

A general approach for population games with application to vaccination

Timothy C. Reluga¹ and Alison P. Galvani²

1. Corresponding Author
Department of Mathematics
Pennsylvania State University
State College, PA 16802
timothy@reluga.org
2. Department of Epidemiology and Public Health
Yale University School of Medicine
New Haven, CT 06520
alison.galvani@yale.edu

April 19, 2010

Abstract

One of the major challenges in the design and implementation of health policies is reconciling the interests of individuals with the interests of the community. In this paper, we present a technique based on a combination of mechanistic population-scale models, Markov decision process theory, and game theory that allows for the valuation of decisions at both individual and community scales. We demonstrate how nonlinear population models can be combined with game theory to inform basic questions concerning the management of infectious diseases. As a demonstration of the technique, we provide solutions to several variants of the simple vaccination game including differential waning rates or imperfect vaccine efficacy. In addition, we show how path-integral approaches can be applied in the study models where strategies are chosen as fixed waiting times rather than by randomized distributions. Our method can be widely applied to decision problems dependent on population dynamics.

keywords: population games, epidemiology, vaccination, infectious diseases, Markov decision processes

1 Introduction

There has been a long-standing dichotomy in our approaches to the control of disease. Epidemiology and public health, have focused on the best interests of communities and populations, implicitly assuming that the interests of individuals parallel those of their community. Medicine, on the other hand, has focused on the well-being of the individual without specific attention to the community. For the most part, we perceive this dichotomy as natural and justified; health impact from the two spheres can be complementary or independent. For example, improvements in sanitation are good for a community because they reduce the risks of disease for everyone in the community. Similarly, curing a patient of an infectious disease is good for both the patient and the community because a cured patient no longer poses a health risk to others.

One consequence of this dichotomy is that analytical methods in these disciplines have developed in independent directions; quantitative analyses often focus exclusively on either aggregate or individual behaviors. Epidemiology models focus on aggregate behaviors to help us understand the transmission and control of epidemic and endemic infectious disease [1, 2]. In these models, individuals are typically treated as particles, obeying prescribed laws rather than as actors with free-will. In contrast, medical models have favored cost-benefit analysis and decision theory approaches that focus on determining the best behaviors for individuals. Such methodologies assist the patient and the doctor in weighing the risks and benefits for various treatment modalities. For instance, the Markov decision process theory (MDPT) uses a stochastic process to describe the possible events in an individual's future [3]. MDPT associates values with events, and calculates the expected value conditional on each choice. The expected value can then be optimized over the different choices available to the individual. Although the risks an individual faces may depend on the overall state of the population, this form of analysis does not account for any costs or risks posed to the population by the individual's choices.

Modern public health and medicine have pushed back the tide of many prominent diseases in wealthy nations. However, as we have pushed back the tide of old problems, we have unveiled new problems. For instance, concerns about vaccine side-effects often become a priority following successful mass-immunization campaigns and drug resistance management gains practical attention only after the drug's efficacy has been widely demonstrated. In some of these new problems, the interests of the individual now diverge from the interests of the community ; the medical and epidemiological modeling plans can give conflicting answers. Such conflicts are known as "market failures" in economics, and "policy resistance" in public health [4].

The maintenance of herd immunity to infectious diseases through mass vaccination is a good example. In the past, the high prevalence of many vaccine-preventable infectious diseases made the choice of whether-or-not to vaccinate a simple one. When the risks from infection were significant and there was a cheap and safe vaccine available, vaccination was a worth-while measure both

for the community and for the individuals. Now, however, the dramatic success of vaccines has greatly reduced the risk of infection for many diseases, and individuals find the decision less clear-cut even though community benefits remain large.

New mathematical approaches are needed to study these conflicts between individuals and communities. One valuable approach is rational-expectation population game theory: if we understand the epidemiological processes, and assume that individuals have broad and reliable information about the state of the world and themselves, we may be able to generate a function based on the epidemiology that predicts the value gained by each individual depending on their own decisions and the decisions of others. This approach does not capture the full complexity of the decision process, but it gives us a useful point of reference. The first serious attempt to using game theory in studying vaccination was motivated by concerns over the safety of pertussis vaccines [5]. This was followed by studies of the economics of compulsory vaccination programs [6] and disease eradication efforts [7, 8]. A simplified approach to the game theory of vaccination was introduced by Bauch et al. [9]. The authors describe individual risk which is not strictly needed when only considering population-scale dynamics. Generalization of this observation leads to the class of models we study here. A growing number of papers expand on these ideas, and investigating how various factors can favorably or unfavorably influence vaccination practices [10, 11, 12, 13, 14, 15, 16, 17].

In this paper, we describe in detail the mathematics of a population-game approach to address problems in health-care. Our approach has been used in some of our previous research [18, 19], but without the detailed exposition or generality appropriate to facilitate further applications. We demonstrate here how population-scale dynamic models can be coupled to individual scale decision-models. This approach allows us to formulate population games that describe the utility of decisions to the individual and to the population through the use of novel closed-form algebraic results. We also review standard conditions for equilibria in population games. In addition, we illustrate the approach using simple vaccination problems. Our approach is simple to use and often completely solvable. Where other methods often require approximations or numerical solution, our method frequently yields exact analytic formulas that can be easily interpreted in the context of biological and economic parameters. As proof-of-principle, we derive analytic solutions to two versions of the vaccination game. Our first version is a compartmental model in which individuals' strategies are fully randomized. Our second version uses McKendrick–von Foerster–like partial differential equations to describe a vaccination game where strategies are based on fixed-time delays. Despite differences in the mathematics, both lead to similar results. In basic scenarios, there is a unique Nash equilibrium for behavior with global invasion potential that converges to the community's preferred policy for very inexpensive and very expensive vaccines. But consideration of generalizations of the compartmental model shows that biological complications like limited vaccine efficacy differences in immunity waning can lead to more complex equilibrium structures.

2 Methods

Our approach has three parts: 1) a model of changes in an individual state analyzed using stochastic process theory, 2) a model of the population's state analyzed using dynamical systems theory, and 3) a model of the value of choices analyzed using decision theory and game theory.

2.1 General formulation

A population-game model contains both population-scale components that describe the states of the populations under study and individual-scale components that describe the states of individuals within these populations. Suppose we have a single population, and individuals within this population occupy one of many possible states. The state of an individual in the population at time t is represented by a probability density $p(t)$ over the possible states. We refer to $p(t)$ as the individual's state.

At the scale of individuals, dynamics are fundamentally stochastic— we can use observational data to infer the likelihoods of certain events occurring at specific points in an individual's life, but no particular individual's future can be predicted with certainty. The extent of uncertainty in future events is important in decision processes, so we adopt a Markov-process framework. Suppose that the state of an individual \mathbf{p} changes according to a continuous-time Markov process describing changes in the probability that an individual will be found in any given state using a linear differential equation. The transition rates in this Markov process may depend on many things including the population's state \mathbf{x} , the environment's state \mathbf{e} , time t , and a vector representing the individual's decisions/strategy $\boldsymbol{\pi}$. Therefore,

$$\frac{d\mathbf{p}}{dt} = \mathbf{Q}(\mathbf{x}, \mathbf{e}, t, \boldsymbol{\pi})\mathbf{p}, \quad (2.1)$$

where \mathbf{Q} is a matrix of transition rates.

In order to use Eq. (2.1), we must also specify a description of the population's dynamics and the environment's dynamics. The population-scale and individual-scale dynamics are driven by separate but related processes. There are several possible ways to describe the dynamics of the population's state. Suppose the population exists in an environment with state vector $\mathbf{e}(t)$. One of the simplest and most convenient approaches is to postulate that the population-scale dynamics satisfy a system of differential equations where the rates of change in the states of the population depend on the current states \mathbf{x} , the environmental state \mathbf{e} , the average behavior of individuals in the population $\bar{\boldsymbol{\pi}}$, and the time;

$$\frac{d\mathbf{x}}{dt} = G_{\mathbf{x}}(\mathbf{x}, \bar{\boldsymbol{\pi}}, \mathbf{e}, t). \quad (2.2)$$

The rate of change in the environment's state given by:

$$\frac{d\mathbf{e}}{dt} = G_{\mathbf{e}}(\mathbf{x}, \bar{\boldsymbol{\pi}}, \mathbf{e}, t). \quad (2.3)$$

also depends on the state of the population, the resident strategies, the environment's current state, and time. The vector $\bar{\boldsymbol{\pi}}$ representing the effective average behavior of individuals is often referred to as the “resident” strategy, as it is most easily specified in scenarios where almost all resident individuals in the population are using exactly this strategy $\bar{\boldsymbol{\pi}}$. We assume the populations are sufficiently large that the population dynamics are not significantly affected by the few individuals with “invading” strategies that differ from the resident strategy.

We develop a framework to investigate how active decision-making on the part of individuals can alter dynamics at the population-scale, and hopefully compare the quality of the outcomes. One possibility is that individuals choose their behaviors $\boldsymbol{\pi}$ with the goal of optimizing a known function, called a utility function $U(\boldsymbol{\pi})$. Suppose we know an individual's entire future, in the sense that there is a function $\rho(t)$ that, given an individual's strategy $\boldsymbol{\pi}$, returns an individual's future state at time t . Then the utility of the strategy $\boldsymbol{\pi}$ is the sum of all gains attributable to that strategy over the course of their life –

$$\int_0^{\infty} e^{-ht} \phi(\rho(t), \rho'(t), \boldsymbol{\pi}) dt, \quad (2.4)$$

where $\phi(\rho, \rho')$ is the rate of utility gain per unit time as a function of the state ρ and the change in state ρ' . But as we commented above, the future of each person is uncertain; $\rho(t)$ is an unknown random variable and the utility is a random variable that depends on the distribution of different ρ 's.

There is no universally accepted answer as to how a rational individual should interpret the distribution of their potential utility when choosing their best strategies [20]. Individuals might maximize their worst-case payoff or minimize their risk given some constraints on acceptable payoffs. For our purposes, we will use the common assumption that individuals maximize their “expected” utility. We can use the governing stochastic process to construct a probability measure on the space of all possible state paths $\rho(t)$ and calculate an expected utility by integrating over all possible state paths. The expected utility conditional on the individuals' state $\rho(t_0)$ at time t_0 is

$$\int_{\rho(t_0)} \int_0^{\infty} e^{-ht} \phi(\rho(t), \rho'(t), \boldsymbol{\pi}) dt \mathcal{D}\rho(\boldsymbol{\pi}), \quad (2.5)$$

where $\mathcal{D}\rho$ denotes the infinitesimal probability of each possible state path starting at $\rho(t_0)$, depending on the strategy $\boldsymbol{\pi}$. If individual's state-paths are governed by Eq. (2.1), \mathbf{f} is a vector of gains per unit time for residence of each state, \mathbf{F} is a matrix of instantaneous utility gains associated with each transition, and \circ

represents the Hadamard product of two matrices ¹ then:

$$U(\boldsymbol{\pi}, \bar{\boldsymbol{\pi}}) = \int_{t_0}^{\infty} e^{-h(t-t_0)} [\boldsymbol{f}^T + \mathbf{1}^T (\mathbf{F} \circ \mathbf{Q})] \boldsymbol{p}(t) dt \quad (2.6)$$

where h is the discount rate and t_0 is the initial time. We write the expected utility as a function of the resident strategy $\bar{\boldsymbol{\pi}}$ and an invading strategy $\boldsymbol{\pi}$ to emphasize that it is specifying a game between competing strategies.

Equation (2.6) is equivalent to a path integral over all an individual's possible state-paths and can be derived from the Markov decision process theory developed by Howard [21]. Let the vector $\boldsymbol{u}(t)$ represent the expected present value of each state at time t relative to the initial time t_0 . Generalizing the Bellman equation for a continuous-time Markov process to the case of non-autonomous processes, the total expected utility at the initial time t_0 is

$$U = \boldsymbol{u}^T(t_0) \boldsymbol{p}(t_0), \quad (2.7)$$

where $\boldsymbol{p}(t_0)$ is the probability measure over initial states, and

$$\frac{d\boldsymbol{u}^T}{dt} = \boldsymbol{u}^T (h\mathbf{I} - \mathbf{Q}) - [\boldsymbol{f}^T + \mathbf{1}^T (\mathbf{F} \circ \mathbf{Q})]. \quad (2.8)$$

Eq. (2.8) is solved backwards from ∞ to initial time t_0 rather than forwards from t_0 to ∞ , because the expected utility today is calculated from the expected utility tomorrow, not from the expected utility yesterday. Integration of Eq. (2.8) with terminal condition $\boldsymbol{u}(\infty) = \mathbf{0}$ yields Eq. (2.6). Readers familiar with optimal-control theory may alternatively recognize \boldsymbol{u} as the vector of present-value shadow prices.

Although population games can be solved and studied in the form we have so far presented, it is often useful to focus on cases where the population-scale dynamics have a simple attractor. If all dynamics are autonomous and $(\boldsymbol{x}^*, \boldsymbol{e}^*)$ is a stationary solution of the population-scale dynamics, the transition rate matrix

$$\mathbf{Q}^* = \mathbf{Q}(\boldsymbol{x}^*, \boldsymbol{e}^*, \boldsymbol{\pi}) \quad (2.9)$$

is also stationary. If there is a positive discount rate $h > 0$ or the Markov process is transient e.g. the dominate eigenvalue $\lambda_0(\mathbf{Q}^*) < 0$, then as the time horizon of the utility calculation becomes infinitely long, \boldsymbol{u}^T will approach steady state. At steady state, the expected utility has closed form, such that:

$$U = [\boldsymbol{f}^T + \mathbf{1}^T (\mathbf{F} \circ \mathbf{Q}^*)] (h\mathbf{I} - \mathbf{Q}^*)^{-1} \boldsymbol{p}(t_0). \quad (2.10)$$

However in some special cases where the Markov process is ergodic (e.g. the largest eigenvalue of \mathbf{Q}^* is 0) and there is no discounting ($h = 0$), Eq. (2.8) will never

¹The components of the Hadamard product are the product of the components of \mathbf{F} and \mathbf{Q} , $(\mathbf{F} \circ \mathbf{Q})_{jk} = \mathbf{F}_{jk} \mathbf{Q}_{jk}$.

converge to a steady state and Eq. (2.10) cannot be calculated, because the matrix inversion is singular. In these cases, it is more appropriate to study the asymptotic growth rate of the expected utility. We can show that if $\hat{\mathbf{p}}$ is the equilibrium distribution of the Markov process ($\mathbf{Q}^* \hat{\mathbf{p}} = 0$ and $\sum_j \hat{p}_j = 1$), then

$$\lim_{t_f \rightarrow \infty} \frac{dU}{dt} = [\mathbf{f}^T + \mathbf{1}^T (\mathbf{F} \circ \mathbf{Q}^*)] \hat{\mathbf{p}}. \quad (2.11)$$

This is particularly useful in theoretical work where closed-form solutions are desirable.

There are naturally occurring problems where the attractor of the population-scale dynamics is periodic, leading to cases where the transition rate matrix \mathbf{Q}^* is also periodic. In elementary cases where $h > 0$ or $\lambda_0(\mathbf{Q}^*(t)) < 0$ for all t , calculation of the expected utility over long horizons under periodic oscillations can then be reduced to the solution of Eq. (2.8) with the periodic boundary condition:

$$\mathbf{u}(t_0) = \mathbf{u}(t_0 + t_p) e^{ht_p}, \quad (2.12)$$

where t_p is the period of oscillation. It is not clear what approximation methods are appropriate in cases where the macroscopic dynamics have quasiperiodic or chaotic attractors.

2.2 Population game equilibria

The game-theoretic aspect of our analysis focuses on the properties of the expected utility $U(\boldsymbol{\pi}, \bar{\boldsymbol{\pi}})$ as a function of the individual's behavior $\boldsymbol{\pi}$ and the population's overall average behavior $\bar{\boldsymbol{\pi}}$.

Strategies used in a population game are commonly classified in terms of their ability to replace and to be replaced by neighboring strategies [22, 23]. A strategy π^* is a global Nash equilibrium if no other alternative strategy has a higher utility than the strategy itself when invading a population using the strategy π^* ; i.e.

$$\forall \pi, \quad U(\pi, \pi^*) \leq U(\pi^*, \pi^*). \quad (2.13)$$

A strategy π^* is a strict Nash equilibrium if every alternative strategy has less utility than the strategy π^* ; i.e.

$$\forall \pi \neq \pi^*, \quad U(\pi, \pi^*) < U(\pi^*, \pi^*). \quad (2.14)$$

Equally important is whether a strategy can successfully invade populations with differing resident strategies. A strategy π^* has global invasion potential if it never does worse than any resident strategy played against itself; i.e.

$$\forall \bar{\boldsymbol{\pi}}, \quad U(\pi^*, \bar{\boldsymbol{\pi}}) \geq U(\bar{\boldsymbol{\pi}}, \bar{\boldsymbol{\pi}}). \quad (2.15)$$

A strategy π^* has strict global invasion potential if equality only holds when the strategy π^* is itself the resident strategy, i.e.

$$\forall \bar{\pi} \neq \pi^*, \quad U(\pi^*, \bar{\pi}) > U(\bar{\pi}, \bar{\pi}). \quad (2.16)$$

Those Nash equilibria with invasion potential where at least one of the conditions is strict are the evolutionarily stable strategies (ESS's) of the population game [24, 25]. **[TCR: inconsistent with more population usage?]**

The definitions above place no restrictions on any set of strategies or on the utility function being studied. They can be applied in general. However, they do not supply an efficient approach for identifying equilibria. In the simple applications presented here, the strategy space is composed of real numbers, and the utility function is continuous and differentiable almost everywhere. Therefore, it is convenient to define local forms of the Nash and invasion conditions that can be tested for using differential calculus and Taylor series. In a single population with a 1-parameter strategy space, a strategy π^* is a critical strategy of a utility function $U(\pi, \bar{\pi})$ if the following equation holds

$$\frac{dU}{d\pi}(\pi^*, \pi^*) = 0. \quad (2.17)$$

A critical strategy is a strict local Nash equilibrium if

$$\frac{d^2U}{d\pi^2}(\pi^*, \pi^*) < 0. \quad (2.18)$$

A critical strategy has strict local invasion potential if

$$\frac{d^2U}{d\pi^2}(\pi^*, \pi^*) + 2\frac{d^2U}{d\pi d\bar{\pi}}(\pi^*, \pi^*) < 0. \quad (2.19)$$

If we define a normalized utility function of the form

$$\tilde{U}(\pi, \bar{\pi}) = \frac{U(\pi, \bar{\pi})}{U(\bar{\pi}, \bar{\pi})}, \quad (2.20)$$

the local invasion potential condition may be expressed in the simplified form

$$\frac{d^2\tilde{U}}{d\bar{\pi}^2}(\pi^*, \pi^*) > 0. \quad (2.21)$$

Note that Eq. (2.19) differs from the more-common convergent-stability condition [?]

$$\frac{d^2U}{d\pi^2}(\pi^*, \pi^*) + \frac{d^2U}{d\pi d\bar{\pi}}(\pi^*, \pi^*) < 0, \quad (2.22)$$

which implies that small improvements in strategy will lead to the critical point. A local Nash equilibrium with local invasion potential is a local ESS, and these conditions imply that the equilibrium is convergently stable. However, a critical point need not have local invasion potential to be convergently stable.

These differential conditions can be generalized to vector strategies, but the remainder of this paper is confined to scalar strategies. For a discussion of generalizations to multiple types of actors, see Reluga [26].

3 Vaccination Games

We will now focus on the application of these mathematical methods to the analysis of vaccination behavior in a couple of simple epidemic models. The focus here is on the mathematical solutions, as an illustration of the population-game framework in action as applied to a simple problem of practical importance.

One of the biggest questions in the formulation of population games is how to model the strategy-space of individual decisions. In general, individuals can choose mixed strategies with arbitrary time dependence conditional on any subset of the system state information. The complexity of such strategy space models is usually unsolvable and may mask general properties of the games in questions. We present here examples which allow for strategies that can be represented in terms of a single parameter. In the first model, we assume that individuals chose the time they wait before vaccination from an exponential distribution. We refer to this as a “hazard-strategy” because it is equivalent to individuals having a constant hazard per unit time of vaccination. In the second model, individuals choose a fixed time to wait between the loss of immunity and their next vaccination. We refer to this as a “delay-strategy” model because the resulting mathematics is closely related to that of delayed differential equations. The population-game framework described above can be employed for both of these models. Although the solution-techniques have technical differences, the final results are very similar.

3.1 Hazard-strategies in a Vaccination Game

Consider the homogeneous population of size N afflicted by a flu-like illness for which there is a costly and risky vaccine available. We describe this with a compartment model where individuals in the population can occupy susceptible (S), infected (I), resistant (R), and vaccinated (V) states. Individuals die and are replaced at rate μ , recover from infection at rate γ , and lose immunity at rate a . The population-scale dynamics are described by

$$\frac{dS}{dt} = \mu N + a(R + V) - \lambda(I)S - (\bar{\pi} + \mu)S, \quad (3.1a)$$

$$\frac{dI}{dt} = \lambda(I)S - (\gamma + \mu)I, \quad (3.1b)$$

$$\frac{dR}{dt} = \gamma I - (a + \mu)R, \quad (3.1c)$$

$$\frac{dV}{dt} = \bar{\pi}S - (a + \mu)V, \quad (3.1d)$$

where $\lambda(I)$ is the force of infection, $\bar{\pi}$ is the population’s resident vaccination rate, and the population size

$$N = S + I + R + V \quad (3.1e)$$

is constant. The force of infection $\lambda(I)$ is assumed to be an increasing concave function of the number of infections I . Initially, we assume that infection and

vaccination immunity are complete and wane at equal rates. This model is sufficiently general to include special cases of life-long immunity with demographic turnover ($a = 0, \mu > 0$) and temporary immunity without demographic turnover ($a > 0, \mu = 0$).

Although the population's state may evolve deterministically, each individual's state evolves according to the Markov process described by Eq. (2.1). At steady state, the stationary forward transition rate matrix is

$$\mathbf{Q}^* = \begin{bmatrix} -\lambda(I^*) - \pi - \mu & 0 & a & a \\ \lambda(I^*) & -\gamma - \mu & 0 & 0 \\ 0 & \gamma & -a - \mu & 0 \\ \pi & 0 & 0 & -a - \mu \end{bmatrix}, \quad (3.2)$$

where π is the individual's chosen strategy of a daily probability whether-or-not to get vaccinated. Every individual entering the population enters in the susceptible state, so

$$\mathbf{p}(0) = [1 \ 0 \ 0 \ 0]^T. \quad (3.3)$$

When we amalgamate the states of all individuals into one process, we get the mean-field equations described by System (3.1).

The parameter we are interested in is the typical vaccination rate of the population $\bar{\pi}$, and its relation to the individual's hazard of vaccinating per unit time, π . Individuals incur costs from both infection and vaccination. We assume that the individual aims to minimize the total cost they incur given the population dynamics. The vaccination cost to the individual occurs in the transitions between states. The infection cost accumulates as individuals reside in the infected state. As these are the only two costs, and there are no explicit gains,

$$\mathbf{F} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ -c_V & 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{f} = [0 \ -c_I \ 0 \ 0]. \quad (3.4)$$

The expected change in utility per unit time for residence in each state is given by the vector

$$\mathbf{f} + \mathbf{1}^T (\mathbf{F} \circ \mathbf{Q}^*) = [-\pi c_V \ -c_I \ 0 \ 0] \quad (3.5)$$

It follows from Eq. (2.10) that the utility of strategy π to an individual in a population at equilibrium with strategy $\bar{\pi}$ is

$$U(\pi, \bar{\pi}) = [\mathbf{f} + \mathbf{1}^T (\mathbf{F} \circ \mathbf{Q}^*)] (h\mathbf{I} - \mathbf{Q}^*)^{-1} \mathbf{p}(0) \quad (3.6a)$$

$$= \frac{-(a + \mu + h) [c_V \pi (\gamma + \mu + h) + c_I \lambda(I^*)]}{(\mu + h) [\lambda(I^*) (\gamma + a + \mu + h) + (\gamma + \mu + h) (\pi + a + \mu + h)]}, \quad (3.6b)$$

where I^* is determined by the steady-state condition

$$\lambda(I^*) = \frac{I^*/N (\gamma + \mu) (\bar{\pi} + a + \mu)}{(a + \mu) - I^*/N (\gamma + a + \mu)}. \quad (3.7)$$

Since $\lambda(I)$ is increasing and concave, $I^*(\bar{\pi})$ must be a decreasing function of the resident strategy $\bar{\pi}$. If the individual never gets vaccinated,

$$U(0, \bar{\pi}) = \frac{-(a + \mu + h) c_I \lambda(I^*)}{(\mu + h) [\lambda(I^*) (\gamma + a + \mu + h) + (\gamma + \mu + h) (a + \mu + h)]}. \quad (3.8)$$

In the limit of infinitely fast vaccination,

$$U(\infty, \bar{\pi}) = \frac{-(a + \mu + h) c_V}{(\mu + h)}. \quad (3.9)$$

Differentiating the utility with respect to the individuals behavior generates:

$$\begin{aligned} \frac{\partial U}{\partial \pi} &= \frac{-(\gamma + \mu + h)(a + \mu + h)}{(\mu + h)} \\ &\times \frac{\{[(\gamma + \mu + h)(a + \mu + h) + \lambda(I^*)(\gamma + a + \mu + h)] c_V - \lambda(I^*) c_I\}}{[\lambda(I^*)(\gamma + a + \mu + h) + (\gamma + \mu + h)(\pi + a + \mu + h)]^2}. \end{aligned} \quad (3.10)$$

The numerator is independent of π . In this model, and in many other simple models, the utility is independent of the individual's strategy at a critical point. When this is true, critical strategies can be determined by setting utilities at the extremes of individual strategies equal to each other. Here, we find that critical strategies satisfy the condition

$$U(0, \pi^*) = U(\pi, \pi^*) = U(\infty, \pi^*), \quad (3.11)$$

so at the critical strategy,

$$c_V = \frac{c_I}{a + \gamma + \mu + h + \frac{(\gamma + \mu + h)(a + \mu + h)}{\lambda(I^*)}}. \quad (3.12)$$

In some simple models, boundary conditions like $U(0, \pi^*) = U(\infty, \pi^*)$ can be used to identify equilibria, but they can not be used in general without additional conditions ensuring that the utility is also flat in between the boundaries.

Eq. (3.12) says that the critical point, the expected cost per unit time of vaccinating as often and quickly as possible, is equal to the expected cost per unit time of infection if one never receives a vaccination. Solving for the force of infection generates:

$$\lambda(I^*(\bar{\pi} = \pi^*)) = \frac{(\gamma + \mu + h) (a + \mu + h)}{\frac{c_I}{c_V} - (\gamma + a + \mu + h)} \quad (3.13)$$

This condition may not have a solution, depending on the costs of infection and vaccination. A necessary condition for the critical strategy to be a positive vaccination rate is $c_I/c_V - (\gamma + a + \mu + h) > 0$. If the cost of vaccination is so large that even when there is no vaccination ($\bar{\pi} = 0$)

$$c_V > \frac{c_I}{a + \gamma + \mu + h + \frac{(\gamma + \mu + h)(a + \mu + h)}{\lambda(I^*(\bar{\pi}=0))}} \quad (3.14)$$

then the critical strategy is for nobody to get vaccinated, $\pi^* = 0$. If the cost of vaccination c_V is very small relative to the minimal force of infection ($\bar{\pi} = \infty$),

$$c_V < \frac{c_I}{a + \gamma + \mu + h + \frac{(\gamma + \mu + h)(a + \mu + h)}{\lambda(I^*(\bar{\pi}=\infty))}} \quad (3.15)$$

then the critical strategy is for every individual to get vaccinated as quickly as possible, $\pi^* = \infty$. In each scenario, the respective extreme point is a Nash equilibrium because no alternative strategy generates greater utility for the individual.

In between, $I^*(\bar{\pi})$ is uniquely determined by Eq. (3.7) because the force of infection is a continuously increasing function of the incidence, implying that there is a unique critical strategy. The critical strategy satisfies the condition for a Nash equilibrium, because

$$\frac{d^2U}{d\pi^2}(\pi, \pi^*) = 0 \quad (3.16)$$

for all π . But the condition for a strict Nash equilibrium is not satisfied. We can now show invasion potential by proving

$$\frac{d^2U(\pi, \bar{\pi})}{d\pi d\bar{\pi}} = \left(\frac{d}{d\pi} \frac{\partial U}{\partial \lambda} \right) \times \frac{\partial \lambda}{\partial I^*} \frac{dI^*}{d\bar{\pi}} + \frac{\partial U}{\partial \lambda} \times \frac{d}{d\pi} \left(\frac{\partial \lambda}{\partial I^*} \frac{dI^*}{d\bar{\pi}} \right) < 0. \quad (3.17)$$

Both λ and I^* are independent of the individual strategy π , so their derivative with respect to the individual strategy is 0. From our assumptions about the force of infection, we also know $\frac{\partial \lambda}{\partial I^*} > 0$ and $\frac{dI^*}{d\bar{\pi}} < 0$. Finally,

$$\frac{d}{d\pi} \frac{\partial U}{\partial \lambda} = \frac{c_V (h + \mu + a) \left[\frac{c_I}{c_V} - (\gamma + a + \mu + h) \right]^3}{(h + \mu + \gamma)(h + \mu) \left[\pi^* (\mu + a + h + \gamma) - \frac{c_I}{c_V} (a + h + \pi + \mu) \right]^2} > 0 \quad (3.18)$$

provided that

$$\frac{c_I}{\gamma + a + \mu + h} > c_V. \quad (3.19)$$

Thus,

$$\frac{d^2U(\pi, \bar{\pi})}{d\pi d\bar{\pi}} < 0, \quad (3.20)$$

so Eq (2.19) holds. The critical point is a Nash equilibrium with invasion potential.

Once equilibria have been identified, we can evaluate their sensitivity to the parameters. The larger the cost of infection per unit time c_I , the lower the force of infection allowed at the Nash equilibrium. The larger the cost of vaccination, the higher the force of infection allowed at Nash equilibrium. By inspection of Eq. (3.6), we see that as the discount rate h increases, the effective cost of infection decreases as vaccination decreases. A higher force of infection is allowed at equilibrium. The effects of the other parameters are less obvious because the endemic disease prevalence I^* is not only a function of the average vaccination rate $\bar{\pi}$, but also an implicit function of the recovery, waning, and death parameters γ , a , and μ .

3.1.1 Exact Solutions for Standard Incidence

To get a more complete understanding of the relationship between parameter values and the Nash equilibrium vaccination rate, it is helpful to consider a specific model. Suppose the force of infection is modeled using standard incidence

$$\lambda(I) = \beta I/N, \quad (3.21)$$

while demographic turnover and discounting are sufficiently slow so that they can be neglected ($\mu = 0$ and $h = 0$).

Then the reproductive number at the disease-free equilibrium is

$$\mathcal{R} = \left(\frac{\beta}{\gamma}\right) \left(\frac{a}{a + \bar{\pi}}\right). \quad (3.22)$$

If $\mathcal{R} < 1$, the disease-free equilibrium is stable. An average vaccination rate of

$$\bar{\pi} \geq \hat{\pi} = a \left(\frac{\beta}{\gamma} - 1\right) \quad (3.23)$$

is sufficient to achieve herd immunity and thus to ensure eradication.

The endemic-disease equilibrium solution when $\mathcal{R} > 1$ is

$$S^* = N\frac{\gamma}{\beta}, I^* = N\frac{\beta a - \gamma a - \bar{\pi}\gamma}{\beta(a + \gamma)}, R^* = N\frac{\gamma}{a}\frac{\beta a - \gamma a - \bar{\pi}\gamma}{\beta(a + \gamma)}, V^* = N\frac{\bar{\pi}\gamma}{\beta a}. \quad (3.24)$$

In the absence of discounting or mortality, we apply Eq. (2.8), and find the asymptotic expected utility gain per unit time

$$U' = \lim_{t_f \rightarrow \infty} \frac{dU}{dt} = \frac{-a(c_I \beta I^*/N + c_V \pi \gamma)}{(a + \gamma)\beta I^*/N + \gamma(\pi + a)} \quad (3.25)$$

Using equilibrium infection prevalence I^* ,

$$U'(\pi, \bar{\pi}) = \frac{a[(\beta a - \gamma a - \bar{\pi}\gamma)c_I + \gamma\pi(a + \gamma)c_V]}{(a + \gamma)(\bar{\pi}\gamma - \pi\gamma - \beta a)} \quad (3.26)$$

The same of U' is illustrated in Figures 1 and 2.




Figure 1: Asymptotic expected utility gain per unit time $U'(\bar{\pi}, \bar{\pi})$ for an individual using the population's average strategy $\bar{\pi}$. There is a corner in the utility gain at vaccination rates of about $\bar{\pi} = .25$ just sufficient to ensure eradications. This is the social utopia, as faster or slower rates yield slower gains to the average individual. Parameter values: $\gamma = 1$, $\beta = 6$, $a = .05$, $c_I = 4$, $c_V = 2$.

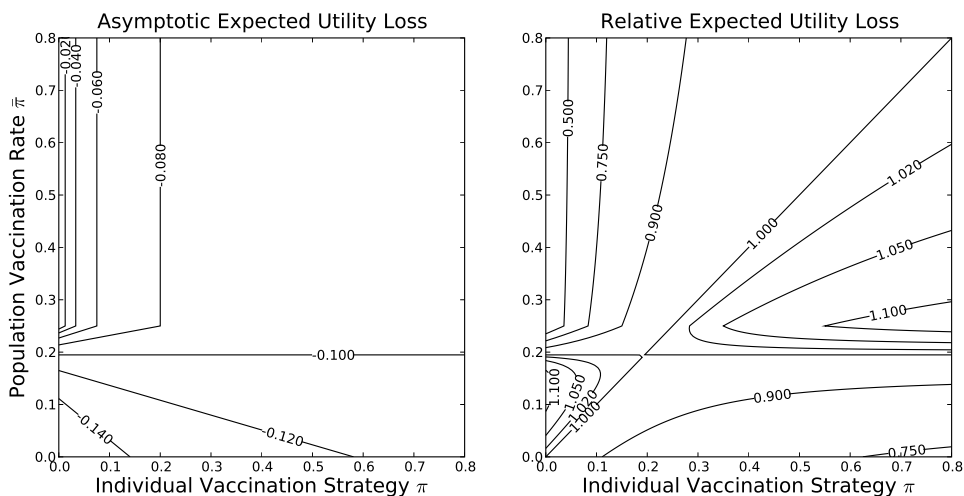


Figure 2: An individual's asymptotic expected utility loss $U'(\pi, \bar{\pi})$ as a function of the individual's strategy and the population's average vaccination rate for the same parameters as Figure 1. The absolute utility loss (left) is minimized for different individual choices, depending on the populations average behavior (All contour values are negative; smaller values represent higher utilities). The relative utility (right), calculated as the ratio of the absolute utility lose divided by the average utility lose, shows that vaccination rates around .195 are self-consistent in the sense that nobody can unilaterally improve on their situation by deviating from the average behavior.

For a given population behavior $\bar{\pi}$, the individual wishes to choose their own vaccination rate π to maximize Eq. (3.26). The best response correspondence² $\pi_{\text{best}}(\bar{\pi})$ is a set-valued mapping that returns the set of optimal replies to a population strategy $\bar{\pi}$. Thus, the best response correspondence is

$$\pi_{\text{Best}}(\bar{\pi}) = \begin{cases} 0 & \text{if } \bar{\pi} > \frac{a\beta}{\gamma} + \frac{a}{c\gamma+ca-1} \\ [0, \infty) & \text{if } \bar{\pi} = \frac{a\beta}{\gamma} + \frac{a}{c\gamma+ca-1} \\ \infty & \text{if } \bar{\pi} < \frac{a\beta}{\gamma} + \frac{a}{c\gamma+ca-1} \end{cases} \quad (3.27)$$

where $c = c_V/c_I$ is the relative cost of vaccination compared to infection. We find that there is a unique Nash equilibrium

$$\pi^* = \begin{cases} 0 & \text{if } c \geq \frac{\beta-\gamma}{\beta(a+\gamma)}, \\ \frac{a\beta}{\gamma} - \frac{a}{1-c(\gamma+a)} & \text{if } 0 \leq c < \frac{\beta-\gamma}{\beta(a+\gamma)}, \end{cases} \quad (3.28)$$

that is both an improvement on every alternative strategy the population may adopt and cannot be improved upon when it is adopted by the population. Only in the limit of $c = 0$ does the equilibrium strategy eradicate disease ($\pi^* = \hat{\pi}$). For all positive costs c , the equilibrium vaccination hazard is less than that needed for eradication ($\pi^* < \hat{\pi}$). When the population adopts π^* and $\beta > \gamma$, the disease incidence

$$I^* = N \frac{a}{\beta} \min \left\{ \frac{c\gamma}{1-c(a+\gamma)}, \frac{\beta-\gamma}{a+\gamma} \right\} \quad (3.29)$$

and the individual's expected asymptotic utility gain is

$$U'(\pi^*, \pi^*) = \max \left\{ -ac_V, -\frac{c_I a(\beta-\gamma)}{\beta(a+\gamma)} \right\} \quad (3.30)$$

For comparison, at the herd-immunity threshold,

$$U'(\hat{\pi}, \hat{\pi}) = - \left(1 - \frac{\gamma}{\beta} \right) c_V a, \quad (3.31)$$

which is strictly greater than the expected asymptotic utility gain at the game equilibrium.

Taken together, our analyses indicate that the recovery rate and the waning rates affect the relative equilibrium vaccination rate. When the waning rates are equal, faster waning decreases the equilibrium vaccination rate. In general, the longer the duration of natural resistance, the less vaccination at equilibrium. Higher vaccination costs decrease vaccination rate. More rapid recovery also decreases vaccination. Infinitely fast vaccination is never an equilibrium in this

²A correspondence is different from a function in that there may be many different outputs for a given input. Mathematically, correspondences are often replaced by functions mapping to the domain power set of the correspondences target. However, such mathematical consideration obscures the important fact that best-responses are not always unique.

special case [9], because all risks to the individual vanish under the standard-incidence assumption (Eq. (3.21)) as the vaccination rate approaches the eradication threshold given in Eq. (3.23). Infinitely fast vaccination rates may be equilibria if there is an external source of risk [27]. The inclusion of discounting h should slow vaccination, but the inclusion of demographic turnover ($\mu > 0$) may increase or decrease the vaccination rate, depending on the other parameter values.

3.1.2 Differential waning of immunity

These results can easily be extended to a case where vaccine-based resistance wanes at rate a_V , whereas naturally-acquired resistance wanes at rate a_R . The population-scale dynamics satisfy

$$\frac{dS}{dt} = a_R R + a_V V - \lambda(I)S - \bar{\pi}S, \quad (3.32a)$$

$$\frac{dI}{dt} = \lambda(I)S - \gamma I, \quad (3.32b)$$

$$\frac{dR}{dt} = \gamma I - a_R R, \quad (3.32c)$$

$$\frac{dV}{dt} = \bar{\pi}S - a_V V, \quad (3.32d)$$

Then, using the corresponding individual-scale model, with the same initial condition and costs we have already described, Nash equilibrium behavior is given by

$$\pi^* = \begin{cases} 0 & \text{if } c \geq \frac{(\beta-\gamma)a_R}{\beta a_V(a_R+\gamma)}, \\ \frac{\beta a_V}{\gamma} - \frac{a_V a_R}{a_R - c a_V(a_R+\gamma)} & \text{otherwise.} \end{cases} \quad (3.33)$$

When $a_V = a_R = a$, we recover Eq. (3.28). The algebraic calculations needed to reach this result are analogous to those of equal waning rates.

The response to changes in the waning rate of vaccine-induced immunity is not monotone. When vaccine waning is very slow, small increases also increase the critical vaccination rate. However, near the critical cost threshold, faster waning decreases the vaccination rate (Figure 3). Vaccination is only useful if its cost-per-time $c_V a_V$ is less than that of natural immunity $c_I \left(1 - \frac{\gamma}{\beta}\right) \frac{a_R}{a_R+\gamma}$.

3.1.3 Imperfect immunity

A more interesting extension is the case where immunity is imperfect. This case has been studied previously by Kremer, Snyder, and Williams (Snyder, personal communication) and by Chen [14]. The following discussion provides an alternative mathematical approach to some of their results. Our analysis indicates that the parameter region where multiple game equilibria can be found is vary narrow, and perhaps rarely encountered in practice.




Figure 3: The equilibrium vaccination rate π^* as a function of the waning rates a_R and a_V . The shaded region indicates parameter values where $\pi^* = 0$. Increases in the waning rate of natural immunity a_R increase Nash equilibrium vaccination. However, the response to changes in the waning rate of vaccine immunity a_V may either increase or decrease equilibrium vaccination. If natural immunity lasts longer when compared to vaccine immunity, the equilibrium is to refuse vaccination. Parameter values $\beta = 4$, $\gamma = 1$, $c = .1$.

Let us assume that vaccination confers partial immunity, and infection confers full immunity, with the condition that immunity wanes at the same rate in both cases. Under these modifications, System (3.1) becomes

$$\frac{dS}{dt} = a_R R + a_V V - \lambda S - \bar{\pi} S, \quad (3.34a)$$

$$\frac{dI}{dt} = \lambda(S + \sigma V) - \gamma I, \quad (3.34b)$$

$$\frac{dR}{dt} = \gamma I - a_R R, \quad (3.34c)$$

$$\frac{dV}{dt} = \bar{\pi} S - \sigma \lambda V - a_V V. \quad (3.34d)$$

When we construct the corresponding individual-scale model and solve for the utility, we find the utility of an individual strategy

$$U(\pi, \bar{\pi}) = \frac{-a_R \lambda^* (\lambda^* \sigma + a_V + \pi \sigma) c_I - a_R \pi \gamma (a_V + \lambda^* \sigma) c_V}{\sigma (\gamma + a_R) \lambda^{*2} + \lambda^* [(\gamma + a_R) (\pi \sigma + a_V) + a_R \sigma \gamma] + a_R \gamma (a_V + \pi)}. \quad (3.35)$$

If $\lambda = \beta I/N$, then the steady-state force of infection λ^* is either 0 or solves the quadratic polynomial

$$0 = \sigma (a_R + \gamma) \lambda^2 + (\gamma a_R \sigma + \gamma \bar{\pi} \sigma + a_R a_V + a_R \bar{\pi} \sigma + a_V \gamma - \beta a_R \sigma) \lambda + a_V \gamma a_R - \beta a_R a_V + \bar{\pi} a_R \gamma - \beta a_R \bar{\pi} \sigma \quad (3.36)$$

If $\sigma\beta/\gamma > 1$, even instant vaccination of every individual will not eradicate the disease.

The utility is still a linear-fractional transform in terms of the individual strategy π , so we expect all strategies to yield the same utilities at Nash equilibria. Local Nash equilibria can be identified in the standard manner (see Figure 4). No vaccination ($\pi^* = 0$) is a Nash equilibrium if

$$c > \frac{a_R(\beta - \gamma)(1 - \sigma)}{\beta(a_R\beta\sigma + a_Ra_V + \gamma a_V - \sigma\gamma a_R)} \quad (3.37)$$

where the relative vaccine cost is $c = c_V/c_I$. Thus, if $\sigma = 1$, vaccination is useless. If everyone is vaccinated as soon as they lose immunity ($\bar{\pi} = \infty$), the equilibrium force of infection is

$$\lambda^* = \max \left\{ 0, \beta \frac{I^*}{N} \right\} = \max \left\{ 0, \beta \left(1 - \frac{\gamma}{\sigma\beta} \right) \frac{a_R}{a_R + \gamma} \right\} \quad (3.38)$$

Instant vaccination by susceptible individuals ($\pi^* = \infty$) is a Nash equilibrium so long as $\sigma\beta/\gamma > 1$ and

$$c < \frac{a_R(\beta\sigma - \gamma)(1 - \sigma)}{(a_Ra_V + \sigma\beta a_R + a_V\gamma - a_R\gamma)(\beta\sigma + \sigma\gamma - \gamma)}. \quad (3.39)$$

In the special case of $a_R = a_V = a$, if $\mathcal{R}_0 = \beta/\gamma > 2$ and $0 < a < \mathcal{R}_0(\mathcal{R}_0 - 2)$ and $\frac{\mathcal{R}_0 + a}{\mathcal{R}_0(\mathcal{R}_0 - 1)} < \sigma < 1$, then there is an interval of relative costs c for which both vaccination and no vaccination are Nash equilibria (Figure 4). This interval can be bounded in terms of Eq. (3.37), Eq. (3.39), and the condition

$$(\gamma - \sigma\gamma + a)^2 c^2 - 2(1 - \sigma)(a + \gamma + \sigma\gamma)c + (1 - \sigma)^2 = 0 \quad (3.40)$$

representing the location of a fold bifurcation in the utility function (see Figure 5 for details). Further study of the utility function shows that both of these equilibria have local invasion potential, and are separated by a third Nash equilibrium that never has invasion potential. Plots and numerical experiments suggest that the regions of non-uniqueness are small, and probably narrower than the variances in population heterogeneity and parameter estimates when waning rates are equal. This suggests that the multiple-equilibrium phenomena identified by Chen and Kremer et al. may be of limited practical importance, but further systematic analysis, particularly in immunological contexts [28, 29], is warranted.

3.2 Delay Strategies

A strategy-space composed of constant hazards may not appeal to all researchers. Individuals probably do not choose waiting times at random from exponential distributions. An alternative formulation might be based on a decision to delay for a fixed time between the loss of immunity and the next vaccination.




Figure 4: The Nash equilibria vaccination rates π^* depending on the relative cost of vaccination ($c = c_V/c_I$). The Nash equilibrium is unique when $c < .393$ or $c > .426$. If $.393 < c < .426$, there are three Nash equilibria. Parameter values $\gamma = 1$, $\beta = 6$, $a_R = a_V = 0.05$, $\sigma = 0.15$.

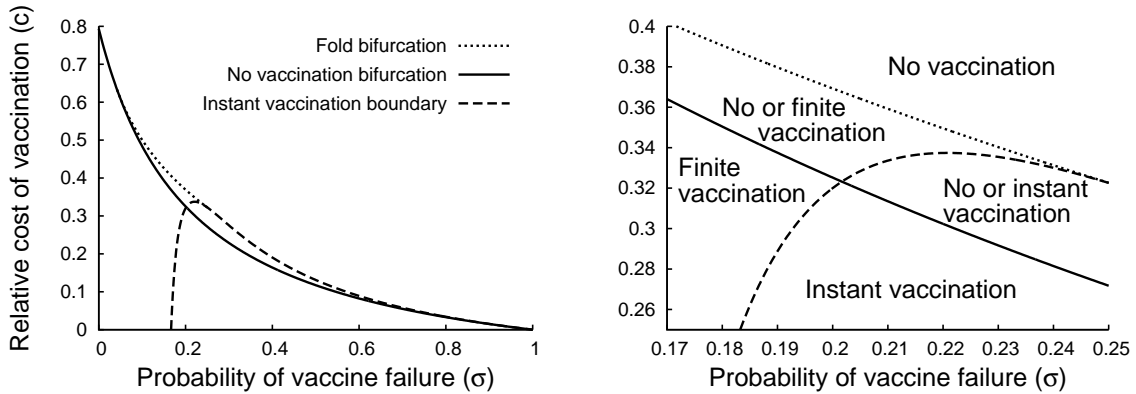


Figure 5: Parameter-space diagrams of the bifurcation structure in the Nash equilibria of Eq. (3.35) as functions of the relative probability of infection σ and the relative cost of vaccine $c = c_V/c_I$. The right plot is a magnification of the intersecting region of the left plot. The curves represent sets of parameters where bifurcations occur. The plot is divided into 5 regions. If the cost of vaccine is very high, no vaccination ($\pi^* = 0$) is the only equilibrium. If the cost of vaccine is very small and the relative probability of infection under vaccination is small, there is a unique equilibrium $0 < \pi^* < \infty$. If the relative probability of infection under vaccination is larger, individuals get vaccinated instantly on entering the susceptible compartment ($\pi^* = \infty$). For vaccine costs near the threshold for no vaccination, there may be two locally evolutionarily stable equilibria: No or instant vaccination ($\pi^* = 0$ and $\pi^* = \infty$) if the relative probability of infection is sufficiently high, and no or a finite vaccination rate if the relative probability of infection is small enough. Parameter values $\gamma = 1$, $\beta = 6$, $a_R = a_V = 0.05$.

The model we described here can easily be adapted to study this problem as well. Although the analysis is more difficult, the results are very similar to those for the constant-hazard strategy space.

The switch from hazard to delay strategies is equivalent to a switch from exponential distributions to delta-function distributions for the wait time before vaccination. To formulate an epidemic model with fixed wait times, we must switch gears from ordinary differential equations to integro-differential equations. The susceptible population will now be structured in terms of the amount of time that the individual has waited since becoming susceptible. This can be represented mathematically with a McKendrick–von Forrester style partial differential equation where the “age” dimension is interpreted as the wait time. After individuals have been susceptible for a certain time, they get vaccinated. The longer an individual remains susceptible, the less often they have to pay vaccination costs but the greater their accumulated risk of infection. This factor can be introduced as a removal hazard proportional to the force of infection. The number of newly infected individuals per unit time is the total of susceptible individuals becoming infected over all ages. Therefore, individuals optimize the tradeoff between vaccine costs and infection risks.

Other than the adjustments described above, the dynamics of the infected, vaccinated, and resistant states are the same as those described in System 3.1 when the mortality rate $\mu = 0$. At the population scale, then, the dynamics satisfy

$$\frac{dS}{dt} + \frac{dS}{dw} = -\lambda(I)S, \quad S(0, t) = aR(t) + aV(t), \quad w \in [0, \bar{w}], \quad (3.41a)$$

$$\frac{dI}{dt} = \lambda(I) \int_0^{\bar{w}} S(t, w)dw - \gamma I, \quad (3.41b)$$

$$\frac{dR}{dt} = \gamma I - aR, \quad (3.41c)$$

$$\frac{dV}{dt} = S(t, \bar{w}) - aV, \quad (3.41d)$$

with the total population size

$$N = \int_0^{\bar{w}} Sdw + I + R + V \quad (3.41e)$$

staying constant. Here, the wait time before vaccination, \bar{w} , is the population’s resident strategy.

The steady-state analysis of System (3.41) shows that there is always a disease-free stationary solution with $I^* = 0, V^* = N/(1 + a\bar{w})$, and $S^*(a) = aN/(1 + a\bar{w})$. There is also a endemic disease stationary solution with $I^* > 0$ if and only if

$$\lim_{I \rightarrow 0} \frac{\lambda(I)}{\gamma I} \frac{a\bar{w}N}{1 + a\bar{w}} > 1. \quad (3.42)$$

In the special case of $\lambda(I) = \beta I/N$, the critical vaccination delay

$$\hat{w} = \frac{\gamma}{a\beta(\beta - \gamma)}. \quad (3.43)$$

Analogous to Eq. (3.7), the endemic disease incidence is the unique positive solution of the transcendental equation [**TCR: This looks very wrong!... Nope, it's mostly right, but I forgot to divide one term by N**]

$$\gamma\lambda(I^*) = a[\beta(1 - I^*/N) - \gamma] [1 - e^{-\lambda(I^*)\bar{w}}] \quad (3.44)$$

[**TCR: I think the correct thing is...**]

$$\lambda(I^*) = \frac{a\gamma I^* (1 - e^{-\lambda(I^*)\bar{w}})}{a(N - I^*) (1 - e^{-\lambda(I^*)\bar{w}}) - \gamma I^*} \quad (3.45)$$

Note that this corresponds with (3.7) in the absence of vaccination and population turnover.

Calculation of an individual's utility is significantly more complicated than in the hazard-strategy model. It is not straight-forward to write down the forward generator for the Markov process describing an individual's state evolution. The forward generator can be constructed using generating functionals, but the explanation and application of these methods is outside the scope of this paper. Instead, we use a recursive decomposition of the path-integral formulation of the expected utility given by Eq. (2.5) to reduce the necessary calculations to matrix arithmetic.

The instantaneous utility gain of each state other than the infected state is 0. When individuals are infected, they lose utility at a rate of c_I per unit time. In addition, there is an instantaneous cost c_V paid every time an individual enters the vaccinated class. From these observations, we can construct recursive formulas for evaluating the expected utility of various initial states. Let U_I , U_S , U_V , and U_R be the utilities conditional on initial infected, susceptible, vaccinated, and resistant states respectively. Suppose an individual is initially in the infected state. Under the current model, there is a constant hazard γ per unit time of leaving the infected state and entering the recovered class. The total time t_1 that the individual spends in the infected state has distribution $\gamma e^{-\gamma t_1}$. The individual loses utility at rate $\phi(I) = -c_I$ for every day they are infected. The expected utility, then, should be the expectation of $-c_I t_1$ plus the expected utility of entering the recovered class at time t_1 , appropriately discounted. Mathematically, the expected utility is evaluated

as follows:

$$U_I = \int_I \int_0^\infty e^{-ht} \phi(x(t)) dt \mathcal{D}x \quad (3.46a)$$

$$= \int_0^\infty \gamma e^{-\gamma t_1} \left[\int_0^{t_1} e^{-ht} \phi(I) dt + \int_R \int_{t_1}^\infty e^{-ht} \phi(x(t)) dt \mathcal{D}x \right] dt_1 \quad (3.46b)$$

$$= \int_0^\infty \gamma e^{-\gamma t_1} \left[\left(\frac{1 - e^{-ht_1}}{h} \right) (-c_I) + e^{-ht_1} \int_R \int_0^\infty e^{-ht} \phi(x(t)) dt \mathcal{D}x \right] dt_1 \quad (3.46c)$$

$$= \int_0^\infty \gamma e^{-\gamma t_1} \left[\left(\frac{e^{-ht_1} - 1}{h} \right) c_I + e^{-ht_1} U_R \right] dt_1 \quad (3.46d)$$

$$= \frac{-c_I + \gamma U_R}{\gamma + h} \quad (3.46e)$$

The expected utilities of the other initial states are calculated in a similar fashion.

$$U_S = \int_{s(0)} \int_0^\infty e^{-ht} \phi(x(t)) dt \mathcal{D}x \quad (3.47a)$$

$$= \int_0^w \lambda(I^*) e^{-\lambda(I^*) t_1} \left[\int_0^{t_1} e^{-ht} \phi(S) dt + \int_I \int_{t_1}^\infty e^{-ht} \phi(x(t)) dt \mathcal{D}x \right] dt_1 \quad (3.47b)$$

$$+ e^{-\lambda(I^*) w} \int_V \int_w^\infty e^{-ht} \phi(x(t)) dt \mathcal{D}x$$

$$= \int_0^w \lambda(I^*) e^{-\lambda(I^*) t_1} \left[\left(\frac{1 - e^{-ht_1}}{h} \right) \phi(S) + e^{-ht_1} U_I \right] dt_1 + e^{-(\lambda(I^*)+h)w} U_V \quad (3.47c)$$

$$= \int_0^w \lambda(I^*) e^{-(\lambda(I^*)+h)t_1} U_I dt_1 + e^{-(\lambda(I^*)+h)w} U_V \quad (3.47d)$$

$$= \lambda(I^*) \left(\frac{1 - e^{-(\lambda(I^*)+h)w}}{\lambda(I^*) + h} \right) U_I + e^{-(\lambda(I^*)+h)w} U_V \quad (3.47e)$$

$$U_V = -c_V + \int_V \int_0^\infty e^{-ht} \phi(x(t)) dt \mathcal{D}x \quad (3.48a)$$

$$= -c_V + \int_0^\infty a e^{-at_1} \left[\int_0^{t_1} e^{-ht} \phi(V) dt + \int_S \int_{t_1}^\infty e^{-ht} \phi(x(t)) dt \mathcal{D}x \right] dt_1 \quad (3.48b)$$

$$= -c_V + \frac{a}{a+h} U_S \quad (3.48c)$$

$$U_R = \int_R \int_0^\infty e^{-ht} \phi(x(t)) dt \mathcal{D}x \quad (3.49a)$$

$$= \int_0^\infty a e^{-at_1} \left[\int_0^{t_1} e^{-ht} \phi(R) dt + \int_S \int_{t_1}^\infty e^{-ht} \phi(x(t)) dt \mathcal{D}x \right] dt_1 \quad (3.49b)$$

$$= \frac{a}{a+h} U_S \quad (3.49c)$$

Thus, the conditional expected utility for each of the four initial conditions can be expressed as a linear equation in terms of the other conditional expected utilities. In matrix form,

$$\begin{bmatrix} U_S \\ U_I \\ U_R \\ U_V \end{bmatrix} = \begin{bmatrix} 0 & \lambda(I^*) \left(\frac{1 - e^{-(\lambda(I^*)+h)w}}{\lambda(I^*)+h} \right) & 0 & e^{-(\lambda(I^*)+h)w} \\ 0 & 0 & \frac{\gamma}{\gamma+h} & 0 \\ \frac{a}{a+h} & 0 & 0 & 0 \\ \frac{a}{a+h} & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} U_S \\ U_I \\ U_R \\ U_V \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{-c_I}{\gamma+h} \\ 0 \\ -c_V \end{bmatrix} \quad (3.50)$$

This matrix equation can be solved to determine the conditional expected utilities. The unconditional expected utility is then determined by weighing each conditional expected utility and the probability of the corresponding initial condition.

However, in the case of slow discounting (small h), it is useful to study the asymptotic behavior. Let the asymptotic utility gain rate

$$U^+ = \lim_{h \rightarrow 0} h U_{x(0)} = \frac{-a\lambda(I^*) [c_I (e^{\lambda(I^*)w} - 1) + c_V \gamma]}{[(\gamma + a)\lambda(I^*) + a\gamma] e^{\lambda(I^*)w} - a [\gamma + \lambda(I^*)]}. \quad (3.51)$$

Note that U^+ is independent of the initial condition $x(0)$.

An example of U^+ is plotted in Figure 6. If a player chooses a very long ($w = \infty$) or a very short ($w = 0$) delay, the asymptotic utility gain

$$U^+(\infty, \bar{w}) = \frac{-a\lambda(I^*)c_I}{(\gamma + a)\lambda(I^*) + a\gamma}, \quad \text{and} \quad U^+(0, \bar{w}) = -ac_V. \quad (3.52)$$

As in the hazard-strategy example, $\partial U^+/\partial w$ vanishes for all individual strategies w when the force of infection λ satisfies

$$c_V = \frac{c_I}{a + \gamma + \frac{a\gamma}{\lambda(I^*)}}. \quad (3.53)$$

Therefore, if there is a non-negative vaccination strategy w^* in which the force of infection satisfies Eq. (3.53) and (3.45) when $\bar{w} = w^*$, then w^* is a Nash equilibrium. When

$$\frac{c_I}{a + \gamma + \frac{a\gamma}{\lambda(I^*(\bar{w}=0))}} < c_V < \frac{c_I}{a + \gamma + \frac{a\gamma}{\lambda(I^*(\bar{w}=\infty))}}, \quad (3.54)$$

Eq. (2.19) is negative because $d^2U^+/dw^2 = 0$ and

$$\frac{d^2U^+}{d\bar{w}dw} (w = w^*, \bar{w} = w^*) = \frac{a\gamma c_V^2 e^{\lambda w} [c_V(\gamma + a) - c_I]}{[c_I(e^{\lambda w} - 1) + c_V\gamma]^2} \frac{d\lambda}{d\bar{w}} < 0 \quad (3.55)$$

since $c_V(\gamma + a) < c_I$ by Eq. (3.54). Thus, the critical strategy also has local invasion potential. If

$$\frac{c_I}{a + \gamma + \frac{a\gamma}{\lambda(I^*(\bar{w}=\infty))}} < c_V, \quad (3.56)$$

then the equilibrium shifts to no vaccination. This condition is automatically satisfied when $c_V > c_I/(\gamma + a)$. If

$$c_V < \frac{c_I}{a + \gamma + \frac{a\gamma}{\lambda(I^*(\bar{w}=0))}}, \quad (3.57)$$

then the equilibrium would be to get vaccinated as soon as the individual loses immunity. But $\lambda(I^*(\bar{w} = 0)) = 0$, so the right hand side diverges and this inequality is never satisfied under the current model.

While the force of infection is the same in both models for the equilibrium strategy, the individual utility gain at the critical delay-strategy is

$$U^+(w^*, w^*) = \max \left\{ -ac_V, \frac{-ac_I}{a + \gamma + \frac{a\gamma}{\lambda(I^*)}} \right\}. \quad (3.58)$$

This is consistent with the hazard-strategy model.




Figure 6: An individual’s asymptotic expected absolute (left) and relative (right) utility losses calculated from Eq. (3.51) as a function of the individual’s strategy and the population’s resident strategy. Parameter values are the same as those used Figures 1 and 2.

4 Discussion

We have described a general method for the calculation of utilities of strategies to individuals in population games with applications to epidemiology, public health, and many other fields. Our method offers a framework for integrating individual and aggregate behaviors. In contrast to population-scale dynamics, individual-scale dynamics are fundamentally stochastic. By combining deterministic models of macroscopic population dynamics with stochastic models of an individual’s dynamics in an optimal-control setting, we have created a powerful reformulation of population games for the study of decision making. We employ a continuous-time Markov decision process theory to capture the probabilistic events for an individual. Each individual’s state evolves according to a Markov process. The resulting transition rate matrix contains elements derived from the population-scale epidemiological model, and is used to calculate the evolution of an individual’s utility through time and at equilibrium.

Our method has wide applicability for population games. For example, our method facilitates the integration of models for social and biological systems, game theory and public policy. Our method also permits the calculation of utilities from relatively complex population games. Consequently, our method allows the expansion of the complexity of population game models in order to take into account realism that could not be incorporated previously. For example, our utility calculation method can be formulated to accommodate population heterogeneity, such as with regard to risk aversion or age [18].

Our previous work on vaccination against seasonal influenza epidemics [18] and age-dependent virulence [19] spurred our interest in applying the models

described in this paper for vaccination. At the level of the individual, a vaccination decision is made to minimize future risks associated with infection. An individual's perceived risk of infection depends on several epidemiological and clinical parameters included in our epidemiological model. Our model simultaneously captures the macroscopic behavior of the population and the microscopic behavior of individuals by incorporating all the necessary parameters to obtain a complete description of vaccination decision dynamics. Our model shows that the amount of time an individual will spend specifically in the susceptible, infected, recovered, and vaccinated compartments is stochastic. Decision utilities depend not only on that individual's current state, but on the probabilistic trajectory of future states.

We showed that the utilities of vaccination decisions can be determined using Markov decision process theory by applying our method to two vaccination models that differ in vaccination timing. We demonstrated the application of these utilities to determine the invisibility of strategies and to calculate Nash equilibria. We also used our method to characterize the utility landscape surrounding critical strategies, revealing that the relationship between the waning of vaccine-induced immunity and the level of vaccination at the Nash equilibrium is not monotone. When vaccine waning is very slow, small increases in waning rate elevate expected vaccination demand. Around the critical cost threshold, however, more rapid waning decreases vaccination.

On-going research includes the analysis of differential games, trembling-hand equilibria conditions, and extensions to populations including age, space, and contact network structures. Alternative ways to specify the coupling between individual's actions and the population dynamics are also an important topic. In cases where actors have good information, unlimited computing power, clear motives, and act independently, game-theory approaches like that studied here are justifiable. However, human behavior frequently diverges from these assumptions, and many alternative approaches have not yet been fully explored [30, 31]. We are also further investigating the importance of immunological dynamics for these decision-problems.

In conclusion, we described a general method for the calculation of the utilities of strategies to individuals in population games. This method offers a framework for integrating individual and aggregate behaviors. We exemplified its applicability by employing our method to determine the invisibility of strategies and to calculate Nash equilibria for vaccination decisions against seasonal influenza. We showed that the utilities of vaccination decisions can be determined using Markov decision process theory. Unlike previous works, our method can be widely applied for population games and permits calculation of utilities from relatively complex systems.

Acknowledgements

This research was supported by NIH grant R01AI072706 (APG and TCR), and Bill and Melinda Gates Foundation Grant Number 49276 (TCR).

References

- [1] R. Ross, H. P. Hudson, An application of the theory of probabilities to the study of a priori pathometry - part II, *Proceedings of the Royal Society London A* 43 (1917) 212–225.
- [2] W. O. Kermack, A. G. McKendrick, Contributions to the mathematical-theory of epidemics, *Proceedings of the Royal Society of London* 115 (1927) 700–721.
- [3] F. Sonnenberg, J. Beck, Markov models in medical decision making: a practical guide, *Medical Decision Making* 13 (1993) 322.
- [4] J. D. Sterman, Learning from evidence in a complex world, *American Journal of Public Health* 96 (2006) 505–514.
- [5] P. E. M. Fine, J. A. Clarkson, Individual versus public priorities in the determination of optimal vaccination policies, *American Journal of Epidemiology* 124 (1986) 1012–1020.
- [6] D. L. Brito, E. Sheshinski, M. D. Intriligator, Externalities and compulsory vaccinations, *Journal of Public Economics* 45 (1991) 69–90.
- [7] P.-Y. Geoffard, T. Philipson, Disease eradication: Private versus public vaccination, *The American Economic Review* 87 (1997) 222–230.
- [8] T. Philipson, Economic epidemiology of infectious disease, in *Handbook of health economics*, A. J. Culyer, J. P. Newhouse, eds., Elsevier, New York, NY, 2000, vol. 1B, 1761–1797.
- [9] C. T. Bauch, A. P. Galvani, D. J. D. Earn, Group interest versus self-interest in smallpox vaccination policy, *Proceedings of the National Academy of Sciences* 100 (2003) 10564–10567.
- [10] P. J. Francis, Optimal tax/subsidy combinations for the flu season, *Journal of Economic Dynamics and Control* 28 (2004) 2037–2054.
- [11] F. H. Chen, Rational behavioral response and the transmission of stds, *Theoretical Population Biology* 66 (2004) 307–316.
- [12] C. T. Bauch, D. J. D. Earn, Vaccination and the theory of games, *Proceedings of the National Academy of Sciences* 101 (2004) 13391–13394.
- [13] C. T. Bauch, Dynamic games with imitation predict vaccinating behavior, *Proceedings of the Royal Society of London, Series B* 272 (2005) 1669–1675.
- [14] F. H. Chen, A susceptible-infected epidemic model with voluntary vaccinations, *Journal of Mathematical Biology* 53 (2006) 253–272.
- [15] F. H. Chen, On the transmission of hiv with self-protective behavior and preferred mixing, *Mathematical Biosciences* 199 (2006) 141–159.

- [16] M. Kremer, C. M. Snyder, H. Williams, Which vaccines deserve the largest subsidies? an integrated economic and epidemiological model 2006. Working paper.
- [17] F. H. Chen, Modeling the effect of information quality on risk behavior change and the transmission of infectious diseases, *Mathematical Biosciences* 217 (2009) 125–133.
- [18] A. P. Galvani, T. C. Reluga, G. Chapman, Long-standing influenza vaccination policy is in accord with individual self-interest but not with the utilitarian optimum, *Proceedings of the National Academy of Sciences* 104 (2007) 5692–5697.
- [19] T. C. Reluga, J. Medlock, E. Poolman, A. P. Galvani, Optimal timing of disease transmission in an age-structured population, *Bulletin of Mathematical Biology* 69 (2007) 2711–2722.
- [20] R. Rockafellar, Coherent approaches to risk in optimization under uncertainty, *Tutorials in Operations Research INFORMS* (2007) 38–61.
- [21] R. A. Howard, *Dynamic Programming and Markov Processes*, MIT Press, Cambridge, MA, 1960.
- [22] S. A. H. Geritz, E. Kisdi, G. Meszena, J. A. J. Metz, Evolutionary singular strategies and the adaptive growth and branching of the evolutionary tree, *Evolutionary Ecology* 12 (1998) 35–57.
- [23] F. Christiansen, V. Loeschke, Evolution and intraspecific exploitative competition I. One-locus theory for small additive gene effects, *Theoretical Population Biology* 18 (1980) 297–313.
- [24] B. Thomas, On evolutionarily stable sets, *Journal of Mathematical Biology* 22 (1985) 105–115.
- [25] M. Mesterton-Gibbons, *An Introduction to Game-Theoretic Modelling*, Addison-Wesley, Redwood City, CA, 1992.
- [26] T. C. Reluga, An sis game with two subpopulations, *Journal of Biological Dynamics* 3 (2009) 515–531.
- [27] T. C. Reluga, C. T. Bauch, A. P. Galvani, Evolving public perceptions and stability in vaccine uptake, *Mathematical Biosciences* 204 (2006) 185–198.
- [28] T. C. Reluga, J. Medlock, A. Perelson, Backward bifurcations and multiple equilibria in epidemic models with structured immunity, *Journal of Theoretical Biology* 252 (2008) 155–165.
- [29] J. Heffernan, M. Keeling, Implications of vaccination and waning immunity, *Proceedings of the Royal Society B* 276 (2009) 2071–2080.

- [30] G. Gigerenzer, P. M. Todd, Simple Heuristics That Make Us Smart, Oxford University Press, 2000.
- [31] D. Kahneman, Maps of Bounded Rationality: Psychology for Behavioral Economics, American economic review 93 (2003) 1449–1475.