Controlling the evolution of antibiotic resistance

A thesis presented for the degree of

Doctor of Philosophy of Imperial College

by

Rafael Pena-Miller

Department of Mathematics

Imperial College

180 Queen's Gate, London SW7 2BZ

JANUARY 21, 2011

I certify that this thesis, and the research to which it refers, are the product of my own work, and that any ideas or quotations from the work of other people, published or otherwise, are fully acknowledged in accordance with the standard referencing practices of the discipline.

Rafael Pena-Miller
Department of Mathematics
Imperial College London

Copyright

Abstract

Modelling antibiotic resistance evolution is inherently a multiscale problem: from the physical interactions between drug molecules and their targets to the epidemiology of drug resistance in clinical settings. Although predicting the evolution of resistance is a difficult and ongoing problem, it is known that pathogens are continually adapting to our drug prescription patterns. For this reason, as well as the continual downturn in the discovery of new drugs, the question of how to best deploy antibiotics has never been more pressing.

The purpose of this thesis is to use tools from control and systems theory to ask the following fundamental question: How can we design rational antibiotic deployment strategies that do not promote the evolution of antimicrobial resistance?

By re-examining epidemiological models from the literature, in the first part of this thesis we show that the optimal drug deployment protocol has a universal structure not determined by biological detail. This class of epidemiological models, however, provide insight into the underlying mechanisms that influence the spread of disease at the population level but fail to capture the complex molecular interactions between different antibiotics and bacteria, as well as to provide an experimental system to test the efficacy of different treatment protocols. Therefore in the second part of the thesis we pose an evolutionary model of an experimental microbial system that allows us to study drug interactions and the effect that combination treatments have on the evolution of multidrug resistance. Again, using optimal control theory we design drug deployment protocols that minimise conditions promoting the evolution of antimicrobial resistance in a single host.

Finally, in the last part of the thesis we propose an epidemiological model where patients are considered as individual agents receiving antimicrobial treatment in a clinical setting. This stochastic and spatially explicit model allows us the possibility to evaluate the efficacy of different drug usage strategies. We conclude with a general principle: the best performing drug usage policies utilise the highest quality of available information.

On Exactitude in Science

In that Empire, the Art of Cartography attained such Perfection that the map of a single Province occupied the entirety of a City, and the map of the Empire, the entirety of a Province. In time, those Unconscionable Maps no longer satisfied, and the Cartographers Guilds struck a Map of the Empire whose size was that of the Empire, and which coincided point for point with it. The following Generations, who were not so fond of the Study of Cartography as their Forebears had been, saw that that vast Map was Useless, and not without some Pitilessness was it, that they delivered it up to the Inclemencies of Sun and Winters. In the Deserts of the West, still today, there are Tattered Ruins of that Map, inhabited by Animals and Beggars; in all the Land there is no other Relic of the Disciplines of Geography.

Suarez Miranda,Viajes de varones prudentes, Libro IV,Cap. XLV, Lerida, 1658.

From Jorge Luis Borges, Collected Fictions.

Acknowledgements

First of all I would like to thank my supervisor, Rob, for the countless beers he has bought me whilst explaining me the rules of cricket. Working with Rob has been a highlight of my life and I cannot overstate how much I have learned from his enthusiastic, creative and innovative approach towards research.

Many thanks to Martin whose friendly and helpful character encouraged many inspirational conversations throughout my PhD. I would also like to thank Hinrich, David and Gunther for the daata and for being so patient with the crazy english mathematician and his mexican student gatecrashing their lab.

In addition, I would like to thank the members of the BioMaths group at Imperial College for all the lively discussions we've had, in particular to Ivana and Vahid. I also wish to acknowledge the contributions, advice and suggestions of Alex Hall, Gabe Perron, Craig MacLean and Sylvain Gandon.

Because not everything in life is science (there is also football), I would like to thank everybody who shared a football pitch with me during my stay in London, either a friendly kickabout in Hyde Park or a competitive match under the A4... *¡Que bien que juega el Acapulco Tropical!*

I also want to thank all my mates from the Department of Mathematics that made Imperial College and its grim grey corridors a lovely place to work. In particular I would like to thank David, Ed, Flor, Kahn, Ewan, Chris and Jack. My *carnales* from the photonics lab deserve a special mention too: Pavel, Mac, Ren and Matt. The mexican crowd has also been very helpful on those long and dark London nights: Ana, Ioscani and Pablo, but specially Aldo and Nata that managed to make me feel at home from the moment I arrived to Heathrow.

Of course, a special acknowledgement to my family, and in particular to José and JoAnn, that have been very caring and supportive not only during my PhD but throughout my entire life. This thesis belongs to them as much as it does to me.

But above all I would like to thank Ayari for sharing this journey with me. Every morning I consider myself lucky to wake up next to someone that encourages and understands me the way she does.

(and both CONACYT and SEP for funding my PhD)

Table of contents

List of Figures

List of Tables

Introduction

In 1940, Albert Alexander, a 48 year old London policeman cut himself while shaving. It was only a minor cut but his case had a major impact in the history of medicine when he developed septicemia and was treated with an antimicrobial substance Howard Florey and Ernest Chain had just recently purified. As with many of the great advances in medicine, the treatment was a success but the patient died. The following attempts to treat blood poisoning, however, were successful. This promising antimicrobial agent was discovered accidentally in Saint Mary's Hospital by Alexander Fleming in 1928, when he observed in a contaminated dish a mould that prevented the growth of *Staphylococcus aureus* (a picture of the famous plate is reproduced in Figure 1a). The mould was later identified as *Penicillium notatum* and therefore the antibacterial substance it produced was called *penicillin*.

A few years later, in addition to pneumonia and septicemia, at the time the major causes of death in hospitals, also syphilis, gonorrhoea, meningitis, tonsillitis, rheumatic fever, scarlet fever, diphtheria and many other diseases were successfully treated with penicillin. But not everything was good news with this new "wonder drug". Only a few years after it was introduced in clinical settings, 25% of the Staphylococcal strains were penicillin-resistant and by the 1970's resistance was present in more than 80% of the isolates, a trend shown in Figure 1b. Nowadays, not only almost none of the staphylococcal infections are susceptible to penicillin, but they also present high levels of resistance to other antibiotics like *methicillin* and *vancomycin*.

Unfortunately Staphylococcus is not an exception, most of bacterial pathogens present a remarkable ability to evolve drug resistance to multiple antimicrobial classes. For example,

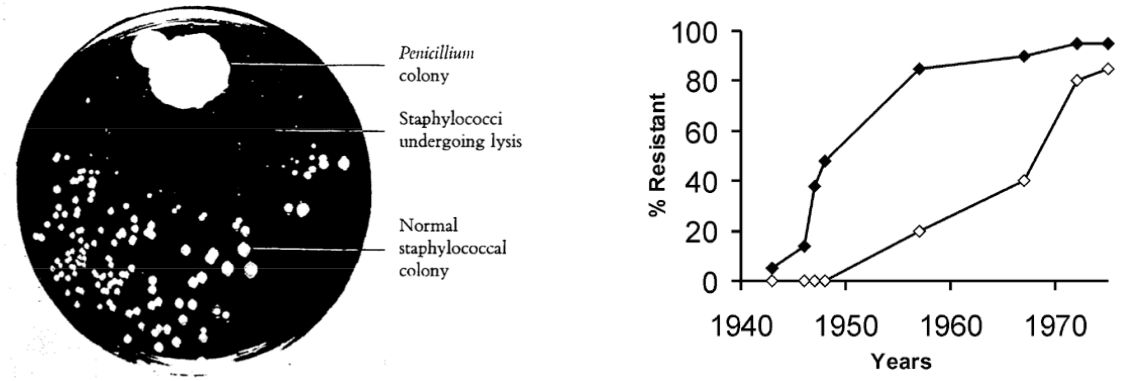

(a) **(b)**

Figure 1: a) Photograph of a culture-plate showing the dissolution of staphylococcal colonies in the neighbourhood of a *Penicillium* colony, reproduced from [40]. b) Rates for penicillin-resistant and meticillin-susceptible strains of *Staphylococcus aureus* in hospitals (closed symbols) and the community (open symbols), figure extracted from [22].

Mycobacterium tuberculosis has been reported to be resistant to as many as eight different antibiotics, making some individuals with tuberculosis practically incurable [14]. Actually, it is known that there is not a single antibiotic class for which bacteria have not evolved drug-resistance mechanisms [62].

For decades, the antibiotic resistance problem was addressed with a world-wide effort in finding new antimicrobial agents, but unfortunately the pharmaceutical industry has basically stopped developing novel antimicrobial substances, focusing their efforts on long-term treatment of chronic conditions [56, 79]. For instance, the United States Food and Drug Administration approval of new antibacterial agents decreased 56% from 1983 to 2002 [97].

The paucity of new antimicrobial agents has precipitated what the authors of [44] call "*the search for synergy*", referring to the efforts of the pharmaceutical and medical communities to find antimicrobial agents that increase their efficacy when used in combination, a drug interaction known as synergism. For example, the standard treatment protocol for tuberculosis in the 1980s was based on a three-drug combination cocktail: *isoniazid*, *rifampin* and *pyrazinamide* [51]. This regimen was recommended by the American Thoracic Society and the CDC when most patients were infected with a susceptible pathogen [1]. A few

years later, however, this treatment protocol was no longer adequate because of the high incidence of drug-resistant infections [2]. This observation highlights an important issue when designing drug combinations: they have to be dynamic in time.

Another possible strategy is to modify drug consumption patterns in health care centres, for instance by prioritising and restricting different antimicrobial classes according to a predefined hospital-wide policy. A central question, and the one we address in this thesis, is how should two antibiotics be utilised to treat infected hosts in a hospital or community so as to minimise the prevalence of antibiotic-resistant pathogens? Clinical trials and theoretical analyses have been conducted to assess the impact different antibiotic deployment protocols have upon the evolution and spread of drug resistance, moreover previous theoretical work on this important problem has created a library of epidemiological models we can explore.

An intervention that has received special attention is antibiotic cycling, a policy in which two or more antibiotic classes are alternated in time [80]. In a published article [9], we argue that by relaxing the periodicity constraint of antibiotic cycling, it is theoretically possible to design drug rotation protocols that are arbitrarily close to the optimum. This result contradicts previous theoretical research that concluded that "*cycling antibiotics may not be good for your health*" [61], claiming instead that antibiotic deployment protocols based on so-called antibiotic mixing are optimal [13, 15]. In mixing protocols, two or more antimicrobial drugs are used simultaneously for a particular infection and different patients are treated with different drugs. We will demonstrate later in this thesis that these strategies based on the random, per-patient allocation of drugs can only be optimal in numerically rare, non-generic theoretical cases that are not of broad-scale relevance.

Our techniques of analysis make our approach, and subsequent conclusion, distinct from previous theoretical research. In the first part of this thesis we apply both classical theory and bespoke optimisation algorithms to determine near-optimal controls for models from the literature. It is important to mention, however, that in order to exactly determine the optimal treatment or even the optimal mixing protocol, one would need a model with a perfect understanding of the future. But is it possible to determine effective treatment strategies, perhaps suboptimal, based only on current and past patterns of susceptibility

and resistance? In response to this we propose the development of feedback controls that take information directly from current observations to inform and adjust future treatment policies. We believe that this rationale, if applied to public health, could prove invaluable in the design of effective antibiotic therapies. Of course, we are not promoting the use of such simple rules for controlling hospitals, as that statement would be far beyond the scope of this thesis. But we argue that designing such heuristics in theoretical models of drug deployment provides a tool to probe the surveying effort required to achieve near-optimal policies.

In response to our paper, the authors of [16] state that even though our mathematical analysis is correct, non-periodic rotation protocols computed using optimisation techniques are impracticable policies that rely on the availability of a model that perfectly describes the patient population dynamics. We agree with Bonhoeffer et al. that without a model calibrated against clinical data, it is of limited interest to compute the theoretical optimum and therefore we only use it as a baseline to evaluate the efficacy of other practicable strategies. Moreover, we believe that although epidemiological models provide insight into the underlying mechanisms that influence the spread of disease at the population level, they fail to capture the complex molecular interactions between different antibiotics and their targets, as well as to provide an experimental system to test the efficacy of different treatment protocols. Hence, in the second part of this thesis we will pose an evolutionary model of an experimental microbial system that will allow us to control the input of different antibiotics in time, in order to study the ecological and evolutionary implications of different drug deployment policies within a single-host.

In particular, we are interested in using this evolutionary model of a microbial microcosm to design effective multidrug combination therapies. We will show, again using results from control theory, that although the commensal microbiota may be outcompeted by an invasive and rapidly evolving pathogen in any fixed multidrug environment, we may still be able to design drug-deployment protocols based on antibiotic rotation that exclude pathogens from the system.

It is noteworthy that this result holds in spite of using two drugs that interact to increase each other's efficacy; even in this extreme situation, the optimal protocol is not to use them in combination but to deploy them sequentially. Furthermore, we will show that even suboptimal but dynamically changing protocols can succeed where any fixed-ratio combination treatment fails. This result suggests that in order to remove an evolving pathogen in the presence of a commensal microbiota, increasing the efficacy of the antibiotics used is not as important as designing rational deployment strategies.

Interestingly, treating every patient optimally could be interpreted at a population level as a mixing strategy, which we have already stated to be suboptimal for the community. The apparent discrepancy between the optimal treatment in a community and the optimal therapy for an individual is a consequence of different paradigms in health care. Should we put society first and preserve the efficacy of antibiotics for the community? Or shall we treat every patient optimally?

To start addressing these questions, in the last part of this thesis we will reconstruct an epidemiological framework where patients are considered as individual agents interacting through the transmission of pathogens. This stochastic and spatially explicit model of a hospital ward will allow us to simulate *in silico* different hospital-wide strategies of drug deployment and study the effect that the community has on the disease dynamics of an evolvable pathogen. This computational framework, in addition to the analysis performed on the standard epidemiological models and the microbial model system discussed in the first two parts of the thesis, will allow us to conclude with a general, model-independent principle: the best performing policies are always dynamic, adaptable and designed using the highest quality of available information.

Part I

Epidemiology of drug resistance

Chapter 1

Emergence and spread of antibiotic resistance

Although newsworthy, the problem of antibiotic drug resistance is not new [40]. It is without exaggeration to say that the continued evolution and spread of antibiotic resistance in clinically relevant pathogens could prove to be a long-term public health menace. Indeed, the rate at which bacteria evolve resistance to antibiotics [28, 45, 49] and the continual downturn in the rate of discovery of new antimicrobials [79] should be sufficient to force us to continually re-examine our attitudes towards antimicrobial therapy.

Different stewardship programs have been proposed to control antibiotic resistance in health care centres [67, 94], with the objective of reducing the overall use of antibiotics, slowing down the transmission rate of resistant strains and minimising conditions that promote the evolution of drug resistance. For instance by introducing strict hygiene guidelines in health care workers, by modifying drug prescription patterns in patients or in clinical settings [91] and in general by continually re-evaluating basic dose adjustment principles [88].

In this thesis we are interested in studying policies based on prioritising and restricting different antimicrobial classes, as well as strategies designed to combine multiple antibiotics with the purpose of increasing drug efficacy while selecting against resistant phenotypes. With this in mind, later we will pose different single-host and epidemiological models that

will allow us to compare the relative merits of different drug usage strategies, but first we would like to discuss an intervention that has received special attention from the medical [17, 80] and theoretical [9, 13, 15, 63] communities: antibiotic rotation.

This policy is based on periodically restricting and prioritising antibiotics from different functional classes, with the intention of decreasing selective pressures to any one antimicrobial class. Although the outcome of some clinical trials support their efficacy [3, 70, 71, 84], other studies argue in the opposite direction [91, 106, 109, 111]. We could cite many studies both for and against cycling, but suffice to say that there is no clear and definitive evidence to support the implementation of drug rotation as a universal principle, as it is further discussed in the following section.

1.1 To rotate, or not to rotate?

Antibiotic rotation was proposed over a decade ago as a way of reducing the incidence of antibiotic-resistant infections. This view, articulated by Niederman in the editorial *Is "Crop Rotation" of Antibiotics the Solution to a "Resistant" Problem in the ICU?* (see [80]) states

> "The 'crop rotation' theory of antibiotic use has suggested that if we routinely vary our 'go to' antibiotic in the ICU (intensive care unit), we can *minimize* the emergence of resistance..."

In the intervening decade, a number of theoretical studies have espoused a different viewpoint in proposing that the heterogenous, random deployment of antibiotics in an ICU unit or hospital can slow the evolution and spread of drug-resistant pathogens. These theoretical studies have concentrated on a direct comparison between cycling and mixing protocols, concluding that mixing reduces the prevalence of resistance optimally [13, 15, 61].

Clinical trials designed to evaluate the efficacy of protocols that increase drug heterogeneity across both hospitals and surgical wards, a property associated with random mixing, have had both beneficial [104] and insignificant outcomes [105]. With respect to cycling protocols, the medical community has tacitly proposed that the nature of antibiotic cycling

could be relaxed when developing trials and conducting theoretical studies. For example, in contrast to previous studies that assume the same period of time for each antibiotic to be deployed, the authors of [17] pointedly ask *'What is the optimal duration of each cycle?'*. In this thesis we address this question by applying classical results from mathematical systems theory, and demonstrate that the optimal treatment is neither antibiotic mixing nor cycling. Instead, a form of drug rotation is optimal. This protocol has some features of cycling in the sense that only one antibiotic should be deployed throughout the population at a time, but the optimal therapy does not cycle periodically; the duration of the cycles must be allowed to vary.

It is important to emphasise there is no discrepancy between the theoretical findings of [13, 15] and the results presented in this thesis; the apparent difference between the two sets of results rests in the interpretation of what *antibiotic rotation* means. The citation of Niederman hints at a scheduled and cyclical rotation that exchanges one drug for another periodically, where that period is fixed at the start of a clinical trial, say, just as a *crop rotation* might only change the crop with each new season. The results presented in [13, 15] show that this idea may not work for standard epidemiological, SI models of drug deployment.

We agree with [13, 15] that there is no theoretical basis to support the optimality of scheduled antibiotic rotation. However, as we will discuss later in this thesis, it is equally true that the random allocation of drugs to each patient is not optimal either.

In terms of empirical evidence for and against the cycling of antibiotics, some studies support rotation [70, 84] but others either advocate against it or at least indicate indifference [106, 109, 111]. For example, the rotation of linezolid and vancomycin in an ICU has been called a 'promising method to reduce infections with MRSA' [96] but rotation has also been implicated as a possible cause of an outbreak of multi-drug resistant *Pseudomonas aeruginosa* [47]. Empirical studies evaluating the efficacy of antibiotic rotation prior to 2005 have also been criticised for 'multiple methodological flaws and a lack of standardization', a particular criticism being the lack of repetition of cycles within rotational protocols [17].

In order to place our analysis into an empirical context we will implement the idea that '...prescription patterns balancing the use of different antimicrobials should be promoted to reduce selection pressure' from [91] to create a *feedback control strategy* that balances the use of different antibiotics. To design the rules for this controller we distill the following observation taken from [3] into a mathematical form:

> "A non-premeditated change of antibiotics in empirical therapy, on the basis of detected resistance patterns, provided promising results in reducing some antimicrobial resistance rates."

We interpret this quotation as a maxim that can be employed to control the spread of resistance in theoretical models of antibiotic use, this maxim states: *if the observed level of resistance to an antibiotic is too* high*, exchange it for a different antibiotic*.

It is the resultant *adaptive rotation* of antibiotics based on the observation, or even partial observation, of those dynamics that may lead to the optimal protocol and minimise selection for drug-resistant pathogens. We arrive at this theoretical result by observing that the optimal protocol will exchange one antibiotic for another across the theoretical ICU unit or hospital, not routinely or randomly, but in a manner commensurate with the epidemiological and evolutionary dynamics observed in each context. For instance, the rotational protocols as they are modelled in [13, 15] switch between the prioritisation of two drugs in such a way that one of them is designated the 'go-to' drug at every moment in time. As we explain later, this form of antibiotic protocol can be written as a bang-bang function which allows us to apply standard control-theoretic results (see [41], for example) and deduce the theoretical optimality, or at least near-optimality, of rotational protocols.

Using optimisation techniques we will compute the optimal rotation protocol for specific models of drug deployment, and illustrate by example that the implementation of very simple heuristics can outperform the random allocation of drugs, but first let us be precise about the differences between antibiotic cycling, antibiotic rotation and antibiotic mixing protocols.

Throughout this thesis, unless stated otherwise, we will consider that we can deploy into the system two antibiotics labelled ‘drug A’ and ‘drug B’.

Definition 1.1.1 *If we denote with $A(t)$ and $B(t)$ the fraction of the population receiving each drug at time t, and we assume that no patient is left untreated, a property expressed mathematically as $B(t) + A(t) = 1$, then a* rotation protocol *satisfies*

$$A(t)B(t) = 0$$

for almost all t within the time interval $[0, T]$. The function $A(t)$ so-defined tells us when to treat everyone *with drug A and hence the number ‘1’ in the must-treat constraint. Consequently if $A(t)$ is zero, everyone in the population is treated with drug B.*

Note that the terms alternating protocol *and* sequential protocol *are used synonymously for the term* antibiotic rotation *in the remainder of the thesis.*

For illustration purposes, we shall use a barcode graphic to denote the deployment of two different antibiotics as part of a rotational protocol, as shown in Figure 1.1. This graphic shows that all patients are treated initially with drug A, before a switch is invoked at time T_1 to drug B.

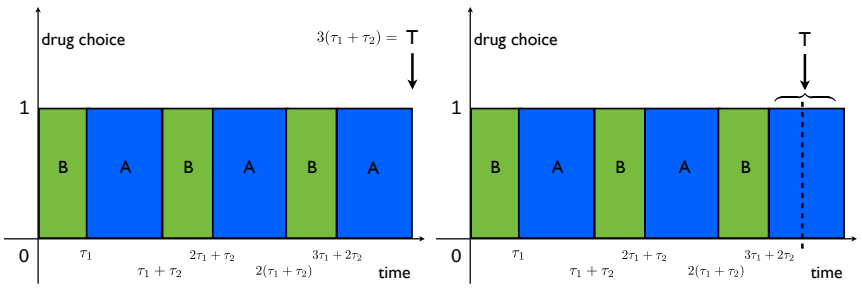

(a) **(b)**

Figure 1.1: Barcodes are used to represent the timing of switches between antibiotics A and B: in a rotational protocol each antibiotic is either deployed at its maximum rate, or is not deployed at all. (a) A congruent cycling protocol whereby $T = 3(\tau_1 + \tau_2)$. (b) A non-congruent cycling protocol: it is possible that the observation time, T, could fall mid-cycle for the best protocols; this possibility is not analysed here.

Definition 1.1.2 Cycling protocols *can be described by any measurable, almost-everywhere periodic function $A(t)$ that can take values 0 or 1, denoting that either none of the patients are being treated with drug A or everyone is being prescribed with this drug. The must-treat constraint $B(t) = 1 - A(t)$ also holds here, so if $A(t) = 0$ nobody is receiving treatment with drug A and drug B is being prescribed across the entirety of the hospital or community.*

Cycling protocols can be characterised using two parameters, τ_1 and τ_2, and a function $A(\tau_1, \tau_2)(t)$, defined for all t within the time interval $[0, T]$ such that

$$A(\tau_1, \tau_2)(t) = \begin{cases} 1 & : \quad 0 \leq t \leq \tau_1, \\ 0 & : \quad \tau_1 < t \leq \tau_1 + \tau_2, \end{cases}$$

and extended so that $A(\tau_1, \tau_2)$ has period $\tau_1 + \tau_2$.

Definition 1.1.3 *If $A(t)$ is constant and $0 \leq A(t) \leq 1$ almost everywhere, then we call this strategy a* mixing protocol. *The epidemiological interpretation of $A(t)$ being a constant is that a fixed fraction of the population is being treated with drug A, while the remaining individuals are prescribed with antibiotic B.*

Any protocol whereby

$$\int_0^T A(t)dt = \int_0^T B(t)dt$$

will be described using the prefix '50-50'.

Therefore a '50-50 mixing' protocol is interpreted as half of the population is receiving drug A while the other half is treated with B. Note how the so-called 'random deployment of drugs' throughout a population satisfies in average with this definition.

1.2 Mathematical models of antibiotic use

From a theoretical perspective, previous studies have investigated the efficacy of different drug deployment protocols both in the community [15] and in health care centres [13, 63], the minimisation of treatment duration [30], the evolution of multiple resistance [15, 29], cost-effectiveness of infection control policies [112] and heterogeneous spatial effects [32].

In the debate on whether periodic cycling, mixing or some other form of antibiotic deployment regimen provides for the most appropriate form of treatment, a number of theoretical studies have indicated that periodically cycling antibiotics may not be the best choice [13, 15, 61]. The modus operandi used within these studies can be described as follows. A mathematical model describes the dynamics of a population supplied with infected individuals from a community who are then treated according to a given hospital-wide policy, where different antibiotic treatment policies are represented by different mathematical functions that may vary in time. By simulating the model with different treatment policies, one can compare their respective outcomes and so make a decision as to which treatment was the most appropriate from all of those simulated. The utility of such a '*what if?*' modus operandi is clear, simply design different drug usage strategies and compare them numerically. This approach, however, cannot begin to address questions of generality or robustness of the most appropriate regimen of those tested.

In the following sections we will briefly describe the most important epidemiological models of antibiotic use published in the literature [13, 15, 27]. These deterministic mathematical models are designed to study the deployment of two antibiotics into a community or a hospital, and do not explicitly incorporate the spatial dynamics of pathogen transmission.

1.2.1 Bonhoeffer, Lipsitch and Levin (1997)

The first theoretical model that described epidemiology of antimicrobial resistance in a community was proposed by Bonhoeffer et al. in [15]. A particular case of this model (Case III) whereby multiply resistant bacteria are not present is illustrated by the compartments shown in Figure 1.2, and its dynamics can be described with the following set of differential

equations:

$$\begin{aligned}
\dot{x} &= \lambda - dx - b(y_w + y_a + y_b)x + r_w y_w + r_a y_a + r_b y_b + ... \qquad (1.1a)\\
&\qquad ... + h(1-s)((f_a + f_b)y_w + f_a y_b + f_b y_a),\\
\dot{y}_w &= (bx - c - r_w - h(f_a + f_b))y_w, \qquad (1.1b)\\
\dot{y}_a &= (bx - c - r_a - hf_b)y_a + hsf_a y_w, \qquad (1.1c)\\
\dot{y}_b &= (bx - c - r_b - hf_a)y_b + hsf_b y_w, \qquad (1.1d)
\end{aligned}$$

whereby the state variable is given by $\boldsymbol{s} := (x, y_w, y_a, y_b)$; x denotes the density of uninfected hosts in a hospital or intensive care unit, y_w is the density of hosts infected by wild-type bacterial strain, y_a are hosts infected with A-resistant bacteria and y_b are hosts infected with B-resistant strains. There are no multidrug-resistant bacterial strains in this model, although that case is discussed in [15].

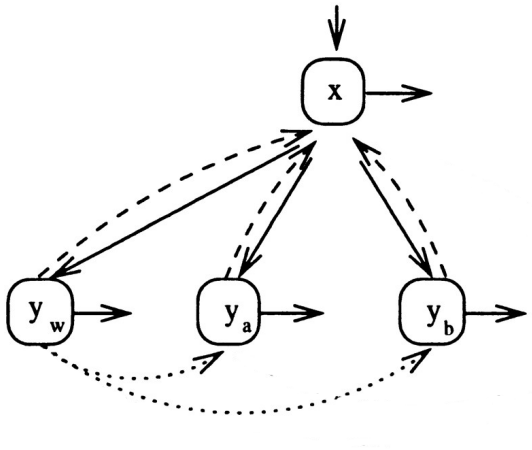

Figure 1.2: Transmission dynamics of the epidemiological model proposed in [15, Case III]; x denotes the density of uninfected individuals, y_w, y_a and y_b hosts infected with susceptible, A-resistant and B-resistant bacteria respectively.

The epidemiological interpretation of the parameters used in (1.1) is described in Table 1.1. It is important to mention that there are two different mechanisms by which patients can be infected with resistant pathogens: through the epidemic spread of resistant phenotypes or by the acquisition of 'de novo' resistance, the latter only occurring in treated hosts.

Table 1.1: Description of the parameters used in the model of drug deployment in the community proposed in [15].

parameter	meaning
f_a, f_b	the fraction of patients treated with antibiotic A and B
r_w, r_a, r_b	recovery rates of wild-type, A-res and B-res infected hosts
b	transmission rate of infection
h	maximum rate at which patients are treated
s	fraction of patients that acquire resistance when treated
d	per capita death rate of uninfected hosts
λ	arrival rate of uninfected hosts
c	infected hosts' death rate

Recovery from infection coincides with the termination of transmission of pathogens between patients, and the fitness cost associated with drug resistance is expressed by a higher recovery rate of infected hosts with resistant bacteria relative to those infected with susceptible pathogens, a property expressed by the inequality $r_r > r_w$. This model does not include superinfection of *wt*-infecteds by resistant pathogens.

In (1.1), f_a is a variable that may depend on time and denotes the proportion of infected hosts treated with antibiotic A, f_b is the proportion of hosts treated with a second antibiotic B, we shall invoke a *must-treat everyone* constraint meaning that $f_a(t) + f_b(t) = 1$ for all times $t \geq 0$.

As discussed by Bonhoeffer et al., the optimal antibiotic policy would be the treatment protocol that preserves and restores antibiotic effectiveness, a goal that could be achieved by maximising the number of uninfected hosts or by minimising the number of infected patients. Another possible objective could be to maximise the time before resistant bacteria constitute some fixed fraction of all bacteria of a given species. The analysis of [15] attempts to balance the short-term benefit due to the use of the antibiotic while trying to preserve drug efficacy in the long term. Mathematically this optimality criterion can be represented by minimising the number of infected hosts integrated over time, that is:

$$\int_0^T (y_w(t) + y_a(t) + y_b(t))\, dt.$$

The purpose of the study presented in [15] is to evaluate the efficacy of different antibiotic usage patterns in preventing the evolution of antibiotic resistance. To achieve this goal, different single-drug treatments and multi-drug combination strategies were compared through straightforward numerical simulations. In particular, periodic cycling protocols were compared with a 50-50 mixing strategy whereby fixed fractions of the population receive each drug at any given moment in time. Combination treatments, where drugs are prescribed simultaneously to each infected host, were also studied.

Definition 1.2.1 *The* 50-50 mixing protocol *for (1.1) is defined by taking a constant value for the treatment protocol f_a, namely*

$$f_a(t) = 1/2$$

for all $t \geq 0$. The interpretation of this condition is that exactly half of all infected hosts are treated with drug A, half with drug B so that $f_b(t) = 1/2$ too. As the model (1.1) does not track individual treatments, so this strategy corresponds to the random allocation of the two drugs per infected patient.

The main results of the study presented in [15] can be summarised as follows:

- The performance of different single-drug strategies is independent of the patterns of antibiotic use, although there is a slight increase in benefit if the drug is used heavily in the early stages.
- Faster rates of antibiotic treatment accelerate the emergence of resistance.
- When treatment is withdrawn, resistant infections reverse to susceptible more slowly than resistance is acquired under treatment.
- When more than one antibiotic is employed, multidrug combinations are in most cases the optimal treatment strategy.
- **50-50 treatments are always superior to antibiotic cycling, regardless of how frequently the drugs are cycled.**

1.2.2 Bergstrom, Lo and Lipsitch (2004)

A different model was proposed by Bergstrom et al. in [13] to study antimicrobial cycling in a hospital setting. This model can be represented by compartments like the ones illustrated in Figure 1.3a, where X denotes uncolonised patients, S denotes the fraction of patients in an ICU unit colonised with an antibiotic susceptible pathogen, R_1 and R_2 represent the fraction of patients colonised by bacteria resistant to antibiotic 1 and 2 respectively. If we denote the state of the system as $\boldsymbol{s} = (S, R_1, R_2, X)$, then the model can be written as:

$$\begin{aligned}
\dot{S} &= \mu(m - S) - (\tau_1 + \tau_2 + \gamma)S + \beta SX + \sigma\beta(c_1R_1 + c_2R_2)S && \text{(1.2a)}\\
\dot{R}_1 &= \mu(m_1 - R_1) - (\tau_2 + \gamma)R_1 + \beta(1 - c_1)R_1X - ... && \text{(1.2b)}\\
&\qquad ... - \sigma\beta(c_1S + (c_1 - c_2)R_2)R_1 &&\\
\dot{R}_2 &= \mu(m_2 - R_2) - (\tau_1 + \gamma)R_2 + \beta(1 - c_2)R_2X && \text{(1.2c)}\\
&\qquad ... - \sigma\beta(c_2S + (c_2 - c_1)R_1)R_2 &&\\
\dot{X} &= \mu(1 - m - m_1 - m_2 - X) + (\tau_1 + \tau_2 + \gamma)S + (\tau_2 + \gamma)R_1... && \text{(1.2d)}\\
&\qquad ... + (\tau_1 + \gamma)R_2 - \beta X(S + (1 - c_1)R_1 + (1 - c_2)R_2). &&
\end{aligned}$$

The interpretation of the parameters used in this model are given in Table 1.2.

Note that [13] also imposes the *must-treat* constraint that $\tau_1(t) + \tau_2(t) = \tau_{\max}$ for all $0 \leq t \leq T$, where $\tau_{\max}$ is a fixed parameter that determines the maximum rate of drug use, and that the *50-50 mixing protocol* is defined by taking a constant value for the treatments: $\tau_1(t) = \tau_{\max}/2$ for all t.

The main difference between this model and the community-based model posed by Bonhoeffer et al. is that here individuals enter and leave the ICU unit at a very high rate, and therefore a hospital is considered to be an open system, while a community is relatively closed. In addition, patients enter the hospital asymptomatically colonised with the species responsible for the nosocomial infection. Moreover, antimicrobial drugs are in principle used at much higher rate in clinical settings than in communities.

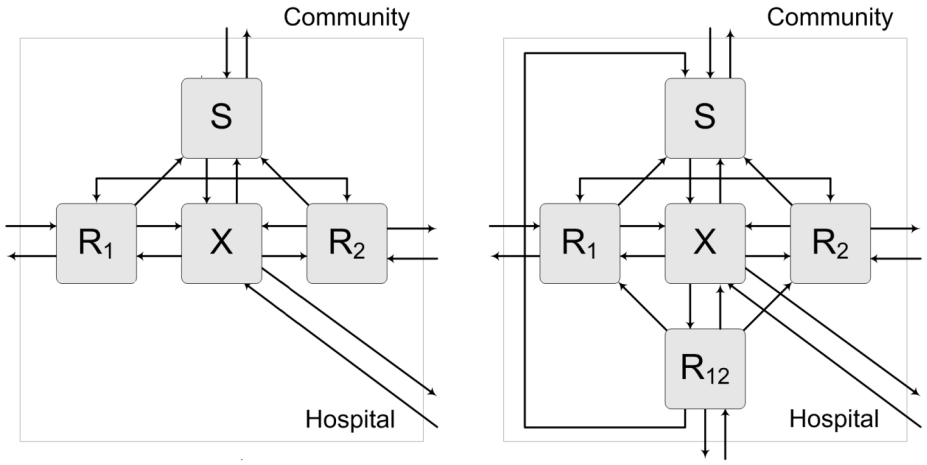

Figure 1.3: (a) Graphical illustration of the model described by equations (1.2). The variables and parameters are explained in the main text. (b) Illustration of the model (1.3) whereby the model (1.2) is extended to include another compartment representing infections by multidrug resistant bacteria.

The purpose of the mathematical study presented by Bergstrom et al. is to illustrate that a simple ecological explanation underlies why, in most of their numerical simulations, antibiotic cycling is outperformed by strategies in which each patient under treatment is prescribed with a random antimicrobial drug.

In order to compare the efficacy of different strategies, the objective in the context of this model can be represented as minimising the frequency of patients colonised by a resistant strain during an observation interval of length T, a goal that can be achieved by minimising the following expression:

$$\int_0^T (R_1(t) + R_2(t))\, dt.$$

The main conclusion drawn from this study is that the prevalence of drug resistance is greater for cycling of any period than 50-50 mixing, and in fact increases monotonically with the cycle period.

Table 1.2: Meaning of parameters used in the model of an ICU unit posed in [13].

parameter	meaning
τ_1, τ_2	rate of use of drugs 1 and 2 per unit time (days)
m, m_1, m_2	patients enter hospital in states S, R_1 and R_2 at rates $\mu m, \mu m_1$ and μm_2 resp.
c_1, c_2	fitness cost of resistance to bacteria
σ	relative rate of secondary colonization to primary colonization
β	rate constant for colonization of uncolonized individuals
γ	untreated patients colonized by susceptible bacteria remain colonized $1/\gamma$ days on average
μ	rate of patient turnover in the hospital
α	represents physician compliance with cycling program

1.2.3 Chow, Wang and Castillo-Chavez (2007)

In order to consider multidrug resistance in the model of an ICU unit described by equations (1.2), an unpublished study by Chow et al. [27] extends this model to include the possibility of individuals being colonised by a bacteria that is resistant to both antibiotic 1 and 2.

The underlying transmission dynamics of this model can be represented by the compartments illustrated in Figure 1.3b, and written mathematically as:

$$\dot{S} = \mu(m - S) - (\tau_1 + \tau_2 + \gamma)S + \beta SX + \sigma\beta(c_1R_1 + c_2R_2 + c_{12}R_{12})S \tag{1.3a}$$

$$\dot{R}_1 = \mu(m_1 - R_1) - (\tau_2 + \gamma)R_1 + \beta(1 - c_1)R_1X - ... \tag{1.3b}$$
$$... - \sigma\beta(c_1S + (c_1 - c_2)R_2 - c_{12}R_{12})R_1$$

$$\dot{R}_2 = \mu(m_2 - R_2) - (\tau_1 + \gamma)R_2 + \beta(1 - c_2)R_2X + ... \tag{1.3c}$$
$$... - \sigma\beta(c_2S + (c_2 - c_1)R_1 - c_{12}R_{12})R_2$$

$$\dot{R}_{12} = \mu(m_{12} - R_{12}) - \gamma R_{12} + \beta(1 - c_{12})R_{12}X + ... \tag{1.3d}$$
$$... - \sigma\beta c_{12}(S + (1 - c_1)R_1 + (1 - c_2)R_2)R_{12}$$

$$\dot{X} = \mu(1 - m - m_1 - m_2 - m_{12} - X) + (\tau_1 + \tau_2 + \gamma)S + (\tau_2 + \gamma)R_1 + ... \tag{1.3e}$$
$$... + \gamma R_{12} + (\tau_1 + \gamma)R_2 - \beta X(S + (1 - c_1)R_1 + (1 - c_2)R_2 + (1 - c_{12})R_{12}).$$

Table 1.3: The parameters used in the extension of the model of an ICU unit to include colonisation by multidrug resistant pathogens.

parameter	meaning
R_{12}	patients colonised by bacteria that are both antibiotic 1 and 2-resistant
m_{12}	patients enter hospital in states R_{12} at rates μ_{12}
c_{12}	fitness cost of multidrug resistance to bacteria

The interpretation of the parameters used in addition to the ones described in Table 1.2 are given in Table 1.3. Note that we also have to modify the treatment payoff in order to minimise both single-drug and multidrug-resistant infections in an observation interval of duration T:

$$\int_0^T (R_1(t) + R_2(t) + R_{12}(t))\, dt.$$

Chapter 2

A control-theoretic approach

Here we take tools from the theory of optimal control and re-examine the epidemiological models presented in the previous chapter, with the purpose of answering the following fundamental questions. What is the optimal drug deployment protocol? Is the optimal policy robust to modelling and parametric uncertainties? How easy is it to design and to implement? How effective is it?

In this chapter we will show that the optimal treatment protocol has a universal form that is model-independent and a classical engineering solution: a bang-bang control whose structure is independent of the underlying biology. Moreover, this solution has a clear epidemiological interpretation: to achieve optimality one should deploy a single antibiotic throughout the entire population, eventually replacing it with a different one. We will also show that through the design of feedback controllers it is possible to decide based on local patterns of resistance when it is best to switch antibiotics. Finally, we will demonstrate that antibiotic mixing protocols are only optimal in rare circumstances that are not robust to parametric uncertainties.

It is important to mention that we are not the first to propose that optimal control theory could be employed to probe either problems in drug delivery or the problem of antibiotic deployment. For example, [19] finds bang-bang solutions to the question of how to use principles of optimality to aid immunotherapy, while [59] applies control theory to the

chemotherapy of HIV. A discussion can be found in [86] of how to use an optimal control approach to decide how to best deploy antibiotics, but a claim is made that the failure of certain convexity requirements prevents the construction of analytic solutions of problems like the ones we study, a conclusion based on a discussion in [55, 90]. Of course, we agree with these studies that optimal control theory could be an aid of unprecedented power in the design and analysis of antibiotic deployment policies.

2.1 Parameter values for simulations: the importance of asymmetry

A crucial biological assumption is used in [13, 15] to simplify the modelling problem, namely *antibiotic symmetry*. This assumption is not benign. It is a mathematical degeneracy and we will use it to prove that antibiotic rotation is optimal whenever such a symmetry property is not present in a mathematical, epidemiological model of antibiotic use.

Throughout the thesis the term *parameter-initial condition set* (PICS) will be used for the set of epidemiological parameters and initial conditions defined within the models (1.1) and (1.2), note that each element of a PICS forms a pair that we shall write throughout as $(p, \boldsymbol{s}_0)$. The following important definition makes explicit the term *symmetric* as it is used in [13] and [15].

Definition 2.1.1 *If $(p, \boldsymbol{s}_0)$ denotes a PICS for (1.1), we say it is symmetric if*

$$r_a = r_b \text{ and } y_a(0) = y_b(0).$$

If $(p, \boldsymbol{s}_0)$ denote a PICS for (1.2), we say it is symmetric if

$$c_1 = c_2, m_1 = m_2 \text{ and } R_1(0) = R_2(0).$$

The two PICSs $(p^{(1)}, \boldsymbol{s}_0^{(1)})$ and $(p^{(2)}, \boldsymbol{s}_0^{(2)})$ so-defined are asymmetric in the sense of Definition 2.1.1, contrasting with the values chosen for numerical simulations in [13, 15] where symmetric values are used.

Mathematical models that have symmetric parameter values and initial conditions can be thought of as descriptions of antibiotic deployment problems in which the fundamental epidemiological properties of the drugs are identical. This means that there are equal fitness costs of antibiotic resistance to the pathogenic bacteria, equal transmission rates of those pathogens or equal prevalence of resistant phenotypes at the beginning of an observation period. However, while it is natural to support antibiotic symmetry on the grounds of numerical parsimony, we claim it is unlikely that two antibiotics will exert precisely the same selection pressures on bacterial pathogens. As a result we have chosen to use slightly different parameter sets for our illustrative simulations given later in the Section 2.6 from those found in [13, 15] in order to mimic the deployment of two antibiotics from distinct *functional groups* as defined, for example, in the sense of [116].

Furthermore, we make the claim that both models (1.1) and (1.2) must reflect this fundamental property on biological grounds too. For instance consider two antibiotics, *rifampicin* and *sorangicin A*, that have the same mode of action and bind to the same residue on their common target protein, inhibiting the synthesis of mRNA by binding to the β subunit of RNA polymerase. Rifampicin causes the bacterial cell to abort transcription at the elongation phase, as does sorangicin, albeit with slightly different abortive transcripts and the gene *rpoB* controls resistance mutations to both antibiotics. However, it is known [20] that mutations in *rpoB* conferring resistance to rifampicin need not confer resistance to sorangicin because of the greater flexibility of the sorangicin A molecule (also see [115]); thus functionally identical antibiotics may be different from an evolutionary perspective.

As a result we argue that we should seek to understand the structure of the optimal solutions for all parameter sets, whether symmetric or asymmetric, and in the remainder of this chapter we will discuss the mathematical reasons why the symmetric case is so special.

2.2 Statement of the control problem

A control problem is to determine an admissible control such that the response of the system is driven towards a given target in such a way that a certain optimality criterion is achieved, expressed by minimising some measure of cost or by maximising a given payoff functional. If there exists at least one successful control that satisfies the corresponding constrained optimisation problem, then we say that the system is controllable.

For example, the optimal control problem for (1.1) is to determine the protocol $f_a(t)$ that minimises the observed prevalence of resistance over a given time period of length T:

$$\text{Problem 1} \qquad \begin{cases} \min \int_0^T (y_w(t) + y_a(t) + y_b(t))\, dt & \text{subject to constraints} \\ 0 \leq f_a(t) \leq 1, \;\; f_b(t) = 1 - f_a(t) \text{ and equation } (1.1), \end{cases}$$

while the optimal treatment problem for (1.2) is to minimise the observed prevalence of antibiotic-resistant infections subject to treating at the maximum rate possible. We state this mathematically as follows:

$$\text{Problem 2} \qquad \begin{cases} \min \int_0^T (R_1(t) + R_2(t))\, dt & \text{subject to constraints} \\ 0 \leq \tau_1(t) \leq \tau_{\max}, \tau_2(t) = \tau_{\max} - \tau_1(t) \quad \text{and} \quad \text{equation } (1.2). \end{cases}$$

Remark 1 *In* Problem 1 *the units of the treatment payoff functional*

$$\int_0^T (y_w(t) + y_a(t) + y_b(t))\, dt$$

is the total number of patients infected over the period observed. In Problem 2 *the treatment payoff*

$$\int_0^T (R_1(t) + R_2(t))\, dt$$

must be multiplied by the total population size in the hospital (some fixed and unknown constant) in order to represent the total number of patients infected with antibiotic-resistant pathogens over the period observed. So, $\int_0^T (R_1(t) + R_2(t))\, dt/T$ *is the per unit time, mean fraction of patients infected with drug-resistant pathogens; it is unimportant whether*

or not we divide by T when minimising the treatment payoff as T is a fixed parameter.

In order to solve Problems 1 and 2, first note that the differential equations can both be written in the abstract form

$$\dot{\boldsymbol{s}} = f(\boldsymbol{s}, p) + A(t)\, g(p) \cdot \boldsymbol{s} + B(t)\, G(p) \cdot \boldsymbol{s}, \qquad \boldsymbol{s}(0) = \boldsymbol{s}_0 \in \mathbb{R}^k. \tag{2.1}$$

Equation (1.1) can be written in the form (2.1) as follows: first set $\boldsymbol{s} = (x, y_w, y_a, y_b)$ and then

$$\begin{aligned} f(\boldsymbol{s}, p) = (\lambda - dx - b(y_w + y_a + y_b)x + r_w y_w + r_a y_a + r_b y_b, ... \\ (bx - c - r_w)y_w, (bx - c - r_a)y_a, (bx - c - r_b)y_b) \end{aligned}$$

so that

$$g(p) = \begin{bmatrix} 0 & h(1-s) & 0 & h(1-s) \\ 0 & -h & 0 & 0 \\ 0 & hs & 0 & 0 \\ 0 & 0 & 0 & -h \end{bmatrix} \text{ and } G(p) = \begin{bmatrix} 0 & h(1-s) & h(1-s) & 0 \\ 0 & -h & 0 & 0 \\ 0 & 0 & -h & 0 \\ 0 & hs & 0 & 0 \end{bmatrix}.$$

For equation (1.2) we have $\boldsymbol{s} = (S, R_1, R_2, X)$ and

$$\begin{aligned} f(\boldsymbol{s}, p) \;=\; & (\mu(m - S) - \gamma S + \beta SX + \sigma\beta(c_1 R_1 + c_2 R_2)S, \\ & \mu(m_1 - R_1) - \gamma R_1 + \beta(1 - c_1)R_1 X - \sigma\beta(c_1 S + (c_1 - c_2)R_2)R_1, \\ & \mu(m_2 - R_2) - \gamma R_2 + \beta(1 - c_2)R_2 X - \sigma\beta(c_2 S + (c_2 - c_1)R_1)R_2, \\ & \mu(1 - m - m_1 - m_2 - X) + \gamma S + \gamma R_1 + ... \\ & \qquad ... + \gamma R_2 - \beta X(S + (1 - c_1)R_1 + (1 - c_2)R_2) \end{aligned}$$

with

$$g(p) = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 1 & 0 & 1 & 0 \end{bmatrix} \text{ and } G(p) = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \end{bmatrix}.$$

The fact that (1.1) and (1.2) can both be written in the form of (2.1) allows us to deduce properties of these two specific models by deducing properties from the more general and structural form of (2.1).

Now, equation (2.1) is a differential equation on a four-dimensional state-space Σ of non-negative vectors, so $\boldsymbol{s}(t) \in \Sigma$ for all t, where the parameter vector p lies in a space $\mathcal{J}$ of positive parameter values and so a PICS, $(p, \boldsymbol{s}_0)$ say, is an element of $\mathcal{J} \times \Sigma$. In Problem 1 we have $\boldsymbol{s} = (x, y_w, y_a, y_b)$ whereas in Problem 2 we write $\boldsymbol{s} = (S, R_1, R_2, X)$. The parameter-dependent linear maps $g(p)$ and $G(p)$ describe how the different rates of input of each antibiotic into the system drive the epidemiological dynamics of that system.

The optimality criteria in Problems 1 and 2 can now be written in an abstract form by defining a weight vector, call it $\boldsymbol{w}$, setting $A(t) + B(t) \equiv \mathsf{C}$, the latter a fixed constant, and then seeking a protocol $A(t) \in L^\infty(0, T)$ that achieves

$$\text{Problem A} \quad \begin{cases} \min \int_0^T \langle \boldsymbol{w}, \boldsymbol{s}(t) \rangle dt & \text{subject to constraints} \\ 0 \le A(t) \le \mathsf{C}, A(t) + B(t) \equiv \mathsf{C} & \text{and equation } (2.1); \end{cases}$$

the optimal protocol that solves Problem A will be denoted throughout by $A^*(t)$. Note that both Problems 1 and 2 have the same form as Problem A and so any statement made of Problem A regarding the structure of $A^*(t)$ has immediate consequences for both Problem 1 and Problem 2.

For each measurable control or deployment function A satisfying $0 \le A(t) \le \mathsf{C}$, the corresponding solution $\boldsymbol{s}_A$ obtained by solving the differential equation (3.2) yields a value of the functional

$$\mathcal{J}(A) := \int_0^T \langle \boldsymbol{w}, \boldsymbol{s}_A(t) \rangle dt$$

that will be denoted $\mathcal{J}(A)$ throughout the remainder and called the *treatment objective*. The

function of t, $\langle \boldsymbol{w}, \boldsymbol{s}_A(t) \rangle$, will be called the *running objective* associated with A. Moreover, for Problems 1 and 2 the weight vectors are

$$\boldsymbol{w} = (0, 1, 1, 1) \text{ and } \boldsymbol{w} = (0, 1, 1, 0),$$

respectively.

Implementing the must-treat constraint $A(t) + B(t) = \mathsf{C}$ in equation (2.1) yields

$$\begin{aligned} \dot{\boldsymbol{s}} &= f(\boldsymbol{s}, p) + A(t)\, g(p) \cdot \boldsymbol{s} + (\mathsf{C} - A(t))\, G(p) \cdot \boldsymbol{s}, \\ &= f(\boldsymbol{s}, p) + \mathsf{C} \cdot G(p) \cdot \boldsymbol{s} + A(t)(g(p) - G(p)) \cdot \boldsymbol{s}, \end{aligned} \tag{2.2}$$

and so we define, here and throughout,

$$\mathcal{G}(p) := g(p) - G(p) \text{ and } \mathcal{F}(\boldsymbol{s}, p) := f(\boldsymbol{s}, p) + \mathsf{C} \cdot G(p) \cdot \boldsymbol{s}.$$

Thus, if there is any parameter value p' for which $g(p') = G(p')$ then the set

$$\{(p', \boldsymbol{s}_0) : \boldsymbol{s}_0 \in \Sigma\}$$

must lie in the mixing PICS because the independence of equation (3.2) of A in this case renders the treatment objective identical for all deployment protocols. This is a trivial form of degeneracy that causes the mixing PICS $\mathcal{M}$ to be non-empty; we discuss less trivial examples below.

The Lagrangian of Problem A is

$$\mathcal{L}(\boldsymbol{s}, \boldsymbol{\lambda}, A) = \int_0^T \langle \boldsymbol{w}, \boldsymbol{s} \rangle + \langle \boldsymbol{\lambda}, -\dot{\boldsymbol{s}} + \mathcal{F}(\boldsymbol{s}, p) + A \cdot \mathcal{G}(p)\boldsymbol{s} \rangle \, dt,$$

and the Hamiltonian H is

$$H(\boldsymbol{s}, \boldsymbol{\lambda}, A) = \langle \boldsymbol{w}, \boldsymbol{s} \rangle + \langle \boldsymbol{\lambda}, \mathcal{F}(\boldsymbol{s}, p) \rangle + \langle \boldsymbol{\lambda}, \mathcal{G}(p)\boldsymbol{s} \rangle \cdot A,$$

finally, the adjoint variable $\boldsymbol{\lambda}$ satisfies the final-value problem

$$-\dot{\boldsymbol{\lambda}} = \boldsymbol{w} + (\mathcal{F}_{\boldsymbol{s}}(\boldsymbol{s}, p)^T + A \cdot \mathcal{G}(p)^T)\boldsymbol{\lambda}, \qquad \boldsymbol{\lambda}(T) = 0. \tag{2.3}$$

As is well-known, the *Hamiltonian* associated with (3.2-2.3) is maximised at all times along an optimal solution $(\boldsymbol{s}^*, \boldsymbol{\lambda}^*, A^*)$ of Problem A with respect to the control variable A:

$$H(\boldsymbol{s}^*(t), \boldsymbol{\lambda}^*(t), A^*(t)) = \max_{0 \le a \le \mathsf{C}} H(\boldsymbol{s}(t), \boldsymbol{\lambda}(t), a).$$

Now $\max\{H(\boldsymbol{s}^*(t), \boldsymbol{\lambda}^*(t), a) | 0 \le a \le \mathsf{C}\}$ occurs when $a = \mathsf{C}$ if $\langle \boldsymbol{\lambda}^*(t), \mathcal{G}(p)\boldsymbol{s}^*(t) \rangle > 0$ and when $a = 0$ if $\langle \boldsymbol{\lambda}^*(t), \mathcal{G}(p)\boldsymbol{s}^*(t) \rangle < 0$, if $\langle \boldsymbol{\lambda}^*(t), \mathcal{G}(p)\boldsymbol{s}(t) \rangle = 0$ then $A(t)$ is said to be **singular**. The solution of the optimal control problem Problem A is therefore a bang-bang function $A^*(t)$ taking *only* the values 0 and C unless t takes values in an interval where the switching function $(\boldsymbol{\lambda}^*(t), \mathcal{G}(p)\boldsymbol{s}^*(t))$ is zero:

$$A^*(t) = \begin{cases} \mathsf{C} & \text{if} \quad \langle \boldsymbol{\lambda}^*(t), \mathcal{G}(p)\boldsymbol{s}^*(t) \rangle > 0 \\ 0 & \text{if} \quad \langle \boldsymbol{\lambda}^*(t), \mathcal{G}(p)\boldsymbol{s}^*(t) \rangle < 0 \\ \text{something else} & \text{if} \quad \langle \boldsymbol{\lambda}^*(t), \mathcal{G}(p)\boldsymbol{s}^*(t) \rangle = 0. \end{cases} \tag{2.4}$$

Bang-bang controls correspond precisely to the antibiotic rotational protocols of Problem A and from the form of the optimal control A^* given in (2.4) we deduce that Problem A may only have a solution that is a mixing protocol when

$$\sigma(t) := \langle \boldsymbol{\lambda}^*(t), \mathcal{G}(p)\boldsymbol{s}^*(t) \rangle = 0 \tag{2.5}$$

for almost all t between 0 and T.

2.3 When is mixing optimal?

Now, we will use condition (2.5) to rule out mathematical models within Problem A for which mixing outperforms antibiotic rotation, by deducing the technical conditions on $\mathcal{F}$ and $\mathcal{G}$ under which there can be *no* solution of Problem A representing an antibiotic mixing protocol.

Theorem 2.3.1 *Suppose that $\omega^* \in (0, \mathsf{C})$ is a fixed constant and that*

$$\langle \boldsymbol{w}, \mathcal{G}(p)\boldsymbol{s}_0 \rangle \neq 0, \tag{2.6}$$

then there is a $\overline{T} > 0$ such that for no $T \in (0, \overline{T})$ is $A^*(t) \equiv \omega^*$ a mixing solution of Problem A. *However, if $\langle \boldsymbol{w}, \mathcal{G}(p)\boldsymbol{s}_0 \rangle = 0$ and either*

$$\langle \boldsymbol{w}, \mathcal{G}(p)(\mathcal{F}(\boldsymbol{s}_0, p) + \omega^* \mathcal{G}(p)\boldsymbol{s}_0) \rangle \neq 0, \qquad \textit{or} \tag{2.7a}$$

$$\langle \boldsymbol{w}, (\mathcal{F}_{\boldsymbol{s}}(\boldsymbol{s}_0, p) + \omega^* \mathcal{G}(p))\mathcal{G}(p)\boldsymbol{s}_0 \rangle \neq 0, \tag{2.7b}$$

then there is a $\overline{T} > 0$ such that for no $T \in (0, \overline{T})$ is the constant function $A^*(t) \equiv \omega^*$ a solution of Problem A.

Proof Begin by defining a new time-scale $\tau := t/T$ and re-writing the Euler-Lagrange equations of Problem A, namely (3.2-2.3), in the form

$$\dot{\boldsymbol{s}} = T(\mathcal{F}(\boldsymbol{s}, p) + A(t)(g(p) - G(p))), \qquad \boldsymbol{s}(0) = \boldsymbol{s}_0, \tag{2.8}$$

$$-\dot{\boldsymbol{\lambda}} = T(\boldsymbol{w} + (\mathcal{F}_{\boldsymbol{s}}(\boldsymbol{s}, p)^T + A \cdot \mathcal{G}(p)^T)\boldsymbol{\lambda}), \qquad \boldsymbol{\lambda}(1) = 0. \tag{2.9}$$

Now set $\boldsymbol{m} := \boldsymbol{\lambda}/T$ so that

$$-\dot{\boldsymbol{m}} = \boldsymbol{w} + T(\mathcal{F}_{\boldsymbol{s}}(\boldsymbol{s}, p)^T + A \cdot \mathcal{G}(p)^T)\boldsymbol{m}), \qquad \boldsymbol{m}(1) = 0. \tag{2.10}$$

To complete the proof we shall need the following auxiliary lemma.

Lemma 2.3.2 *Suppose that* $\boldsymbol{s}(t) = (s_1(t), ..., s_k(t))$ *is any continuous function defined on* $[0,1]$ *such that* $\boldsymbol{s}(0) = \boldsymbol{s}_0, m \leq s_i(t) \leq M$ *for all* $1 \leq i \leq k$ *and that* $A : \mathbb{R}^k \to L(\mathbb{R}^k)$ *is a continuous map. If* $\boldsymbol{w} \in \mathbb{R}^k$ *is any vector then the solution* $\boldsymbol{\lambda} \in C^1([0,1], \mathbb{R}^k)$ *of*

$$\dot{\boldsymbol{\lambda}}(t) = T(A(\boldsymbol{s}(t))\boldsymbol{\lambda} + \boldsymbol{w})$$

with $\boldsymbol{\lambda}(1) = \mathbf{0}$ *satisfies*

$$\|\boldsymbol{\lambda}(t)\|_\infty \leq \|\boldsymbol{w}\|_2(e^{T\rho} - 1)\rho^{-1},$$

where $\rho = \max\{\|A(\boldsymbol{s})\|_2 : m \leq \boldsymbol{s} \leq M\}$. *Hence* $\boldsymbol{m}(t) := \boldsymbol{\lambda}(t)/T$ *satisfies* $\|\boldsymbol{m}(t)\|_\infty \leq \|\boldsymbol{w}\|_2(e^{T\rho} - 1)/(T\rho)$.

Proof Let $\Phi(t)$ be the smooth, one-parameter family of matrices that satisfies

$$\dot{\Phi}(t) = T \cdot A(\boldsymbol{s}(t))\Phi(t), \qquad \Phi(0) = I.$$

If $\boldsymbol{u} \in \mathbb{R}^n$ is any vector then $|(A(\boldsymbol{s})\boldsymbol{u}, \boldsymbol{u})| \leq \|A(\boldsymbol{s})\|_2\|\boldsymbol{u}\|_2^2$ and so

$$\begin{aligned}\frac{d}{dt}\|\Phi(t)\boldsymbol{u}\|_2^2 &= 2T(A(\boldsymbol{s}(t))\Phi(t)\boldsymbol{u}, \Phi(t)\boldsymbol{u}) \leq \|A(\boldsymbol{s}(t))\|_2 \cdot 2T\|\Phi(t)\boldsymbol{u}\|_2^2 \\ &\leq \max_{m \leq \boldsymbol{s} \leq M} \|A(\boldsymbol{s})\|_2 \cdot 2T\|\Phi(t)\boldsymbol{u}\|_2^2\end{aligned}$$

and so $\|\Phi(t)\boldsymbol{u}\|_2 \leq e^{\rho Tt}\|\Phi(0)\boldsymbol{u}\|_2 = e^{\rho Tt}\|\boldsymbol{u}\|_2$ from where $\|\Phi(t)\|_2 \leq e^{\rho Tt}$. Now $\boldsymbol{\lambda}(t) = T\int_1^t \Phi(t - t')\boldsymbol{w}dt'$ and so

$$\|\boldsymbol{\lambda}(t)\|_\infty \leq \|\boldsymbol{\lambda}(t)\|_2 \leq T\int_0^1 \|\Phi(t - t')\|_2\|\boldsymbol{w}\|_2 dt' \leq T\|\boldsymbol{w}\|_2 \int_0^1 e^{T\rho(t-t')}dt'$$

and the result follows. □

The proof of Theorem 2.3.1 follows immediately below and to reduce notational clutter we assume without loss of generality that the constant C defined in Problem A equals one.

Suppose that a parameter value T, that we label T^*, exists for which Problem A has optimal control $A^* \equiv \omega^* \in (0,1)$ with treatment objective $\mathcal{J}(\omega^*)$ and so we may suppose $(\boldsymbol{s}^*, \boldsymbol{m}^*)$

is a solution of the re-scaled Euler-Lagrange equations (2.8-2.9). If we now define

$$\boldsymbol{S}^*(t) := \boldsymbol{s}^*(t) - \boldsymbol{s}_0,$$

we may re-write the Euler-Lagrange equations associated with Problem A as a nonlinear operator equation that we denote $\mathcal{E}(\boldsymbol{S}, \boldsymbol{m}, T, \omega) = 0$, where

$$\mathcal{E}(\boldsymbol{S}, \boldsymbol{m}, T, \omega) := \begin{pmatrix} -\dot{\boldsymbol{S}} + T(\mathcal{F}(\boldsymbol{s}_0 + \boldsymbol{S}, p) + \omega \cdot \mathcal{G}(p)(\boldsymbol{s}_0 + \boldsymbol{S})) \\ \dot{\boldsymbol{m}} + T(\mathcal{F}_{\boldsymbol{s}}^T(\boldsymbol{s}_0 + \boldsymbol{S}, p) + \omega \cdot \mathcal{G}(p)^T)\boldsymbol{m} + \boldsymbol{w} \end{pmatrix}.$$

Hence $\mathcal{E}(\boldsymbol{S}^*, \boldsymbol{m}^*, T^*, \omega^*) = \boldsymbol{0}$ for $(\boldsymbol{S}^*, \boldsymbol{m}^*) \in U \times V$ where

$$U := \{\boldsymbol{S} \in C^1([0,1], \mathbb{R}^k) \ : \ \boldsymbol{S}(0) = \boldsymbol{0}\} \ \text{ and } \ V := \{\boldsymbol{m} \in C^1([0,1], \mathbb{R}^k) \ : \ \boldsymbol{m}(1) = \boldsymbol{0}\}$$

are Banach spaces when endowed with standard C^1 norms and

$$\mathcal{E} : U \times V \times \mathbb{R} \to C^0([0,1], \mathbb{R}^k) \times C^0([0,1], \mathbb{R}^k)$$

is an everywhere continuously Fréchet differentiable nonlinear mapping.

Define the following isomorphism of Banach spaces $D : U \times V \to C^0([0,1], \mathbb{R}^k) \times C^0([0,1], \mathbb{R}^k)$ given by the differential operator

$$D(\boldsymbol{S}, \boldsymbol{m}) = \frac{d}{dt}(-\boldsymbol{S}, \boldsymbol{m}).$$

Being an isomorphism, D is a linear operator of Fredholm index 0 but then

$$\partial_{\boldsymbol{S},\boldsymbol{m}}\mathcal{E} : U \times V \to C^0([0,1], \mathbb{R}^k) \times C^0([0,1], \mathbb{R}^k)$$

is also a linear, Fredholm mapping of index-0 because it is a compact perturbation of D. Thus, $\partial_{\boldsymbol{S},\boldsymbol{m}}\mathcal{E}(\boldsymbol{S}^*, \boldsymbol{m}^*, T^*, \omega^*)$ is an isomorphism if and only if it is injective and so, seeking a null-space of the linear operator $\partial_{\boldsymbol{S},\boldsymbol{m}}\mathcal{E}(\boldsymbol{S}^*, \boldsymbol{m}^*, T^*, \omega^*)$ we must solve the linear differential equation

$$\partial_{\boldsymbol{S},\boldsymbol{m}}\mathcal{E}(\boldsymbol{S}^*, \boldsymbol{m}^*, T^*, \omega^*)[X, Y] = [\boldsymbol{0}, \boldsymbol{0}]$$

for X and Y.

Computing the form of the derivative matrix $\partial_{\boldsymbol{S},\boldsymbol{m}}\mathcal{E}(\boldsymbol{S}^*,\boldsymbol{m}^*,T^*,\omega^*)[X,Y]$ gives

$$\partial_{\boldsymbol{S},\boldsymbol{m}}\mathcal{E}(\boldsymbol{S}^*,\boldsymbol{m}^*,T^*,\omega^*)[X,Y]=\begin{pmatrix} -\dot{X}+T^*(\mathcal{F}_{\boldsymbol{s}}+\omega^*\mathcal{G})X \\ \dot{Y}+T^*((\mathcal{F}_{\boldsymbol{s}}^T+\omega^*\mathcal{G}^T)Y+... \\ \qquad\qquad ...+\mathcal{F}_{\boldsymbol{ss}}(\boldsymbol{s}_0+\boldsymbol{S}^*)^T[\boldsymbol{m}^*,X] \end{pmatrix},$$

where the entries in this matrix are evaluated at $(\boldsymbol{S}^*,\boldsymbol{m}^*,T^*,\omega^*)$ and we deduce

$$-\dot{X}+T^*(\mathcal{F}_{\boldsymbol{s}}+\omega^*\mathcal{G})X=\mathbf{0}$$

for some function $X\in U$. Standard uniqueness theorems for non-autonomous ordinary differential equations now yield $X(t)\equiv 0$ as $X(0)=0$, but then $Y=\mathbf{0}$ immediately follows and so $\partial_{\boldsymbol{S},\boldsymbol{m}}\mathcal{E}(\boldsymbol{S}^*,\boldsymbol{m}^*,T^*,\omega^*)$, having been shown to be injective, is an isomorphism for all $(\boldsymbol{S}^*,\boldsymbol{m}^*,T^*,\omega^*)\in U\times V\times\mathbb{R}^2$ with $\mathcal{E}(\boldsymbol{S}^*,\boldsymbol{m}^*,T^*,\omega^*)=\mathbf{0},\omega^*\in[0,1]$ and $T^*>0$. As a consequence, we can apply the implicit function theorem to solve $\mathcal{E}(\boldsymbol{S},\boldsymbol{m},T,\omega)=\mathbf{0}$ near to any given solution $(\boldsymbol{S}^*,\boldsymbol{m}^*,T^*,\omega^*)$ to provide a locally unique, two-dimensional solution surface on which one can write $\boldsymbol{S}=\boldsymbol{S}(T,\omega)\in U,\boldsymbol{m}=\boldsymbol{m}(T,\omega)\in V$ such that

$$\mathcal{E}(\boldsymbol{S}(T,\omega),\boldsymbol{m}(T,\omega),T,\omega)\equiv\mathbf{0},$$

where $\boldsymbol{S}(T^*,\omega^*)=\boldsymbol{S}^*$ and $\boldsymbol{m}(T^*,\omega^*)=\boldsymbol{m}^*$. Denote the common domain of definition of $\boldsymbol{S}(T,\omega)$ and $\boldsymbol{m}(T,\omega)$ as provided above by the implicit function theorem by Ω' and then define Ω to be $\Omega'\cap(0,T^*]\times(0,1)$.

We shall call the two-parameter function $(\boldsymbol{S}(T,\omega),\boldsymbol{m}(T,\omega))$ the *mixing surface* of Problem A for it contains every possible small-T mixing solution of this optimal control problem. We can extend the domain of this surface, currently Ω, to the entire rectangular domain $[0,T^*]\times[0,1]$ using Lemma 2.3.2 and the implicit function theorem, but we shall only sketch the argument as follows.

First fix $\omega = \omega^*$. If

$$\inf\{\overline{T} : (\boldsymbol{S}(T,\omega^*), \boldsymbol{m}(T,\omega^*), T, \omega^*) : (\overline{T}, T^*] \to U \times V \times [0,\infty) \text{ such that } \mathcal{E}(\boldsymbol{S}(T,\omega^*), \boldsymbol{m}(T,\omega^*), T, \omega^*) = \mathbf{0}\} > 0$$

then we can find a sequence $(\boldsymbol{S}_n^*, \boldsymbol{m}_n^*, T_n^*, \omega^*) \in U \times V \times [0, T^*]$ to form this infimum. By Lemma 2.3.2 this sequence is C^0-bounded, but from the form of $\mathcal{E}(\boldsymbol{S}_n^*, \boldsymbol{m}_n^*, T_n^*, \omega^*) = \mathbf{0}$ we can bootstrap to readily obtain C^2 bounds on this same sequence and so extract C^1-convergent subsequences that we do not relabel that converge to a solution of $\mathcal{E}(\boldsymbol{S}, \boldsymbol{m}, T, \omega^*) = \mathbf{0}$. We can then apply the implicit function theorem using the fact that $\partial_{\boldsymbol{S},\boldsymbol{m}}\mathcal{E}(\boldsymbol{S}, \boldsymbol{m}, T, \omega^*)$ is an isomorphism at this point to further extend the definition of $(\boldsymbol{S}(T), \boldsymbol{m}(T), T, \omega^*)$ to a lower value of T. This is a contradiction which ensures that

$$\inf\{\overline{T} : (\boldsymbol{S}(T,\omega^*), \boldsymbol{m}(T,\omega^*), T, \omega^*) : (\overline{T}, T^*] \to U \times V \times [0,\infty) \text{ such that } \mathcal{E}(\boldsymbol{S}(T,\omega^*), \boldsymbol{m}(T,\omega^*), T, \omega^*) = \mathbf{0}\} = 0.$$

With a further application of the implicit function theorem at each point $(T, \omega^*) \in [0, T^*] \times \{\omega^*\}$, we can extend the domain of definition of the mixing surface in an entirely analogous manner to a rectangular strip $[0, T^*] \times (\omega^* - \eta, \omega^* + \eta)$, for some $\eta > 0$, that contains the line $[0, T^*] \times \{\omega^*\}$. Lemma 2.3.2 can then be used to bootstrap and so continuously extend the domain of definition of the mixing surface to the strip $[0, T^*] \times [\omega^* - \eta, \omega^* + \eta]$. Further applications of this bootstrapping process and the implicit function theorem then allow one to extend this domain from a thin strip to the entire rectangle $[0, T^*] \times [0, 1]$.

Although the mixing surface is now defined on $[0, T^*] \times [0, 1]$, because mixing solutions must be totally singular in the construction of the optimal control (2.4) this surface only contains *mixing solutions* of Problem A when the switching function σ defined in (2.5) equals zero as a function in $C^0[0, 1]$ when evaluated on that surface. In other words,

$$\sigma(T,\omega)(t) := \langle \mathcal{G}(p)(\boldsymbol{s}_0 + \boldsymbol{S}(T,\omega)(t)), \boldsymbol{m}(T,\omega)(t)\rangle = 0$$

must be satisfied for all t between 0 and 1 in order for the mixing protocol ω to be a solution of Problem A.

Our goal now is to use the infinite-dimensional version of Taylor's theorem to determine conditions that must be satisfied by $\mathcal{F}$ and $\mathcal{G}$ under the assumption that mixing is optimal. From this working assumption, the optimal control of Problem A is $A^*(t) \equiv \omega^*$ identically in t which is a constant and so smooth function. Accordingly we can apply the infinite-dimensional version of Taylor's theorem and write, for $0 \leq t \leq 1$ and fixed $\omega > 0$,

$$\sigma(T,\omega)(t) = \sigma(0,\omega)(t) + T\partial_T\sigma(0,\omega)(t) + O(T^2),$$

where the $O(T^2)$ term here is measured in the C^0-norm. Solving $\mathcal{E}(\boldsymbol{S},\boldsymbol{m},T,\omega) = (0,0)$ when $T = 0$ and ω is arbitrary yields the unique solution

$$\boldsymbol{S}(t) = \mathbf{0}, \;\; \boldsymbol{m}(t) = (1-t)\boldsymbol{w}.$$

Continuing with the application of the Taylor's theorem and expanding the solution locus of $\mathcal{E}(\boldsymbol{S},\boldsymbol{m},T,\omega) = (0,0)$ locally as a Taylor series, we therefore obtain

$$\boldsymbol{S}(T,\omega)(t) = \mathbf{0} + T\partial_T\boldsymbol{S}(0,\omega)(t) + O(T^2), \;\; \boldsymbol{m}(t) = (1-t)\boldsymbol{w} + T\partial_T\boldsymbol{m}(0,\omega)(t) + O(T^2).$$

Let us now compute the T-derivative $\partial_T\boldsymbol{S}(T,\omega)(t)$ that we denote by $\boldsymbol{S}_T \in U$, for the derivative $\partial_T\boldsymbol{m}(T,\omega)(t)$ we shall write $\boldsymbol{m}_T \in V$. On differentiating the equation $\mathcal{E}(\boldsymbol{S},\boldsymbol{m},T,\omega) = \mathbf{0}$ with respect to T we find

$$\frac{d}{dt}\boldsymbol{S}_T = \mathcal{F}(\boldsymbol{s}_0,p) + \omega\cdot\mathcal{G}(p)\boldsymbol{s}_0 \tag{2.11a}$$

$$-\frac{d}{dt}\boldsymbol{m}_T = (\mathcal{F}_{\boldsymbol{s}}(\boldsymbol{s}_0,p)^T + \omega\cdot\mathcal{G}(p)^T)\boldsymbol{m} \tag{2.11b}$$

where $\boldsymbol{m}(t) = (1-t)\boldsymbol{w}$. Solving (2.11 a-b) and incorporating boundary conditions we obtain, for $0 \leq t \leq 1$ and at $T = 0$,

$$\boldsymbol{S}_T(t) = t(\mathcal{F}(\boldsymbol{s}_0,p) + \omega\cdot\mathcal{G}(p)\boldsymbol{s}_0), \qquad \boldsymbol{m}_T(t) = \frac{1}{2}(1-t)^2(\mathcal{F}_{\boldsymbol{s}}(\boldsymbol{s}_0,p)^T + \omega\cdot\mathcal{G}(p)^T)\boldsymbol{w}.$$

But the following expression for the switching function σ is identically zero in ω, T and t:

$$\begin{aligned}
\mathbf{0} = \sigma(T,\omega)(t) &= \langle \mathcal{G}(p)(\boldsymbol{S}(T,\omega)(t) + \boldsymbol{s}_0), \boldsymbol{m}(T,\omega)(t)\rangle \\
&= \langle \mathcal{G}(p)\boldsymbol{s}_0 + \mathcal{G}(p)\boldsymbol{S}(T,\omega)(t), (1-t)\boldsymbol{w} + T\boldsymbol{m}_T(0,\omega)(t) + O(T^2)\rangle \\
&= \langle \mathcal{G}(p)\boldsymbol{s}_0 + \mathcal{G}(p)(\boldsymbol{S}(0,\omega)(t) + T\boldsymbol{S}_T(0,\omega)(t) + O(T^2)\rangle, \\
&\qquad\qquad (1-t)\boldsymbol{w} + T\boldsymbol{m}_T(0,\omega)(t) + O(T^2)\rangle \\
&= \langle \mathcal{G}(p)\boldsymbol{s}_0 + T\mathcal{G}(p)(\boldsymbol{S}_T(0,\omega)(t) + O(T)), \\
&\qquad\qquad (1-t)\boldsymbol{w} + T\boldsymbol{m}_T(0,\omega)(t) + O(T^2)\rangle \\
&= (1-t)\overbrace{\langle \mathcal{G}(p)\boldsymbol{s}_0, \boldsymbol{w}\rangle}^{\text{compare with (2.6)}} + T(1-t)\langle \mathcal{G}(p)\boldsymbol{S}_T(0,\omega)(t), \boldsymbol{w}\rangle + ... \\
&\qquad\qquad ... + T\langle \mathcal{G}(p)\boldsymbol{s}_0, \boldsymbol{m}_T(0,\omega)(t)\rangle + O(T^2)
\end{aligned} \tag{2.12}$$

In order for the mixing constant ω^* to be the optimal solution of Problem A from the $O(1)$ terms in (2.12) we require $\langle \mathcal{G}(p)\boldsymbol{s}_0, \boldsymbol{w}\rangle = 0$, but the $O(T)$ terms must also be identically zero in t. Hence, the quadratic expression in t

$$t(1-t)\langle \mathcal{G}(p)[(\mathcal{F}(\boldsymbol{s}_0,p) + \omega\cdot\mathcal{G}(p)\boldsymbol{s}_0)], \boldsymbol{w}\rangle + \frac{1}{2}(1-t)^2\langle \mathcal{G}(p)\boldsymbol{s}_0, (\mathcal{F}_{\boldsymbol{s}}(\boldsymbol{s}_0,p)^T + \omega\cdot\mathcal{G}(p)^T)\boldsymbol{w}\rangle$$

must be zero for all $t \in [0,1]$ and all (T,ω) in the domain of σ, concluding the proof. □

Theorem 2.3.1 is a negative result in the sense that it does not help us find solutions of Problem A, but it can be used to tell us when antibiotic mixing is *not* a solution of Problem 1 and Problem 2 in concrete cases. In particular, we have the following two corollaries which state that Problem 1 and Problem 2 have optimal controls that are mixing protocols *only* when their respective sets of parameters and initial conditions (PICSs) are symmetric.

Corollary 1 *Suppose that system parameters (given by the vector p) and initial conditions (given by the vector $\boldsymbol{s}_0 = (x(0), y_w(0), y_a(0), y_b(0))$) are non-negative in (1.1) with $h > 0$ and suppose also that* Problem 1 *has an optimal mixing treatment $f_a^*(t)$ that we denote by the constant $\omega^* \in (0,1)$. If $y_w(0) > 0$ and $h > 0$, then the PICS $(p, \boldsymbol{s}_0)$ is necessarily*

symmetric:

$$\omega^* = \frac{1}{2}, r_a = r_b \ \textit{and} \ \ y_a(0) = y_b(0), \tag{2.13}$$

and so $y_a(t) = y_b(t)$ *for all* $t \geq 0$.

Proof On setting $\boldsymbol{w} = (0,1,1,1), \boldsymbol{s}_0 = (x(0), y_w(0), y_a(0), y_b(0))$ and using the functions $\mathcal{F}$ and $\mathcal{G}$ to represent the system (1.1), we find

$$\mathcal{G}(p) = \begin{bmatrix} 0 & 0 & -h(1-s) & h(1-s) \\ 0 & 0 & 0 & 0 \\ 0 & hs & h & 0 \\ 0 & -hs & 0 & -h \end{bmatrix}.$$

Applying condition (2.6) and using

$$\mathcal{G}(p)\boldsymbol{s}_0 = ((y_a(0) - y_b(0))h(1-s), 0, hsy_w(0) + hy_a(0), -hsy_w(0) - hy_b(0))$$

we obtain

$$\begin{aligned} (\boldsymbol{w}, \mathcal{G}(p)\boldsymbol{s}_0) &= (0,1,1,1) \cdot ((y_a(0) - y_b(0))h(1-s), 0, hsy_w(0) + hy_a(0), ... \\ &\qquad ..., -hsy_w(0) - hy_b(0)) \qquad (2.14) \\ &= h(y_a(0) - y_b(0)) = 0 \qquad (2.15) \end{aligned}$$

But then

$$h(y_a(0) + hsy_w(0))(r_b - r_a + h(2\alpha - 1)) = 0 \tag{2.16}$$

and

$$h(y_a(0)(r_b - r_a) + h(2\alpha - 1)(y_a(0) + hsy_w(0))) = 0 \tag{2.17}$$

follow from conditions (2.7 a-b) of Theorem 2.3.1. Assuming $h > 0$ and $y_w(0) > 0$, the first part of the corollary (2.13) follows on solving the three algebraic relations (2.15), (2.16) and (2.17).

The last part of the statement of this corollary follows by noting that if $\delta(t) := y_a(t) - y_b(t)$, along solutions of (1.1) the function δ satisfies $\delta(0) = 0$ and $\dot{\delta}(t) \equiv 0$ if the restriction (2.13) is imposed from where we deduce that δ is identically zero as required. □

The following shows that a similar statement can be made for Problem 2.

Corollary 2 *Suppose that system parameters (given by the vector p) and all initial conditions (given by the vector $(S(0), R_1(0), R_2(0), X(0)$) are non-negative in (1.2) and suppose* Problem 2 *has a mixing optimal control τ_1^* that we denote by the constant $\omega^* \in (0, \tau_{\max})$, then*

$$R_1(0) = R_2(0),\ c_2 = c_1 + \frac{\tau_{\max} - 2\omega^*}{\beta(X(0) + \sigma S(0))},\ m_1 = m_2 + \frac{2(2\omega^* - \tau_{\max})\sigma R_1(0)^2}{\lambda(X(0) + \sigma S(0))}. \tag{2.18}$$

As a result, if the 50-50 mixing protocol $\omega^* = \frac{\tau_{\max}}{2}$ is optimal then $c_1 = c_2$ and $m_1 = m_2$, therefore $R_1(t) = R_2(t)$ for all $t \geq 0$.

Proof On setting $\boldsymbol{w} = (0, 1, 1, 0), \boldsymbol{s}_0 = (S(0), R_1(0), R_2(0), X(0))$ and using the functions $\mathcal{F}$ and $\mathcal{G}$ to represent the system (1.2), we find that

$$\mathcal{G}(p) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & -1 & 1 & 0 \end{bmatrix}.$$

Using the fact that $\mathcal{G}(p)\boldsymbol{s}_0 = (0, R_1(0), -R_2(0), R_2(0) - R_1(0))$ we obtain

$$\langle \mathcal{G}(p)\boldsymbol{s}_0, \boldsymbol{w} \rangle = R_1(0) - R_2(0) \tag{2.19}$$

and the first part of (2.18) follows from condition (2.6). The remaining two conditions of Theorem 2.3.1, (2.7 a) and (2.7 b), yield the following algebraic relations for elements

within the mixing PICS:

$$\begin{aligned} \mu\, m_1 - R_1\, \tau - \beta\, R_1\, X c_1 - \sigma\, \beta\, S c_1\, R_1 - 2\,\sigma\,\beta\, {R_1}^2 c_1 + 2\,\sigma\,\beta\, {R_1}^2 c_2 \;\; + \ldots \qquad (2.20) \\ \ldots + 2\,\alpha\, R_1 - \mu\, m_2 + \beta\, R_1\, X c_2 + \sigma\,\beta\, S c_2\, R_1 \;\; = \;\; 0 \end{aligned}$$

and

$$-\, R_1\, \tau - \beta\, R_1\, X c_1 - \sigma\,\beta\, S c_1\, R_1 + 2\,\alpha\, R_1 + \beta\, R_1\, X c_2 + \sigma\,\beta\, S c_2\, R_1 \;\; = \;\; 0, \qquad (2.21)$$

where the initial condition of each variable, S, R_1, R_2 and X, has been omitted for clarity, so that S denotes $S(0)$ and similarly for the other variables. On solving the relations (2.20) and (2.21) for (m_1, c_2) in terms of all the other variables, equation (2.18) results.

Finally, if $\omega^* = \frac{\tau_{\max}}{2}$ is optimal and we define the function $\delta(t) := R_1(t) - R_2(t)$, δ can be shown to satisfy $\delta(0) = 0$ and $\dot{\delta}(t) \equiv 0$ along solutions of (1.2) when one imposes the restriction that $c_1 = c_2$ and $m_1 = m_2$ from (2.18). The result now follows. □

Corollaries 1 and 2 represent analogous statements in terms of Problems 1 and 2 that may be summarised as follows: we do not yet know whether antibiotic mixing protocols are optimal for Problems 1 and 2, but if mixing is optimal for one of these models at some parameter value, the parameters and initial conditions within that model *must be symmetric* in the sense of Definition 2.1.1. These two results form the essence of our argument, ensuring as they do that many biologically interesting parameter values exist for which antibiotic mixing is not the optimal protocol. Indeed, these corollaries show that mixing may only be optimal in mathematically rare cases.

2.4 Optimal protocols: bang-bang controls

The results of the previous section are entirely negative and give no clue as to what the optimal deployment protocols might actually be for a given mathematical model. So, we now apply standard control-theoretic results to establish the epidemiological result that alternating protocols are optimal for Problem A, or at least 'ϵ-suboptimal' in a sense described below.

The set of admissible controls $\mathcal{U}$ for Problem A is the set of measurable functions taking values almost everywhere between 0 and C:

$$\mathcal{U} = \{\phi \in L^\infty[0,T] : 0 \leq \phi(t) \leq \mathsf{C} \quad \text{for almost all } t \in [0,T]\},$$

we are interested in conditions under which a solution of Problem A exists and lies in $\mathcal{U}$. The set of bang-bang functions $\mathcal{B}$ is contained within $\mathcal{U}$ and is defined by

$$\mathcal{B} = \{\phi \in \mathcal{U} : \exists \text{ partition } 0 = t_1 < ... < t_n = T : \ \phi(t) \in \{0, \mathsf{C}\} \ \forall\, t \in (t_k, t_{k+1})\}\,.$$

It is important to note that bang-bang functions $\mathcal{B}$ exactly describe the rotational protocols of equation (3.2) because the range of a function $\phi \in \mathcal{B}$ can only contain the two values 0 and C. In terms of Problem A, if $A(t) = \phi(t)$ and $B(t) = \mathsf{C} - \phi(t)$ then A and B represent a rotational protocol that is completely described by ϕ.

The following basic existence theorem tells us that an optimal control exists for *Problem A* provided (3.2) has a natural control-independent, point-dissipative bound. More importantly, it shows that the optimal deployment protocol can be approximated arbitrarily closely in a suitable sense by functions that rotate between the two antibiotics.

Theorem 2.4.1 *Suppose that there is a finite constant C depending on C, p, T and $\boldsymbol{s}_0$ such that for any function $A \in \mathcal{U}$, the solution $\boldsymbol{s}$ of (3.2) with $\boldsymbol{s}(0) = \boldsymbol{s}_0$ satisfies, for any norm $\|\cdot\|$,*

$$\sup_{0 \leq t \leq T} \|\boldsymbol{s}(t)\| \leq C(\mathsf{C}, p, T, \boldsymbol{s}_0). \tag{2.22}$$

Then Problem A *has at least one solution $A^* \in \mathcal{U}$ with corresponding state response $\boldsymbol{s}^*$ which satisfies (3.2) with $A = A^*$. For each $\epsilon > 0$ there is a function $A_\epsilon \in \mathcal{B}$ such that if $\boldsymbol{s}_\epsilon(t)$ is obtained by setting $A = A_\epsilon$ in (3.2), then A_ϵ is ϵ-suboptimal in the sense that*

$$\int_0^T \langle \boldsymbol{w}, \boldsymbol{s}_\epsilon(t) \rangle dt < \int_0^T \langle \boldsymbol{w}, \boldsymbol{s}^*(t) \rangle dt + \epsilon.$$

Proof Suppose that the sequence $(\boldsymbol{s}_n, A_n)$ provides the infimum

$$\mathcal{J}^* := \inf \{ \mathcal{J}(A) : A \in \mathcal{U} \},$$

then we may assume that there is an $A_{\mathsf{inf}} \in \mathcal{U}$ such that $A_n \overset{*}{\rightharpoonup} A_{\mathsf{inf}}$ in $L^\infty(0,T)$ as $n \to \infty$ because $\mathcal{U}$ is compact with respect to the weak* topology on L^∞. Without the loss of any generality, let us shift the initial datum to zero in equation (3.2) by assuming that $\boldsymbol{s}_n$ satisfies

$$\dot{\boldsymbol{s}}_n = \mathcal{F}(\boldsymbol{s}_0 + \boldsymbol{s}_n, p) + A_n \cdot \mathcal{G}(p)(\boldsymbol{s}_0 + \boldsymbol{s}_n), \qquad \boldsymbol{s}_n(0) = 0,$$

instead of (3.2). We obtain the bound $\left\| \frac{d}{dt} \boldsymbol{s}_n \right\|_\infty \leq \|\mathcal{F}(\boldsymbol{s}_0 + \boldsymbol{s}_n)\|_\infty + \mathsf{C} \|\mathcal{G}(p)\|_1 \|\boldsymbol{s}_0 + \boldsymbol{s}_n\|_\infty$, but $\|\boldsymbol{s}_n\| \leq C(\mathsf{C}, p, T, \boldsymbol{s}_0)$ and as all finite-dimensional norms are equivalent it follows that the sequence $(\boldsymbol{s}_n) \subset W_0^{1,\infty}((0,T), \mathbb{R}^k)$ is bounded (the space $W_0^{1,\infty}((0,T), \mathbb{R}^k)$ appropriately incorporates the zero boundary condition at $t = 0$). As a result $(\boldsymbol{s}_n)$ has a weak* convergent subsequence that we do not relabel, converging to $\boldsymbol{s}_{\mathsf{inf}} \in W_0^{1,\infty}((0,T), \mathbb{R}^k)$. As the nonlinear mapping $\mathcal{N} : W_0^{1,\infty}(0,T) \times L^\infty(0,T) \to L^\infty(0,T)$ given by

$$\mathcal{N}(\boldsymbol{s}, A) = -\frac{d}{dt}\boldsymbol{s} + \mathcal{F}(\boldsymbol{s}_0 + \boldsymbol{s}, p) + A \cdot \mathcal{G}(p)\boldsymbol{s}$$

is continuous with respect to weak* convergence in $W_0^{1,\infty}(0,T) \times L^\infty(0,T)$, we see that the limiting pair $(\boldsymbol{s}_0 + \boldsymbol{s}_{\mathsf{inf}}, A_{\mathsf{inf}})$ satisfies (3.2), that is $\mathcal{N}(\boldsymbol{s}_{\mathsf{inf}}, A_{\mathsf{inf}}) = 0$ and the result follows on setting $A^* = A_{\mathsf{inf}}$. □

2.5 The optimal mixing protocol

As pointed out in Appendix B3 of [15], the idea of an optimal mixing protocol is meaningful in the context of asymmetric antibiotic deployment problems whereby asymmetric PICS values are used. In such a case, the optimal mixing protocol has to be adjusted from the 50-50 value of $\omega = 1/2$ to account for their different evolutionary and epidemiolgical properties.

So, let $\boldsymbol{s}_\omega(t)$ be the solution of the differential equation

$$\dot{\boldsymbol{s}} = f(\boldsymbol{s}, p) + \mathsf{C} \cdot G(p)\boldsymbol{s} + \omega(g(p) - G(p)) \cdot \boldsymbol{s}, \qquad \boldsymbol{s}(0) = \boldsymbol{s}_0$$

which sees the constant deployment of two antibiotics at some rate $\omega \in [0, C]$. The *optimal mixing protocol* for equation (3.2) is found by solving a one-dimensional optimisation problem which asks for the single value ω between 0 and C, denoted ω^*, for which the treatment objective

$$\mathcal{J}(\omega) := \int_0^T \langle \boldsymbol{w}, \boldsymbol{s}_\omega(t) \rangle dt$$

is minimal. It is clear that the optimal mixing protocol is suboptimal in the context of (3.2) because

$$\mathcal{J}(\omega^*) = \min_{\substack{0 \le \omega \le \mathsf{C} \\ \omega \text{ constant}}} \mathcal{J}(\omega) \geq \min_{\substack{0 \le A(t) \le \mathsf{C} \\ A \text{ measurable}}} \mathcal{J}(A) = \mathcal{J}(A^*), \tag{2.23}$$

by definition. Note that we have already proven in Corollaries 1 and 2 of the previous section that equality is possible in (2.23) for Problem 1 and Problem 2 *only* when the parameters and initial conditions used within those problems are symmetric.

Theorem 2.4.1 can be applied to Problems 1 and 2 to provide the main mathematical result presented here as a corollary.

Corollary 3 Problems 1 *and* 2 *have optimal controls* $f_a^*(t) \in L^\infty(0, T)$ *and* $A^*(t) \in L^\infty(0, T)$ *respectively. If their respective PICSs are asymmetric then there are infinitely many antibiotic rotation protocols that outperform antibiotic mixing in terms of the performance measure* $\mathcal{J}(A)$.

Proof From Theorem 2.4.1 we only have to establish the existence of a dissipative bound of the form (2.22) for equations (1.1) and (1.2), the result then follows from the second part of Theorem 2.4.1.

1. From equation (1.1) define the strictly positive vector $\boldsymbol{v} = (d, c, c + hsf_b, c + hsf_a)$ and $v_{\min} = \min\{\boldsymbol{v}\}$ which is either c or d if $hs > 0$ and $f_a + f_b = 1$.

   Define the state vector $\boldsymbol{s} = (X, y_w, y_a, y_b)$, the vector of initial conditions $\boldsymbol{s}_0 = (X(0), y_w(0), y_a(0), y_b(0))$ and the vector $\mathbf{1} = (1, 1, 1, 1)$. Also define the parameter vector $p = (\lambda, d, c, h, r_w, s, r_a, r_b, b)$ for completeness.

   Equation (1.1) is point dissipative in the sense that if $f_a \in L^\infty(0, T)$ is any measurable function with $0 \leq f_a(t) \leq 1$, $f_b(t) = 1 - f_a(t)$ and $hs > 0$ then

$$\langle \mathbf{1}, \boldsymbol{s}(t) \rangle \leq \frac{\lambda}{v_{\min}} + e^{-v_{\min} t} \left(\langle \mathbf{1}, \boldsymbol{s}_0 \rangle - \frac{\lambda}{v_{\min}} \right)$$

   for all $t \geq 0$. To see this define $n(t) := \langle \mathbf{1}, \boldsymbol{s}(t) \rangle - \frac{\lambda}{v_{\min}}$, then

$$\frac{d}{dt} n = \langle \mathbf{1}, \dot{\boldsymbol{s}}(t) \rangle = \lambda - \boldsymbol{v}^T \boldsymbol{s}(t) = \left(\frac{\lambda - \langle \boldsymbol{v}, \boldsymbol{s}(t) \rangle}{v_{\min}} \right) v_{\min},$$

   but $\frac{-\langle \boldsymbol{v}, \boldsymbol{s}(t) \rangle}{v_{\min}} < -\langle \mathbf{1}, \boldsymbol{s}(t) \rangle$ and so the following differential inequality results

$$\frac{d}{dt} n < \left(\frac{\lambda}{v_{\min}} - \langle \mathbf{1}, \boldsymbol{s}(t) \rangle \right) v_{\min} = -n v_{\min}.$$

   Integration of the latter inequality when $n(0) > 0$ implies $n(t) < e^{-v_{\min} t} n(0)$, the result now follows because if $n(0) < 0$, then $n(t)$ can never be positive and so

$$\sup_{0 \leq t \leq T} \|\boldsymbol{s}(t)\|_1 \leq \frac{\lambda}{v_{\min}} + e^{-v_{\min} T} \left(\langle \mathbf{1}, \boldsymbol{s}_0 \rangle - \frac{\lambda}{v_{\min}} \right) =: C(p, T, \boldsymbol{s}_0). \tag{2.24}$$

2. Now consider (1.2) and define the state vector $\boldsymbol{s} = (S, R_1, R_2, X)$, the vector of initial conditions $\boldsymbol{s}_0 = (S(0), R_1(0), R_2(0), X(0))$, the vector $\mathbf{1} = (1, 1, 1, 1)$ and the vector of parameters $p = (\mu, \sigma, m, m_1, m_2, \gamma, \beta, \alpha, \tau_{\max}, c_1, c_2)$.

Equation (1.2) is point dissipative in the sense that if $\tau_1 \in L^\infty(0,T)$ is any measurable function with $0 \leq \tau_1(t) \leq \tau_{\max}, \tau_2(t) = \tau_{\max} - \tau_1(t)$ then

$$\langle \mathbf{1}, \boldsymbol{s}(t)\rangle = 1 - \mu e^{-\mu t}\langle \mathbf{1}, \boldsymbol{s}_0\rangle,$$

for all $t \geq 0$. To see this define $n = (1-(S+R_1+R_2+X))/\mu = (1-\langle \boldsymbol{s}, \mathbf{1}\rangle)/\mu$ and a short calculation shows that $\dot{n} = -\mu n$. As a result $n(t) = e^{-\mu t}n(0)$ and therefore

$$\langle \mathbf{1}, \boldsymbol{s}(t)\rangle = X(t) + R_1(t) + R_2(t) + S(t) = 1 - \mu e^{-\mu t}n(0), \qquad t \geq 0$$

and so

$$\sup_{0\leq t\leq T} \|\boldsymbol{s}(t)\|_1 \leq 1 - e^{-\mu T} + e^{-\mu T}\langle \mathbf{1}, \boldsymbol{s}_0\rangle =: C(p,T,\boldsymbol{s}_0). \tag{2.25}$$

The bounds (2.24) and (2.25) ensure that Theorem 2.4.1 can be applied to Problems 1 and 2 and the result follows. □

The following theorem illustrates that when condition (2.6) of Theorem 2.3.1 applies, the optimal antibiotic deployment protocol cannot be antibiotic mixing. Indeed, within the optimal protocol there is a time interval over which one of the drugs should not be deployed and the analysis immediately below tells us that this is because condition (2.6) can be thought of as telling us when the prevalence of resistance to one of the antibiotics is *too high*. We formalise this idea in the following theorem.

Theorem 2.5.1 *Suppose that there is a finite constant C depending on $\mathsf{C}, \boldsymbol{s}_0, T$ and p (but not A) such that for any function $A \in \mathcal{U}$, the solution $\boldsymbol{s}$ of (3.2) with $\boldsymbol{s}(0) = \boldsymbol{s}_0$ satisfies $\|\boldsymbol{s}\| \leq C(\mathsf{C}, p, T, \boldsymbol{s}_0)$. Also assume that condition (2.6) holds:*

$$\langle \boldsymbol{w}, \mathcal{G}(p)\boldsymbol{s}_0\rangle \neq 0$$

and write $\boldsymbol{s}^*(t)$ for the solution of (3.2) corresponding to an optimal control $A^*(t)$ of Problem A. *As a result, to each T we can associate at least one optimal control A^*_T by Theorem 2.4.1.*

Under these restrictions there exists uncountably many $T > 0$ for which $A^*_T(\cdot)$ takes either the value 0 or C *on a non-trivial sub-interval of $[0, T]$ of the form $[0, \tau)$ and so* cannot *be a mixing protocol.*

Proof

Let (T_n) be any positive sequence of times converging to zero and let $A^*_{T_n}$ be an optimal solution of Problem A associated with these times; such a sequence is well-defined from the conditions of the theorem. Now define the switching function $\sigma_n(t) := \langle \boldsymbol{m}_n(t), \mathcal{G}(p)\boldsymbol{s}_n(t)\rangle$ where $\boldsymbol{s}_n$ and $\boldsymbol{m}_n$ provides a solution of the re-scaled Euler-Lagrange equations given by the pair (2.8) and (2.10) when the function $A(t)$ in those equations is given by the optimal control $A^*_{T_n}$. (The rescaling alluded to changes the time interval of the problem from $[0, T]$ to $[0, 1]$ and so this will be assumed in the remainder of the proof.)

As $T_n \to 0$ in the Euler-Lagrange equations (2.8) and (2.10), the associated solutions $(\boldsymbol{s}_n, \boldsymbol{m}_n)$ with control $A_n := A^*_{T_n}$ satisfies

$$\boldsymbol{s}_n \to \boldsymbol{s}_0 \quad \text{and} \quad \boldsymbol{m}_n \to (1-t)\boldsymbol{w},$$

as $n \to \infty$, where the convergence is strong in $W^{1,\infty}(0, 1)$, as can be seen by bootstrapping on the assumption of the existence of the *a-priori* bound $\|\boldsymbol{s}_n\| \leq C(p, T, \boldsymbol{s}_0)$. Thus, the corresponding sequence of switching functions (as given in (2.5) but now with $\boldsymbol{m}(t)/T$ replacing $\boldsymbol{\lambda}(t)$)

$$\sigma_n(t) := \frac{1}{T}\langle \mathcal{G}(p)\boldsymbol{s}_n(t), \boldsymbol{m}_n(t)\rangle \qquad \text{satisfies} \qquad \sigma_n(t) \to \frac{(1-t)}{T}\langle \boldsymbol{w}, \mathcal{G}(p)\boldsymbol{s}_0\rangle,$$

strongly in $W^{1,\infty}(0, 1)$ as $n \to \infty$.

However, the affine function of t, $(1-t)\langle \boldsymbol{w}, \mathcal{G}(p)\boldsymbol{s}_0\rangle$ defined for $0 \leq t \leq 1$ is non-zero on $[0, 1)$ by assumption and has a transverse zero at $t = 1$. As a result, by the properties of uniform convergence, there is a sequence τ_n converging to 1 from below such that for all large enough n the function $\sigma_n(t)$ is non-zero in $[0, \tau_n)$.

Now let $A^0_n(t)$ denote any measurable function bounded below by 0 and above by C. From

(2.4) the optimal control A_n has the form

$$A_n(t) = \begin{cases} \mathsf{C} & : \ \sigma_n(t) > 0, \\ 0 & : \ \sigma_n(t) < 0, \\ A_n^0(t) & : \ \sigma_n(t) = 0, \end{cases}$$

for each n, it follows for sufficiently large n that A_n has the form

$$A_n(t) = \begin{cases} \mathsf{C} & : \ 0 \leq t \leq \tau_n, \\ A_n^0(t) & : \ \tau_n < t \leq 1, \end{cases}$$

if we assume that $\langle \boldsymbol{w}, \mathcal{G}(p)\boldsymbol{s}_0 \rangle > 0$. If, on the other hand $\langle \boldsymbol{w}, \mathcal{G}(p)\boldsymbol{s}_0 \rangle < 0$, then

$$A_n(t) = \begin{cases} 0 & : \ 0 \leq t \leq \tau_n, \\ A_n^0(t) & : \ \tau_n < t \leq 1, \end{cases}$$

completing the proof. □

Applying Theorem 2.5.1 to Problems 1 and Problems 2 gives the following natural condition on the form of the optimal controls. From equation (2.15) in the case of Problem 1, condition (2.6) can be written

$$h(y_a(0) - y_b(0)) \neq 0$$

whereas from equation (2.19) in the case of Problem 1 this abstract condition becomes

$$R_1(0) - R_2(0) \neq 0.$$

We can see from an epidemiological perspective that the abstract condition (2.6) has a very simple and practical interpretation: if resistance to one of the antibiotic is greater than to the other, do not use that antibiotic.

We now ask what happens when we take the idea hinted at in the previous paragraph of deploying only one antibiotic when the situation demands, for example use only drug 2 if if $R_1(t) > R_2(t)$, and extrapolate it as a deployment rule into the future. While this

protocol will not usually produce an optimal policy, in the next section we show that it can produce effective rotational protocols that are superior to antibiotic mixing.

2.6 Numerical examples

Throughout this chapter we shall use the parameter set for Problem 1 defined by

$$p^{(1)} := \Big\{\lambda = 100,\ d = 1,\ c = \frac{3}{2},\ h = 1,\ r_w = 0,\ s = \frac{1}{10}, ... \\ ...,\ r_a = \frac{9}{10},\ r_b = \frac{1}{10},\ b = \frac{4}{100}\Big\}$$

with the following initial conditions

$$\boldsymbol{s}_0^{(1)} := \left\{x(0) = (c + rw)/b,\ y_w(0) = \frac{\lambda}{c} - \frac{d}{b} - \frac{dr_w}{bc},\ y_a(0) = 0,\ y_b(0) = 0\right\}.$$

This set of parameters differs from the values given in [15] as they used $s = 1/1000, ra = 1/10$ and the parameter b does not appear to have a defined value in this reference. The only parameter that we allow to change in this section, unless otherwise explicitly stated, is the length of time used in the numerical experiment which we denote, both here and throughout, by the parameter T.

Let us clarify the reason why we consider a different parameter set than [15] is because we believe that there is no biological reasons on why to consider drug symmetry, and the mathematical importance of breaking this symmetry has already been discussed in (2.1). For these reasons we will make explicit the notion that the two antibiotics A and B are chosen to come from different functional groups in the sense, for example, of [116]. This leads us to enforce the restriction on the parameters that

$$r_a \neq r_b.$$

When working with Problem 2 we shall use the numerical parameter set

$$p^{(2)} := \Big\{\mu = \frac{1}{10}, \sigma = \frac{1}{4}, m = \frac{7}{10}, m_1 = \frac{1}{20}, m_2 = \frac{1}{20}, \gamma = \frac{3}{100}, \beta = 1, ... \\ ..., \alpha = \frac{4}{5}, \tau_{\max} = \frac{1}{2}, c_1 = \frac{35}{100}, c_2 = \frac{1}{20}\Big\}$$

with initial conditions

$$\boldsymbol{s}_0^{(2)} := \left\{S(0) = \frac{1}{5}, R_1(0) = \frac{3}{10}, R_2(0) = \frac{1}{10}, X(0) = \frac{2}{5}\right\}.$$

2.6.1 Drug rotation outperforms mixing

The first numerical example, illustrated in Figure 2.1, provides a comparison of equation (1.1) for symmetric and asymmetric parameter sets, where optimal mixing is compared with a sequence of cycling protocols. In the symmetric case of Figure 2.1a where 50-50 mixing provides the optimal mixing protocol, the protocols that cycle between the two antibiotics are inferior to optimal mixing; note that the optimal protocol itself is not known for these parameters so this figure is a comparison of several sub-optimal protocols. In Figure 2.1c where asymmetric parameters are used (the values in $(p^{(2)}, \boldsymbol{s}_0^{(2)})$) and 50-50 mixing performs poorly as a result, a range of cycling protocols biased to one of the drugs outperform optimal mixing provided each cycle occurs sufficiently quickly.

The purpose of this computation is to show that cycling and mixing protocols cannot be compared in any definitive sense: cycling can beat mixing and vice versa, the precise nature of the comparison depends on the structure of the cycling itself and on the numerical parameters used in the mathematical model.

Figure 2.2 shows the result of a numerical computation that deploys an optimisation algorithm to determine the best rotational protocols where the asymmetric parameter set $(p^{(1)}, \boldsymbol{s}_0^{(1)})$ has been used to parameterise the model (1.1). While both antibiotic rotation protocols outperform optimal mixing, if only by relatively small amount with less than 1% difference, the dynamics of antibiotic rotation shown as black lines exhibit spikes whereby drug resistance can increase sharply after the introduction of a new antibiotic regime. Nev-

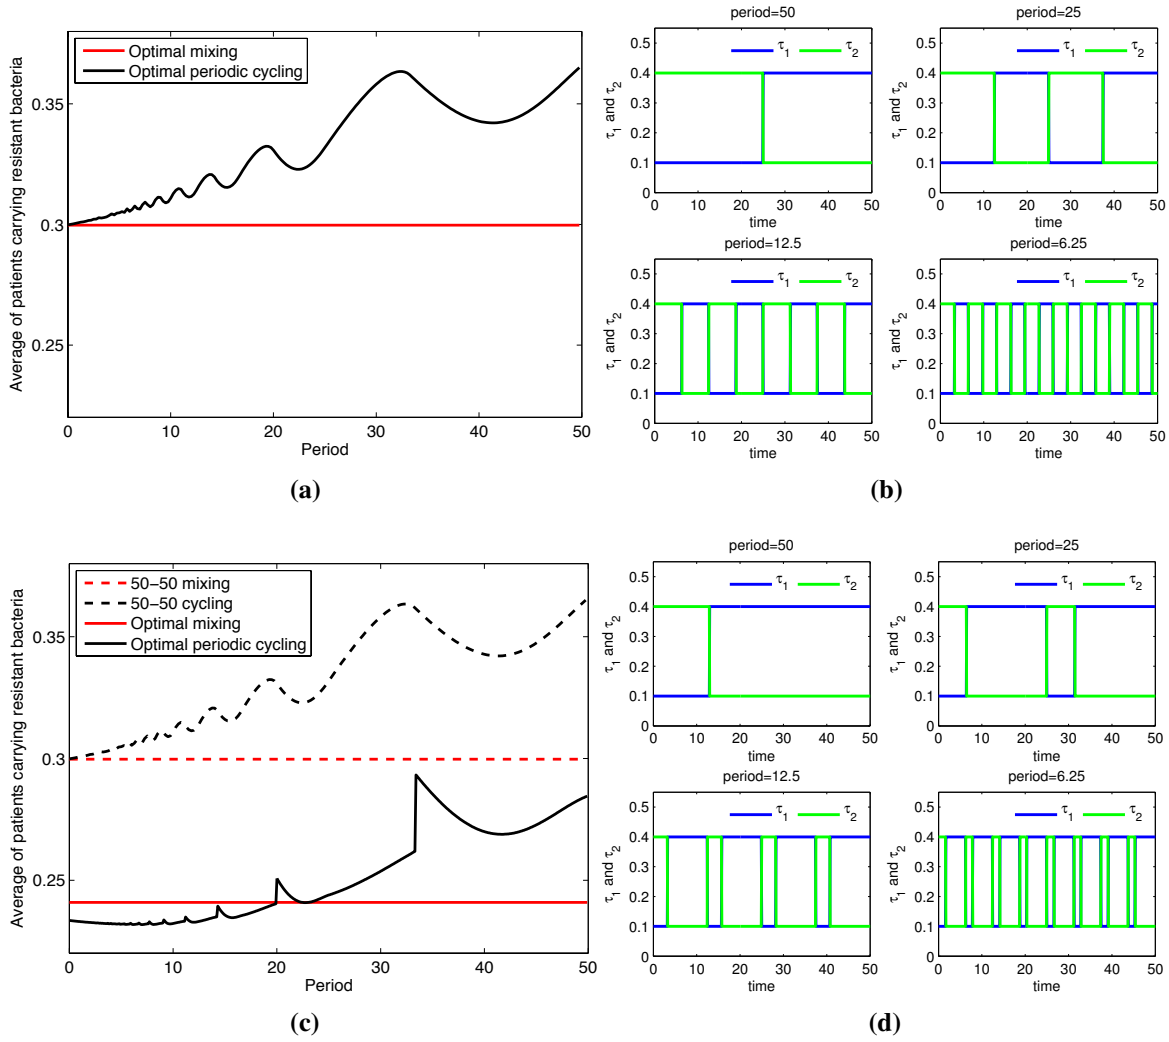

Figure 2.1: Two different parameter sets, one symmetric and one asymmetric, are used in Problem 2 to compute the response to the cycling protocols shown in the right-hand column and mixing protocols: in (A-B) the symmetric parameter values are taken from [13] but in (C-D) we used the asymmetric set $(p^{(2)}, \boldsymbol{s}_0^{(2)})$ defined in Appendix A, taking $T = 50$ in both cases. The (red) mixing and (black/solid) cycling lines in the two figures illustrate that cycling protocols may be outperformed by the optimal mixing protocol and vice versa (the symmetric case (A) and the asymmetric case (C), respectively). (The dashed lines in (C) are a reproduction of the data from (A); the cycling protocols used in (D) are biased towards more frequent use of one of the drugs whereas the cycling protocols in (B) may be described as 50-50.

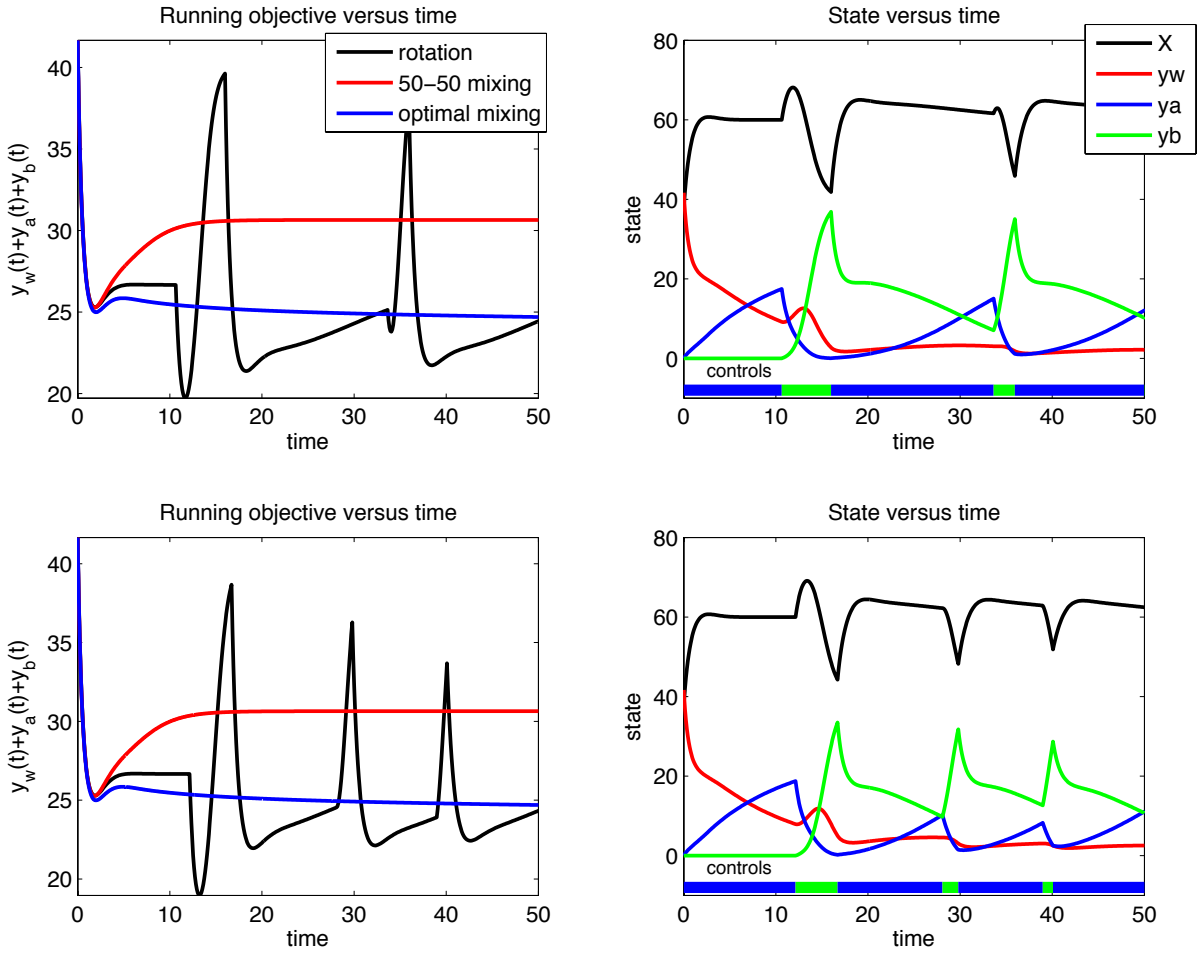

Figure 2.2: For a treatment duration of $T = 50$, the PICS used to simulate (1.1) is $(p^{(1)}, \boldsymbol{s}_0^{(1)})$: the best 5 and 6-switch controls achieve treatment objectives of 1259 and 1257 patients, respectively, optimal mixing is worse at 1260 and 50-50 mixing is worse still at 1506 patients.

ertheless, it is with rotational protocols, and not through mixing protocols, that we can minimise the performance measure defined in [15].

2.6.2 Adaptive rotation protocols

With such a large body of theoretical and numerical evidence that optimal mixing can be outperformed not only by the difficult-to-determine optimal treatment but by very many different rotating protocols, how might we find just one of the latter without employing

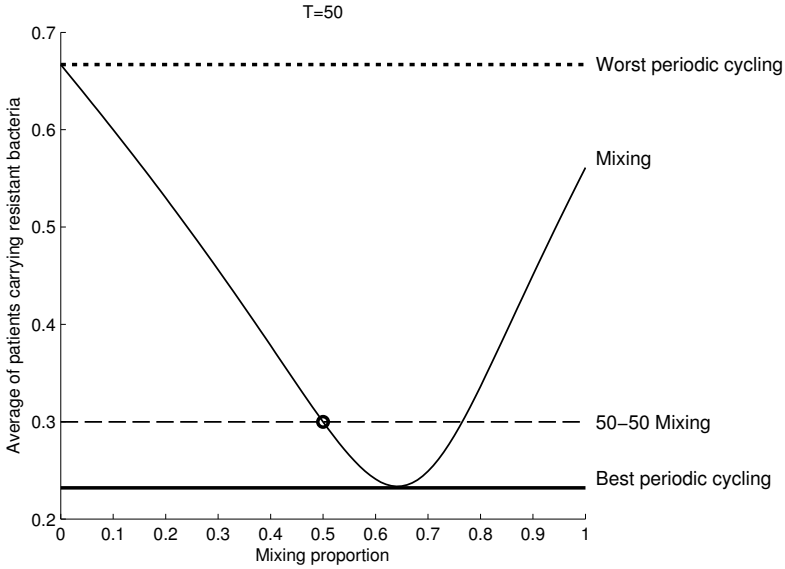

Figure 2.3: Average of patients carrying resistant bacteria for different mixing proportions in model (1.2) when $T = 50$. Note that 50-50 mixing is very ineffective and that the best mixing proportion is as good as the best periodic cycling. The opposite is also true, as the worst mixing protocol is as ineffective as the worst periodic cycling, in this case to the deployment of a single-drug.

searching algorithms of any kind? From a practical perspective it is important to know if we can simply inspect the current state of the system to find patterns of resistance and susceptibility of both current and historically observed pathogens and subsequently use that information to decide when to invoke a switch of antibiotics?

In short, is possible to design suboptimal, but effective *closed-loop feedback controls* based on practicable heuristics? To probe this issue, consider the following feedback control laws:

Rule 1: in Problem 1 use antibiotic A until $y_a(t) > y_b(t)$ and then switch to antibiotic B. Similarly, if $y_b(t) > y_a(t)$ switch to A.

Rule 2: in Problem 2 use antibiotic 2 if $R_1(t) > R_2(t)$, otherwise use 1.

One further concept needed to complete the definition of the feedback controls is the idea of a *sample time*. The variable t in Rule 1 and Rule 2 may refer to all instances of time or t could be a sample time whereby the control decision is taken periodically or at some other prescribed instants in time. In the numerical examples of the next section we take the latter

approach due to its practical relevance to managing antibiotic use in hospitals and ask how often must the system be sampled so that the feedback rules outperform antibiotic mixing? This can be interpreted in the sense of how much information do we need so that a protocol based on exploiting that information outperforms protocols founded on no information at all, like cycling and mixing.

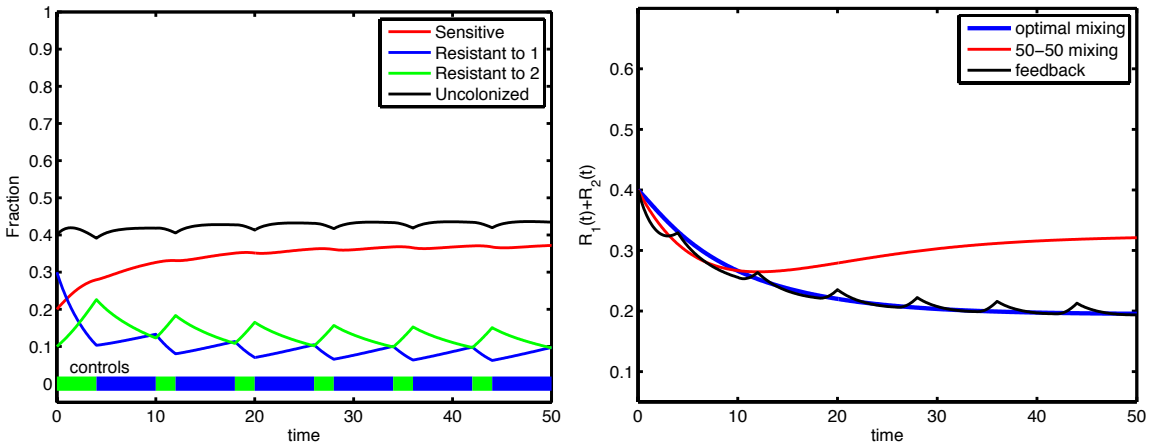

Figure 2.4: Rule 2 applied to (1.2) when $T = 50$ and 25 maximum possible switches. (A) The state response obtained Rule 2. (B) The running treatment objective (the function $R_1(t) + R_2(t)$) obtained using 50-50 mixing (shown in red, with treatment objective equal to 15), optimal mixing (blue, treatment objective close to 11.7) and the rule-based feedback (black, treatment objective close to 11.61).

Figure 2.5 shows one outcome of applying Rule 1 to Problem 1 using the same asymmetric parameter values as Figure 2.2 where it is evident that the rule-based control measure is superior to optimal mixing even though the rule only implements seven switches of antibiotic. Figure 2.4 is an analogous computation that implements Rule 2 on Problem 2 using parameters $(p^{(2)}, \boldsymbol{s}_0^{(2)})$. Similarly, the rule-based controller produces rotational protocols that outperform optimal mixing.

Finally, in Figure 2.6 both Rules 1 and 2 have been applied to equations (1.1) and (1.2) in the search for suboptimal rotational protocols that outperform antibiotic mixing. With a time parameter T of fifty units, no more than N switches of antibiotic were allowed on any given simulation and the dynamical systems (1.1) and (1.2) were sampled T/N time units

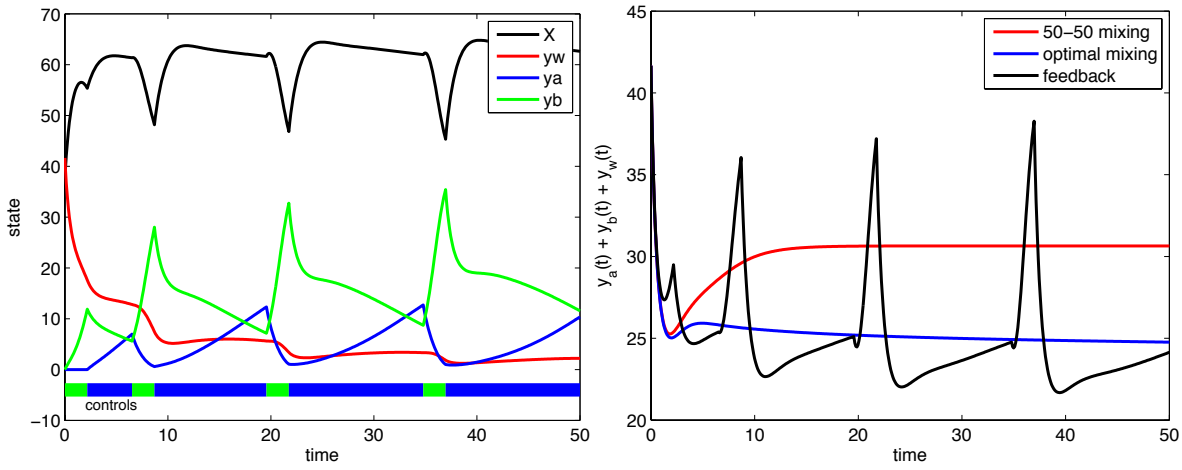

Figure 2.5: Rule 1 applied to (1.1) with $T = 50$ and 23 maximum possible switches. (A) The state obtained using Rule 1 has a performance of 1252 patients, less than optimal mixing strategy (1261) and 50-50 mixing (1506). (B) Comparison of the running treatment objectives (the function $y_a(t) + y_b(t) + y_w(t)$) of optimal mixing (blue line) and 50-50 mixing (red line) with the rule-based feedback treatment (black line). As the black line is lowest on average, the feedback outperforms all mixing protocols.

apart to make the deployment decision as to which antibiotic would be used until the next sample. The sampling parameter N is shown along the horizontal axis in Figure 2.6 where it is labelled as *sampling points* and both diagrams in the figure show that the performance of these rule-based controls (as plotted on the vertical axis) improves dramatically with increasing N, although not monotonically. In both cases a value of N is reached above which the feedback rules Rule 1 and Rule 2 outperform optimal mixing. We deduce from this computation that there are *infinitely many* alternating protocols superior to optimal mixing.

2.6.3 Releasing the must-treat constraint

We have shown that even simple heuristic-based controls are effective in controlling the prevalence of single-drug resistance, but what happens if we consider the possibility that bacteria can evolve resistance to multiple antibiotics? In order to take into account this pos-

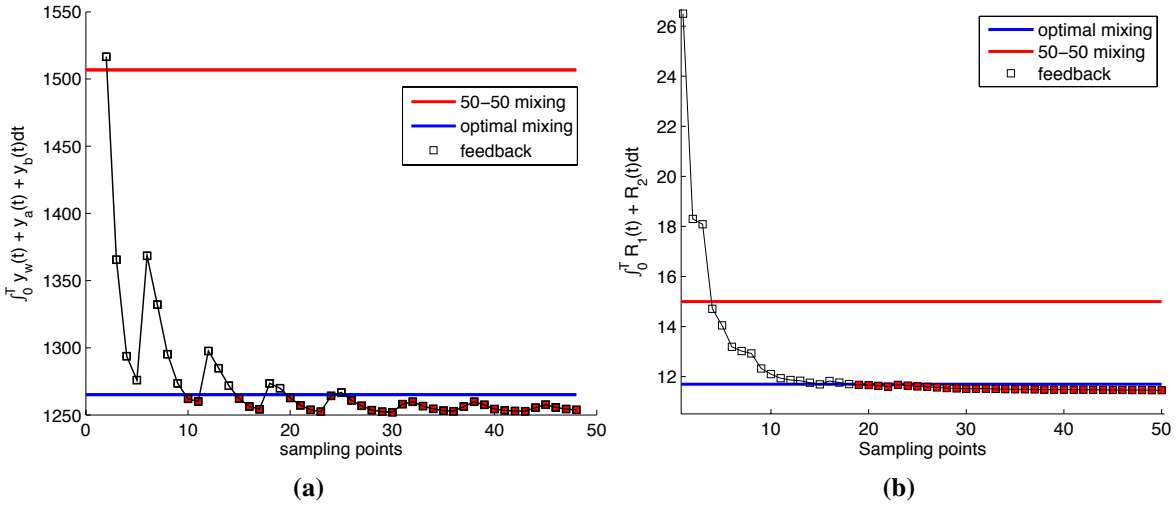

Figure 2.6: Comparing the performance of optimal mixing (blue), 50-50 mixing (red) with Rule 1 and Rule 2 (boxes). The filled boxes illustrate the number of sampling points for which the rules-based controllers outperform optimal mixing. The asymmetric parameter sets used for these simulations are defined in the text and $T = 50$ for both models (1.1) and (1.2), any number from 1 to 50 sampling points were used with at most one sample per unit time. Diagram (A) shows results obtained for (1.1) and (B) are results for (1.2).

sibility, we numerically simulated an extension of the model (1.2) described by equations (1.3).

It is clear from Figure 2.7 that by constantly deploying antibiotics into the system, we are always selecting in favour of antibiotic-resistant phenotypes, either single-drug resistant or multi-drug resistant. If such selective pressures occur independently of the patterns of drug usage, it is not clear that it may be possible to design policies that are effective in controlling the evolution of drug resistance.

A possible strategy would be to reduce the overall antimicrobial use in clinical settings, and thus decrease the selective pressures in favour of resistance, and indeed such an intervention has been trialled and reported to be partially effective [39]. From a control-theoretic perspective, it is clear that the must-treat constraint, expressed in Problem 2 as $\tau_2(t) + \tau_1(t) \equiv \tau_{\max}$, is a suboptimal constraint that if relaxed would yield policies that perform least as good as the previously estimated optimal.

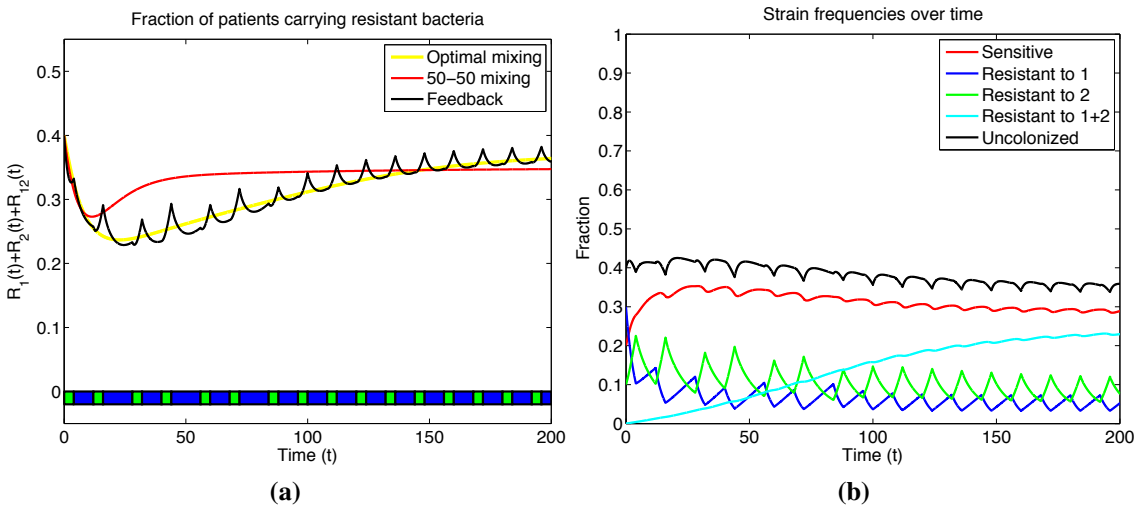

Figure 2.7: (a) Fraction of patients carrying resistant bacteria under different treatment policies in a model that includes multidrug resistance. The yellow line represents the optimal mixing protocol, the red line the 50-50 mixing and in black the feedback heuristic defined in Rule 2. (b) Frequencies of each strain under the feedback protocol. Note how as time increases there is a steady increase in the frequency of hosts colonised with a bacteria resistant to antibiotics 1 and 2.

In clinical settings releasing this constraint would introduce a moral dilemma: some patients would remain untreated. This predicament was discussed in [42], where it is argued that the ultimate solution to the antibiotic resistance problem would require us to put society before the individual and allow some patients to go untreated.

Of course, it would be difficult to trial such a strategy in practice, but in a theoretical model of antibiotic deployment like the ones discussed in this thesis it is an easy task to implement. For example, we could solve the following optimisation problem:

$$\text{Problem 3} \qquad \begin{cases} \min \ \int_0^T \left(R_1(t) + R_2(t) + R_{12}(t)\right) dt & \text{subject to constraints} \\ 0 \leq \tau_1(t) \leq \tau_{\max}, \tau_2(t) + \tau_1(t) \leq \tau_{\max} & \text{and equation } (1.3). \end{cases}$$

in order to find the optimal control that minimises multidrug resistance. But instead of synthesising the theoretical optimum, for illustration purposes only, we implement the following simple rule-based heuristic:

Rule 3: in Problem 3 use antibiotic 1 if $R_1(t) > R_2(t)$ and $(R_1 + R_2) < R_{12}$, deploy antibiotic 2 if $R_1(t) > R_1(t)$ and $(R_1 + R_2) < R_{12}$. If $(R_1 + R_2) \geq R_{12}$ do not deploy any drugs, that is $\tau_1 = 0$ and $\tau_2 = 0$.

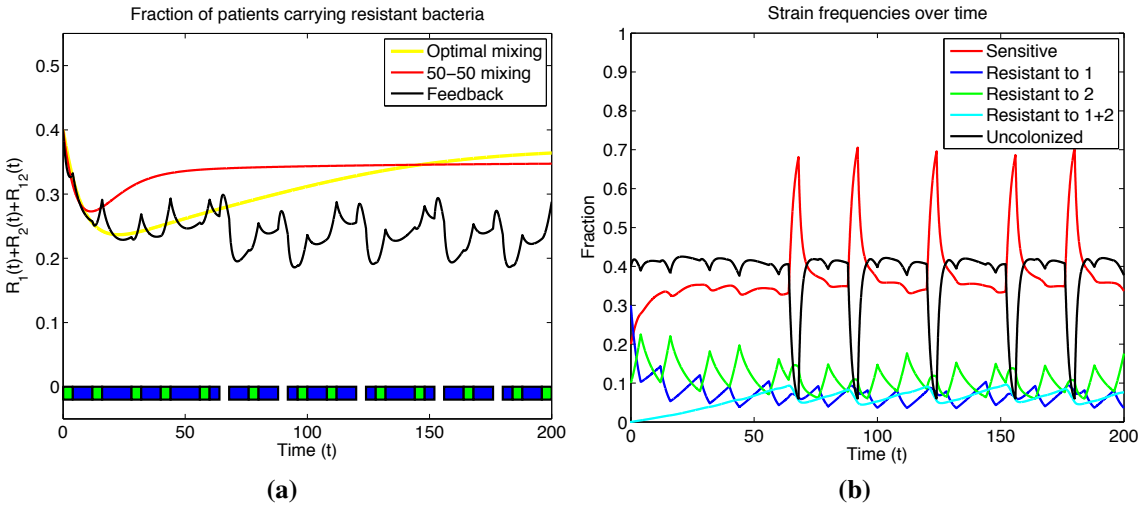

Figure 2.8: Releasing the must-treat constraint is effective in controlling drug resistance. However, the intervals of time where no drug is deployed are responsible of a sharp decrease of uncolonised patients and a dramatic increase in patients infected by a sensitive strain. (a) Fraction of patients carrying resistant bacteria in time. The barcode illustrates the drug used at each moment in time under the feedback control: blue corresponds to drug 1, green to drug 2 and a white box represents time intervals when no drug is deployed. (b) Strain frequencies over time for the feedback control.

Figure 2.8 shows the implementation of Rule 3 in an exemplar situation of model (1.3). A feature to notice is that only short intervals of no drug deployment have a huge impact in the prevalence of multidrug resistance. However, it is important to notice that these sudden decreases in levels of resistance are accompanied by a very sharp increase on the fraction of the population that is colonised by the sensitive strain, a feature illustrated in Figure 2.8b.

This observation highlights the risk of considering hospital-wide strategies that include periods of time where no antibiotics are used within a population. It may be possible, however, to decide on a per-patient basis which individuals can be left without treatment without a significant risk to their health state, but the SI models discussed so far do not allow us the possibility of tracking the health state of each individual in time. With this in

mind, in the last part of this thesis we will pose an agent-based model of a hospital ward that will allow us to design individualised strategies of drug deployment.

Chapter 3

In general, which is optimal: mixing or cycling?

In this chapter we ask the question *Which antibiotic deployment protocols select best against drug-resistant microbes: mixing or periodic cycling?* and demonstrate that the statistical distribution of the performances of both sets of protocols, mixing and periodic cycling, must have overlapping supports. In other words, it is a general, mathematical result that there must be mixing policies that outperform cycling policies and vice versa.

It is important to mention that in order to illustrate that 50-50 mixing is not *always* superior to antibiotic cycling as claimed in [13, 15], in the previous chapter we constructed counterexamples whereby a specific period of rotation between drugs from different functional classes and with different costs of resistance outperformed even the best mixing protocol. We did not claim that *any* periodic cycling protocol could outperform *all* antibiotic mixing protocols.

What we can assert, and prove in this chapter, is a general property of theoretical models of drug deployment as proposed in [13, 15] that, mathematically speaking, the best periodic cycling protocols are as good as the best mixing strategies. More precisely, we prove that the statistical support of the distributions of performances of both classes of protocol must overlap and believe that this property might be invoked as a tentative explanation of the

inconclusiveness of recent clinical trials designed to evaluate the efficacy of cycling [17].

3.1 Optimal cycling and mixing: two impractical protocols

The criticism in [16] that protocols determined using optimal control theory cannot be implemented in practise is valid. While the randomised allocation of drugs in a population (or 50-50 mixing) described in [13, 15] are practicable (also see [104]), we argue that the optimal mixing protocol defined in [16] as *optimally dividing the drugs between appropriately sized patient groups* is also an impractical, theoretical concept.

Moreover, even if it may be possible in theory to determine the best period of cycling, the best mixing proportion or even the optimal protocol, doing so certainly does not answer the much harder question

> *What is the likelihood that a given rotation will outperform a given mixing protocol?*

The difficulties in providing a meaningful answer to this are not to be underestimated. One would need a mathematical model of the epidemiological dynamics at play in the hospital capable of making hindcasts and predictions calibrated against known data. This would be akin to predicting the weather, only potentially harder because of the absence of physical laws. Attempts have been made at using simple mathematical models to produce forecasts of antibiotic resistance evolution over the coming decades [69], but criticism of the lack of biological depth followed [50]. The best one can ever hope for is to provide an ensemble of possible futures, as in [29]. As a result, it would be churlish to assert, on the basis of any one epidemiological model, that a specific protocol achieves optimality in real, clinical contexts or even just for a large class of mathematical models.

The purpose of the previous chapter was to provide general results, really little more than a reformulation of known mathematical ideas, which are subsequently applied to two specific mathematical models taken from [13, 15]. These applications illustrate that a wide variety of conditions are met whereby the best non-periodic rotations can outperform the best mix-

ing protocol and such conditions are not rare but arise when so-called *drug symmetries* are broken. The negative answer to Brown and Nathwani's question posed in the clinical review [17], *Does resistance to different antibiotics develop at the same rate?* provides one example of such an asymmetry.

The difficulty of determining optimal treatment protocols in practice is real and so we proposed adaptive rotation protocols as a practicable strategy that relies only on observations of local patterns of susceptibility and resistance to decide which drug, A or B, to give. Note, we *do not* claim that our adaptive feedback rule is anywhere near to being optimal nor a good approximation for such. It is a merely a tool for determining rotational protocols that outperform optimal mixing and we find, by example, that even this simple heuristic can achieve this goal for particular exemplar models. We do claim that the adaptive protocols discussed in Section 2.6.2 that seek to exploit more information by sampling the patient cohort for drug-resistant pathogens more frequently, in turn, perform better. We give examples to show that such protocols can, given enough information, outperform optimal mixing; [16] does not provide any evidence that this statement is false. Furthermore, it is absolutely true that the difference in performance between optimal rotation and optimal mixing found in models may be marginal, but this difference will be specific to the parameters implemented within any particular model and specific to the model itself.

We would like to emphasize that the analysis provided in the previous chapter and in [9] hold for a large class of epidemiological models* and independent of any particular parameter values or initial conditions. In particular, our only assumption over the interval of observation, T, is that it has to be finite, for our arguments to work it cannot be infinite. For this reason, we disagree with [16] that the arguments presented in [9] are only valid in general for short intervals of observation and are solely due to 'transients'. One must also bear in mind that as T is finite it is also a parameter in the problem and that optimal mixing, optimal cycling and the theoretically optimal protocol may well all depend on T itself. Hence there is no general rationale to support the idea that an optimal protocol found at $T = 50$, say, will also be optimal when extended to $T = 51$, let alone to $T = \infty$.

*There are many antibiotic deployment problems that one could pose outside of the mathematical class that we analyse in this article, but this class does encompass some of the models presented in [13, 15].

Instead of seeking for a particular set of parameters where periodic cycling outperforms mixing in specific models, in the following section we will provide statistical arguments that may help understand the difficulties that arise when trying to quantify the differences between cycling and mixing, and we will demonstrate that it is always possible to find cycling protocols that outperform mixing, and vice versa.

3.2 Cycling versus mixing: a theorem concerning their statistics

Rather than delve into the details of any particular epidemiological scenario or mathematical model, we shall keep this analysis general and work for the moment in an abstract framework. First of all, we shall need the vector of all the different patient classes. Following previous approaches we could write

$$\boldsymbol{y}(t) = (S(t), R_A(t), R_B(t), X(t)),$$

say, where each entry represents a time-series of the proportion of patients in different classes: drug-susceptible bacterial infections, drug-A resistant infections, drug-B resistant infections and uncolonised patients, respectively.

In principle we would need a mathematical model to describe how the vector $\boldsymbol{y}(t)$ changes in time. Let us assume that the hospital will operate at its maximal capacity at all times so that the sum of elements in $\boldsymbol{y}(t)$ must be a constant value. We may assume without the loss of any generality that this value is unity in the remainder, hence the entries of $\boldsymbol{y}(t)$ will contain the relative proportions of each of the different patient classes.

Let us now suppose the existence of a dynamical model, specifying none of its detail, of the form

$$\frac{d}{dt}\boldsymbol{y} = F(\boldsymbol{y}) + A(t)G_a(\boldsymbol{y}) + B(t)G_b(\boldsymbol{y}); \tag{3.1}$$

the models in [13, 15] are of this form. In (3.1) the functions G_a and G_b describe how the dynamics of the ICU unit or hospital ward are affected when we use one of the drugs, drug A or drug B. For example the use of drug A might well be positively correlated with

selection for drug A-resistant pathogens, thus increasing the value of $\frac{d}{dt}R_A(t)$. For equation (3.1) to make any sense we also need an initial condition. We could say $\boldsymbol{y}(0) = (1, 0, 0, 0)$ so that all patients are infected with the drug-susceptible pathogen to begin with, but with no drug-resistant pathogens. Clearly there are many other choices that could also be used and we will not restrict attention to any one in particular. Let us also suppose that policy ensures that everyone in the hospital is treated, according to our assumptions this means that $A(t) + B(t) = 1$ for all times.

Under mild conditions, the resulting equation

$$\frac{d}{dt}\boldsymbol{y} = F(\boldsymbol{y}) + A(t)G_a(\boldsymbol{y}) + (1 - A(t))G_b(\boldsymbol{y}). \tag{3.2}$$

can be shown to define a mathematical dependancy or mapping in the sense that to each function $A(t)$ representing a drug-deployment policy, we can associate a solution $\boldsymbol{y}(t)$ that depends on A; we will write $\boldsymbol{y}(A)$ or $\boldsymbol{y}(A)(t)$ for this dependancy. Provided that there exists a constant K independent of A such that for any solution of (3.2) there results

$$\limsup_{t\to\infty} \|\boldsymbol{y}(t)\| \leq K, \tag{3.3}$$

the vector-valued function $\boldsymbol{y}(A)(t)$ can be extended to the entire interval $[0, T]$. As we tacitly assumed earlier in this discussion that the entries of $\boldsymbol{y}(A)(t)$ sum to unity for all t and for any bounded function A between 0 and 1, we may assume (3.3) holds throughout our discussion.

We want to decide upon a drug deployment policy, $A(t)$, that minimises some performance criterion. So, we now define a weight vector $\boldsymbol{w}$ that gives the relative importance of each component of $\boldsymbol{y}(A)(t)$. A mathematical optimisation procedure can determine the optimal policies by locating a function or family of functions $A(t)$ between 0 and 1 such that the measure of the performance of the protocol A defined by

$$\mathcal{J}(A) := \frac{1}{T}\int_0^T \langle \boldsymbol{w}, \boldsymbol{y}(A)(t)\rangle dt \tag{3.4}$$

is minimal.

So, for example, if $\boldsymbol{w} = (0, -1, -1, 0)$, then the performance of our policy is measured by the number of observed drug-resistant infections per unit time: $\langle \boldsymbol{w}, \boldsymbol{y}(t) \rangle = -R_A(t) - R_B(t)$; throughout the remainder, the latter will be our performance criterion. Likewise, if $\boldsymbol{w} = (0, 0, 0, 1)$ then $\langle \boldsymbol{w}, \boldsymbol{y}(t) \rangle = X(t)$ and maximising $\mathcal{J}(A)$ corresponds with increasing the number of uncolonised patients seen over the observation period.

Notice that the performance measure can be thought of as a functional: to each protocol $A(t)$, a function defined on the interval $[0, T]$ taking values in the interval $[0, 1]$, we can associate the value $\mathcal{J}(A)$. Our performance measure $\mathcal{J}(\cdot)$ makes sense as a mapping when we supply it with a domain, so let us define the following spaces of functions:

Let Per denote the set of all periodic functions bounded between 0 and 1, defined on the observation interval $[0, T]$ and with period T/n for some integer $n \geq 1$; these functions are so constrained because we cannot possibly treat more than all the patients in the hospital or ward. Let BB be the set of functions bounded between 0 and 1 such that for all t between 0 and T one of either $\alpha(t) = 0$ or $\alpha(t) = 1$ must be true, where α represents any function in BB; thus BB represents the set of protocols that only deploy one drug at a time. Now define Cyc $:=$ BB $\cap$ Per which is the set of all congruent, periodic cycling protocols. The set of mixing protocols, Mix, is much smaller and defined by the constant functions, so that a mixing protocol α satisfies $\alpha(t) = c$ for $0 \leq t \leq T$ and some constant c between 0 and 1; the value $c = 1/2$ corresponds to the so-called '50-50' or random mixing protocol. The following mathematical results are key to our argument.

We say that a sequence of functions $(u_n) \subset L^\infty([0, T], \mathbb{R}^n)$ converges weak* to u if

$$\lim_{n\to\infty} \int_0^T \varphi(t) \cdot u_n(t) dt = \int_0^T \varphi(t) \cdot u(t) dt$$

for all $\varphi \in L^1([0, T], \mathbb{R}^n)$ and we write $u_n \overset{*}{\rightharpoonup} u$ as $n \to \infty$.

Assume that $F(\cdot), G_a(\cdot)$ and $G_b(\cdot)$ are smooth functions and that there is a K independent of A such that for any solution of (3.2), there results (in any finite dimensional norm) $\limsup_{t\to\infty} \|\boldsymbol{y}(t)\| \leq K$.

Theorem 3.2.1 *The mapping $\mathcal{J} : L^\infty([0,T],\mathbb{R}^n) \to \mathbb{R}$ defined in (3.4) is a continuous functional with respect to weak* convergence and for each $m \in \mathsf{Mix}$ there is a sequence $(c_n) \subset \mathsf{Cyc}$ such that $c_n \overset{*}{\rightharpoonup} m$ as $n \to \infty$. There is an optimal mixing protocol and a worst mixing protocol, namely values $\underline{m}, \overline{m} \in [0,1]$ such that*

$$\mathcal{J}(\underline{m}) = \min_{0\le m\le 1} \mathcal{J}(m) \qquad \textit{and} \qquad \mathcal{J}(\overline{m}) = \max_{0\le m\le 1} \mathcal{J}(m).$$

Proof The first part follows almost immediately from, for example, [66, Theorem 1], the weak* continuity of $\mathcal{J}(\cdot)$ with respect to $A \in L^\infty([0,T],\mathbb{R}^n)$ stemming from the form of (3.2) as a smooth differential equation defined affinely with respect to A. The last part follows because the map $m \mapsto \mathcal{J}(m)$ is a continuous function defined on the interval $[0,1]$ and so achieves its extreme values. □

Corollary 4 *The best possible performance of the set of congruent antibiotic cycling protocols is at least as good as the best performance of the mixings in the sense that*

$$\sup_{A\in\mathsf{Cyc}} \mathcal{J}(A) \ge \sup_{A\in\mathsf{Mix}} \mathcal{J}(A).$$

Proof Suppose, seeking a contradiction, that $\sup_{A\in\mathsf{Cyc}} \mathcal{J}(A) < \sup_{A\in\mathsf{Mix}} \mathcal{J}(A) = \mathcal{J}(m)$ where $m \in [0,1]$ is a constant. However, by Theorem 3.2.1 of the appendix there is a sequence $c_n \overset{*}{\rightharpoonup} m$ as $n \to \infty$ and therefore

$$\mathcal{J}(m) = \lim_{n\to\infty} \mathcal{J}(c_n) \le \sup_{A\in\mathsf{Cyc}} \mathcal{J}(A) < \sup_{A\in\mathsf{Mix}} \mathcal{J}(A) = \mathcal{J}(m)$$

because $\mathcal{J}$ is continuous with respect to weak* convergence. This is a contradiction and the result follows. □

Let us examine more closely what this mathematical result means theoretically for our ambiguously phrased question of whether 'cycling or mixing is best'. Firstly, because of the form of $\mathcal{J}$ (it is continuous with respect to weak* convergence) a type of robustness with respect to the cycles follows: if there is at least one cycling protocol that outperforms optimal mixing, there are infinitely many nearby cycling and acyclical drug rotation protocols that also outperform optimal mixing.

Secondly, if the optimal mixing protocol, m say, performs better than all the congruent cycling protocols and yet $\epsilon > 0$ is any number giving a measure of sub-optimality, then there is a congruent cycling policy A_ϵ^{cyc} and a corresponding state response $\boldsymbol{y}(A_\epsilon^{cyc})$ such that

$$\mathcal{J}(m) - \epsilon \;\; < \;\; \mathcal{J}(\boldsymbol{y}(A_\epsilon^{cyc})).$$

This tells us that if the optimal antibiotic mixing performs better than all the cycling protocols in our model, if we want a cycling policy that performs within 99% or 99.9% of the performance of the best mixing, there are cycling policies in Cyc that will achieve this.

From the perspective of seeking theoretical support for antibiotic cycling this appears to represent a positive outcome but, particularly in light of the accompanying critique [16], we will discuss the following result. We believe it may have important practical implications.

Corollary 5 *The worst possible performance of the set of congruent antibiotic cycling protocols is at least as bad as the worst performance of the mixings in the sense that*

$$\inf_{A \in \mathsf{Cyc}} \mathcal{J}(A) \leq \inf_{A \in \mathsf{Mix}} \mathcal{J}(A).$$

Proof The proof is almost identical to Corollary 4. □

Let us now define the performances of the worst and best mixing protocols:

$$\underline{p} := \min_{0 \leq m \leq 1} \mathcal{J}(m) \qquad \text{and} \qquad \overline{p} := \max_{0 \leq m \leq 1} \mathcal{J}(m).$$

If we draw the graph of the performance surface $\mathcal{J}(A(\tau_1, \tau_2))$ as a function of (τ_1, τ_2) then Corollaries 4 and 5 predict that this two-dimensional graph will either just touch, or lie above $\overline{p}$ in one region of the (τ_1, τ_2) plane. The same theory predicts it will also either just touch, or lie below $\underline{p}$ in another region of the plane. Figure 3.1 illustrates a particular instance of this idea, it was obtained using a model from [13] whereby the performance surface of the cycling protocols does in fact pass below the performance of the optimal mixing protocol (by around 1% or less); Figure 3.1(b) illustrates where this property is satisfied in the (τ_1, τ_2)-plane.

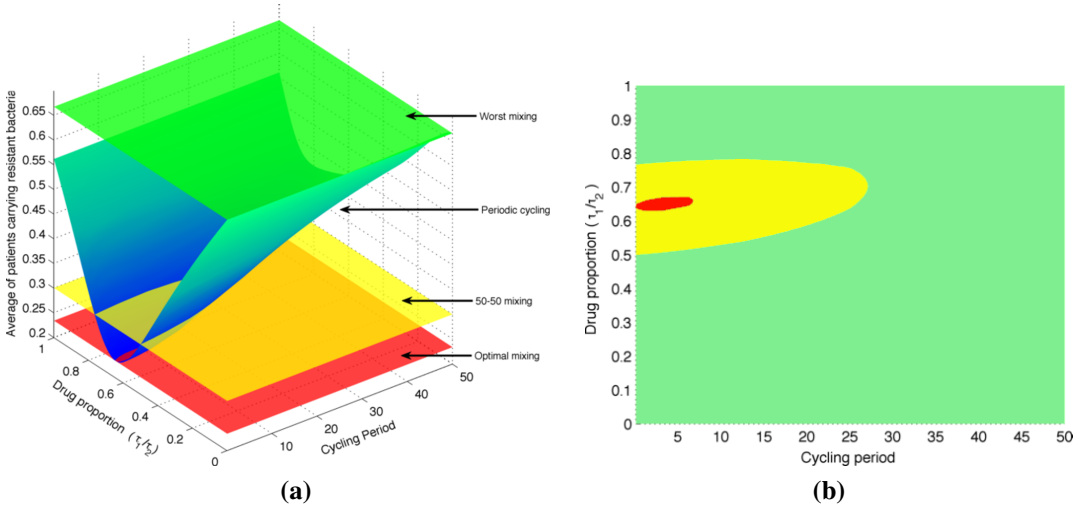

(a) **(b)**

Figure 3.1: (a) The performance measure $\mathcal{J}$ defines a performance surface formed from a sample of congruent cycling protocols (shown as a blue-green surface); red shows the performance of optimal mixing ($\underline{p}$), yellow of 50-50 mixing and green the worst possible mixing ($\overline{p}$). (b) The red region shows the neighbourhood from (A) where the best cycling protocols outperform the best mixing, the yellow region shows where cycling outperforms 50-50 mixing. In both (A) and (B) we used $T = 50$.

Let us phrase this slightly differently: for each performance measure $p \in (\underline{p}, \overline{p})$ lying anywhere between the extreme performances of the best and worst mixing protocols, there is a congruent cycling protocol defined by (τ_1, τ_2) such that $p = \mathcal{J}(A(\tau_1, \tau_2))$. From this we deduce what might be called an interlacing property: for any model of the form (3.2) there exist congruent cycling protocols, labelled (τ_{11}, τ_{12}) and (τ_{21}, τ_{22}), and two mixing protocols, m_1 and m_2, such that

$$\mathcal{J}(m_1) < \mathcal{J}(A(\tau_{11}, \tau_{12})) < \mathcal{J}(m_2) < \mathcal{J}(A(\tau_{21}, \tau_{22})). \tag{3.5}$$

In light of this, when we ask *Which is optimal: mixing or cycling?* it is impossible to answer unless further clarification is given as to what meaning is really intended by that question. In more statistical language, (3.5) means that the support of the two distributions of all possible performances of the mixing protocols and the congruent cycling protocols *must* overlap, with the distribution of performances of the cyclings having the potentially larger support because of Corollaries 4 and 5.

This universal result within the model class (3.2) points to one possible reason for the inconclusive nature of clinical trials performed over the last decade when evaluating the efficacy of cycling. If such a trial were to be mimicked computationally by simulating a model such as (3.2), if model parameters are sampled sufficiently widely one must find an inequality like (3.5) within the simulated data. Indeed, this property is observed in Figure 3.2 computed using a model from [13] where one can clearly see that the empirical distributions of performances of cycling and mixing protocols have almost identical supports.

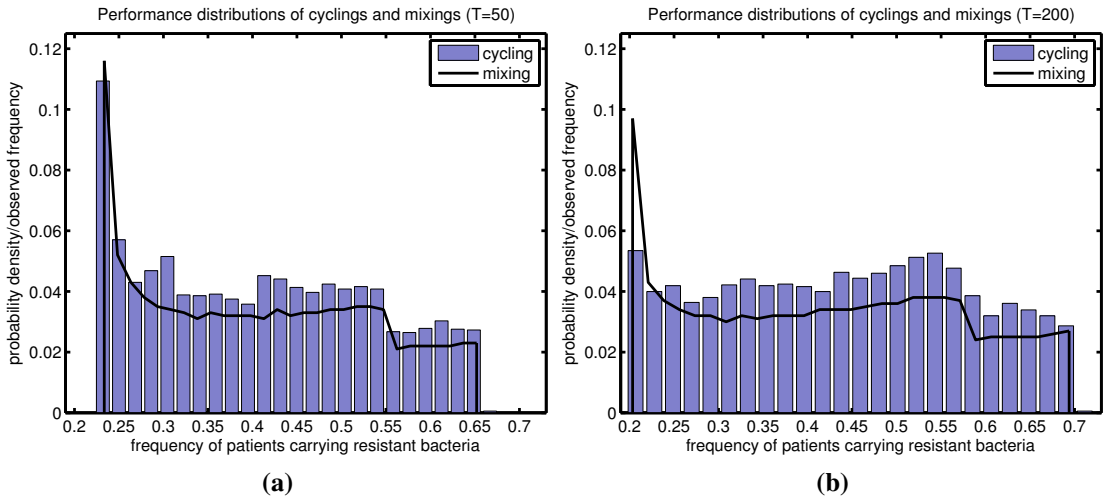

Figure 3.2: Two histograms comparing the distribution of performances of mixing protocols and a sample of the congruent cycling protocols determined using the mathematical model given in [13]; the same parameters are used in both plots except that (A) T = 50 and (B) T= 200 time units. Note, consistent with the theoretical results given in the text, the support of the mixing distribution lies either on or within the cycling distribution in both cases.

In order to generate Figure 3.2 we constructed an approximation of the probability density function of the performances of mixing protocols numerically directly from the performance locus $\{\mathcal{J}(m) : 0 \leq m \leq 1\}$. For the cycling protocols we first defined a set of *periods* $\rho := T/N$ where $N = 1, ..., 30$ are integers, we then defined $\tau_1 := t\rho$ and $\tau_2 := (1-t)\rho$ where $t = 0/M, 1/M, 2/M, ..., M/M = 1$, so that $\tau_1 + \tau_2 = \rho$, we finally set $M = 100$. The performances of the 3,030 different congruent cycling protocols: $\mathcal{D} := \{\mathcal{J}(A(\tau_1, \tau_2)) : \tau_1 = t\rho, \tau_2 = (1-t)\rho,\ \rho = T/N, N = 1, ..., 30, t = 0, 1/100, ..., 1\}$

were then binned (using 25 bins) with the `histc` command in Matlab. If the choice of time units as *days* and other parameter values given in [13, 16] could be justified from empirical data, some of the cycling protocols used to form $\mathcal{D}$ could be eliminated as clinically impractical.

Part II

Optimal antibiotic deployment into a single-host

Chapter 4

Controlling the evolution of resistance

Since the discovery of antibiotics conventional wisdom has recommended that bacterial infections should be treated aggressively [37]. This principle was even stated as a moral by Alexander Fleming himself: *"if you use penicillin, use enough"* [40]. The hit-early and hit-hard prescription paradigm has impinged upon our understanding of antimicrobial treatment to the point that nowadays it is widely believed that shortening treatment duration may be responsible for promoting the evolution of drug resistance. Indeed, high doses of antimicrobials increase drug efficacy and thus reduce the length of treatment, allowing less time for resistance to emerge. But high drug doses also impose stronger selective pressures in favour of drug-resistant bacteria.

The potential for promoting drug resistance, however, is not the only issue when designing treatment protocols. It is known that broad-spectrum antibiotics can be associated with harmful side effects [31] and more importantly, can be responsible for disrupting an innate protective mechanism known as *colonisation resistance*, by shifting the balance between commensal microbiota and enteropathogenic bacteria in favour of the pathogens, increasing the risk of infections [34, 98, 100].

Here we address the following fundamental questions: how do we deploy a single, bacteriostatic antibiotic into a single-host in order to remove optimally the pathogen and its antibiotic resistant mutants? Can we design drug deployment protocols that support the

commensal microbiota and therefore promote colonisation resistance? To begin to address these questions, the purpose of this chapter is to use optimal control theory to show that the optimal drug deployment protocol is model-independent and a classical engineering solution, interpreted in this context as intermittent periods of maximum drug deployment.

4.1 A possible control-theoretic formulation

The task of constructing a biologically-realistic mathematical model describing the disease dynamics of a host organism is difficult and ongoing. Later in this thesis we will propose a model of a microbial microcosm that will allow us to design and evaluate the efficacy of different antibiotic usage strategies, but in the search of general results, first we shall work within a core modelling framework that comprises a large class of evolutionary models of bacteria growing under resource limitation and under the effect of an antimicrobial agent.

This abstract model contains n different commensal bacterial phenotypes competing for limited resources with m different types of rapidly-evolving pathogens. Let us represent with vectors C and P the densities of each phenotype of commensal and pathogenic bacteria respectively. If we denote with S the concentration of available resource and with A the concentration of antibiotic present in the environment, then we can represent the state of the system with the variable $\boldsymbol{x} = (S, C, P, A)$.

Therefore the dynamic equations that describe the growth and subsequent evolution of the state of the system will be given by $\dot{\boldsymbol{x}} = \mathcal{F}(\mathbf{x})$ suitably augmented with initial conditions. As our goal is to control the input of antibiotic in order to to drive the state of the system towards a pathogen-free state, let us write the system of equations in a more convenient form

$$\dot{\boldsymbol{x}} = \mathcal{F}(\boldsymbol{x}, \delta) = \mathcal{G}(\boldsymbol{x}) + \delta(t) A_0 \mathbf{e}, \qquad x(0) \in \mathbb{R}^{1+n+m+1} \tag{4.1}$$

where $\delta(t)$ represents our control variable, A_0 denotes an external supply of antibiotics and $\mathbf{e} = (0, 0, ..., 0, 1) \in \mathbb{R}^{1+n+m+1}$.

In order to ensure our abstract model possesses the gravity of the situation facing users of antibiotics, we quite purposefully impose an ominous set of restrictions on the physiological and evolutionary responses of the commensal and pathogen populations to the antibiotic. In particular, we impose the restriction that there are no environmental conditions under which the commensal has the highest growth rate. At low antibiotic concentrations the wild-type pathogen colonises the host, and at high antibiotic concentrations the resistant pathogens outcompetes every other bacterial phenotype. In this case we say that the pathogen population has *complete competitive advantage*, a property that can be described mathematically in the nonlinear function $\mathcal{G}(\boldsymbol{x})$.

4.1.1 Optimality of antibiotic pulsing

The question we would like to address is how can we design drug usage strategies to ensure the prevalence of the commensal population in such a severe situation? In answer to this, the nature of the control problem we wish to solve can be stated as follows.

Find a control strategy $\delta(t)$ on some time interval $[0, T]$ for $T \leq \infty$ such that pathogen density is reduced while growth of commensal bacteria is still supported. We could achieve this with a feedback stabilisation approach by writing $\delta = K\boldsymbol{x}$ and seeking a matrix K to render the pathogen-free steady-state of equation (4.1) stable, although real-time control decisions are not generally available to biological systems such as this. Instead we are going to use tools from optimal control theory [102] to determine the control δ that maximises a given payoff functional that illustrates our objective.

Of course, the precise details of the optimal control will depend on the objective functional we choose to consider. For instance, in the numerical examples presented later in this thesis we will use a saddle point functional that maximises the overall commensal population while minimising the density of pathogens, but we could define simpler objectives like for example minimising pathogen density in an experiment of duration T, or maximising the commensal population. In any case, the payoff functional would be linear with respect to both the control and state variables, so we can write it down in terms of a sign-indefinite

weight vector $w \in \mathbb{R}^{1+n+m+1}$, *ie.*

$$\mathcal{J}(\delta) = \int_0^T \langle w, \boldsymbol{x}(t) \rangle \, dt, \tag{4.2}$$

where angle brackets denote the standard inner product. Now, if we denote with the fixed constant $\delta_{max} \geq 0$ the maximum rate at which we can introduce antibiotics into the system, then the set of admissible controls is the set of measurable functions taking values almost everywhere between 0 and δ_{max},

$$\Omega = \{\delta \in L^\infty[0,T] \, : \, 0 \leq \delta(t) \leq \delta_{max} \text{ for almost all } t \in [0,T]\}.$$

Note that because we want to control an experiment of fixed duration T, then this control problem is referred to as a fixed time, free end-point problem, and therefore it is immediate that $\Omega \neq \emptyset$ for $T > 0$. Then our constrained optimisation problem is to determine an admissible control $\delta^*(t)$ such that

$$\mathcal{J}(\delta^*) = \max_{\delta \in \Omega} \mathcal{J}(\delta)$$

subject to the system of ordinary differential equations defined in (4.1).

In order to show that there exists an admissible control $\delta^* \in \Omega$ such that $\mathcal{J}(\delta^*) \geq \mathcal{J}(\delta)$ for all $\delta \in \Omega$, let us suppose that

$$\sup_{\delta \in \Omega} \mathcal{J}(\delta) = \lim_{k \to \infty} \mathcal{J}(\delta_k)$$

for $(\delta_k)_{k=1}^\infty$ a supremising sequence of admissible controls with associated responses $(\boldsymbol{x}_k)_{k=1}^\infty \in W^{1,\infty}([0,T])$.

Note that the objective functional $\mathcal{J} : L^\infty \to \mathbb{R}$ is continuous with respect to the weak* topology. Moreover Ω is a closed and convex subset of L^∞ and therefore Ω is compact with respect to the weak* topology on L^∞. Then because $0 \leq \delta_k \leq d$ almost everywhere, we can assume without loss of generality that $\delta_k \overset{*}{\rightharpoonup} \delta^*$ as $k \to \infty$.

Also, because the system is dissipative, $\boldsymbol{x}(t)$ is uniformly bounded and there $\exists \, M > 0$ such

that $|\boldsymbol{x}(t)| < M$ for all $0 \leq t \leq T$. Then we can obtain a $W^{1,\infty}([0,T])$ bound on the sequence $(\boldsymbol{x}_k)$ with a weak convergence $x_k \overset{*}{\rightharpoonup} x^*$ as $k \to \infty$. Furthermore,

$$\lim_{k\to\infty} \mathcal{J}(\delta_k) = \lim_{k\to\infty} \int_0^T \langle w, x_k \rangle \, d = \int_0^T \langle w, \lim_{k\to\infty} x_k \rangle dt = \int_0^T \langle w, x^* \rangle \, dt = \mathcal{J}(\delta^*).$$

Hence, $\mathcal{J}(\delta^*) = \sup\{\mathcal{J}(\delta) \, : \, \delta \in \Omega\}$ and therefore δ^* is optimal.

Now by using standard results from control theory we can synthesise the optimal control by forming control theory Hamiltonian

$$\begin{aligned} \mathcal{H}(\boldsymbol{x}, \boldsymbol{\lambda}, \delta) : &= \langle \boldsymbol{\lambda}, \mathcal{F}(\boldsymbol{x}, \delta) \rangle - \langle w, \boldsymbol{x} \rangle \\ &= \langle \boldsymbol{\lambda}, \mathcal{G}(\boldsymbol{x}) \rangle + \langle \boldsymbol{\lambda}, \delta A_0 \mathbf{e} \rangle - \langle w, \boldsymbol{x} \rangle \, , \end{aligned}$$

where $\boldsymbol{\lambda}$ is called the adjoint variable and satisfies the final-value problem

$$-\dot{\boldsymbol{\lambda}} = w + (\nabla_{\boldsymbol{x}} \mathcal{G}(\boldsymbol{x}))^T \boldsymbol{\lambda}, \qquad \boldsymbol{\lambda}(T) = 0. \tag{4.3}$$

Well-known results from control theory state that the Hamiltonian associated with the constrained optimisation problem determined by equations (4.1-4.3) is maximised at all times along the optimal solution $(\boldsymbol{x}^*, \boldsymbol{\lambda}^*, \delta^*)$,

$$\mathcal{H}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*, \delta^*) = \max_{\delta \in \Omega} \{\mathcal{H}(\boldsymbol{x}, \boldsymbol{\lambda}, \delta)\} \, .$$

Furthermore, as the Hamiltonian is linear with respect to the control variable $\delta(t)$, then the condition $\max\{\mathcal{H}(\boldsymbol{x}^*, \boldsymbol{\lambda}^*, \delta) \, | \, 0 \leq \delta \leq \delta_{max}\}$ holds when $\delta(t) = \delta_{max}$ and $\langle \boldsymbol{\lambda}^*, A_0 \, \mathbf{e} \rangle > 0$, or $\delta(t) = 0$ and $\langle \boldsymbol{\lambda}^*, A_0 \, \mathbf{e} \rangle < 0$, a control known as bang-bang.

In the context considered here, where a single antibiotic is deployed into a host, bang-bang controls can be interpreted as time intervals of maximum drug deployment with intervals where no drug is deployed. This antibiotic deployment strategy is what is described in the medical literature as an *intermittent* or *pulsing* dosage schedule [36, 53, 54]. However, if $\langle \boldsymbol{\lambda}^*, A_0 \, \mathbf{e} \rangle = 0$ then we say that the control is *singular* and in this case we cannot make any claims about the optimality of pulsating controls.

But we argue that pulsating protocols are, in general, superior to ones that define a fixed drug dose throughout an experiment of duration T. To show this property, let us define a partition of the interval $[0, T]$ as $\mathcal{I}_N = (t_i)_{i=1}^N$ such that

$$0 = t_0 < t_1 < \cdots < t_{N-1} < t_N = T,$$

and let $PC(\mathcal{I}_N)$ represent the space of piecewise constant functions on $\mathcal{I}_N$ taking values 0 or δ_{max} on each subinterval (t_i, t_{i-1}). Then the set of admissible pulsating protocols is defined as

$$\Omega_p = \{\delta \in L^\infty[0, T] \, : \exists N, \mathcal{I}_N \text{ and } d \in PC(\mathcal{I}_N) \text{ such that } \delta(t) = d(t)\}.$$

Because Ω_p is a weak* dense subset in Ω (see [103]), therefore

$$\max_{\delta \in \Omega} \mathcal{J}(\delta) = \sup_{\delta \in \Omega_p} \mathcal{J}(\delta).$$

4.2 Other therapeutically relevant controls

The solution of any such control problem will only be of limited interest if we cannot implement its solution in practise. Indeed, it must be stated that there is little expectation that an optimal control computed using optimisation algorithms will turn out to be optimal in real experimental systems. Moreover, we wish to make clear that we are not proposing to compute optimal pulsating schedules for therapeutic purposes, as that statement would be far beyond the scope of this thesis. Our only purpose when determining the optimal pulsating strategy is to provide us with a baseline that can be used to compare theoretically optimal treatments with other strategies easier to implement in clinical and experimental systems, like the ones described below.

4.2.1 Optimal duration of antibiotic therapy

The first question we could ask when designing drug usage strategies, is how much antibiotic should we use? Or if we fix the drug dosage, we then could ask what is the optimal treatment duration? Surprisingly, there is no consensus on the medical and pharmaceutical communities on how to answer these simple questions.

High-doses of antibiotic are believed to be more effective clearing infections, but they are also responsible for promoting the evolution of drug resistance, and thus decreasing drug efficacy for future treatments. For instance, long-term treatments, usually preferred in clinical settings, have been associated with unnecessary side effects [74] and with a considerable increase in treatment costs [81]. Furthermore, it has also been shown that patients that receive unnecessarily long treatments were significantly at more risk of acquiring nosocomial infections [82].

Interestingly, there is clinical evidence that shorter duration of antibiotic therapy has the same clinical effect than the current guidelines that prefer long-term high-dose treatments. For example, a 3-day antibiotic therapy does not lead to inferior clinical results when treating community acquired pneumonia than a standard 8-day treatment [38, 89]. Similarly, it has been shown that an 8-day course of antibiotic therapy for ventilator-associated pneumonia was equally effective as a 15-day course [23].

While in practise the optimal treatment duration may always remain elusive, computing the optimal stopping time in a theoretical model of drug deployment is a simple task. To achieve this, let us assume that we will deploy antibiotics at maximum rate δ_{max} during θ units of time, where $0 \leq \theta \leq T$. Then our control can be written as

$$\delta(t;\theta) = \begin{cases} \delta_{max} & \text{if} \quad t < \theta \\ 0 & \text{if} \quad t \geq \theta. \end{cases}$$

The optimal stopping time that we will denote θ^* can be computed by numerically max-

imising the objective functional defined in (4.2) for each value of $\theta \in [0, T]$, that is

$$\mathcal{J}(\delta(t; \theta^*)) = \max_{\theta \in [0,T]} \mathcal{J}(\delta(t; \theta)).$$

4.2.2 Tapering dosage strategies

There is clinical evidence that support *high-dose*, *short-course* therapies to treat bacterial infections like pneumonia [92], but it is also known that early cessation of antibiotic therapy can be associated with increased mortality in critically ill patients [85]. So instead of stopping treatment completely, it has been proposed that treatment should start with a high dose of antibiotics and gradually reduce drug dosage to decrease side-effects due to long-term drug treatment.

This strategy is called *drug tapering* and it has been reported to be an effective method to reduce recurrent episodes of *Clostridium difficile* disease attributed to the disruption of the intestinal microflora as well as to the persistence of non-vegetative spores [73, 101].

In general, tapered drug deployment protocols consist of a stepwise or a continuous decrease of drug dosage over a period of time. As the statement of our control problem allows us to continuously control the input of antibiotic into the system, then for the purpose of this paper we will consider the latter, and thus define this treatment protocol as follows:

- Start treatment with maximum dosage, a property expressed as $\delta(0) = \delta_{max}$.
- Gradually decrease drug input with a rate of descent given by a parameter $\alpha \in [0, \pi/2]$, so the control is defined as $\delta(t; \alpha) = \tan(\pi - \alpha)t + \delta_{max}$. Then $\alpha = 0$ corresponds to deploying antibiotic at maximum dose at all times, while $\alpha = \pi/2$ denotes a situation where no antibiotic is used at all.
- Let $t = t_\alpha$ be the first moment in time when $\delta(t) = 0$. Do not deploy antibiotics afterwards, that is $t_\alpha \leq t \leq T$.

Of course, every rate of descent α will produce different controls $\delta(t;\alpha)$ with corresponding payoffs $\mathcal{J}(\delta(t;\alpha))$. We will call the optimal tapering parameter, α^*, the parameter such that $\mathcal{J}(\delta(t;\alpha^*)) = \max \mathcal{J}(\delta(t;\alpha))$ for all $\alpha \in [0, \pi/2]$.

4.2.3 Adaptive pulsing

It is important to notice that in order to precisely determine the optimal pulsing treatment or even the optimal tapering strategy, one would need a calibrated model with perfect understanding of the future. But is it possible to determine effective treatment strategies, perhaps suboptimal, based only on current and past patterns of susceptibility and resistance? In response to this we propose the development of *feedback controls* that take information directly from current observations to inform and adjust future treatment policies.

The difficulty with this approach is the appropriate design of feedback rules. While it is theoretically possible to use linear feedback laws to stabilise the pathogen-free equilibrium of equation (6.2), we prefer to consider feedbacks that might be implemented in a practical, empirical setting.

In the following chapter we will illustrate using numerical examples that we can construct antibiotic deployment strategies based on feedback controls that are effective in supporting the commensal population. These controls have been proposed as *adaptive management strategies* in ecological settings, and are predicated on the very simple maxim, *if antibiotic resistance is high, cease its use immediately*. As we are considering the scenario where we only have one available antibiotic this must mean that controls will oscillate between periods of deployment and non-deployment, and thus we call them *adaptive pulsing protocols*.

Note that feedback controls may be costly in the sense that in order to be implemented, it may be necessary to have complete information about the state of the system at all instants in time, represented in this case as having detailed genotypic information of every bacteria present in the system. While it may be possible to have rapid genotypic identification of resistance genes [11], these tests are not yet widely available in clinical settings. As a result we seek feedbacks based on much coarser information, for instance by making observations

at discrete times every ϑ units apart. Then the control $\delta(t)$ during the time interval $[\vartheta_i, \vartheta_{i+1}]$ would be determined in response to observations on the state of the system at time $t = \vartheta_i$ by setting $\delta(t) = \overline{\delta}(x(\vartheta_i))$ for some suitably-chosen feedback law defined within the function or rule $\overline{\delta}$.

4.3 Colonisation resistance

A dense microbial community inhabits the intestine, so for invasive pathogens to colonise the gut they first have to outcompete the commensal bacteria. This protective mechanism is called colonisation resistance [110], and although the molecular basis of the interactions between host, pathogens and microbiota remains largely unknown, it is believed that pathogenic bacteria have developed different strategies to overcome colonisation resistance. For example by triggering the host's immune response and in consequence changing the population structure of the gut ecosystem, allowing the pathogens to colonise the host [98].

It is known as well that broad-spectrum antibiotic treatment can decrease the overall microbial population in the gut considerably. Therefore a side effect of chemotherapy can be to disrupt colonisation resistance by shifting the balance between microbiota and enteropathogenic bacteria in favour of the pathogens, increasing the risk for gut infections [99, 100].

Furthermore, it has been reported that although there is day-to-day variability in the diversity patterns of the human microbiota, the average community composition remains stable over several months in the absence of perturbations. However, the effect of antibiotics on the gut microbiota is profound and very rapid, with a loss of diversity and a shift in the community composition occurring a few days after drug treatment is initiated. Although the microbiota eventually returns to a stable state after the end of the treatment course, it often shifts to a different stable average state [34]. The health consequences of this new ecological equilibrium are fundamentally unknown [33], but there is evidence that suggests that disrupting these co-evolved microbial communities can be associated with health

problems like inflammatory bowel disease.

This counterproductive outcome of antibiotic therapy highlights the necessity of studying the complex interactions between microbial evolution and human disease, and in particular the ecological and evolutionary implications of antimicrobial treatment within-host dynamics. For this reason, the main objective of the mathematical models discussed in the remaining of this thesis is to design antibiotic deployment protocols that support the commensal microbiota while driving the pathogenic population to extinction.

Chapter 5

Modelling bacteriostatic antibiotics

In order to compare the efficacy of the theoretical optimal control with the practicable treatment strategies discussed in the previous chapter, here we will pose a mathematical model of a proxy for a host organism that takes into consideration the competition for resources between different bacterial subpopulations and under the effect of a bacteriostatic antimicrobial agent. This model is based on the mass action law and the inhibitory effect by the antimicrobial agent is constructed from kinetic interactions between drug molecules and their targets.

5.1 The Growth Inhibition Coefficient $\gamma(A)$

Antibiotics are termed *bactericidal* if they kill bacteria, for example by inhibiting cell wall biosynthesis and therefore enhancing the likelihood of cell *lysis*, or *bacteriostatic* if they suppress bacterial growth rate by inhibiting DNA replication, RNA transcription or by interfering with protein synthesis and other aspects of cellular metabolism. For example, the bacteriostatic antibiotic *Rifampicin* inhibits the function of RNA polymerase during transcription by binding to its β subunit [107], while *Streptomycin* targets the 30S ribosomal subunit and inhibits translation [83]. Resistance to antibiotics can therefore arise through mutations that alter the structure of the protein targeted by the drug, although such muta-

tions may come at a cost of reduced fitness in antibiotic-free environments [6, 43, 60].

With *rifampicin* in mind and following [4] we continue by developing a mathematical model of the inhibition of transcription by an RNA polymerase-binding antibiotic with the following kinetic model. In this model σ denotes the concentration of free sigma factors in the cell that promote transcription and P represents the concentration of free RNA polymerase that must bind to a gene's promoter region to initiate transcription of DNA into mRNA, so that P_σ is the concentration of promoters bound to polymerase. The variable R will denote the concentration of unbound gene promoter regions so that P_σ^R is a complex that can initiate transcription of mRNA, the latter at a concentration denoted M. The variable A represents the concentration of an antibiotic molecule that binds to P to form a non-functional polymerase-antibiotic complex with concentration denoted P_A. We describe this process as follows

$$\begin{aligned}
&\sigma + P \underset{k_{-1}}{\overset{k_1}{\rightleftharpoons}} P_\sigma, \\
&P + A \underset{k_{-2}}{\overset{k_2}{\rightleftharpoons}} P_A, \\
&\sigma + P_A \underset{k_{-3}}{\overset{k_3}{\rightleftharpoons}} P_\sigma^A, \\
&P_\sigma + R \underset{k_{-4}}{\overset{k_4}{\rightleftharpoons}} P_\sigma^R \xrightarrow{k_5} R + P_\sigma + M.
\end{aligned}$$

A consequence of producing the inhibited polymerase complex P_A is a reduction in the availability of free polymerase which then reduces the rate of RNA production, where the latter is given by $k_5 P_\sigma^R$ as P_σ^R denotes the fraction of gene promoters bound to a sigma factor-polymerase complex.

$$\frac{d}{d\tau}\sigma = -k_1\sigma \cdot P + k_{-1}P_\sigma - k_3\sigma \cdot P_A + k_{-3}P_\sigma^A, \tag{5.1a}$$
$$\frac{d}{d\tau}A = k_{-2}P_A - k_2P \cdot A, \tag{5.1b}$$
$$\frac{d}{d\tau}P = -k_1\sigma \cdot P + k_{-1}P_\sigma - k_2A \cdot P + k_{-2}P_A, \tag{5.1c}$$
$$\frac{d}{d\tau}P_A = k_2P \cdot A - k_{-2}P_A - k_3P_A \cdot \sigma + k_{-3}P_\sigma^A, \tag{5.1d}$$
$$\frac{d}{d\tau}P_\sigma = k_1\sigma \cdot P - k_{-1}P_\sigma + (k_5 + k_{-4})P_\sigma^R - k_4P_\sigma \cdot R, \tag{5.1e}$$
$$\frac{d}{d\tau}P_\sigma^A = k_3P_A \cdot \sigma - k_{-3}P_\sigma^A, \tag{5.1f}$$
$$\frac{d}{d\tau}P_\sigma^R = k_4P_\sigma \cdot R - (k_5 + k_{-4})P_\sigma^R, \tag{5.1g}$$
$$\frac{d}{d\tau}R = -k_4P_\sigma \cdot R + (k_5 + k_{-4})P_\sigma^R, \tag{5.1h}$$
$$\frac{d}{d\tau}M = k_5P_\sigma^R. \tag{5.1i}$$

For the remainder of this thesis the timescale pertinent to cell growth will be denoted with t, so the variable τ denotes a fast timescale at which the production of a single transcript happens. Note that the total concentration of bound and unbound polymerase does not change over time and so we may define the following constant

$$P_{tot} := P + P_A + P_\sigma + P_\sigma^A + P_\sigma^R. \tag{5.2}$$

Now let us assume that the production of RNA is a much slower process that the binding and unbinding of other complexes. Therefore we can make the following quasi-steady-state assumption:

$$\frac{d}{d\tau}R = \frac{d}{d\tau}P_\sigma^R = \frac{d}{d\tau}P_\sigma^A = \frac{d}{d\tau}P_A = \frac{d}{d\tau}P_\sigma = 0. \tag{5.3}$$

Then from equation (5.1b) we deduce

$$P_A = \frac{k_2P \cdot A}{k_{-2}} \tag{5.4}$$

Similarly from (5.1f) and (5.1g),

$$P_\sigma^A = \frac{k_3 P_\sigma \cdot A}{k_{-3}} \tag{5.5}$$

$$P_\sigma^R = \frac{k_4 P_\sigma \cdot R}{(k_5 + k_{-4})} \tag{5.6}$$

Also from (5.1e)

$$P_\sigma(k_{-1} + k_4 R) = k_1 \sigma \cdot P + (k_{-4} + k_5) P_\sigma^R$$

and using (5.6)

$$P_\sigma(k_{-1} + k_4 R) = k_1 \sigma \cdot P + k_4 P_\sigma \cdot R.$$

Subtracting $k_4 P_\sigma \cdot R$ from both sides:

$$P_\sigma = \frac{k_1 \sigma P}{k_{-1}} \tag{5.7}$$

Finally, defining $\kappa_1 = k_1/k_{-1}$, $\kappa_2 = k_2/k_{-2}$ and $\kappa_3 = k_3/k_{-3}$ and using (5.2) we derive

$$\begin{aligned} P_{tot} &= P + P_A + P_\sigma + P_\sigma^A + P_\sigma^R \\ &= P + \kappa_2 P \cdot A + \kappa_1 P \cdot \sigma + \kappa_1 \kappa_3 P \cdot \sigma \cdot A + \frac{\kappa_1 k_4}{k_{-4} + k_5} P \cdot \sigma \cdot R \end{aligned}$$

and the rate or *velocity* of mRNA transcription, $\frac{d}{d\tau} M$, is given by $v := k_5 P_\sigma^R$ and so we define

$$v = v(A) := k_5 P_\sigma^R = \frac{\kappa_1 k_4 k_5}{k_5 + k_{-4}} \cdot \frac{R \cdot \sigma \cdot P_{tot}}{1 + \kappa_1 \sigma + \kappa_{145} R \cdot \sigma + A(\kappa_2 + \kappa_1 \kappa_3 \sigma)}, \tag{5.8}$$

where $\kappa_{145} := \kappa_1 k_4 / (k_5 + k_{-4})$.

Note that the rate of RNA production, $v(A)$, in (5.8) tends to zero as $A \to \infty$ leading the cell ever-closer to *complete inhibition*. Then the rate of production of RNA in the absence

of antibiotic is given by

$$v(0) = k_5 P_\sigma^R = \frac{\kappa_1 k_4 k_5}{k_5 + k_{-4}} \cdot \frac{R \cdot \sigma \cdot P_{tot}}{1 + \kappa_1 \sigma + \kappa_{145} R \cdot \sigma}$$

and therefore the relative decrease in transcription rate due to the presence of antibiotic takes the form

$$\frac{v(A)}{v(0)} = \frac{1}{1 + A\left(\frac{\kappa_2 + \kappa_1 \kappa_3 \sigma}{1 + \kappa_1 \sigma + \kappa_{145} R \cdot \sigma}\right)} =: \frac{1}{1 + \kappa A}, \tag{5.9}$$

for some parameter $\kappa = \kappa(\sigma, R) \geq 0$ that can be thought of as the phenotype of each cell. Note that κ is a single-cell measure of antibiotic efficacy in the sense that if A_{50} is the antibiotic concentration required to halve the transcription rate, then $A_{50} = 1/\kappa$.

We used the one-parameter freedom in (5.9) to fit the binding affinity data of *rifampicin* and the related antibiotic molecule *rifabutin* to RNA polymerase using data taken from [115]. The result is shown in Figure 5.1 where the R^2 values of the least-squares fitting procedure are given in the legend of the figure. While this gives good agreement, it does not demonstrate that transcription rate decreases with antibiotic concentration in the same manner as (5.9), but is consistent with this hypothesis.

While this is a vast over-simplification of the true molecular biology that does not include features like RNA degradation, for example, this model will at least provide us with some broad insights into how an antibiotic like rifampicin slows transcription and so inhibits cell growth. We could go further in the model and include the fact that rifampicin-bound polymerase can bind to gene promoters to produce short RNA oligomers (trimers) [115]. This more complex model, however, leads to a form for $\gamma(A)$ that closely resembles the one presented above in (5.9) but whose derivation is somewhat more lengthy, so we have omitted it for the sake of brevity.

In order to determine the growth rate of each cell phenotype we shall assume that the concentration of antibiotic in the environment and inside the cell are the same and that

$$\text{cell growth rate} = \text{constant} \times \text{resource uptake rate} \times \text{transcription rate}. \tag{5.10}$$

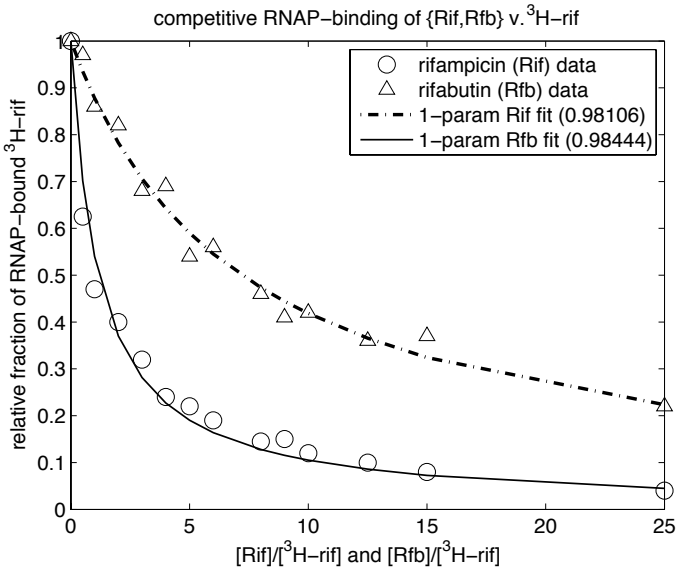

Figure 5.1: Effects of rifampicin (Rif) and rifabutin (Rfb) on the binding of ^3H rifampin to wild-type RNAP. This data was obtained using the standard method of competing each antibiotic against a radioactively marked Rif molecule ^3H-rif with a fixed concentration of 250nM and a reduction in the bound radioactive molecule is assumed to be due the binding of the non-radioactive drug. (Data taken from [115, Figure 3] where no error bars are given, R^2 values are provided in the figure legend.) This illustrates that the fraction of Rif and Rfb-free RNAP as a function of the concentration of each drug follows a curve of the form $1/(1 + p \cdot A)$ where A is the drug concentration and p is a parameter used in the datafit.

Hence the growth rate of each bacterial type is not only determined by the resource availability that we shall denote by S, or $S(t)$ at time t, but also by the concentration of the antibiotic present in the environment that we denote by A and $A(t)$ respectively. If we do not consider stress responses like down regulation of transcription at low-resource environments, then bacterial growth rate can be written as a standard monotonic and saturating Monod function in S multiplied by an antibiotic-dependent conversion constant $\boldsymbol{c}(A)$ that converts units of the limiting resource to biomass and describes the efficiency of cell production per unit resource. Thus the per-cell, per-unit time growth rate can be written in the form

$$G(S, A) = \frac{V_{\max} S}{K + S} \cdot \boldsymbol{c}(A)$$

where $V_{\max}$ represents a maximal resource uptake rate and K is a half-saturation constant also known as the cell's *affinity* for the limiting resource.

Working under the assumption that (5.10) holds for each cell then, if the antibiotic molecule has no effect on the uptake of the limiting resource so that $V_{\max}$ and K are independent of A, it follows that the antibiotic must reduce the efficiency of the cell in this simple scenario. Indeed, in order to compensate for a reduction in growth rate in the presence of antibiotic, a cell could produce more polymerase but this has the net effect of increasing the resources needed to create a single cell and so leads to a decrease in cell efficiency. Therefore using (5.9) and (5.10) we can write $\boldsymbol{c}(A)$ as a product of the cell conversion rate in an antibiotic-free environment $\boldsymbol{c} := \boldsymbol{c}(0)$ and a dimensionless inhibition coefficient, $\gamma(A)$, such that

$$\gamma(A) := \frac{G(S,A)}{G(S,0)} = \frac{\boldsymbol{c}(A)}{\boldsymbol{c}} = \frac{v(A)}{v(0)} = \frac{1}{1+\kappa A}. \tag{5.11}$$

We then can model a population of bacteria growing under resource limitation and under the effect of a bacteriostatic antibiotic with the following set of differential equations

$$\dot{S} = u(S)B, \tag{5.12a}$$

$$\dot{B} = G(S,A)B, \tag{5.12b}$$

$$\dot{A} = -aA\,B, \tag{5.12c}$$

with initial conditions $\boldsymbol{x}(0) = (S(0), B(0), A(0))$ where $S(t)$ is the concentration of resource in the environment at time t, $A(t)$ the concentration of antibiotic and $B(t)$ the density bacteria. $u(S)$ represents a resource uptake function defined as

$$u(S) = \frac{V_{\max} S}{K+S} \tag{5.13}$$

and therefore $G(S,A) = c\,u(S)\gamma(A)$. Also, c and a are constants that denote conversion rates for resource and antibiotic respectively.

In summary, although different antibiotic classes may have different mechanisms of action, in some cases it may be possible to describe their inhibitory effect on bacterial growth by a

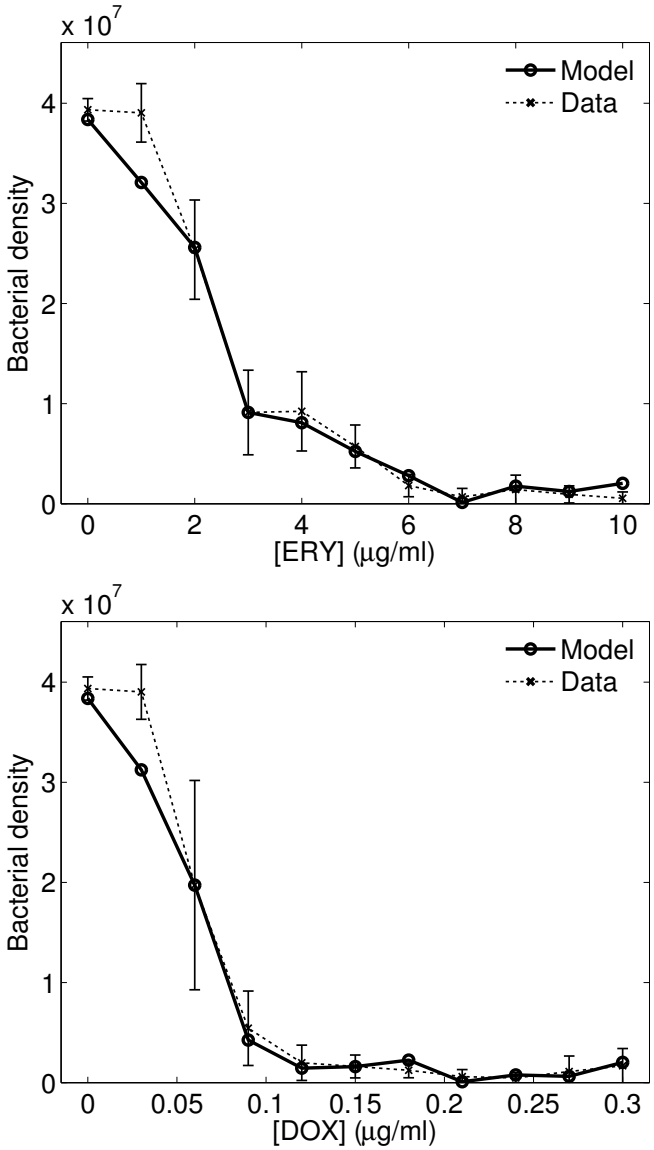

Figure 5.2: Dose-response curves of *Escherichia coli* exposed to different antibiotic concentrations. Dotted lines represent experimental data with standard error bars and solid lines predictions of the model described by equations (5.12a-c) with growth parameters: $V_{max} = 0.000251\mu$g / cell / h, $K = 0.62$ μg/ml, $c = 1.8 \cdot 10^4$ cells/μg. (left) *Erythromycin* is a macrolide that binds to the 50s ribosomal subunit inhibiting protein synthesis, with inhibition parameters given by $\kappa_1 = 3.46$ ml/μg and $\kappa_2 = 0.079$ ml/μg. The fit has an $R^2 = 0.973$. (right) *Doxycycline* a tetracycline that binds to the 30S subunit and inhibits binding of aminoacyl-t-RNA to the acceptor site on the 70S ribosome. Fitted inhibition parameters are $\kappa_1 = 0.068$ ml/μg and $\kappa_2 = 0.198$ ml/μg with an $R^2 = 0.966$. Numerical fit performed using a constrained optimisation algorithm and experimental methods described in Appendix D.

growth inhibition function that satisfies the following properties:

I1. $\gamma(0) = 1$,

I2. $\gamma(A) \geq 0$ for all $A \geq 0$ and

I3. $\gamma'(A) \leq 0$ for all $A \geq 0$

For the numerical examples presented in this thesis we shall extend the definition of the function in (7.5) by writing $\gamma(A) = (1 + \kappa A)^{-1} = 1 - \kappa A(1 + \kappa A)^{-1}$ and we shall permit a two-parameter dependence of the form

$$\overline{\gamma}(A) = 1 - \frac{\kappa_1 A}{1 + \kappa_2 A}.$$

In this case $\lim_{A\to\infty} \overline{\gamma}(A) = 1 - \kappa_1/\kappa_2$ and so, provided $\kappa_1 < \kappa_2$, $\overline{\gamma}(A)$ represents a growth inhibition function that cannot lead to complete inhibition for finite A.

We used this two-parameter inhibition function to model the effect of two bacteriostatic antibiotics that inhibit protein synthesis: *erythromycin* that binds to the 50s subunit of the bacterial 70s rRNA complex and *doxycycline* that binds to the 30S ribosomal subunit. Notably, in both cases we were able to predict the dose-response curve of an isogenic population of *E. coli* growing under resource limitation and exposed to various concentrations of each antibiotic, as shown in Figure 5.2.

5.2 Evolutionary dynamics

In a clinical context the effect that the evolution of antimicrobial resistance has on the efficacy of treatment protocols is fundamental. For example, consider the case of a lethal infection by enteropathogenic bacteria, by this we mean that the invasive pathogen has a higher growth rate than the commensals in an antibiotic-free environment so they will ultimately colonise the host. But at high antibiotic concentrations, pathogenic bacteria are outcompeted by commensals because the latter are less susceptible to the antibiotic. It is clear from this example that in order to drive the pathogens to extinction, the best treatment

protocol is to deploy as much antibiotic as soon as possible. But what happens if we consider an evolutionary system where the invasive pathogen can evolve resistance to the antibiotic used? Can we design drug usage strategies that minimise conditions that promote the evolution of resistant pathogens? In order to address these questions and to understand the evolutionary implications of antimicrobial therapy, for the remainder of this thesis we shall consider that pathogenic bacteria can evolve resistance mechanisms.

There are several genetic and biochemical mechanisms that may confer antimicrobial resistance upon bacteria [5]: *i)* by changing the structure of target molecules to prevent antibiotics from binding, *ii)* by inactivating antibiotics through enzymatic degradation, *iii)* by excluding the entry of antibiotics through the cell wall or *iv)* by actively effluxing them from the cell. These resistance mechanisms are illustrated in Figure 5.3.

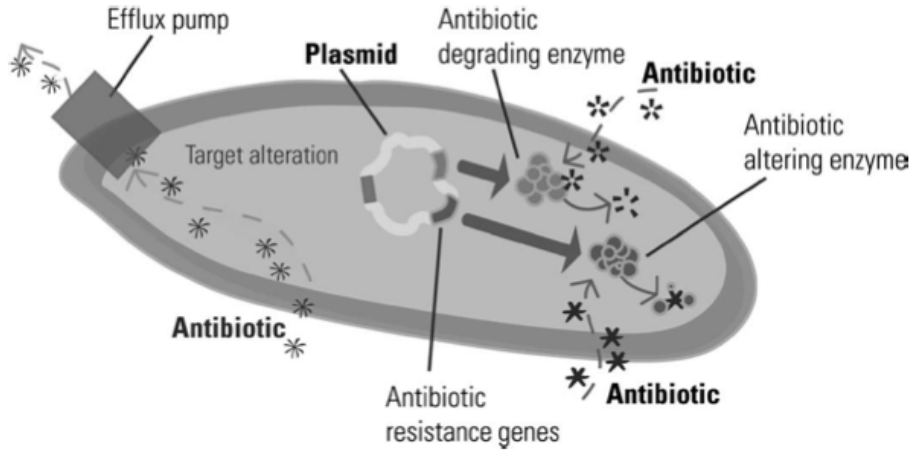

Figure 5.3: Resistance mechanisms can be acquired by structurally altering target molecules to prevent antibiotic binding, excluding antibiotics from cell entry, inactivation of antibiotics through enzymatic degradation, or pumped out of the cell. Figure reproduced from [5].

Interestingly, the cellular targets of antimicrobial agents are of a very limited nature [120]. While some drugs, like beta-lactamases, attack bacteria by inhibiting cell wall biosynthesis and causing lysis, other substances inhibit DNA replication, RNA synthesis or protein synthesis by binding to the ribosome or the RNA polymerase. This limitation on the modes of action of most antimicrobial substances allow bacteria to evolve resistance mechanisms through a single-point mutation or the accumulation of multiple mutations. In some cases

mutations can confer resistance to several antibiotics simultaneously, what is known as cross-resistance.

Although it is known that horizontal gene transfer, within and between species, play an important role in the prevalence and spread of resistance mechanisms, for the purpose of this thesis we will consider that resistance mechanisms are acquired vertically, through a phenotypic mutation occurring with probability $\epsilon > 0$ during cell division.

If we denote with the vector $\boldsymbol{P} = (P_1, P_2, \ldots, P_n)$ the densities of n different bacterial types, then the rate of mutation from bacterium type j into type i can be represented by the ij-th entry of a mutation matrix M. Hence the evolutionary dynamics of the system is described by the stochastic mutation matrix

$$\mathcal{M}_\epsilon = (1-\epsilon)I + \epsilon M,$$

where I is the identity matrix representing clonal reproduction, $0 < \epsilon \ll 1$ the probability of point mutation and M the n-by-n matrix where the m_{ij} entry represents the mutation rate of bacterial type i into phenotype j.

Note that if $\mathbf{1}$ denotes a vector of ones of the appropriate dimension, $\mathbf{1} = (1, 1, ..., 1)$, then

$$\mathbf{1}^T \mathcal{M}_\epsilon = \mathbf{1} \text{ and } \mathbf{1}^T M = \mathbf{1}$$

will all be assumed throughout and $\mathcal{M}_\epsilon$ will be assumed to be irreducible in the sense that there is a number p such that no entry of $\mathcal{M}_\epsilon^p$ is zero.

5.3 Competitive advantage and fitness cost of resistance

The derivation of the previous section allows us to model mutations that may occur when a cell divides and that confer antibiotic resistance. However, in the case of rifampicin this mutation entails a configurational change to the β-subunit of the RNA polymerase complex, so a mutation that proves beneficial in the presence of antibiotic may well have an associated cost at low concentrations of that antibiotic. In order to describe this situation

mathematically we suppose that a bacterium, the wild-type say, has the growth rate (in an environment (S, A) of limiting resource concentration S and antibiotic concentration A)

$$G^{wt}(S, A) = \boldsymbol{c}^{wt} \cdot \frac{V^{wt}_{\max} S}{K^{wt} + S} \cdot \gamma^{wt}(A),$$

and that a mutant bacterium has the growth rate

$$G^{mut}(S, A) = \boldsymbol{c}^{mut} \cdot \frac{V^{mut}_{\max} S}{K^{mut} + S} \cdot \gamma^{mut}(A).$$

Antibiotic resistance of the mutant phenotype is then expressed through the property that

$$\lim_{A\to\infty} G^{wt}(S, A) < \lim_{A\to\infty} G^{mut}(S, A),$$

but the *cost* of that resistance mutation is incurred through the property that

$$G^{wt}(S, 0) > G^{mut}(S, 0),$$

meaning that a reduction in bacterial fitness occurs in an antibiotic-free environment. We will call such a wild-type bacterium *antibiotic susceptible* and the mutant will be called *antibiotic resistant*, even though strictly speaking both types are susceptible. Then the patterns of susceptibility and resistance of each subpopulation of bacteria are contained within the growth inhibition curves, like the ones described in Figure 5.4.

Now suppose that both previously discussed bacterial types, drug resistant and drug susceptible, are pathogens and that there is third bacterium of an entirely different species that we deem the *commensal*, suppose also that its growth response is given by the function $G^{com}(S, A)$. For the purpose of this paper we are considering that pathogens are fitter than commensals, an assumption based on the idea that commensal microbiota would be doomed to extinction if no antibiotic treatment were given to the host or if the antibiotic were over-deployed. In our model we say that the pathogen population has *complete competitive advantage* if the following condition holds for all $A, S \geq 0$:

$$G^{com}(S, A) < \max\left\{G^{wt}(S, A), G^{mut}(S, A)\right\}$$

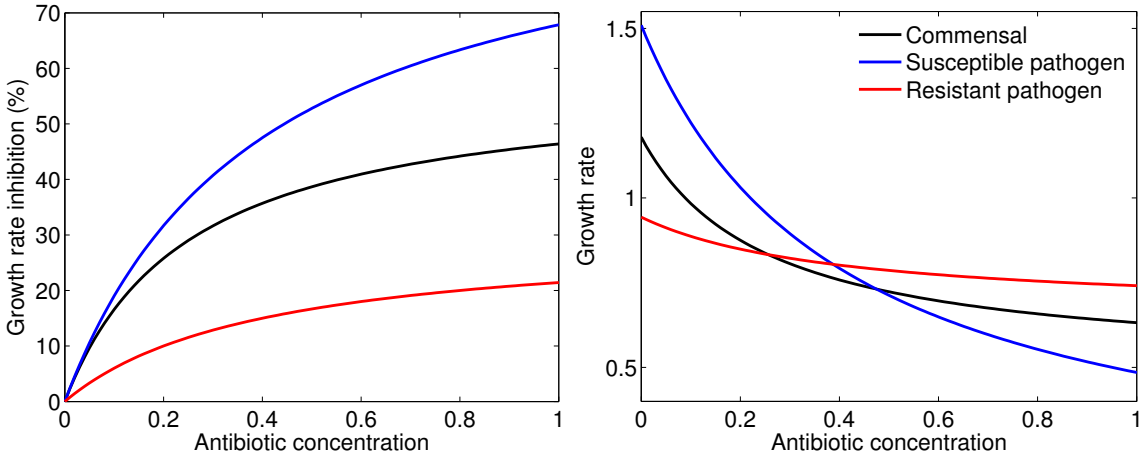

Figure 5.4: (a) Percentage of growth rate inhibition as a function of antibiotic concentration. (b) Growth rate as a function of antibiotic concentration in the environment. Here we can see that the resistant strain has a higher growth rate when the antibiotic concentration is high, while the fitness cost associated with resistance is expressed by having a lower growth rate when the antibiotic concentration is low. Note that for any concentration of antibiotic, commensal bacteria are outcompeted by at least one type of pathogenic bacteria.

In other words, in all abiotic environments and at all antibiotic concentrations the pathogen has a phenotype that is fitter than the commensal bacterium and in such a scenario we might reasonably expect the pathogen to out-compete the commensal for resources irrespective of how an antibiotic is deployed, a situation illustrated in Figure 5.4. In the following chapter we will use optimal control theory and optimisation algorithms to show that we can support the commensal microbiota in such dire circumstances by appropriately fluctuating the concentration of drug in the environment.

Chapter 6

Designing effective single-drug deployment protocols

Using optimal control theory as the basic theoretical tool, the purpose of this chapter is to investigate the efficacy of different antibiotic treatment protocols in the direst of circumstances described as follows. The commensal bacteria resident in an experimental proxy of a host organism compete for resources with a rapidly evolving and fitter pathogen, the latter so-named because the commensals would be doomed to extinction if no antibiotic treatment were given to the host or if the antibiotic were over-deployed.

As it is difficult to study pharmacodynamics and pharmacokinetics in a mammalian host in the presence of a pathogen and in real time, although perhaps not impossible with advances in imaging techniques [118], the experimental microbial system in which we are going to study antibiotic deployment strategies are continuous culture devices. In particular, chemostats provide ideal experimental systems to test different patterns of drug use because they can be mechanically adapted to control the supply of different antibiotics dynamically throughout the course of an experiment of duration T hours.

In this chapter we will show that it is possible to design optimal drug deployment protocols that support the commensal microbiota and drive the pathogens to extinction by appropriately fluctuating the concentration of drug in the environment. Furthermore, we will discuss

the efficacy of different antibiotic usage strategies and show that dynamically changing antimicrobial therapies may be effective in clearing a bacterial infection even when every 'static' monotherapy fails.

6.1 Chemostat model

Bioreactors have been used for decades to model the complex ecological interactions between hosts and pathogens in a controlled environment. In particular, continuous culture devices like chemostats offer the possibility to grow bacteria with a constant input of resource, maintaining bacterial growth rate at steady state. We could also adapt a second supply vessel to the chemostat to dynamically control the input of antibiotics into the system while maintaining resource input constant, as illustrated in Figure 6.1.

In a three week chemostat experiment, roughly 200 generations of bacteria can be grown, so it is possible to study the evolution of microbes at a large time scale. Also, the concentration of bacterium per unit of volume is so high and uniformly distributed, that it is possible to describe the ecological interactions between different bacterial populations with deterministic models based on the mass action law as the one described below.

Let us denote with $S(t)$ the concentration of limiting resource in the chemostat at time t, $A(t)$ the concentration of antibiotic, $C(t)$ a variable representing the density of commensal bacteria and $\boldsymbol{P}(t)$ a vector of n pathogenic types. Note that we are imposing the condition that commensal bacteria are isogenic and immutable. The reason behind this strong biological assumption is that we are considering that commensal bacteria are composed of many different species, and while resistance might be able to evolve, in some of these species it would not. By making the assumption that commensal bacteria are immutable we are saying that the commensal community as a whole will behave as if it is mostly composed of susceptible bacteria.

Then the state of the system can be represented by $\boldsymbol{x}(t) := (S(t), C(t), \boldsymbol{P}(t), A(t)) \in \mathbb{R}^{1+1+n+1}$. Our goal is to find antibiotic deployment strategies that minimise pathogen densities while still supporting the commensal population, so we shall utilise the following

saddle-point objective functional

$$\mathcal{J}(\delta) := \int_0^T \left(C(t) - \sum_{i=1}^{n} P_i(t) \right) dt. \tag{6.1}$$

Note that this payoff functional can be written in the form $\mathcal{J}(\delta) = \int_0^T \langle w, \boldsymbol{x}(t) \rangle dt$, where w is a vector of the sign-indefinite form

$$w = (0, 1, -1, -1, ..., -1, 0) \in \mathbb{R}^{1+1+n+1}.$$

Then our control problem would be to find a δ^* that maximises (6.1) subject to the following equations that describe the evolutionary and ecological dynamics of the system:

$$\dot{S} = d(S_0 - S) - u_c(S) \cdot C - \langle u_p(S), P \rangle \,, \tag{6.2a}$$

$$\dot{C} = G_c(S, A) \cdot C - dC, \tag{6.2b}$$

$$\dot{P} = \mathcal{M}_\epsilon \left(G_p(S, A) \cdot P \right) - dP, \tag{6.2c}$$

$$\dot{A} = \delta A_0 - dA - A(a_c C + \langle \mathbf{1}, a_p P \rangle), \tag{6.2d}$$

with initial conditions $\boldsymbol{x}(0) = (S(0), C(0), P(0), A(0))$ where $A(0) = 0$ and $S(0) = 0$ are the initial concentration of antibiotic and limiting, abiotic resource in the chemostat. The parameter S_0 is the supply concentration of abiotic resource and d the washout rate of the chemostat. The maximum input of antibiotic is bounded by the chemostat's dilution rate, and therefore for the remainder of this paper we will consider $\delta_{max} = d$.

Functions $u_c(S)$ and $u_p(S)$ denote uptake rates of the limiting resource by commensal and pathogenic bacteria, respectively, growth rates $G_p^*(S, A)$ are vectors whose entries have the form $c^* \cdot \gamma^*(A) \cdot u_p^*(S)$ where a superscript asterisk (*) here represents a placeholder for any of the pathogenic bacterial phenotypes. The symbol $\mathbf{1}$ denotes a vector of ones, $(1, 1, ..., 1)$ of the appropriate dimension and brackets standard inner products. Parameters a_c and a_p are the binding rates of commensal and pathogen to the antibiotic molecule and A_0 is the concentration of antibiotic held in a second supply vessel, as shown in Figure 6.1.

The irreducible matrix $\mathcal{M}_\epsilon$ represent phenotypic mutation processes, namely changes in

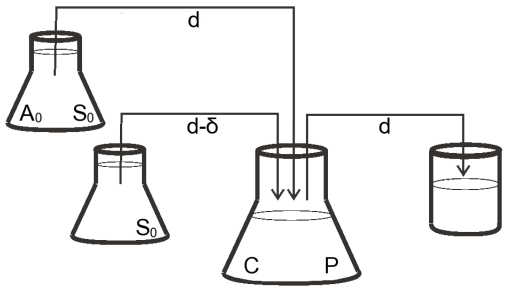

Figure 6.1: Diagram of a chemostat with dilution rate d and nutrient input S_0 adapted to supply antibiotic A_0 at a rate δ and to maintain the volume and resource input rate constant.

phenotype that arise due to errors that occur during cell division and these matrices can be further decomposed into the form

$$\mathcal{M}_\epsilon = I + \epsilon(M - I)$$

where the rate of mutations is the same for all phenotypes and equals ϵ. The matrix $M = (m_{ji})$, where i and j represent labels of two different pathogenic phenotypes, contains the probabilities that the phenotype of an offspring cell is i, given that the phenotype of the parent cell is j at division and a mutation occurs that changes the phenotype.

The dynamics of (6.2a-d) could be complex. Note that when $\epsilon = 0, A(0) = 0$ and $A_0 = 0$, a situation that represents an absence of mutations and with no antibiotic agent present in the environment, equation (6.2) obeys a competitive exclusion principle in the sense that only one type of either commensal or pathogenic bacteria will persist. Moreover, if the pathogen has complete competitive advantage over the commensal, one of the pathogenic phenotypes will go to fixation in the chemostat as time increases. However, if one forces ϵ to take on some positive but small value, we can deduce from the implicit function that a steady-state $O(\epsilon)$-close to the competitive exclusion state will then be locally stable and support the long-term dynamics of this system; one might say that this new steady-state is in mutation-selection balance. Global stability of this balanced state is not easy to prove, however, and indeed it is known that competitive exclusion, or indeed a small perturbation of it, is not necessarily expected in the presence of a growth inhibitor due to the possible existence of cycles in population densities [95]. The article [114] demonstrates that a model with a

form of inhibition closely related to the one used in (6.2) obeys the competitive exclusion principle.

Interestingly, the model defined by equations (6.2) is able to capture some of the population-level effects of the constant deployment of antibiotic at different concentrations. For example, it is observed experimentally that there exists a critical antibiotic concentration such that the population of bacteria decreases suddenly, a feature illustrated in Figure 6.2c. This threshold is known as the *Minimum Inhibitory Concentration* (MIC) and it is used as a quantitative pharmacodynamic parameter that characterises every antibiotic [64, 108]. This model shows that the apparent MICs can be explained as a population-level effect consequence of the experimental setup (length of the experiment or initial densities of each strain, for example), and that it is not necessary to assume thresholds in protein activity or in the growth inhibition curve to produce the plateaus and sharp drops observed experimentally when measuring optical densities.

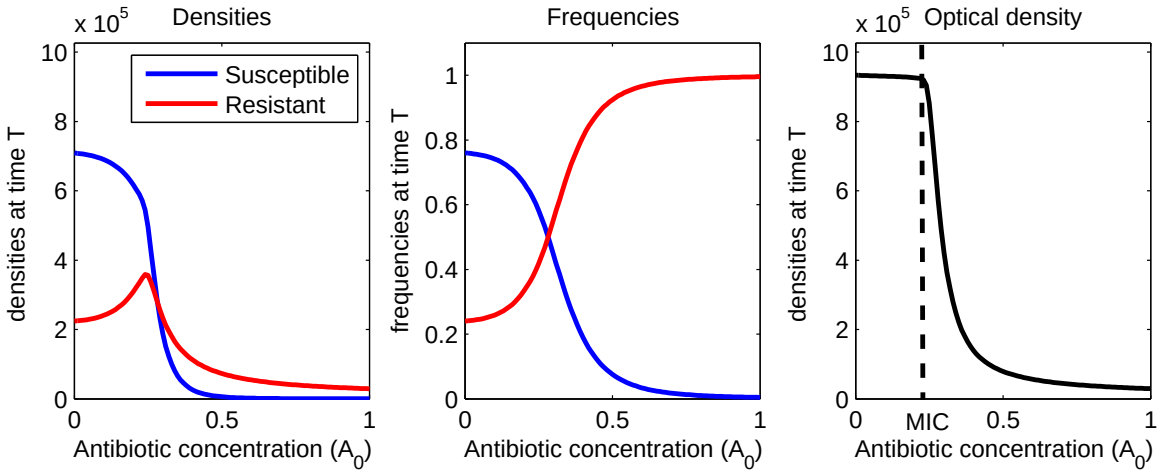

Figure 6.2: a) Densities of bacteria at time T as a function of the antibiotic concentration in the supply vessel with a constant input rate. b) Frequencies at time T as a function of A_0. c) Optical density (susceptible and resistant bacteria) at time T as a function of the input of antibiotic. Note the plateau and subsequent sharp drop in final densities, this feature is observed experimentally and is used to characterise antibiotic efficacy and drug resistance.

Note the model given by equations (6.2) can be written in the abstract form $\dot{\boldsymbol{x}} = \mathcal{G}(\boldsymbol{x}) + \delta(t)A_0\mathbf{e}$, where $\boldsymbol{x} = (S, C, \boldsymbol{P}, A)$ and $\boldsymbol{x}(0) \in \mathbb{R}^{1+1+n+1}$. Also, we can write equation (6.1) in the form $\mathcal{J}(\delta) = \int_0^T \langle w, \boldsymbol{x}(t)\rangle \, dt$, where w is has the following form

$$w = (0, 1, -1, -1, ..., -1, 0) \in \mathbb{R}^{1+1+n+1},$$

and from the discussion in the previous section we can prove that there exists a pulsating drug deployment protocol that is optimal. With this in mind, the purpose of the following section is to present numerical examples in order to compute the theoretical optimal control and compare it with other treatment protocols of therapeutic interest.

6.2 Evaluating the efficacy of different strategies

For the purpose of this section let us now explicitly consider that pathogens can evolve drug resistance to the antibiotic used, and so we set $n = 2$ and write $\boldsymbol{P} = (P_S, P_R)$, where P_S is the susceptible wild-type pathogen and P_R the antibiotic-resistant first-order mutant. The mutation matrix for this two-phenotype system is

$$\mathcal{M}_\epsilon = \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix} = (1-\epsilon)I + \epsilon \overbrace{\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}}^{M}.$$

To evaluate the efficacy of different drug deployment protocols, first we compute the theoretical optimal control by solving a regularised version of the optimisation problem (see Appendix B). The switching function, illustrated in Figure 6.3, determines precisely the moments in time where it is necessary to switch *on* or *off* the antibiotic input in order to maximise the payoff functional defined in (6.1). For example, the near-optimal control for an experiment of duration $T \approx 42$ and parameters defined in Table A.1 consists on one long interval of drug deployment followed by a short pulse, as shown in Figure 6.4a.

We now compare this theoretical optimal control with a simple feedback rule designed to stop the treatment when resistant pathogens start to colonise the host, a heuristic that can

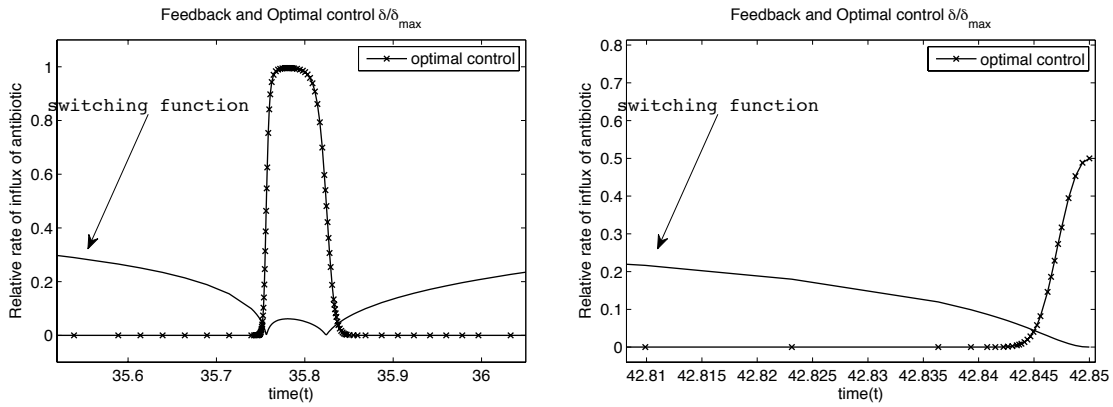

Figure 6.3: Two zooms of the switching function and the optimal control computed for the set of parameters defined in Table A.1: note the non-uniformity of the computational mesh needed to resolve spiking in the optimal control. Note also that the boundary spike near $T = 42.85$ arises as an artefact of the regularisation discussed in Appendix B.

be interpreted as

$$\mathsf{FB}: \qquad \overline{\delta}(\mathsf{x}) = \begin{cases} \delta_{max} & \text{if} \quad P_R < C + P_S \\ 0 & \text{otherwise} \end{cases}$$

Figure 6.4b) shows an implementation of the rule FB to the same set of parameters and sampling the system every $\vartheta = 2$ units of time. Note that the resulting pulsating control is similar to the optimal control, consisting on a large interval of drug deployment and followed by short pulses.

Now we proceed by computing the optimal tapering control, illustrated in Figure 6.4c. In this example, the optimal tapering control decreases antibiotic input at a rate $\alpha^* = 0.72$ and stops completely drug input when $T \approx 28$. Similarly, we can compute the optimal stopping time and obtain that the best single-pulse control stops the treatment when $\theta = 24.2$.

The responses to these controls are illustrated in Figure 6.5, and show that although the single-pulse, feedback and tapering controls perform poorly relative to the optimal control, they still manage to support the prevalence of commensal bacteria, a property that none of the fixed-dose protocols achieves in the complete competitive advantage scenario we are considering. Indeed, Figure 6.6 shows that the payoff of the tapering and feedback

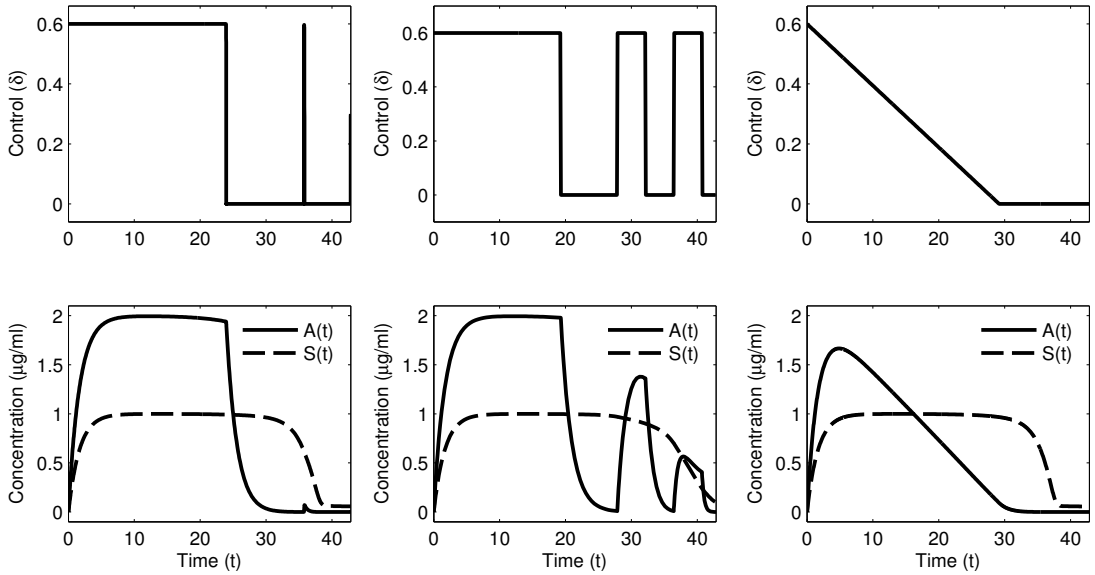

Figure 6.4: The figures in the top illustrate three different control strategies: (left) Optimal control (centre) Feedback control (right) Tapering control. The bottom figures show the corresponding concentration of resource of antibiotic inside the chemostat as a function of time.

controls outperform both the maximum deployment of antibiotic and the null treatment, but certainly not the optimal control.

A feature to notice in Figure 6.7a is that some tapering parameters produce controls with a very poor performance and that support either the susceptible or the drug-resistant pathogenic population, as shown in Figure 6.8. Only a small set of tapering parameters are able to outperform the feedback heuristic or even the maximum deployment of antibiotic. For example the tapering control illustrated in Figure 6.4c) corresponds to the optimal tapering parameter $\alpha^* = 0.72$, a value determined numerically by maximising the payoff functional for different values of $\alpha \in [0, \frac{\pi}{2}]$. The same argument applies to single-pulse protocols; although the neighbourhood of the optimal stopping time $\theta = 24.2$ produce controls that outperform the feedback control, the majority of stopping times fail to control the pathogens, a property shown in Figure 6.7b.

It is important to notice that the optimal tapering parameter as well as the optimal stopping time depend both in the length of the experiment T and the initial conditions of the system, and therefore are very difficult to estimate in practice and are not robust to modelling and

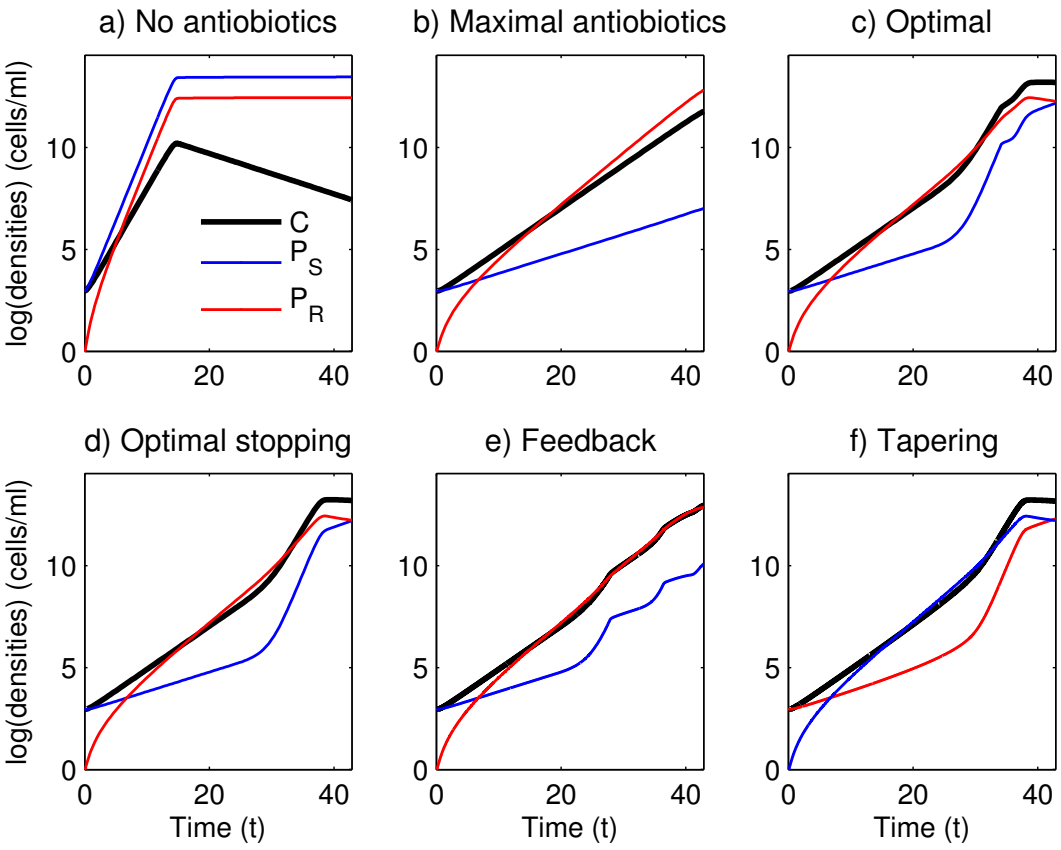

Figure 6.5: Bacterial densities of pathogens and commensals for different antibiotic deployment protocols. a) No drug is deployed, $\delta(t) = 0$, b) maximum antibiotic deployment, $\delta(t) = \delta_{max}$, c) near-optimal control, δ^*, d) single-pulse with optimal stopping time $\theta = 24.2$, e) feedback control with measurements every $\vartheta = 2$ units of time, and f) optimal tapering control with rate of descent $\alpha = 0.72$. Notice that pathogens outcompete commensals if no drug is used or if it is over-deployed, but commensals have the highest densities at the end of the experiment with the dynamic strategies implemented.

parametric uncertainties. Moreover, the optimal pulsing control is not only difficult to determine in a practical scenario, but it is also difficult to synthesise in a theoretical model like the one we are considering. One of the complications arises from the difficulty to compute numerically the switching function as time increases and consequently the number of pulses. The feedback control, however, does not increase in computational complexity as we increase the length of the experiment, because only observations on the current state of the system are used to determine the deployment or not of antibiotics. For example, Figure 6.9a shows a long-term experiment of duration $T = 400$ where the adaptive pulsing

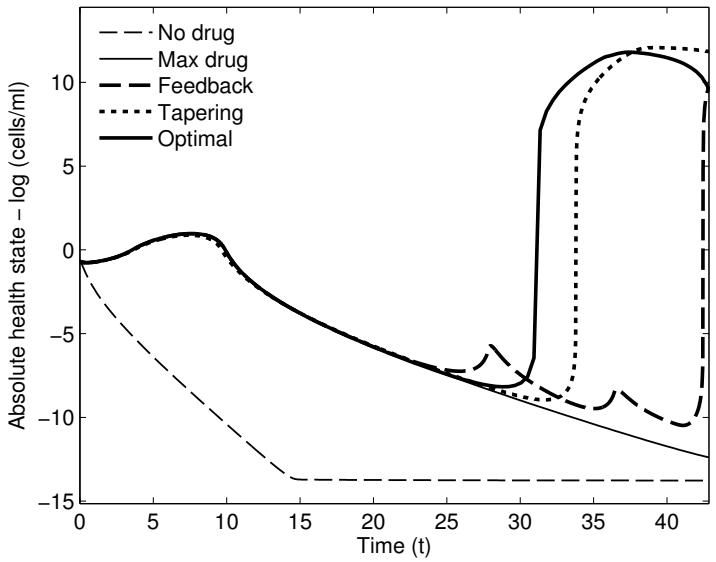

Figure 6.6: A comparison of the running payoff of the system obtained using different antibiotic deployment strategies: tapering and feedback controls are suboptimal relative to optimal control but still ensure that the commensal colonises the chemostat, a property represented by having a positive health state at the end of the experiment.

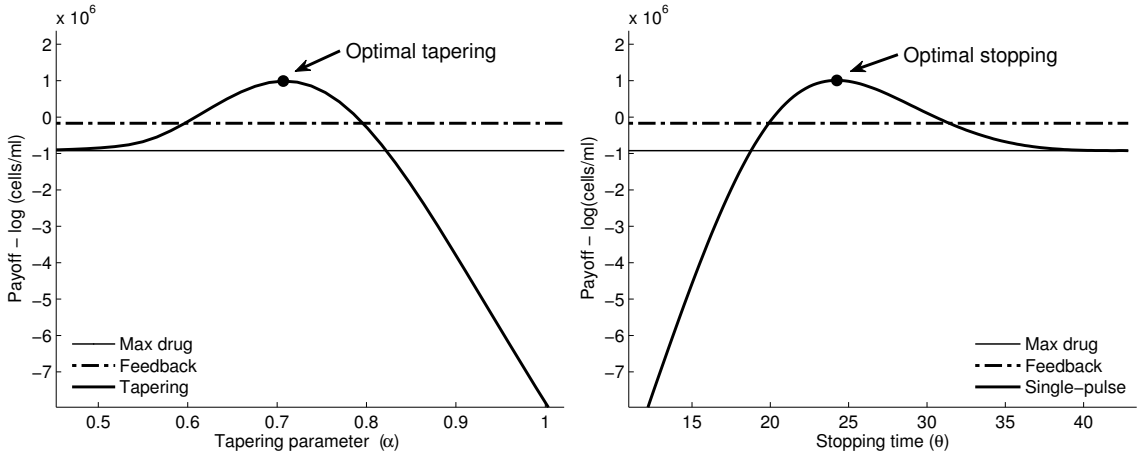

Figure 6.7: (a) Payoff $\mathcal{J}$ of the system for tapered controls with different parameters $\alpha \in [0, \pi/2]$. (b) Payoff corresponding to single-pulse controls with different stopping times $\theta \in [0, T]$. Note how in both cases the range of parameters that outperforms the feedback-based control is very small.

strategy results in a series of eventually periodic pulses of antibiotic into the chemostat: a pulse suppresses the growth of the antibiotic-susceptible pathogenic phenotype and the

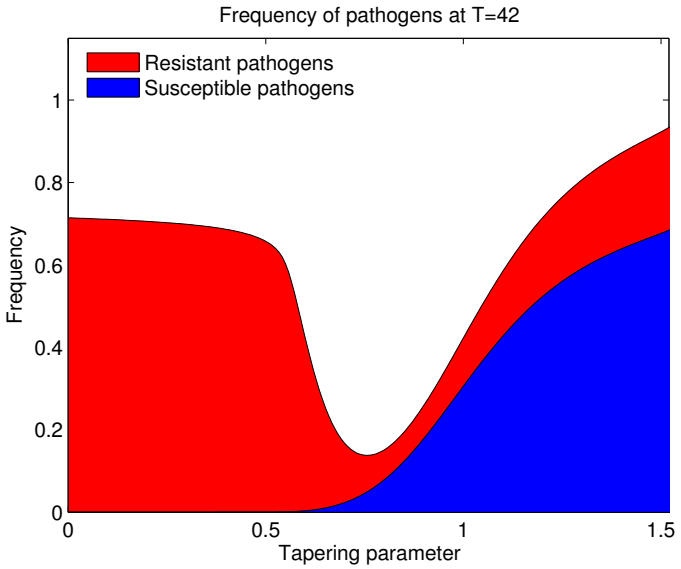

Figure 6.8: Frequency of resistance as a function of the tapering parameter. If the tapering parameter is low and therefore the antibiotic is over-deployed the resistant pathogens colonise the host. Conversely, when the tapering parameter is high and therefore not enough antibiotic is used the susceptible phenotype has the higher frequency. The tapering parameter that minimises pathogen density in an experiment of duration $T = 42.85$ is $\alpha^* = 0.72$ and is illustrated by the dotted line.

non-deployment of antibiotic suppresses the resistant pathogen. We claim that with the correct timings of these pulses it appears possible to support the commensal population.

Finally, the feedback heuristic FB also seems to be robust to changes on the maximum antibiotic dose, as shown in Figure 6.9b. This is an important feature, because in fixed-dose regimes increasing the concentration of antibiotic corresponds to an increase in the frequency of resistant pathogens at the end of the experiment. While the adaptive pulsing strategy yields the same coexistence pattern between pathogens and commensals independent of the drug concentration A_0.

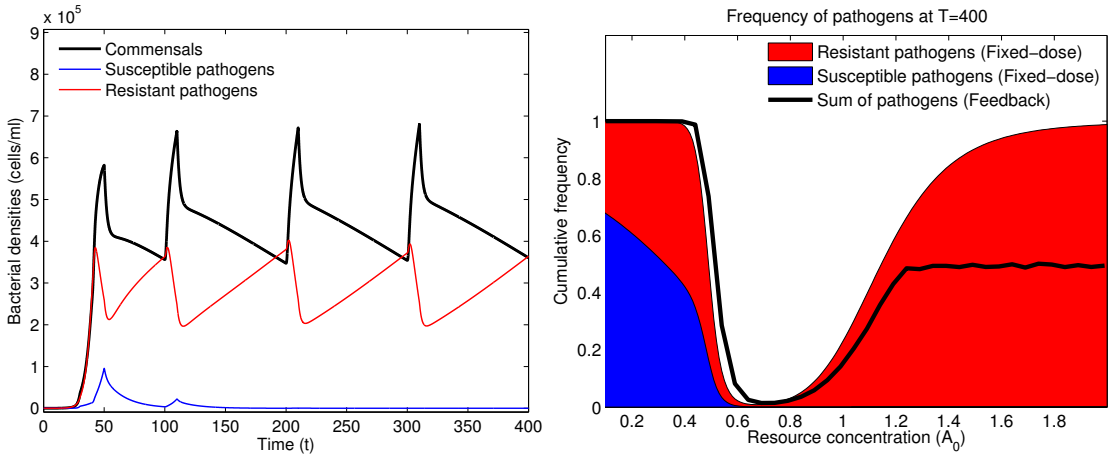

(a) **(b)**

Figure 6.9: (a) Bacterial densities under the adaptive pulsing protocol for $T = 400$ hours; note that feedback control FB with samples taken every $\vartheta = 20$ units of time leads to an eventually-periodic input of antibiotics sufficient to ensure that the commensal persists in an oscillatory fashion with the pathogen. (b) Frequency of pathogens at the end of the experiment as a function of the antibiotic concentration in the supply vessel, A_0. The shaded areas represent the cumulated frequencies of susceptible and drug resistant pathogens under a fixed-dose protocol. High doses of antibiotic select for resistant pathogens, and low doses select for susceptible pathogens. Circles represent the frequency of pathogens, both susceptible and resistant, under the adaptive pulsing strategy.

Chapter 7

Antibiotic interactions

The emergence of antimicrobial resistance is a consequence of the selective pressure imposed by the prescription of that antibiotic. This reasonable hypothesis suggests that multidrug combination treatments could help reducing the prevalence of resistant phenotypes. However, a single point mutation or the accumulation of multiple mutations may confer on bacteria resistance to several antibiotics, what is known as cross-resistance.

For this reason, in order to reduce the prevalence of multidrug resistance, we argue that it is important to understand the molecular interactions between antibiotics and to study the effect that different combination protocols have on the evolution of antimicrobial resistance. The purpose of the chapter is to characterise the profiles of interaction between two bacteriostatic antibiotics from the kinetic interactions between drug molecules and their targets, with the purpose of studying the ecological and evolutionary implications of different multidrug combinations.

7.1 Classifying drug interactions

Drug combinations can be classified as *synergistic* if they interact to increase each other's effect, *antagonistic* if their combined effect is less than the most effective drug used individually or *additive* when they do not interact. Although this intuitive classification of drug

interactions seems straightforward, the difficulty of rigorously defining the *null* interaction has been a matter of controversy for decades and the subject of a series of misconceptions [10, 44].

The most widely accepted test to quantify drug interactions are isobolograms as proposed by Loewe [65] to measure synergy. This experiment measures *in vitro* the susceptibility of a population of bacteria growing in a bidimensional array of different combinations of antibiotics. Optical densities are measured at the end of the experiment and from the dose-response surface so obtained (see Figure D.1 for an example), the interaction can be classified by lines of equal drug effect that are known as isoboles.

For the purpose of this thesis, instead of modelling population-level dose-response surfaces directly, we are going to derive them as an emergent property of a model constructed from the competitive inhibition of an essential metabolite at a single-cell level, a method discussed in [24–26].

First, let us consider that at any given moment in time, two antibiotics are present in the environment at concentrations A and B and measured in micrograms per millilitre. If the basal concentrations for each drug are denoted by A_0 and B_0 respectively, and the fixed effective dose is Δ, then any drug combination can be represented by the proportion between both drugs, a property described by the parameter $\sigma \in [0, 1]$ in the expression

$$\Delta = \sigma A_0 + (1 - \sigma) B_0. \tag{7.1}$$

Using (7.1) we can define the optimal drug ratio σ^*, where $0 \leq \sigma^* \leq 1$, as the proportion σ that minimises the growth inhibition coefficient $\gamma(\sigma A_0, (1 - \sigma) B_0)$ for σ in $[0, 1]$,

$$\gamma(\sigma^* A_0, (1 - \sigma^*) B_0) = \min_{0 \leq \sigma \leq 1} \gamma(\sigma A_0, (1 - \sigma) B_0).$$

Any treatment that deploys a fixed combination of drugs at a constant ratio, σ say, will be called a *combination treatment* in the remainder.

As mentioned before, two drugs act synergistically when their combined effect is larger than the effect of each drug used separately, a property represented by σ^* being inside the

interval $(0, 1)$. Conversely, if the optimal drug combination is to use only one of the drugs, that is $\sigma^* = 0$ or $\sigma^* = 1$, then we say their interaction is *antagonistic*. If $\gamma(\sigma A_0, (1-\sigma)B_0)$ is constant as a function of σ and so σ^* could be any value between 0 and 1, we then say that either the drugs do not interact with each other, or their interaction is additive; we consider the latter two statements to be synonymous.

7.2 Synergistic interaction between two antibiotics with the same target

Motivated by the interaction between *rifampicin* and closely related drugs like *rifabutin* and *sorangicin A*, consider the following kinetic model:

$$A + P \underset{k_{-1}}{\overset{k_1}{\rightleftharpoons}} P_A, \qquad B + P \underset{k_{-2}}{\overset{k_2}{\rightleftharpoons}} P_B, \qquad \sigma + P \underset{k_{-3}}{\overset{k_3}{\rightleftharpoons}} P_\sigma, \tag{7.2a}$$

$$\sigma + P_A \underset{k_{-4}}{\overset{k_4}{\rightleftharpoons}} P_\sigma^A, \qquad \sigma + P_B \underset{k_{-5}}{\overset{k_5}{\rightleftharpoons}} P_\sigma^B, \qquad \sigma + P_{AB} \underset{k_{-10}}{\overset{k_{10}}{\rightleftharpoons}} P_\sigma^{AB}, \tag{7.2b}$$

$$B + P_A \underset{k_{-8}}{\overset{k_8}{\rightleftharpoons}} P_{AB} \underset{k_9}{\overset{k_{-9}}{\rightleftharpoons}} P_B + A, \tag{7.2c}$$

$$P_\sigma + R \underset{k_{-6}}{\overset{k_6}{\rightleftharpoons}} P_\sigma^R \overset{k_7}{\to} M + P + \sigma + R. \tag{7.2d}$$

Equation (7.2) describes the concentration of RNA polymerase P, an antibiotic pair A and B that bind to P and the final product mRNA that we denote by M. Messenger RNA is transcribed when P binds to a sigma-factor, here denoted σ and the P_σ complex then binds to the promoter region R of some gene on the bacterial chromosome. With the consumption of ATP, not modelled here, the eventual transcription of M results whereupon the promoter region of the gene, the sigma factor and RNA polymerase unbind from the transcription unit of the gene. In the model (7.2) we have encoded the assumption that both antibiotics A and B block the binding of the mature transcription unit, which here is just an RNA polymerase-sigma factor complex, to the promoter region of the gene.

However, this does not really describe how the drugs rifampin (rif) and sorangicin A work. To more closely respect the details of the mode of action of these molecules we would need

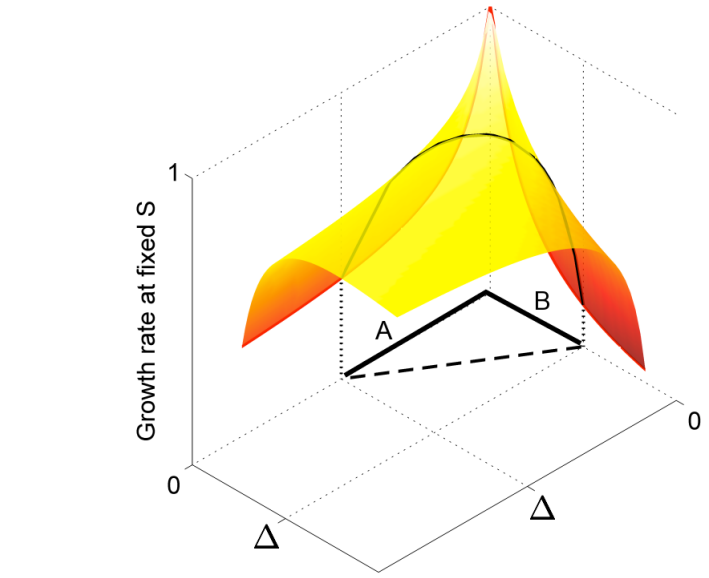

(a) ***Antagonism***

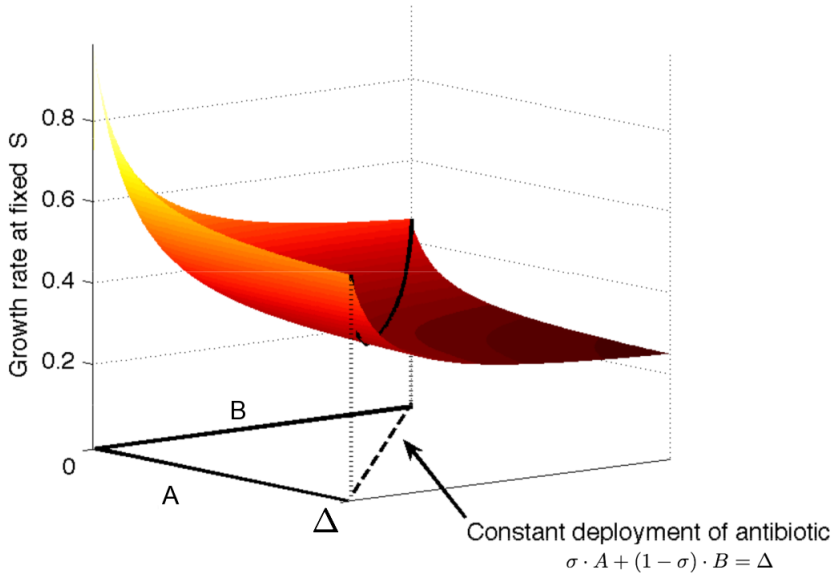

(b) ***Synergism***

Figure 7.1: Growth rate inhibition surfaces as a function of the concentration of two bacteriostatic antibiotics. The black line illustrates the constraint that the total concentration of antibiotics in the environment is fixed. (a) We say they have an antagonistic interaction if growth rate is inhibited maximally by using just a single-antibiotic. (b) Conversely, if their interaction is synergistic, then the optimal mixing protocol is a combination of both antibiotics

to include a cascade of additional terms of the form

$$P_\sigma^A + R \to P_\sigma^{AR} \to M_3 + P + \sigma + R \text{ and } P_\sigma^B + R \to P_\sigma^{BR} \to M_5 + P + \sigma + R.$$

Here M_3 and M_5 represent oligomers just a handful of nucleotides long which are the abortive mRNA-like transcripts produced when rifampicin-like drugs are bound to RNA polymerase [115]. However, we do not believe that this more complicated model provides any additional theoretical insight into the interactions of drugs targeting the same enzyme, equation (7.2) is already sufficient to have non-trivial, synergistic drug interactions as we now aim to show.

Assuming mass-action kinetics of (7.2) we obtain the following model for the time-course of concentrations of each of the agents involved in the production of mRNA:

$$\frac{d}{d\tau}P_\sigma^{AB} = k_{10}\sigma \cdot P_{AB} - k_{-10}P_\sigma^{AB}, \tag{7.3a}$$

$$\frac{d}{d\tau}P_{AB} = k_8 P_A \cdot B - k_{-8}P_{AB} + k_9 A \cdot P_B - k_{-9}P_{AB} - k_{10}\sigma \cdot P_{AB} + k_{-10}P_\sigma^{AB}, \tag{7.3b}$$

$$\frac{d}{d\tau}P_\sigma^{A} = k_4 P_A \cdot \sigma - k_{-4}P_\sigma^{A}, \tag{7.3c}$$

$$\frac{d}{d\tau}P_\sigma^{B} = k_5 P_B \cdot \sigma - k_{-5}P_\sigma^{B}, \tag{7.3d}$$

$$\frac{d}{d\tau}P_A = -k_8 P_A \cdot B + k_{-8}P_{AB} - k_4\sigma \cdot P_A + k_{-4}P_\sigma^A + k_1 P \cdot A - k_{-1}P_A, \tag{7.3e}$$

$$\frac{d}{d\tau}P_B = -k_9 P_B \cdot A + k_{-9}P_{AB} - k_5\sigma \cdot P_5 + k_{-5}P_\sigma^B + k_2 P \cdot B - k_{-2}P_B, \tag{7.3f}$$

$$\frac{d}{d\tau}P_\sigma^{R} = k_6 P_\sigma \cdot R - (k_7 + k_{-6})P_\sigma^R, \tag{7.3g}$$

$$\frac{d}{d\tau}P_\sigma = -k_6 P_\sigma \cdot R + k_{-6}P_\sigma^R + k_3 P \cdot \sigma - k_{-3}P_\sigma, \tag{7.3h}$$

$$\frac{d}{d\tau}R = -k_6 P_\sigma \cdot R + (k_7 + k_{-6})P_\sigma^R, \tag{7.3i}$$

$$\frac{d}{d\tau}\sigma = -k_3 P \cdot \sigma + k_{-3}P_\sigma - k_4 P_A \cdot \sigma + k_{-4}P_\sigma^A - k_5 P_B \cdot \sigma + k_{-5}P_\sigma^B + \ldots \tag{7.3j}$$

$$\ldots - k_{10}\sigma \cdot P_{AB} + k_{-10}P_\sigma^{AB} + k_7 P_\sigma^R,$$

$$\frac{d}{d\tau}M = k_7 P_\sigma^R, \tag{7.3k}$$

where A and B are assumed to be held at constant values during transcription.

Equation (7.3) has the following constants, namely the total available RNA polymerase

$$P_{tot} := P + P_\sigma + P_A + P_B + P_{AB} + P_\sigma^A + P_\sigma^B + P_\sigma^{AB} + P_\sigma^R$$

and the totality of gene promoters $R_{tot} := R + P_\sigma^R$. Let us now assume that all process have equilibrated so that M is produced at a constant rate or velocity in (7.3k), then

$$P_\sigma^{AB} = k_{10}/k_{-10} \cdot \sigma \cdot P_{AB} \text{ and } P_{AB} = \frac{k_8 B \cdot P_A + k_9 A \cdot P_B}{k_{-8} + k_{-9}}.$$

Now, $P_\sigma^A = \frac{k_4}{k_{-4}} P_A \cdot \sigma$ and $P_\sigma^B = \frac{k_5}{k_{-5}} P_B \cdot \sigma$ so that with (7.3e) and (7.3f) in equilibrium, we obtain

$$\begin{aligned} 0 = & -k_8 P_A \cdot B + k_{-8} P_{AB} - k_4 \sigma \cdot P_A + k_{-4} P_\sigma^A + k_1 P \cdot A - k_{-1} P_A, \\ 0 = & -k_9 P_B \cdot A + k_{-9} P_{AB} - k_5 \sigma \cdot P_5 + k_{-5} P_\sigma^B + k_2 P \cdot B - k_{-2} P_B, \end{aligned}$$

that can be rewritten using the preceding algebraic relationships for P_A and P_B as a function of P, A and B:

$$\begin{pmatrix} k_{-1} + B\frac{k_8 k_{-9}}{k_{-8}+k_{-9}} & -\frac{k_9 k_{-8}}{k_{-8}+k_{-9}} A \\ -\frac{k_8 k_{-9}}{k_{-8}+k_{-9}} B & k_{-2} + A\frac{k_9 k_{-8}}{k_{-8}+k_{-9}} \end{pmatrix} \begin{pmatrix} P_A \\ P_B \end{pmatrix} = P \begin{pmatrix} k_1 A \\ k_2 B \end{pmatrix}. \tag{7.4}$$

The linear equation defined by (7.4) can be solved to give

$$P_A = P \cdot A \cdot \kappa_A(A, B) \text{ and } P_B = P \cdot B \cdot \kappa_B(A, B)$$

where κ_A and κ_B are rational functions of the form

$$\kappa(A, B) = \frac{p_1 + p_2 A + p_3 B}{1 + p_4 A + p_5 B},$$

where each of the five parameters $p_1, ..., p_5$, and with one set for each of the two antibiotics, are various combinations of the rate parameters in (7.2).

Continuing with the equilibrium assumptions, $P_\sigma^R = P_\sigma \cdot R \cdot k_6/(k_7 + k_{-6})$ follows from (7.3g) but then adding the latter to (7.3h) yields $P_\sigma = k_3\sigma \cdot P/(k_{-3} + R\frac{k_6k_7}{k_7+k_{-6}}) =: \sigma \cdot P \cdot n(R)$, whence

$$P_\sigma^R = \sigma \cdot P \cdot \frac{k_3 R}{k_{-3} + R\frac{k_6k_7}{k_7+k_{-6}}} \frac{k_6}{k_7 + k_{-6}} =: \sigma \cdot P \cdot m(R).$$

By defining $\kappa_{AB}(A, B) := (k_8\kappa_A(A, B) + k_9\kappa_B(A, B))(k_{-8} + k_{-9})^{-1}$ we can write

$$\begin{aligned} P_{tot} &= P \cdot \left(1 + n(R)\sigma + A\kappa_A + B\kappa_B + AB\kappa_{AB} + \frac{k_4\sigma}{k_{-4}}A\kappa_A + \frac{k_5\sigma}{k_{-5}}B\kappa_B + m(R)\sigma\right), \\ &= P \cdot (1 + (n(R) + m(R))\sigma + A\kappa_A(1 + \ell_4\sigma) + B\kappa_B(1 + \ell_5\sigma) + AB\kappa_{AB}) \end{aligned}$$

and so

$$\frac{d}{d\tau}M = k_7\sigma \cdot P \cdot m(R) = \frac{k_7\sigma \cdot m(R) \cdot P_{tot}}{1 + (n(R) + m(R))\sigma + A\kappa_A(1 + \ell_4\sigma) + B\kappa_B(1 + \ell_5\sigma) + AB\kappa_{AB}},$$

a function that we denote by $v(A, B)$. If we finally define $\gamma(A, B)$ to be the dimensionless reduction in the production rate of mRNA due to A and B then $\gamma(A, B) = v(A, B)/v(0, 0)$ which equals

$$\begin{aligned} \gamma(A, B) &= \frac{1 + (n(R) + m(R))\sigma}{1 + (n(R) + m(R))\sigma + A\kappa_A(1 + \ell_4\sigma) + B\kappa_B(1 + \ell_5\sigma) + AB\kappa_{AB}}, \\ &= \frac{1}{1 + A \cdot q_A(R, \sigma)\kappa_A + B \cdot q_B(R, \sigma)\kappa_B + AB \cdot q_{AB}(R, \sigma)\kappa_{AB}}. \end{aligned} \tag{7.5}$$

We have written $q_A(R, \sigma)$ in place of $(1 + \ell_4\sigma)/(1 + (n(R) + m(R))\sigma)$ for brevity, the functions q_B and q_{AB} are similarly defined as $(1+\ell_5\sigma)/(1+(n(R)+m(R))\sigma)$ and $1/(1+(n(R) + m(R))\sigma)$ respectively.

The simpler expression than (7.5)

$$\gamma(A, B) = \frac{1}{1 + A\kappa_A + B\kappa_B + AB\kappa_{AB}} \tag{7.6}$$

is used for convenience in the numerical simulations presented in this thesis, where κ_A, κ_B

and κ_{AB} are assumed to be constant and all three strictly positive. To see that the function $\gamma(A, B)$ so-defined yields a *synergistic* interaction let σ lie strictly between 0 and 1 and consider the expression

$$\begin{aligned}\frac{1}{\gamma(A,B)} - \frac{1}{\gamma(\sigma A,(1-\sigma)B)} &= (1-\sigma)\cdot A\kappa_A + \sigma\cdot B\kappa_B + (1-\sigma+\sigma^2)\cdot AB\kappa_{AB}\\ &> 0.\end{aligned}$$

As a result, $\gamma(\sigma A,(1-\sigma)B) < \gamma(A,B)$ whenever $0 < \sigma < 1$ and so the drugs A and B synergise because they are most effective at reducing the rate of mRNA synthesis when used in combination.

We note that interactions between drugs targeting different enzymes in a pathway do not necessarily lead to such simple functional forms for the drug inhibition surface as the one described in (7.6), see [58, 119] for example, explaining our restriction to drugs targeting the same enzyme. However, we argue that in some cases this simple model may be able to capture the interaction between two antibiotics that bind to non-overlapping sites of the same target and whose interaction is known to be synergistic, for example the drug pair studied in [48] and illustrated in Figure 7.2.

7.3 Predicting the dose-response surface

The derivation of the previous section allows us to quantify the inhibitory effect of a drug combination of two bacteriostatic antibiotics with a synergistic interaction. The obtained growth inhibition surface (7.6) describes the inhibition of the production of an essential metabolite at a single-cell level, but does not necessarily explain the observed population-level inhibition. Here we ask if we can predict an experimentally obtained population-level dose-response surface from the single-cell inhibition surface.

Let us assume that at any given moment in time two antibiotics are present in the environment at concentrations A and B. As discussed in Chapter 5, if we consider that bacterial growth rate can be described by a function $G(S, A, B)$ that depends on the concentrations

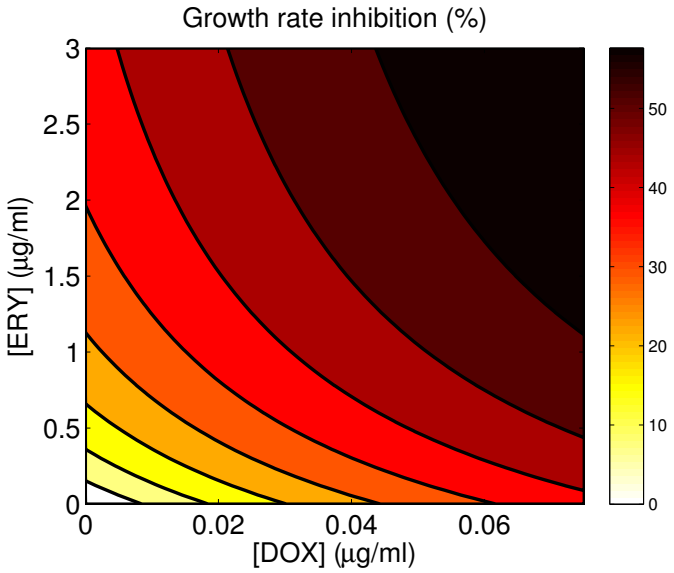

Figure 7.2: The shape of the isoboles of equal effect illustrates that doxycycline and erythromycin have a synergistic interaction, as lower-dose drug combinations have the same effect as high-dose monotherapies.

of both drugs and upon the concentration of a limiting resource S in the environment, then this growth function can be modelled as a standard Michaelis-Menten term multiplied by a growth inhibition coefficient $\gamma(A, B)$ that depends on the concentration of both drugs,

$$G(S, A, B) = c \cdot u(S) \cdot \gamma(A, B). \tag{7.7}$$

Again, c denotes a resource conversion rate and $u(S)$ is the resource uptake function defined in equation (5.13) and $\gamma(A, B)$ a function such that $\gamma(A, 0) = \gamma_A(A)$ and $\gamma(0, B) = \gamma_B(B)$, where $\gamma_A(A)$ and $\gamma_B(B)$ are growth inhibition functions that characterise the single, separate use of each antibiotic.

Therefore if we represent with $P(t)$ the density of an isogenic population of pathogenic bacteria growing under resource limitation and under the effect of two bacteriostatic antibiotics of concentrations $A(t)$ and $B(t)$, and we assume that the concentration of bacterium per unit of volume is so high and uniformly distributed in space, then it is possible to describe the ecological interactions between different bacterial populations with deterministic

models based on the mass action law [95], like the one described below

$$\dot{S} = u(S) \cdot P, \tag{7.8a}$$

$$\dot{P} = G(S, A, B) \cdot P, \tag{7.8b}$$

$$\dot{A} = -aA \cdot P, \tag{7.8c}$$

$$\dot{B} = -bB \cdot P, \tag{7.8d}$$

with initial conditions $\boldsymbol{x}(0) = (S(0), P(0), A(0), B(0))$. Constants a and b are parameters that denote the rate of binding of antibiotic molecules to their targets.

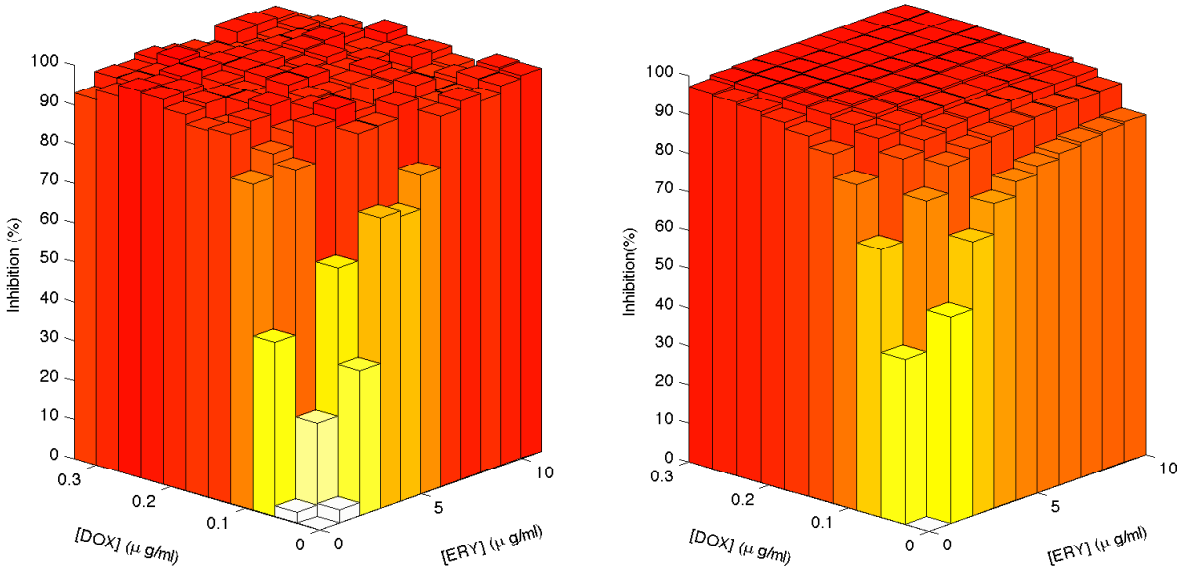

Figure 7.3: Population-level inhibition of E. Coli growing in a glucose-limited environment and under the effect of different concentrations of doxycycline and erythromycin. (a) Experimental data (b) Model predictions

To test our model, we measured experimentally the inhibitory effect on an isogenic population of E. Coli growing under glucose limitation and under the effect of two bacteriostatic antibiotics that target different sites in the ribosome, and thus inhibit protein synthesis: doxycycline that binds to the 30s subunit and erythromycin that binds to the 50s subunit. Using the simple model described by equations (7.8a-d) and the growth inhibition surface for synergistic interactions derived in the previous section and illustrated by the

isobolograms shown in Figure 7.2, we were able to predict with quantitative accuracy the population-level inhibition of bacteria growing at different drug concentrations, as illustrated in Figure 7.3. Details of the experiment and the numerical fit are described in Appendix D. In the following chapter we will use this calibrated model to design optimal drug combinations in a microbial model system that incorporates the evolution of resistance.

Chapter 8

Ongoing work: Evolution of resistance in multidrug environments

A common approach to dealing with the evolution of antibiotic resistance and to increase the efficacy of antimicrobial treatments is to use multidrug combination therapies [75–77]. With this in mind, in this chapter we turn to a simple lab-based microbial system and consider perhaps the simplest possible experimental model into which one can deploy antibiotics: two antibiotics are used to target an isogenic bacterial inoculate in shaken flasks with batch transfers. The term *flask* could also be interpreted as wells on a shaken microtiter plate, but as we work with units of bacterial cells per millilitre, the distinction is not too important. This is a standard experimental system often used in microbial evolution experiments [18] and has the advantage of removing a great deal of the chemical complexity commonly found in the environmental niches occupied by microbes. This simplicity can be exploited by the modeller to make predictive and testable statements regarding the short-term microevolution of the bacteria.

The results presented in this chapter are based on a series of experiments performed in Kiel in September of 2010 and designed to study the evolutionary consequences of different multidrug combination protocols. Although the experiments are still in a preliminary stage, we were able to calibrate our models with growth and inhibition parameters fitted from

these experiments, and therefore instead of presenting general statements based on abstract models, in this chapter we will attempt to make predictions that we aim to validate in future experiments.

8.1 Optimal drug proportion

For decades the medical and pharmacological communities have sought antimicrobial drugs that increase their effect when used in combination, an interaction known as *synergism* [44]. It is known, however, that synergistic drug combinations may select for multidrug resistance [48] or have synergistic side effects [57]. For this reason, it has been proposed that antagonistic interactions, usually dismissed in clinical settings, may in fact help to select against the evolution of resistant phenotypes [21, 117]. For the purpose of this thesis, however, instead of concentrating on how to design drug interactions that minimise selective pressures in favour of resistant phenotypes, we are going to assume that the profile of interaction between both antibiotics is given and focus on studying the ecological and evolutionary implications of different drug usage strategies. Our aim is to understand whether there are circumstances in which other types of therapy might be favoured over therapies based on fixed-dose multidrug combinations.

Hence, for the remainder of this thesis we will consider the drug interaction preferred in clinical contexts, where drugs increase their efficacy when used in combination. In particular, we will consider two bacteriostatic antibiotics of different functional classes with a strong synergistic interaction: erythromycin, a macrolide that binds to the 50s ribosomal subunit with a concentration that we will denote with ERY and doxycycline, a tetracycline that binds to the 30S ribosomal subunit at a concentration denoted by DOX.

Assuming that we can model their interaction as mutually non-exclusive competitive inhibitors, the derivation of the previous chapter allows us to describe their interaction through the following growth rate inhibition coefficient:

$$\gamma(DOX, ERY) = \gamma_{dox}(DOX) \cdot \gamma_{ery}(ERY). \tag{8.1}$$

$\gamma_{dox}(DOX)$ and $\gamma_{ery}(ERY)$ are the drug inhibition coefficients that describe the single use of each antibiotic, functions that can be written as

$$\gamma_{dox}(DOX) = 1 - \left(\frac{\kappa_1^{dox}\, DOX}{\kappa_2^{dox} + DOX} \right),$$
$$\gamma_{ery}(ERY) = 1 - \left(\frac{\tilde{\kappa}_1^{ery}\, ERY}{\tilde{\kappa}_2^{ery} + ERY} \right),$$

where κ_2^{dox} and κ_2^{ery} are the corresponding antibiotic affinities to their targets, while κ_1^{ery} and κ_1^{dox} represent the level of maximum growth inhibition achieved by each antibiotic.

Another assumption of this chapter is that the rate of input of antibiotics is fixed at a concentration Δ. This fixed-dose dosage assumption allows us to describe drug combinations through a mixing parameter $\sigma \in [0, 1]$.

Then if DOX_0 and ERY_0 are the basal concentrations of each drug, we can describe any drug cocktail with the expression

$$\sigma \cdot DOX_0 + (1 - \sigma) \cdot ERY_0 = \Delta. \tag{8.2}$$

Note that this expression allows us to define the optimal drug ratio, denoted σ^*, as the drug proportion that minimises the inhibition coefficient $\gamma(\sigma DOX_0, (1-\sigma)ERY_0)$ for σ in $[0, 1]$, a property expressed by the following condition:

$$\gamma(\sigma^* DOX_0, (1-\sigma^*)ERY_0) = \min_{0 \le \sigma \le 1} \{\gamma(\sigma DOX_0, (1-\sigma)ERY_0)\}.$$

Now, if we inoculate a flask with an isogenic population of susceptible bacteria, for example wild-type *E. Coli*, and for the moment we assume that they cannot acquire antibiotic resistance to any of the drugs used, then in order to reduce bacterial density maximally it is clear that it would be optimal to deploy the drug cocktail that maximises the synergistic effect.

Indeed, Figure 8.1 shows that the the drug ratio between erythromycin and doxycycline at sub-inhibitory concentrations that maximises cellular-level inhibition, in this case $\sigma^* =$

0.72, corresponds roughly with the drug proportion that minimises the density of bacteria in a short-time experiment ($T = 24$ hours). Details of the experiment can be found in Appendix D.

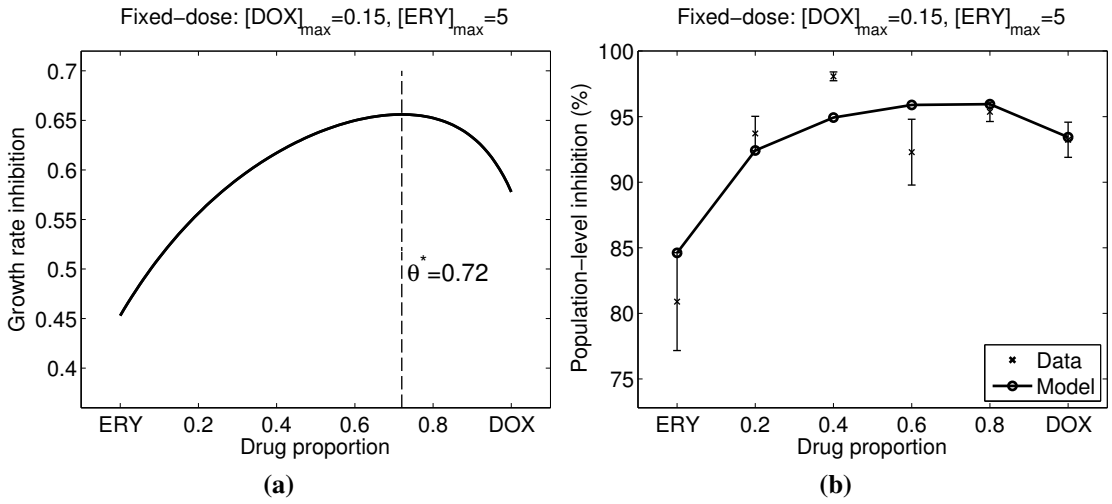

Figure 8.1: (a) Growth rate inhibition at a single-cell level of a susceptible bacteria under the effect of a fixed drug dose described by a proportion parameter $\sigma \in [0, 1]$, where $\sigma = 0$ corresponds to deploying erythromycin at maximum dose ($5\mu g/ml$) and $\sigma = 1$ the maximum dose of doxycycline ($0.15\mu g/ml$). The optimal drug proportion is $\sigma^* = 0.72$ and the parameter used are detailed in Table A.3. (b) Population-level growth inhibition of *E. coli* growing at a fixed concentration of resource ($S = 2000\mu g/ml$) and exposed to different drug proportions in a 24-hour experiment. Experimental data is represented with standard error bars, while solid lines represent predictions of the model using the growth inhibition curve defined in (a). Note how the drug proportion that maximises the synergistic effect corresponds roughly with the drug ratio that maximises population-level inhibition inhibition.

But what happens in a longer experiment where bacteria have enough time to evolve drug-resistance? In this case the population structure would be changing over time and therefore the inhibition curve predicted for susceptible bacteria would not be able to describe the inhibitory effect of the antibiotic on the resistant sub-population present in the system. Consequently, for long-term experiments there is no reason to expect a direct correspondence between the predicted optimal single-drug proportion σ^* and the optimal population-level inhibition. So we ask, in an evolutionary model what is the optimal drug combination? In order to address this question, in the following section we model an evolvable bacteria growing under resource limitation in a serial transfer experiment.

8.2 Modelling serial transfers

To study the evolution of drug resistance in the presence of a given drug combination, we will consider that a shaken flask with liquid growth medium is inoculated with an initial density of isogenic bacteria. Cells are grown for a fixed amount of time, after which some of the abiotic resources needed for cell growth are consumed, and therefore the population density has reached an equilibrium. Then a small and fixed volume is taken from the flask and transferred to a second flask which contains fresh liquid growth medium. Repeating this process N times defines a serial transfer experiment.

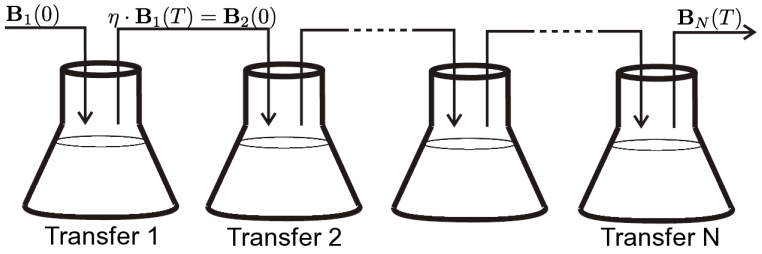

Figure 8.2: Diagram illustrating a serial transfer experiment. The system is inoculated with an isogenic type of bacteria and allowed to grow under resource limitation for T units of time. Then a fraction of the population is transferred into a new flask with fresh medium.

Now, in each transfer we introduce a combination of antibiotics with basal concentrations DOX_0 and ERY_0, and thus we can describe any treatment protocol by a vector of drug proportions

$$\boldsymbol{\sigma} = (\sigma_1, \sigma_2, ..., \sigma_{N+1})$$

where the relative concentration of drugs used in the j-th transfer is σ_j, a value between 0 and 1. The actual concentration of drug deployed in the j-th flask is then $\sigma_j \cdot DOX_0\, \mu g/ml$ of doxycycline and $(1-\sigma_j) \cdot ERY_0\, \mu g/ml$ of erythromycin; we cannot choose *not* to deploy the drug in any of the flasks because we only have freedom to choose σ_j and not the overall doses.

For the purpose of this chapter we will denote the vector containing the densities of each phenotype of bacteria with the variable $\boldsymbol{P}(t) = (P^s, P^d, P^e, P^{de})$, where P^s represents

an antibiotic-susceptible bacteria, P^d and P^e resistant phenotypes to doxycycline and erythromycin respectively, and P^{de} the multidrug-resistant bacterial phenotype.

Furthermore, if we denote with the variable $S(t)$ the concentration of the limiting resource and we run the system for N transfers and denote each one as $j \in \{1, N+1\}$, then the initial density of bacteria of transfer j is determined by the terminal condition of the previous transfer. If we assume that every transfer has a duration of T hours, then our temporal variable can be represented by $t \in [0, T]$, and the initial condition of each transfer would be

$$\boldsymbol{P}_j(0) = \eta \cdot \boldsymbol{P}_{j-1}(T)$$

where $0 < \eta \ll 1$ is a dilution parameter. Note that although shaken flasks remove spatial structure, the transfer process introduces a temporal seasonality in the ecological dynamics of the system that has been shown to be relevant in some circumstances [68].

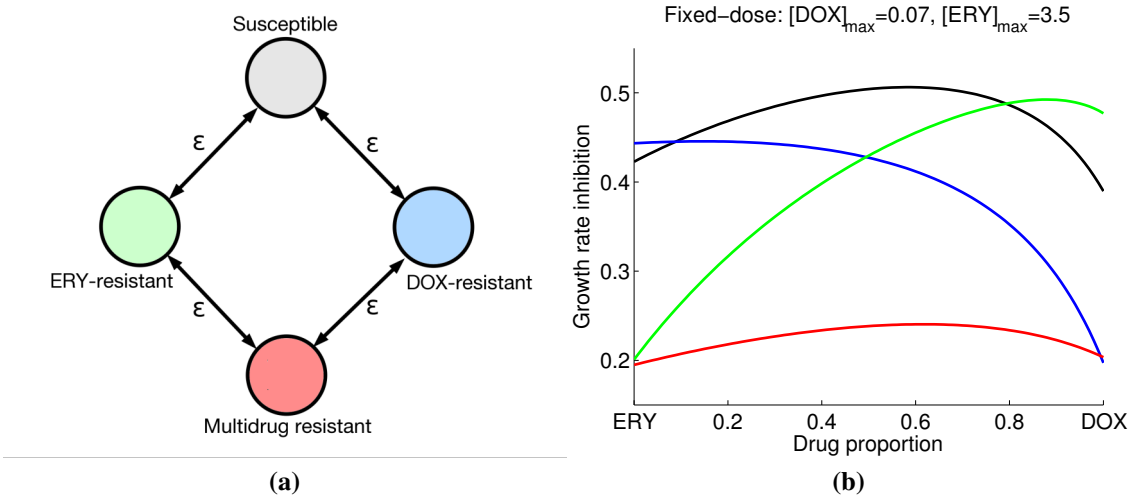

(a) **(b)**

Figure 8.3: (a) The graph associated with the mutation process $\mathcal{M}_\epsilon$ illustrating that with a probability ϵ bacteria can evolve resistance to any of the antibiotics used or have a compensatory mutation. The accumulation of two mutations can confer bacteria with multidrug resistance. (b) Growth rate inhibition curves for each of the phenotypes present in the system; the black line denotes susceptible bacteria, green and blue lines denote resistant to erythromycin and doxycycline respectively, and red represents multidrug resistance. Maximum drug doses are given by $3.5\mu g/ml$ for erythromycin and $0.07\mu g/ml$ for doxycycline.

Now let us assume that bacteria can evolve resistance to each one of the antibiotics used through a single point mutation occurring at a rate $\epsilon > 0$, and that the accumulation of two mutations can confer multidrug resistance upon a susceptible bacteria. Of course, in reality the path to drug resistance can follow different mutational trajectories, but as discussed in [113] most of these trajectories are inaccessible through natural selection. Then for simplicity we will consider the evolutionary dynamics illustrated in Figure 8.3 and we will describe it mathematically by the mutation matrix $\mathcal{M}_\epsilon = (1-\epsilon)I + \epsilon M$, where I is the identity matrix representing clonal reproduction, ϵ the probability of mutation and M is a n-by-n matrix where the m_{ij} entry represents the mutation rate of bacterial type i into bacterial type j,

$$\mathcal{M}_\epsilon = \begin{pmatrix} 1-\epsilon & \frac{1}{2}\epsilon & \frac{1}{2}\epsilon & 0 \\ \frac{1}{2}\epsilon & 1-\epsilon & 0 & \frac{1}{2}\epsilon \\ \frac{1}{2}\epsilon & 0 & 1-\epsilon & \frac{1}{2}\epsilon \\ 0 & \frac{1}{2}\epsilon & \frac{1}{2}\epsilon & 1-\epsilon \end{pmatrix} = (1-\epsilon)I + \epsilon \overbrace{\begin{pmatrix} 0 & \frac{1}{2} & \frac{1}{2} & 0 \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} \\ 0 & \frac{1}{2} & \frac{1}{2} & 0 \end{pmatrix}}^{M}.$$

As discussed in Section 5.3, bacterial phenotypes can be defined using growth rates and antibiotic susceptibility and resistance patterns, properties that are contained within the set of parameters $(V_{max}, K, \kappa_1, \kappa_2)$. For example, the susceptible phenotype, P^s, has a growth rate described by the function

$$G^s(S,A) = c\left(\frac{V^s_{max}S}{K^s+S}\right)\left(1-\frac{\kappa^s_1 A}{1+\kappa^s_2 A}\right),$$

and the *ERY-resistant* phenotype would have, by definition, a higher growth rate than *antibiotic-susceptible* bacteria at high concentrations of erythromycin. More formally, this means that if the growth rate of a resistant phenotype is given by the function

$$G^e(S,A) = c\left(\frac{V^e_{max}S}{K^e+S}\right)\left(1-\frac{\kappa^r_1 A}{1+\kappa^r_2 A}\right)$$

then there exists of a critical concentration of erythromycin $ERY^* > 0$ for which $G^e(S,0,ERY) > G^s(S,0,ERY)$ for all $ERY \geq ERY^*$.

It is well-known [7, 43] that antibiotic-resistant phenotypes may encounter a reduction in fitness in an antibiotic-free environment, a so-called *fitness cost of resistance*. In our modelling framework, this property can be realised by imposing a lower resource uptake rate $u^e(S) < u^s(S)$ and therefore a lower growth rate at low antibiotic concentrations that results in $G^e(S, 0, 0) < G^s(S, 0, 0)$.

An analogous statement can be made for the *DOX-resistant* pathogen P^d and for the multidrug resistant pathogen P^{de}. The growth inhibition surfaces for each phenotype are illustrated in Figure 8.4.

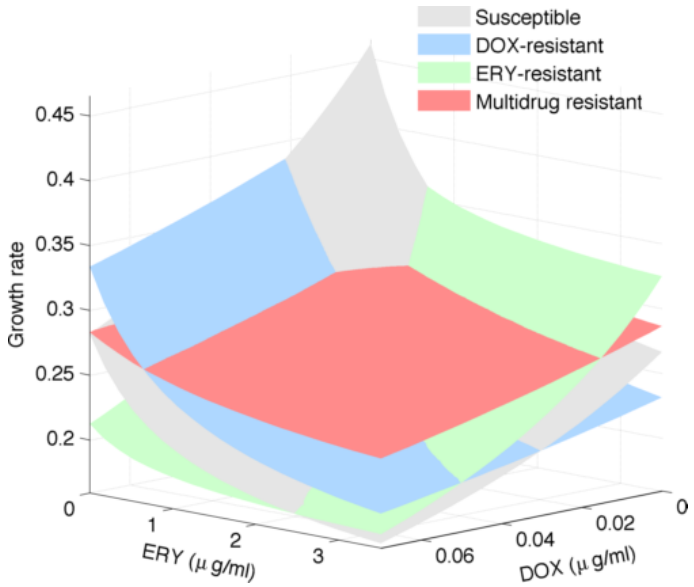

Figure 8.4: Growth inhibition surfaces for susceptible, single-drug resistant and multidrug resistant bacteria. Note how in multidrug environments the multidrug resistant phenotype has the higher growth rate, but it pays a cost of being resistant having a lower growth rate at low antibiotic concentrations. These inhibition surfaces were computed using the parameters described in Table A.3 and at fixed resource concentration $S = 2000\mu g/ml$.

Then if we denote the state of the system for transfer j with the variable

$$\boldsymbol{x_j}(t) := (S_j(t), \boldsymbol{P}_j(t), DOX_j(t), ERY_j(t)),$$

the system of differential equations governing the evolution of the state in each transfer is

given by $\dot{\boldsymbol{x}}_{\boldsymbol{j}} = \mathcal{F}(\boldsymbol{x}_j, \sigma_j)$, and $\mathcal{F}$ can be written as

$$\begin{aligned}
\frac{d}{dt}S &= -\langle u(S), \boldsymbol{P}\rangle, && \text{(8.3a)}\\
\frac{d}{dt}\boldsymbol{P} &= \mathcal{M}_\epsilon(G(S)\cdot \boldsymbol{P}), && \text{(8.3b)}\\
\frac{d}{dt}DOX &= -aDOX\cdot(P^s + P^e + P^d + P^{de}), && \text{(8.3c)}\\
\frac{d}{dt}ERY &= -bERY\cdot(P^s + P^e + P^d + P^{de}), && \text{(8.3d)}
\end{aligned}$$

with each transfer having an initial condition $\boldsymbol{x}_{\boldsymbol{j}}(0) = (S(0), \boldsymbol{P}_j(0), \sigma_j \cdot DOX_0, (1-\sigma_j)\cdot ERY_0)$. Parameters a and b denote the binding of antibiotic molecules to their targets.

8.3 Designing optimal multidrug combinations

A fundamental problem when solving optimisation problems is defining an objective functional. Is it more important to minimise conditions that select for resistant pathogens? Or do we simply care about minimising the overall bacterial density? Of course, the nature of the optimal solution would depend on the objective we choose to consider, and therefore it is a fundamental decision to make. For instance, we could consider the following objectives

- **Minimise drug-resistant bacteria.** A possible optimality criterium would be to minimise the average of resistance at the end of each transfer during an experiment of N transfers, a payoff functional that can be written as

$$\mathcal{J}_1(\boldsymbol{x};\sigma) = -\sum_{j=1}^{N+1} R_j(T). \tag{8.4}$$

  where $R_j(t) = P_j^d(t) + P_j^c(t) + P_j^{dc}(t)$ represents the density of resistant bacteria at the end of transfer j. Also, we could try to minimise the final density of bacteria at

the end of the experiment. An objective functional defined as

$$\mathcal{J}_2(\boldsymbol{x};\sigma) = -R_{N+1}(T). \tag{8.5}$$

- **Minimise bacterial density.** Similarly, we could try to minimise the overall average of bacteria at the end of each transfer. An objective represented in the expression

$$\mathcal{J}_3(\boldsymbol{x};\sigma) = -\sum_{j=1}^{N+1} \boldsymbol{P}_j(T), \tag{8.6}$$

or the density of bacteria at the end of the experiment,

$$\mathcal{J}_4(\boldsymbol{x};\sigma) = -\boldsymbol{P}_{N+1}(T). \tag{8.7}$$

- **Minimise maximum bacterial density.** A common approach in optimisation problems is to minimise the maximum value of a given set of costs. In this context, this *minimax* optimisation criterium could be interpreted as minimising the bacterial density of the phenotype with the highest density at the end of each transfer, an objective functional that could be written as

$$\mathcal{J}_5(\boldsymbol{x};\sigma) = -\max_{j\in\{1,N+1\}}(\boldsymbol{P}_j) \tag{8.8}$$

- **Minimise the fitness of the population.** For each transfer, let us define by $\Delta P = \boldsymbol{P}(T) - \boldsymbol{P}(\mathbf{0})$ the total increase in bacterial density. The *time of adaptation*, denoted t_{adapt}, is defined as the interpolated time at which the population of bacteria has reached half its maximum value. Then we say that the *fitness* of the bacterial population to this particular environment is

$$\phi = \frac{\Delta P}{2 * t_{adapt}}. \tag{8.9}$$

This measure is similar to the one proposed in [48], as the *rate of adaptation* is a proxy of the overall fitness of the population as it adapts to a given environmental

condition. This definition allows us to write down the following payoff functional

$$\mathcal{J}_6(\boldsymbol{x};\sigma) = -\sum_{j=1}^{N+1} \phi_j. \tag{8.10}$$

Maximising this objective could be interpreted as reducing optimally the ability of a population of bacteria to adapt to a given environment.

From the different objectives discussed above, it is clear that there is not a single definitive quantity that we can rely upon in order to define the optimal drug proportion, as each objective would produce a different optimal solution. For instance, if our purpose was simply to minimise drug resistant bacteria, it is clear that the optimal treatment is never to use antibiotics. As we already discussed in Section 2.6.3, it has even been postulated in [42] that we should put society before the individual, and to preserve the efficacy of antibiotics for the community some patients should remain untreated. However, this is an unreasonable strategy from the perspective of a single-host, because without treatment the wild-type pathogen would colonise the host.

We also argue that goals defined in terms of “average of resistance" are more suitable for our purposes than the ones based on end-point minimisation. The reason being that the optimal strategy that minimises bacterial density at the end of the last transfer could be to deploy a single drug, for example ERY, during the first $N-1$ transfers and then, when the population consists mainly on ERY-resistant pathogens, switch to DOX in the last transfer. This strategy may decrease the final density of bacteria, but would not guarantee that low densities are maintained throughout the experiment. The minimax objective represented by equation (8.8) would certainly take into account this possibility, but for simplicity we would consider for our numerical simulations the objective defined in (8.6) and we will denote it $\mathcal{J}(\boldsymbol{x};\sigma)$ hereafter.

Now, having defined our goal, computing the optimal fixed-proportion drug deployment protocol, $\boldsymbol{\sigma}^*_{fixed}$, can be achieved by numerically solving the following one-dimensional

optimisation problem:

$$\mathcal{J}(\boldsymbol{x}, \boldsymbol{\sigma}^*_{fixed}) = \max_{\boldsymbol{\sigma}\in[0,1]} \{\mathcal{J}(\boldsymbol{x}, \boldsymbol{\sigma}) : 0 \leq \sigma_i \leq 1 \text{ and } \sigma_i = \sigma_j \text{ for all } i, j \in \{1, N+1\}\} \tag{8.11}$$

Figure 8.5 illustrates an example where the optimal drug ratio that maximises the payoff functional defined in (8.6) corresponds to a drug combination described by the mixing parameter $\boldsymbol{\sigma}^*_{fixed} = 0.57$. Note how this drug cocktail is very effective inhibiting bacterial growth rate during the first few transfers, but at the end of the experiment its efficacy has decreased considerably because the population is composed mainly of multidrug resistant bacteria, growing without any problems until reaching stationary phase.

It is important to mention that the optimal fixed-dose proportion $\boldsymbol{\sigma}^*_{fixed} = 0.57$ was computed for a 12-transfer simulated experiment. Of course, there is no reason to believe that this optimal drug cocktail would continue to be optimal for a long-term experiment where the population structure presents different patterns of resistance and susceptibility. In fact, our simulations show that the optimal drug proportion is indeed a dynamic property of the system that depends, among other factors, on the duration of the experiment. This property is illustrated in Figure 8.6: for short-term experiments the optimal drug proportion corresponds roughly with the drug ratio that maximises the synergistic effect, but this correspondence does not hold for an experiment with more dilutions. In this particular example the predicted optimal drug proportion seems to contain more doxycycline as time increases.

Moreover, the optimisation problem defined in (8.11) assumes that the optimal drug combination is fixed for the duration of the experiment, a property expressed by considering that $\sigma_i = \sigma_j$ for all $i, j \in \{1, N+1\}$. Although this is a standard assumption in the literature (for example, see [48]), the experimental system under consideration allows us the possibility of relaxing this suboptimal constraint, and hence define a different optimisation problem:

$$\mathcal{J}(\boldsymbol{x}, \boldsymbol{\sigma}^*_{opt}) = \max_{\boldsymbol{\sigma}\in[0,1]} \{\mathcal{J}(\boldsymbol{x}, \boldsymbol{\sigma}) : 0 \leq \sigma_i \leq 1 \text{ for all } i \in \{1, N+1\}\}. \tag{8.12}$$

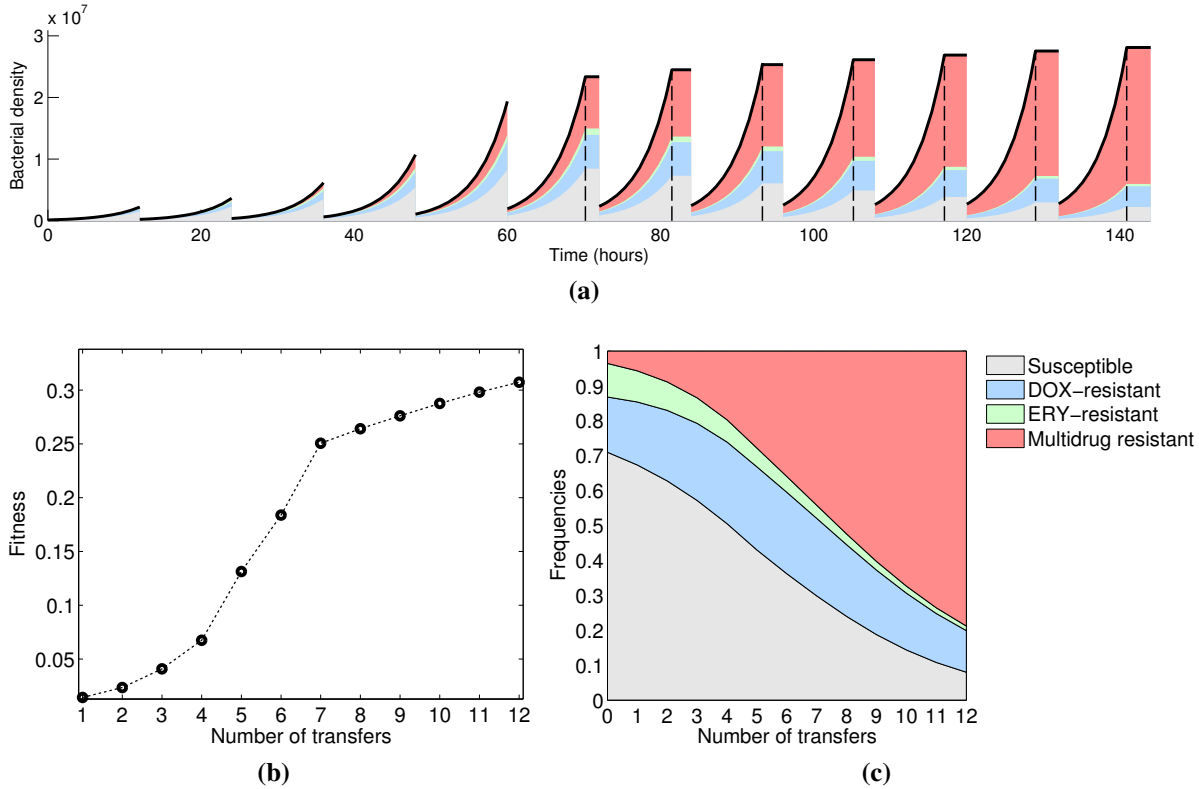

Figure 8.5: Simulations of a serial transfer experiment designed to test the optimal combination protocol. (a) Optical densities of a population of bacteria treated with a fixed combination of erythromycin and doxycycline. At the beginning of each transfer fresh medium is deployed at concentration $S = 2000\mu g/ml$, as well as a fixed dose of antibiotic with a drug ratio of $\boldsymbol{\sigma}^*_{fixed} = 0.57$. Cells are grown for 12 hours before propagating a fraction of the end-population into the next transfer. The dotted line illustrates the moment in time that the population enters the stationary phase and shaded areas represent the densities of each bacterial phenotype. (b) Fitness of the overall bacterial population as a function of the number of transfers. Note how the fitness increases monotonously as the population adapts to the fixed-drug environment. (c) Frequencies of resistance to each drug in the population as a function of the number of transfers performed. At the end of the first transfer the population is composed mainly by susceptible bacteria but at the end of the experiment more than eighty percent of the population is multidrug resistant.

Solving numerically (8.12) is not a difficult task either. We could use optimisation techniques like the genetic algorithm described in Appendix C, but if the number of transfers is not too large, it can even be solved by evaluating every 2^N possibility and identifying the control that maximises the payoff. In the example presented here, the optimal 12-transfer

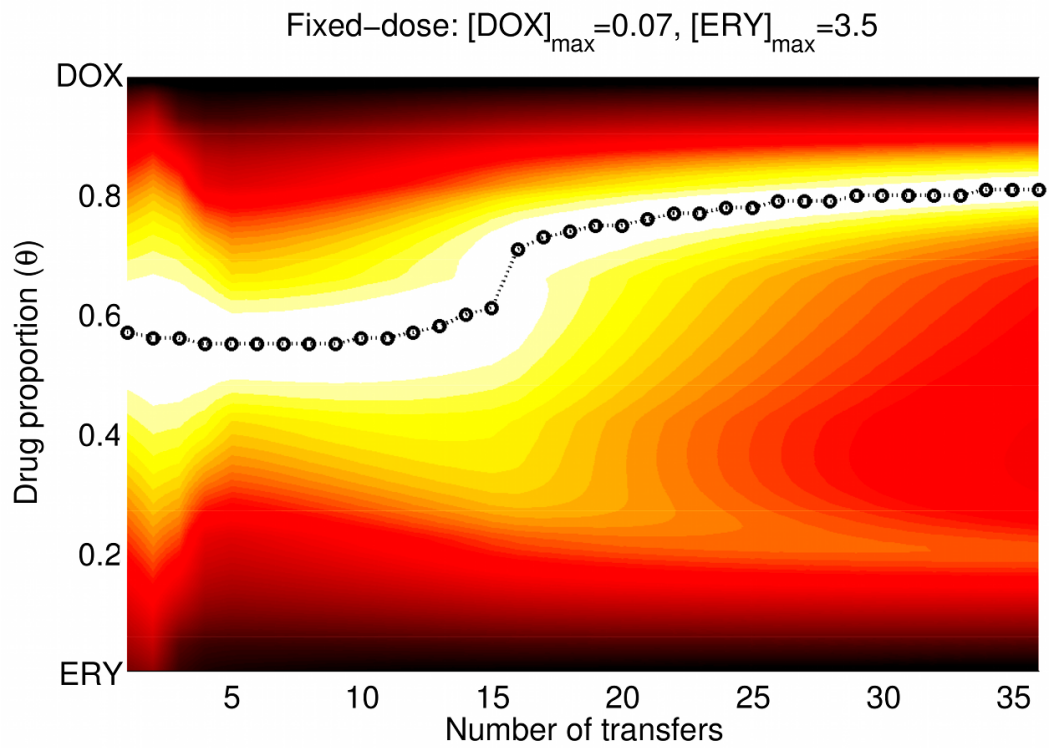

Figure 8.6: Optimal drug proportion as a function of the number of transfers performed. Darker colours represent a lower payoff while lighter colours denote that in average there is a low pathogen density at the end of each transfer. Circles denote the optimal drug proportion for an experiment with the corresponding number of transfers. Note that for short-term experiments the optimal drug proportion corresponds roughly with the drug ratio that maximises the synergistic effect. But as the number of transfers in the experiment increases and consequently bacteria start adapting to the fixed environment and therefore the population structure changes, the optimal drug ratio also changes. This figure illustrates that the optimal drug combination is a dynamic property of the system that depends on the population structure and the length of the experiment.

control is

$$\sigma_{opt} = (0.59, 0.54, 0.56, 0.55, 0.51, 0.99, 0.72, 0.61, 0.60, 0.53, 0.55, 0.48).$$

It is clear from the value of σ^* that computing the optimal control gives us very little insight into how these strategies could be implemented in practise. As discussed before, the theoretical optimal control is not robust to parametric and modelling uncertainties and hence there is no reason to believe that it would still be optimal in a *real* experimental setup.

We argue, however, that the difference in payoffs obtained from solving problems (8.12) and (8.11) is so considerable, that releasing the fixed-proportion constraint seems to be an effective strategy. In the interest of simplicity, we could, for example, allow the drug ratio to change over time, but imposing the constraint that we can only deploy a single antibiotic in each transfer. This *optimal rotation protocol*, that we will denote σ_{rot}, can be computed by solving the following optimisation problem:

$$\mathcal{J}(\boldsymbol{x}, \boldsymbol{\sigma}^*_{rot}) = \max_{\boldsymbol{\sigma}\in[0,1]} \{\mathcal{J}(\boldsymbol{x}, \boldsymbol{\sigma}) : 0 \leq \sigma_i \leq 1 \text{ and } \sigma_i = \{0, 1\} \text{ for all } i \in \{1, N+1\}\} \tag{8.13}$$

Note that because the set of admissible rotation protocols is much smaller than the set of admissible controls, then the optimal rotation protocol is suboptimal with respect to the optimal dynamic multidrug combination.

It is important to emphasise that ERY and DOX have a synergistic interaction, and that under a rotation protocol bacteria are exposed just to a single-drug. Then it is not a surprise that at the beginning of the experiment the optimal rotation protocol seems to be very ineffective at reducing bacterial growth rate, a feature illustrated in Figure 8.7. However, because of the fitness costs of resistance and the fact that rotation protocols do not impose strong selective pressures in favour of multidrug resistant pathogens, bacterial density at the end of the experiment is lower than the optimal fixed-dose multidrug combination protocol, as shown in Figure 8.8.

The conclusion we can draw from the exemplar situation presented in this chapter is that there is a universal structure to all optimal multidrug treatments: they are dynamic. In the following chapter we will expand on this idea and re-visit the colonisation resistance problem discussed in Chapter 6 with the purpose of designing optimal multidrug deployment strategies that support the commensal microbiota while driving the pathogens towards extinction in a continuous culture device.

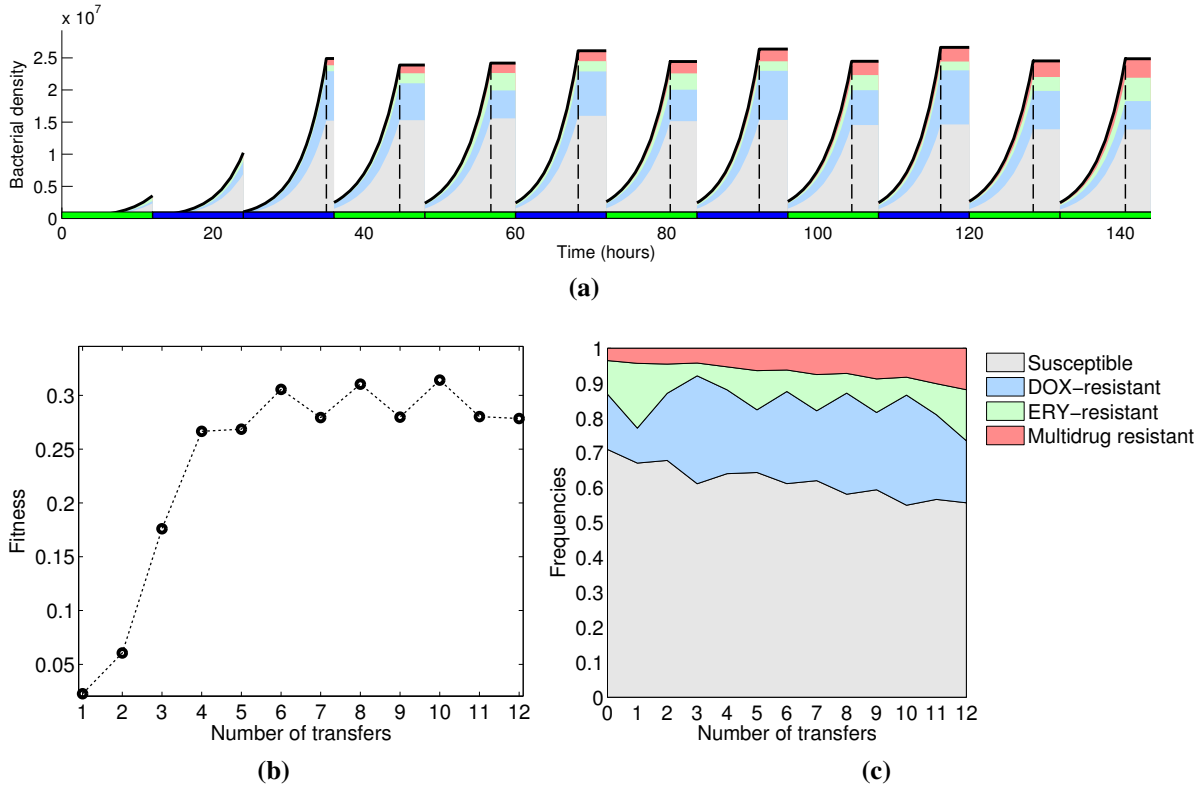

Figure 8.7: Example of a serial transfer experiment where only one drug is deployed at the beginning of each transfer. (a) The antibiotic used in each transfer is denoted by the colour code at the bottom of the figure; blue denotes that doxycycline was used at a concentration $0.7\mu g/ml$ while green denotes erythromycin at $3.5\mu g/ml$. The optimal rotation protocol illustrated here can be represented with the vector $\boldsymbol{\sigma}^*_{rot} = (0,1,1,0,0,1,0,1,0,1,0,0)$ and was determined after simulating all possible 12-transfer rotations. (b) Fitness of the population as a function of the number of transfers. It is important to notice that switches between antibiotics are associated with a reduction in the overall fitness of the population. (c) Frequencies of resistance as a function of the number of transfers. As opposed to the fixed-dose combination treatment that selected for multidrug resistance, in this case even after several dilutions the population is still mainly composed of susceptible bacteria.

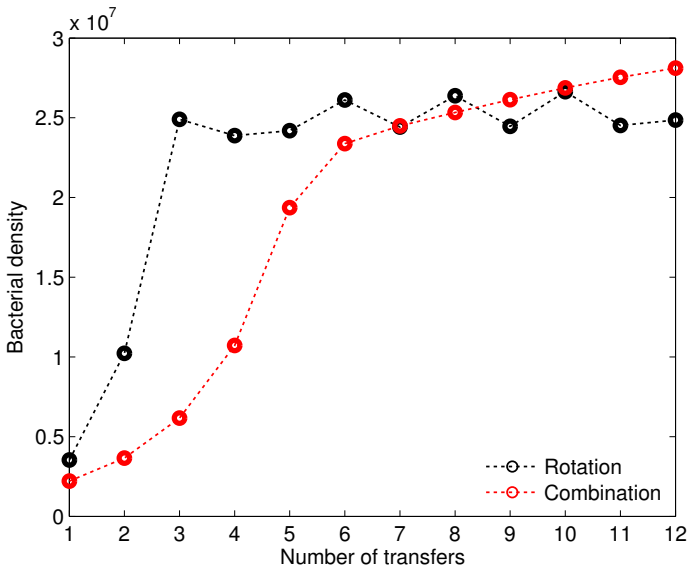

Figure 8.8: Optimal densities at the end of each transfer under two different drug deployment regimes: the rotation protocol illustrated in Figure 8.7 and the fixed-proportion combination protocol shown in Figure 8.5. At the beginning of the experiment the combination protocols minimises the density of bacteria considerably more than the rotation protocol. Interestingly, this trend is reversed at the end of the experiment, when the treatment protocol that minimises bacterial density is the one that deploys both synergistic antibiotics sequentially and not in combination.

Chapter 9

Optimal deployment of synergistic antibiotics into a chemostat

The purpose of this chapter is to pose an evolutionary model of commensal and pathogenic bacteria competing for a single limiting resource in a continuous culture device that allows us to control the input of two synergistic antibiotics in time. Exploiting well-known results from optimal control theory, we use an evolutionary algorithm to design antibiotic deployment protocols that support commensal bacteria while minimising the density of pathogens. The main result presented in this chapter can be stated thus: the optimal deployment of synergistic antibiotics to remove a pathogen in the presence of commensal bacteria in our model system occurs not in *combination*, but by deploying them sequentially.

9.1 Incorporating antibiotics into a continuous culture device

A basic chemostat like the one described in Chapter 6 could be readily adapted to have two supply vessels with different antibiotics, as shown in Figure 9.1, whereby each supply vessel contains the same concentration of abiotic resources so that the deployment of antibiotics does not unduly effect the underlying resource competition.

We can extend the model defined in equations (6.2) in order to consider the input of two

two bacteriostatic antibiotics of concentration A and B. Then the state of the system is now represented by the variable $\boldsymbol{x}(t) := (S(t), C(t), \boldsymbol{P}(t), A(t), B(t))$ and the evolutionary model for a population of commensal and an evolvable pathogenic bacteria competing in a single resource-limited environment and subject to the inhibiting effect of two bacteriostatic antibiotics of concentration A and B can be written as

$$\begin{aligned}
\dot{S} &= \overbrace{(\alpha+\beta)}^{d} S_0 - dS - u_c(S)\cdot C - \langle u_p(S), \boldsymbol{P}\rangle\,, && \text{(9.1a)}\\
\dot{C} &= G_c(S,A,B)\cdot C - dC, && \text{(9.1b)}\\
\dot{\boldsymbol{P}} &= \mathcal{M}_\epsilon(G_p(S,A,B)\cdot \boldsymbol{P}) - d\boldsymbol{P}, && \text{(9.1c)}\\
\dot{A} &= \alpha A_0 - dA - aA(C + \langle \mathbf{1}, \boldsymbol{P}\rangle), && \text{(9.1d)}\\
\dot{B} &= \beta B_0 - dB - bB\,(C + \langle \mathbf{1}, \boldsymbol{P}\rangle), && \text{(9.1e)}
\end{aligned}$$

with initial conditions $\boldsymbol{x}(0) = (S(0), C(0), \boldsymbol{P}(0), A(0), B(0))$. Parameters a and b represent the binding rates of each bacterial phenotype to each of the corresponding antibiotic molecules and $\mathbf{1}$ denotes a vector of ones $(1, 1, \ldots, 1)$ of the appropriate dimension. In the remainder, the system (9.1a-e) will be written

$$\dot{\boldsymbol{x}} = \mathcal{F}(\boldsymbol{x}; \alpha, \beta)$$

for brevity, with initial datum given by a non-negative vector $\boldsymbol{x}(0) = \boldsymbol{x}_0$. We want to control the input of both antibiotics into the system throughout an experiment of duration T, so our control variables will be $\alpha(t)$ and $\beta(t)$, where $0 \leq t \leq T$, and both control variables are constrained to lie between 0 and d.

9.2 Complete competitive advantage

In order to mimic a potentially dire situation for the commensal bacteria we shall be interested in the following situation. We shall suppose to begin with that the wild-type susceptible pathogen (denoted P_s) *outcompetes* the commensal bacteria (C) in the absence of antibiotics. We then assume that in the presence of antibiotics and at sufficiently high

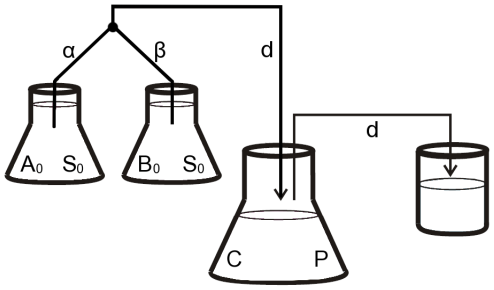

Figure 9.1: Diagram of a chemostat adapted to supply antibiotics A at a rate α and antibiotic B at a rate β, while maintaining the volume and supply rate of limiting resources at a constant rate. In this diagram, the constraint $\alpha + \beta = d$ must apply.

concentrations, the commensal bacteria are able to outcompete the wild-type pathogen. However, P_s will then be able to evolve resistance to the drug A, say, and the A-resistant pathogenic phenotype denoted (P_a) will outcompete the commensal and similarly, if drug B is high in concentration the B-resistant pathogenic strain (P_b) will also outcompete the commensal.

Moreover, in a multidrug environment we shall impose strong selective pressures in favour of those pathogens who are resistant to both antibiotics and thus the multi-resistant strain (P_{ab}) will outcompete C and may eventually colonise the host when both A and B are sufficiently high in concentration. Note that we are not assuming that commensals can evolve drug resistance, the reason being that we consider the commensals to be a community composed of many different species, and while resistance might evolve in some of these species, we are making the assumption that the commensal microbiota as a whole will behave as if it is mostly composed of susceptible bacteria.

Then in the situation just described, there is no value for the concentration vector of the two drugs (A, B) for which the commensals have the highest growth rate, as illustrated in Figure 9.2. Therefore, for any fixed environment, there exist at least one pathogenic strain that outcompetes the commensal type and we therefore say that pathogens have *complete competitive advantage*. In this case, we pose the following questions: is there a drug deployment policy that can make the commensals persist as time increases? And if one exists, can we design the optimal treatment protocol that minimises the density of pathogens?

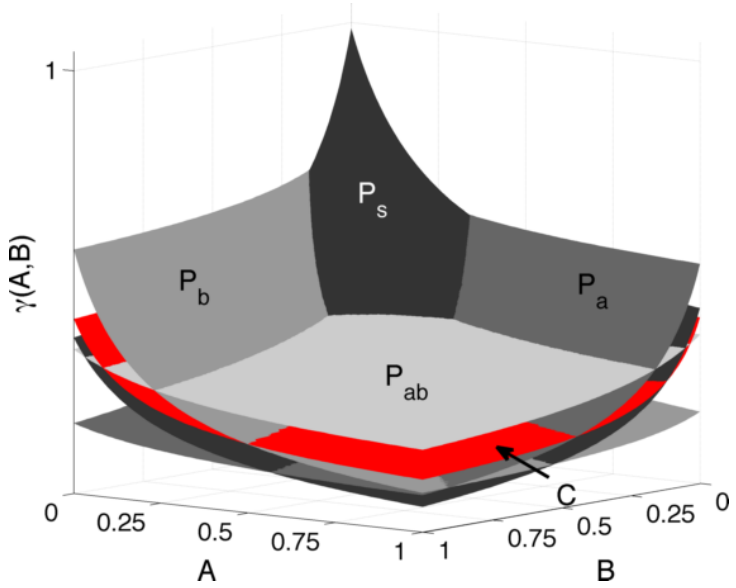

Figure 9.2: Growth rate inhibition surfaces for each bacterial phenotype as a function of the concentration of each antibiotic. Note that commensals are not the fittest bacteria in any possible fixed environment (and here S is held at a fixed value).

9.3 Controlling the chemostat optimally

To minimise the evolution of antibiotic resistance, it is evident that the optimal strategy is simply never use antibiotics, but this is suboptimal from the perspective of a single host, because without treatment the wild-type pathogen will colonise the host. Conversely, drug overdose selects for resistant phenotypes and as a consequence decreases the longer-term efficacy of the antibiotic. Then, for a drug deployment policy to be considered optimal, it has to be able to drive the pathogens to extinction whilst preserving its efficacy.

Taking into consideration the idea that we would like to support the commensal microbiota, we define the *healthy state* of the system whereby the commensal bacteria are in equilibrium and there are no pathogens present. Ideally we would drive the system towards the healthy state in such a way that we minimise antibiotic use but our present question is simpler: do there exist treatment protocols that drive the pathogen density to zero and the commensal to equilibrium even when the former has complete competitive advantage? Motivated by results and techniques from optimal control, we shall use a genetic algorithm to show numerically that this is possible by optimising the treatment protocol with respect

to the following objective functional:

$$\mathcal{J}(\alpha,\beta)=\int_0^T\left(C(t)-\sum_{j=1}^{n}P_i(t)\right)dt \quad \text{where} \quad \dot{\boldsymbol{x}}=\mathcal{F}(\boldsymbol{x};\alpha,\beta),\boldsymbol{x}(0)=\boldsymbol{x}_0. \tag{9.2}$$

We shall assume throughout the remainder that the maximum input rate of antibiotic is the same as the dilution rate of the chemostat, so that

$$0\leq\alpha,\ 0\leq\beta,\ \alpha+\beta\leq d.$$

Now, it is a well-defined problem to seek functions α and β that satisfy

$$\mathcal{J}(\alpha^*,\beta^*)=\max\left\{\mathcal{J}(\alpha,\beta)|\alpha\geq 0,\beta\geq 0,\alpha+\beta\leq d\right\}, \tag{9.3}$$

where the set of admissible controls for this problem is

$$\Omega:=\{(\alpha,\beta)\in L^\infty(0,T)\times L^\infty(0,T)\,|\,\alpha\geq 0,\beta\geq 0,\alpha+\beta\leq d\ \ a.e.\}.$$

However, in order to simplify computational aspects of the problem we also impose the exogenous, suboptimal constraint that we must deploy antibiotics at a constant rate, d, and so the control problem is to then determine only what proportion of both antibiotics we should use at any given moment in time. We define another set of admissible controls accordingly by

$$\Omega':=\{(\alpha,\beta)\in L^\infty(0,T)\times L^\infty(0,T)\,|\,\alpha\geq 0,\beta\geq 0,\ \alpha+\beta=d\ \ a.e.\}$$

and our second constrained optimisation problem is to determine the admissible control α^* and β^* such that the payoff is maximised:

$$\mathcal{J}(\alpha^*,\beta^*)=\max_{(\alpha,\beta)\in\Omega'}\mathcal{J}(\alpha,\beta). \tag{9.4}$$

The experimental protocol corresponding to the optimal control problem (9.4) is illustrated in Figure 9.1.

The dependence of equation (9.1a-e) on the control variables is affine in the sense that there is an smooth mapping $\mathcal{F}_0 : \mathbb{R}^{1+n+3} \to \mathbb{R}^{1+n+3}$ such that the function $\mathcal{F}(\boldsymbol{x}; \alpha, \beta) = \mathcal{F}_0(\boldsymbol{x}) + \alpha\boldsymbol{v} + \beta\boldsymbol{u}$ and the objective functional $\mathcal{J}(\alpha, \beta)$ in (9.2) is linear in the state $\boldsymbol{x}$ as it can be written in terms of a weight w vector,

$$\mathcal{J}(\alpha, \beta) = \int_0^T \langle w, \boldsymbol{x}(t)\rangle \, dt. \tag{9.5}$$

As a result, $\mathcal{J} : L^\infty \times L^\infty \to \mathbb{R}$ is continuous with respect to weak*, L^∞-convergence. Moreover, both Ω and Ω' are closed, convex subsets of $L^\infty \times L^\infty$ and so are compact with respect to the weak* topology on $L^\infty \times L^\infty$.

The solutions of (9.1a-e) satisfy a control-independent, dissipative bound of the following form: for each initial condition $x(0)$ there is a t' depending on $\boldsymbol{x}(0)$ such that

$$S(t) + \frac{1}{c}\left(C(t) + \langle \mathbf{1}, \boldsymbol{P}(t)\rangle\right) \leq S_0 + 1, \tag{9.6a}$$

$$A(t) \leq A_0 + 1, B(t) \leq B_0 + 1, \tag{9.6b}$$

for all $t > t'$. To see this, define $\Sigma := S_0 - S - \frac{1}{c}(C(t) + \langle \mathbf{1}, \boldsymbol{P}(t)\rangle)$ and note that the differential inequality $\frac{d}{dt}\Sigma \leq -d\Sigma$ follows from (9.1a-e), the bound in (9.6a) now follows. The bounds in (9.6b) are obtained from the inequalities $\frac{d}{dt}A \leq d(A_0 - A)$ and $\frac{d}{dt}B \leq d(B_0 - B)$ that satisfied by positive solutions of (9.1a-e) when $\max(\alpha, \beta) \leq d$, $\min(\alpha, \beta) \geq 0$.

Theorem 9.3.1 *There exists an optimal control* $(\alpha^*, \beta^*) \in \Omega'$ *such that* $\mathcal{J}(\alpha^*, \beta^*) \geq \mathcal{J}(\alpha, \beta)$ *for all* $(\alpha, \beta) \in \Omega'$.

Proof Seeking an optimal control that maximises the payoff functional defined in (9.5) subject to the system of differential equations given in (9.1a-e), suppose that

$$\sup_{(\alpha,\beta)\in\Omega} \mathcal{J}(\alpha, \beta) = \lim_{k\to\infty} \mathcal{J}(\alpha_k, \beta_k)$$

for a sequence $\{(\alpha_k, \beta_k)\} \subset \Omega'$ where $\beta_k = d - \alpha_k$, a supremising sequence of admissible controls. This sequence has an associated sequence of state responses $\{\boldsymbol{x}_k\} \subset$

$W^{1,\infty}((0,T),\mathbb{R}^{1+n+3})$ that satisfies $\frac{d}{dt}\boldsymbol{x}_k = \mathcal{F}(\boldsymbol{x}_k,\alpha_k,\beta_k)$ and $\boldsymbol{x}_k(0)=\boldsymbol{x}_0$ for each k.

Because $0 \leq \alpha_k + \beta_k \leq d$ almost everywhere, we can assume without loss of generality that

$$(\alpha_k,\beta_k) \overset{*}{\rightharpoonup} (\alpha^*,\beta^*) \in \Omega$$

as $k \to \infty$ because Ω' is compact with respect to the weak* topology in L^∞.

Due to the dissipative bound (9.6), there is an $M > 0$ independent of k such that $\|\boldsymbol{x}_k\|_\infty \leq M$ and so, using $\frac{d}{dt}\boldsymbol{x}_k = \mathcal{F}(\boldsymbol{x}_k;\alpha_k,\beta_k), \boldsymbol{x}_k(0) = \boldsymbol{x}_0$, we obtain a k-independent $W^{1,\infty}$-bound on $\boldsymbol{x}_k$. We may therefore assume without the loss of generality that $\boldsymbol{x}_k \overset{*}{\rightharpoonup} \boldsymbol{x}^*$ in $W^{1,\infty}$ and so, basic compact embedding results allow us to take the latter convergence in $C^0([0,T],\mathbb{R}^{1+n+3})$. As a result, the continuity of the nonlinear mapping $\mathcal{F}$ can be used to deduce that $\frac{d}{dt}\boldsymbol{x}^* = \mathcal{F}(\boldsymbol{x}^*;\alpha^*,\beta^*)$ and $\boldsymbol{x}^*(0)=\boldsymbol{x}_0$. Furthermore,

$$\begin{aligned}\lim_{k\to\infty}\mathcal{J}(\alpha_k,\beta_k) &= \lim_{k\to\infty}\int_0^T \langle w,\boldsymbol{x}_k\rangle dt = \int_0^T \langle w, \lim_{k\to\infty}\boldsymbol{x}_k\rangle dt \\ &= \int_0^T \langle w,\boldsymbol{x}^*\rangle dt = \mathcal{J}(\alpha^*,\beta^*)\end{aligned}$$

and so $\mathcal{J}(\alpha^*,\beta^*) = \sup\{\mathcal{J}(\alpha,\beta)\,|\,(\alpha,\beta)\in\Omega'\}$. □

We are interested in the idea of *antibiotic rotation*, defined as follows. Two admissible antibiotic deployment protocols α and β are said to be *in rotation* if only one of the antibiotics A and B is used at any one time:

$$\alpha(t)\cdot\beta(t) = 0 \tag{9.7}$$

for almost all t between 0 and T. If we define the set of admissible rotational protocols

$$\Omega_{rot} := \{(\alpha,\beta)\in\Omega' \mid \alpha(t)\cdot\beta(t) = 0\ a.e.\}$$

then Ω_{rot} is a weak* dense subset of Ω'. To see this, the following construction is useful.

Let $\mathcal{I}_N = (t_k)_{k=0}^N$ be a partition of the interval $[0,T]$ so that $t_0 = 0 < t_1 < ... < t_{N-1} <$

$t_N = T$ and let $PC(\mathcal{I}_N)$ be the space of real-valued, piecewise-constant functions on $\mathcal{I}_N$ taking only one of the values either 0 or 1 on each interval of the form (t_k, t_{k+1}). The set of admissible bang-bang controls

$$\Omega_{bb} := \{(\alpha, \beta) \in \Omega' \mid \exists N, \mathcal{I}_N \text{ and } a \in PC(\mathcal{I}_N) \text{ s.t. } \alpha(t) = d \cdot a(t), \beta(t) = d \cdot (1 - a(t))\}$$

is a subset of, and weak* dense in Ω' (see [103]) but $\Omega_{bb} \subset \Omega_{rot}$ and so Ω_{rot} is also weak* dense in Ω'.

The outcome of applying these rather standard control theoretic ideas is the result that rotational protocols can do as well as any within Ω':

$$\max_{(\alpha,\beta)\in\Omega'} \mathcal{J}(\alpha, \beta) = \sup_{(\alpha,\beta)\in\Omega_{rot}} \mathcal{J}(\alpha, \beta).$$

This means that antibiotic deployment protocols that rotate or alternate in turn between two antibiotics can perform as well as any protocol in Ω'. In particular, since Ω' contains the entire spectrum of combination therapies, from drug A-only to drug-B only, we see that rotational protocols are, in general, superior to ones that define a fixed proportion of each drug throughout the term of the experiment.

The set of rotational protocols is not weak* closed as it contains the non-rotational 50%-A, 50%-B combination therapy $\psi(t) = d/2$ in its weak* closure and the existence of an optimal control within Ω_{rot} cannot be established in general. Notice also that the optimal deployment protocol in Ω is not a rotation in general because Ω_{rot} is a much smaller set than Ω itself.

It follows from the fact that $\Omega_{bb} \subset \Omega'$ is weak* dense that

$$\max_{(\alpha,\beta)\in\Omega'} \mathcal{J}(\alpha, \beta) = \sup_{(\alpha,\beta)\in\Omega_{bb}} \mathcal{J}(\alpha, \beta).$$

This property in turn informs us that to find an approximate optimal deployment protocol $(\alpha, \beta) \in \Omega'$ we only need search within $PC(\mathcal{I}_N)$ for large enough N, this property is convenient for the numerical calculations performed using a genetic algorithm to locate

near-optimal controls in the next section.

9.4 Supporting commensal bacteria with rotational protocols

Let us consider a dynamic whereby pathogenic bacteria and commensal bacteria are introduced into the chemostat at $t = 0$. In order to colonise the chemostat, the pathogens have to outcompete the commensal bacteria and to mimic a colonisation resistance experiment we shall assume that the initial population of commensals and wild-type pathogens are at equal densities when $t = 0$ with few or no drug-resistant mutants. Recalling our working assumption that pathogens have a greater fitnesses than commensals in an antibiotic-free environment, the pathogens will colonise the host in the absence of antibiotics. In effect, we are forced into treating the chemostat host with antibiotics in order for the commensal to survive.

Now suppose that pathogens can evolve resistance to the two different antibiotics of concentrations through point mutations that occur at mutation rate $0 < \epsilon \ll 1$ and the accumulation of two mutations will confer pathogens with multidrug resistance. To reflect this, the mutation matrix for the population of pathogenic bacteria $\boldsymbol{P} = (P_s, P_a, P_b, P_{ab})$ is given by the irreducible, stochastic matrix $M_\epsilon = (1 - \epsilon)I + \epsilon M_P$, where I is the identity matrix representing clonal reproduction, ϵ the probability of mutation and M_P is a n-by-n matrix where the m_{ij} entry represents the mutation rate of bacterial type i into bacterial type j.

To make the situation particularly stringent on the commensal bacteria we shall impose the restriction that the commensal bacteria cannot evolve resistance to any of the antibiotics used. As mentioned before, this strong biological assumption is based on the idea of considering the commensal as a community composed mainly of susceptible bacteria. But also because we want to consider the worst possible scenario for the commensal population and, for this reason, we will also assume that the pathogen population has *complete competitive advantage* meaning there are no drug combinations for which the commensals are not outcompeted by at least one pathogenic phenotype.

As already discussed, by controlling the input of antibiotics A and B at rates $\alpha(t)$ and

$\beta(t)$ respectively, we will attempt to drive the pathogens to extinction while promoting the commensal microbiota. If we now denote the state of this system by

$$\boldsymbol{x} = (S, C, P_s, P_a, P_b, P_{ab}, A, B),$$

we aim to do this by taking advantage of the theoretical results of the previous section and numerically maximising the objective functional (9.5) within the space of rotational protocols where $w = (0, 1, -1, -1, -1, -1, 0, 0)$.

In the scenario of complete competitive advantage of the pathogen over the commensal considered here, some numerical simulations are shown in Figure 9.3 to illustrate the dynamics of different drug deployment mechanisms. The single deployment of one antibiotic will be ineffective and while the best combination strategy is very effective at suppressing pathogens, it imposes a severe cost on the commensal population. Indeed, the constant deployment of any two-drug cocktail at any fixed drug ratio will have as a consequence that one of the pathogens will colonise the system (see Figure 9.3). It is clear, therefore, that in order to find a drug deployment strategy that can allow the commensals to persist, we have to allow the ratio between drugs to change over time.

Although optimal treatments can be obtained for the treatment objective $\mathcal{J}(\alpha, \beta)$ either with an optimisation algorithm or by solving the Euler-Lagrange equations associated with this functional, in practice this can be a challenging computational task, particularly when T is large. However, the optimal control for $\mathcal{J}$ under the constraint that we must deploy antibiotic at all times is a rotational protocol that can be described as follows. Deploy one antibiotic at the beginning of the experiment until some as yet unknown moment in time when a switch to a different will be needed; a repetition of this process will then define a rotation protocol. While computation of near-optimal controls could be achieved using a variety of different optimisation methods, we chose the genetic algorithm described in Appendix C.

Figure 9.4 shows the results of applying this genetic algorithm to the model (9.1) where the parameter values are given in Table A.2, the result is a rotational protocol whereby the commensals have the highest density at the end of the experiment obtained by numerically

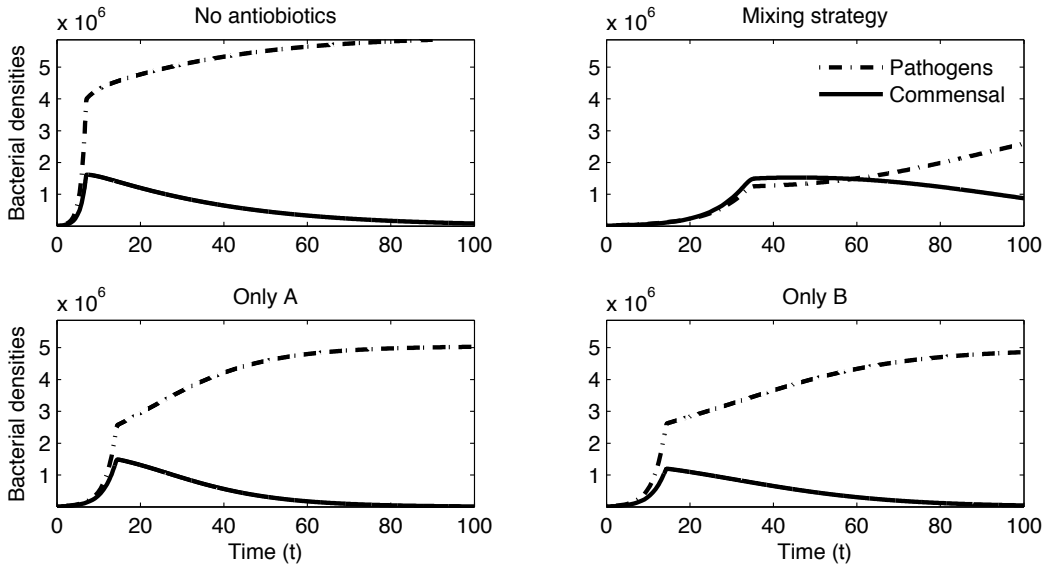

Figure 9.3: Numerical simulations with the parameter values given in Table A.4 showing the densities of pathogens and commensals throughout an experiment of duration $T = 100$ hours. (top-left) Bacterial densities when no antibiotics are used and drug-susceptible pathogens colonise the system. (top-right) Because the two antibiotics are synergistic, a 50-50 mixing strategy (equal concentrations of each drug) is very effective at suppressing the density of pathogens. However, commensals are also suppressed, eventually allowing the multidrug-resistant pathogen to colonise the host. (bottom-left) and (bottom-right) Single-drug protocols are particularly ineffective, allowing resistant pathogens to colonise the system.

maximising the treatment objective $\mathcal{J}$. The rotational protocol so-obtained is just one from an infinite family of suboptimal controls that outperform all the fixed-ratio combination protocols with respect to the treatment measure $\mathcal{J}$, as shown in Figure 9.4b).

It is important to note that to determine the optimal control, or even a near-optimal control, we essentially need to have access to complete information about the future dynamics of the system which are not likely to be available in practise. Determining the temporal dynamics of drug resistance would be very difficult, necessitating vast amounts of genomic data to be processed. An alternative to searching for the optimal protocol could be to implement rotational treatment protocols based on on partial information of the present state of the system as part of a feedback control. Depending on our precise goal, we could implement such a feedback as part of a heuristic that allows us to decide upon the best drug usage strategy at each moment in time, based on the information we currently have access

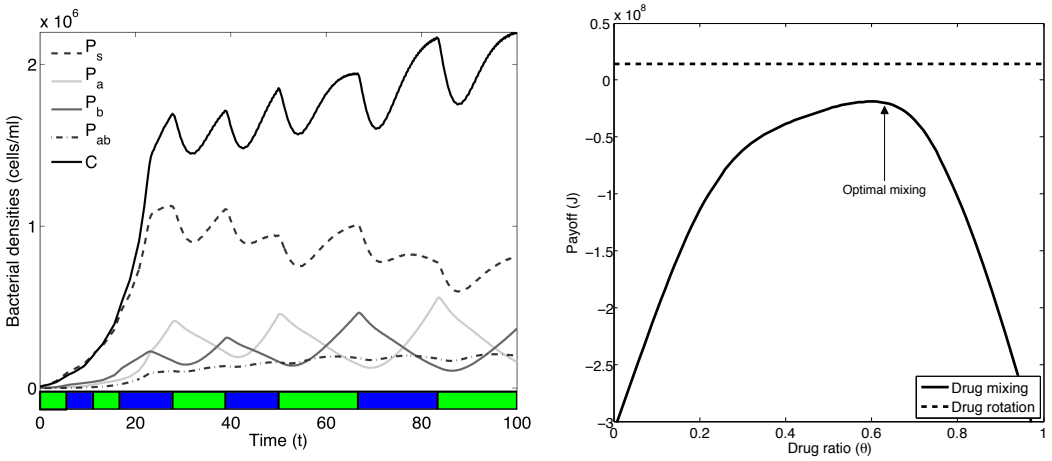

Figure 9.4: (a) Bacterial densities under an antibiotic deployment regime determined using a genetic algorithm (see Appendix C for details). The rotation protocol is represented by the barcode in the bottom of the figure, a blue (or green) box represents the time that drug A (or B) is being deployed at maximal dose during that time interval. Note that commensal bacteria have the highest density of all the types present, despite never being the fittest cell type. (b) The solid line represents the treatment objective $(\mathcal{J})$ of every possible combination with a fixed ratio between drugs (drug ratio, called σ in the text). The dashed line shows the payoff of the rotational protocol illustrated in (left). Although it is not the $\mathcal{J}$-optimal treatment, it outperforms even the optimal combination strategy.

to. For example, consider the simple feedback controller contained in the rule described immediately below:

1. deploy a single antibiotic into the chemostat;
2. after τ units of time, measure the levels of resistance to the current antibiotic used, we call τ the sampling time;
3. if the density of pathogens resistant to that antibiotic is larger than the sum of the densities of the other phenotypes, switch to the other drug. Otherwise, continue the deployment of the same drug for another τ units of time. Repeat the process from rule (2).

Surprisingly, this feedback heuristic can be sufficient to allow the commensals to persist and so drive the pathogens to extinction (see Figure 9.5). This illustrates that even suboptimal

rotational protocols can be effective at treating the presence of the pathogen in situations where even the best combination treatment will not work, as shown in Figure 9.3(top right).

However, the practical implementation of this feedback rule is contingent upon knowing the strain responsible for a bacterial infection and then knowing the relative densities of strains carrying resistance alleles or plasmids. This is biologically detailed information, although the increasing availability of tools to determine such information [11] would make the implementation of such a feedback a practical possibility.

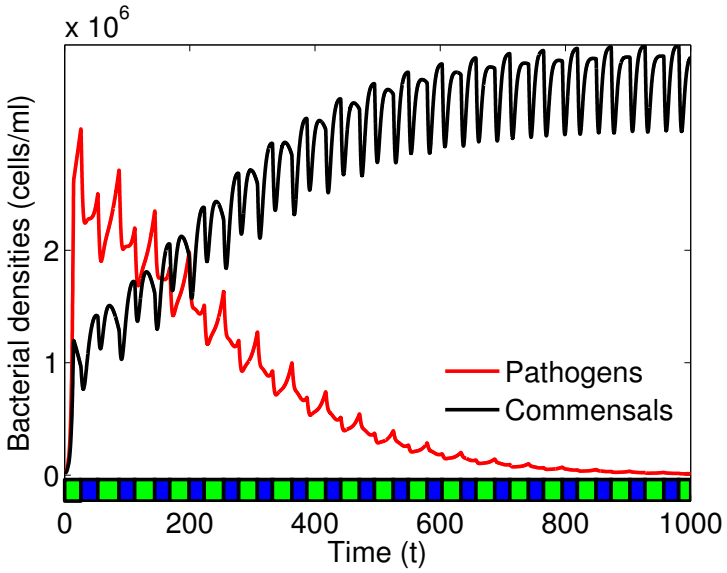

Figure 9.5: Numerical simulations of a long-time experiment ($T = 1000$) of a complete competitive advantage scenario where the numerical values of the parameters used are given in Table A.4. In this example, commensals are able to persist with a drug rotation protocol designed using a feedback control heuristic. The rotational protocol is illustrated with the dark boxes denoting drug B and the light boxes denoting drug A at the bottom of the figure, the dashed line represents the sum over all pathogenic phenotypes while the solid line represents the density of commensals through the duration of the experiment.

Part III

Selecting against antibiotic resistance in clinical settings

Chapter 10

An individual-based model of a hospital ward

The purpose of the last part of this thesis is to pose an individual-based model of a hospital ward that will allow us to test *in silico* different antibiotic deployment protocols used in clinical trials that have taken place over the last decade. Unfortunately there is no standard form of hospital interventions upon which we can focus to develop this model, but the following brief survey highlights features of policies that have been trialled in practise, and that we will be able to implement in our computational framework and thence consider their relative merits.

10.1 Stewardship programs to control resistance in health care centres

The scheduled rotation of one antibiotic for another discussed in the first part of this thesis is not the only type of policy trialled in hospitals. A striking example of this can be seen in the number of different policies trialled in the treatment of ventillator-associated pneumonia as described in [91]. First a ten-month, patient-specific strategy was implemented using multiple drug treatment, a twenty-four month period of scheduled rotation based on

four-month prioritisation and restriction cycles then followed, ending with a ten-month period of so-called *antibiotic mixing* whereby the first-line antibiotic was changed in each consecutive patient.

The so-called PAMS methodology (periodic antibiotic monitoring and supervision) [105] is not based on predetermined schedules of rotation but, as stated in [91], the basis of PAMS is the adaptation of observed prescription patterns aimed at '... balancing the use of different antimicrobials ... to reduce selection pressure'. The authors of [91, 105] achieve this by seeking to maximise the heterogeneity of the drugs prescribed. This essentially means that antibiotics used in the recent past must have a lower probability of being prescribed in the near future and a numerical index (a so-called *Peterson index*) is used to measure the heterogeneity of drug prescription and a committee was given the task of ensuring that prescription patterns maximise this index as the trial proceeds. While this rationale is ignorant of the patterns of resistance that actually emerge during the trial, it is responsive to the patterns of prescription that might be correlated with increased resistance and like the theoretical antibiotic mixing protocols proposed in [13, 15], it does maximise the heterogeneity of drug use.

While the PAMS methodology does seek to exploit information about previous drug use to inform future policy, some studies have gone further in seeking collated, drug susceptibility data as the basis for that policy. For example, patterns of drug prioritisation and withdrawal have been set in response to previously observed levels of susceptibility and resistance in an ICU unit [3]. Of their rotation-based *microbiological surveillance* methodology, the authors of this article stated that 'A non-premeditated change of antibiotics in empirical therapy, on the basis of detected resistance patterns, provided promising results in reducing some antimicrobial resistance rates.'

The individualisation of treatments can militate against the inappropriate use of antibiotics, important because the wrong initial choice of antibiotics in empirical treatment has been shown to increase the likelihood of patient death, for example in the case of intubation-associated pneumonia [35]. As a result, this study recommended effecting a change in antibiotics based on individual susceptibility measurements using patient-specific microbiological data, data that became available between two and three days after admittance to

the unit. However, even a delay of this length in providing a basic understanding of the pathogens responsible for infection is known to be associated with an increase in patient mortality across a range of multidrug resistant bacteria [87]. Furthermore, initially inappropriate therapy for non-nosocomial methicillin-resistant *Staphylococcus aureus* has been observed in as many as 77% of patients admitted [93].

Dissatisfied with such delays, the authors of [11] proposed the exploitation of detailed microbiological information based on the rapid, within one hour, DNA-based identification of microorganisms and the resistance genes they carry. It is known that the evolution of antibiotic resistance can occur *in vivo* in a matter of weeks, so rapidly that over thirty resistance mutations were observed in isolates of MRSA from a single patient and identified using whole-genome sequencing [78]. The sequenced mutations were correlated with changes in the MICs for multiple drugs, including rifampin, vancomycin, oxacillin and daptomycin. As the rapidity of sequencing technology increases and costs decrease, rather than wait until treatment fails to assess the evolutionary path taken by the infection, why not use this technology as a diagnostic tool to exert some control over that path?

This brief assessment of the clinical literature suggests at least five classes of antibiotic deployment protocols that have been used in clinical trials and that we can implement in a mathematical model. Collated according to the increasing degree of information required for their implementation in an ICU unit or hospital, they are:

- antibiotic mixing* based on random drug prescription (see [13, 15]),
- scheduled, rotating cycles of antibiotic prioritisation and restriction (see [47, 70, 84]),
- antibiotic mixing based on prior patterns of prescription (see [91, 105]),
- surveillance-based rotation determined on previously observed, emergent patterns of resistance (see [3]) and
- diagnosis-based treatments determined using patient-specific susceptibility, assessing the appropriateness of treatment at least once (see [11, 93]).

*The term *mixing* will refer to any protocol whereby two patients simultaneously diagnosed with infection by the same pathogen may be purposefully prescribed different drugs.

This list is not meant to be exhaustive but has been collated to test the hypothesis that policies based on more information and which increase overall drug appropriateness as a result are necessarily of greater efficacy. This list does not encompass all the protocols used in practise, for example the de-escalation protocol used in the treatment of pneumonia [46] uses an initially broad-spectrum empirical treatment in advance of receiving the results of microbiological cultures. Nor, for example, does our list include the patient-by-patient rotation of the first-line antibiotic implemented in [91].

Note that although our modelling framework allows us the possibility of considering that single patients can be treated with a cocktail of multiple antibiotics like the ones discussed in Part II of this thesis, we quite purposely have not included this treatment protocol in our assessment of the clinical literature, but we will discuss this strategy later.

10.2 Modelling patients as single-hosts

As opposed to classic epidemiological models where the population is subdivided in compartments and the dynamics of the system is given by the transition rates between compartments, here we will consider explicitly patients as individual hosts interacting through the transmission of pathogens.

Individuals can be colonised both by commensal microbiota and by an evolvable pathogen. In order to clear the pathogens from the host, patients can be treated with a combination of two different antibiotics of concentrations that we will denote with the variables $A(t)$ and $B(t)$; the details of the treatment protocol will be determined by the hospital-wide strategy of drug usage.

Let C denote the density of commensal bacteria and $\boldsymbol{P} = (P_s, P_a, P_b, P_{ab})$ a vector containing the densities of different phenotypes of pathogenic bacteria, where P_s represents a susceptible pathogen, P_a and P_b pathogens resistant to antibiotics A and B respectively and P_{ab} a multidrug resistant pathogen. Phenotypic mutation rates during clonal reproduction for pathogenic bacteria are described by the irreducible and stochastic mutation matrix $\mathcal{M}_\epsilon = (1-\epsilon)I + \epsilon M$, where I is the identity matrix representing clonal reproduction,

$0 < \epsilon \ll 1$ the probability of point mutation and M is a 4-by-4 matrix where the m_{ij} entry represents the mutation rate of pathogenic type i into phenotype j:

$$\mathcal{M}_\epsilon = \begin{pmatrix} 1-\epsilon & \frac{1}{2}\epsilon & \frac{1}{2}\epsilon & 0 \\ \frac{1}{2}\epsilon & 1-\epsilon & 0 & \frac{1}{2}\epsilon \\ \frac{1}{2}\epsilon & 0 & 1-\epsilon & \frac{1}{2}\epsilon \\ 0 & \frac{1}{2}\epsilon & \frac{1}{2}\epsilon & 1-\epsilon \end{pmatrix}$$

Growth rate of each bacterium type is not only determined by the concentration of available resource, denoted by S, but also by the concentration of antibiotics present in the host. Therefore growth rate of bacterial type i can be modelled as a resource uptake function multiplied by an inhibition term that depends on the concentration of both drugs:

$$G_i(S, A, B) = u_i(S) \cdot \gamma_i(A, B), \tag{10.1}$$

where $\gamma_i(A, B)$ is the *growth inhibition function* derived from kinetic properties of the interaction between bacterium type i and both antibiotics, as discussed in Chapter 7, and $u_i(S)$ is the standard Monod function with a maximal growth rate V^i_{max} and a half-saturation constant K_i:

$$u_i(S) = \frac{V^i_{max}\, S}{K_i + S}, \tag{10.2}$$

We will assume that both drugs have a *synergistic* interaction and that we can model them as non-exclusive competitive inhibitors. This implies that the pharmacodynamics of each drug is acting on independent metabolic pathways or different parts of the cell division process. Then the multidrug growth inhibition function can be written as

$$\gamma_i(A, B) = \left(1 - \frac{\kappa^1_i\, A}{\kappa^2_i + A}\right) \cdot \left(1 - \frac{\tilde{\kappa}^1_i\, B}{\tilde{\kappa}^2_i + B}\right), \tag{10.3}$$

where κ_i^2 and $\tilde{\kappa}_i^2$ represent the affinity of the cell to antibiotics A and B respectively, while κ_i^1 and $\tilde{\kappa}_i^1$ control the level of maximal growth inhibition by each drug.

As discussed in Section 5.3, each bacterial phenotype has different profiles of drug resistance, a property determined by the set of parameters $(V_{max}^i, \kappa_i^1, \tilde{\kappa}_i^1, \kappa_i^2, \tilde{\kappa}_i^2)$. For instance, we say that P_a is resistant to drug A because $0 \leq \gamma_s(A, \cdot) < \gamma_a(A, \cdot) \leq 1$ when $A \gg 0$. Fitness cost associated with drug resistance at low antibiotic concentrations is also considered, as we can choose parameters in such a way that $G_s > G_a$ at low concentrations of A. Analogous statements can be made about the B-resistant and the AB-resistant phenotypes.

It is important to notice that we are considering that antibiotics can be introduced into the host continuously, as opposed to a discrete dosage schedule. The reason being that critically ill patients in ICU units are usually treated with intravenous continuous infusion. Then we will denote with $A(t)$ and $B(t)$ the concentration of each drug in a given patient at time t and with $x = (S, C, \boldsymbol{P}, A, B)$ the state of the system, the evolutionary model for a population of commensal and pathogenic bacteria competing for limited resources and subject to the inhibitory effects of both antibiotics can be written as:

$$\begin{aligned}
\dot{S} &= m(S_0 - S) - c\left(u_c(S)C + \langle u(S), P\rangle\right) \\
\dot{C} &= G_c(S, A, B)C - dC \\
\dot{\boldsymbol{P}} &= \mathcal{M}_\epsilon(G(S, A, B) \cdot \boldsymbol{P}) - d\boldsymbol{P} \\
\dot{A} &= \alpha A_0 - mA - aA(C + \sum_i P_i) \\
\dot{B} &= \beta B_0 - mB - bB(C + \sum_i P_i)
\end{aligned} \tag{10.4}$$

with initial conditions $x(0) = (S_0, C_0, P_0, A_0, B_0)$. Parameters a and b represent the binding rates of each bacterial phenotype to each of the corresponding antibiotic molecules, c denotes a resource conversion constant and m the constant influx and efflux of resources from the system. Bacteria are also cleared from the system at a constant rate d.

The treatment protocol is represented by the control functions $\alpha(t)$ and $\beta(t)$, that represent the input rate of antibiotics A and B into the host. Both control functions are bounded, $0 \leq \alpha(t) \leq 1$ and $0 \leq \beta(t) \leq 1$, and A_0 and B_0 denote the maximum dose allowed for

each drug. Drugs are also washed out of the system at a constant rate m.

Now, if we define the *health state* of a given patient as:

$$h(t) = \frac{C(t)}{C(t) + \sum_i P_i(t)} \tag{10.5}$$

then our goal will be to treat the patient until there are almost no pathogens present and the commensal population has re-colonised the host. Let us define $t = T$ to be the first time when $h(t) > \delta$ for a given $\delta > 0$, and we then say that the Length of Stay (LoS) of this patient is T units of time.

10.3 A non-deterministic, spatially explicit model of a hospital ward

Let us consider that a hospital ward is composed of a linear array of B beds, each one denoted by b_j where $j \in \{1, B\}$. Now let $\mathbf{Q} = \{p_1, p_2, \ldots, p_n\}$ represent a randomised queue of n patients, where each individual p_i is colonised by initial densities of commensal and pathogenic bacteria.

When admitted into the hospital each patient is assigned to an available bed, where it will receive treatment until the health state defined in (10.5) is above a certain threshold and then the patient is discharged from the hospital. Then the next patient from the queue is assigned to the available bed where it will also receive treatment until healthy. When every patient from the queue has been discharged from the hospital ward, we then calculate the Mean Length of Stay (MLoS) of this ordered queue associated with this specific treatment protocol.

Also, every τ units of time a proportion β of bacteria present in each patient are transmitted between hosts of neighbouring beds. The reason we are assuming a discrete transmission dynamics over a continuous within-host dynamics is because we are considering that most of the transmission of bacteria between hosts is due to movement of patients and health care workers from bed to bed.

Note that there are two main sources of stochasticity within our model: the initial frequency

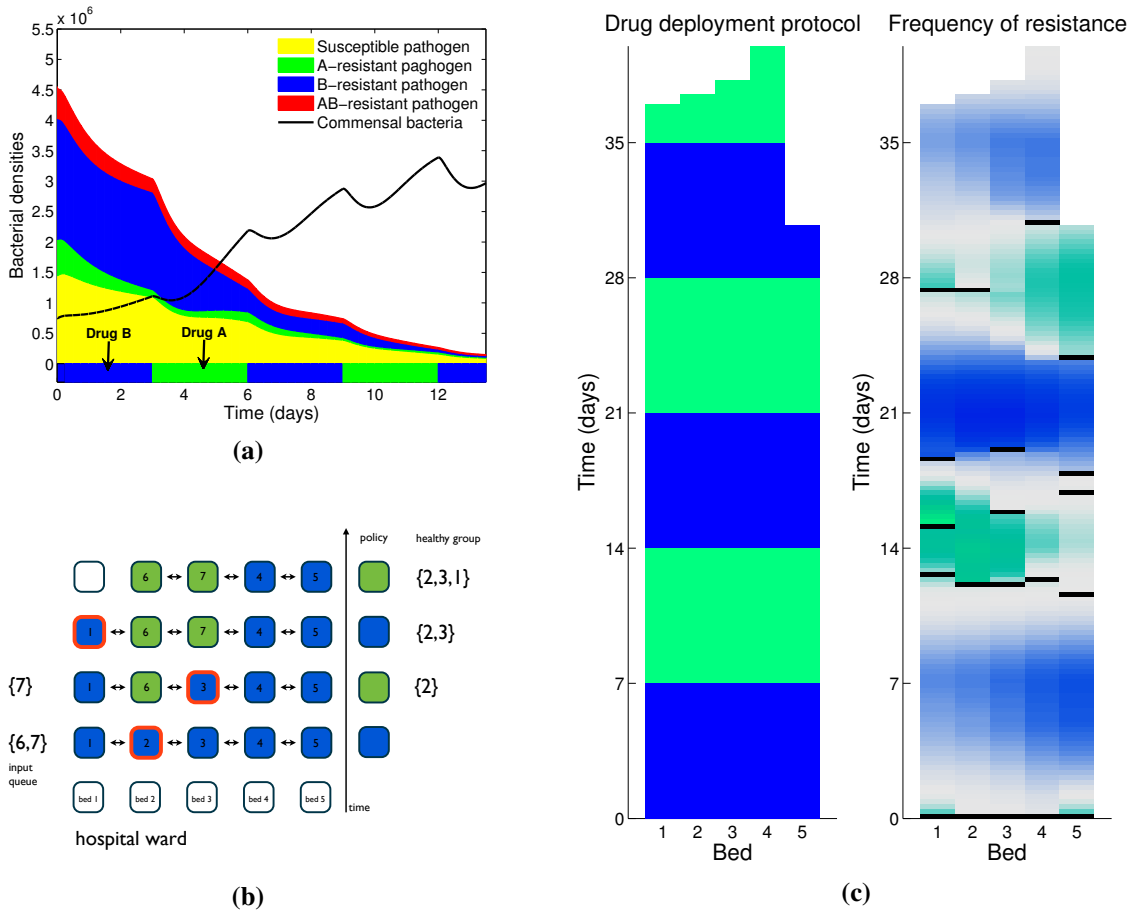

Figure 10.1: (a) Health state of a single patient treated with a cycling protocol. Coloured areas show the density of each pathogenic phenotype as a function of time, while the black line represents the density of commensal bacteria. (b) Diagram of a hospital ward with 5 beds where a queue of seven patients (here labelled $\{1, 2, 3, 4, 5, 6, 7\}$ for convenience) is to be treated (c) This Manhattan plot illustrates one realisation of the model using a queue of 20 patients and 5 beds where a rapid 20-day cycling protocol between the two drugs has been implemented. The left-hand illustration show the protocol used where the green and blue boxes denote patients (beds) treated each day with one of the two available drugs and a red outline denotes the discharge of a recovered patient followed by the arrival of a new member from the treatment queue. The relative frequency of each of the two single-drug resistant pathogens in each patient is shown in the right-most diagram and illustrates that drug resistance is correlated with the current deployment policy; the grey boxes in this plot show patients carrying equal fractions of two different single-drug resistant pathogens.

of resistance of each patient of the queue and the order of the queue itself. The former is very important as it represents the level of community-acquired resistance, but the latter is

also relevant because the position of the patient in that queue can lead to a change in the treatment administered to that patient. This implies that by re-ordering the queue of patients and implementing the same drug usage strategy multiple times, we can obtain meaningful statistics of the MLoS, allowing us to compare this antibiotic deployment strategy with others implemented to the same queue of patients.

In summary, the key assumptions of our model are the following: (1) patients are colonised both by commensal and pathogenic bacteria, if the pathogen density is very low (below a measurable threshold) the patient is discharged from the ward; (2) pathogens can evolve resistance to any drug used through a single-point mutation and the accumulation of multiple mutations confers pathogens with multi-drug resistance. There is a fitness cost associated with drug resistance expressed by resistant pathogens having a lower growth rate than susceptible bacteria in the absence of that antibiotic; (3) there is no infection-induced mortality and, if left untreated, all patients will eventually recover, although more slowly than if treated; (4) patients are always treated with one of two possible drugs, the drug prescribed to each patient at any moment in time is decided based on the hospital-wide strategy under consideration; (5) this model explicitly considers the transmission of pathogens between patients of neighbouring beds.

The transmission dynamics of pathogens within the ward has a significant impact on the efficacy of different drug usage strategies. For example Figure 10.1c) shows how epidemic waves of drug-resistant pathogens pass through the ward as they are transmitted rapidly from one bed to its neighbour under a cycling protocol. This dynamic forces the model to exhibit the following property common to the theoretical models used in the previous section: from the moment one drug is restricted and the other prioritised at the end of the cycle, the rate of ascent of resistance to the new drug is greater than the rate of descent of resistance to the previous drug.

10.4 Evaluating the efficacy of different hospital-wide usage strategies

The strength of this modelling approach is that even though we may not be able to determine exactly the optimal drug deployment, we can simulate specific drug usage strategies and obtain statistics about their performance, from hospital-wide protocols that restrict and prioritise drugs based on predefined schedules, to individual-based strategies that decide which drug to give to each patient based on local patterns of susceptibility and resistance.

We simulated $N = 1000$ replicates of different hospital interventions for a randomly generated queue of $n = 100$ patients receiving treatment in an 8-bed hospital. The protocols simulated are the following:

- **Random sequential treatment:** each patient receives a random drug each day.
- **Empirical treatment:** a random drug is allocated to the patient and determined the moment the patient arrives at the hospital.
- **Periodic cycling:** scheduled, rotating cycles of antibiotic prioritisation and restriction are fixed before any patients have entered the hospital.
- **Periodic antibiotic monitoring and supervision:** the next patient to arrive at the hospital will be treated with the drug that maximises the heterogeneity of drugs currently used within the hospital.
- **Surveillance-based rotation:** drug-resistance assays are regularly conducted on samples from the patient population and the same drug, namely the one estimated to have the lowest incidence of resistance, is prescribed to all patients.
- **Dynamic, diagnosis-based treatments:** an initial and rapid assessment of the bacterial responsible for each patient's infection is made the moment a patient arrives at the hospital, a reassessment may also be conducted later during the treatment.

Our computational results are summarised in the histograms of Figure 10.3a which shows the per-patient mean length of stay (MLoS) distribution. One treatment is said to be more

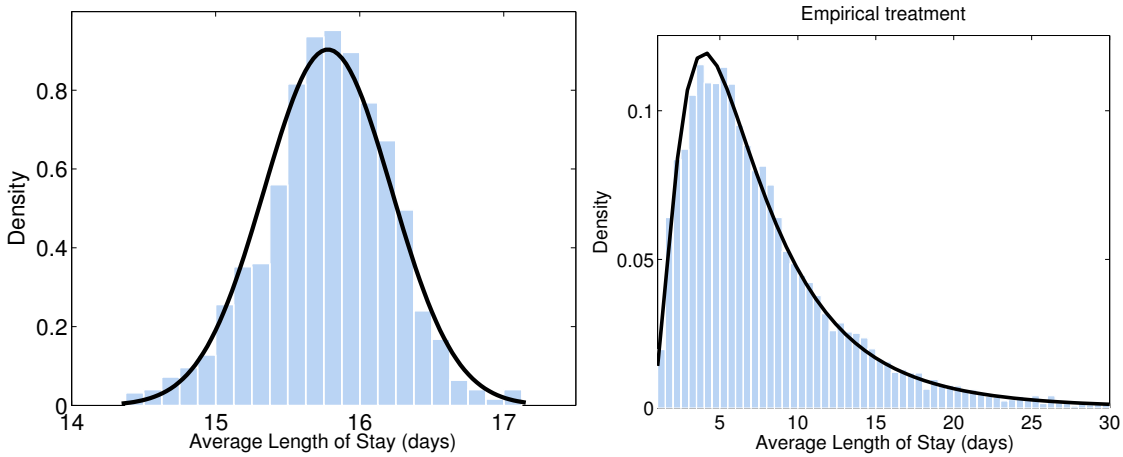

Figure 10.2: (a) The mean length-of-stay distribution for the null protocol applied to multiple re-arrangements of the patient queue: a near-normal distribution with mean close to 16 days. (b) The length of patient stay of the empirical treatment follows a log-normal distribution.

effective than another for a given patient queue if the mean length of stay of patients in that queue is smaller. As a matter of completeness, we note the following: if we do not treat any patients from the queue when they enter the hospital (the null protocol), the mean length of stay is normally distributed with respect to different arrangements of the queue order, as shown in Figure 10.2a; moreover, the distribution of length of patient stay in the simulations follows an approximately log-normal distribution for each queue, as illustrated in Figure 10.2b.

There are three main features of note in our computational results illustrated in Figure 10.3a: (i) The empirical treatment distribution has a very long tail and so this protocol can be very ineffective which is largely due to the number of patients receiving inappropriate therapy in this protocol. The property that inappropriate treatment leads to longer MLoS distributions is also illustrated in Figure 10.3b and Figure 10.3c. (ii) Diagnosis-based treatments whereby the appropriate drug is prescribed to each patient each day perform better than all the other treatments simulated. (iii) Surveillance-based rotation requires less information to implement than diagnosis-based treatments but it performs better than the antibiotic mixing protocols.

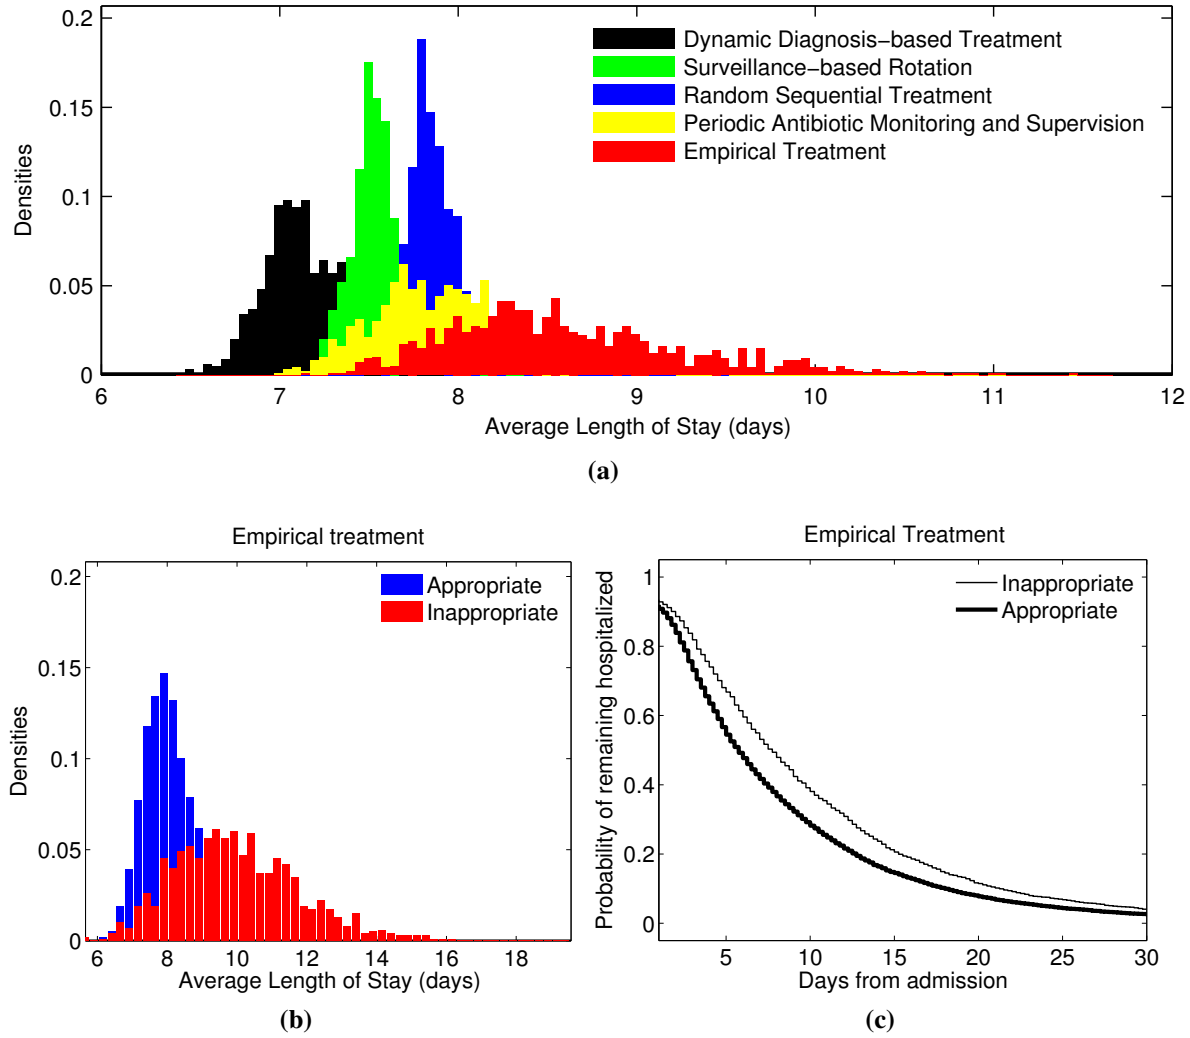

Figure 10.3: (a) The length of stay distributions of different scheduled treatment policies: diagnosis-based treatment, surveillance-based rotation, random sequential treatment, PAMS and an empirical-only control treatment. (b-c) Mean length-of stay and the inappropriate use of antibiotics: (b) the mean LoS per queue is normally distributed for both classes, with higher mean and variance for the inappropriate drug class. (c) A Kaplan-Meier plot illustrating the length of stay for each class.

And what of antibiotic cycling? Upon simulating strictly cyclical periods of drug rotation across the hospital we found that appropriately chosen rapid cycles could prove as effective as some of the patient-based diagnosis treatments. However, protocols based on cycles of longer duration that might be implemented in practise, up to one month say, had the effect of doubling the mean length of stay for the patient queue, as illustrated in Figure 10.4.

This makes the scheduled rotation of different antibiotics an impractical and potentially dangerous solution because it would be essential to know in advance of implementing such a protocol which particular duration would be appropriate for each particular pathogen-drug combinations [8].

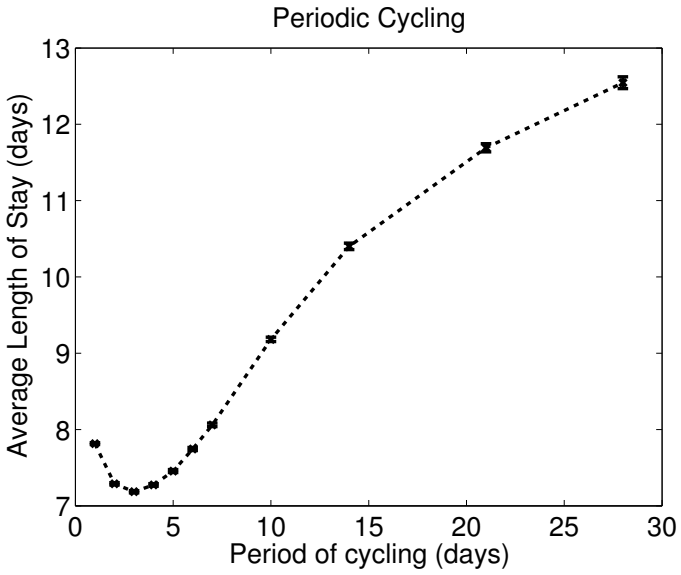

Figure 10.4: MLoS of antibiotic cycling as a function of the cycling period. Scheduled rotation may work, but only when cycling at the correct cadence: very rapid cycling performs as well as a random sequential treatment policy, longer cycles lead to a rapid increase in the expected length of stay. Without the benefit of hindsight, these results do not support the use of this protocol.

The fact that random drug deployment is easily outperformed by patient-specific, diagnosis-based treatments is perhaps not a surprise, as any drug deployment policy that increases the appropriateness of drug usage is sure to decrease the time to recovery from infection. However, more interesting is the question of how much information do we actually need in order to maintain the effectiveness of the diagnosis-based protocol?

To probe this question we incorporated into our simulation the following delay in determining drug-susceptibility information: we assumed that from the moment each patient is admitted to hospital, results are only available from the bacteriology lab some time later, from one to fourteen days. Before this time, an empirical treatment is given, after this time the most appropriate drug upon admission is then prescribed. Figure 10.5a shows that there

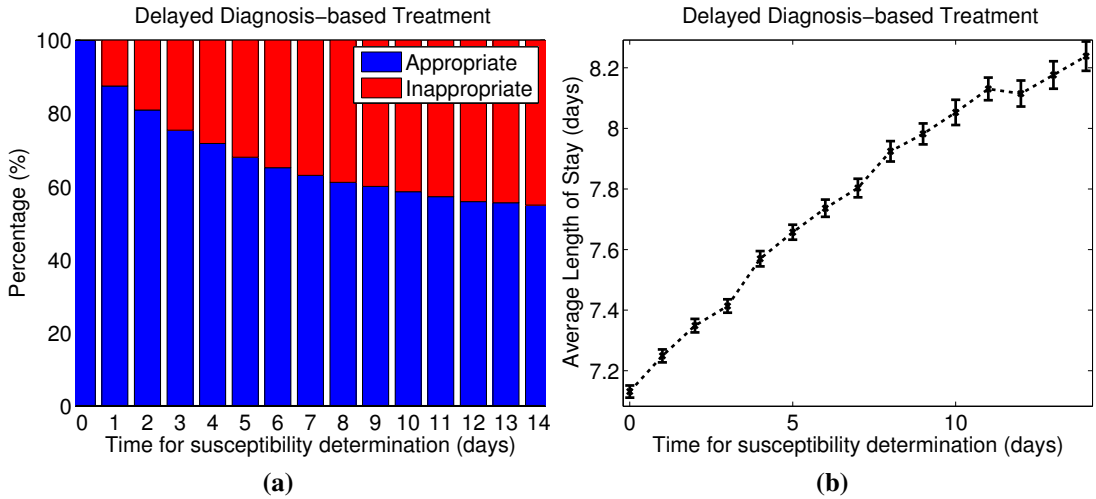

Figure 10.5: a) Delaying the availability of drug susceptibility information decreases drug appropriateness to no better than random b) delays for susceptibility determination increase the expected length of stay.

is a correlation between the time for this susceptibility determination and the percentage of patients receiving appropriate drugs, Figure 10.5b then shows that an increase in this delay is responsible for an almost linear increase in the observed MLoS.

Interestingly, Figure 10.9e shows that multiple diagnoses have little effect on the performance of the diagnosis-based protocol: although drug-resistance pathogens could be transmitted through the hospital and go on to infect and colonise a patient already receiving treatment for a different bacterial strain, in our simulation this was not an important factor. Much more important was that the correct diagnosis of the strain responsible for infection was determined quickly after the time each patient was admitted to hospital. Just one genetic test performed at that moment in time was sufficient to yield a treatment more effective in average than any other mixing or rotational protocol.

Our computational framework allows us to treat every host with a combination of both available antibiotics, so we evaluated the effect that a drug combination has on the disease dynamics of our theoretical hospital ward. As expected, because the drugs considered have a strong synergistic interaction, patients recovered considerably faster than with a 50-50 drug combination than with any monotherapy. It is important to notice, however, that this

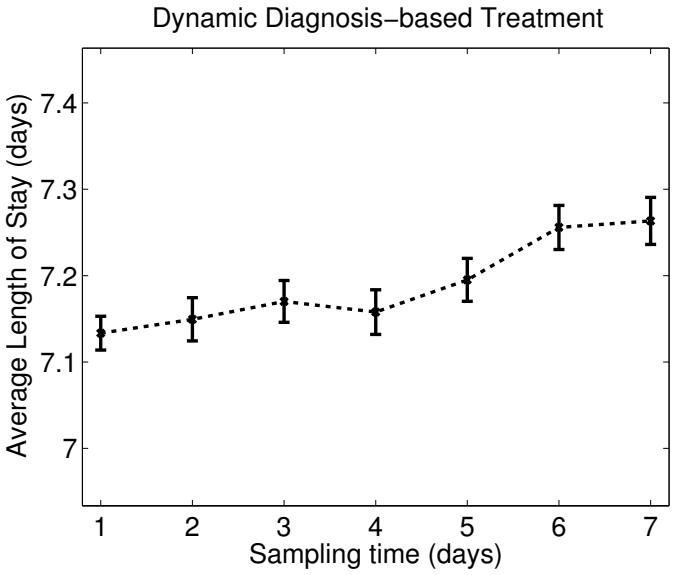

Figure 10.6: Under a dynamic diagnosis-based treatment protocol, if the initial rapid diagnosis is performed immediately, subsequent diagnoses as long as one week later may have little effect on the mean length of stay.

treatment protocol imposes strong selective pressures in favour of multidrug resistance and in consequence decreases the long-term drug efficacy. We also found that the efficacy of combination treatments depends on the levels of drug resistance in the community, a property discussed in the following section.

10.5 Community structure and the efficacy of treatment protocols

The variation between nosocomial and community-acquired resistance is very high and with a significant impact on the epidemiology of drug resistance. For instance, countries that implement infection control programs report that methicillin-resistant *Staphilococus aureus* (MRSA) represent only 1% of the over-all *Staphilococus aureus* bacteramia, while the incidence of MRSA is above 50% in countries where there are no infection control policies [112]. Consequently, the efficacy of different strategies would depend not only in the hospital dynamics but also in the community structure. In our computational framework we can study the effect that the levels of drug resistance in the community have on the effi-

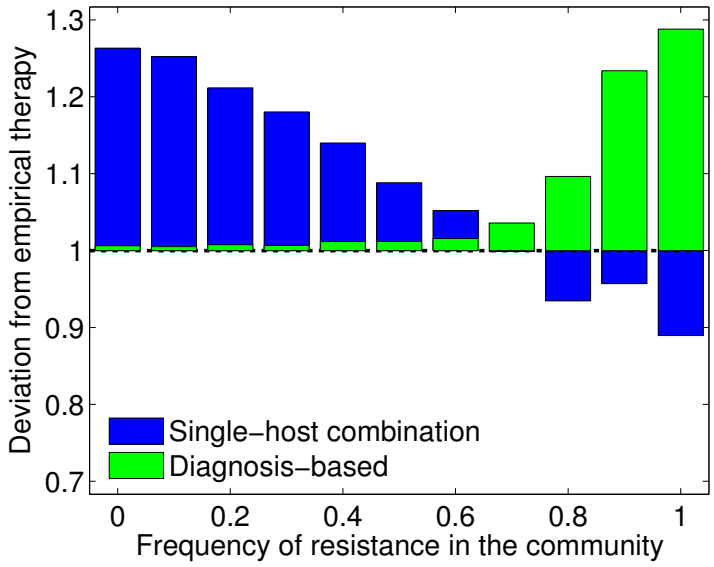

Figure 10.7: As the frequency of resistance in the community increases, so does the relative merits of diagnosis-based treatment. Single-host multidrug combination is very effective when resistance levels are low, but performs even worse than empirical treatments when community-acquired resistance is high.

cacy of different hospital interventions by varying a parameter, $\rho \in [0, 1]$, that denotes the probability that a patient admitted into the ward is infected by a given resistant pathogen. If $\rho = 0$ then every drug resistant pathogen in the system is nosocomial-acquired, while $\rho = 1$ denotes the case when every patient in the queue is colonised with drug-resistant pathogens when admitted into the hospital ward.

When levels of drug resistance in the community are low and thus patients are admitted into the hospital ward infected mainly with susceptible pathogens, then A and B are similarly effective in clearing the infection. In this case, the empirical treatment performs as well as the diagnosis-based protocol. Conversely, when levels of resistance in the community are high there is a considerable benefit obtained from the correct diagnosis and the consequent prescription of the appropriate drug. In this case a diagnosis-based treatment performs significantly better than the empirical therapy, as shown in Figure 10.7. Moreover, as discussed in the previous section, when antibiotics have a synergistic interaction multidrug combination treatments are very effective, in particular when patients are colonised mainly with susceptible pathogens. Interestingly when patients are admitted into the ward with

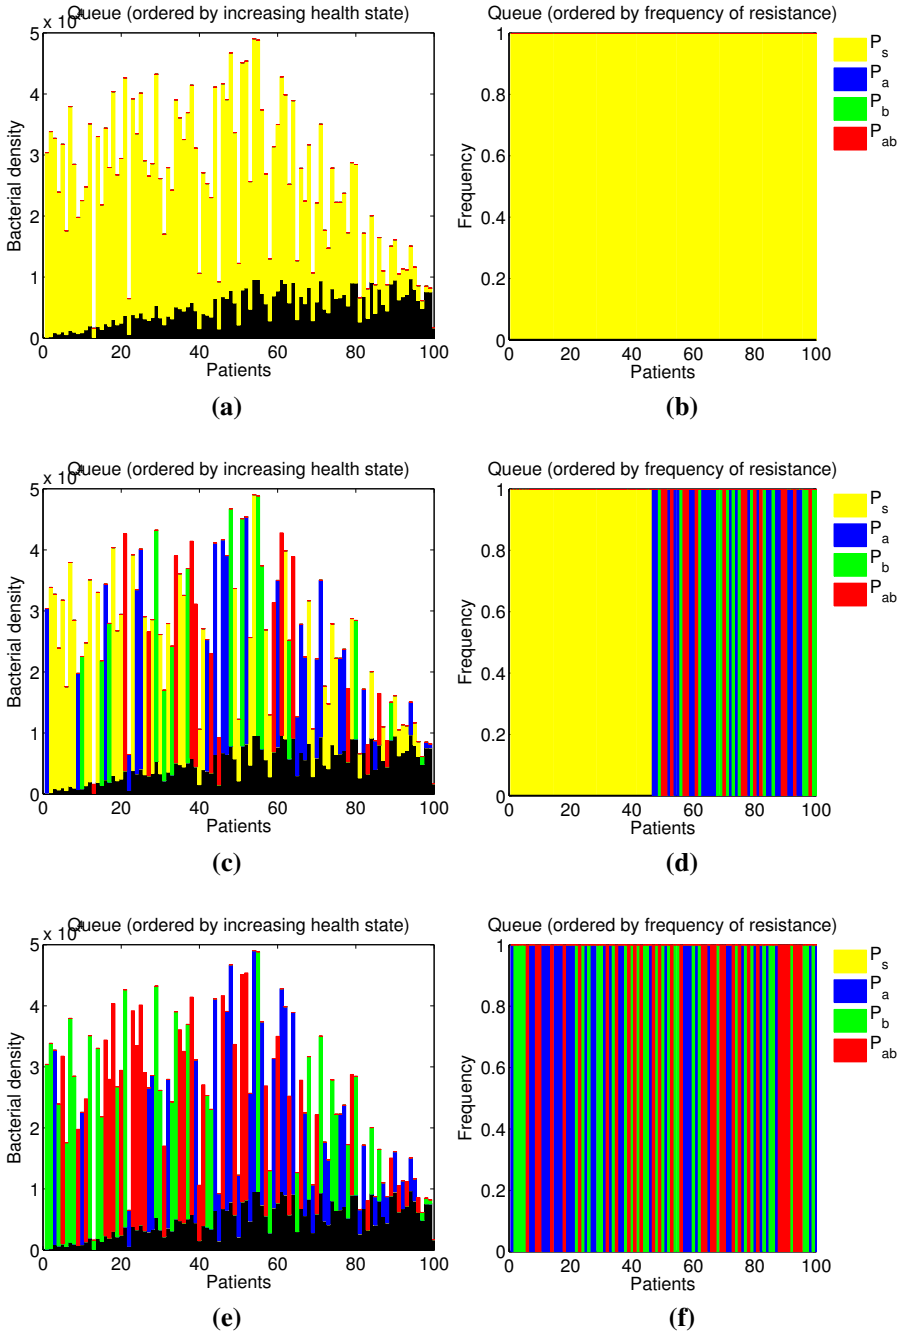

Figure 10.8: Replacement scenario. Patients are colonised with random densities of commensal and pathogenic bacteria. The probability of pathogens being single-drug resistant is given by parameter $0 < \rho_r < 1$.
(left) Initial bacterial densities of patients in the queue. (right) Frequencies of single-drug and multidrug resistant pathogens. (a-b) No drug resistance in the community, $\rho_r = 0$. (c-d) Intermediate levels of resistance in the community, $\rho_r = 0.5$. (e-f) High levels of resistance in the community, $\rho_r = 1$.

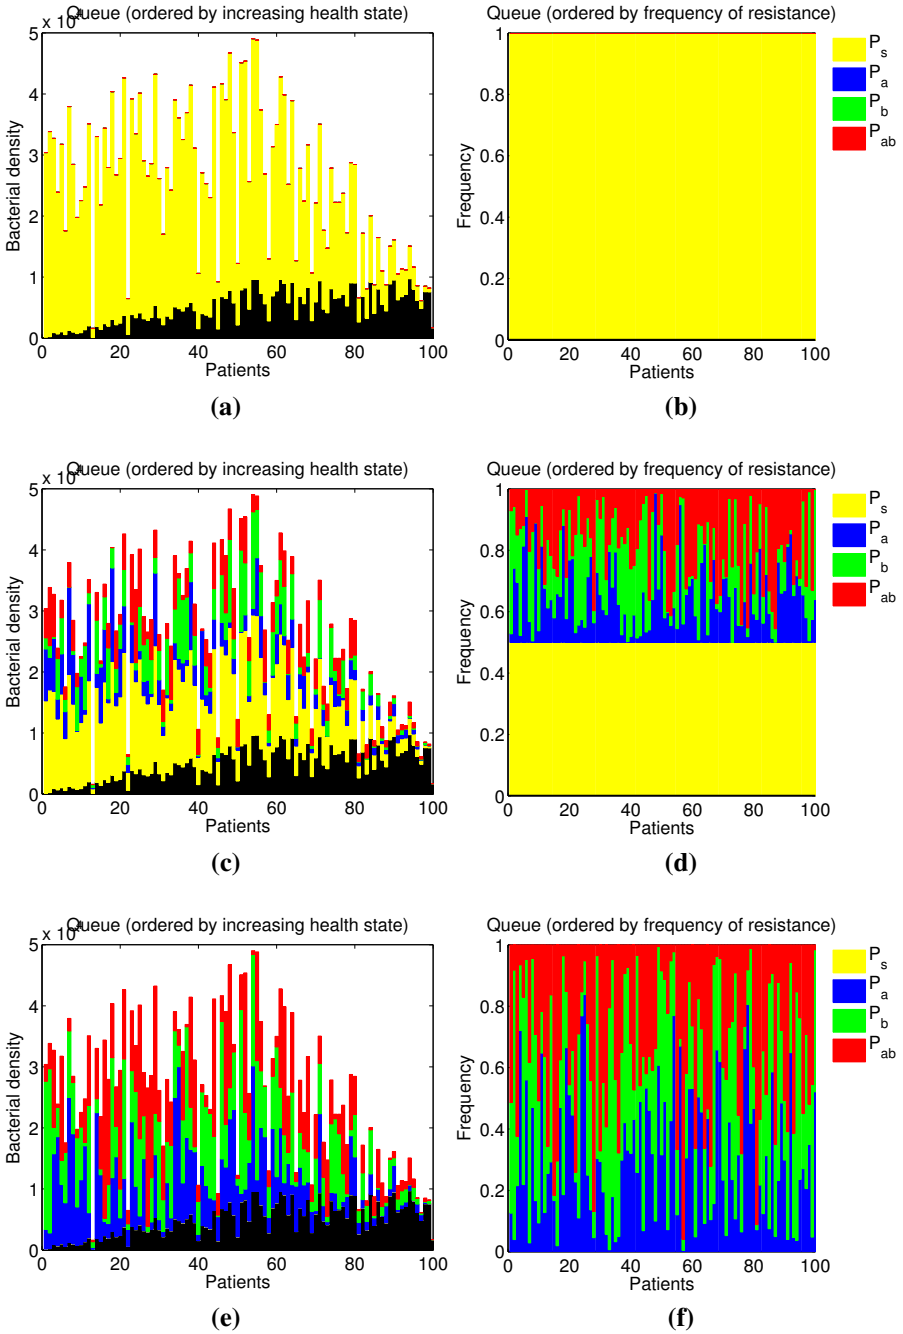

Figure 10.9: Addition scenario. Patients are colonised with random densities of commensal bacteria, as well as infected by susceptible, single-drug and multi-drug resistant pathogens. Frequency of resistance is given by parameter $0 < \rho_a < 1$.
(left) Initial bacterial densities of patients in the queue. (right) Frequencies of single-drug and multidrug resistant pathogens. (a-b) No drug resistance in the community, $\rho_a = 0$. (c-d) Intermediate levels of resistance in the community, $\rho_a = 0.5$. (e-f) High levels of resistance in the community, $\rho_a = 1$.

high levels of drug resistance, then multidrug resistance spreads very quickly throughout the hospital ward and this treatment protocol turns out to be very ineffective, performing even worse than the empirical treatment.

In addition to the community structure, the distribution of susceptible and resistant pathogens within hosts is also relevant [112]. The *addition scenario* illustrated in Figure 10.9 refers to the case where resistant and susceptible pathogens coexist within a single-host. In contrast, Figure 10.8 represents a queue of patients generated under a *replacement scenario*, so-called because patients are infected by a single pathogenic phenotype, either resistant or susceptible. For brevity we have only included in this chapter simulations that consider the addition scenario, and although the results presented would differ quantitatively, the main conclusions presented here seem to be independent of the single-host patterns of resistance.

Conclusion

The many different physical scales relevant to modelling antibiotic resistance evolution, from the physical chemistry of drug molecules and their targets to the clinical and epidemiological effects of drug usage in patients and communities, may mean that a strongly data-driven predictive model across those scales may simply be too much to ask for. While we would like to re-iterate the limitations of our study and suggest the necessity of developing calibrated models designed for specific pathogens and specific antimicrobials, we argue that the models presented in this thesis may still be able to explain some features observed in practise.

In particular, the numerical simulations of the last chapter present one feature in common with clinical trials: over-lapping distributions. We believe this is an important issue when comparing the performance of different interventions, because if we were to test just two simulated trials and compare them, it is possible that our single samples would not reflect the true nature of the distributions from which they were drawn. In our computational study, we can overcome this problem by forming an empirical distribution from repeated simulations, but this exercise cannot be carried out in clinical trials. The same conclusion can be drawn from standard SI models of drug deployment, where it can be shown that the statistical distribution of the performances of different sets of protocols, for example mixing and periodic cycling, must have overlapping supports. In other words, it is a general, mathematical result that there must be mixing policies that outperform cycling policies and vice versa.

Another conclusion that seems to hold across the different modelling frameworks discussed in this thesis is that strategies that increase drug appropriateness are the most effective in the long-term. This leads us to postulate that the one feature common to all successful treatments in our theoretical models is information: the more information we have, the better the performance of the best treatments, although a caveat commensurate with the law of diminishing returns is appropriate.

In summary, it is our contention that the search for effective antibiotic usage strategies should not be reduced to a straightforward comparison between mixing and rotation protocols. More recent treatment paradigms place the individual patient at the core of the treatment and, for example, seek to maximise the appropriateness of the drugs prescribed while minimising the duration of empirical therapy [72]. With these goals in mind, the development of DNA-based rapid sequencing tools that aim to reduce the delay in obtaining microbiology results, as discussed in [12], may prove to be a successful strategy.

We conclude by agreeing with [16] that the last word on the nature of the optimal antibiotic deployment protocol has not been spoken. This is a difficult theoretical and practical problem and we hope that no avenue of investigation will be set aside when searching for measures to combat the scourge of antibiotic resistance.

Appendix

Appendix A

Model parameters

Table A.1: Model parameters for the scenario in which a single-drug is deployed into a population of pathogens with complete competitive advantage over commensals.

Parameter	Description	Value
d	Chemostat dilution rate	$0.6h^{-1}$
ϵ	Rate of phenotypic mutations	0.1
a_p & a_c	antibiotic-cell binding rates	$4 \cdot 10^{-6} \mu g/cell/hr$
$\boldsymbol{c}$	Resource-cell conversion rate in antibiotic-free environment	$1 \times 10^{6} cell/\mu g$
K	Half-saturation constant for all bacterial types (com, wt and mut)	$0.06 \mu g/ml$
$V^*_{\max}$	Maximal resource uptake rate of bacterial type $*$, $V^*_{\max} = \mu^*_{\max}/\boldsymbol{c}$	$\mu^{com}_{\max} = 1.2/h$, $\mu^{wt}_{\max} = 1.6/h$ $\mu^{mut}_{\max} = 1.0/h$
S_0	Resource supply concentration	$1 \mu g/ml$
A_0	Antibiotic supply concentration	$2 \mu g/ml$
κ^*_1	Affinity for antibiotic of bacterial type $*$	$\kappa^{com}_1 = \kappa^{wt}_1 = \kappa^{mut}_1 = 0.05 ml/\mu g$
κ^*_2	Maximal growth inhibition of bacterial type i	$\kappa^{com}_2 = 0.71 ml/\mu g, \kappa^{wt}_2 = 0.5 ml/\mu g,$ $\kappa^{mut}_2 = 0.9 ml/\mu g$

Table A.2: Model parameters for the example where two synergistic antibiotics are deployed to preserve the commensal population while driving a fitter pathogen to extinction.

Parameter	Description	Value
d	Chemostat dilution rate	$0.6h^{-1}$
ϵ	Rate of point mutations	0.01
c	Resource conversion rate	$10^8 cell/\mu g$
K	Bacterial half-saturation constant	$0.06\mu g/ml$
μ^i	Maximal growth rate of bacterial type i per hour (note $c \cdot V^i_{max} = \mu^i$)	$\mu^c = 1.16$, $\mu^s = 1.3$, $\mu^a = 1$, $\mu^b = 1$, $\mu^{ab} = 0.87$
S_0	Resource supply concentration	$6\mu g/ml$
A_0	Antibiotic A supply concentration	$1\mu g/ml$
B_0	Antibiotic B supply concentration	$1\mu g/ml$
a	Antibiotic A binding rate	$1 \times 10^{-8}\mu g/cell$
b	Antibiotic B binding rate	$1 \times 10^{-8}\mu g/cell$
κ_1^i	Maximal growth inhibition of bacterial type i by antibiotic A	$\kappa_1^c = 0.59$, $\kappa_1^s = 0.45$ $\kappa_1^a = 0.81$, $\kappa_1^b = 0.41$ $\kappa_1^{ab} = 0.69$
$\tilde{\kappa}_1^i$	Maximal growth inhibition of bacterial type i by antibiotic B	$\tilde{\kappa}_1^c = 0.56$, $\tilde{\kappa}_1^s = 0.47$ $\tilde{\kappa}_1^a = 0.38$, $\tilde{\kappa}_1^b = 0.79$, $\tilde{\kappa}_1^{ab} = 0.71$
κ_2^i	Affinity for antibiotic A of bacterial type i	$\kappa_2^c = 0.15$, $\kappa_2^s = 0.2$, $\kappa_2^a = 0.3$, $\kappa_2^b = 0.2$, $\kappa_2^{ab} = 0.4$
$\tilde{\kappa}_2^i$	Affinity for antibiotic B of bacterial type i	$\tilde{\kappa}_2^c = 0.18$, $\tilde{\kappa}_2^s = 0.25$, $\tilde{\kappa}_2^a = 0.2$, $\tilde{\kappa}_2^b = 0.4$, $\tilde{\kappa}_2^{ab} = 0.39$

Table A.3: Parameters estimated from the experiments described in Appendix D

Parameter	Description	Value
c	Resource conversion rate	1.8 x $10^4 cells/\mu g$
K	Bacterial half-saturation constant	$0.62 \mu g/ml$
V_{max}	Maximal growth rate	$0.000251 \mu g/cell/h$
κ_d	Maximal growth inhibition by doxycyclin	0.068
κ_e	Maximal growth inhibition by erythromycin	3.46
$\tilde{\kappa}_d$	Affinity for doxycyclin	0.198
$\tilde{\kappa}_e$	Affinity for erythomycin	0.079
S_0	Resource supply concentration	$2000 \mu g/ml$
DOX_0	Maximum doxycyclin concentration	$0.7 \mu g/ml$
ERY_0	Maximum erythomycin concentration	$3.5 \mu g/ml$
η	Fraction of the population transferred	0.1
ϵ	Rate of point mutations	0.001
κ_1^i	Maximal growth inhibition of bacterial type i by doxycyclin	$\kappa_d^d = 0.285$, $\kappa_d^e = 0.0767$, $\kappa_d^{de} = 0.274$
$\tilde{\kappa}_1^i$	Maximal growth inhibition of bacterial type i by erythromycin	$\kappa_e^d = 1.12$, $\kappa_e^e = 0.8$, $\kappa_e^{de} = 2.156$
κ_2^i	Affinity for doxycyclin of bacterial type i	$\tilde{\kappa}_d^d = 0.000000214$, $\tilde{\kappa}_d^e = 0.000000219$, $\tilde{\kappa}_d^{de} = 0.0000016$
$\tilde{\kappa}_2^i$	Affinity for erythomycin of bacterial type i	$\tilde{\kappa}_e^d = 0.4147$, $\tilde{\kappa}_e^e = 0.753$, $\tilde{\kappa}_e^{de} = 0.685$

Table A.4: Model parameters considered for the stochastic and spatially explicit model of a hospital ward.

Parameter	Description	Value
N	Number of hospitals simulated (unless specified otherwise)	1000
n	Number of patients in the queue	100
B	Number of beds in hospital	8
β	Rate of transmission between neighbouring beds	0.1
δ	Ratio of commensals and pathogens required in order to patients to be discharged	0.9
m	Rate of input and output of resources	$0.4h^{-1}$
d	Rate of clearance of bacteria	$0.4h^{-1}$
ϵ	Rate of point mutations	0.01
c	Resource conversion rate	$10^8 cell/\mu g$
K	Bacterial half-saturation constant	$0.06\mu g/ml$
μ^i	Maximal growth rate of bacterial type i per hour (note $c \cdot V^i_{max} = \mu^i$)	$\mu^c = 1.5$, $\mu^s = 1.3$, $\mu^a = 1$, $\mu^b = 1$, $\mu^{ab} = 0.87$
S_0	Resource supply concentration	$6\mu g/ml$
A_0	Antibiotic A supply concentration	$1\mu g/ml$
B_0	Antibiotic B supply concentration	$1\mu g/ml$
a	Antibiotic A binding rate	$1 \times 10^{-8}\mu g/cell$
b	Antibiotic B binding rate	$1 \times 10^{-8}\mu g/cell$
κ_1^i	Maximal growth inhibition of bacterial type i by antibiotic A	$\kappa_1^c = 0.59$, $\kappa_1^s = 0.45$ $\kappa_1^a = 0.81$, $\kappa_1^b = 0.41$ $\kappa_1^{ab} = 0.69$
$\tilde{\kappa}_1^i$	Maximal growth inhibition of bacterial type i by antibiotic B	$\tilde{\kappa}_1^c = 0.56$, $\tilde{\kappa}_1^s = 0.47$ $\tilde{\kappa}_1^a = 0.38$, $\tilde{\kappa}_1^b = 0.79$, $\tilde{\kappa}_1^{ab} = 0.71$
κ_2^i	Affinity for antibiotic A of bacterial type i	$\kappa_2^c = 0.15$, $\kappa_2^s = 0.2$, $\kappa_2^a = 0.3$, $\kappa_2^b = 0.2$, $\kappa_2^{ab} = 0.4$
$\tilde{\kappa}_2^i$	Affinity for antibiotic B of bacterial type i	$\tilde{\kappa}_2^c = 0.18$, $\tilde{\kappa}_2^s = 0.25$, $\tilde{\kappa}_2^a = 0.2$, $\tilde{\kappa}_2^b = 0.4$, $\tilde{\kappa}_2^{ab} = 0.39$

Appendix B

Entropic Regularisation

It will be convenient for the purpose of numerically calculating near-optimal controls with respect to the functional $\mathcal{J}(\delta)$ defined in (6.1) to regularise it by introducing an *entropic* term of the form

$$e(\delta) := \delta \ln \delta + (d - \delta) \ln(d - \delta)$$

and then defining

$$\mathcal{J}_\eta(\delta) := -\eta \int_0^T e(\delta) dt + \mathcal{J}(\delta),$$

where η is a temperature-like parameter. We then seek a solution of the unconstrained optimal problem

$$\text{find } \delta^*(\eta) \in L^\infty(0,T) \text{ such that } \mathcal{J}_\eta(\delta^*(\eta)) = \sup\{\mathcal{J}_\eta(\delta) : \delta \in L^\infty(0,T)\}. \tag{B.1}$$

It is easy to prove from the convexity properties of the function $e(\cdot)$, on defining an optimal control δ^* by

$$\delta^* \in L^\infty(0,T) \text{ satisfies } \mathcal{J}(\delta^*) = \sup\{\mathcal{J}(\delta) : \delta \in L^\infty(0,T), 0 \le \delta \le d\} \tag{B.2}$$

that $\delta^*(\eta) \overset{*}{\rightharpoonup} \delta^*$ in L^∞ as $\eta \to 0$ where $\eta > 0$.

The effect of this regularisation is to ensure that minimisers of $\mathcal{J}_\eta$, members of $L^\infty(0,T)$

denoted $\delta^*(\eta)(t)$, are in fact analytic functions of t and it also ensures that convergence to the maximal entropy state $\delta^*(\eta)(t) \to \frac{d}{2}$ as occurs in an appropriate strong topology as $\eta \to \infty$.

The value of the regularisation is that it allows us to readily determine $\delta^*(\eta)$ numerically, largely by smoothing singularities in the optimal solution, but it also allows us to use a numerical continuation procedure akin to annealing to compute small-η optimal controls. This continuation procedure starts with η at a high value and with an initial guess of $\delta^*_{\mathsf{guess}} = d/2$ and then successively reduces η. By coupling this continuation algorithm to an adaptive meshing strategy ensures that we are able to resolve the switches that we anticipate to be present in the optimal control δ^* between the deployment and non-deployment of antibiotic. We used the Matlab code `bvp4c` [52] to implement this numerical continuation strategy.

To understand the effect of the regularisation in more detail, a necessary condition for the optimality of $\delta^*(\eta)$ with respect to $\mathcal{J}$ states, using the Pontryagin Maximum Principle, that the associated Hamiltonian

$$\mathcal{H}(\boldsymbol{x}, \lambda, \delta) := \langle w, \boldsymbol{x} \rangle + \langle \lambda, \mathcal{G}(\boldsymbol{x}; d, S_0) + \delta A_0 \mathbf{e} \rangle$$

is maximised along the optimal trajectory. Here $\lambda = (\lambda_S, \lambda_C, \lambda_P, \lambda_A) \in \mathbb{R}^{1+n+m+1}$ is an adjoint variable that satisfies the adjoint equation (B.3a) below

$$-\dot{\lambda} = w + (\nabla_{\boldsymbol{x}} \mathcal{G}(\boldsymbol{x}; d, S_0))^T \lambda, \qquad \lambda(T) = 0, \tag{B.3a}$$

$$\dot{\boldsymbol{x}} = \mathcal{G}(\boldsymbol{x}; d, S_0) + \delta A_0 \mathbf{e}, \qquad \boldsymbol{x}(0) \text{ is given.} \tag{B.3b}$$

The regularised payoff functional $\mathcal{J}_\eta$ requires a suitably modified Hamiltonian, namely

$$\tilde{\mathcal{H}}(\boldsymbol{x}, \lambda, \delta; \eta) = \mathcal{H}(\boldsymbol{x}, \lambda, \delta) - \eta e(\delta). \tag{B.4}$$

Observe that the only terms in $(B.4)$ that depend on the control variable δ are

$$\langle \lambda, A_0 \delta \boldsymbol{e} \rangle - \eta e(\delta) = A_0 \delta(\lambda, \boldsymbol{e}) - \eta e(\delta) = A_0 \delta \lambda_A - \eta e(\delta) \tag{B.5}$$

and so we define

$$h(\delta) := \lambda_A A_0 \delta - \eta e(\delta).$$

As $\eta > 0$, the modified Hamiltonian $\tilde{\mathcal{H}}$ is maximised with respect to δ when $0 = h'(\delta) = \lambda_A A_0 - \eta e'(\delta)$, a condition that holds if and only if $\log\left(\frac{\delta}{d-\delta}\right) - \frac{A_0 \lambda_A}{\eta} = 0$ from where we find an expression for δ in terms of the regularisation parameter η and the adjoint variable λ_A:

$$\delta = \delta\left\langle \eta, \lambda_A \right\rangle := d(1 + e^{-A_0 \lambda_A \eta^{-1}})^{-1}.$$

If we now define $\Lambda := (\Lambda_S, \Lambda_C, \Lambda_P, \Lambda_A) = \lambda/\eta$ then Λ satisfies the final-value equation

$$-\dot{\Lambda} = w\eta^{-1} + (\nabla_{\boldsymbol{x}} \mathcal{G}(\boldsymbol{x}; d, S_0))^T \Lambda, \qquad \Lambda(T) = 0, \tag{B.6}$$

and the control associated with it can be written $\delta(t) = d(1 + e^{-A_0 \Lambda_A(t)})^{-1}$, illustrating that (B.3) is singularly perturbed with respect to η near $\eta = 0$. The nature of this singular perturbation will result in the appearance of internal layers where the optimal control $\delta^*(\eta)$ switches between the two extreme states, $\delta = 0$ and $\delta = d$. As the behaviour of $\lambda_A(t)$ decides when the control switches from one regime to the other, this function is called the *switching function*. However, (B.3) is regularly perturbed about $\eta = \infty$ in the sense that we can formally substitute $\eta = \infty$ in (B.6) and find a solution of the resulting differential equation. The resulting solution is analytic, unique and satisfies $\Lambda(t) \equiv (0, 0, 0, 0)$ from where the corresponding optimal control is given by the maximum entropy state. This occurs at $\max_{0 \le \delta \le d}(-e(\delta))$ namely when $\delta = d/2$ and we therefore conclude that

$$\lim_{\eta \to \infty} \delta^*(\eta) = \frac{d}{2}$$

in a C^1 sense. This argument can be made rigorous using the implicit function theorem and it provides an initial guess (the function $\delta^*_{\mathsf{guess}} = d/2$) for the numerical annealing procedure used to numerically determine optimal controls.

As $\Lambda_A(T) = 0$ it follows that $\delta^*(\eta)(T) = d(1 + e^0)^{-1} = d/2$ and so the regularisation introduces an artefact into the $\mathcal{J}_\eta$-optimal control which forces this control to equal $d/2$ at the end-point of the time domain.

Appendix C

Genetic algorithm

In order to compute near-optimal antibiotic deployment protocols we used a genetic algorithm to exploit the fact that bang-bang controls are optimal in a certain sense. The algorithm searches for near-optimal rotational protocols through the space of bang-bang controls using the following simple-minded strategy. Consider a partition of the time interval $[0, T]$ into n regular sub-intervals, a single drug must flow into the chemostat at maximum dose on each sub-interval but we allow the possibility of switching antibiotics at the end of each sub-interval.

The set of all potential switches $\sigma := \{\sigma_1, \sigma_2, \ldots, \sigma_n\}$ can be expressed as a binary string of length n and this represents the *chromosome*, $\mathcal{C}$, of a given treatment. Each *gene*, σ_i, can be expressed by *allele* 0 meaning that the same drug is going to be deployed at the next sub-interval, or by allele 1 meaning that a switch to the other drug is necessary. We can therefore associate each chromosome $\mathcal{C} \in \{0, 1\}^n$ with a bang-bang control strategy $\alpha(t; \mathcal{C})$ and from this we can compute its fitness that we define to be $J(\alpha(t; \mathcal{C}), \beta(t; \mathcal{C}))$ where $\beta(t; \mathcal{C}) = d - \alpha(t; \mathcal{C})$, J is computed using an accurate time-stepping algorithm to solve the differential equation (9.1).

The number of different rotation protocols to search through is 2^n and the following algorithm located a near-optimal control with a reasonable amount of computer time:

1. Fix a population size N and choose a random initial population of N chromosomes,

$\mathcal{PC}^0 = \{\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_N\}$
Given $\mathcal{PC}^k$ define $\mathcal{PC}^{k+1}$ as follows.

2. Compute the numerical fitness $J(\alpha(t;\mathcal{C}_i), \beta(t;\mathcal{C}_i))$ where $\beta(t;\mathcal{C}_i) = d - \alpha(t;\mathcal{C}_i)$ for every chromosome $\mathcal{C}_i$ of the current population.

3. Select the fittest chromosomes and reproduce them pair-wise. The chromosome of the offspring is constructed by recombination, inheriting each allele from only one of the parents, where these are chosen randomly for each gene.

4. Mutate a fixed proportion of genes of the population by randomly bit-flipping the alleles of some proportion of chromosome, also randomly chosen.
The population $\mathcal{PC}^{k+1}$ then contains only the fittest N members of the previous generation, their offspring and the mutated chromosomes.

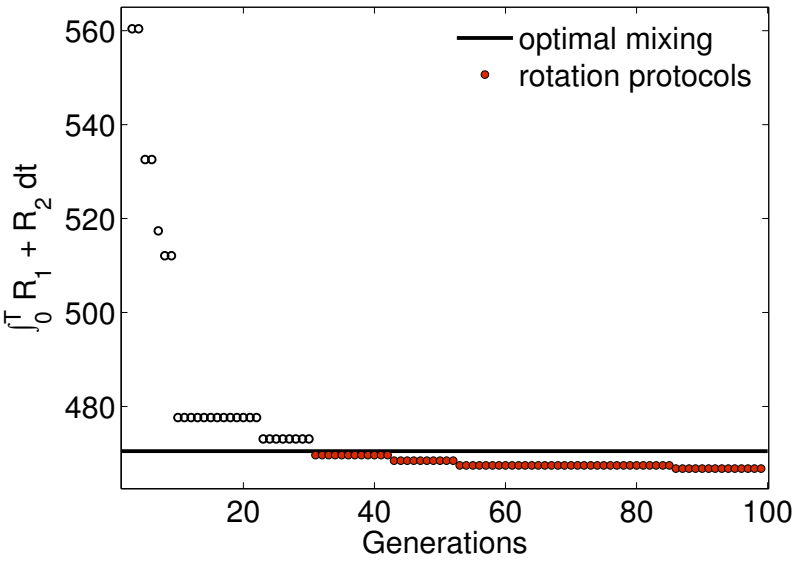

Figure C.1: Illustration of the genetic algorithm applied to Problem 2. Black circles represent the rotation protocols that have a lower fitness than the optimal mixing protocol (yellow line), while red circles represent protocols that outperform the best mixing strategy.

Appendix D

Experimental methods

Bacterial Strain. *E. coli* Strain MC4100 (ordered from the Coli Genetic Stock Center: http://cgsc2.biology.yale.edu/Strain.php?ID=9973) was used. Stocks were made from a single colony picked for an overnight LB-culture and frozen for further use.

M9 Growth Medium. The experiment was then performed in M9 Growth Medium prepared as follows: Part A: 350 $\frac{g}{L}$ K_2HPO_4, 100 $\frac{g}{L}$ KH_2HPO_4; Part B: 29.4 $\frac{g}{L}$ Trisodium citrate, 50 $\frac{g}{L}$ $(NH_4)_2SO_4$, 5 $\frac{g}{L}$ $MgSO_4$. Parts A and B were 50x stock solutions in water, that were sterilised by autoclaving. For M9 Minimal Medium they were diluted in water accordingly. 0.2 % Glucose and 0.1 % Casamino Acid were added as nutrients for the M9 Growth Medium.

Antibiotics used. Antibiotics used were Erythromycin (from Aldrich, Product # 856193) and Doxycycline (Sigma, Product # D9891), two antibiotics known to synergise [48]. Liquid stocks were prepared from powder stocks at 50 $\frac{mg}{ml}$. For this purpose, Erythromycin hydrate was dissolved in Ethanol and Doxycycline hyclate in sterile water (afterwards filter sterilised).

Maximal antibiotic concentrations were determined for each antibiotic in preliminary trials as concentrations leading to full inhibition of bacterial growth for 24 hours using just one of the antibiotics at a time. For Erythromycin 10 $\frac{\mu g}{ml}$ was the maximal concentration,

for Doxycycline 0.3 $\frac{\mu g}{ml}$. For each, 11 dilutions in M9 Growth Medium ranging from no antibiotic to maximal concentration.

Setup. A grid of all combinations of these concentrations was replicated 6 times and spread in a systematically randomised way over three 384-well plates at 50 μl per well. Extra wells were filled with growth medium but not inoculated with bacteria for background OD measurements. For inoculation, the wells of a 96-well plate were filled with 100 μl of MC4100 overnight culture and then a very small subsample was taken by dipping a full 96-rack of 10 μl tips first into the 96-well plate and then into one of the four 96-well subsets of the 384-well plate.

The three plates were sealed with a transparent plastic foil, the liquid was centrifuged down with a short spin and the OD was measured in a Tecan genios plate reader. They were then grown in a 37°C incubator and shaken at 130 rpm for 24 hours. For OD measurements, they were taken out of the incubator every 30 minutes.

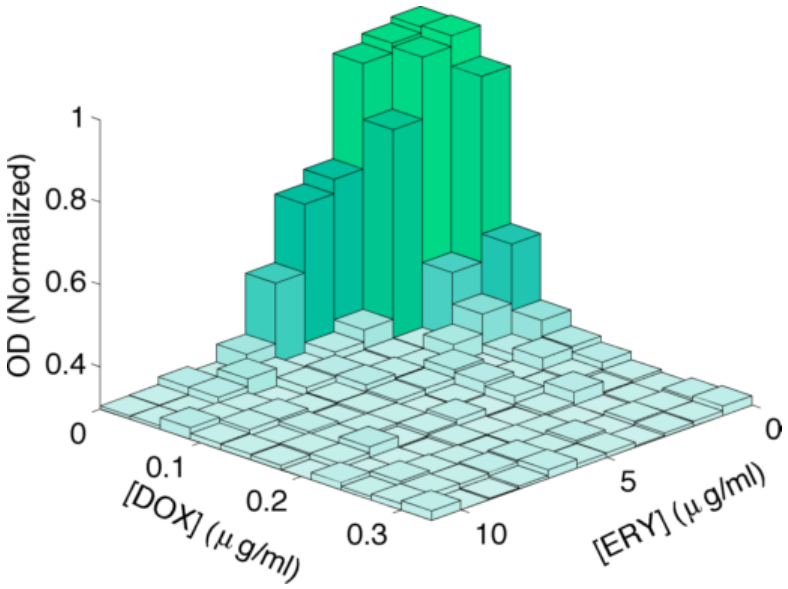

Figure D.1: Dose-response grid of E. Coli growing under glucose limitation and under the effect of different concentrations of docycyclin and erythromycin.

References

[1] American thoracic society. medical section of the american lung association: Treatment of tuberculosis and tuberculosis infection in adults and children. *Am Rev Respir Dis*, 134(2):355–63, Aug 1986.

[2] Initial therapy for tuberculosis in the era of multidrug resistance. recommendations of the advisory council for the elimination of tuberculosis. *MMWR Recomm Rep*, 42(RR-7):1–8, May 1993.

[3] B. Allegranzi, R. Luzzati, A. Luzzani, F. Girardini, L. Antozzi, R. Raiteri, G. Di Perri, and E. Concia. Impact of antibiotic changes in empirical therapy on antimicrobial resistance in intensive care unit-acquired infections. *J Hosp Infect*, 52(2):136–40, Oct 2002.

[4] U. Alon. *An Introduction to Systems Biology*. Chapman and Hall, 2006.

[5] D. I. Andersson. The ways in which bacteria resist antibiotics. *The International Journal of Risk and Safety in Medicine*, 17:111–116, 2005.

[6] D. I. Andersson and D. Hughes. Antibiotic resistance and its cost: is it possible to reverse resistance? *Nat Rev Microbiol*, 8(4):260–71, Apr 2010.

[7] D. I. Andersson and B. R. Levin. The biological cost of antibiotic resistance. *Curr Opin Microbiol*, 2(5):489–93, Oct 1999.

[8] R. E. Beardmore and R. Peña-Miller. Antibiotic cycling versus mixing: The difficulty of using mathematical models to definitively quantify their relative merits. *Math Biosci Eng*, 7(4):923–33, Oct 2010.

[9] R. E. Beardmore and R. Peña-Miller. Rotating antibiotics selects optimally against antibiotic resistance, in theory. *Math Biosci Eng*, 7(3):527–52, Jul 2010.

[10] M. Berenbaum. What is synergy? *Pharmacol. Rev.*, (41):93–141, 1989.

[11] M. Bergeron and M. Ouellette. Preventing antibiotic resistance through rapid genotypic identification of bacteria and of their antibiotic resistance genes in the clinical microbiology laboratory. *Clinical Microbiology*, 36(8):2169–2172, 1998.

[12] M. G. Bergeron. Revolutionizing the practice of medicine through rapid (< 1h) dna-based diagnostics. *Clin Invest Med*, 31(5):E265–71, 2008.

[13] C. T. Bergstrom, M. Lo, and M. Lipsitch. Ecological theory suggests that antimicrobial cycling will not reduce antimicrobial resistance in hospitals. *Proc Natl Acad Sci U S A*, 101(36):13285–90, Sep 2004.

[14] B. Bloom and C. Murray. Tuberculosis: Commentary on a reemergent killer. *Science*, 257(5073):1055–64, 1992.

[15] S. Bonhoeffer, M. Lipsitch, and B. R. Levin. Evaluating treatment protocols to prevent antibiotic resistance. *Proc Natl Acad Sci U S A*, 94(22):12106–11, Oct 1997.

[16] S. Bonhoeffer, P. zur Wiesch, and R. D. Kouyos. Rotating antibiotics does not minimize selection for resistance. *Math Biosci Eng*, 7(4):919–922, Oct 2010.

[17] E. M. Brown and D. Nathwani. Antibiotic cycling or rotation: a systematic review of the evidence of efficacy. *J Antimicrob Chemother*, 55(1):6–9, Jan 2005.

[18] A. Buckling, R. Craig Maclean, M. A. Brockhurst, and N. Colegrave. The beagle in a bottle. *Nature*, 457(7231):824–9, Feb 2009.

[19] T. Burden, J. Ernstberger, and K. Fister. Optimal control applied to immunotherapy. *Discrete and Continuous Dynamical Series B*, 4:134–146, 2004.

[20] E. A. Campbell, O. Pavlova, N. Zenkin, F. Leon, H. Irschik, R. Jansen, K. Severinov, and S. A. Darst. Structural, functional, and genetic analysis of sorangicin inhibition of bacterial rna polymerase. *EMBO J*, 24(4):674–82, Feb 2005.

[21] R. Chait, A. Craney, and R. Kishony. Antibiotic interactions that select against resistance. *Nature*, 446(7136):668–671, 04 2007.

[22] H. Chambers. The changing epidemiology of staphylococcus aureus? *Emerg Infect Dis*, 7(2):178–82, 2001.

[23] J. Chastre, M. Wolff, J.-Y. Fagon, S. Chevret, F. Thomas, D. Wermert, E. Clementi, J. Gonzalez, D. Jusserand, P. Asfar, D. Perrin, F. Fieux, S. Aubas, and PneumA Trial Group. Comparison of 8 vs 15 days of antibiotic therapy for ventilator-associated pneumonia in adults: a randomized trial. *JAMA*, 290(19):2588–98, Nov 2003.

[24] T. Chou and P. Talalay. Quantitative analysis of dose-effect relationships: the combined effect of multiple drugs or enzyme inhibitors. *Advances in Enzyme Regulation 1984; 22: 27-55*, 1984.

[25] T. C. Chou. Drug combinations: From laboratory to practice. *J Lab Clin Med*, 132(1), 1998.

[26] T. C. Chou and P. Talalay. Generalized equations for the analysis of inhibitions of michaelis-menten and higher-order kinetic systems with two or more mutually exclusive and nonexclusive inhibitors. *Eur J Biochem*, 115(1):207–16, Mar 1981.

[27] K. C. Chow, X. Wang, and C. Castillo-Chavez. A mathematical model of nosocomial infection and antibiotic resistance: evaluating the efficacy of antimicrobial cycling programs and patient isolation on dual resistance. Technical Report MTBI-04-05M, Mathematical and Theoretical Biology Institute, Arizona State University, 2007.

[28] M. L. Cohen. Epidemiology of drug resistance: implications for a post-antimicrobial era. *Science*, 257(5073):1050–5, Aug 1992.

[29] T. Cohen and M. Murray. Modeling epidemics of multidrug-resistant m. tuberculosis of heterogeneous fitness. *Nat Med*, 10(10):1117–21, Oct 2004.

[30] E. M. C. D'Agata, M. Dupont-Rouzeyrol, P. Magal, D. Olivier, and S. Ruan. The impact of different antibiotic regimens on the emergence of antimicrobial-resistant bacteria. *PLoS One*, 3(12):e4036, 2008.

[31] S. J. Dancer. How antibiotics can make us sick: the less obvious adverse effects of antimicrobial chemotherapy. *Lancet Infect Dis*, 4(10):611–9, Oct 2004.

[32] F. Débarre, T. Lenormand, and S. Gandon. Evolutionary epidemiology of drug-resistance in space. *PLoS Comput Biol*, 5(4):e1000337, Apr 2009.

[33] L. Dethlefsen, M. McFall-Ngai, and D. A. Relman. An ecological and evolutionary perspective on human-microbe mutualism and disease. *Nature*, 449(7164):811–8, Oct 2007.

[34] L. Dethlefsen and D. A. Relman. Microbes and health sackler colloquium: Incomplete recovery and individualized responses of the human distal gut microbiota to repeated antibiotic perturbation. *Proc Natl Acad Sci U S A*, Sep 2010.

[35] H. Dupont, H. Mentec, J. P. Sollet, and G. Bleichner. Impact of appropriateness of initial antibiotic therapy on the outcome of ventilator-associated pneumonia. *Intensive Care Med*, 27(2):355–62, Feb 2001.

[36] H. Eagle, R. Fleischman, and M. Levy. "continuous" vs. "discontinuous" therapy with penicillin; the effect of the interval between injections on therapeutic efficacy. *N Engl J Med*, 248(12):481–8, Mar 1953.

[37] P. Ehrlich. Chemotherapeutics: scientific principles, methods and results. *Lancet*, pages 445–451, 1913.

[38] R. el Moussaoui, C. A. J. M. de Borgie, P. van den Broek, W. N. Hustinx, P. Bresser, G. E. L. van den Berk, J.-W. Poley, B. van den Berg, F. H. Krouwels, M. J. M. Bonten, C. Weenink, P. M. M. Bossuyt, P. Speelman, B. C. Opmeer, and J. M. Prins. Effectiveness of discontinuing antibiotic treatment after three days versus eight days in mild to moderate-severe community acquired pneumonia: randomised, double blind study. *BMJ*, 332(7554):1355, Jun 2006.

[39] R. G. Finch, J. P. Metlay, P. G. Davey, L. J. Baker, and International Forum on Antibiotic Resistance colloquium. Educational interventions to improve antibiotic use in the community: report from the international forum on antibiotic resistance (ifar) colloquium, 2002. *Lancet Infect Dis*, 4(1):44–53, Jan 2004.

[40] A. Fleming. *Penicillin*. Nobel Lecures, Physiology or Medicine 1942-1962. Elsevier Publishing Company, 1964.

[41] W. Fleming and R. Rishel. *Deterministic and Stochastic Optimal Control*. Springer, 1975.

[42] K. R. Foster and H. Grundmann. Do we need to put society first? the potential for tragedy in antimicrobial resistance. *PLoS Med*, 3(2):e29, Feb 2006.

[43] S. Gagneux, C. D. Long, P. M. Small, T. Van, G. K. Schoolnik, and B. J. M. Bohannan. The Competitive Cost of Antibiotic Resistance in Mycobacterium tuberculosis. *Science*, 312(5782):1944–1946, 2006.

[44] W. Greco, G. Bravo, and J. Parsons. The search for synergy: a critical review from a response surface perspective. *Pharmacol. Rev.*, (47):331–85, 1995.

[45] H. Grundmann, L. Aires-de Sousa, J. Boyce, and E. Tiemersma. Emergence and resurgence of meticillin-resistant staphylococcus aureus as a public-health threat. *The Lancet*, 368(9538):874–885, September 2006.

[46] D. R. Guay. Guidelines for the management of adults with health care-associated pneumonia: implications for nursing facility residents. *Consult Pharm*, 21(9):719–25, Sep 2006.

[47] T. L. Hedrick, A. S. Schulman, S. T. McElearney, R. L. Smith, B. R. Swenson, H. L. Evans, J. D. Truwit, W. M. Scheld, and R. G. Sawyer. Outbreak of resistant pseudomonas aeruginosa infections during a quarterly cycling antibiotic regimen. *Surg Infect (Larchmt)*, 9(2):139–52, Apr 2008.

[48] M. Hegreness, N. Shoresh, D. Damian, D. Hartl, and R. Kishony. Accelerated evolution of resistance in multidrug environments. *Proc Natl Acad Sci U S A*, 105(37):13977–81, Sep 2008.

[49] S. Holmberg, S. Solomon, and P. Blake. Health and economic impacts of antimicrobial resistance. *Rev Infect Dis*, 9(6):1065–1078, November 1987.

[50] P. Huovinen. Mathematical model–tell us the future! *J Antimicrob Chemother*, 56(2):257–8; discussion 431, Aug 2005.

[51] M. D. Iseman. Treatment of multidrug-resistant tuberculosis. *N Engl J Med*, 329(11):784–91, Sep 1993.

[52] J. Kierzenka and L. F. Shampine. A BVP solver based on residual control and the Maltab PSE. *ACM Trans. Math. Softw.*, 27(3):299–316, 2001.

[53] W. M. Kirby and W. A. Craig. Theory and applications of pulse dosing: a summary of the symposium. *Rev Infect Dis*, 3(1):1–3, 1981.

[54] C. M. Kunin. Dosage schedules of antimicrobial agents: a historical review. *Rev Infect Dis*, 3(1):4–11, 1981.

[55] R. Laxminarayan. Bacterial resistance and the optimal use of antibiotics. *Resources for the Future*, 2001.

[56] M. Leeb. Antibiotics: a shot in the arm. *Nature*, 431(7011):892–3, 2004.

[57] J. Lehar, A. S. Krueger, W. Avery, A. M. Heilbut, L. M. Johansen, E. R. Price, R. J. Rickles, G. F. Short III, J. E. Staunton, X. Jin, M. S. Lee, G. R. Zimmermann, and A. A. Borisy. Synergistic drug combinations tend to improve therapeutically relevant selectivity. *Nat Biotech*, 27(7):659–666, 07 2009.

[58] J. Lehar, G. Zimmermann, A. S. Krueger, R. A. Molnar, J. T. Ledell, A. M. Heilbut, G. F. S. III, L. C. Giursti, G. P. Nolan, O. A. Magid, M. S. Lee, A. A. Borisy, B. R. Stockwell, and C. T. Keith. Chemical combination effects predict connectivity in biological systems. *Molecular Systems Biology*, 3(80), 2007.

[59] K. F. S. Lenhart and J. S. McNally. Optimizing chemotherapy in an hiv model. *Electronic J Differential Equations*, 1998.

[60] R. E. Lenski. Bacterial evolution and the cost of antibiotic resistance. *Int Microbiol*, 1(4):265–70, Dec 1998.

[61] B. R. Levin and M. J. M. Bonten. Cycling antibiotics may not be good for your health. *Proc Natl Acad Sci U S A*, 101(36):13101–2, Sep 2004.

[62] A. Levy and B. Marshall. Antibacterial resistance worldwide: causes, challenges and responses. *Nature Medicine*, 10(12), 2004.

[63] M. Lipsitch, C. T. Bergstrom, and B. R. Levin. The epidemiology of antibiotic resistance in hospitals: paradoxes and prescriptions. *Proc Natl Acad Sci U S A*, 97(4):1938–43, Feb 2000.

[64] Y. Q. Liu, Y. Z. Zhang, and P. J. Gao. Novel concentration-killing curve method for estimation of bactericidal potency of antibiotics in an in vitro dynamic model. *Antimicrob Agents Chemother*, 48(10):3884–91, Oct 2004.

[65] S. Loewe. The problem of synergism and antagonism of combined drugs. *Arzneimittelforschung*, 3(6):285–290, 1953.

[66] D. Lukkassen and P. Wall. On weak convergence of locally periodic functions. *J. Nonlinear Math. Phys*, 9(1):47–57, 2002.

[67] C. MacDougall and R. E. Polk. Antimicrobial stewardship programs in health care systems. *Clinical Microbiology Reviews*, 18(4):638–656, October 2005.

[68] R. C. MacLean and I. Gudelj. Resource competition and social conflict in experimental populations of yeast. *Nature*, 441(7092):498–501, May 2006.

[69] J. T. Magee. The resistance ratchet: theoretical implications of cyclic selection pressure. *J. of Antimicrobial Chemotherapy*, 2005.

[70] J. Martinez, J. Nicolas, F. Marco, J. Horcajada, G. Garcia-Segarra, A. Trilla, C. Codina, A. Torres, and J. Mensa. Comparison of antimicrobial cycling and mixing strategies in two medical intensive care units. *Crit Care Med*, 34(2):329–336, February 2006.

[71] R. G. Masterton. Antibiotic cycling: more than it might seem? *J Antimicrob Chemother*, 55(1):1–5, Jan 2005.

[72] R. G. Masterton. The new treatment paradigm and the role of carbapenems. *Int J Antimicrob Agents*, 33(2):105–110, Feb 2009.

[73] L. McFarland, G. Elmer, and C. Surawicz. Breaking the cycle: treatment strategies for 163 cases of recurrent clostridium difficile disease. *Am J Gastroenterol*, 97(7):1769–75, 2002.

[74] M. Michael, E. M. Hodson, J. C. Craig, S. Martin, and V. A. Moyer. Short versus standard duration oral antibiotic therapy for acute urinary tract infection in children. *Cochrane Database Syst Rev*, (1):CD003966, 2003.

[75] M. Michea-Hamzehpour, J. Pechere, B. Marchou, and R. Auckenthaler. Combination therapy: a way to limit emergence of resistance? *The American Journal of Medicine*, 80(6, Supplement 2):138 – 142, 1986.

[76] D. Milatovic and I. Braveny. Development of resistance during antibiotic therapy. *Eur. J. Clin. Microbiol.*, (6), 1987.

[77] A. Montero, J. Ariza, X. Corbella, A. Domenech, C. Cabellos, J. Ayats, F. Tubau, C. Borraz, and F. Gudiol. Antibiotic combinations for serious infections caused by carbapenem-resistant Acinetobacter baumannii in a mouse pneumonia model. *J. of Antimicrobial Chemotherapy*, (54), 2004.

[78] M. M. Mwangi, S. W. Wu, Y. Zhou, K. Sieradzki, H. de Lencastre, P. Richardson, D. Bruce, E. Rubin, E. Myers, E. D. Siggia, and A. Tomasz. Tracking the in vivo evolution of multidrug resistance in staphylococcus aureus by whole-genome sequencing. *Proc Natl Acad Sci U S A*, 104(22):9451–6, May 2007.

[79] C. Nathan. Antibiotics at the crossroads. *Nature*, 431(7011):899–902, Oct 2004.

[80] M. S. Niederman. Is "crop rotation" of antibiotics the solution to a "resistant" problem in the icu? *Am J Respir Crit Care Med*, 156(4 Pt 1):1029–31, Oct 1997.

[81] B. C. Opmeer, R. El Moussaoui, P. M. M. Bossuyt, P. Speelman, J. M. Prins, and C. A. J. M. de Borgie. Costs associated with shorter duration of antibiotic therapy in hospitalized patients with mild-to-moderate severe community-acquired pneumonia. *J Antimicrob Chemother*, 60(5):1131–6, Nov 2007.

[82] J. Pépin, N. Saheb, M.-A. Coulombe, M.-E. Alary, M.-P. Corriveau, S. Authier, M. Leblanc, G. Rivard, M. Bettez, V. Primeau, M. Nguyen, C.-E. Jacob, and L. Lanthier. Emergence of fluoroquinolones as the predominant risk factor for clostridium difficile-associated diarrhea: a cohort study during an epidemic in quebec. *Clin Infect Dis*, 41(9):1254–60, Nov 2005.

[83] J. Poehlsgaard and S. Douthwaite. The bacterial ribosome as a target for antibiotics. *Nat Rev Microbiol*, 3(11):870–81, Nov 2005.

[84] D. Raymond, S. Pelletier, T. Crabtree, T. Gleason, L. Hamm, T. Pruett, and R. Sawyer. Impact of a rotating empiric antibiotic schedule on infectious mortality in an intensive care unit. *Crit Care Med*, 29(6):1101–1108, June 2001.

[85] J. Rello and E. Diaz. Optimal use of antibiotics for intubation-associated pneumonia. *Intensive Care Med*, 27(2):337–339, Feb 2001.

[86] T. C. Reluga. Simple models of antibiotic cycling. *Math Med Biol*, 22(2):187–208, Jun 2005.

[87] L. B. Rice. Emerging issues in the management of infections caused by multidrug-resistant gram-negative bacteria. *Cleve Clin J Med*, 74 Suppl 4:S12–20, Aug 2007.

[88] J. Roberts and J. Lipman. Pharmacokinetic issues for antibiotics in the critically ill patient. *Crit. Care Med.*, 37(3):840–851, 2009.

[89] B. M. Roede, P. Bresser, R. El Moussaoui, F. H. Krouwels, B. T. J. van den Berg, P. M. Hooghiemstra, C. A. J. M. de Borgie, P. Speelman, P. M. M. Bossuyt, and J. M. Prins. Three vs. 10 days of amoxycillin-clavulanic acid for type 1 acute exacerbations of chronic obstructive pulmonary disease: a randomised, double-blind study. *Clin Microbiol Infect*, 13(3):284–90, Mar 2007.

[90] R. Rowthorn and G. M. Brown. *Battling resistance to antibiotics and pesticides*, chapter Using Antibiotics when Resistance is Renewable, pages 42–75. 2002.

[91] A. Sandiumenge, E. Diaz, A. Rodriguez, L. Vidaur, L. Canadell, M. Olona, M. Rue, and J. Rello. Impact of diversity of antibiotic use on the development of antimicrobial resistance. *J Antimicrob Chemother*, 57(6):1197–204, Jun 2006.

[92] J. Segreti, H. R. House, and R. E. Siegel. Principles of antibiotic treatment of community-acquired pneumonia in the outpatient setting. *Am J Med*, 118 Suppl 7A:21S–28S, Jul 2005.

[93] A. F. Shorr, S. T. Micek, and M. H. Kollef. Inappropriate therapy for methicillin-resistant staphylococcus aureus: resource utilization and cost implications. *Crit Care Med*, 36(8):2335–40, Aug 2008.

[94] J. D. Siegel, E. Rhinehart, M. Jackson, L. Chiarello, and and the Healthcare Infection Control Practices Advisory Committee. Management of multidrug-resistant organisms in health care settings, 2006. *Am J Infect Control*, 2007.

[95] H. Smith and P. Waltman. *The Theory of the Chemostat*. Cambridge University, 1995.

[96] R. L. Smith, H. L. Evans, T. W. Chong, S. T. McElearney, T. L. Hedrick, B. R. Swenson, W. M. Scheld, T. L. Pruett, and R. G. Sawyer. Reduction in rates of methicillin-resistant staphylococcus aureus infection after introduction of quarterly linezolid-vancomycin cycling in a surgical intensive care unit. *Surg Infect (Larchmt)*, 9(4):423–31, Aug 2008.

[97] B. Spellberg, J. Powers, E. P. Brass, L. Miller, and J. Edwards. Trends in antimicrobial drug development: implications for the future. *Clin Infect Dis*, 38(9):1279–86, 2004.

[98] B. Stecher and W.-D. Hardt. The role of microbiota in infectious disease. *Trends Microbiol*, 16(3):107–114, Mar 2008.

[99] B. Stecher, R. Robbiani, A. W. Walker, A. M. Westendorf, M. Barthel, M. Kremer, S. Chaffron, A. J. Macpherson, J. Buer, J. Parkhill, G. Dougan, C. von Mering, and W.-D. Hardt. Salmonella enterica serovar typhimurium exploits inflammation to compete with the intestinal microbiota. *PLoS Biol*, 5(10):2177–89, Oct 2007.

[100] A. Sullivan, C. Edlund, and C. Nord. Effect of antimicrobial agents on the ecological balance of human microflora. *Lancet Infect Dis*, 1(2):101–114, Sep 2001.

[101] C. M. Surawicz. Treatment of recurrent clostridium difficile-associated disease. *Nat Clin Pract Gastroenterol Hepatol*, 1(1):32–8, Nov 2004.

[102] H. T. Sussman. Thirty years of optimal control: was the path unique? In *Modern Optimal Control*, pages 359–375. Marcel Dekker Inc, 1989.

[103] H. T. Sussmann. The bang-bang problem for certain control systems in GL(n,r). *SIAM J. Control*, 1972.

[104] Y. Takesue, K. Nakajima, K. Ichiki, M. Ishihara, Y. Wada, Y. Takahashi, T. Tsuchida, and H. Ikeuchi. Impact of a hospital-wide programme of heterogeneous antibiotic use on the development of antibiotic-resistant gram-negative bacteria. *J Hosp Infect*, 75(1):28–32, May 2010.

[105] Y. Takesue, H. Ohge, M. Sakashita, T. Sudo, Y. Murakami, K. Uemura, and T. Sueda. Effect of antibiotic heterogeneity on the development of infections with antibiotic-resistant gram-negative organisms in a non-intensive care unit surgical ward. *World J Surg*, 30(7):1269–76, Jul 2006.

[106] P. Toltzis, M. J. Dul, C. Hoyen, A. Salvator, M. Walsh, L. Zetts, and H. Toltzis. The effect of antibiotic rotation on colonization with antibiotic-resistant bacilli in a neonatal intensive care unit. *Pediatrics*, 110(4):707–11, Oct 2002.

[107] V. Trinh, M.-F. Langelier, J. Archambault, and B. Coulombe. Structural perspective on mutations affecting the function of multisubunit rna polymerases. *Microbiol Mol Biol Rev*, 70(1):12–36, Mar 2006.

[108] K. I. Udekwu, N. Parrish, P. Ankomah, F. Baquero, and B. R. Levin. Functional relationship between bacterial cell density and the efficacy of antibiotics. *J Antimicrob Chemother*, 63(4):745–57, Apr 2009.

[109] H. van Loon, M. R. Vriens, A. Fluit, A. Troelstra, C. van der Warken, J. Verhoef, and M. J. M. Bonten. Antibiotic rotation and development of gram-negative antibiotic resistance. *Am J Respir Crit Care Med*, 171:480–487, 2005.

[110] E. J. Vollaard and H. A. Clasener. Colonization resistance. *Antimicrob Agents Chemother*, 38(3):409–14, Mar 1994.

[111] D. K. Warren, H. A. Hill, L. R. Merz, M. H. Kollef, M. K. Hayden, V. J. Fraser, and S. K. Fridkin. Cycling empirical antimicrobial agents to prevent emergence of antimicrobial-resistant gram-negative bacteria among intensive care unit patients. *Crit Care Med*, 32(12):2450–6, Dec 2004.

[112] M. W. M. Wassenberg, G. A. de Wit, B. A. van Hout, and M. J. M. Bonten. Quantifying cost-effectiveness of controlling nosocomial spread of antibiotic-resistant bacteria: the case of mrsa. *PLoS One*, 5(7):e11562, 2010.

[113] D. M. Weinreich, N. F. Delaney, M. A. Depristo, and D. L. Hartl. Darwinian evolution can follow only very few mutational paths to fitter proteins. *Science*, 312(5770):111–4, Apr 2006.

[114] G. S. K. Wolkowicz and Z. Lu. Global dynamics of a mathematical model of competition in the chemostat: General response functions and differential death rates. *SIAM Journal on Applied Mathematics*, 52(1):222–233, 1992.

[115] M. Xu, Y. N. Zhou, B. P. Goldstein, and D. J. Jin. Cross-resistance of Escherichia coli RNA polymerases conferring rifampin resistance to different antibiotics. *J Bacteriol*, 187(8):2783–2792, Apr 2005.

[116] P. Yeh, A. I. Tschumi, and R. Kishony. Functional classification of drugs by properties of their pairwise interactions. *Nat Genet*, 38(4):489–94, Apr 2006.

[117] P. J. Yeh, M. J. Hegreness, A. P. Aiden, and R. Kishony. Drug interactions and the evolution of antibiotic resistance. *Nat Rev Microbiol*, 7(6):460–6, Jun 2009.

[118] M. Zhao, M. Yang, E. Baranov, X. Wang, S. Penman, A. R. Moossa, and R. M. Hoffman. Spatial-temporal imaging of bacterial infection and antibiotic response in intact animals. *PNAS*, 98(17):9814–9818, 2001.

[119] G. R. Zimmermann, J. Lehár, and C. T. Keith. Multi-target therapeutics: when the whole is greater than the sum of the parts. *Drug Discov Today*, 12(1-2):34–42, Jan 2007.

[120] S. Zlitni and E. D. Brown. Drug discovery: Not as fab as we thought. *Nature*, 458(7234):39–40, Mar 2009.