
- [HBR96] M. Huxham, S. Beaney, and D. Raffaelli. “Do parasites reduce the chances of triangulation in a real food web?” *Oikos*, pp. 284–300, 1996.
- [HD04] C. Hauert and M. Doebeli. “Spatial structure often inhibits the evolution of cooperation in the snowdrift game.” *Nature*, **428**(6983):643–646, 2004.
- [Het00] H.W. Hethcote. “The mathematics of infectious diseases.” *SIAM review*, **42**(4):599–653, 2000.
- [His] North Carolina Digital History. “Smallpox.” <http://www.learnnc.org/lp/editions/nchist-twoworlds/1696>.
- [HK02] P. Holme and B.J. Kim. “Growing scale-free networks with tunable clustering.” *Physical review E*, **65**(2):026107, 2002.

- [HS98] J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.
- [IKD04] M. Ifti, T. Killingback, and M. Doebeli. “Effects of neighbourhood size and connectivity on the spatial continuous Prisoner’s Dilemma.” *Journal of theoretical biology*, **231**(1):97–106, 2004.
- [JMB01] H. Jeong, S.P. Mason, A.L. Barabasi, and Z.N. Oltvai. “Lethality and centrality in protein networks.” *Nature*, **411**(6833):41–42, 2001.
- [JTA00] H. Jeong, B. Tombor, R. Albert, Z.N. Oltvai, and A.L. Barabási. “The large-scale organization of metabolic networks.” *Nature*, **407**(6804):651–654, 2000.
- [KD02] T. Killingback and M. Doebeli. “The continuous prisoner’s dilemma and the evolution of cooperation through reciprocal altruism with variable investment.” *American Naturalist*, **160**(4):421–438, 2002.
- [KDH10] T. Killingback, M. Doebeli, and C. Hauert. “Diversity of cooperation in the tragedy of the commons.” *Biological Theory*, **5**(1):3–6, 2010.
- [KDK99] T. Killingback, M. Doebeli, and N. Knowlton. “Variable investment, the continuous prisoner’s dilemma, and the origin of cooperation.” *Proceedings of the Royal Society of London. Series B: Biological Sciences*, **266**(1430):1723, 1999.
- [KKK02] D. Kempe, J. Kleinberg, and A. Kumar. “Connectivity and inference problems for temporal networks.” *Journal of Computer and System Sciences*, **64**(4):820–842, 2002.
- [Kle07] J. Kleinberg. “Computing: The wireless epidemic.” *Nature*, **449**(7160):287–288, 2007.
- [Kol98] P. Kollock. “Social dilemmas: The anatomy of cooperation.” *Annual review of sociology*, pp. 183–214, 1998.
- [KRF02] B. Kerr, M.A. Riley, M.W. Feldman, B.J.M. Bohannan, et al. “Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors.” *Nature*, **418**(6894):171–174, 2002.
- [KSC97] J.O. Kephart, G.B. Sorkin, D.M. Chess, and S.R. White. “Fighting computer viruses.” *Scientific American*, **277**(5):56–61, 1997.

- [LCS07] E.A. Leicht, G. Clarkson, K. Shedden, and M.E.J. Newman. “Large-scale structure of time evolving citation networks.” *The European Physical Journal B*, **59**(1):75–83, 2007.
- [LEA03] F. Liljeros, C.R. Edling, and L.A.N. Amaral. “Sexual networks: Implications for the transmission of sexually transmitted infections.” *Microbes and Infection*, **5**(2):189–196, 2003.
- [LHN95] M. Lipsitch, E.A. Herre, and M.A. Nowak. “Host population structure and the evolution of virulence: a law of diminishing returns.” *Evolution*, pp. 743–748, 1995.
- [LHN05] E. Lieberman, C. Hauert, and M.A. Nowak. “Evolutionary dynamics on graphs.” *Nature*, **433**(7023):312–316, 2005.
- [Lib] The Navy Department Library. “Influenza of 1918 (Spanish Flu) and the US Navy.” http://www.history.navy.mil/library/online/influenza_main.htm.
- [LKF07] J. Leskovec, J. Kleinberg, and C. Faloutsos. “Graph evolution: Densification and shrinking diameters.” *ACM Transactions on Knowledge Discovery from Data (TKDD)*, **1**(1):2, 2007.
- [LM94] R.E. Lenski and R.M. May. “The evolution of virulence in parasites and pathogens: reconciliation between two competing hypotheses.” *Journal of Theoretical Biology*, **169**(3):253–265, 1994.
- [LM01] A.L. Lloyd and R.M. May. “How viruses spread among computers and people.” *Science*, **292**(5520):1316–1317, 2001.
- [LS88] C.C. Lin and L.A. Segel. *Mathematics Applied to Deterministic Problems in the Natural Sciences*, volume 1. Society for Industrial Mathematics, 1988.
- [LS08] K. Leyton-Brown and Y. Shoham. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2008.
- [Lus03] D. Lusseau. “The emergent properties of a dolphin social network.” *Proceedings of the Royal Society of London. Series B: Biological Sciences*, **270**(2):186–188, 2003.
- [MA84] R.M. May and R.M. Anderson. “Spatial heterogeneity and the design of immunization programs.” *Mathematical Biosciences*, **72**(1):83–111, 1984.

- [MA87] R.M. May and R.M. Anderson. “Transmission dynamics of HIV infection.” *Nature*, **326**(6109):137–142, 1987.
- [Mar91] N.D. Martinez. “Artifacts or attributes? Effects of resolution on the Little Rock Lake food web.” *Ecological Monographs*, pp. 367–392, 1991.
- [MB07] B.J. McGill and J.S. Brown. “Evolutionary game theory and adaptive dynamics of continuous traits.” *The Annual Review of Ecology, Evolution, and Systematics*, **38**:403–435, 2007.
- [MH97] M.S. Mooring and B.L. Hart. “Reciprocal allogrooming in wild impala lambs.” *Ethology*, **103**(8):665–680, 1997.
- [Mil67] S. Milgram. “The small world problem.” *Psychology Today*, **2**(1):60–67, 1967.
- [Mil87] M. Milinski. “Tit for tat in sticklebacks and the evolution of cooperation.” *Nature*, **325**(6103):433–435, 1987.
- [MKC04] N. Madar, T. Kalisky, R. Cohen, D. Ben-Avraham, and S. Havlin. “Immunization and epidemic dynamics in complex networks.” *The European Physical Journal B*, **38**(2):269–276, 2004.
- [ML01] R.M. May and A.L. Lloyd. “Infection dynamics on scale-free networks.” *Physical Review E*, **64**(6):066112, 2001.
- [Moo02] J. Moody. “The importance of relationship timing for diffusion.” *Social Forces*, **81**(1):25–56, 2002.
- [MOS92] H. Matsuda, N. Ogita, A. Sasaki, and K. Sato. “Statistical mechanics of population.” *Progress of Theoretical Physics*, **88**(6):1035–1049, 1992.
- [MRL04] C.E. Mills, J.M. Robins, and M. Lipsitch. “Transmissibility of 1918 pandemic influenza.” *Nature*, **432**(7019):904–906, 2004.
- [MSK06] M. Milinski, D. Semmann, H.J. Krambeck, and J. Marotzke. “Stabilizing the Earth’s climate is not a losing game: Supporting evidence from public goods experiments.” *Proceedings of the National Academy of Sciences*, **103**(11):3994–3998, 2006.
- [Mye97] R.B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997.
- [Nas50] J.F. Nash. “Equilibrium points in n-person games.” *Proceedings of the National Academy of Sciences*, **36**(1):48–49, 1950.

- [New01] M.E.J. Newman. “The structure of scientific collaboration networks.” *Proceedings of the National Academy of Sciences*, **98**(2):404–409, 2001.
- [New02] M.E.J. Newman. “Assortative mixing in networks.” *Physical Review Letters*, **89**(20):208701, 2002.
- [New03] M.E.J. Newman. “The structure and function of complex networks.” *SIAM review*, **45**(2):167–256, 2003.
- [New05] M.E.J. Newman. “Power laws, Pareto distributions and Zipf’s law.” *Contemporary physics*, **46**(5):323–351, 2005.
- [New10] M.E.J. Newman. *Networks: An Introduction*. Oxford University Press, 2010.
- [Nis07] N. Nisan. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [NMK07] J. von Neumann, O. Morgenstern, H.W. Kuhn, and A. Rubinstein. *Theory of Games and Economic Behavior (Commemorative Edition)*. Princeton Classic Editions. Princeton University Press, 2007.
- [Nou07] A. Nouweland. “Rock-Paper-Scissors: A New and Elegant Proof.” Department of Economics - Working Papers Series 1003, The University of Melbourne, 2007.
- [Now06a] M.A. Nowak. *Evolutionary Dynamics: Exploring the Equations of Life*. Belknap Press, 2006.
- [Now06b] M.A. Nowak. “Five rules for the evolution of cooperation.” *Science*, **314**(5805):1560–1563, 2006.
- [NS92] M.A. Nowak and K. Sigmund. “Tit for tat in heterogeneous populations.” *Nature*, **355**(6357):250–253, 1992.
- [NS93] M. Nowak and K. Sigmund. “A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner’s Dilemma game.” *Nature*, **364**(6432):56–58, 1993.
- [NS98] M.A. Nowak and K. Sigmund. “Evolution of indirect reciprocity by image scoring.” *Nature*, **393**(6685):573–577, 1998.
- [NS05] M.A. Nowak and K. Sigmund. “Evolution of indirect reciprocity.” *Nature*, **437**(7063):1291–1298, 2005.
- [NS07] M.A. Nowak and K. Sigmund. “How populations cohere: five rules for cooperation.” *Theoretical Ecology*, pp. 7–16, 2007.

- [NST04] M.A. Nowak, A. Sasaki, C. Taylor, and D. Fudenberg. “Emergence of cooperation and evolutionary stability in finite populations.” *Nature*, **428**(6983):646–650, 2004.
- [OHL06] H. Ohtsuki, C. Hauert, E. Lieberman, and M.A. Nowak. “A simple rule for the evolution of cooperation on graphs and social networks.” *Nature*, **441**(7092):502–505, 2006.
- [ON06a] H. Ohtsuki and M.A. Nowak. “Evolutionary games on cycles.” *Proceedings of the Royal Society B: Biological Sciences*, **273**(1598):2249–2256, 2006.
- [ON06b] H. Ohtsuki and M.A. Nowak. “The replicator equation on graphs.” *Journal of Theoretical Biology*, **243**(1):86, 2006.
- [Org] The World Health Organization. “International Classification of Diseases (ICD).” <http://www.who.int/classifications/icd/en/>.
- [Org95] The World Health Organization. “The World Health Report.”, 1995. http://www.who.int/whr/1995/en/whr95_ch1_en.pdf.
- [Org04] The World Health Organization. “The World Health Report.”, 2004. http://www.who.int/whr/2004/annex/topic/en/annex_2_en.pdf.
- [OSH07] J.P. Onnela, J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A.L. Barabási. “Structure and tie strengths in mobile communication networks.” *Proceedings of the National Academy of Sciences*, **104**(18):7332–7336, 2007.
- [Ost90] E. Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990.
- [PA93] J.F. Padgett and C.K. Ansell. “Robust action and the rise of the Medici, 1400-1434.” *American Journal of Sociology*, pp. 1259–1319, 1993.
- [Pri65] D.S. Price. “Networks of scientific papers.” *Science*, **149**(3683):510–515, 1965.
- [Pri76] D.S. Price. “A general theory of bibliometric and other cumulative advantage processes.” *Journal of the American Society for Information Science*, **27**(5):292–306, 1976.
- [PV01a] R. Pastor-Satorras and A. Vespignani. “Epidemic dynamics and endemic states in complex networks.” *Physical Review E*, **63**(6):066117, 2001.

- [PV01b] R. Pastor-Satorras and A. Vespignani. “Epidemic spreading in scale-free networks.” *Physical review letters*, **86**(14):3200–3203, 2001.
- [PV01c] R. Pastor-Satorras and A. Vespignani. “Epidemic spreading in scale-free networks.” *Physical review letters*, **86**(14):3200–3203, 2001.
- [PV02a] R. Pastor-Satorras and A. Vespignani. “Epidemic dynamics in finite size scale-free networks.” *Physical Review E*, **65**(3):035108, 2002.
- [PV02b] R. Pastor-Satorras and A. Vespignani. “Immunization of complex networks.” *Physical Review E*, **65**(3):036104, 2002.
- [RF02] M. Ripeanu and I. Foster. “Mapping the Gnutella network: Macroscopic properties of large-scale peer-to-peer systems.” *Peer-to-Peer Systems*, pp. 85–93, 2002.
- [Rus59] B. Russell. *Common Sense and Nuclear Warfare*. Allen & Unwin, 1959.
- [SA94] D. Stauffer and A. Aharony. *Introduction To Percolation Theory*. Taylor & Francis, 1994.
- [Sad72] D.S. Sade. “Sociometrics of *Macaca mulatta*. I. Linkages and cliques in grooming matrices.” *Folia Primatologica*, **18**(3-4):196–223, 1972.
- [SCM80] R.M. Seyfarth, D.L. Cheney, and P. Marler. “Vervet monkey alarm calls: semantic communication in a free-ranging primate.” *Animal Behaviour*, **28**(4):1070–1094, 1980.
- [SDC03] P. Sen, S. Dasgupta, A. Chatterjee, P.A. Sreeram, G. Mukherjee, and S.S. Manna. “Small-world properties of the Indian railway network.” *Physical Review E*, **67**(3):036106, 2003.
- [SF07] G. Szabo and G. Fáth. “Evolutionary games on graphs.” *Physics Reports*, **446**(4-6):97–216, 2007.
- [SG84] L.D. Sailer and S.J.C. Gaulin. “Proximity, sociality, and observation: the definition of social groups.” *American Anthropologist*, **86**(1):91–98, 1984.
- [Sim55] H.A. Simon. “On a class of skew distribution functions.” *Biometrika*, **42**(3/4):425–440, 1955.
- [Sky03] B. Skyrms. *The Stag Hunt and the Evolution of Social Structure*. Cambridge University Press, 2003.

- [SL96] B. Sinervo and C.M. Lively. “The rock-paper-scissors game and the evolution of alternative male strategies.” *Nature*, **380**(6571):240–243, 1996.
- [SM76] R. Sibly and D. McFarland. “On the fitness of behavior sequences.” *The American Naturalist*, **110**(974):601–617, 1976.
- [Smi74] J.M. Smith. “The theory of games and the evolution of animal conflicts.” *Journal of Theoretical Biology*, **47**(1):209–221, 1974.
- [Smi82] J.M. Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.
- [SN99] K. Sigmund and M.A. Nowak. “Evolutionary game theory.” *Current Biology*, **9**(14):503–505, 1999.
- [SP73] J.M. Smith and G.R. Price. “The logic of animal conflict.” *Nature*, **246**:15, 1973.
- [SPN04] C.P. van Schaik, G.R. Pradhan, and M.A. van Noordwijk. “Mating conflict in primates: infanticide, sexual harassment and female sexuality.” *Sexual Selection in Primates: New and Comparative Perspectives*, pp. 131–150, 2004.
- [SR83] S.R. Schulman and D.I. Rubenstein. “Kinship, need, and the distribution of altruism.” *American Naturalist*, pp. 776–788, 1983.
- [SR04] D. Stutzbach and R. Rejaie. “Characterizing today’s Gnutella topology.” Technical report, Department of Computer Science, University of Oregon, 2004.
- [SS97] J.M. Smith and E. Szathmary. *The Major Transitions in Evolution*. Oxford University Press, 1997.
- [Str94] S.H. Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Westview Press, 1994.
- [Str01] S.H. Strogatz. “Exploring complex networks.” *Nature*, **410**(6825):268–276, 2001.
- [TC99] P.E. Turner, L. Chao, et al. “Prisoner’s dilemma in an RNA virus.” *Nature*, **398**(6726):441, 1999.
- [TJ78] P.D. Taylor and L.B. Jonker. “Evolutionary stable strategies and game dynamics.” *Mathematical Biosciences*, **40**(1):145–156, 1978.

- [TM69] J. Travers and S. Milgram. “An experimental study of the small world problem.” *Sociometry*, pp. 425–443, 1969.
- [TN06] A. Traulsen and M.A. Nowak. “Evolution of cooperation by multilevel selection.” *Proceedings of the National Academy of Sciences*, **103**(29):10952–10955, 2006.
- [Tri71] R.L. Trivers. “The evolution of reciprocal altruism.” *Quarterly Review of Biology*, pp. 35–57, 1971.
- [TWH05] J.R. Tyler, D.M. Wilkinson, and B.A. Huberman. “E-mail as spectroscopy: Automated discovery of community structure within organizations.” *The Information Society*, **21**(2):143–153, 2005.
- [VMC09] B. Viswanath, A. Mislove, M. Cha, and K.P. Gummadi. “On the evolution of user interaction in Facebook.” In *Proceedings of the 2nd ACM SIGCOMM Workshop on Social Networks (WOSN’09)*. ACM, 2009.
- [Web07] J.N. Webb. *Game Theory: Decisions, Interaction and Evolution*. Springer undergraduate mathematics series. Springer, 2007.
- [Wei81] R.M. Weigel. “The distribution of altruism among kin: a mathematical model.” *American Naturalist*, pp. 191–201, 1981.
- [WF94] S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, 1994.
- [Wil84] G.S. Wilkinson. “Reciprocal food sharing in the vampire bat.” *Nature*, **308**(5955):181–184, 1984.
- [Wil90] G.S. Wilkinson. “Food sharing in vampire bats.” *Scientific American*, **262**(2):76–82, 1990.
- [WM92] C.H. Watts and R.M. May. “The influence of concurrent partnerships on the dynamics of HIV/AIDS.” *Mathematical Biosciences*, **108**(1):89–104, 1992.
- [Wor99] N.C. Wormald. “Models of random regular graphs.” *London Mathematical Society Lecture Note Series*, pp. 239–298, 1999.
- [WS98] D.J. Watts and S.H. Strogatz. “Collective dynamics of ‘small-world’ networks.” *Nature*, **393**(6684):440–442, 1998.

- [WST86] J.G. White, E. Southgate, J.N. Thomson, S. Brenner, J.G. White, E. Southgate, J.N. Thomson, and S. Brenner. “The structure of the nervous system of the nematode *Caenorhabditis elegans*.” *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, **314**(1165):1–340, 1986.
- [WT04] J. Wallinga and P. Teunis. “Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures.” *American Journal of Epidemiology*, **160**(6):509–516, 2004.
- [WY22] J.C. Willis and G.U. Yule. “Some statistics of evolution and geographical distribution in plants and animals, and their significance.” *Nature*, **109**(2728):177–179, 1922.
- [XS04] R. Xulvi-Brunet and I.M. Sokolov. “Reshuffling scale-free networks: From random to assortative.” *Physical Review E*, **70**(6):066102, 2004.
- [Zac77] W.W. Zachary. “An information flow model for conflict and fission in small groups.” *Journal of Anthropological Research*, pp. 452–473, 1977.
- [Zee80] E. Zeeman. “Population dynamics from game theory.” *Global theory of dynamical systems*, pp. 471–497, 1980.
- [Zip49] G.K. Zipf. “Human behavior and the principle of least effort.” *Addison-Wesley Press*, 1949.