Local replicator dynamics and the evolution of bacterial populations

Seeking an application of models of evolutionary game theory in the biomedical domain

Cecile Wolfs - i6010084
University College Maastricht
2011/2012
Capstone

Dr. Frank Thuijsman
Dr. Petra Wolffs
Philippe Uyttendaele, MSc.

June 20, 2012

Abstract

Evolutionary game theory is an interdisciplinary field that combines game theory and evolutionary biology. Since it's emergence in the 1970s it has been extensively studied and it developed further more recently. The local replicator dynamics are the result of this development. Being relatively new, there is a search for novel applications of the local replicator dynamics. One of the areas of interest is the field of microbiology. This thesis provides an introduction to the models of evolutionary game theory, as well as an overview of literature from the biomedical domain of possible interest for application of the models. Furthermore, two articles are explored more in depth. They serve as a case study for the application of the local replicator dynamics. Finally, conclusions concerning the applicability of these models to the biomedical domain are derived and discussed. It appears that it is possible to apply the local replicator dynamics to a phenomenon found in the literature, but future research is needed to see if the applicability can be broadened.

Table of contents

Introduction

Since the 1970s evolutionary game theory is extensively studied. Combining game theory and evolutionary biology, evolutionary game theory is an interdisciplinary field that is applicable to many situations. The development of models that take the location of individuals in a population into account, the local replicator dynamics, leads to the possibility of modelling evolutionary processes even more realistically. However, since these local models are relatively new, there is still a search for more phenomena that they can be applied to. One of the fields that is currently being considered is the field of microbiology. The research question of this thesis is therefore: can local models of evolutionary game theory be applied to biomedical phenomena?

From communication with a microbiologist, it became clear that certain phenomena are observed, but there is not enough knowledge about underlying mathematical models. The mathematicians, on the other hand, have these relatively new local models of evolutionary game theory and are looking for more applications (P. Wolffs & F. Thuijsman, personal communication, March 21, 2012). Knowing the models that underlie certain phenomena will help understanding them better, since mathematical models can help to describe, explain and make predictions concerning the situation at hand. Moreover, if one knows the dynamics driving a particular process, it is often possible to control it by manipulating certain characteristics of the process. Therefore, it would be favourable to have mathematical models explaining particular observations and phenomena in the biomedical domain. This thesis is meant to provide such a link between evolutionary game theory and the biomedical domain. The aim is to find a microbiological phenomenon in the literature to which the local replicator dynamics can be applied, and which thus establishes a bridge between the fields of mathematics and microbiology.

To go into slightly more detail, this thesis focuses on bacteria and not on cells and viruses. Personal communication (P. Wolffs, April 2, 2012) led to this decision. Since bacteria multiply and reproduce quickly, it would be easy to do experiments and monitor their evolution in a laboratory and see how they relate to the mathematical models that are available. At first viruses where also

considered, which is why there is also some literature studied that focuses on those pathogens, but it turned out that the evolution and virulence of viruses were too complicated and time-consuming to monitor for this thesis.

In order to answer the research question and reach the goal of finding a link between evolutionary game theory and the biomedical domain, this thesis is set up in the following way. The first chapter elaborates on the mathematical framework on which this thesis is built. Classical game theory as well as evolutionary game theory and the local replicator dynamics are discussed and illustrated by examples. After that, the second chapter provides a short introduction into relevant concepts and observations in the biomedical domain. The third chapter is the literature study, which gives an overview of the literature that was possibly of interest for the research question and that was provided by the microbiologist. In this chapter, the last two articles are discussed in detail, because they were of particular interest for the aim of this thesis. The fourth and last chapter takes the form of a case study, in which is explained why the articles of interest are relevant and how the local replicator dynamics can be adjusted in order to be applicable to the phenomena specified in these articles. Finally, in the discussion the results of the literature study and the case study are discussed and a conclusion is formed. It can be derived that there are opportunities for applying the local replicator dynamics to the biomedical domain. However, this thesis only shows one such application. Future research is necessary to broaden the application of evolutionary game theory in microbiology.

Chapter 1: Game Theoretical Background

In this chapter the models used as a basis for this thesis are explained in detail. The first section concerns basic game theoretical concepts, such as games, strategies and the Nash equilibrium, which are illustrated with an example in the form of the game of Chicken. The second section takes this basic model a step further, and explains the models of evolutionary game theory. Again, the concepts explained here, such as Evolutionarily Stable Strategy and fitness matrix, are illustrated by elaborating upon an example. In this case, the Hawk-

Dove game is evaluated. The third and last section describes the latest development in evolutionary game theory: the local replicator dynamics. The basics of these local models are explained in order to be able to build on them in the rest of this thesis.

1.1 Classical game theory

1.1.1 Games, strategies and the Nash equilibrium

In the field of game theory, mathematical models are used to provide insight into the decisions actors make in a competitive environment. These mathematical models take the form of games. A game has several important features, which essentially make it a mathematical model of a competitive situation. First of all, there is a set of players. The players can be any kind of individual or group of individuals. Secondly, each player has a set of actions he can take at his disposal, and needs to make a decision about which of these actions would be most beneficial to perform. This leads to the last characteristic of a game: taking a certain action will lead the player to get a particular payoff. This payoff can be either positive or negative; it just means that the player gets "something" by choosing a certain action. It is usually assumed that all players act rational and try to maximize their own payoff (Osborne, 2004).

Such games can be displayed as matrices, in which each entry represents the payoff for the players. The actions player 1 can take correspond to the rows of the matrix and are indexed by i, while the actions player 2 can take correspond to the columns of the matrix and are indexed by j. A bimatrix game is a game for which the payoff matrix contains the payoffs for two players. In general, a bimatrix can be expressed as (A, B), with A representing a payoff matrix for player 1 and B a payoff matrix for player 2. Entry (i, j) of the matrix then represents a payoff (a_{ij}, b_{ij}) where a_{ij} constitutes the payoff for row player 1 and b_{ij} the payoff for column player 2 (Osborne, 2004).

Furthermore, each player can devise a strategy, representing the actions he or she is planning to take when playing the game. It can be seen as a plan of action that specifies how the player plans to respond at each possible stage of the game. A strategy can be either pure or mixed. This means that on the one hand

the player has the option to choose an action with probability 1 in case of the pure strategy, which basically means that he just chooses one action to play. On the other hand, in case of the mixed strategy, the player can specify probabilities with which he is choosing each action, meaning that he does not pick one action for sure, but there is a chance that he might pick an action. To represent this mathematically, consider a game with two players, in which player 1 has m actions and player 2 has n actions. The payoff matrices A and B then both have size $m \times n$. The mixed strategy for player 1 is a probability distribution represented by the vector $p = (p_1, p_2, \dots, p_m)$ with the probability of choosing action i being $p_i \geq 0$, and $\sum_{i=1}^{m} p_i = 1$. The mixed strategy for player 2 can be represented by a similar vector $q = (q_1, q_2, \dots, q_n)$. The expected payoff for both players can be calculated by multiplying the probability vector of a player by the payoff matrix of that player and the probability vector of the other player. This results in an expected payoff of $pAq = \sum_i \sum_j p_i a_{ij} q_j$ for player 1 and an expected payoff of $pBq = \sum_i \sum_j p_i b_{ij} q_j$ for player 2 (Osborne, 2004).

Knowing this, it is possible to establish what a Nash equilibrium is. As Nash himself wrote, a Nash equilibrium is a set of actions, "such that each player's mixed strategy maximizes his payoff if the strategies of the others are held fixed. Thus, each player's strategy is optimal against those of the others." (p. 287, 1951). Basically, when two actions form a Nash equilibrium, it means that there is no incentive for either player to deviate from playing this strategy, because he or she cannot improve his or her payoff by using a different strategy. In terms of the notation used before, this means that a pair of strategies (p^*, q^*) constitutes a Nash equilibrium if $p^*Aq^* \geq pAq^*$ for every mixed strategy p and $p^*Bq^* \geq p^*Bq$ for every mixed strategy q. This does not only hold for bimatrix games, but for any n-player game with a finite set of actions (Nash, 1951).

1.1.2 The game of Chicken

The story behind the game of Chicken is as follows: at the ends of a street, two cars are waiting, directly facing each other. When a signal is given, the two cars start racing towards one another. While racing, each car has two options: swerve

- and thus chicken out - or not swerve - and thus taking the risk of crashing into the other car. A payoff matrix of this game looks as follows:

$$\begin{array}{rcc} & \text{Swerve} & \text{Not swerve} \\ \text{Swerve} & 3,3 & 1,4 \\ \text{Not swerve} & 4,1 & 0,0 \end{array}$$

There are two Nash equilibria to be found when pure strategies are used. That is, when each of the players chooses an action with probability 1, there are two situations in which they both play a best response against the action of the other. These Nash equilibria are found as follows. Assume that the column player chooses to swerve; he plays the strategy (1,0). The row player's best response would then be to not swerve, which means playing (0,1), since that will yield him a payoff of 4 instead of 3. Furthermore, when the column player plays (0,1) and thus chooses not to swerve, the row player will get the highest payoff - 1 instead of 0 - when he chooses to swerve and thus play (1,0). In a similar way the best responses for the column player can be computed, and it turns out that these are the same as for the row player: swerve when the other does not, and not swerve when the other does. Thus, the strategy pairs (p^*, q^*) that constitute a Nash equilibrium in this game are $((0,1),(1,0))$ and $((1,0),(0,1))$ (Rapoport & Chammah, 1966).

However, there is also a Nash equilibrium in mixed strategies. This equilibrium is $\left(\left(\frac{1}{2},\frac{1}{2}\right),\left(\frac{1}{2},\frac{1}{2}\right)\right)$, meaning that each player plays each action with probability $\frac{1}{2}$. That this mixed strategy is a Nash equilibrium can be seen by calculating the payoffs for the two players. Starting with player 1, the row player, its expected payoff for choosing to swerve is $\frac{1}{2}\times 3\times\frac{1}{2}+\frac{1}{2}\times 1\times\frac{1}{2}=1$, and its expected payoff for choosing not to swerve is $\frac{1}{2}\times 4\times\frac{1}{2}+\frac{1}{2}\times 0\times\frac{1}{2}=1$. Since the game is symmetrical, these expected payoffs are exactly the same for player 2, the column player. Looking at the expected payoffs, it does not matter if the players choose to swerve or not to swerve, both actions lead to an expected payoff of 1. Moreover, since each action is equally good, playing $\left(\frac{1}{2},\frac{1}{2}\right)$ means playing a best response against the other player. Therefore, $\left(\left(\frac{1}{2},\frac{1}{2}\right),\left(\frac{1}{2},\frac{1}{2}\right)\right)$ is a Nash equilibrium as well, because both players are playing a best response against the other player.

1.2 Evolutionary game theory

1.2.1 The basics

The basic idea of evolutionary game theory is that it is an application of the previous explained game theoretical models to evolutionary biology. It studies the development of populations, which consist of a finite amount of different types of individuals that are interacting with each other. It is assumed that the different types meet randomly and frequently, so that the population is essentially playing against itself. In the context of evolutionary game theory, the payoff matrix is called the fitness matrix, and each entry thus represents the fitness for a certain type instead of a payoff. An easy way of thinking about these fitness entries is to consider them as being the amount of offspring that will be of the same type as the parent type. Every generation, each type acts exactly once as the row player, establishing a random interaction with another type, which will lead to offspring that is of the same type as the type acting as row player. The offspring then constitutes the new population. It is important to note that an individual does not decide if it is the row player or the column player; when two individuals meet one is randomly assigned to be the row player and the other to be the column player (Maynard Smith, 1982).

1.2.2 Evolutionarily stable strategy and replicator dynamics

As mentioned before, evolutionary game theory deals with the development of populations. When a game is played repeatedly in time, the population distribution changes going from one generation to the next. If a type is doing better than average, its proportion of the total population will increase, while the proportion of a type that is doing worse than average will decrease. The stabilization of the population distribution in time, meaning that each type surviving is playing a best response against that population distribution, can be considered to be the equivalent of a Nash equilibrium in the context of evolutionary game theory: an Evolutionarily Stable Strategy (ESS) (Maynard Smith, 1982).

However, an Evolutionarily Stable Strategy is not just a Nash equilibrium; an extra condition needs to be satisfied. If we consider the fitness matrix A for a population of n types and the population distribution $x = (x_1, x_2, \dots, x_n)$ over the

types $1, 2, \dots, n$, then the following two conditions need to be in place in order for x to be an ESS.

- $xAx \geq yAx$ for all y
- For each $y \neq x$: if $yAx = xAx$, then $xAy > yAy$

The first condition basically states that x should be a best response to x, which means that (x, x) is a Nash equilibrium for the bimatrix game (A, A^T). Here, A^T is the transpose of A, such that the game (A, A^T) represents the payoffs for the different types of the populations playing against each other. The essence of the second condition is that the population distribution should move back towards x, would it by chance at some point in time be a different population distribution y. The idea is that when a type has the same fitness as an average member of the population, its fraction of the total population could slowly increase. However, the second condition makes sure that as soon as individuals of this type meet, their fitness is below average, which makes their fraction of the population decrease again, such that the population distribution moves back to x (Maynard Smith & Price, 1973).

An ESS can be also be defined by using the concept of invasion. The idea is that if a population distribution x cannot by invaded by any mutant $y \neq x$, then this population distribution is an Evolutionarily Stable Strategy. Replacing a fraction $\varepsilon > 0$ of a population x by mutants would lead to a new population, characterized by

$$x_{\varepsilon y} = (1 - \varepsilon)x + \varepsilon y,$$

where y is the distribution of the mutant types and is not equal to x. It turns out that if the mutants have higher fitness than the original members of the population, then they can invade the population. In mathematical terms, invasion can occur if for small $\varepsilon > 0$, it holds that $yAx_{\varepsilon y} \geq xAx_{\varepsilon y}$ (Maynard Smith & Price, 1973).

Another concept closely related to the ESS and the fitness matrix are the replicator dynamics. The replicator dynamics are a system of differential equations that look as follows:

$$\dot{x}_i = x_i(e_iAx - xAx) \text{ for } i = 1, 2, \dots, n.$$

What they describe is how the population changes over time according to the fitness matrix, where the rate of growth of the fraction of type i is proportional

to how many individuals of type i are in the population, denoted by x_i, and to the difference in the average fitness of type i compared to the fitness of an average member of the population. The term e_iAx denotes the average fitness for an individual of type i, whereas the term xAx denotes the average fitness for any population member. If e_iAx is larger than xAx, then the term $e_iAx - xAx$ will be positive, so $\dot{x}_i$ will also be positive and the fraction of type i increases. Likewise, if e_iAx is smaller than xAx, then the term $e_iAx - xAx$ will be negative, so $\dot{x}_i$ will be negative as well and the fraction of type i decreases. A fixed or dynamically stable point of these dynamics is found when all the derivatives of $\dot{x}_i$ are 0, meaning that the population distribution is not changing anymore. An ESS is a specific type of such a dynamically stable point, namely an asymptotically stable one. A point x being asymptotically stable means that if x_0 is the starting point somewhere in an open neighbourhood of x, then the dynamical process converges to x (Taylor & Jonker, 1978).

1.2.3 The Hawk-Dove game

The Hawk-Dove game is essentially the same as the game of Chicken. In this interpretation, Hawk is seen as the aggressive type (i.e. the one who does not swerve) and Dove is the peaceful type (i.e. the one who does swerve). The population considered consists of individuals that can be either type Hawk or type Dove. The way individuals behave directly influences their fitness. Since the population is playing against itself, only one player's fitness matrix needs to be taken into consideration. This fitness matrix is the matrix A, displayed below, with the actions D and H, standing for Dove and Hawk respectively.

$$\begin{array}{c} \\ D \\ H \end{array} \begin{array}{c} \begin{array}{cc} D & H \end{array} \\ \begin{pmatrix} 3 & 1 \\ 4 & 0 \end{pmatrix} \end{array}$$

As an example, it will be shown that the mixed strategy $x = (\frac{1}{2}, \frac{1}{2})$, a Nash equilibrium in the game of Chicken, constitutes an ESS. In order to do this, it is necessary to make sure that any population distribution entering the game will perform worse than the one that is already there. In this case, an invader consisting only of individuals of type Dove, $y = (1, 0)$, and an invader consisting only of Hawks, $z = (0, 1)$, are taken into consideration. This allows for seeing

whether there are any distributions doing better than a population in which half of the individuals are Doves and the other half are Hawks, $x = (\frac{1}{2}, \frac{1}{2})$. Calculating the expected payoffs of these three populations playing against x as explained before gives the following outcomes:

$$xAx = 2 \qquad yAx = 2 \qquad zAx = 2$$

Thus, $xAx = yAx = zAx$. The first condition of the ESS theory states that xAx should be greater than or equal to $\tilde{x}Ax$ for all invaders $\tilde{x}$, which means that in this game xAx should be greater than or equal to both yAx and zAx. Since this is the case, the first condition is satisfied for x. However, because the expected payoffs are all equal, the second condition needs to be assessed as well in order to validate that x is indeed an ESS. This second condition states that $xA\tilde{x}$ should be greater than $\tilde{x}A\tilde{x}$ if it is the case that $\tilde{x}Ax$ is equal to xAx. For this game, this means that xAy and xAz should be strictly greater than yAy and zAz respectively. The expected payoffs for these strategies are as follows.

$$xAy = 3\frac{1}{2} \qquad yAy = 3 \qquad xAz = \frac{1}{2} \qquad zAz = 0$$

As can be seen, $xAy > yAy$ and $xAz > zAz$, meaning that the second condition holds for x as well. Therefore, the conclusion that the mixed strategy $x = (\frac{1}{2}, \frac{1}{2})$ is an Evolutionarily Stable Strategy can be drawn.[1]

1.3 Local models

The local models of evolutionary game theory were developed in two theses, the first one by Chantal Bendermacher (2009) and the second one by Mandy Tak (2012), in order to overcome a limitation of the global models, namely the assumption that the different types meet randomly and that their location does not influence what other types they meet. However, in real life one's own location does have an influence on which other individuals one meets and interacts with. For this reason, the local replicator dynamics are devised in such a way that they take the location of the types into account, so that the portion of the population that an individual can interact with is also influenced.

[1] The two Nash equilibria in pure strategies from the game of Chicken cannot be played here, since they do not represent a population playing against itself.

How does this work in practice? In order to be able to model the location of the individuals within a population, a playing field is used. As can be seen in figure 1.1, this field consists of cells that have a hexagonal shape. Important to note is that the playing field is not flat as displayed in the picture, but the left and right side are connected, as well as the top and bottom. This way, it is ensured that every cell has the same number of cells surrounding it: 6 in this case.

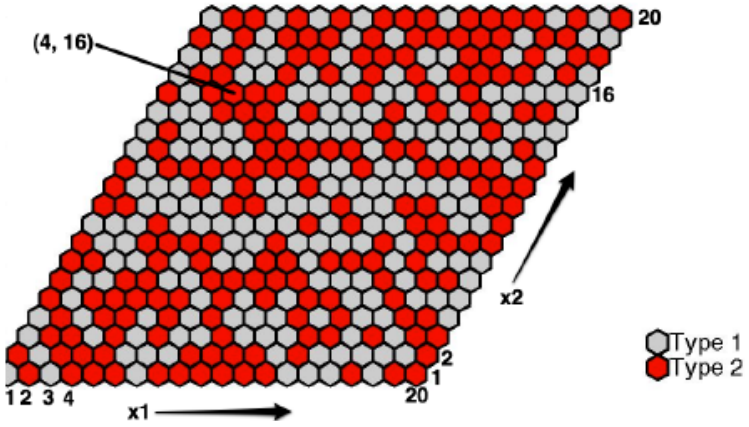

Figure 1.1 The playing field (Tak, 2012)

Every cell can only be occupied by one type, but it can also be empty. Locality comes into play by the fact that every cell can only interact with its neighbourhood. The neighbourhood consists of the cells surrounding a particular central cell *c*. The size of the neighbourhood can be denoted by all cells that are at most *M* cells away from *c*. *M* represents the maximum distance, and with this concept of distance, rings can be defined as well. As shown in figure 1.2, ring 0 is considered to be *c* itself, the cells directly neighbouring *c* constitute ring 1, the cells adjacent to the cells of ring 1 form ring 2, and so on (Bendermacher, 2009; Tak, 2012).

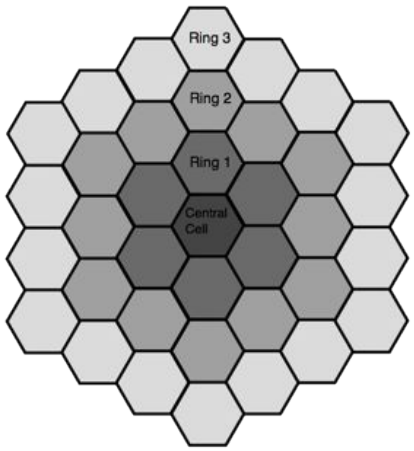

Figure 1.2 A central cell with three rings surrounding it (Tak, 2012)

In order to model the evolution of the population occupying the field, the field has to be updated. Each update round can be interpreted as the types playing their game once, leading to offspring. An update of the field consists of two main steps. First of all, for each cell on the field the fitness has to be calculated. The fitness of a cell depends on several factors, for instance on its neighbouring cells and the distance between the neighbouring cells and the central cell. Cells that are close to the cell for which the fitness is calculated have a greater impact on its fitness than cells that are further away. The second step consists of determining which type will occupy a cell in the next generation. This depends on the fitness of the types in that cell's neighbourhood. For each type, the fitness in the neighbourhood is summed up. The type with the highest summed fitness is the type that will occupy the cell in the next generation. Important to note here is that both the calculation of the fitness and the determination of which type will occupy the cell next generation is done simultaneously for each cell. After that, the whole field is updated at once (Bendermacher, 2009; Tak, 2012).

There are a number of aspects of the before mentioned model of local replicator dynamics that can be adjusted. For instance, in the basic model described above it is assumed that the fitness matrix is the same and fixed across the entire field. It is possible, however, to have a varying fitness matrix, both in space and time. For more variations and possible alterations, please refer to Tak (2012).

Chapter 2: Microbiological background

The second chapter gives a short overview of concepts of microbiology and the biomedical domain that will be encountered later on in the literature review. This serves as an introduction to the area of microbiology that is of interest for this thesis. It also gives a summary of some observations made by professionals working in the field of microbiology and in areas related to it.

2.1 Pathogenesis and drug-resistance

A pathogen is any microorganism that is able to cause disease. Whereas some microorganisms are highly pathogenic, and thus cause disease very often, others

are less pathogenic, and rarely cause disease. Virulence is a concept that is used to quantitatively measure this pathogenicity. It is usually measured by the amount of microorganisms necessary to cause disease, so it can also be seen as a measure of the aggressiveness of a certain microbe. Virulence factors are physical properties of a microorganism that influence the aggressiveness of that organism. Whether a microbe is able to survive host defence systems, such as acid in the host's stomach, is an example of such a virulence factor (Levinson, 2006).

Pathogens cause infectious disease when they overpower the host's defence system. Not all hosts are equally susceptible to pathogens, because the defence systems of all hosts have differences and are not equally strong. Two determinants that are of great importance for a microorganism when overpowering the host are the amount of microorganisms as well as their virulence. If there is a higher amount of organisms, it is more likely that an infection will occur. However, the virulence is very important as well, since a highly virulent organism only needs a small population in order to cause disease, whereas a less virulent organism needs many more individuals to get the same effect (Levinson, 2006).

Just as not all hosts are equally susceptible to pathogens, not all pathogens are equally sensitive to medication that is designed to kill them. A pathogen population usually contains a few organisms that are naturally resistant to the administered drugs. For a population of bacteria, the process of developing resistance goes as follows: the population contains both cells that are sensitive and cells that are resistant to drugs. The drug-sensitive cells constitute the majority of the population; there are only a few cells that are naturally resistant. When all cells are exposed to an antimicrobial drug, the drug-sensitive cells are inhibited. This facilitates multiplication of drug-resistant cells, since they do not have to compete with the drug-sensitive cells; a competition they lose in the absence of drugs because drug-resistant strains naturally multiply less quickly. The drug-resistant cells have the opportunity to take over, which eventually happens and leads to the development of a resistant strain (Bauman, 2006).

2.2 Observations

During meetings with microbiologists and other professionals working in fields related to the biomedical domain, some observations were mentioned that might serve as a basis for future research. For the time course of this thesis it was not possible to evaluate these in depth, but for completion they are still mentioned here. One observation was that patients in a hospital sometimes acquire a certain strain of bacteria, but the development of the bacteria differed according to whether it was acquired prior to or during hospitalization. This difference in development could not be explained, and the question was if it would be possible to investigate and explain this phenomenon with models of evolutionary game theory. Moreover, spread of diseases was mentioned multiple times. The general consensus was that there is a need to be able to predict the spread of several diseases, such that precautions can be taken at an early stage and less people will be infected. A last observation was that certain types of bacteria increase their resistance to antibiotics when being brought in contact with it once. It was suggested to investigate this phenomenon with evolutionary game theory as well, in order to be able to understand the processes better (P. Wolffs et al., personal communication, June 4, 2012).

Chapter 3: Literature Review

This chapter gives an overview of some of the literature in the biomedical domain that could possibly relate to the models of evolutionary game theory, provided by a microbiologist. The first section will give the general overview and provide summaries and evaluations of the articles that were not of direct interest. The second section gives a detailed description of the articles that were of interest, and which were chosen as the basis for a case study.

3.1 Overview

There is some literature relating to evolutionary dynamics and game theory. However, not everything is relevant for the aim of this thesis. The following articles were not of interest, for different reasons that are given in each review. The articles are assessed in chronological order.

3.1.1 *A general theory for the evolutionary dynamics of virulence* - Day and Proulx (2004)

Day and Proulx (2004) questioned the applicability of the game-theoretical approach towards the evolutionary dynamics of virulence. According to them, game theory falls short in this respect, since it does not allow for predicting the evolution of virulence. Instead, they argue that there are other methods and techniques that are more advantageous than game theory and that can help to understand this evolutionary process. The method they are elaborating on is an approach that is analogous to quantitative genetics. The authors state that their reasons for choosing this approach are that they believe that it is of the same level of difficulty as game-theoretical models and, more importantly, that it separates evolutionary and epidemiological processes.

The model that Day and Proulx (2004) come up with is a system of three differential equations, of which two represent the epidemiological model of interest and the last one represents the evolution of the mean level of virulence. Parameters and variables that are taken into account in the epidemiological part are the density of the infected hosts, a vector of remaining state variables (such as the density of competitors, the density of resources, the density of predators, etc.), the phenotype of the parasite of interest (which can denote the virulence of that parasite), and a vector for the dynamics of the state variables mentioned earlier. The evolutionary part of the model includes variables and parameters denoting the variation in strain type, the effect on evolution of mutation within a host and the effect on evolution of secondary infection.

It is quite obvious why this article is not suitable to form a link between the local models of evolutionary game theory and the biomedical domain. The article is only shortly about game theory and then elaborates on a different approach of modelling evolutionary dynamics. Although it might have been possible to try to prove Day's and Proulx' (2004) assumptions on game theory wrong by means of the local replicator dynamics, there is not enough elaboration on the perceived shortcomings of game theory to be able to build on that.

3.1.2 *Game theory of social distancing in response to an epidemic* - Reluga (2010)

In case of an epidemic, people can lower their risk of infection relatively easily by reducing the rate of contact with individuals who might infect them. However, Reluga (2010) argues, this social distancing usually comes at a cost (i.e. time, money, convenience, etc.). In this article Reluga (2010) analyses how individuals can best use social distancing and other behaviour to protect themselves during epidemics.

The analysis consists of two parts. In the first part, an epidemiological-economics model is formulated. This model again consists of differential equations. It accounts for the transmission and removal rate, which are used in more general epidemiological models as well, but next to that also considers direct costs of the epidemic, such as the cost of infection per day, the cost of vaccination and the investment in social distancing. The second part then employs differential game theory in order to identify the strategies that constitute an equilibrium. In other words, the differential game represents how individuals should employ social distancing in relation to the behaviour of the rest of the population. The basic idea is that at each point in the game, an individual has two choices: either pay a cost that is associated with social distancing and reduce the risk of infection, or not. This part of the analysis thus attempts to discover how individuals could best respond and employ social distancing, with respect to the behaviour of others and possible costs (Reluga, 2010).

Although the article by Reluga (2010) elaborates extensively on the technical details of the model and game in the analysis, it is not of interest for this thesis. The main reason why this is not the case, is that the analysis concerns the behaviour of people with regards to epidemics, but it does not mention the evolution of bacteria or viruses whatsoever. Since the article is not specifically about the evolution on the microbiological level, this article will not be of help in establishing a link between evolutionary game theory and microbiology.

3.1.3 *Epidemiological game-theory dynamics of chickenpox vaccination in the USA and Israel* - Liu et al. (2012)

In this article by Liu et al. (2012) the epidemiological dynamics of the vaccination against chickenpox are investigated. Contrary to most other articles, this one does not investigate the spread of the actual virus, but how this spread is contained over time by vaccinations. Liu et al. (2012) take two different approaches, namely the approach of the Nash equilibrium in which the decision to vaccinate is driven by self-interest on the one hand, and the so-called utilitarian equilibrium in which the decision to vaccinate is driven by externalities that influence the society as a whole on the other hand. Both of the approaches make use of game theory to assess the perceived costs and benefits and decide upon the most favourable action to take.

The model itself is an epidemic model of the vaccination and transmission of chickenpox that is structured according to age, infection, immunity and vaccination. Liu et al. (2012) use five parameters to determine vaccine efficacy. These parameters are the probability that the vaccination failed, the probability that the protection against the virus is only temporarily, the rate of decline of immunity, the probability of infection without direct symptoms and the probability of developing chickenpox. Besides that, the costs of vaccination and infection were computed as well. After that, the Nash equilibrium was determined by identifying the percentage of the population that is vaccinated such that an individual under consideration does not have an incentive to deviate from its current strategy. The utilitarian equilibrium, on the other hand, was found by determining at which level of vaccination coverage the total expected costs to society were minimized (Liu et al., 2012).

Again, there were a few reasons why this article was not directly useful. As the article before, Liu et al. (2012) do not investigate the virus itself, but people's behaviour with regards to this pathogen. Although in this article people's behaviour influences how the virus spreads and evolves, this is not the focal point of the article and is not elaborated on. Furthermore, since the area of interest is the biomedical domain and in particular the evolution and antibiotic resistance of bacteria, this article about viruses cannot contribute to the purpose of this thesis.

3.1.4 *New insights into virulence evolution in multigroup hosts* - Williams (2012)

Williams (2012) states that host populations are not always homogeneous, as is often assumed. Instead, he argues that individuals in the host population usually are not equally susceptible and vulnerable to pathogens. Such a heterogeneous host population can occur for a number of reasons, e.g. because of the use of different vaccines that are not all as effective. The differences in susceptibility and vulnerability in the host population can be influential in the evolution of the pathogens, as well as in the evolution of their virulence.

The model Williams (2012) proposes to support his stance consists of game theoretic equilibrium expressions. These expressions serve to show clearly what effect heterogeneity of the host population has on the evolution of the pathogen. Unfortunately, at the time of writing, this article was not fully available yet, which means that the details of the game theoretical model are not known and cannot be related to the local replicator dynamics. Besides, it is not clear if the evolution of the host population also influences the evolution of the pathogens, which would make adjusting the model more complicated. Therefore, this article is not considered for serving the purpose of this thesis either.

3.2 Articles of choice

Both Kerr et al. (2002) and Kirkup and Riley (2004) researched the evolution of populations of the *E. coli* bacteria. They were all interested in three different strains of these bacteria that behave according to rock-paper-scissors dynamics, in the sense that the first strain beats the second, the second beats the third and the third beats the first. These three strains are a strain that produces colicins (C), a strain that is sensitive to these colicins (S) and a strain that is resistant to them (R). Colicins are a type of toxin, in particular a type of antibiotic that this strain of the *E. coli* bacteria produces itself. When the three strains are put together, they satisfy the rock-paper-scissors relationship. S is able to displace R, because the growth rate of S is higher than the growth rate of R, R can displace C, because of a higher growth rate of R with respect to C, and C displaces S, because S is sensitive to the colicins produced by C.

3.2.1 *Antibiotic-mediated antagonism leads to a bacterial game of rock-paper-scissors in vivo* - Kirkup and Riley (2004)

Kirkup and Riley (2004) did their experiment *in vivo*, meaning that they placed the bacteria populations in living organisms, mice in this case. They used 36 mice, placing a population of one strain of the *E. coli* bacteria in each mouse colon; 12 mice got the colicin-producing strain, 12 received the sensitive strain and in the other 12 mice they placed a population of the resistant strain. After that, they placed 3 mice together in a cage, such that each strain was represented in every cage, making each cage a model community. For 12 weeks they measured the amounts of the different strains of bacteria once a week. The results of this experiment are shown in figure 3.1.

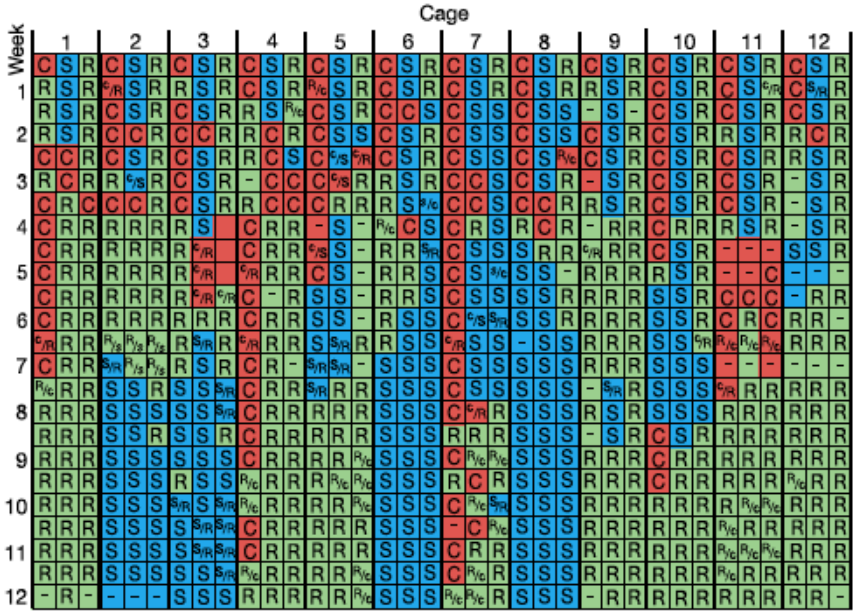

Figure 3.1 The results of the *in vivo* experiment (Kirkup & Riley, 2004)

As can be seen, in 8 out of the 12 cages the resistant strain took over and populated all three mice. In the other 4 cages the sensitive strain did the same.

One thing that was puzzling for Kirkup and Riley (2004) was the fact that some strains showed up again after they seemed to have disappeared. An example of this can be seen in cage 10, where each of the strains at a certain point is not represented anymore, but suddenly shows up again some weeks later. They generated three hypotheses that could possibly explain this reappearance. The first one was that the strains of bacteria that disappeared could also be found in the environment and that the mice reacquired those

strains from their surroundings. The second hypothesis was that at a certain point the population of a particular strain was so small that it fell below the detection threshold. The strain was still there but could not be detected. The last hypothesis stated that mutations were the reason behind strains reappearing. Kirkup and Riley (2004) concluded that the second hypothesis was most likely to accurately explain the reappearance, since their data suggested that the other two hypotheses were highly unlikely to be the cause.

The main conclusion of the experiment was that "colicins are effective antagonistic agents within *E. coli* populations in an animal host" (Kirkup & Riley, p. 413, 2004). The experiments also showed that the colicins produced by this strain of *E. coli* bacteria play a significant role in the strain dynamics by taking the role of mediator. Another observation was that strain diversity within a cage was lost very quickly. In nearly all cages only one strain was detected in the last week (Kirkup & Riley, 2004).

3.2.2 *Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors* - Kerr et al. (2002)

According to research done by Kerr et al. (2002) the maintenance of strain diversity in this system of *E. coli* bacteria is dependent on the existence of local populations. As mentioned before, their experiments concerned the same strains of bacteria as the experiments of Kirkup and Riley (2004). The difference is that Kerr et al. (2002) did not perform their experiments *in vivo*, but *in silico* and *in vitro*, meaning that they first made predictions with a computer simulation (*in silico*) and afterwards did experiments with bacteria outside of living organisms (*in vitro*).

The computer simulation made use of a 250×250 lattice consisting of squares. Initially, random assignment made sure that every lattice point, a point where four squares touch, was occupied by S, R, C or the empty state. To update the lattice, focal points were randomly selected and their states were changed probabilistically. The new state of the focal point was dependent on both its current state and the states of the points of its neighbourhood. The neighbourhood could either be local, consisting of the 8 points directly surrounding the focal point, or global, consisting of every point on the grid

outside of the focal point. The updating process is the same for either neighbourhood. If the focal point is an empty point, the probability that it will be filled with type i, with $i \in \{S, C, R\}$, is equal to f_i, the fraction of the neighbourhood that is currently occupied by i. If, on the other hand, the focal point is already occupied by type i, this type will be killed with probability Δ_i. Whereas Δ_R and Δ_C are fixed values, Δ_S is dependent on the number of C cells surrounding the focal point with state S. The value of Δ_S can be computed by taking $\Delta_{S,0}$ (the probability of death of an S cell without any C cells in the neighbourhood) and adding $\tau(f_C)$, which measures the toxicity of any neighbouring C cells (Kerr et al., 2002).

With the above-mentioned model, Kerr et al. (2002) were able to make some predictions about the evolution of communities consisting of the three strains of the *E. coli* bacteria (so-called C-S-R-communities). Their predictions are shown in figure 3.2.

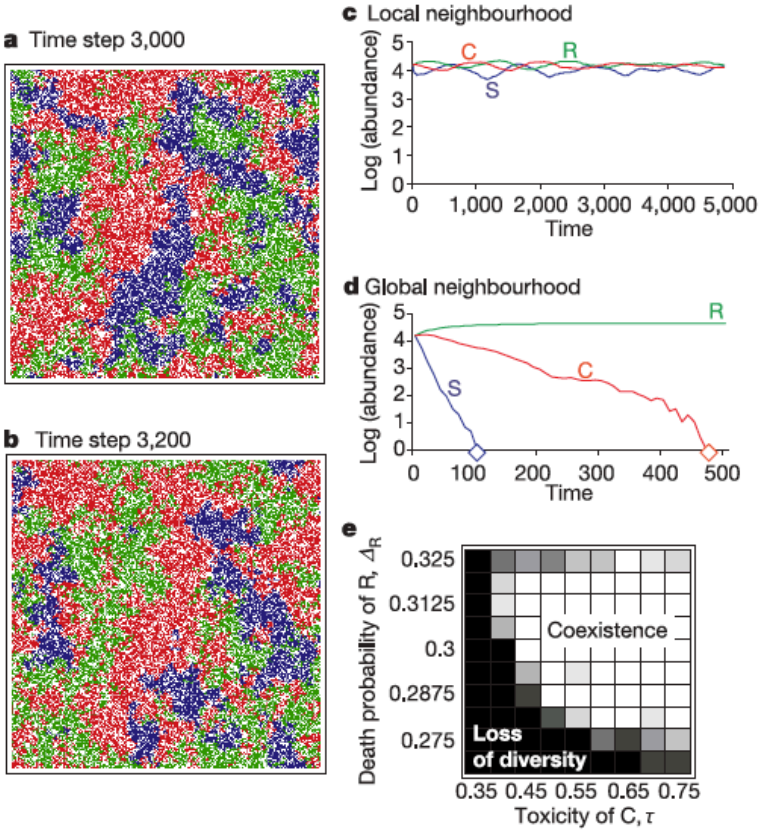

Figure 3.2 The predictions of the *in silico* experiment (Kerr et al., 2002)

What they observed was that when the interactions were local, types started to clump together, as can be seen in figure 3.2a and 3.2b. Different clumps then chased one another according to the rock-paper-scissors dynamics, so clumps of type C displaced the ones of type S, who invaded clumps of type R, who in turn displaced the ones of type C. This 'balanced chase' was the reason for the maintenance of strain diversity. On the contrary, when the interactions were global, this balance was not found. In this situation they observed that C redistributed continuously, which led to extinction of S. With S gone, R took its chance and could outcompete C. The graphs in figure 3.2c and 3.2d illustrate these observations. From these graphs it can also be derived that when interactions between strains are local, it is possible for the strains to coexist, whereas R quickly takes over the environment in case of global interactions (Kerr et al., 2002).

To test if their predictions corresponded to reality, Kerr et al. (2002) also did an *in vitro* experiment. They compared the evolution of C-S-R-communities in three different environments: in a flask, on a static plate and on a mixed plate. The flask constituted a well-mixed environment. It contained a liquid medium and the flasks were shaken when growing the bacteria, such that every bacterium could get in contact with every other bacterium. Therefore, the neighbourhood of each bacterium in the well-mixed flask environment was global. The static plate, on the other hand, constituted an environment in which interaction was primarily local. In this case the bacteria were grown in Petri plates, on the surface of a solid medium. Every 24 hours the bacteria were transferred by pressing the plate onto a platform covered in sterile velveteen cloth. By doing this, a small sample was transferred and the spatial pattern that developed on the previous plate could be preserved. The last environment, the mixed plate, was identical to the static plate environment, except for the fact that while transferring the bacteria onto new plates, the old plate was pressed against the velvet multiple times. Each time the plate was rotated at a random angle, such that the spatial pattern of the previous plate got lost.

After the experiments were finished, Kerr et al. (2002) again made graphs representing the evolution of the C-R-S-communities. These graphs turned out to be highly similar to the predictive graphs from the computer simulation, as can be seen when comparing figure 3.3 to figure 3.2.

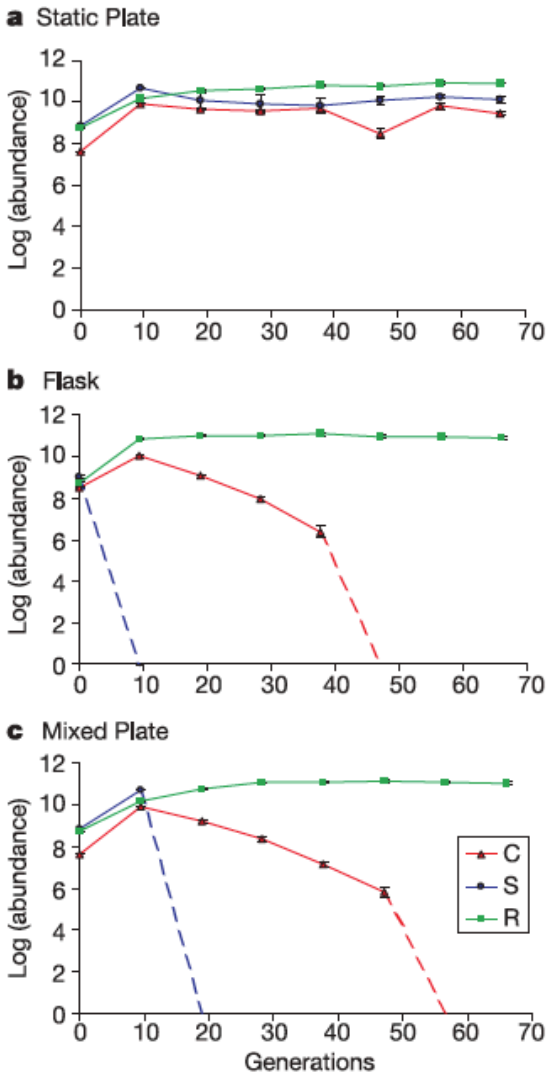

Figure 3.3 The results of the *in vitro* experiment (Kerr et al., 2002)

Moreover, photographs of the static plate environment indicated that indeed the strains clumped together and R chased C, C chased S and S chased R. The conclusions Kerr et al. (2002) derived from these experiments were that the three strains of the *E. coli* bacteria could coexist without absolute spatial isolation. Besides, interactions between the strains took place at the boundaries of the clumps of bacteria of the same type, supporting the prediction that strain diversity is maintained by means of 'balanced chasing'. In contrast, in the flask environment, the levels of S and C strains dropped below the detection threshold

before the experiments were completed, meaning that R could take over. These dynamics were also found in the mixed plate environment. Since all other factors were kept constant, this difference in dynamics between the static plate environment and the flask and mixed plate environments could only be because in the one case the interactions were local and in the other they were global. These results supported the predictions, and led to the main conclusion that diversity can be promoted by non-hierarchical relationships, such as the rock-paper-scissors dynamics, but in this case interactions need to be localized, because otherwise diversity gets lost.

Chapter 4: Case study - Bacterial populations playing a game of rock-paper-scissors

In the previous chapter the last two articles were elaborated on in detail. The first section of this chapter evaluates shortly why these articles were chosen for further investigation by means of the models of evolutionary game theory, both global and local. Afterwards, the second section shows how the game theoretical models can be adjusted, in order to provide a clear link between the phenomena observed in the biomedical domain and the models used in evolutionary game theory.

4.1 Motivation articles of choice

At first, only the article by Kirkup and Riley (2004) was considered. This article is interesting, since it concerns a game of rock-paper-scissors. In evolutionary game theory, this game has been evaluated, and an important outcome of this game theoretical evaluation of rock-paper-scissors is that it has been observed that the game does not have an Evolutionarily Stable Strategy. For instance, the matrix displayed below is a typical rock-paper-scissors matrix.

$$\begin{array}{c} \\ \text{Rock} \\ \text{Scissors} \\ \text{Paper} \end{array} \begin{array}{c} \begin{array}{ccc} \text{Rock} & \text{Scissors} & \text{Paper} \end{array} \\ \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \end{array}$$

Looking at what a player of this game could do, playing a best response would in this case mean playing the mixed strategy $(\frac{1}{3},\frac{1}{3},\frac{1}{3})$. When playing this strategy, the

population distribution would be stable; each type would constitute $\frac{1}{3}$ of the population. However, this population distribution is not robust to mutants, which is why the mixed strategy is not an ESS. If, for instance, paper starts mutating into rock, the population distribution is not $(\frac{1}{3},\frac{1}{3},\frac{1}{3})$ anymore but it starts to behave according to an oscillating pattern, because there are suddenly a lot more rocks that win from scissors. When the fraction of scissors decreases, however, paper gets the chance to flourish. When this happens, paper still wins from rock, which makes the fraction of rocks diminish again. However, when the fraction of rocks goes down, the fraction of scissors that wins from paper can increase. This pattern keeps repeating in time. For the local replicator dynamics, starting with a random playing field does not lead to a stable population distribution either. It has been observed that when starting out with a random field and going from one generation to the next, the field stays random (F. Thuijsman, personal communication, May 10, 2012).

What was curious in the article by Kirkup and Riley (2004) was that they describe how the strains of the *E. coli* bacteria play a game of rock-paper-scissors, but game theoretically speaking this does not seem to be the case. Their experiments showed that in 8 out of 12 cages of mice the resistant strain R took over, and in the other 4 cages the sensitive strain S did the same. If the strains truly play rock-paper-scissors, this would not happen. Either the population distribution would remain equal (every strain represents $\frac{1}{3}$ of the total population), so all of the strains remain present, or - if strains mutate - an oscillating pattern would occur with at one point in time one of the strains being overrepresented, but at a another point in time a different strain being overrepresented. However, since Kirkup and Riley (2004) mention that mutations are unlikely to have taken place, from a game theoretical point of view it would have been expected that every strain would have been represented equally at every stage of the experiment.

Kerr et al. (2002) on the other hand, find that strain diversity is maintained, with all strains represented roughly equally when interactions are local. This implies that on a local level, indeed a game of rock-paper-scissors is taking place. However, both their *in vitro* experiments and their computer

simulation show that strains of the same type form clumps and group together, something that is not found in the application of local replicator dynamics to the game of rock-paper-scissors. The difference between the simulation model used by Kerr et al. (2002) and the local replicator dynamics, is that the local replicator dynamics take the fitness of the individuals into account, while the simulation described earlier does not. The question is now, how can the local replicator dynamics be adjusted such that they can model the behaviour of these three strains of bacteria?

4.2 Applying the local replicator dynamics

In order to model the outcomes of the experiments by Kerr et al. (2002) and Kirkup and Riley (2004) with models of evolutionary game theory, an implementation of the local replicator dynamics in MATLAB was used. The aim was to replicate the processes that are shown in figures 3.2 and 3.3. Several alterations of the basic models were tried out, but not all led to the desired result.

Firstly, a so-called realization was implemented. The idea behind this realization is that the field is not updated at once. Instead, each round two cells are randomly selected and those cells are the only ones playing that round. However, this did not lead to an outcome similar to the lattice and graphs that Kerr et al. (2002) found. When starting with the same playing field and the same rock-paper-scissors matrix in three different runs of the program, each time a different type took over the complete field. Secondly, an implementation in which multiple individuals of the same type were allowed to occupy one cell was experimented with. Yet, this method did not mimic the processes observed in the articles either.

It turns out that the underlying fitness matrix is not a typical rock-paper-scissors matrix like the one presented before. The fitness matrix corresponding to the processes observed by Kerr et al. (2002) and Kirkup and Riley (2004) is a matrix looking like the one displayed below.

$$\begin{array}{c} \\ \text{Type 1} \\ \text{Type 2} \\ \text{Type 3} \end{array} \begin{array}{c} \begin{array}{ccc} \text{Type 1} & \text{Type 2} & \text{Type 3} \end{array} \\ \begin{pmatrix} 0 & 1 & 0 \\ 3 & 0 & 0 \\ -2 & 2 & 0 \end{pmatrix} \end{array}$$

This matrix displays the observed dynamics between the three strains of the *E. coli* bacteria. It can be seen that type 2 wins from type 1, because type 2 playing as row player against type 1 as column player yields a payoff of 3, while type 1 as row player against type 2 as column player only leads to a payoff of 1. Following the same reasoning it can be derived that type 3 beats type 2 (payoff of 2 against 0) and type 1 wins from type 3 (payoff of 0 against -2). When looking at the global dynamics for this fitness matrix, type 3 will take over the population. This happens because type 2 quickly overrules type 1. Since the fraction of 1 decreases fast, at a certain point type 3 wins more against 2 than it loses against 1. This process was observed in the results of the experiments by Kerr et al. (2002) and Kirkup and Riley (2004) as well. In particular, the graphs in figures 3.2d, 3.3b and 3.3c clearly show these dynamics, with type 1 corresponding to S, type 2 to C and type 3 to R. Looking at the local dynamics, however, all types remain. Kerr et al. (2002) found this result too, as can bee seen in figures 3.2c and 3.3a. The fitness matrix above thus correctly models the findings from Kerr et al. (2002) as well as those from Kirkup and Riley (2004). This analysis of the local replicator dynamics in relation to the interaction of three strains of the *E. coli* bacteria shows that it is possible to apply models of evolutionary game theory to phenomena that are observed in the biomedical domain.

Discussion and conclusion

As can be derived from this thesis, it is indeed possible to apply evolutionary game theory to observations and phenomena in microbiology. It has been shown that the local models of evolutionary game theory can replicate the results of biomedical experiments. In particular, from a literature review a situation of interest was chosen, after which a fitness matrix for three strains of the *E. coli* bacteria behaving according to rock-paper-scissors dynamics has been found.

From the application of the local replicator dynamics an important issue came forward. Both the articles by Kerr et al. (2002) and Kirkup and Riley

(2004) stated that the *E. coli* bacteria played the game of rock-paper-scissors. Game theoretically speaking, however, this is not the case since the fitness matrix is not a typical rock-paper-scissors one. It is important to distinguish between organisms behaving in a certain way and determining which game they actually play. It is true that the three strains of the *E. coli* bacteria interact just like rock, paper and scissors do in their game. That is, it has been shown that indeed one strain beats the second, the second beats the third and the third again wins from the first. However, stating that bacteria, or any other organism for that matter, play a game of rock-paper-scissors cannot be concluded on this basis. When looking at the fitness matrix of the bacteria it can be seen that they indeed play a different game. Therefore, researchers should be careful when stating that microorganisms play a certain game.

Furthermore, this thesis provides opportunities for further research. As mentioned before, from meetings with professionals in the biomedical domain and fields related to that, it became clear that there are observations made for which no explanation can be found or for which predictions can be useful. Models of evolutionary game theory can possibly be applied to these situations. Future research can focus on real-life situations, instead of literature, which was the focus of this thesis. With input from professionals in microbiology and related domains, it will be possible to find a new range of applications.

Bibliography

Bauman, R. W. (2006). *Microbiology*. San Francisco, CA: Pearson Benjamin Cummings.

Bendermacher, C. H. M. (2009). *The evolution of populations using local replicator dynamics.* Maastricht University, Maastricht.

Day, T., & Proulx, S. R. (2004). A general theory for the evolutionary dynamics of virulence. *The American Naturalist, 163*(4), E40-E63.

Kerr, B., Riley, M. A., Feldman, M. W., & Bohannan, B. J. M. (2002). Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors. *Nature, 418*(6894), 171-174.

Kirkup, B. C., & Riley, M. A. (2004). Antibiotic-mediated antagonism leads to a bacterial game of rock-paper-scissors in vivo. *Nature, 428*, 412-414.

Levinson, W. (2006). *Review of medical microbiology and immunology*. New York, NY: The McGraw-Hill Companies.

Liu, J., Kochin, B. F., Tekle, Y. I., & Galvani, A. P. (2012). Epidemiological game-theory dynamics of chickenpox vaccination in the USA and Israel. *Journal of The Royal Society Interface, 9*(66), 68-76.

Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge, UK: Cambridge University Press.

Maynard Smith, J., & Price, G. (1973). The logic of animal conflict. *Nature, 246*, 15.

Nash, J. (1951). Non-cooperative games. *The Annals of Mathematics, 54*(2), 286-295.

Osborne, M. J. (2004). *An introduction to game theory*. Oxford: Oxford University Press.

Rapoport, A., & Chammah, A. M. (1966). The game of chicken. *American*

Behavioral Scientist, 10(3), 10-28.

Reluga, T. C. (2010). Game theory of social distancing in response to an epidemic. *PLoS Computational Biology, 6*(5).

Tak, M. J. W. (2012). *Experiments in local evolutionary dynamics.* Maastricht University, Maastricht.

Taylor, P. D., & Jonker, L. B. (1978). Evolutionary stable strategies and game dynamics. *Mathematical Biosciences, 40*(1), 145-156.

Williams, P. D. (2012). New insights into virulence evolution in multigroup hosts. *The American Naturalist, 179*(2), 228-239.