

Selecting Against Antibiotic-Resistant Pathogens: Optimal Treatments in the Presence of Commensal Bacteria

Rafael Peña-Miller · David Lähnemann ·
Hinrich Schulenburg · Martin Ackermann ·
Robert Beardmore

Received: 12 May 2011 / Accepted: 6 October 2011
© Society for Mathematical Biology 2011

Abstract Using optimal control theory as the basic theoretical tool, we investigate the efficacy of different antibiotic treatment protocols in the most exacting of circumstances, described as follows. Viewing a continuous culture device as a proxy for a much more complex host organism, we first inoculate the device with a single bacterial species and deem this the ‘commensal’ bacterium of our host. We then force the commensal to compete for a single carbon source with a rapidly evolving and fitter ‘pathogenic bacterium’, the latter so-named because we wish to use a bacteriostatic antibiotic to drive the pathogen toward low population densities. Constructing a mathematical model to mimic the biology, we do so in such a way that the commensal would be eventually excluded by the pathogen if no antibiotic treatment were given to the host or if the antibiotic were over-deployed. Indeed, in our model, all fixed-dose antibiotic treatment regimens will lead to the eventual loss of the commensal from the host proxy. Despite the obvious gravity of the situation for the commensal bac-

R. Peña-Miller (✉) · R. Beardmore
Department of Mathematics, Imperial College London, London, SW7 2A7, UK
e-mail: R.Pena-Miller@exeter.ac.uk

Present address:
R. Peña-Miller · R. Beardmore
Biosciences, University of Exeter, EX4 4SB, Exeter, UK

D. Lähnemann
Institute of Evolution and Ecology, University of Tübingen, 72076, Tübingen, Germany

H. Schulenburg
Department of Evolutionary Ecology and Genetics, University of Kiel, 24118, Kiel, Germany

M. Ackermann
Department of Environmental Sciences, ETH Zurich, Zurich, Switzerland

M. Ackermann
Department of Environmental Microbiology, Eawag, Dübendorf, Switzerland

terium, we show by example that it is possible to design drug deployment protocols that support the commensal and reduce the pathogen load. This may be achieved by appropriately fluctuating the concentration of drug in the environment; a result that is to be anticipated from the theory optimal control where *bang-bang solutions* may be interpreted as intermittent periods of either maximal and minimal drug deployment. While such ‘antibiotic pulsing’ is near-optimal for a wide range of treatment objectives, we also use this model to evaluate the efficacy of different antibiotic usage strategies to show that dynamically changing antimicrobial therapies may be effective in clearing a bacterial infection even when every ‘static monotherapy’ fails.

Keywords Evolution of antimicrobial resistance · Control theory

1 Introduction


Almost since the very discovery of antibiotics, a conventional wisdom has established which states that bacterial infections should be treated aggressively (Ehrlich 1913). This principle was given as a maxim by Alexander Fleming himself: ‘*if you use penicillin, use enough*’ (Fleming 1964). Such a *hit early and hard* paradigm has impinged upon our understanding of antimicrobial treatment to the point that today it has become the convention to advocate that shortening the duration of a treatment can promote the evolution of drug resistance in pathogens (Beers and Fletcher 2004).

The reasoning goes that a very high drug dose in excess of the *mutant-prevention concentration* reduces treatment duration and so permits less time and fewer bacteria from which resistant genotypes may emerge. However, one might equally argue that higher doses of antibiotic impose stronger selective pressures in favour of drug-resistant bacteria (Yeh et al. 2009). As a result, multi-drug combination therapies are now commonly used to both preserve drug efficacy and to supposedly retard the evolution of resistance (Michel et al. 2008; Zimmermann et al. 2007; Torella et al. 2010).

Despite an ongoing tendency for increasing the complexity of anti-bacterial treatments, itself at least partially resulting from a lack of novel antibiotics (Boucher et al. 2009), we contend that there remains a lack in our understanding of how to deploy just one antibiotic into a single host in order to select best against drug resistant pathogens. Indeed, the issue of how competition for resources of a pathogen with a commensal bacterium mediates the optimal treatment is one that, to our knowledge, is yet to be addressed in the theoretical literature.

It is increasingly clear that broad-spectrum antibiotics may be associated with harmful side effects (Dancer 2004) and can disrupt an innate protective mechanism known as *colonisation resistance*. By shifting the balance between the commensal micro-biota and entero-pathogenic bacteria in favour of the pathogens, antibiotics may even be responsible for increasing the likelihood of infection (Stecher and Hardt 2008; Sullivan et al. 2001; Dethlefsen and Relman 2011). This highlights the importance of considering not only the evolutionary dynamics at play during treatment but, in time, it may become essential to account for the ecological shift within a single host that takes place after the administration of a drug to a patient.

Fig. 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



The purpose of this paper is to deploy techniques of mathematical analysis on a theoretical model calibrated against an experimental dataset to address the following questions: How do we deploy a single, bacteriostatic antibiotic into a host in order to optimally displace the pathogen and suppress its antibiotic resistant mutants? What practicable drug deployment protocols can be used to support a commensal bacterium in its struggle with a fitter pathogen, and so use an antibiotic to promote colonisation resistance?

2 A Possible Control-Theoretic Formulation

The task of constructing a biologically-realistic mathematical model describing the population and evolutionary dynamics of a pathogen in a host organism is both difficult and ongoing. Later in this paper we will propose a very simple model of a microbial microcosm, one that is calibrated against known experimental data that will allow us to design and evaluate the efficacy of different antibiotic usage strategies. But first, in a search for general results, we will stay within a core modelling framework that comprises a large class of genetical models of bacteria growing under both resource limitation and the inhibitory effect of a bacteriostatic anti-microbial agent.

2.1 A Class of Abstract Population Genetics Models

Bio-reactors 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 permit the growth of bacteria in near-constant resource conditions, maintaining bacterial growth rate at a steady-state. Thus, for our proxy of a single host organism, we adapt the standard chemostat whereby a second supply vessel is attached to the chemostat to dynamically control the input of antibiotics into the microcosm while maintaining the overall resource input concentration at a constant value, as illustrated in Fig. 1.

Let us now suppose that this device contains n different commensal bacterial phenotypes competing for the limited resource with m different types of rapidly-evolving pathogens. Let us use vectors $C \in \mathbb{R}^n$ and $P \in \mathbb{R}^m$ to represent the densities of each genotype of the commensals and pathogens, respectively. If we denote the concentration of a single, limiting carbon source by S and use A for the concentration of a bacteriostatic antibiotic present in the growth medium, then we can represent the

state of the host using the variable $\mathbf{x}(t) := (S(t), C(t), P(t), A(t))$, where $t \geq 0$ is time.

We shall concentrate on the following model in the remainder of the paper:

$$\frac{d}{dt}S = d(S_0 - S) - \langle U_c(S), C \rangle - \langle U_p(S), P \rangle, \tag{1a}$$

$$\frac{d}{dt}C = G_c(S, A) \cdot C - dC, \tag{1b}$$

$$\frac{d}{dt}P = \mathcal{M}_\epsilon(G_p(S, A) \cdot P) - dP, \tag{1c}$$

$$\frac{d}{dt}A = \delta(t)A_0 - dA - A(\langle \mathbf{a}_c, C \rangle + \langle \mathbf{a}_p, P \rangle), \tag{1d}$$

where non-negative initial conditions

$$\mathbf{x}(0) = (S(0), C(0), P(0), A(0)) \geq (0, \mathbf{0}, \mathbf{0}, 0)$$

are given with $S(0) \leq S_0$ and $A(0) \leq A_0$, while vectors \mathbf{a}_c and \mathbf{a}_p represent mass-action binding rates of different phenotypes of commensal and pathogen bacteria to the antibiotic molecule. The angled brackets denote standard inner products and A_0 is the concentration of that antibiotic held in a second supply vessel, as illustrated in Fig. 1.

The initial vector $P(0)$ will represent an *isogenic bacterial population*, meaning that $P(0)$ will reside in the span of vectors of the form $(1, 0, \dots, 0)$. The parameter S_0 is the supply concentration of abiotic resource and d the washout rate of the chemostat. The control variable, $\delta(t)$, must be bounded by the dilution rate d and, therefore, for the remainder of this paper we will assume δ is a bounded, measurable function with $0 \leq \delta(t) \leq d$ almost everywhere.

The functions $U_c(S)$ and $U_p(S)$ denote the uptake rates of the limiting resource by commensal and pathogenic bacteria, respectively, and are positive, monotone increasing functions with linear upper bounds. The growth rates $G_c(S, A)$ and $G_p(S, A)$ are vectors which must satisfy the defining property that growth rates are reduced when antibiotic is introduced:

$$G_c(S, A) < G_c(S, 0) \quad \text{and} \quad G_p(S, A) < G_p(S, 0). \tag{2}$$

The non-negative, irreducible matrix \mathcal{M}_ϵ represents a mutation processes, namely the de novo changes in drug-resistance profiles of bacteria that arise due to DNA replication errors. This matrix will be decomposed into the form

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

where the rate of mutations is the same for all phenotypes and equals $\epsilon \in (0, 1)$. Here M is an irreducible stochastic matrix with zeros on its main diagonal, whence $M^T \mathbf{1} = \mathbf{1}$ and so $\mathcal{M}_\epsilon^T \mathbf{1} = \mathbf{1}$ for all $\epsilon > 0$, also. It is important to note in this model that we only allow resistance to be acquired through point mutations occurring during cell division and are not considering other acquisition mechanisms, horizontal gene

transfer (Chan et al. 2011), or gene duplication and amplification (Sandegren and Andersson 2009) for example.

Our goal is to find antibiotic deployment strategies that select against the pathogen while supporting the commensal population. To achieve this, we propose that the following saddle-point objective functional is a useful theoretical device:

$$\mathcal{J}(\delta) := \int_0^T \langle \mathbf{1}, C \rangle - \langle \mathbf{1}, P \rangle dt, \tag{3}$$

where $\mathcal{J} : L^\infty(0, T) \rightarrow \mathbb{R}$ is a well-defined functional. This functional can be written in the form $\mathcal{J}(\delta) = \int_0^T (w, \mathbf{x}(t)) dt$, where w is the following sign-indefinite weight vector

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

We can now state our control problem: *determine a function δ^* that maximises $\mathcal{J}(\delta)$ in (3) subject to the constraint (1a)–(1d).*

Remark 1 The dynamics of (1a)–(1d) can be non-trivial even when δ is a constant function. For example, 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, (1) obeys a competitive exclusion principle in the sense that only one type of either commensal or pathogenic bacteria can persist in the device. It is known that competitive exclusion is not necessarily expected in the presence of a growth inhibitor due to the possible existence of cycles in population densities (Smith and Waltman 1995). However, the article (Wolkowicz and Lu 1992) demonstrates that a chemostat model with a form of inhibition similar to the one used in (1) may indeed obey the competitive exclusion principle.

Definition 1 To ensure the abstract model (1a)–(1d) possesses the full gravity of the situation facing current and future users of antibiotics, we quite purposefully impose a set of restrictions on the physiological and evolutionary responses of the commensal and pathogen populations to the antibiotic. In particular, we will impose the restriction that there are no environmental conditions under which the commensal has the highest growth rate. Thus, if

$$G_c(S, A) = (G_c^1, G_c^2, \dots, G_c^n) \quad \text{and} \quad G_p(S, A) = (G_p^1, G_p^2, \dots, G_p^m)$$

then for each $S \geq 0, A \in [0, A_0]$ there results

$$\max_{1 \leq j \leq n} G_c^j(S, A) < \max_{1 \leq j \leq m} G_p^j(S, A).$$

In this case we say that the pathogen has *complete competitive advantage* over the commensal.

In the remainder of the paper we will show, aided by the \mathcal{J} -optimal controls that the commensal can be held at greater densities than the pathogen despite the latter having complete competitive advantage.

2.2 Control Existence: Near-Optimality of Antibiotic Pulsing

The dynamical equations (1a)–(1d) may be represented, in complete generality, by a smooth non-linear mapping, \mathcal{F} , whereby

$$\frac{d}{dt} \mathbf{x} = \mathcal{F}(\mathbf{x}) + \delta(t) \cdot A_0 \mathbf{e}, \quad \mathbf{x}(0) \in \mathbb{R}^{1+n+m+1}. \tag{4}$$

Here, $\delta(t)$ represents our bounded and measurable control variable to be determined, A_0 denotes the supply concentration of antibiotic to the host and $\mathbf{e} = (0, 0, \dots, 0, 1) \in \mathbb{R}^{1+n+m+1}$ is a fixed vector.

Let us define the set of admissible controls; this is the set of measurable functions taking values almost everywhere between 0 and d ,

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

Theorem 1 *We say that \mathcal{F} is control dissipative if there is a constant $M > 0$ such that for all $\delta \in \Omega$, if $\mathbf{x}(\delta) \in W^{1,\infty}(0, T)$ is the corresponding solution of (1a)–(1d), then $\|\mathbf{x}(\delta)\|_{L^\infty} \leq M$ and $\mathbf{x}(t) \geq \mathbf{0}$ for (almost) all t . There is an optimal control $\delta^* \in \Omega$ such that*

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

if \mathcal{F} is control dissipative.

Proof Suppose that $\sup_{\delta \in \Omega} \mathcal{J}(\delta) = \lim_{k \rightarrow \infty} \mathcal{J}(\delta_k)$ for $(\delta_k)_{k=1}^\infty$ a ‘supremising’ sequence of admissible controls with associated state responses $(\mathbf{x}_k)_{k=1}^\infty \in W^{1,\infty}(0, T)$.

Note that the objective functional $\mathcal{J} : L^\infty \rightarrow \mathbb{R}$ is continuous with respect to the weak* topology on L^∞ as \mathcal{J} is linear with respect to \mathbf{x} and (4) is affine with respect to δ . 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 $\delta_k \xrightarrow{*} \delta^* \in \Omega$ as $k \rightarrow \infty$.

Now, as \mathcal{F} is assumed to be control dissipative, (\mathbf{x}_k) is uniformly bounded in L^∞ and therefore also in $W^{1,\infty}$ by a simple bootstrap argument. Hence, there exists an $M' > 0$ independent of k such that $\|\mathbf{x}_k\|_{W^{1,\infty}} < M'$, and so a further subsequence (\mathbf{x}_k) can be found that we do not relabel for notational convenience, such that $\mathbf{x}_k \xrightarrow{*} \mathbf{x}^*$ as $k \rightarrow \infty$. Furthermore,

$$\begin{aligned} \lim_{k \rightarrow \infty} \mathcal{J}(\delta_k) &= \lim_{k \rightarrow \infty} \int_0^T \langle w, \mathbf{x}_k \rangle dt = \int_0^T \langle w, \lim_{k \rightarrow \infty} \mathbf{x}_k \rangle dt = \int_0^T \langle w, \mathbf{x}^* \rangle dt \\ &= \mathcal{J}(\delta^*), \end{aligned}$$

so that $\mathcal{J}(\delta^*) = \sup\{\mathcal{J}(\delta) : \delta \in \Omega\}$ and, therefore, δ^* is the optimal control with state response \mathbf{x}^* . □

Theorem 2 *An antibiotic deployment regimen $\delta_P(\cdot) \in \Omega$ is said to be a pulsing strategy if $\delta_P(t)$ equals 0 or d for each $t \in [0, T]$. For each $\eta > 0$, there is a pulsing strategy $\delta_P(\cdot)$ such that $\mathcal{J}(\delta_P) > \max_{\delta \in \Omega} \mathcal{J}(\delta) - \eta$.*

Proof Let us define a partition of $[0, T]$ by $\mathcal{I}_N = (t_i)_{i=0}^N$ where $0 = t_0 < t_1 < \dots < t_{N-1} < t_N = T$ and let $PC(\mathcal{I}_N)$ represent the space of piecewise constant functions on \mathcal{I}_N taking *only* the values 0 or d on each sub-interval $(t_i, t_{i+1}]$. Then the set of admissible pulsing protocols is defined as

$$\Omega_P = \{ \delta_p \in \Omega : \exists N \in \mathbb{N}, \mathcal{I}_N \text{ such that } \delta_p \in PC(\mathcal{I}_N) \}.$$

As Ω_P is a weak* dense subset of Ω , see Sussmann (1972), it follows that $\max_{\delta \in \Omega} \mathcal{J}(\delta) = \sup_{\delta \in \Omega_P} \mathcal{J}(\delta)$ and the result follows. \square

Corollary 1 Equation (1a)–(1d) has an optimal control with respect to the functional \mathcal{J} defined in (3). Moreover, solutions of (1) are non-negative for any admissible control δ and both $0 \leq S(t) \leq S_0$ and $0 \leq A(t) \leq A_0$ hold for all $t \geq 0$.

Proof Defining $\Sigma = S_0 - S - \langle \mathbf{1}, C \rangle - \langle \mathbf{1}, P \rangle$, it follows from assumption (2) that $\frac{d}{dt} \Sigma \geq -d \Sigma$ for all $t \geq 0$ and so $\Sigma(t) \leq e^{-dt} \Sigma(0)$, or

$$S + \langle \mathbf{1}, C \rangle + \langle \mathbf{1}, P \rangle \leq S_0 + O(e^{-dt}).$$

Moreover, it is straightforward to prove the non-negativity of solutions of (1) for non-negative initial data, independently of the (admissible) control, and from there it follows that $\frac{d}{dt} A \leq \delta(t) A_0 - dA$. As a result, $A(t) \leq A_0 + O(e^{-dt})$. These inequalities are sufficient to deduce that the nonlinear mapping \mathcal{F} defined by the right-hand side of (1) is control dissipative and so the first part follows.

As $S(0) \leq S_0$ and $A(0) \leq A_0$ were assumed in the discussion following the presentation of (1a)–(1d), a straightforward examination of (1a) and (1d) shows that the final claim of the corollary must be true. \square

Remark 2 Pulsing deployment strategies are described in the medical literature where they are also known as *intermittent dosage schedules* (Eagle et al. 1953; Kirby and Craig 1981; Kunin 1981).

Some pulsing strategy might not, in fact, be optimal for the following reason. Using standard control theory, we can write down a boundary value problem satisfied by the optimal control by forming the Hamiltonian

$$\mathcal{H}(\mathbf{x}, \boldsymbol{\lambda}, \delta) := \langle \boldsymbol{\lambda}, \mathcal{F}(\mathbf{x}) + \delta A_0 \mathbf{e} \rangle + \langle w, \mathbf{x} \rangle,$$

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

$$-\dot{\boldsymbol{\lambda}} = w + \nabla \mathcal{F}(\mathbf{x})^T \boldsymbol{\lambda}, \quad \boldsymbol{\lambda}(T) = 0. \tag{5}$$

A standard result states that the Hamiltonian associated with the constrained optimisation problem determined by (4)–(5) is maximised at all times along the optimal solution $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \delta^*)$:

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

As the Hamiltonian \mathcal{H} is affine with respect to the control δ , then the maximum of $\mathcal{H}(\mathbf{x}^*, \boldsymbol{\lambda}^*, \delta)$ for $0 \leq \delta \leq d$ occurs when $\delta(t) = d$ if $\langle \boldsymbol{\lambda}^*(t), \mathbf{e} \rangle > 0$ or else $\delta(t) = 0$ if $\langle \boldsymbol{\lambda}^*(t), \mathbf{e} \rangle < 0$. As a result of this, $\langle \boldsymbol{\lambda}^*(t), \mathbf{e} \rangle$ is called the *switching function*. If $\langle \boldsymbol{\lambda}^*(t), \mathbf{e} \rangle = 0$ holds for t in some interval of non-zero measure then the control is *singular*; in this situation, we may make no claim about the optimality of antibiotic pulsing strategies.

We would like to emphasise at this point that although it can be shown theoretically that there exist admissible pulsating protocols arbitrarily close to the true \mathcal{J} -optimal control, such a control determined for a specific mathematical model will only be of limited interest if we are unable to implement that control in practice. Indeed, it must be stated that we have little expectation that an optimal control computed numerically using optimisation techniques will turn out even to be a useful concept for real-world experimental or clinical systems. We would like to make clear that we are not proposing to compute optimal pulsating schedules for therapeutic purposes; such a claim would be far beyond the scope of this paper. There are two rationales for the consideration of pulsating strategies. First, they provide a baseline that can be used to compare a theoretical optimum with other strategies that are easier to implement in real-world systems, such as those described in the following section. Second, if one were somehow able to determine the best possible pulsing strategy in a real-world system, the arguments of this section indicate that such an object cannot be too far from the true optimal.

3 Other Therapeutically Relevant Controls

3.1 Optimal Single-Pulse Therapies

The first question one might ask when designing drug usage strategies is *how much antibiotic should be used?* Or, if we fixed the drug dosage, we might then ask *what would the optimal treatment duration be?* Surprisingly, perhaps, there is no consensus in the medical and pharmaceutical communities on how to answer these simple questions.

High doses of antibiotics are believed to be more effective for clearing infections, but they are also responsible for promoting the evolution of drug resistance, thus decreasing drug efficacy for future treatments. For instance, long-term treatments, usually preferred in clinical settings, have been associated with unnecessary side effects (Michael et al. 2003) and with a considerable increase in treatment costs (Opmeer et al. 2007). Furthermore, it has also been shown that patients that receive unnecessarily long treatments may suffer a significantly greater risk of acquiring nosocomial infections (Pépin et al. 2005).

There is clinical evidence that therapies of shorter duration may have the same clinical effect as guidelines that recommend longer-term and higher-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 (el Mousaoui et al. 2006; Roede et al. 2007). 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 (Chastre et al. 2003). While our rather mathematically-defined optimal treatment will almost certainly always remain an elusive concept in practice, computing the *optimal stopping time* in theoretical models is a very simple task.

So, let us deploy antibiotics at the maximum rate d during θ units of time, where $0 \leq \theta \leq T$. Then our single-pulse control redolent of Fleming's 'hit early maxim' may be written as follows:

$$\delta_{sp}(\theta)(t) := \begin{cases} d, & \text{if } 0 \leq t < \theta, \\ 0, & \text{if } T \geq t \geq \theta. \end{cases}$$

The optimal stopping time, denoted θ^* , is computed by numerically maximising the objective functional defined in (3) for each value of $\theta \in [0, T]$. That is, θ^* satisfies

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

3.2 Tapering Dosage Strategies

There is clinical evidence that support *high-dose, short-course* therapies to treat bacterial infections like pneumonia (Segreti et al. 2005), but it is also known that early cessation of antibiotic therapy can be associated with increased mortality in critically ill patients (Rello and Diaz 2001). So, instead of stopping treatment completely, it has been proposed that treatment could start with a high dose of antibiotics whose dosage is gradually reduced with time in order to decrease the side-effects associated with long-term treatments. In clinical settings, drug tapering strategies have been reported to be effective against the recurrent episodes of *Clostridium difficile* infection that are attributed to both the disruption of the intestinal micro-flora and to the persistence of spores (McFarland et al. 2002; Surawicz 2004).

Tapered treatment protocols consist of either a stepwise or continuous decrease in dosage over time. As the statement of our control problem (1) allows the continuous input of an antibiotic into the system, then for the purpose of this paper we will consider the latter case, and thus define a tapering protocol as follows: begin treatment at the maximum dosage, a property expressed through $\delta(0) = d$ and gradually decrease the input rate of the drug with a rate of descent given by a parameter $\alpha \in [0, \pi/2]$. So, mathematically, the control is defined by $\delta_{tap}(\alpha)(t) = \tan(\pi - \alpha)t + d$. As a result, the case where $\alpha = 0$ corresponds to deploying antibiotic at its maximum possible dose at all times, while $\alpha = \pi/2$ denotes the situation where no antibiotic is used at all. Let $t = t_\alpha$ be the first moment in time when $\delta_{tap}(\alpha)(t) = 0$ and we do not deploy any antibiotic after that point, namely when $t_\alpha \leq t \leq T$.

Each rate of descent, α , will produce a different control $\delta_{tap}(\alpha)(t)$ with a corresponding payoff given by $\mathcal{J}(\delta_{tap}(\alpha))$. We will call the optimal tapering parameter, $\alpha^* \in [0, \pi/2]$, that parameter value for which

$$\mathcal{J}(\delta_{tap}(\alpha^*)) = \max_{\alpha \in [0, \pi/2]} \mathcal{J}(\delta_{tap}(\alpha)).$$

3.3 Adaptive Pulsing

In order to precisely determine a near-optimal pulsing treatment or even an optimal tapering strategy in practice, one would need a highly-calibrated mathematical model obtained using refined clinical, genomic, and metabolic data with a detailed understanding of the response of the host to the treatment. This is a daunting task. So, we ask for the following instead: Is it possible to determine effective treatment strategies, although perhaps sub-optimal, based on the patterns of drug-susceptible and drug-resistant pathogens observed over time in the host? In response to this, we propose the development of *feedback controls* that take information directly from observations of the host to adjust the future course of treatment.

The difficulty with this approach is the appropriate design of the feedback. While it may be possible to use linear feedbacks to stabilise the pathogen-free equilibrium of (1a)–(1d), we prefer feedbacks predicated on a 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 necessarily leads to controls that oscillate between periods of deployment and non-deployment, thus we call them *adaptive pulsing protocols*.

Such feedback controls may be costly to implement. They require a continual supply of information regarding the acquisition of drug-resistance genes by the pathogen and the response of the pathogen load to treatment must be tracked. However, it may soon be possible to effect the rapid genotyping of the drug-resistance genes found in a host during infection (Bergeron and Ouellette 1998). So, here we will implement feedback based on a coarse information about the condition of the host and its infection by making observations at discrete time intervals ϑ units apart.

We first choose a feedback law by defining a function or rule that is applied to the current state $\mathbf{x}(t)$ which tells us simply whether to use the drug or not, denote this binary rule by $\delta_f(\cdot)$. The adaptive pulsing control $\delta_{ap}(\vartheta)(t)$ is then given for t in the near-future time interval $(j \cdot \vartheta, (j + 1) \cdot \vartheta)$, where j is an integer, and is determined in response to observations made on the state of the system at the timepoint $\vartheta \cdot j$ by setting $\delta_{ap}(\vartheta)(t) = \delta_f(\mathbf{x}(j \cdot \vartheta))$. The larger ϑ , therefore, the fewer the number of observations that are made on the state of the host, and so the coarser the information utilised when determining the response of the infection to the feedback rule $\delta_f(\cdot)$. Of course, we would like to determine \mathcal{J} -optimal feedback rules, but it is not clear that they should exist for non-linear systems like (1).

4 A Single-Host Model

In order to numerically compare the efficacy of theoretical, optimal drug deployment protocols with the other treatment strategies defined in the previous section, we now complete the specification of the mathematical model (1). We begin by invoking a working assumption that the concentrations of antibiotic both outside and inside each cell are the same. Moreover, we will also assume that per unit time growth rate is simultaneously proportional to transcription rate and to the uptake rate of the limiting carbon source in the chemostat:

$$\text{cell growth rate} \propto \text{resource uptake rate} \times \text{transcription rate.} \quad (6)$$

Hence, the growth rate of each bacterial type is not only determined by the resource availability, S , but also by the concentration of the antibiotic, A . The veracity or predictive power of such an assumption must be tested as it ignores several biological responses. For example, any stress response that sees the down regulation of transcription due to low resource availability cannot be captured by this model.

However, let us continue and represent resource uptake as a monotonic, saturating Monod function in S , that is multiplied by an antibiotic-dependent conversion constant $c(A)$; this converts units of resource into units of bio-mass 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

$$G(S, A) = \frac{V_{\max} S}{K + S} \cdot c(A) \tag{7}$$

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 (6) holds for each cell and assuming that 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 reduces growth rate by reducing the efficiency of the cell in this simple scenario. For example, if the cell were to compensate for the reduction in growth rate suffered in the presence of antibiotic, it might, hypothetically, up-regulate the synthesis of RNA polymerase. But this would have the effect of increasing the resource needed to create a cell and so would lead to a decrease in efficiency.

We then write $c(A)$ as a product of the cell conversion rate in an antibiotic-free environment, $c := c(0)$. Thus, all cell population growth rates defined in (1) will have the form $G(S, A) = c \cdot \frac{V_{\max} S}{K + S} \cdot \gamma(A)$, where $\gamma(A)$ is a dimensionless inhibition coefficient derived from kinetic properties between drug molecules and their targets (as discussed in Appendix A) that satisfies the following key properties:

$$\gamma(0) = 1, \quad \gamma(A) \geq 0 \quad \text{and} \quad \gamma'(A) \leq 0 \quad \text{for all } A \geq 0.$$

In particular, for the numerical examples that appear later in the paper, we shall extend the definition of the growth inhibition function derived in Appendix A by writing $\gamma(A) = (1 + \kappa A)^{-1} = 1 - \kappa A(1 + \kappa A)^{-1}$ and then permitting a two-parameter dependence of the form

$$\bar{\gamma}(A) = 1 - \frac{\kappa_1 A}{1 + \kappa_2 A}. \tag{8}$$

In this case, $\lim_{A \rightarrow \infty} \bar{\gamma}(A) = 1 - \kappa_1/\kappa_2$ and so, provided $\kappa_1 < \kappa_2$, $\bar{\gamma}(A)$ represents a growth inhibition function that is uniformly bounded above complete inhibition.

In order to validate the core growth inhibition assumptions discussed above, we experimentally verified the outcome of a simple model describing the dynamics of an isogenic population of bacteria growing under resource limitation *in a single flask* (not the chemostat, note) in the presence of a bacterio-static antibiotic. This model is closely related to the chemostat model described in (1) and defined by the following




Fig. 2 Dose-response curves of *E. coli* exposed to different antibiotic concentrations. Dotted lines represent optical densities at 24 hours measured experimentally with standard error bars and solid lines are predictions of the model described by (9a)–(9c) with growth parameters $V_{max} = 0.0000251 \mu\text{g}/\text{cell}/\text{h}$, $K = 0.62 \mu\text{g}/\text{ml}$, $c = 1.8 \cdot 10^4 \text{ cells}/\mu\text{g}$. (Left) Erythromycin: inhibition parameters are $\kappa_1 = 3.46 \text{ ml}/\mu\text{g}$ and $\kappa_2 = 0.079 \text{ ml}/\mu\text{g}$; $R^2 = 0.973$. (Right) Doxycycline: inhibition parameters are $\kappa_1 = 0.068 \text{ ml}/\mu\text{g}$ and $\kappa_2 = 0.198 \text{ ml}/\mu\text{g}$; $R^2 = 0.966$

set of differential equations

$$\frac{d}{dt}S = U(S) \cdot B, \tag{9a}$$

$$\frac{d}{dt}B = G(S, A) \cdot B, \tag{9b}$$

$$\frac{d}{dt}A = -aA \cdot B, \tag{9c}$$

with given initial conditions $(S(0), B(0), A(0))$. As before, $S(t)$ is the concentration of resource in the environment at time t , $A(t)$ the concentration of antibiotic, and $B(t)$ the density of the clonal population. Here, $U(S)$ denotes the resource uptake function $U(S) = \frac{V_{max}S}{K+S}$ and we set $G(S, A) = c \cdot U(S) \cdot \gamma(A)$.

So, as above, c is a constant that denotes a conversion rate for resource into biomass and a is an antibiotic binding constant. The latter rests on the assumption that the antibiotic adsorbs to each cell commensurate with the law of mass-action, although the resulting model ignores known mechanisms of avoiding the uptake of antibiotics, with efflux pumps for example.

Following the experimental methods detailed in Hegreness et al. (2008) and using two translation-inhibiting antibiotics, erythromycin and doxycycline, we obtained the dose-response curves for a strain of *E. coli* exposed to different antibiotic concentrations shown in Fig. 2. Dotted lines represent experimental data with standard error bars and solid lines are predictions of the model described by (9a)–(9c).

The targets and the modes of action of both drugs are different: erythromycin is a macrolide that binds to the 50 s ribosomal sub-unit inhibiting protein synthesis and doxycycline is a tetracycline that binds to the 30S sub-unit and inhibits binding of aminoacyl-t-RNA to the acceptor site on the 70S ribosome. However, the model

(9a)–(9c) is sufficient to capture the dose-response data and we remark that (1a)–(1d) are a natural, multiple-phenotype extension of (9a)–(9c) to a chemostat environment. The remainder of the paper now returns to an analysis of the controllability properties of the former.

4.1 Complete Competitive Advantage

The discussion of the previous section indicates that simple population-level models may be able to capture the inhibited growth kinetics of different bacterial mutants with different antibiotic resistance profiles. In the case of *rifampicin*, those mutants gain configurational changes to their β -subunits of RNA polymerase with respect to the wild-type configuration; moreover, a mutation that may prove beneficial if it arises in the presence of antibiotic may well have an associated cost at low concentrations of that same antibiotic.

So, in order to describe this situation mathematically, we suppose that a bacterial clone, the wild-type say, has the following growth rate profile (as usual, in an environment (S, A) of limiting resource concentration S and antibiotic concentration A)

$$G^w(S, A) = e^w \cdot \frac{V_{\max}^w S}{K^w + S} \cdot \gamma^w(A),$$

and that a mutant bacterium has the growth rate profile

$$G^m(S, A) = e^m \cdot \frac{V_{\max}^m S}{K^m + S} \cdot \gamma^m(A).$$

We define the term ‘*antibiotic resistance*’ of the mutant phenotype through the assumed property that

$$\lim_{A \rightarrow \infty} G^w(S, A) < \lim_{A \rightarrow \infty} G^m(S, A),$$

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

$$G^w(S, 0) > G^m(S, 0);$$

the latter means that a mutant bacterium sees a reduction in fitness in an antibiotic-free environment.

We will call this wild-type bacterium *antibiotic susceptible* and the mutant will be called *antibiotic resistant*, even though strictly speaking both types are susceptible to the drug. A phenomenological or phenotypic mutation rate, ϵ , controls the rate at which the offspring cells undergo the transition from susceptible wild-type to resistant mutant per cell division, per unit time.

For the purposes of this paper, it is our aim to engineer a model in which the pathogen population is fitter than the commensal population; an assumption that is based on the idea that a niche occupied by a commensal would be invaded by the pathogen if no antibiotic treatment were given to the host or if the antibiotic were over-deployed. So, suppose that both of the above bacterial types, the wild-type drug-resistant and mutant drug-susceptible, are pathogenic strains and that there is third




Fig. 3 (Color online) The pathogen population has complete competitive advantage: (left) the percentage of growth inhibition as a function of the antibiotic concentration present in the environment; (right) the bacterial growth rate at a fixed resource concentration when the complete competitive advantage property holds. The commensal has a lower growth rate than at least one pathogen in all environments but no one bacterial clone is the fittest in all environments. (Growth rate is quoted per hour)

bacterium of an entirely different species that we deem to be the *commensal*. Suppose further that its growth response is given by the function $G^c(S, A)$. We say that the pathogen population has *complete competitive advantage* if the following condition holds for all $A \geq 0$ and $S \geq 0$:

$$G^c(S, A) < \max\{G^w(S, A), G^m(S, A)\}.$$

Thus, in all abiotic environments and at all antibiotic concentrations, the pathogen has a phenotype that is fitter than the commensal bacterium. In this scenario we might reasonably expect the pathogen to out-compete the commensal for resources, irrespective of how the antibiotic is deployed; this situation is illustrated in Fig. 3.

5 Numerical Examples

We now explicitly impose the condition that the commensal bacteria evolve slowly, so slowly in fact that they form an isogenic and immutable lineage; for definiteness we set $n = 1$. The pathogen can evolve drug resistance, however, and so we set $m = 2$ and write $P = (P_S, P_R)$, where P_S is the susceptible wild-type pathogen and P_R is the antibiotic-resistant, first-order mutant. The pathogen’s mutation matrix, therefore, is for this two-phenotype system is set to

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

The rationale supporting the strong and biologically unrealistic assumption of setting $n = 1$ is the following. While the micro-biota is composed from many different species, so that we should really set $n \gg 1$ and we should allow commensals to evolve

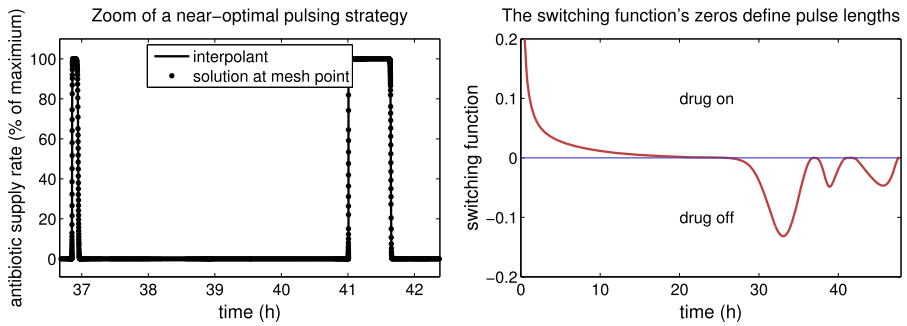


Fig. 4 (Color online) (Left) A zoom of the switching function and the \mathcal{J} -optimal control computed for the set of parameters defined in Table 1: note the non-uniformity of the computational mesh needed to resolve a spike in a near-optimal pulsing control. (Right) The switching function associated with the left-hand plot

drug resistance, we also require a model that has a stable ecology that is adapted to the environment and its carbon sources in the absence of the pathogen. As we only have one carbon source for reasons of modelling simplicity and with the competitive exclusion principle in mind, we set $n = 1$. A useful modelling side-effect of this assumption is that it makes the task of removing the pathogen yet more difficult: we thus have a highly evolvable pathogenic strain in competition with an unresponsive commensal species.

Given this biological modelling construct, in order to evaluate the efficacy of different drug deployment protocols using (1a)–(1d) we will compute the approximations to \mathcal{J} -optimal controls by solving a regularised version of that optimisation problem using `bvp4c` in Matlab (see Appendix B for details of the regularisation).

The switching function of a zoom into a typical solution of this problem, as illustrated in Fig. 4 (right), touches zero precisely at the moments in time when it is necessary to switch the antibiotic supply *on* or *off* in order to maximise the payoff functional \mathcal{J} . Note, as shown in the left-hand diagram of this figure, that `bvp4c` uses a non-uniform computational mesh (shown as crosses superimposed upon the solution itself). The non-uniform meshing property of the computational algorithm underpinning `bvp4c` is almost a necessity when computing the spiking behaviour of optimal pulsing strategies.

5.1 A 48-Hour Race to Colonise

We will compare the result of using six different antibiotic usage strategies in the following specific situation: $C(0)$ and $P_S(0)$ are positive but in some sense ‘small’, well below the value they would achieve at carrying capacity in each other’s absence. We will run the chemostat for 48 hours and seek control strategies to ensure that the value of $C(t)$ is greater than both $P_S(t)$ and $P_R(t)$ when $t = 48$. The treatment will terminate at that time.

In addition to the \mathcal{J} -optimal and other controls defined in Sect. 3, we will implement the following simple feedback control rule on this colonisation problem. It is designed to stop the treatment when the resistant pathogens start to colonise the host;

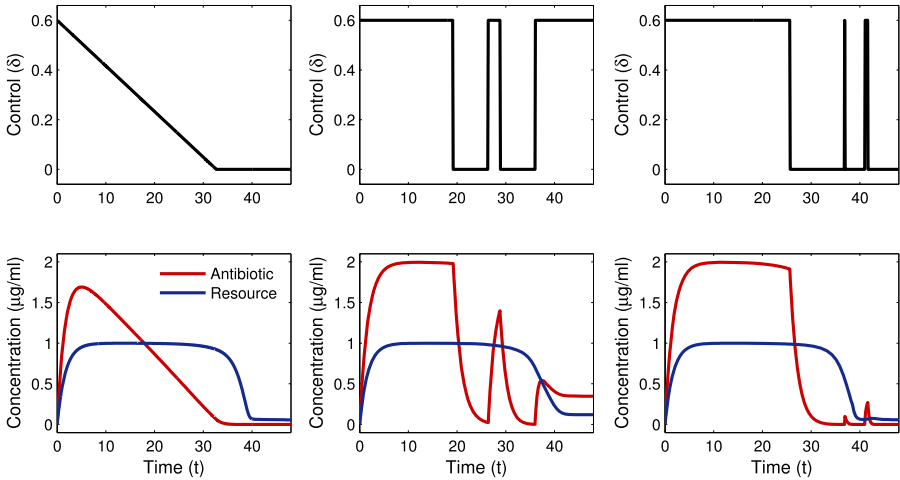


Fig. 5 (Color online) The figures *in the top* illustrate three different control strategies: (*left column*) the optimal tapering control, (*centre column*) the feedback control, (*right column*) the \mathcal{J} -optimal control. *The bottom figures* show the corresponding concentration of resource of antibiotic inside the chemostat as a function of time that result from each control

a heuristic that can be interpreted as saying that the drug should only be used if the pathogen load is sufficiently low:

$$\text{FB: } \delta_f(\mathbf{x}) = \begin{cases} d, & \text{if } P_R < C + P_S, \\ 0, & \text{otherwise,} \end{cases}$$

where, here, \mathbf{x} is to be interpreted as (S, C, P_S, P_R, A) .

So, Fig. 5 (left column) illustrates the \mathcal{J} -optimal control whereas Fig. 5 (centre column) shows an implementation of the rule FB using the same set of biological parameters but based upon sampling the system every $\vartheta = 2$ hours. Notice that the resulting pulsing control is similar to the optimal control, consisting of a large initial interval of drug deployment followed by a series of short pulses.

Figure 5 (right column) illustrates the optimal tapering control. In this figure, control decreases the supply of antibiotic at a rate given by $\alpha^* = 0.72$ and completely ceases the supply of the drug when $T \approx 28$ h. Similarly, we can compute the optimal single-pulse therapy (not shown) and we find that the optimal stopping time for that treatment occurs when $\theta \approx 24$ h.

The responses to all of these different control strategies are illustrated in Fig. 6. They show that although the single-pulse, feedback, and tapering controls support the commensal less than achieved with the \mathcal{J} -optimal control, they still manage to maintain the commensal in excess of the two pathogen phenotypes; a property that no fixed-dose protocol can achieve in this scenario whereby the pathogen has complete competitive advantage.

A feature of note that may be observed in Fig. 7 (left) is that some tapering parameters produce controls with very poor performances and only a small subset of all the tapering parameters lead to controls that out-perform the feedback heuristic

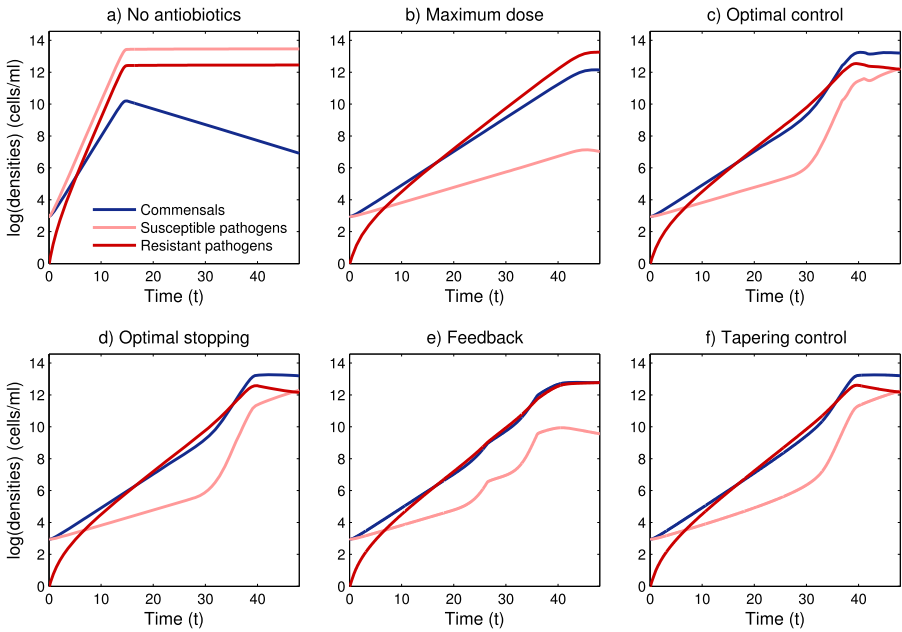


Fig. 6 (Color online) Bacterial densities of pathogens and commensals for different antibiotic deployment protocols lasting 48 h: **(a)** no drug is deployed, **(b)** maximum antibiotic deployment, **(c)** the \mathcal{J} -optimal, **(d)** optimal single-pulse control, **(e)** adaptive pulsing feedback control with measurements every $\vartheta = 2$ h, and **(f)** optimal tapering control. In this example, pathogens out-compete the commensals if no drug is used **(a)** or if it is over-deployed **(b)**, but commensals have the highest final densities with each dynamic strategy implemented **(c)**–**(f)**

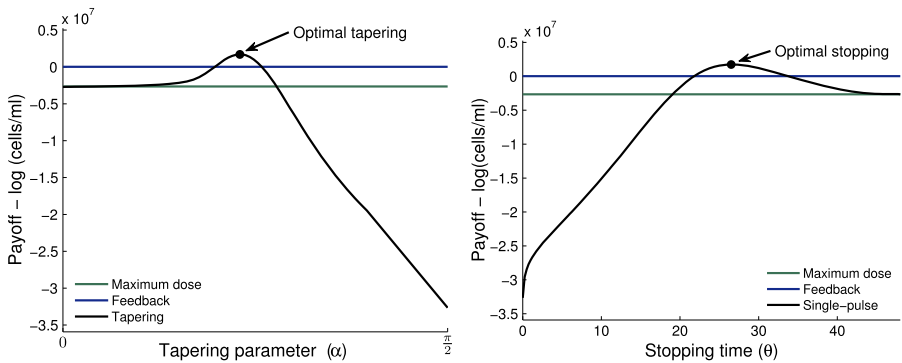


Fig. 7 (Color online) *(Left)* The payoff \mathcal{J} for tapered controls with different tapering parameters $\alpha \in [0, \pi/2]$. *(Right)* The payoffs of single-pulse controls with different stopping times $\theta \in [0, T]$. The range of parameters for which these controls outperform the adaptive pulsing feedback control defined by FB is illustrated in both cases

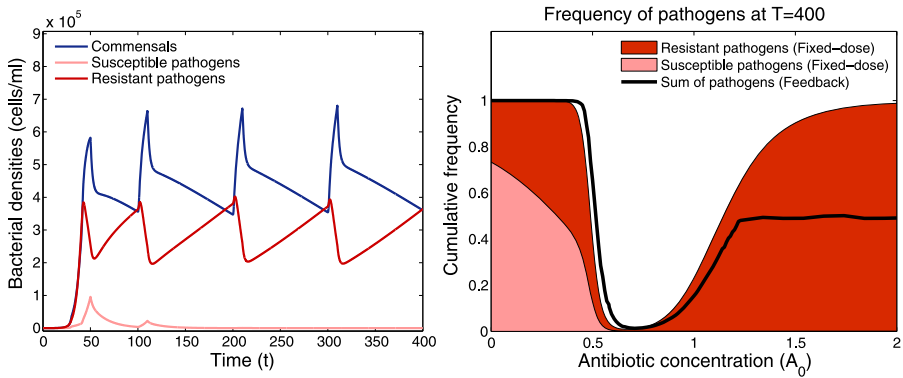


Fig. 8 (Color online) (*Left*) Bacterial densities for the adaptive pulsing strategy at $T = 400$ hours; the strategy FB with observations made every $\vartheta = 10$ hours ensures that the commensal persists in an oscillatory manner with the pathogen. (*Right*) The frequency of pathogens at the end of the experiment as a function of the antibiotic concentration in the supply vessel, A_0 . The shaded areas represent the total frequencies of the susceptible and resistant pathogens with a fixed-dose therapy: high doses of antibiotic select for resistant pathogens, and low doses select for susceptible pathogens. The solid line represents the relative frequency of pathogens at $T = 400$ hours for the adaptive pulsing strategy

or the maximum deployment of antibiotic. An analogous comment applies to single-pulse protocols whereby the majority of stopping times fail to control the pathogens, a property that can be seen in Fig. 7 (right).

6 Conclusions

It is essential to state that the various optimality criteria presented here lead to optimal controls that depend both in the length of the experiment T and also the initial conditions of the system. This makes it very difficult to estimate or implement such optimal solutions in practice, even worse they need not be robust to parametric uncertainties in the model. 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. The feedback controls, however, are probably the most practicable of all the controls we have used. Moreover, it can be effective at supporting the commensal bacterium for quite some time, even when the drug-resistant pathogen fixed in the pathogen population.

To see this, consider Fig. 8 (left) that shows a simulation of an experiment that lasts $T = 400$ hours, where the adaptive pulsing strategy δ_f defined in FB results in an adaptive pulsing protocol that appears from the computations to be eventually periodic (if simulated for an even longer period). Each pulse of drug suppresses the antibiotic-susceptible pathogenic phenotype but the non-deployment of antibiotic suppresses the resistant pathogen due to the costs of resistance that it suffers from. It is clear from this one example that with the correct timings of each pulse, it can be possible to support the commensal population in this simple, chemostat-based model.

6.1 Dosage as a Control Variable

The feedback heuristic FB has two notable features if the maximum possible antibiotic dosage, A_0 , is permitted to vary between 0 and 2 $\mu\text{g/ml}$ for $T = 400$ hours. First, there is an optimal value of the dose, $A_0 \approx 0.7 \mu\text{g/ml}$, at which the frequency of the pathogen population, relative to the commensal, is minimised when the feedback is implemented. Figure 8 (right) compares the result of using the feedback FB with a fixed-dose mono-therapy whereby $\delta(t) = d$ for all t between 0 and T , it shows that the frequency of the pathogens is minimised for both treatments at the same value of A_0 .

However, the feedback has a second notable property: *There is a sense in which its performance is robust to the overuse of antibiotics.* If A_0 is above its optimal value, as can be seen in Fig. 8 (right), the pathogen load cannot exceed the density of the commensal in spite of the pathogen population having complete competitive advantage.

This feature is particularly notable because fixed-dose therapies increase the frequency of resistant pathogens as A_0 increases, whereas the strategy FB yields the same pattern of coexistence between pathogens and commensals independently of the drug concentration A_0 , provided $A_0 > 0.7 \mu\text{g/ml}$ in this particular example.

6.2 Other Optimality Criteria

As a final remark, we propose that even single-drug antibiotic therapies could be successful in clearing pathogens from a host with a single bacteriostatic antibiotic in rather onerous circumstances. However, such therapies are dynamic in time, use observation of the host to inform future course of treatment, and according to theory they may well consist of pulses, namely time intervals both with and without treatment.

There are many possible optimality criteria that could be augmented with a mathematical model like (4) to evaluate different antibiotic treatment strategies. For example, we could include an upper bound on the total amount of antibiotic that can be deployed and seek

$$\max \left\{ \mathcal{J}(\delta) : \delta \in \Omega, \int_0^T \delta(t) \leq \bar{\delta} \right\}$$

where $\bar{\delta}$ is a fixed parameter from which this bound is derived. One can show that antibiotic pulsing provides an effective, near-optimal set of strategies for this problem, too, irrespective of the weight vector w used in the definition of \mathcal{J} .

However, there is absolutely no reason that the treatment length T should not be included in the control variable to be determined, in addition to the antibiotic supply concentration. Hence, if \bar{A} is a physical upper bound on the drug concentration, such as the concentration at which it remains soluble or harmless to the host, we could seek

$$\max \{ \mathcal{J}(\delta) : 0 \leq A_0 \leq \bar{A}, T \geq 0, \delta \in \Omega \}.$$


However, the inclusion of T as an unknown is a substantial change to the mathematical nature of this control problem and it is not clear if pulsing maintains its favourable status for this new formulation.

Appendix A: Growth Inhibition Coefficient

Antibiotics are termed *bactericidal* if they lead to bacterial cell death, for example by inhibiting cell wall biosynthesis and, therefore, enhancing the likelihood of cell lysis. They are said to be *bacteriostatic* if they suppress the bacteria's growth rate by inhibiting DNA replication, RNA transcription, or by interfering with protein synthesis and cellular metabolism but without leading to cell lysis. For example, the bacteriostatic antibiotic *rifampicin* inhibits the function of RNA polymerase during transcription by binding to its β subunit (Trinh et al. 2006), while *streptomycin* targets the 30S ribosomal subunit and inhibits translation (Poehlsgaard and Douthwaite 2005). Resistance to antibiotics may therefore arise through mutations that alter the structure of the protein or subunit targeted by the drug, although such mutations can come at a cost of reduced fitness of the mutant bacterium upon its return to an antibiotic-free environment (Lenski 1998; Andersson and Hughes 2010; Gagneux et al. 2006).

With *rifampicin* in mind and following (Alon 2006), we can describe the inhibition of transcription by an RNA polymerase-binding antibiotic with the following kinetic model. In this highly-simplified model, σ denotes the concentration of free sigma factors in the cell that are needed for the initiation of transcription and P represents the concentration of free RNA polymerase that must bind to a gene's promoter region to initiate the transcription of DNA into mRNA. Hence, P_σ is the concentration of mature transcription complexes; all other transcription factors and phosphorylation processes will be neglected completely from the model. 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 will represent 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 the following kinetic model:



This model assumes that the antibiotic works by producing A -bound polymerase complex P_A and P_σ^A that prevent the binding of the mature transcription complex to the chromosome, thus reducing the rate of RNA production. The latter is given by $k_5 P_\sigma^R$ because P_σ^R is the fraction of gene promoter regions that are bound by a mature transcription complex.

We therefore propose that the mechanism by which an antibiotic like *rifampicin* inhibits RNA transcription may be described phenomenologically with the following

system of equations, suitably augmented with initial conditions:

$$\frac{d}{d\tau}\sigma = -k_1\sigma \cdot P + k_{-1}P_\sigma - k_3\sigma \cdot P_A + k_{-3}P_\sigma^A, \tag{10a}$$

$$\frac{d}{d\tau}A = k_{-2}P_A - k_2P \cdot A, \tag{10b}$$

$$\frac{d}{d\tau}P = -k_1\sigma \cdot P + k_{-1}P_\sigma - k_2A \cdot P + k_{-2}P_A, \tag{10c}$$

$$\frac{d}{d\tau}P_A = k_2P \cdot A - k_{-2}P_A - k_3P_A \cdot \sigma + k_{-3}P_\sigma^A, \tag{10d}$$

$$\frac{d}{d\tau}P_\sigma = k_1\sigma \cdot P - k_{-1}P_\sigma + (k_5 + k_{-4})P_\sigma^R - k_4P_\sigma \cdot R, \tag{10e}$$

$$\frac{d}{d\tau}P_\sigma^A = k_3P_A \cdot \sigma - k_{-3}P_\sigma^A, \tag{10f}$$

$$\frac{d}{d\tau}P_\sigma^R = k_4P_\sigma \cdot R - (k_5 + k_{-4})P_\sigma^R, \tag{10g}$$

$$\frac{d}{d\tau}R = -k_4P_\sigma \cdot R + (k_5 + k_{-4})P_\sigma^R, \tag{10h}$$

$$\frac{d}{d\tau}M = k_5P_\sigma^R. \tag{10i}$$

The variable τ here denotes a very fast timescale relative to cell division that describes the production of a single transcript; the variable t will be used elsewhere to denote the slower timescale pertinent to cell growth. Note that the total concentration of bound and unbound polymerase does not change over time and so we may define the following constant:

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

Let us now assume that the processes are in equilibrium, apart from the production of RNA itself which we assume to occur much more slowly than the binding and unbinding of other complexes. This quasi-steady state assumption imposes

$$\frac{d}{d\tau}M \neq 0 \quad \text{but} \quad \frac{d}{d\tau}R = \frac{d}{d\tau}P_\sigma^R = \frac{d}{d\tau}P_\sigma^A = \frac{d}{d\tau}P_A = \frac{d}{d\tau}P_\sigma = 0.$$

We therefore deduce that $P_A = P \cdot A \cdot k_2/k_{-2}$, and so it will be convenient to define $\kappa_j := k_j/k_{-j}$ for all $j = 1, \dots, 5$. Similarly,

$$P_\sigma^R = \frac{k_4P_\sigma \cdot R}{(k_5 + k_{-4})} \quad \text{and} \quad P_\sigma^A = \kappa_3P_\sigma \cdot A$$

and we now need to determine an expression for P_σ . However, from (10e), $(k_{-1} + k_4R) \cdot P_\sigma = k_1\sigma \cdot P + (k_5 + k_{-4})P_\sigma^R = k_1\sigma \cdot P + k_4P_\sigma \cdot R$ and so we deduce that $P_\sigma = \kappa_1\sigma \cdot P$.

Finally, therefore,

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

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

$$v = v(A) := k_5 P_\sigma^R = \frac{\kappa_1 k_4 k_5}{k_5 + k_{-4}} \cdot \frac{R \cdot \sigma \cdot P_{\text{tot}}}{1 + \kappa_1 \sigma + \kappa_1 \kappa_3 R \cdot \sigma + A(\kappa_2 + \kappa_1 \kappa_3 \sigma)}, \tag{11}$$

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

The rate of RNA production, $v(A)$, in (11) tends to zero as $A \rightarrow \infty$ leading the cell ever-closer to *complete inhibition*. Note that the rate of production of RNA in the absence of antibiotic is given by

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

and, therefore, the relative decrease in transcription rate due to the presence of antibiotic can be described by a dimensionless inhibition function of the form

$$\gamma(A) = \frac{v(A)}{v(0)} = \frac{1}{1 + A \left(\frac{\kappa_2 + \kappa_1 \kappa_3 \sigma}{1 + \kappa_1 \sigma + \kappa_{145} R \cdot \sigma} \right)} =: \frac{1}{1 + \kappa A}, \tag{12}$$

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 (12) to fit the binding affinity data of *rifampicin* and the related antibiotic molecule *rifabutin*, with a similar mode of action, to RNA polymerase using data taken from Xu et al. (2005). The result is shown in Fig. 9 where the R^2 values of the least-squares fitting procedure are given in the legend of the figure. While this gives good and, perhaps, expected agreement between this simple model and antibiotic-target binding data, it does not demonstrate that transcription rate decreases with antibiotic concentration in the same manner as (12), but the data is at least consistent with this hypothesis.

While this is a vast over-simplification of the true molecular biology that does not include features like RNA degradation or the energy cost in terms of the ATP required during transcription, 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) (Xu et al. 2005). This more complex model, however, leads to an inhibition form that closely resembles the one presented above in (12) but whose derivation is somewhat more lengthy, so we have omitted it for the sake of brevity.




Fig. 9 Effects of *Rifampicin* (Rif) and *Rifabutin* (Rfb) on the binding of ^3H rifampin to wild-type RNAP. This data was obtained by competing each antibiotic for radioactively marked Rif, ^3H -rif, with a fixed concentration of 250 nM and a reduction in the bound radioactive molecule is assumed to be due the binding of the non-radioactive drug. (Data taken from Xu et al. 2005, Fig. 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 $1/(1 + \kappa \cdot A)$ where A is the drug concentration and κ is a parameter used in the datafit

Appendix B: An 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 (3) to regularise it by introducing ‘entropy’

$$e(\delta) := \delta \ln \delta + (d - \delta) \ln(d - \delta)$$

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 will then seek a solution of the unconstrained optimal problem

$$\text{find } \delta^*(\eta) \in L^\infty[0, T] \text{ such that } \mathcal{J}_\eta(\delta^*(\eta)) = \sup\{\mathcal{J}_\eta(\delta) : \delta \in L^\infty[0, T]\}. \quad (13)$$

It is possible to prove from the convexity properties of the function $e(\cdot)$, upon defining an optimal control δ^* by

$$\delta^* \in L^\infty[0, T] \text{ satisfies } \mathcal{J}(\delta^*) = \sup\{\mathcal{J}(\delta) : \delta \in L^\infty[0, T], 0 \leq \delta \leq d\}, \quad (14)$$

that $\delta^*(\eta) \xrightarrow{*} \delta^*$ in L^∞ as $\eta \rightarrow 0$, provided $\eta > 0$.

This regularisation allows us to determine $\delta^*(\eta)$ using a numerical continuation procedure akin to annealing in order to compute small- η optimal controls. This procedure starts with η at a high value with an initial guess of $\delta_{\text{guess}}^* = d/2$ and then

reduces η . By coupling this continuation algorithm to an adaptive meshing strategy, we can resolve the switches that are typically present in the optimal control δ^* . We used the Matlab code `bvp4c` (Kierzenka and Shampine 2001) to implement this numerical continuation strategy.

To understand the effect of the regularisation in more detail, a necessary condition for the optimality of δ^* with respect to \mathcal{J} states, using the Pontryagin maximum principle, that the associated Hamiltonian

$$\mathcal{H}(\mathbf{x}, \boldsymbol{\lambda}, \delta) := \langle w, \mathbf{x} \rangle + \langle \boldsymbol{\lambda}, \mathcal{F}(\mathbf{x}) + \delta A_0 \mathbf{e} \rangle$$

is maximised along the optimal trajectory. Here, $\boldsymbol{\lambda} = (\lambda_S, \lambda_C, \lambda_P, \lambda_A) \in \mathbb{R}^{1+n+m+1}$ is an adjoint variable that satisfies the adjoint equation (15a) below

$$-\dot{\boldsymbol{\lambda}} = w + \nabla \mathcal{F}(\mathbf{x})^T \boldsymbol{\lambda}, \quad \boldsymbol{\lambda}(T) = 0, \tag{15a}$$

$$\dot{\mathbf{x}} = \mathcal{F}(\mathbf{x}) + \delta A_0 \mathbf{e}, \quad \mathbf{x}(0) \text{ given.} \tag{15b}$$

The regularised functional \mathcal{J}_η requires a suitably modified Hamiltonian, namely

$$\tilde{\mathcal{H}}(\mathbf{x}, \boldsymbol{\lambda}, \delta; \eta) = \mathcal{H}(\mathbf{x}, \boldsymbol{\lambda}, \delta) - \eta e(\delta). \tag{16}$$

Observe that the only terms in (16) that depend on the control variable δ are

$$\langle \boldsymbol{\lambda}, A_0 \delta \mathbf{e} \rangle - \eta e(\delta) = A_0 \delta \langle \boldsymbol{\lambda}, \mathbf{e} \rangle - \eta e(\delta) = A_0 \delta \cdot \lambda_A - \eta e(\delta) \tag{17}$$

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 can derive an expression for δ in terms of the regularisation parameter η and the adjoint variable λ_A :

$$\delta = \delta(\eta, \lambda_A) := d / \left(1 + e^{-A_0 \lambda_A \eta^{-1}}\right).$$

Substituting this expression for δ into (15) we then obtain a boundary-value differential equation for $(\mathbf{x}, \boldsymbol{\lambda})$ and, from standard regularity results, we then deduce that the regularised optimal control $\delta^*(\eta)(t) = \delta(\eta, \lambda_A(t))$, when it exists, is analytic in t . Moreover, the constraint $0 < \delta(\eta, \lambda_A(t)) < d$ is clearly satisfied for all $t \in [0, T]$, strictly so for $\eta > 0$.

If we now define $\Lambda := (\Lambda_S, \Lambda_C, \Lambda_P, \Lambda_A) = \boldsymbol{\lambda} / \eta$, then Λ satisfies the final-value equation

$$-\dot{\Lambda} = w \eta^{-1} + \nabla \mathcal{F}(\mathbf{x})^T \Lambda, \quad \Lambda(T) = 0, \tag{18}$$

and the control associated with it can be computed from $\delta(\eta)(t) = d / (1 + e^{-A_0 \Lambda_A(t)})$.

This calculation illustrates that (15) is singularly perturbed with respect to η near $\eta = 0$. It is this singular perturbation that results in the appearance of internal layers where the regularised optimal control $\delta^*(\eta)$ switches between the two extreme states,

$\delta = 0$ and $\delta = d$, necessitating the use of non-uniform computational meshes. As the behaviour of $\lambda_A(t)$ decides when the control switches from one regime to the other, this is the *switching function*.

However, (15) is regularly perturbed about $\eta = \infty$ in the sense that we can formally substitute $\eta = \infty$ in (18) 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 $\max_{0 \leq \delta \leq d}(-e(\delta))$, namely when $\delta = d/2$, and we therefore conclude that

$$\lim_{\eta \rightarrow \infty} \delta^*(\eta) = \frac{d}{2}$$

in a C^1 sense. This argument can be made rigorous using the implicit function theorem and it provides an initial guess (the function $\delta_{\text{guess}}^* = d/2$) for the numerical annealing procedure.

Finally, 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 which forces the regularised control to equal $d/2$ at the end-point of the time domain, T .

Appendix C: Model Parameters: Complete Competitive Advantage

Table 1 Model parameters for which the pathogen population has complete competitive advantage over the commensal

Parameter	Description	Value
d	Chemostat dilution rate	$0.6h^{-1}$
ϵ	Rate of phenotypic mutations	0.1 per cell per hour
a_p^* & a_c^*	Antibiotic-cell binding rates for each bacterial types (c , w and m)	$4 \cdot 10^{-6}$ $\mu\text{g}/\text{cell}/\text{hr}$
c	Resource-cell conversion rate in antibiotic-free environment	1×10^6 cell/ μg
K	Half-saturation constant for all bacterial types (c , w and m)	0.06 $\mu\text{g}/\text{ml}$
V_{max}^*	Maximal resource uptake rate of bacterial type *, $V_{\text{max}}^* = \mu_{\text{max}}^*/c$	$\mu_{\text{max}}^c = 1.2/h$, $\mu_{\text{max}}^w = 1.6/h$ $\mu_{\text{max}}^m = 1.0/h$
S_0	Resource supply concentration	1 $\mu\text{g}/\text{ml}$
A_0	Antibiotic supply concentration	2 $\mu\text{g}/\text{ml}$
κ_1^*	Affinity for antibiotic of bacterial type *	$\kappa_1^c = \kappa_1^w = \kappa_1^m = 0.05$ ml/ μg
κ_2^*	Maximal growth inhibition of bacterial type i	$\kappa_2^c = 0.71$ ml/ μg , $\kappa_2^w = 0.5$ ml/ μg , $\kappa_2^m = 0.9$ ml/ μg

References

- Alon, U. (2006). *An introduction to systems biology*. London: Chapman and Hall.
- Andersson, D. I., & Hughes, D. (2010). Antibiotic resistance and its cost: is it possible to reverse resistance? *Nat. Rev., Microbiol.*, *8*(4), 260–271.
- Beers, M. H., & Fletcher, A. J. (2004). *The Merck manual of medical information*. Rahway: Merck. (2nd home ed., online version ed.).
- Bergeron, M., & Ouellette, M. (1998). Preventing antibiotic resistance through rapid genotypic identification of bacteria and of their antibiotic resistance genes in the clinical microbiology laboratory. *Clin. Microbiol.*, *36*(8), 2169–2172.
- Boucher, H. W., Talbot, G. H., Bradley, J. S., Edwards, J. E., Gilbert, D., Rice, L. B., Scheld, M., Spellberg, B., & Bartlett, J. (2009). Bad bugs, no drugs: no escape! an update from the infectious diseases society of America. *Clin. Infect. Dis.*, *48*(1), 1–12.
- Chan, C. X., Beiko, R. G., & Ragan, M. A. (2011). Lateral transfer of genes and gene fragments in staphylococcus extends beyond mobile elements. *J. Bacteriol.*, *193*(15), 3964–3977.
- Chastre, J., Wolff, M., Fagon, J.-Y., Chevret, S., Thomas, F., Wermert, D., Clementi, E., Gonzalez, J., Jusserand, D., Asfar, P., Perrin, D., Fieux, F., & Aubas, S. (PneumA Trial Group) (2003). Comparison of 8 vs 15 days of antibiotic therapy for ventilator-associated pneumonia in adults: a randomized trial. *JAMA*, *290*(19), 2588–2598.
- Dancer, S. J. (2004). How antibiotics can make us sick: the less obvious adverse effects of antimicrobial chemotherapy. *Lancet Infect. Dis.*, *4*(10), 611–619.
- Dethlefsen, L., & Relman, D. A. (2011). Incomplete recovery and individualized responses of the human distal gut microbiota to repeated antibiotic perturbation. *Proc. Natl. Acad. Sci. USA*, *108*(Suppl), 4554–4561. doi:10.1073/pnas.1000087107.
- Eagle, H., Fleischman, R., & Levy, M. (1953). “Continuous” vs. “discontinuous” therapy with penicillin; the effect of the interval between injections on therapeutic efficacy. *N. Engl. J. Med.*, *248*(12), 481–488.
- Ehrlich, P. (1913). Address in pathology on chemotherapeutics: Scientific principles, methods, and results. *Lancet*, *182*, 445–451. doi:10.1016/S0140-6736(01)38705-6.
- el Moussaoui, R., de Borgie, C. A. J. M., van den Broek, P., Hustinx, W. N., Bresser, P., van den Berk, G. E. L., Poley, J.-W., van den Berg, B., Krouwels, F. H., Bonten, M. J. M., Weenink, C., Bossuyt, P. M. M., Speelman, P., Opmeer, B. C., & Prins, J. M. (2006). 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.
- Fleming, A. (1964). Penicillin. In *Nobel lectures, physiology or medicine 1942–1962*. Amsterdam: Elsevier.
- Gagneux, S., Long, C. D., Small, P. M., Van, T., Schoolnik, G. K., & Bohannan, B. J. M. (2006). The competitive cost of antibiotic resistance in mycobacterium tuberculosis. *Science*, *312*(5782), 1944–1946.
- Hegreness, M., Shores, N., Damian, D., Hartl, D., & Kishony, R. (2008). Accelerated evolution of resistance in multidrug environments. *Proc. Natl. Acad. Sci. USA*, *105*(37), 13977–13981.
- Kierzenka, J., & Shampine, L. F. (2001). A BVP solver based on residual control and the Matlab PSE. *ACM Trans. Math. Softw.*, *27*(3), 299–316.
- Kirby, W. M., & Craig, W. A. (1981). Theory and applications of pulse dosing: a summary of the symposium. *Rev. Infect. Dis.*, *3*(1), 1–3.
- Kunin, C. M. (1981). Dosage schedules of antimicrobial agents: a historical review. *Rev. Infect. Dis.*, *3*(1), 4–11.
- Lenski, R. E. (1998). Bacterial evolution and the cost of antibiotic resistance. *Int. Food Microbiol.*, *1*(4), 265–270.
- McFarland, L., Elmer, G., & Surawicz, C. (2002). Breaking the cycle: treatment strategies for 163 cases of recurrent clostridium difficile disease. *Am. J. Gastroenterol.*, *97*(7), 1769–1775.
- Michael, M., Hodson, E. M., Craig, J. C., Martin, S., & Moyer, V. A. (2003). Short versus standard duration oral antibiotic therapy for acute urinary tract infection in children. *Cochrane Database Syst. Rev.*, *1*, CD003966.
- Michel, J.-B., Yeh, P. J., Chait, R., Moellering, R. C. Jr., & Kishony, R. (2008). Drug interactions modulate the potential for evolution of resistance. *Proc. Natl. Acad. Sci. USA*, *105*(39), 14918–14923.
- Opmeer, B. C., El Moussaoui, R., Bossuyt, P. M. M., Speelman, P., Prins, J. M., & de Borgie, C. A. J. M. (2007). Costs associated with shorter duration of antibiotic therapy in hospitalized patients with

- mild-to-moderate severe community-acquired pneumonia. *J. Antimicrob. Chemother.*, 60(5), 1131–1136.
- Pépin, J., Saheb, N., Coulombe, M.-A., Alary, M.-E., Corriveau, M.-P., Authier, S., Leblanc, M., Rivard, G., Bettez, M., Primeau, V., Nguyen, M., Jacob, C.-E., & Lanthier, L. (2005). 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), 1254–1260.
- Poehlsgaard, J., & Douthwaite, S. (2005). The bacterial ribosome as a target for antibiotics. *Nat. Rev. Microbiol.*, 3(11), 870–881.
- Rello, J., & Diaz, E. (2001). Optimal use of antibiotics for intubation-associated pneumonia. *Intensive Care Med.*, 27(2), 337–339.
- Roede, B. M., Bresser, P., El Moussaoui, R., Krouwels, F. H., van den Berg, B. T. J., Hooghiemstra, P. M., de Borgie, C. A. J. M., Speelman, P., Bossuyt, P. M. M., & Prins, J. M. (2007). Three vs 10 days of amoxicillin-clavulanic acid for type 1 acute exacerbations of chronic obstructive pulmonary disease: a randomised, double-blind study. *Clin. Microbiol. Infect.*, 13(3), 284–290.
- Sandegren, L., & Andersson, D. I. (2009). Bacterial gene amplification: implications for the evolution of antibiotic resistance. *Nat. Rev. Microbiol.*, 7(8), 578–588.
- Segreti, J., House, H. R., & Siegel, R. E. (2005). Principles of antibiotic treatment of community-acquired pneumonia in the outpatient setting. *Am. J. Med.*, 118(Suppl 7A), 21S–28S.
- Smith, H., & Waltman, P. (1995). *The theory of the chemostat*. Cambridge: Cambridge University Press.
- Stecher, B., & Hardt, W.-D. (2008). The role of microbiota in infectious disease. *Trends Microbiol.*, 16(3), 107–114.
- Sullivan, A., Edlund, C., & Nord, C. (2001). Effect of antimicrobial agents on the ecological balance of human microflora. *Lancet Infect. Dis.*, 1(2), 101–114.
- Surawicz, C. M. (2004). Treatment of recurrent clostridium difficile-associated disease. *Nat. Clin. Pract. Gastroenterol. Hepatol.*, 1(1), 32–38.
- Sussmann, H. T. (1972). The “bang-bang” problem for certain control systems in $GL(n,r)$. *SIAM J. Control*, 10, 470. doi:10.1137/0310036.
- Torella, J. P., Chait, R., & Kishony, R. (2010). Optimal drug synergy in antimicrobial treatments. *PLoS Comput. Biol.*, 6(6), e1000796.
- Trinh, V., Langelier, M.-F., Archambault, J., & Coulombe, B. (2006). Structural perspective on mutations affecting the function of multisubunit rna polymerases. *Microbiol. Mol. Biol. Rev.*, 70(1), 12–36.
- Wolkowicz, G. S. K., & Lu, Z. (1992). Global dynamics of a mathematical model of competition in the chemostat: general response functions and differential death rates. *SIAM J. Appl. Math.*, 52(1), 222–233.
- Xu, M., Zhou, Y. N., Goldstein, B. P., & Jin, D. J. (2005). Cross-resistance of Escherichia coli RNA polymerases conferring rifampin resistance to different antibiotics. *J. Bacteriol.*, 187(8), 2783–2792.
- Yeh, P. J., Hegreness, M. J., Aiden, A. P., & Kishony, R. (2009). Drug interactions and the evolution of antibiotic resistance. *Nat. Rev. Microbiol.*, 7(6), 460–466.
- Zimmermann, G. R., Lehár, J., & Keith, C. T. (2007). Multi-target therapeutics: when the whole is greater than the sum of the parts. *Drug Discov. Today*, 12(1–2), 34–42.