

Evolutionary dynamics of infectious diseases in finite populations

Jan Humplik^{a,b,c}, Alison L. Hill^{a,d,*}, Martin A. Nowak^{a,e}

^a*Program for Evolutionary Dynamics, Harvard University, One Brattle Square, Cambridge, MA 02138, USA*

^b*Institute of Science and Technology Austria, Am Campus 1, 3400 Klosterneuburg, Austria*

^c*Faculty of Mathematics and Physics, Charles University in Prague, Czech Republic*

^d*Biophysics Program and Harvard-MIT Division of Health Sciences and Technology, Harvard University, Cambridge, MA 02138, USA*

^e*Department of Organismic and Evolutionary Biology, Department of Mathematics, Harvard University, Cambridge, MA 02138, USA*

Abstract

In infectious disease epidemiology the basic reproductive ratio, R_0 , is defined as the average number of new infections caused by a single infected individual in a fully susceptible population. Many models describing competition for hosts between non-interacting pathogen strains in an infinite population lead to the conclusion that selection favors invasion of new strains if and only if they have higher R_0 values than the resident. Here we demonstrate that this picture fails in finite populations. Using a simple stochastic SIS model, we show that in general there is no analogous optimization principle. We find that successive invasions may in some cases lead to strains that infect a smaller fraction of the host population, and that mutually invisable pathogen strains exist. In the limit of weak selection we demonstrate that an optimization principle does exist, although it differs from R_0 maximization. For strains with very large R_0 , we derive an expression for this local fitness function and use it to establish a lower bound for the error caused by neglecting stochastic effects. Furthermore, we apply this weak selection limit to investigate the selection dynamics in the presence of a trade-off between the virulence and the transmission rate of a pathogen.

Keywords: SIS model, stochastic logistic model, basic reproductive ratio, virulence

1. Introduction

Managing infectious diseases of humans, animals and crops requires predicting their dynamics – from short-term spread to long-term evolution. Mathematical models provide a framework for this task (Anderson and May, 1991; Diekmann and Heesterbeek, 2000). The classic modeling approach involves tracking specific compartments of the host population with a system of differential equations. These compartments generally divide individuals into different demographic classes (based on, for example, age or exposure) and different stages of infection (such as susceptible, infected, and recovered) with various pathogen strains. Individuals transition between states based on parameters specific

*To whom correspondence should be addressed.

Email addresses: jhumplik@ist.ac.at (Jan Humplik), alhill@fas.harvard.edu (Alison L. Hill)

to the disease, environment, and mixing patterns. Perhaps the most important insight from such models, when applied deterministically to infinite populations, is the existence of a critical value of the parameters necessary for a disease to cause an epidemic. This threshold is generally represented with the basic reproductive ratio, R_0 , which describes the expected number of secondary infections caused by a single infected individual in a fully susceptible population (reviewed in Heffernan et al. (2005)). For a pathogen-environment-host scenario with $R_0 < 1$, the infection is subcritical and is guaranteed to die out before infecting a substantial fraction of the population, while for $R_0 > 1$, an epidemic can occur. In these deterministic models, $R_0 = 1$ often corresponds to a transcritical bifurcation.

Pathogen populations are genetically diverse due to high mutation rates and large population sizes. Multiple strains are generally competing for the same hosts. Pathogens acquire adaptations as they move from one environment or species to another, and are continually in an evolutionary arms race with their hosts. As well as describing the spread of a disease, in many traditional infectious disease models R_0 is sufficient to describe the evolutionary trajectory of the disease: the selection gradient is in the direction of higher R_0 , and R_0 is the quantity that is maximized by the evolutionarily stable strategy (for example, Anderson and May (1982)).

This strategy of maximizing R_0 to determine the evolutionarily optimal pathogen strategy turns out to have limitations in many systems (Metz et al., 1996, 2008). Violations of this principle have been observed when the simplest models are extended to account for biologically-relevant population dynamics, such as particular types of density-dependent demographic or transmission parameters, frequency-dependent selection, or host-pathogen co-evolution (reviewed in Dieckmann (2002)). More generally, sufficient conditions for R_0 maximization have been derived for one of the most common epidemiological models (the well-mixed SIR model). These include the absence of genotype-by-environment interactions, the presence of only a single transmission pathway, and specific constraints on the density-dependence of mortality (Cortez, 2013). Additionally, the option for pathogens to co-infect hosts or displace resident infections within hosts (“super-infection”) can also result in evolution towards suboptimal R_0 (May and Nowak, 1994; Bonhoeffer and Nowak, 1994; Nowak and May, 1994; May and Nowak, 1995). Due to the convenience of conclusions based on R_0 , exceptions to the “ R_0 maximization” rule remain important to understand. However, most previous work has focused on deterministic models in infinite populations, and the generalizability of these conclusions outside this limit are unknown. In this paper we show that even in the context of the simplest “SIS” disease model, R_0 maximization may fail when we consider finite-sized populations, and the direction of selection may be substantially different in small populations.

The Susceptible-Infected-Susceptible (SIS) Model (Anderson and May, 1991; Dieckmann and Heesterbeek, 2000) is one of the simplest mathematical models describing an infectious process and has been used to describe a variety of phenomena such as communicable diseases, computer viruses and peer-influenced behaviors. The model classifies individuals as either susceptible (healthy) or infected at any point in time. Susceptible individuals can become infected through contact with infected individuals, with a transmission rate β , and once infected, individuals recover at a rate δ . Recovered hosts are once again susceptible to the infection. Individual pathogens are not explicitly tracked, but assumed to only survive in living hosts and be transferred through some form of contact. This model ignores the detailed time-course of the disease within a patient (such as

latent or exposed phases) and any type of long-term immunity. In simplest form, it also assumes that the population consists of randomly and homogeneously mixing individuals, although extensions to network-based contact patterns are common (for example, Eames and Keeling (2002); Cator and Van Mieghem (2013)). Despite these limitations, the small number of parameters of the SIS model (β , δ) means that detailed analysis of the phase space is possible, which often yields results that can be generalized to more complex systems.

While deterministic models are by far the most commonly used in epidemiology, stochastic models are required to answer questions about finite populations and the probability of an epidemic occurring. Analysis of the stochastic SIS model is complicated by the fact that the model contains an absorbing state when all individuals are healthy, which is guaranteed to be reached for all non-zero values of the parameters as time goes to infinity (Nasell, 1995). As a consequence of this absorbing state, there is no longer a critical threshold value of R_0 such that the infection reaches a non-zero equilibrium level, since the only true equilibrium is the zero-infection state. The transition is a bit more subtle and is clear only in the asymptotic limit of large N : for $R_0 < 1$, the disease prevalence decays exponentially, with survival time $\tau = \mathcal{O}(\log(N))$, while for $R_0 > 1$, survival time grows at least as e^{N^α} for some $\alpha > 0$. Also, for large N a long-lived “metastable state” (or “quasi-stationary state”) is reached with an infected level equivalent to the deterministic model (Ovaskainen, 2001; Ganesh et al., 2005; Castellano and Pastor-Satorras, 2010). Various methods have been developed to analyze the quasi-stationary distribution of the stochastic SIS model, by analyzing the limiting behavior of related processes where the absorbing state has been eliminated. One recent approach considers a perturbation to the model to include spontaneous infection of healthy nodes at a rate ϵ and hence eliminates the absorbing state (the ϵ -SIS model). This model has been solved exactly, and for certain small values of ϵ , its threshold and equilibrium behavior approximates that of the SIS model (Van Mieghem and Cator, 2012). Another method considers a related “return process”, in which all absorbing states are re-assigned randomly to transient states, according to an arbitrary “return” distribution. By iteratively finding the equilibrium of this system and using it as the next return distribution, the equilibrium eventually approximates the SIS quasi-stationary distribution (Barbour and Pollett, 2010, 2012). For smaller N , the behavior of the stochastic SIS model is further complicated by the fact that the average number of secondary infections is no longer described by the expression for R_0 in infinite populations, and is consistently smaller (Ross, 2011).

Here we are interested not only in the initial spread but in the evolution of a pathogen obeying SIS dynamics. We will consider the potential for competition between strains with different values of β and δ . In the deterministic, infinite population case, $R_0 = \beta/\delta$, and evolution towards increasing R_0 will tend to increase β and decrease δ when they are uncorrelated. If they are correlated, as is often considered in studies of the evolution of virulence, evolution may proceed towards optimal intermediate values (Antia et al., 1994; Bull, 1994; Lenski and May, 1994; Lipsitch et al., 1996; Levin, 1996; Regoes et al., 2000; Ganusov et al., 2002; Ganusov and Antia, 2003; Ebert and Bull, 2008). We expect that evolutionary outcomes for infectious diseases in finite populations may be different, as important effects have been found in other systems. It is only in finite populations that neutral drift and the fixation probability become meaningful terms. In evolutionary game theory (a framework for studying frequency-dependent selection)(Smith, 1982; Weibull, 1997; Hofbauer and Sigmund, 1998; Nowak and Sigmund, 2004), the conditions for a

strategy to be evolutionarily stable in an infinite population are neither necessary nor sufficient in a finite population (Nowak et al., 2004). Here we demonstrate that similarly interesting results are found for infectious disease evolution in finite populations.

2. The SIS model for two diseases

The model we examine is an extension of the stochastic SIS model that describes the dynamics of two concurrent infections: strain 1, \mathbf{s}_1 , and strain 2, \mathbf{s}_2 . We consider a constant population of N hosts, where each individual can either be in a susceptible state, or be infected with one of the two strains. This model is formally described by a continuous-time birth-death process. The number of individuals infected with the strain \mathbf{s}_i is given by n_i . The number of susceptible individuals is always $N - n_1 - n_2$. The state space of this process consists of pairs of integers $[n_1, n_2]$ such that $n_1 + n_2 \leq N$ (Figure 1A).

A susceptible individual becomes infected with strain \mathbf{s}_i ($i = 1$ or 2) at a rate equal to the product of β_i/N and the number of infected individuals it contacts. The values of β_i will be referred to as the *transmission rates*. We assume that the population is well-mixed, i.e., that all individuals have equally weighted contact with all other individuals at all times. Infected individuals recover from the disease and become susceptible again with a rate δ_i , which we refer to as the *recovery rates*. Equivalently, an individual may die, and - to keep population size constant - immediately be replaced by a new susceptible host. Uninfected hosts die at a rate u . Hosts infected with the strain \mathbf{s}_i die at a rate $u + v_i$, where v_i , the *virulence*, is defined as the increase in the death rate of a host caused by the presence of the infection. The event of either death (with replacement) or recovery of an individual infected with the strain \mathbf{s}_i constitutes a turnover event and occurs at the *turnover rate* $a_i = u + v_i + \delta_i$. The last assumption we make is that individuals may only be infected with a single strain at a time, so that neither superinfection nor coinfection are allowed.

If we define $P_{[n_1, n_2]}(t)$ as the probability that the system is in a state $[n_1, n_2]$ at time t , then we can write a system of differential equations (the *master equations*) describing the time evolution of this process:

$$\begin{aligned} \frac{dP_{[n_1, n_2]}(t)}{dt} = & \left[\frac{\beta_1}{N}(n_1 - 1)P_{[n_1-1, n_2]}(t) + \frac{\beta_2}{N}(n_2 - 1)P_{[n_1, n_2-1]}(t) \right] (N - n_1 - n_2 + 1) \\ & + a_1(n_1 + 1)P_{[n_1+1, n_2]}(t) + a_2(n_2 + 1)P_{[n_1, n_2+1]}(t) \\ & - \left[\left(\frac{\beta_1}{N}n_1 + \frac{\beta_2}{N}n_2 \right) (N - n_1 - n_2) + a_1n_1 + a_2n_2 \right] P_{[n_1, n_2]}(t). \end{aligned} \quad (1)$$

Here we additionally prescribe the function $P_{[n_1, n_2]}(t)$ to be zero outside of the domain $0 \leq n_1 + n_2 \leq N$. Equation(1) requires an initial condition, which is discussed in Section 3. Note that by conditioning on either $n_1 = 0$ or $n_2 = 0$, the model reduces to a standard one-disease stochastic SIS model.

Analogously to the simple SIS model, the state where all individuals are susceptible ($n_1 = n_2 = 0$) is an absorbing state, and since the total number of states is finite, recovery of the entire population is guaranteed as $t \rightarrow \infty$. Consequently, the stationary solution of Eq.(1) is trivial, and different methods are required to understand the important disease dynamics described by the model.

In particular, we are interested in the following scenario. After the resident infection \mathbf{s}_1 has spread among the hosts, a single individual infected with a mutant strain \mathbf{s}_2

emerges in the population (see Section 3 for a precise description). Our goal is to decide whether selection favors the invasion of the mutant strain into the population infected by the resident strain. To answer this question, let us define the *fixation probability of \mathbf{s}_2 invading \mathbf{s}_1* as the probability that strain 1 becomes extinct before strain 2 does. We will say that invasion of strain 2 into strain 1 is favored by selection if its fixation probability is greater than that of a neutral mutant. A mutant strain is neutral if its fixation probability is the same as that of a mutant with parameters identical to those of the resident (i.e., $\beta_2 = \beta_1$, $a_2 = a_1$). In particular, a neutral mutant of \mathbf{s}_1 added into a population containing n individuals infected with \mathbf{s}_1 has a fixation probability $1/(1+n)$. These definitions follow a standard approach to studying adaptation in finite populations (for example Crow and Kimura (1970); Proulx and Day (2002); Nowak et al. (2004)).

To anticipate the results of this analysis, consider a population of only two hosts. One host is infected with strain 1 and the other with strain 2. Since no new infections can occur while both exist in the population, invasion of strain 2 will be favored by selection if its average turnover time is greater than the average turnover time of strain 1 (the turnover times are exponentially distributed random variables). This will happen if $a_1 > a_2$. On the other hand, if the population is infinite, the condition for the invasion of strain 2 to be favored is that $R_1 < R_2$, where $R_i = \beta_i/a_i$ is the basic reproductive ratio. Therefore, selection dynamics depend heavily on the population size, with low turnover rates being more important for smaller populations. One of the goals of this paper is to describe the behavior for population sizes between the two extreme cases of $N = 2$ and $N = \infty$.

The rest of this paper is organized as follows. In Section 3, we precisely define the initial condition for the competition between two strains. In Section 4 we proceed to calculate the fixation probability of a newly introduced strain, and derive a closed form expression in the limit of large values of the basic reproductive ratios. We use these results in Section 5 to analyze which strains are favored to invade different resident strains, and we compare the resulting behavior with known results for infection dynamics in infinite populations. In Section 6, we analyze the weak selection limit to quantify how population size changes the direction of selection and to identify population sizes at which stochastic effects become important. We also use this limit to study the selection dynamics when there is a trade-off between the virulence and the transmission rate.

3. Quasi-stationary distribution as a description of the initial epidemic

We now describe in detail the initial conditions for the invasion of a new disease strain into a population infected with a resident strain at an endemic equilibrium. There are two components to the initial condition: a precise description of the state of the resident infection at this “equilibrium,” and a procedure for the introduction of a mutant strain.

Prior to the introduction of the mutant (strain 2), the dynamics of the resident infection (strain 1) are described by the solution of Eq.(1) conditioned on $n_2 = 0$ (the simple stochastic SIS model). Let us denote the resulting probability distribution $P_{n_1}(t; a_1, \beta_1, N)$. The *quasi-stationary distribution* for this process is defined by conditioning $P_{n_1}(t; a_1, \beta_1, N)$ to non-extinction:

$$Q_{n_1}(R_1, N) = \lim_{t \rightarrow \infty} \frac{P_{n_1}(t; a_1, \beta_1, N)}{1 - P_0(t; a_1, \beta_1, N)}, \quad n_1 \in \{1, \dots, N\}. \quad (2)$$

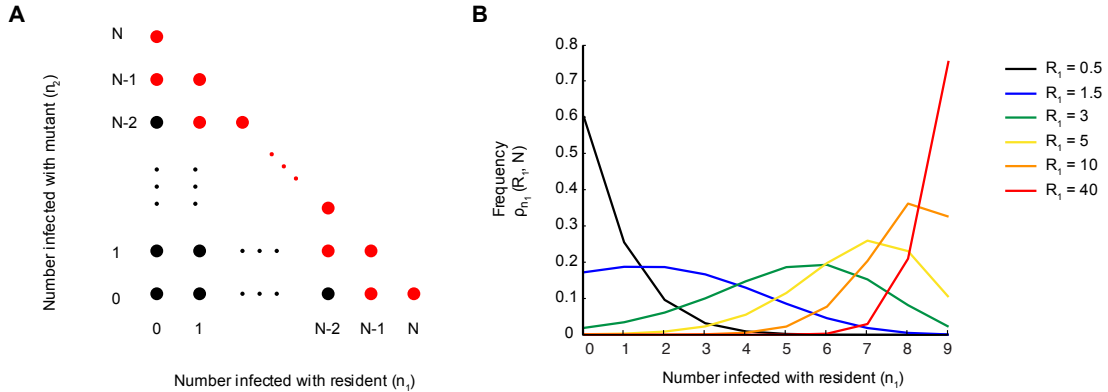


Figure 1: State space and initial conditions for the two-disease stochastic SIS model. A) Each dot represents a possible state of the two infections in a population. Transitions can occur only between neighboring states. Reaching either the lower or left boundary of this state space means that one of the disease strains has gone extinct (and so the remaining strain is said to have “fixed”). The red dots represent the state space in the case of very large basic reproductive ratios (described in Section 4). Restriction of the original Markov process to this state space reduces it to a process on a finite one-dimensional lattice with two absorbing states. B) The initial condition (Eq.(3)) describes the number of hosts infected with the resident strain (n_1) at the instant the mutant strain (\mathbf{s}_2) appears. The mutant strain is introduced when a mutation arises in an infected host. All graphs use $N = 10$. Distributions for several values of the basic reproductive ratio of the resident strain (R_1) are shown. As $R_1 \rightarrow \infty$, the distribution becomes non-zero only for $n_1 = N - 1$.

We assume that strain 1 is in this quasi-stationary state before strain 2 appears. The existence of this limit, and its independence of the initial condition of $P_{n_1}(t; a_1, \beta_1, N)$, was proved in previous work (Mandl, 1960). Although a closed form expression for the parameter dependence of Eq.(2) is not known, many approximations have been derived (Kryscio and Lefèvre, 1989; Nåsell, 2001; Ovaskainen, 2001). Since we could use either a_1 or β_1 to reparametrize time in Eq.(1), the quasi-stationary distribution depends only on R_1 .

Throughout the paper, we will refer to the quantities $R_i = \beta_i/a_i$ as basic reproductive ratios, although, strictly speaking, they are basic reproductive ratios in infinite populations. Basic reproductive ratios in finite populations calculated as the average number of secondary infections differ slightly from the quantities R_i (Ross, 2011).

Even though the quasi-stationary distribution exists for any R_1 , we will mostly restrict our analysis to residents with $R_1 > 1$. Roughly speaking, this ensures that the infection can affect a non-negligible fraction of the population, although the distinction is not as clear as in infinite populations.

We now address the issue of how a mutant is introduced into an infected population. We consider the following scenario: The resident infection \mathbf{s}_1 has reached quasi-stationarity in a population of N individuals, at which point a pathogen in one of the infected individuals undergoes a mutation into a new strain \mathbf{s}_2 . This corresponds to the convention used in evolutionary invasion analysis. In many cases, a weak selection limit is assumed to hold, and the parameters of \mathbf{s}_2 differ only infinitesimally from those of \mathbf{s}_1 .

Under this scenario, after the mutation occurs, the population consists of N hosts in which a single individual is infected with the mutant strain \mathbf{s}_2 and n_1 individuals are infected with the resident strain \mathbf{s}_1 . The variable n_1 is random with the probability

distribution

$$\rho_{n_1}(R_1, N) \equiv Q_{n_1+1}(R_1, N), \quad n_1 \in \{0, \dots, N-1\}. \quad (3)$$

Figure 1B shows several examples of this distribution assuming $N = 10$. Eq.(1) describes a linear system ($d\mathbf{p}/dt = \mathbf{A}\mathbf{p}$ where the elements of \mathbf{p} correspond to $P_{[n_1, n_2]}$) that, in the case of only one disease present, contains one absorbing and N transient states. The quasi-stationary distribution is calculated numerically by evaluating the normalized right eigenvector corresponding to the largest eigenvalue of the matrix \mathbf{A}' , which is obtained from the transition rate matrix (\mathbf{A}) by removing the row and the column corresponding to the absorbing state (Keeling and Ross, 2008).

Although we have chosen a precise definition of the initial condition, there are other alternative definitions that are equally biologically feasible. For example, strain 2 could appear by immigration and a subsequent increase in the total host population size, as opposed to by mutation. This scenario is considered in Section 7.1. Importantly, our results do not change qualitatively with this alternative definition. Many of the consequences of disease competition in finite populations can be inferred from the behavior of strains with very large R_0 , and this behavior is completely independent of the choice of the initial condition (detailed in §7.1).

4. Calculating the fixation probability

The quantity we are interested in is the fixation probability $F(a_2, \beta_2; a_1, \beta_1, N)$ of a mutant. Let $F_{n_1}(a_2, \beta_2; a_1, \beta_1, N)$ denote the conditional fixation probability given that the mutant strain invades a population containing exactly n_1 individuals infected with the resident strain. By the law of total probability, the unconditional fixation probability is given by

$$F(a_2, \beta_2; a_1, \beta_1, N) = \sum_{n_1=0}^{N-1} \rho_{n_1}(R_1, N) F_{n_1}(a_2, \beta_2; a_1, \beta_1, N). \quad (4)$$

The most straightforward way to calculate the conditional fixation probabilities is to consider the embedded Markov chain corresponding to the continuous-time Markov process defined by Eq.(1). This Markov chain has the same state space as the original process, and the transition probabilities between these states are given by

$$\begin{aligned} \mathcal{P}([n_1, n_2] \rightarrow [n_1 + 1, n_2]) &= \frac{\beta_1 n_1 \frac{N-n_1-n_2}{N}}{a_1 n_1 + a_2 n_2 + (\beta_1 n_1 + \beta_2 n_2) \frac{N-n_1-n_2}{N}}, \\ \mathcal{P}([n_1, n_2] \rightarrow [n_1 - 1, n_2]) &= \frac{a_1 n_1}{a_1 n_1 + a_2 n_2 + (\beta_1 n_1 + \beta_2 n_2) \frac{N-n_1-n_2}{N}}, \\ \mathcal{P}([n_1, n_2] \rightarrow [n_1, n_2 + 1]) &= \frac{\beta_2 n_2 \frac{N-n_1-n_2}{N}}{a_1 n_1 + a_2 n_2 + (\beta_1 n_1 + \beta_2 n_2) \frac{N-n_1-n_2}{N}}, \\ \mathcal{P}([n_1, n_2] \rightarrow [n_1, n_2 - 1]) &= \frac{a_2 n_2}{a_1 n_1 + a_2 n_2 + (\beta_1 n_1 + \beta_2 n_2) \frac{N-n_1-n_2}{N}}. \end{aligned} \quad (5)$$

First, let us adjust this Markov chain by changing the state space (see Figure 1A for the graphical representation). We exclude the state $[0, 0]$, and we identify all states with $n_2 = 0$ as one absorbing state, and all states with $n_1 = 0$ as another one. This reduction

in states does not change the outcome of the process since once a strain becomes extinct, it cannot reappear in the population. The next step is to enumerate all the possible states $[n_1, n_2]$ together with these two absorbing states. Constructing the matrix \mathbf{Q} and a vector \mathbf{b} such that $Q_{i,j}$ is the transition probability from a transient state i to a transient state j , and b_i is the transition probability from a transient state i to the absorbing state $n_1 = 0$, we can calculate the vector \mathbf{f} of probabilities of absorption in the state $n_1 = 0$ starting in a state i by solving the linear equation (see Kemeny and Snell (1976))

$$(\mathbf{I} - \mathbf{Q})\mathbf{f} = \mathbf{b}, \quad (6)$$

where \mathbf{I} is the identity matrix.

The conditional fixation probability $F_{n_1}(a_2, \beta_2; a_1, \beta_1, N)$ is equal to the element of \mathbf{f} that corresponds to the state $[n_1, 1]$. Solving Eq.(6) requires the inversion of an $(N^2 - N)/2$ -dimensional matrix and, even for moderate N , the resulting formulas are too complicated to be of any use. Therefore, what remains is to start with certain strain parameters, and numerically solve the Eq.(6). Since the quasi-stationary distribution can be calculated numerically as well, we can also obtain the exact fixation probability $F(a_2, \beta_2; a_1, \beta_1, N)$ for given parameters.

For very large R_i , we can obtain a closed form expression for the fixation probability given by

$$\lim_{R_i \rightarrow \infty} F(a_2, \beta_2; a_1, \beta_1, N) = \frac{(N-1) \left(\frac{\beta_1 a_2}{\beta_2 a_1} - 1 \right)^2}{N - 1 + \frac{a_2}{a_1} - N \frac{\beta_1 a_2}{\beta_2 a_1} + \left(\frac{\beta_1 a_2}{\beta_2 a_1} \right)^N \left(1 - N \frac{\beta_2}{\beta_1} + \frac{a_2}{a_1} (N-1) \right)}. \quad (7)$$

Derivation of this equation is given in Appendix A. It assumes that when taking the limit $R_i \rightarrow \infty$, the ratios a_2/a_1 and β_2/β_1 are kept constant.

Although strains with very large basic reproductive ratios are not necessarily realistic, they demonstrate the effects of finite populations and will be used to present analytic results throughout the paper. There are several important properties of this limit that simplify the analysis. For a single infection in the limit $R_i \rightarrow \infty$, the quasi-stationary distribution is non-zero only for $n_i = N$. The mean time to extinction from quasi-stationarity then follows from a standard result for the SIS model (Ovaskainen, 2001; Kryscio and Lefèvre, 1989)

$$\tau = \frac{1}{a_i} \frac{(N-1)!}{N^N} R_i^{N-1} \left(1 + \mathcal{O} \left(\frac{1}{R_i} \right) \right). \quad (8)$$

Therefore, strains with very large basic reproductive ratio exist in their quasi-stationary state essentially forever, resembling the behavior of infinite populations. Once a second infection is introduced, every individual in the population is always infected with one of the strains. However, the relative prevalence of the two strains can change over time. Even when each individual is infected, the probability of recovery is non-zero for every individual, but, with probability one, any such recovery event is followed by a transmission event. Although the whole population is always infected, every transition between states is accompanied by an infinitesimally short visit to a state with one uninfected host, and so the state space of the population is effectively the set of the red dots in Figure 1A.

5. Evolutionary invasion analysis

5.1. Definitions

Each strain is characterized by its turnover and transmission rates. Since these are positive numbers, the set of all possible mutants, i.e., the phase space of selection, is the first quadrant of \mathbb{R}^2 . For a given resident strain, we will refer to the fixation probability as a function on this phase space as the *fixation probability landscape of the resident*.

The first question we want to ask is the following: given a resident strain 1, what is the set of all strains that are favored by selection to invade it? By definition, this set consists of all strains whose fixation probability is greater than that of strain 1's neutral mutant. The conditional fixation probability of a neutral mutant of \mathbf{s}_1 is equal to $1/(1 + n_1)$. Therefore, the fixation probability of a neutral mutant is

$$F_{\text{neutral}}(R_1, N) = \sum_{n_1=0}^{N-1} \frac{1}{1 + n_1} \rho_{n_1}(R_1, N). \quad (9)$$

The boundary that separates the successful and unsuccessful mutants is given by all the neutral mutants of the strain 1. We will refer to this boundary as the *neutral invasion curve of the resident*. For a fixed a_1 and β_1 , it is defined implicitly in the $a_2 - \beta_2$ plane by the equation

$$F(a_2, \beta_2; a_1, \beta_1, N) = F_{\text{neutral}}(R_1, N). \quad (10)$$

Another question we are concerned with is whether knowing that invasion of *strain 2 into strain 1* is favored by selection implies that invasion of *strain 1 into strain 2* is opposed by selection (and vice versa). For a given resident strain 1, the set of mutants that do not satisfy this property is defined by two boundary curves. One of them is the neutral invasion curve of strain 1 (corresponding to the premise of the implication), and the other is a curve of all mutants such that strain 1 is their neutral mutant (corresponding to the conclusion of the implication). We will refer to this second curve simply as the *second boundary curve*, and, for fixed strain 1 parameters, it is defined implicitly in the $a_2 - \beta_2$ plane by

$$F(a_1, \beta_1; a_2, \beta_2, N) = F_{\text{neutral}}(R_2, N), \quad (11)$$

where R_2 also depends on β_2 and a_2 .

Throughout the rest of this section we carry out this analysis for several distinct regions of the model parameters.

5.2. Invasion analysis if $R_i > 1$ and $N \rightarrow \infty$

In very large populations, strain 1 spreads to eventually reach a steady state in which the fraction of infected individuals is equal to $1 - 1/R_1$. Therefore, when strain 2 enters the population, it sees N/R_1 uninfected hosts. Its basic reproductive ratio with respect to this smaller population is R_2/R_1 , and so it has a non-zero probability of spreading in this population only if $R_2 > R_1$. If satisfied, this probability of spreading is equal to $1 - R_1/R_2$ (derived from the extinction probability for the corresponding birth-death process, see Kendall (1949); Iwasa et al. (2004)). As the number of individuals infected with strain 2 increases, strain 1 will be driven to extinction. Therefore, assuming $R_i > 1$,




Figure 2: Fixation probability landscapes in the following cases: A. $N \rightarrow \infty$; B. $N = 10, R_i \rightarrow \infty$; C. $N = 4, R_i \rightarrow \infty$. Contour plots show the fixation probability (red = 1, blue = 0) for a strain 2 (β_2, a_2) invading strain 1 (β_1, a_1). The strains to the left of the neutral invasion curve (solid white line) are favored to invade strain 1 (fixation probability greater than that of a neutral mutant), while the strains to the right are not. Strains in between the neutral invasion curve and second boundary curve (dashed white line) are not favored to invade strain 1, but strain 1 is also not favored to invade them. In the limit $N \rightarrow \infty$ these two curves coincide, and the strain with larger basic reproductive ratio ($R_i = \beta_i/a_i$) is always favored for invasion. Note that because of the special parameter dependence of the fixation probability in either regime $R_i \rightarrow \infty$ or $N \rightarrow \infty$, these plots are independent of the choice of the resident strain.

the fixation probability in this limit is

$$\lim_{N \rightarrow \infty} F(a_2, \beta_2; a_1, \beta_1, N) = \begin{cases} 0 & \text{if } R_2 < R_1 \\ 1 - \frac{R_1}{R_2} & \text{if } R_2 > R_1 \end{cases}. \quad (12)$$

When $N \rightarrow \infty$ the fixation probability of a neutral mutant is 0, and so the neutral invasion curve is simply given by $\beta_2 = R_1 a_2$. It coincides with the second boundary curve. If strain 2 is favored to invade strain 1, strain 1 is not favored to invade strain 2. The relation $R_1 = R_2$ divides the phase space into equivalence classes, and two strains can coexist in a population only if they belong to the same equivalence class, i.e., if they are neutral mutants of each other. Figure 2A shows a contour plot of the fixation probability landscape of Eq.(12) together with the (overlapping) neutral invasion curve.

5.3. Invasion analysis if N is finite and $R_i \rightarrow \infty$

When the population size is finite but the basic reproductive ratio is very large, the quasi-stationary distribution for the resident strain is non-zero only for $n_1 = N$. As a result, the fixation probability of a neutral mutant is $1/N$ (Eq.(9)), independent of the resident strain. For other mutants, the fixation probability has a closed form expression given by Eq.(7).

Figures 2B,C show contour plots of the fixation probability landscape together with the corresponding neutral invasion and second boundary curves for two different values of N . Comparing these results with the one in Figure 2A, we can immediately infer several new features of the evolutionary competition that are not present if the population is infinite. Firstly, a mutant can be favored to invade a resident even if its basic reproductive ratio is lower than that of the resident (the neutral invasion curves dips below the diagonal in Fig. 2B,C; below line resolution in 2B). Therefore, *the basic reproductive ratio may not be maximized by selection in a finite population*. We will return to this problem in

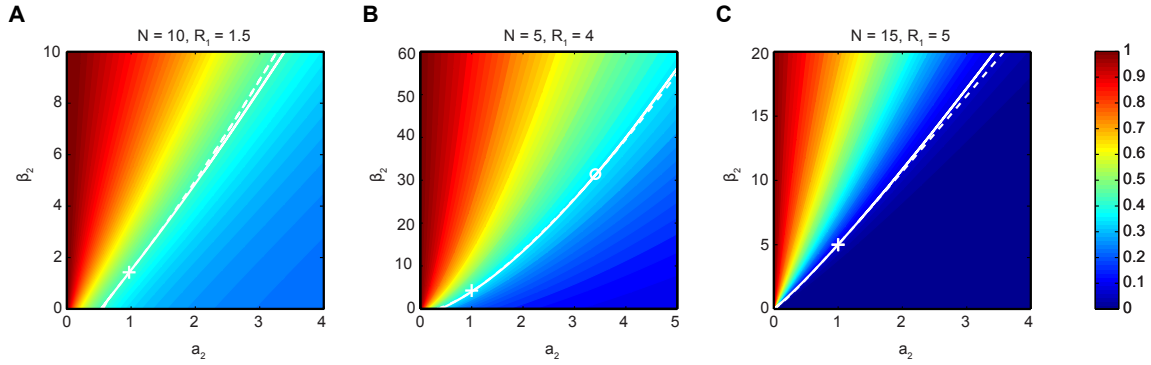


Figure 3: Fixation probability landscape for finite R_i and different population sizes. A. $N = 10$, $R_1 = 1.5$, B. $N = 5$, $R_1 = 4$, C. $N = 15$, $R_1 = 5$. Contour plots show the fixation probability (red = 1, blue = 0) for strain 2 (β_2, a_2) invading strain 1 ($\beta_1 = R_1, a_1 = 1$, marked by white +). The strains to the left of the neutral invasion curve (solid white line) are favored to invade strain 1 (fixation probability greater than that of a neutral mutant), while the strains to the right are not. Only as population size increases is the strain with larger basic reproductive ratio ($R_i = \beta_i/a_i$) always favored for invasion. If the neutral invasion curve is to the right of the second boundary curve (dashed white line), the strains in between them are favored to invade strain 1, and strain 1 is also favored to invade them (in A, above the +; in B, between the + and the O). If the neutral invasion curve is to the left of the second boundary curve, the strains in between them are not favored to invade strain 1, but strain 1 is also not favored to invade them (In B, above the O, in C; above the +). Note that the range of the axes is not the same in each plot.

Section 5.4 when we deal with general values of R_i since in that case, this effect has more profound consequences.

A second emergent effect is that the neutral invasion curve may be located to the left of the second boundary curve. The region in the phase space bounded by these two curves consists of all mutants which are not favored to invade strain 1, nor is strain 1 favored to invade them. This creates a sort of “mutual exclusion” or “status quo” situation between the strains. The existence of such “mutually excluding” strains also implies that, unlike in infinite populations, being a neutral mutant of a strain does not define an equivalence relation on the set of all strains. As a consequence, *it is impossible to assign a value to each strain whose maximization would determine the likely outcome of an invasion in a finite population.*

5.4. Invasion analysis in the case of general N and R_i

5.4.1. Location of the neutral invasion curve and its implications

For general values of the population size and basic reproductive ratio, the fixation probability landscape must be calculated numerically. Several examples, together with the corresponding neutral invasion and second boundary curves, are shown in Figure 3. Unlike in Figure 2 ($R_i \rightarrow \infty$), we do not use the normalized turnover and transmission rates ($a_2/a_1, \beta_2/\beta_1$) because the fixation probability is no longer a function of only these fractions. However, by reparametrizing time in Eq.(1), we can still reduce the number of relevant parameters by one (usually, this is achieved by setting $a_1 = 1$).

In addition to the conclusions made in Section 5.3, the finiteness of R_i is responsible for several new features. Firstly, because the mean number of infected individuals in quasi-stationarity is an increasing function of R_0 , the fact that invasion of a mutant with $R_2 < R_1$ can be favored by selection also means that *selection can favor the invasion*




Figure 4: Contour plot of the mean time to extinction from quasi-stationarity in the SIS model with $N = 5$. Time is measured in multiples of $\frac{1}{a_1} = 1$. Dashed white lines highlight mean extinction time contours that include the strains with basic reproductive ratios 0.7, 2, 3, and 5, and turnover rate $a_1 = 1$. Overlaid solid lines show the neutral invasion curves for these strains. Selection can favor a strain with shorter mean time to extinction.

of a novel strain which is better for the host population. This is in contrast to infinite populations where a favored invader must have a basic reproductive ratio larger than that of the resident, and therefore such a selective sweep must necessarily lead to a higher fraction of infected individuals. The extreme case of this behavior is that, given a resident strain, there may exist mutants which cannot even infect other individuals (i.e., which have $\beta_2 = 0$), and yet still be favored by selection to invade. This is evident from the non-zero x-intercepts in Fig.3A and B. This odd behavior occurs if the timescale of the resident epidemic is sufficiently shorter than the mean turnover time of the mutant. However, this may not be possible in all scenarios because the mean turnover time has an upper bound given by $1/u$ (where u is the natural death rate of hosts).

Secondly, for finite population sizes and finite basic reproductive ratios, selection can also favor faster times to extinction from quasi-stationarity (Figure 4). This is again an example of evolution towards strains that are better for the host population.

5.4.2. Mutual invasibility and mutual exclusion

When R_i is finite, Figure 3 demonstrates that there may exist strains that are favored to invade strain 1, but strain 1 is also favored to invade them (the second boundary curve is left of the neutral invasion curve). Roughly speaking one can imagine this “mutual invasibility” situation as follows. When the infection occurs, there is a pressure towards replacing the old infection in the population. However, when this new infection has almost completely spread, i.e., there is only one individual infected with the old strain, the pressure is reversed towards spreading of the old infection again. Although intuitive, any such classification is complicated in finite populations since the success or failure of an invasion is probabilistic, and the fixation probability cutoff for being “favored by selection” is only a definition relative to the neutral mutant.

Using numerical calculations, we can give a qualitative description of the relative positions of the neutral invasion curve and the second boundary curve, hence characterizing

the mutual invasibility properties of different pairs of strains. For simplicity, we will restrict ourselves to pairs of strains with $R_i > 1$. Without loss of generality, we will also assume that $\beta_2 > \beta_1$ (mutual invasibility is a symmetric relation, and so invasion behavior for pairs of strains with $\beta_2 < \beta_1$ can be inferred by switching the identities of strains 1 and 2). By definition, the neutral invasion curve and the second boundary curve coincide at the resident strain \mathbf{s}_1 where they can either cross (Figure 3A), or touch (Figure 3B,C). In addition to this point of contact, the curves can also cross in the region $\beta_2 > \beta_1$ (Figure 3B). There is at most one such additional crossing, and its presence depends on the value of R_1 . For large values of R_1 , there is no crossing, and all strains with $\beta_2 > \beta_1$ in between the neutral invasion curve and the second boundary curve are such that neither \mathbf{s}_2 nor \mathbf{s}_1 are favored to invade the other (“mutual exclusion”, Figure 3C). Lowering R_1 beyond a certain threshold value, a crossing emerges in the region $\beta_2 > \beta_1$ (Figure 3A (out of plot range), B). Below this crossing but above the resident, there is a set of strains \mathbf{s}_2 such that \mathbf{s}_2 and \mathbf{s}_1 are mutually invisable. Above the crossing, mutual exclusion remains.

6. Weak selection limit

In this section we study long-term evolutionary dynamics in the scenario where each mutation induces only small changes in the strain parameters. This assumption may be violated in particular biologically-relevant scenarios, such as within-host evolution in the presence of strain-specific immune responses, where a point mutation may result in complete immune escape and a large fitness benefit. The ideal framework for studying evolution in this “weak selection” limit is adaptive dynamics (Geritz et al., 1998; Diekmann, 2004), with the fixation probability (Eq.(4)) assuming the role of the invasion fitness (Proulx and Day, 2002).

6.1. Selection gradient

By definition, the selection gradient points from the resident strain (a, β) to the strain with the maximal fixation probability in its infinitesimal neighborhood. Formally, it is given by

$$\vec{g}(a, \beta, N) = \nabla F(a', \beta'; a, \beta, N)|_{(a', \beta')=(a, \beta)}. \quad (13)$$

The operator ∇ represents the gradient with respect to the first two coordinates. We are interested only in the direction of \vec{g} , not its magnitude. The first thing to note is that this gradient depends only on the parameters of the resident strain and on the population size. This allowed us to drop the index “1” used to distinguish the resident strains. Also, as noted before, we can always set $a = 1$ without any loss of generality. Therefore, for a given population size, *strains with the same basic reproductive ratio have the same direction of selection.*

The vector field of Eq.(13) determines a dynamical system that we will refer to as *selection dynamics*. Assuming it is possible, one could integrate $\vec{g}(a, \beta, N)$ to obtain a *local fitness function*. The selection dynamics act to maximize this function. The local fitness function is not unique but its contour lines are (this is also true in an infinite population since the transformation $R_0 \rightarrow f(R_0)$, where $f(x)$ is a strictly monotone function, will not change the outcome of selection). These contour lines satisfy the differential equation

$$\frac{d\beta(a)}{da} = - \frac{\frac{\partial F(a', \beta'; a, \beta(a), N)}{\partial a'}}{\frac{\partial F(a', \beta'; a, \beta(a), N)}{\partial \beta'}} \Bigg|_{(a', \beta')=(a, \beta(a))}. \quad (14)$$

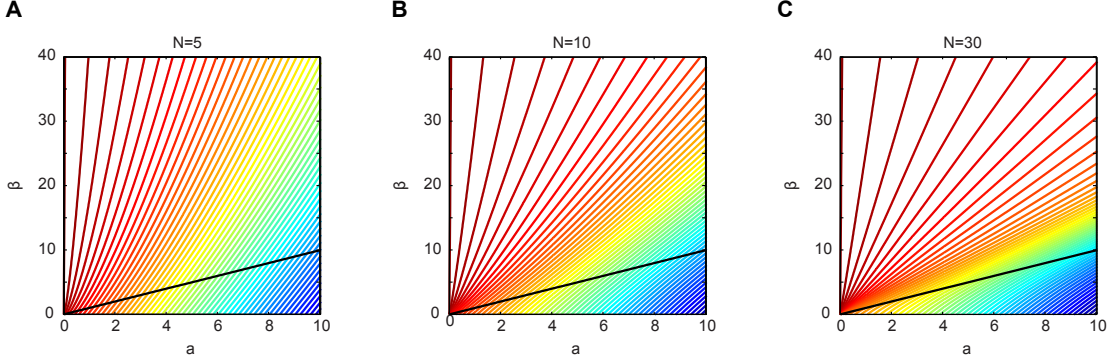


Figure 5: Contour lines for the local fitness function of an infectious disease in a finite population. The direction of selection of a disease strain with parameters a and β is perpendicular to these curves. A. $N = 4$. B. $N = 10$. C. $N = 30$. The selection gradient is determined by the fixation probabilities of neighboring strains in the phase space, and the local fitness function is constructed to have a gradient field pointing in the direction of the selection gradient. The contour lines of this function can be calculated by numerical solution of Eq.(14). The black line represents the boundary $R_0 = 1$ that separates strains that can (above line) from those that cannot (below line) infect a significant fraction of the population.

In the special case $R_0 = \beta/a \rightarrow \infty$, this equation can be solved analytically. Using the expression (7) on the right hand side of Eq.(14), we obtain the differential equation

$$\frac{d\beta(a)}{da} = \left(\frac{N-1}{N-2} \right) \frac{\beta(a)}{a}. \quad (15)$$

This equation can be solved to give

$$\beta(a) = \beta_1 \left(\frac{a}{a_1} \right)^{\frac{N-1}{N-2}}, \quad (16)$$

where we have chosen to describe a contour line that passes through the strain (a_1, β_1) . One can deduce from the form of these contours that when R_0 is very large, a possible choice for the local fitness function is

$$\tilde{R}_0(a, \beta, N) = \frac{\beta}{a^{\frac{N-1}{N-2}}}. \quad (17)$$

As expected, $\lim_{N \rightarrow \infty} \tilde{R}_0(a, \beta, N) = R_0$.

For strains with generic values of R_1 , the corresponding contour lines can be solved for numerically and are more complicated (Figure 5). They do not pass through the origin, and they are not convex.

For any R_0 value, the major difference between the selection dynamics in finite versus infinite populations can be understood as follows. In an infinite population, any contour curve passing through a strain with $R_1 > 1$ is just a ray with a slope equal to R_1 . This is not true for contour curves passing through strains with $R_1 < 1$. The dynamics of these strains can only be described by a stochastic model because the number of infected individuals never reaches macroscopic numbers. The results are that any contour curve that passes through a strain with $R_1 < 1$ does not pass through the origin, is concave, and always lies under the $\beta = a$ line. For finite populations, Figure 5 demonstrates that

the distinction between strains with $R_1 < 1$ and $R_1 > 1$ no longer exists. Any contour contains strains with all possible values of the basic reproductive ratio. These numerical calculations suggest that for any contour there is some a_c (R_c) such that the contour is convex in the region $a > a_c$ (or equivalently for strains with $R > R_c$). The thresholds R_c seem to be close to but not equal to 1.

6.2. Rate of convergence towards the $N \rightarrow \infty$ limit

Because the direction of the selection gradient depends only on the basic reproductive ratio and the population size, it provides a good way to measure the importance of the “finiteness” of the population. For a resident with a basic reproductive ratio $R_0 > 1$, we will denote $\theta(N, R_0)$ the angle between its selection gradient and the a -axis. For an infinite population, this is simply $\arctan(1/R_0)$. To estimate the population sizes for which finite size effects are important, we define the relative difference in selection pressure as compared to infinite population

$$\delta\theta(N, R_0) = \frac{\lim_{N \rightarrow \infty} \theta(N, R_0) - \theta(N, R_0)}{\lim_{N \rightarrow \infty} \theta(N, R_0)} \in [0, 1]. \quad (18)$$

Furthermore, we define the *threshold population size* $N_{\text{th}}(R_0)$ as the population size above which $\delta\theta(N, R_0) < 5\%$ (this choice is arbitrary and was chosen simply to characterize the behavior in regions where an analytic solution is not available). Using Eq.(15), it follows that

$$\lim_{R_0 \rightarrow \infty} \delta\theta(N, R_0) = \frac{1}{N - 1}. \quad (19)$$

Consequently, $\lim_{R_0 \rightarrow \infty} N_{\text{th}}(R_0) = 21$. Numerical calculations confirm that the threshold population size increases with decreasing R_0 . Based on these results we can conclude that, in a population of N hosts, the relative deviation of the selection pressure from the $N \rightarrow \infty$ case is at least $1/(N - 1)$. However, for pathogens with smaller R_0 , this error can be significantly larger. In Figure 6 we show the N dependence of $\delta\theta(N, R_0)$ together with the corresponding threshold population sizes for several different values of R_0 .

6.3. Selection dynamics with a trade-off between the transmission rate and the virulence

In this section we consider the possibility that the virulence (the disease-induced death rate) and the transmission rate of a strain might not be independent parameters. For simplicity, we assume no recovery ($\delta_i = 0$), so that $a_i = u + v_i$. We continue to assume that death is immediately followed by replacement with a susceptible host. The curve that describes all the allowed combinations of virulence and transmission rates will be denoted as $\beta(v)$. The basic condition on $\beta(v)$ is that it is an increasing function. This convention follows from an extensive literature on the evolution of virulence (reviewed in Dieckmann et al. (2002)), and arises from the idea that increasing pathogen numbers within an individual host can lead to increasing virulence (v) but also lead to an increase in transmission rate (β).

Considering again the weak selection limit (and hence infinitesimal mutations), the selection dynamics can be described by restricting the local fitness function to the line $\beta(v)$. Singular points of the selection dynamics correspond to critical points of the local fitness function. Geometrically, a strain with virulence v^* is a singular point if the direction of selection at v^* is perpendicular to the curve $\beta(v)$ at the point v^* . This is equivalent to the

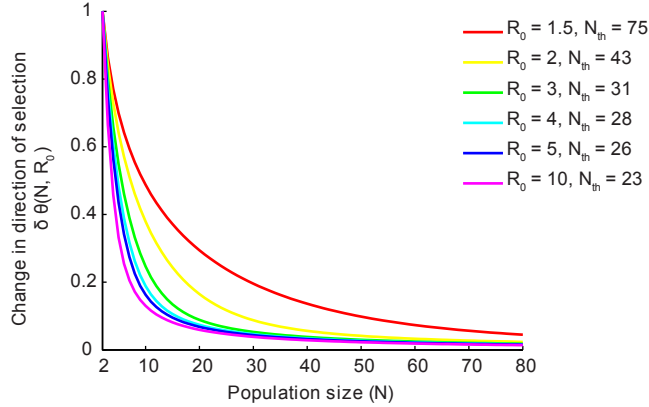


Figure 6: Dependence of the selection gradient on the population size. The direction of selection ($\theta(N, R_0)$) for an infectious disease in a finite population depends on the population size (N) and the basic reproductive ratio of the resident strain (R_0). The angle θ is measured in a counterclockwise direction ($\theta = 0$ corresponds to selection towards lower a). The coefficient $\delta\theta(N, R_0)$ is defined as the relative change in the direction of selection due to the finite size N of the population. The threshold population size (N_{th}) is defined as the population size above which $\delta\theta(N, R_0) < 5\%$.

condition that the curve $\beta(v)$ at the point v^* shares a tangent with the contour curve of the local fitness function that passes through the strain $(v^*, \beta(v^*))$. Stability properties of the singular points follow from relative curvatures of $\beta(v)$ and the contour curve (see Mazancourt and Dieckmann (2004) for a general analysis of this geometrical approach). A short calculation shows that this is equivalent to a standard second derivative test of the local fitness function. Therefore, the local attractors of the selection dynamics (also referred to as convergence stable strategies) correspond to the virulences, v^* , that locally maximize the local fitness function.

A convergence stable strategy can be either an evolutionarily stable strategy, or an evolutionary branching point (Geritz et al. (1998)). In the limit of either $R_0 \rightarrow \infty$ or $N \rightarrow \infty$, mutual invasibility between two neighboring strains is not possible, and so, in these cases, evolutionary branching is ruled out.

In an infinite population, the convergence stable strategies are simply local maxima of $R_0(v) = \beta(v)/(u + v)$ (Figure 7A,B; dashed lines). Therefore, taking first and second derivatives, the necessary conditions for v^* to be a convergence stable strategy are

$$\beta'(v^*) = \frac{\beta(v^*)}{u + v^*} = R_0(v^*), \text{ and } \beta''(v^*) \leq 0. \quad (20)$$

In particular, if $\beta(v)$ has a positive second derivative, the selection pressure either increases the virulence towards infinity or decreases it towards zero, and there cannot be any convergence stable strategy with an intermediate v^* . An example is depicted graphically in Figure 7 (compare the concave constraint in A and B, with one attractor for $N \rightarrow \infty$, to the convex constraint in C and D, with no attractors for $N \rightarrow \infty$).

To quantitatively analyze the situation in a finite population, we assume the limit $R_i \rightarrow \infty$, and use the local fitness function (17) given by

$$\tilde{R}_0(v) = \frac{\beta(v)}{(u + v)^{\frac{N-1}{N-2}}}.$$

By finding the maxima of $\tilde{R}_0(v)$, we determine that the necessary conditions for v^* to be a convergence stable strategy are

$$\beta'(v^*) = \frac{(N-1)}{(N-2)} \frac{\beta(v^*)}{(u+v^*)} = \frac{(N-1)}{(N-2)} R_0(v^*), \text{ and } \beta''(v^*) \leq \frac{(N-1)}{(N-2)^2} \frac{\beta(v^*)}{(u+v^*)^2}. \quad (21)$$

This result introduces two interesting conclusions about the evolution of virulence in finite populations. Firstly, non-zero local attractors can exist even when $\beta(v)$ has a positive curvature (is convex) (Figure 7C,D). Geometrically, the inequality in Eq.(21) means that the second derivative of the constraint $\beta(v)$ at v^* must be less than or equal to the second derivative of the contour curve of the local fitness function passing through $(v^*, \beta(v^*))$. Therefore, a constraint that might continuously increase virulence in an infinite population can lead to a finite convergence stable virulence in a finite population.

Secondly, for constraints $\beta(v)$ with negative curvature, where a local attractor already exists in an infinite population, the location of the attractor changes depending on the population size. We demonstrate this as follows. Denote v^* as the local attractor in an infinite population, i.e., it satisfies the conditions (20). To obtain analytic results, we assume that $\frac{N-1}{N-2} - 1 = \frac{1}{N-2}$ is small. We also assume that the local attractor \bar{v}^* in a finite population is given by a solution of Eq.(21). Since $\frac{1}{N-2}$ is small, it is expected that the new attractor \bar{v}^* is very close to v^* . Linearizing this equation by making a Taylor expansion around v^* on both sides, and keeping only terms linear in $v - v^*$, we demonstrate that

$$\bar{v}^* = v^* - \frac{\beta(v^*)}{\frac{\beta(v^*)}{u+v^*} + (N-2)(u+v^*)|\beta''(v^*)|} \approx v^* - \frac{1}{(N-2)} \frac{R_0(v^*)}{|\beta''(v^*)|}, \quad (22)$$

where we used the condition $\beta''(v^*) < 0$. Therefore, *finite populations shift the convergence stable strategies towards lower virulence*. The geometric interpretation of this result is shown in Figure 7A,B.

The fact that this expectedly small correction is proportional to $R_0(v^*)$, when we consider R_0 to be infinite, is not inconsistent when considering a specific method of attaining this limit. One can think about the $R_0 \rightarrow \infty$ limit as setting $\beta(v) = c\beta_{\text{ref}}(v)$, and taking $c \rightarrow \infty$ (where $\beta_{\text{ref}}(v)$ is some finite reference value of the transmission rate). Then, the rightmost term in Eq.(22) is manifestly finite since it is proportional to $\beta(v^*)/|\beta''(v^*)|$ which is independent of c . One could also mediate the $R_0 \rightarrow \infty$ limit by sending the turnover rate $(u+v)$ to zero. This method is additionally complicated by the fact that the derivative is taken with respect to virulence, however, by employing the same strategy and expressing the turnover rate in terms of some finite reference, one can verify that Eq.(22) is again independent of the scaling factor.

Returning to the case of constraints with positive curvature, we consider what happens to the attractors as the population size is continually reduced from infinity to small values. We determine that for certain constraints, after this local attractor appears at a particular population size, it can then bifurcate into two new local attractors. This feature is demonstrated in Figure 7E using the convex piecewise linear constraint

$$\beta(v) = \begin{cases} 2.6v & \text{if } v < 3.786 \\ -5.3 + 4v & \text{if } v > 3.786 \end{cases}. \quad (23)$$

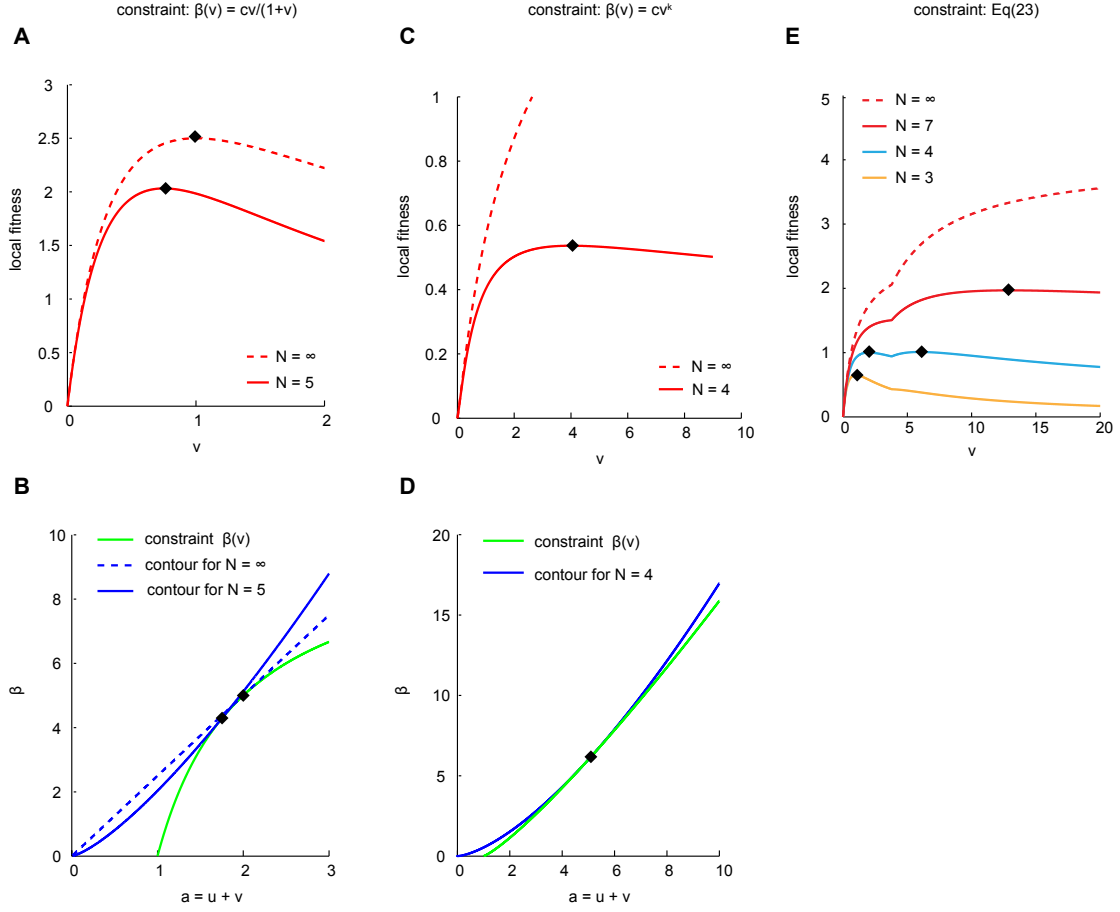


Figure 7: Selection dynamics with a trade-off between the transmission rate and the virulence. In all figures it is assumed that the baseline death rate $u = 1$, and that the recovery rate $\delta = 0$. Black diamonds mark different local attractors. A. The local fitness function (R_0, \tilde{R}_0) versus the virulence, v , for the constraint $\beta(v) = cv/(1+v)$, with $c = 10$. There is at most one local attractor for a concave constraint. The convergence stable virulence v^* shifts to lower values when the population size is decreased from $N = \infty$ (dashed line) to $N = 5$ (solid line). B. The convergence stable virulence occurs at the value of v where the contour line of the local fitness function (blue line) shares a tangent with the constraint $\beta(v)$ (green line), so that the direction of selection is perpendicular to the constraint. C. The local fitness function (R_0, \tilde{R}_0) versus the virulence, v , for the constraint $\beta(v) = cv^k$, with $c = 1.137$ and $k = 1.2$. There is no convergence stable strategy for $N = \infty$ (dashed line). When N decreases, a convergence stable strategy appears (solid line). D. The convergence stable virulence occurs at the value of v where the contour line of the local fitness function (blue line) is parallel to the constraint $\beta(v)$ (green line), so that the direction of selection is perpendicular to the constraint. In finite populations, the point of contact is a local attractor only if the second derivative of the constraint at this point is less than the one of the contour. E. The local fitness function (R_0, \tilde{R}_0) versus the virulence, v , for the convex constraint in Eq.(23). For $N = \infty$, this constraint does not lead to a finite convergence stable virulence. As the population size is decreased, a single local attractor emerges, and when the population size reaches $N = 4$, this attractor bifurcates into two new local attractors. At $N = 3$, one of these disappears again.

While unnecessary for our demonstration, one could smoothen this constraint for example by replacing the sharp step by a Hill function. If the constraint is concave, there can be at most one attractor for any population size, and hence no bifurcation occurs.

Although we have analyzed the finite population case using the limit $R_i \rightarrow \infty$, based on the shape of the contour curves (see Figure 5), we expect similar behavior in the general case as well. However, because mutual invasibility is possible for certain strains, some local attractors might turn out to be evolutionary branching points.

7. Alternative definitions

7.1. Alternative initial conditions

The fixation probability of a mutant is dependent upon the choice of the initial condition for Eq.(1). Throughout the paper, we made the choice (3) that a novel strain was introduced into the population by a process that converted an individual infected with a resident strain to one infected with a mutant strain. Population size is preserved in this process, and it corresponds to conventions used in evolutionary invasion analysis. An example of a different but equally important initial condition is described by the following scenario. The resident infection \mathbf{s}_1 has reached quasi-stationarity in a population of $N - 1$ hosts, and then an individual infected with a novel strain \mathbf{s}_2 migrates to the population. Since the per contact transmission rate of the resident in a population of N individuals is β_1/N , its transmission rate in a population of $N - 1$ individuals is $\beta_1(N - 1)/N$. Therefore, in this scenario, the probability that initially there are n_1 individuals infected with the resident strain is

$$\rho'_{n_1}(R_1, N) \equiv Q_{n_1} \left(R_1 \frac{N - 1}{N}, N - 1 \right), \quad n_1 \in \{1, \dots, N - 1\}. \quad (24)$$

Luckily, there are no qualitative differences in the behavior if one decides to use this initial condition instead of the standard one. In particular, in the limit $R_i \rightarrow \infty$, both initial conditions reduce to the same distribution. This is not surprising. The initial condition with $N - 1$ individuals infected with the resident and 1 individual infected with the mutant is the only meaningful choice when the limit $R_i \rightarrow \infty$ is assumed because any non-infected individual must be immediately infected. That this limit eliminates the ambiguity in the choice of the initial condition is an important fact that further strengthens the view of the large R_i limit as an ideal indicator of the stochastic effects.

7.2. Alternative invasion conditions

Throughout this paper we have been using a definition of neutrality based on the fixation probability of a mutant that is indistinguishable from the resident. While this definition has a clear biological interpretation, it is definitely not the only possible choice. A different choice could be based on a comparison of times the population spends infected with strains 1 and 2 if mutations between the strains occur at some rate u , and the timescale of mutation is much larger than the timescale of invasion. We could say that the strains are neutral mutants of each other if these times are equal. In the limit $R_i \rightarrow \infty$, the total mutation rate of a population is strain-independent because every individual is always infected, and hence is a source of a possible mutation. This allows us to conclude that the times the population spends infected with each disease are equal if and only if

the fixation probability of strain 1 invading strain 2 is the same as the fixation probability of the reverse process. Using Eq.(7), this happens if and only if

$$\frac{\beta_2^{N-2}}{a_2^{N-1}} = \frac{\beta_1^{N-2}}{a_1^{N-1}}. \quad (25)$$

Therefore, at least in the large R_i limit, the “invasion curves” corresponding to this alternative definition of neutrality would coincide with the contour curves of the weak selection limit. In particular, neutrality would define an equivalence relation, and we would not need to worry about the existence of pairs of strains mutually not favored to invade each other. Although this is a nice property for a neutrality definition to satisfy, we believe that our original approach is more fundamental, hence we do not pursue this alternative any further.

8. Discussion

We have analyzed the competition for hosts between two pathogen strains, each of them obeying the dynamics of a stochastic SIS model. We examined whether invasion of a mutant strain into a resident strain at endemic equilibrium is favored by selection, by comparing its fixation probability to that of a neutral mutant. This is one of the simplest models of evolutionary competition of infectious diseases, and yet, our results violate the standard rule of R_0 maximization observed for this same model in infinite populations. We find that strains with R_0 values lower than the resident strain can be favored to invade. In fact, we show that in finite populations, there exist pairs of strains such that neither of them is favored by selection to invade the other (“mutual exclusion”), or that both are favored to invade the other (“mutual invasibility”). This implies that invasion of new pathogens in finite populations cannot be described simply by comparing R_0 values (or values of any other function) of each strain.

Since, for any strain, one can find a mutant with a smaller R_0 value that is favored to invade, evolution can lead to a healthier host population. This effect is impossible in infinite populations for simple SIS models. Though related, these results are not a simple consequence of the dependence of the number of secondary infections on population size (Ross, 2011), but depend on longer-term behavior associated with disease competition.

If we restrict our analysis to mutants whose parameters are infinitesimally close to those of the resident, the situation simplifies considerably. In this weak selection limit, we can locally define a fitness function that is maximized by selection. In the limit of $R_0 \rightarrow \infty$, this local fitness function is equal to $\tilde{R}_0(a, \beta, N) = R_0 \left(\frac{1}{a}\right)^{\frac{1}{N-2}}$ (Eq.(17)). Although the assumption $R_0 \rightarrow \infty$ is not applicable to most pathogens, understanding it is important for three reasons. First, the results in this limit are independent of the choice of the initial condition. Second, it provides a bound for the behavior in the case of general R_0 . Specifically, in a finite population of N hosts, the relative deviation of the selection pressure from the one given by R_0 maximization ($\delta\theta(N, R_0)$) is at least $\frac{1}{N-1}$. Third, the form of the results is very simple, and as such it is useful for a qualitative understanding of the effects of finite populations. If one studies a model based on R_0 maximization, and is interested in what happens when the population size is reduced, replacement of R_0 with \tilde{R}_0 should provide a first insight into the stochastic behavior.

The interpretation of the form of \tilde{R}_0 is simple. The smaller the population size, the

more important it is to have a large turnover time rather than a large transmission rate. This is in opposition to infinite populations, where there is a symmetry between these two attributes, and both are equally important. That turnover time matters more in small populations is also demonstrated by the result that, when the transmission rate depends on the virulence, lowering the population size also lowers the evolutionarily stable virulence. Intuitively, slower turnover (longer lifespan, lower recovery rate, or lower virulence) is more important than a higher transmission rate in small populations because it allows a strain to beat a competitor not only by infecting more individuals, but - because stochastic extinction is inevitable - by “waiting it out.”

A natural generalization of our model is to consider a finite but spatially-structured population. This will be the focus of future work, but intuition suggests that any non-trivial geometry should manifest itself in the exponent in Eq.(17) for the $R_0 \rightarrow \infty$ limit. This in turn would have an effect on the lower bound of the relative deviation from R_0 maximization. The ultimate question is whether for certain population structures this lower bound fails to approach zero as the population size goes to infinity. For example, if one imagines a very large population that consists of many small weakly interacting subpopulations, it is intuitive that some of the features of the small populations should be relevant on the macroscopic scale. Whether and under what conditions this is true still needs to be investigated. A large body of related work has used SIR-type models to study the evolution of virulence in particular types of structured populations (for example, Boots and Sasaki (1999); Haraguchi and Sasaki (2000); Boots et al. (2004); Webb et al. (2013), reviewed in Messinger and Ostling (2009)).

Another necessary step is to incorporate the possibility of continual mutations into the model. Only by studying the resulting mutation-selection balance can one study disease evolution on more general timescales and relate the model to empirical data.

Finally, there is an interesting similarity between our model in the regime $R_i \rightarrow \infty$ and the Moran process with frequency-dependent selection (Nowak et al., 2004). Both models assume a population of N hosts. In the Moran process, every individual carries one of two alleles A or B. In the limit $R_i \rightarrow \infty$ of our model, every individual is infected with one of two strains \mathbf{s}_1 or \mathbf{s}_2 . In the Moran process, every event consists of choosing a random individual for death, and a random individual for reproduction. In our model, every event consists of choosing a random individual for a recovery (or death), and a random individual for infecting the newly recovered host. There are, however, two important differences. First, in the Moran process, all individuals have the same probability of dying. This is not so in our model, where the probability of being chosen for recovery depends on a_i . Second, in the Moran process, the events of choosing an individual for death and choosing an individual for reproduction are independent. In our model, the event of choosing an individual who will infect the newly recovered host depends on whether it was an individual infected with \mathbf{s}_1 or \mathbf{s}_2 who recovered. This second difference is the reason the limit $R_i \rightarrow \infty$ of our model cannot be mapped on a standard Moran process with frequency-dependent selection.

Acknowledgements

The authors thank Daniel Rosenbloom, Benjamin Allen, Oliver Hauser and Gabriel Leventhal for fruitful discussions. JH received support from the Zdenek Bakala Foundation and the Mobility Fund of Charles University in Prague. MAN received support from

the John Templeton Foundation. MAN and ALH received support from the Foundational Questions in Evolutionary Biology Fund.

- Anderson, R. M., May, R. M., 1982. Coevolution of hosts and parasites. *Parasitology* 85 (02), 411–426.
- Anderson, R. M., May, R. M., 1991. *Infectious diseases of humans: dynamics and control*. Oxford University Press, USA.
- Antia, R., Levin, B. R., May, R. M., 1994. Within-host population dynamics and the evolution and maintenance of microparasite virulence. *American Naturalist*, 457–472.
- Barbour, A. D., Pollett, P. K., Dec. 2010. Total variation approximation for quasi-stationary distributions. *Journal of Applied Probability* 47 (4), 934–946, zentralblatt MATH identifier 05835859 , Mathematical Reviews number (MathSciNet) MR2752899.
URL <http://projecteuclid.org/euclid.jap/1294170510>
- Barbour, A. D., Pollett, P. K., Nov. 2012. Total variation approximation for quasi-equilibrium distributions, II. *Stochastic Processes and their Applications* 122 (11), 3740–3756.
URL <http://www.sciencedirect.com/science/article/pii/S030441491200155X>
- Bonhoeffer, S., Nowak, M. A., 1994. Mutation and the evolution of virulence. *Proceedings of the Royal Society of London. Series B: Biological Sciences* 258 (1352), 133–140.
URL <http://rspb.royalsocietypublishing.org/content/258/1352/133.short>
- Boots, M., Hudson, P. J., Sasaki, A., Feb. 2004. Large shifts in pathogen virulence relate to host population structure. *Science* 303 (5659), 842–844, PMID: 14764881.
URL <http://www.sciencemag.org/content/303/5659/842>
- Boots, M., Sasaki, A., Oct. 1999. ‘Small worlds’ and the evolution of virulence: infection occurs locally and at a distance. *Proceedings of the Royal Society of London. Series B: Biological Sciences* 266 (1432), 1933–1938.
URL <http://rspb.royalsocietypublishing.org/content/266/1432/1933.abstract>
- Bull, J. J., Oct. 1994. Perspective: Virulence. *Evolution* 48 (5), 1423.
- Castellano, C., Pastor-Satorras, R., Nov. 2010. Thresholds for epidemic spreading in networks. *Physical Review Letters* 105 (21), 218701.
URL <http://link.aps.org/doi/10.1103/PhysRevLett.105.218701>
- Cator, E., Van Mieghem, P., 2013. Susceptible-infected-susceptible epidemics on the complete graph and the star graph: Exact analysis. *Physical Review E* 87 (1), 012811.
URL <http://pre.aps.org/abstract/PRE/v87/i1/e012811>
- Cortez, M. H., Dec. 2013. When does pathogen evolution maximize the basic reproductive number in well-mixed host–pathogen systems? *Journal of Mathematical Biology* 67 (6-7), 1533–1585.
URL <http://link.springer.com/article/10.1007/s00285-012-0601-2>
- Crow, J., Kimura, M., 1970. *An introduction to population genetics theory*. Burgess Pub. Co.
URL <http://books.google.at/books?id=MLETAQAIAAJ>
- Dieckmann, U., 2002. Adaptive dynamics of pathogen-host interactions. In: Dieckmann, U., Metz, J. A. J., Sabelis, M. W., Sigmund, K. (Eds.), *Adaptive Dynamics of Infectious Diseases: In Pursuit of Virulence Management*. Cambridge University Press, Cambridge, UK, pp. 39–58.
- Dieckmann, U., Metz, J., Sabelis, M., Sigmund, K., 2002. *Adaptive Dynamics of Infectious Diseases: In Pursuit of Virulence Management*. Cambridge Studies in Adaptive Dynamics. Cambridge University Press.
- Dieckmann, O., 2004. A beginners guide to adaptive dynamics. *Banach Center publications* 63, 47–86.
URL <http://igitur-archive.library.uu.nl/math/2006-0727-200409/UUindex.html>
- Dieckmann, O., Heesterbeek, J. A. P., Apr. 2000. *Mathematical Epidemiology of Infectious Diseases: Model Building, Analysis and Interpretation*. John Wiley & Sons.
- Eames, K. T. D., Keeling, M. J., Oct. 2002. Modeling dynamic and network heterogeneities in the spread of sexually transmitted diseases. *Proceedings of the National Academy of Sciences of the United States of America* 99 (20), 13330–13335.
URL <http://www.pnas.org.ezp-prod1.hul.harvard.edu/content/99/20/13330.abstract>
- Ebert, D., Bull, J. J., Jan. 2008. The evolution and expression of virulence. In: Stearns, S. C., Koella, J. C. (Eds.), *Evolution in Health and Disease*, 2nd Edition. Oxford University Press, USA, pp. 153–167.
- Ganesh, A., Massoulié, L., Towsley, D., 2005. The effect of network topology on the spread of epidemics. In: *Proceedings IEEE INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies*. Vol. 2. pp. 1455–1466 vol. 2.

- Ganusov, V. V., Antia, R., Sep. 2003. Trade-offs and the evolution of virulence of microparasites: do details matter? *Theoretical population biology* 64 (2), 211–220, PMID: 12948682.
- Ganusov, V. V., Bergstrom, C. T., Antia, R., Feb. 2002. Within-host population dynamics and the evolution of microparasites in a heterogeneous host population. *Evolution; international journal of organic evolution* 56 (2), 213–223, PMID: 11926490.
- Geritz, S., Kisdi, E., Mesze'NA, G., Metz, J., 1998. Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evolutionary Ecology* 12 (1), 35–57.
- Haraguchi, Y., Sasaki, A., 2000. The evolution of parasite virulence and transmission rate in a spatially structured population. *Journal of Theoretical Biology* 203 (2), 85–96.
- Heffernan, J. M., Smith, R. J., Wahl, L. M., Sep. 2005. Perspectives on the basic reproductive ratio. *Journal of The Royal Society Interface* 2 (4), 281–293, PMID: 16849186.
URL <http://rsif.royalsocietypublishing.org/content/2/4/281>
- Hofbauer, J., Sigmund, K., May 1998. *Evolutionary Games and Population Dynamics*. Cambridge University Press.
- Iwasa, Y., Michor, F., Nowak, M. A., Jan. 2004. Evolutionary dynamics of invasion and escape. *Journal of Theoretical Biology* 226 (2), 205–214.
URL <http://www.sciencedirect.com/science/article/pii/S0022519303003333>
- Keeling, M., Ross, J., 2008. On methods for studying stochastic disease dynamics. *Journal of The Royal Society Interface* 5 (19), 171–181.
URL <http://rsif.royalsocietypublishing.org/content/5/19/171.abstract>
- Kemeny, J., Snell, J., 1976. *Finite Markov Chains: With a New Appendix "Generalization of a Fundamental Matrix"*. Undergraduate Texts in Mathematics. Springer.
- Kendall, D. G., Jan. 1949. Stochastic processes and population growth. *Journal of the Royal Statistical Society. Series B (Methodological)* 11 (2), 230–282.
URL <http://www.jstor.org/stable/2984078>
- Krysicio, R. J., Lefèvre, C., 1989. On the extinction of the S-I-S stochastic logistic epidemic. *Journal of Applied Probability* 26 (4), 685–694.
- Lenski, R. E., May, R. M., Aug. 1994. The evolution of virulence in parasites and pathogens: Reconciliation between two competing hypotheses. *Journal of Theoretical Biology* 169 (3), 253–265.
URL <http://www.sciencedirect.com/science/article/pii/S0022519384711465>
- Levin, B. R., 1996. The evolution and maintenance of virulence in microparasites. *Emerging Infectious Diseases* 2 (2), 93–102, PMID: 8903208 PMID: PMC2639826.
- Lipsitch, M., Siller, S., Nowak, M. A., 1996. The evolution of virulence in pathogens with vertical and horizontal transmission. *Evolution*, 1729–1741.
URL <http://www.jstor.org/stable/10.2307/2410731>
- Mandl, P., 1960. On the asymptotic behaviour of probabilities within groups of states of a homogeneous markov process (in russian). *Časopis pro pěstování matematiky* 85 (4), 448–456.
- May, R. M., Nowak, M. A., 1994. Superinfection, metapopulation dynamics, and the evolution of diversity. *Journal of Theoretical Biology* 170 (1), 95–114.
URL <http://www.cscs.umich.edu/~rlr/CAFI/uploads/Main/mayNovakSuperinfection.pdf>
- May, R. M., Nowak, M. A., Aug. 1995. Coinfection and the evolution of parasite virulence. *Proceedings of the Royal Society B: Biological Sciences* 261 (1361), 209–215, PMID: 7568274.
URL <http://www.ncbi.nlm.nih.gov/pubmed/7568274>
- Mazancourt, C. d., Dieckmann, U., 2004. Trade-off geometries and frequency-dependent selection. *The American Naturalist* 164 (6), pp. 765–778.
URL <http://www.jstor.org/stable/10.1086/424762>
- Messinger, S. M., Ostling, A., Oct. 2009. The consequences of spatial structure for the evolution of pathogen transmission rate and virulence. *The American Naturalist* 174 (4), 441–454.
URL <http://www.jstor.org/stable/27735856>
- Metz, J. A., Mylius, S. D., Dieckmann, O., 1996. When Does Evolution Optimise?: On the Relation Between Types of Density Dependence and Evolutionarily Stable Life History Parameters. IIASA Working Paper WP-96-004. International Institute for Applied Systems Analysis, Laxenburg, Austria.
- Metz, J. A. J., Mylius, S. D., Dieckmann, O., 2008. When does evolution optimize? *Evol. Ecol. Res* 10 (5), 629–654.
URL <http://www.bio.vu.nl/thb/course/ecol/MetzMyli2008.pdf>
- Nasell, I., Jul. 1995. The threshold concept in stochastic epidemic and endemic models. In: Mollison,

- D. (Ed.), *Epidemic Models: Their Structure and Relation to Data*. Cambridge University Press, pp. 71–83.
- Nowak, M. A., May, R. M., 1994. Superinfection and the evolution of parasite virulence. *Proceedings of the Royal Society B: Biological Sciences*, 81–89.
- Nowak, M. A., Sasaki, A., Taylor, C., Fudenberg, D., Apr. 2004. Emergence of cooperation and evolutionary stability in finite populations. *Nature* 428 (6983), 646–650.
URL <http://www.nature.com/nature/journal/v428/n6983/abs/nature02414.html>
- Nowak, M. A., Sigmund, K., Feb. 2004. Evolutionary dynamics of biological games. *Science* 303 (5659), 793–799.
URL <http://www.sciencemag.org/cgi/content/abstract/sci;303/5659/793>
- Nåsell, I., 2001. Extinction and quasi-stationarity in the verhulst logistic model. *Journal of Theoretical Biology* 211 (1), 11 – 27.
URL <http://www.sciencedirect.com/science/article/pii/S0022519301923288>
- Ovaskainen, O., 2001. The quasistationary distribution of the stochastic logistic model. *J. Appl. Probab.* 38 (4), 898–907.
- Proulx, S. R., Day, T., 04 2002. What can invasion analyses tell us about evolution under stochasticity in finite populations? *Selection* 2 (1), 2–15.
URL <http://dx.doi.org/10.1556/Select.2.2001.1-2.2>
- Regoes, R. R., Nowak, M. A., Bonhoeffer, S., Feb. 2000. Evolution of virulence in a heterogeneous host population. *Evolution* 54 (1), 64–71.
URL <http://www.jstor.org/stable/2640474>
- Ross, J. V., Nov. 2011. Invasion of infectious diseases in finite homogeneous populations. *Journal of Theoretical Biology* 289, 83–89.
URL <http://www.sciencedirect.com/science/article/pii/S0022519311004383>
- Smith, J. M., Oct. 1982. *Evolution and the Theory of Games*. Cambridge University Press.
- Van Mieghem, P., Cator, E., Jul. 2012. Epidemics in networks with nodal self-infection and the epidemic threshold. *Physical Review E* 86, 016116.
URL <http://link.aps.org/doi/10.1103/PhysRevE.86.016116>
- Webb, S. D., Keeling, M. J., Boots, M., Mar. 2013. A theoretical study of the role of spatial population structure in the evolution of parasite virulence. *Theoretical population biology* 84, 36–45, PMID: 23274478.
- Weibull, J. W., 1997. *Evolutionary Game Theory*. MIT Press.

Appendix A. Fixation probability for residents and mutants with large basic reproductive ratios

Appendix A.1. Quasi-stationary distribution for large basic reproductive ratio

As a first step towards calculating the fixation probability, we need an expression for the quasi-stationary distribution (2). In the region of large R_1 , it behaves as (Ovaskainen, 2001)

$$Q_{n_1}(R_1, N) = \begin{cases} 1 - \frac{N^2}{(N-1)R_1} + \mathcal{O}\left(\frac{1}{R_1^2}\right) & \text{if } n_1 = N \\ \frac{N^2}{(N-1)R_1} + \mathcal{O}\left(\frac{1}{R_1^2}\right) & \text{if } n_1 = N - 1, \\ \mathcal{O}\left(\frac{1}{R_1^2}\right) & \text{if } n_1 < N - 1 \end{cases} \quad (\text{A.1})$$

Looking at the transition probabilities of the single-disease SIS model (i.e., setting $n_2 = 0$ in Eq.(5)), one can gain an intuitive insight into this result. The condition of large R_1 means that the probability of infecting an uninfected host is much larger than the probability that an infected individual recovers (or dies). Therefore, the population soon reaches a state where all individuals are infected. Every once in a while, an individual is recovered. This happens at a rate Na_1 . However, this newly recovered host is quickly

reinfected with the rate $\beta_1(N-1)/N$. If we observe the system for a long time, the fraction of time it spends in the state $n_1 = N-1$ is

$$\frac{\frac{1}{\frac{\beta_1}{N}(N-1)}}{\frac{1}{Na_1} + \frac{1}{\frac{\beta_1}{N}(N-1)}} \approx \frac{N^2}{(N-1)R_1} \frac{1}{R_1}, \quad (\text{A.2})$$

which, assuming ergodicity, can be identified with the probability of being in the state $n_1 = N-1$, explaining the expression (A.1). Furthermore, this reasoning is valid only if

$$\mathcal{P}(n_1 \rightarrow n_1 + 1) \gg \mathcal{P}(n_1 \rightarrow n_1 - 1) \quad \forall n_1 \in \{1, \dots, N-1\}. \quad (\text{A.3})$$

This is equivalent to $R_1 \gg N$, giving us a more useful description of the region of validity of the approximation (A.1).

Appendix A.2. Derivation of the fixation probability in the $R_i \rightarrow \infty$ limit

Using a similar reasoning as in the case of only one strain present, most of the time every host in the population will be infected. Occasionally, an individual recovers, however, it is quickly reinfected with either one of the two strains. Therefore, in the region of large R_i , the state of the population is constrained to the red dots shown in Figure 1A. More precisely, if we want this to be true, the transition probabilities from states with $n_1 + n_2 = N-1$ must satisfy

$$\begin{aligned} \mathcal{P}([n_1, n_2] \rightarrow [n_1 + 1, n_2]) &\gg \mathcal{P}([n_1, n_2] \rightarrow [n_1 - 1, n_2]), \\ \mathcal{P}([n_1, n_2] \rightarrow [n_1 + 1, n_2]) &\gg \mathcal{P}([n_1, n_2] \rightarrow [n_1, n_2 - 1]), \\ \mathcal{P}([n_1, n_2] \rightarrow [n_1, n_2 + 1]) &\gg \mathcal{P}([n_1, n_2] \rightarrow [n_1 - 1, n_2]), \\ \mathcal{P}([n_1, n_2] \rightarrow [n_1, n_2 + 1]) &\gg \mathcal{P}([n_1, n_2] \rightarrow [n_1, n_2 - 1]), \end{aligned} \quad (\text{A.4})$$

which is equivalent to

$$\begin{aligned} R_1 &\gg N, & R_1 &\gg \frac{a_2}{a_1} N(N-2), \\ R_2 &\gg N, & R_2 &\gg \frac{a_1}{a_2} N(N-2). \end{aligned} \quad (\text{A.5})$$

Note that this is a much more stringent condition than $R_i > N$.

Assuming the conditions (A.5) are satisfied, the original process reduces to a Markov chain on a one-dimensional lattice with $2N-1$ states (the red dots in Figure 1A). We will assign the number $i=0$ to the state $[n_1 = N-1, n_2 = 0]$, $i=1$ to $[n_1 = N-1, n_2 = 1]$, $i=2$ to $[n_1 = N-2, n_2 = 1]$ and so on until $i=2(N-1)$ is assigned to $[n_1 = 0, n_2 = N-1]$. The states $i=0$ and $i=2(N-1)$ are absorbing at $n_2 = 0$ and $n_1 = 0$ respectively. It follows from Eq.(5) that the jump probabilities are

$$\left. \begin{aligned} \mathcal{P}(i \rightarrow i+1) &= \frac{a_1 \left(N - \frac{i+1}{2}\right)}{a_1 \left(N - \frac{i+1}{2}\right) + a_2 \frac{i+1}{2}}, \\ \mathcal{P}(i \rightarrow i-1) &= \frac{a_2 \frac{i+1}{2}}{a_1 \left(N - \frac{i+1}{2}\right) + a_2 \frac{i+1}{2}}, \end{aligned} \right\} \quad \text{if } i \text{ is odd} \quad (\text{A.6})$$

$$\left. \begin{aligned} \mathcal{P}(i \rightarrow i+1) &= \frac{\beta_2 \frac{i}{2}}{\beta_1 \left(N - 1 - \frac{i}{2}\right) + \beta_2 \frac{i}{2}}, \\ \mathcal{P}(i \rightarrow i-1) &= \frac{\beta_1 \left(N - 1 - \frac{i}{2}\right)}{\beta_1 \left(N - 1 - \frac{i}{2}\right) + \beta_2 \frac{i}{2}}. \end{aligned} \right\} \quad \text{if } i \text{ is even} \quad (\text{A.7})$$

In the limit of $R_1 \rightarrow \infty$, the initial condition (3) is non-zero only for $n_1 = N - 1$. The fixation probability is equal to the conditional fixation probability $F_{N-1}(a_2, \beta_2; a_1, \beta_1, N)$, which is the same as the probability of absorption at $i = 2(N - 1)$ starting from the state $i = 1$. Note that the jump probabilities, and therefore also the absorption probabilities, depend only on the ratios a_2/a_1 and β_2/β_1 . Probabilities of absorption from any state satisfy solvable recurrence equations that yield

$$\begin{aligned} F_{N-1}(a_2, \beta_2; a_1, \beta_1, N) &= \frac{1}{1 + \sum_{k=1}^{2N-3} \prod_{i=1}^k \frac{\mathcal{P}(i \rightarrow i-1)}{\mathcal{P}(i \rightarrow i+1)}} \\ &= \frac{(N-1) \left(\frac{\beta_1 a_2}{\beta_2 a_1} - 1 \right)^2}{N-1 + \frac{a_2}{a_1} - N \frac{\beta_1 a_2}{\beta_2 a_1} + \left(\frac{\beta_1 a_2}{\beta_2 a_1} \right)^N \left(1 - N \frac{\beta_2}{\beta_1} + \frac{a_2}{a_1} (N-1) \right)}. \end{aligned} \quad (\text{A.8})$$

Calculating the conditional fixation probability starting at $n_1 = N - 2$, $n_2 = 1$,

$$F_{N-2}(a_2, \beta_2; a_1, \beta_1, N) = \left(1 + \frac{\mathcal{P}(1 \rightarrow 0)}{\mathcal{P}(1 \rightarrow 2)} \right) F_{N-1}(a_2, \beta_2; a_1, \beta_1, N), \quad (\text{A.9})$$

one could also take into account the $1/R_1$ corrections in Eq.(A.1). Using Eq.(4), we find

$$F(a_1, \beta_1; a_2, \beta_2, N) = \frac{\left(N-1 + \frac{N^2}{N-1} \frac{a_2}{a_1} \frac{1}{R_1} + \mathcal{O}\left(\frac{1}{R_i}\right) \right) \left(\frac{\beta_1 a_2}{\beta_2 a_1} - 1 \right)^2}{N-1 + \frac{a_2}{a_1} - N \frac{\beta_1 a_2}{\beta_2 a_1} + \left(\frac{\beta_1 a_2}{\beta_2 a_1} \right)^N \left(1 - N \frac{\beta_2}{\beta_1} + \frac{a_2}{a_1} (N-1) \right)}. \quad (\text{A.10})$$

In addition to higher order terms, the $\mathcal{O}\left(\frac{1}{R_i}\right)$ correction also contains $1/R_i$ terms that come from a possible jump to the states on the $n_1 + n_2 = N - 2$ diagonal in the state space. We do not evaluate these corrections since, for our purposes, the most interesting case is when $R_i \rightarrow \infty$. Then, the fixation probability is simply

$$\lim_{R_i \rightarrow \infty} F(a_2, \beta_2; a_1, \beta_1, N) = \frac{(N-1) \left(\frac{\beta_1 a_2}{\beta_2 a_1} - 1 \right)^2}{N-1 + \frac{a_2}{a_1} - N \frac{\beta_1 a_2}{\beta_2 a_1} + \left(\frac{\beta_1 a_2}{\beta_2 a_1} \right)^N \left(1 - N \frac{\beta_2}{\beta_1} + \frac{a_2}{a_1} (N-1) \right)}. \quad (\text{A.11})$$