

Optimization of Influenza Vaccine Strain Selection

by

Joseph T. Wu

Bachelor of Science in Chemical Engineering
Massachusetts Institute of Technology, 1999

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Operations Research

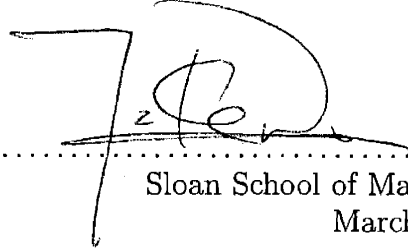
at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2003

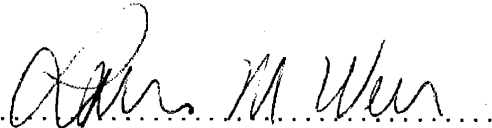
© Massachusetts Institute of Technology 2003. All rights reserved.

Author



Sloan School of Management
March 14, 2003

Certified by




Lawrence M. Wein
Professor of Operations, Information and Technology
Thesis Supervisor

Accepted by



John N. Tsitsiklis
Codirector, Operations Research Center

ARCHIVES



Optimization of Influenza Vaccine Strain Selection

by

Joseph T. Wu

Submitted to the Sloan School of Management
on March 14, 2003, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Research

Abstract

The World Health Organization (WHO) is responsible for making annual vaccine strains recommendation to countries around the globe. However, various studies have found that the WHO vaccine selection strategy has not been effective in some years. This motivates the search for a better strategy for choosing vaccine strains. In this work, we use recent results from theoretical immunology to formulate the vaccine selection problem as a discrete-time stochastic dynamic program with a high-dimensional continuous state space. We discuss the techniques that were developed for solving this difficult dynamic program, and present an effective and robust heuristic policy. We compare the performance of the heuristic policy, the follow policy, and the no-vaccine policy and show that the heuristic policy is the best among the three. After taking the cost of implementation into account, however, we conclude that the WHO policy is a cost-effective influenza vaccine strain selection policy.

Thesis Supervisor: Lawrence M. Wein

Title: Professor of Operations, Information and Technology

Acknowledgements

First, I wish to thank Professor Lawrence Wein for being a very supportive and understanding advisor since the first day I walked into his office. I would also like to thank Dr. Alan Perelson for giving me the opportunity to work with him at Los Alamos National Laboratory. Both Larry and Alan have been really wonderful mentors, and I am very fortunate to be their student. I wish I can elaborate on how much they have done for me, but it's already 5 am... so please forgive me for my terseness. :)

My family and friends have been a continual source of encouragement for my learning. I am forever indebted to my grandma and parents for all the things they have done for me. Thank you Waiduen for sharing my good days and bad days for the last two and a half years. And many thanks to all my friends for making my days at MIT such an unforgettable experience: Ian, Eddy, Kayi, Ching, Monique, Joyce, Addy, Damian, Sanne, Romy, Susan, Michele, and many others. Thank you guys!

It's 5:15 am... and finally, here is the last page of my thesis.

Contents

1	Introduction	11
1.1	Health and Economic Impact of Influenza	11
1.2	Influenza Viruses	12
1.3	Influenza Vaccination	13
1.4	Motivation and Objectives	14
2	Formulation of the Model	17
2.1	The Role of Antibodies in Anti-Flu Immune Response	17
2.2	Shape Space and the Antigenic Distance Hypothesis	18
2.3	Influenza Vaccine Strain Selection as a Discrete-time Control Problem	20
2.3.1	How does the epidemic strain sequence $\{e_t\}$ evolve?	22
2.3.2	How does the immune state change after an immune response?	22
2.3.3	How do we measure the efficacy of a vaccine strain?	26
2.4	Dynamic Programming Formulation	26
3	Analysis of the 1-Period Problem	29
3.1	Some Useful Formulas	29
3.2	Reformulation of the 1-Period Problem	30
3.3	Properties of the Optimal Solution to the Transformed 1-D Problem	32
3.4	Construction of the Approximate Optimal Solution	34
3.5	Testing the Performance of Different Policies in the 1-Period Problem	36
4	Numerical Methods for DP Approximation	43
4.1	Parameter Values for DP	43
4.2	Breaking the Curse of Dimensionality via Features Extraction	44

4.3	Reward-To-Go Approximation at the 1st DP Iteration	45
4.4	Reward-To-Go Approximation at the 2nd DP Iteration	46
4.5	DP Approximation Algorithm	49
5	Discussion	61
5.1	Performance Evaluation of Different Policies	61
5.2	Sensitivity of Policies Performance to Parameter Values	64
5.3	Limitations of the Model and Recommendations on Vaccine Strain Selection	73
A	Proofs and Formulas	77
A.1	Proof of Theorem 3.1	77
A.2	Proof of Theorem 3.2	80
A.3	Closed Form Formulas for Reward-To-Go Approximations	81
A.4	Parameter Values in Case Studies	84

List of Figures

1-1	The basic structure of influenza A and B.	12
2-1	The dynamics of cross-reactions.	20
3-1	Symmetry in the example where $h_N = v_{N-1}$	31
3-2	Plot of the weight function for $\eta = 5$. The weight drops to negligible level when $\ s\ \geq 1$	37
3-3	Frequency plots of R^h and R^f	38
3-4	Frequency plots of r^h , r^f , and r^n	39
3-5	Scatterplot of R^h and R^f versus the level of residual immunity.	40
3-6	Frequency plot of $\frac{F_N^h(v_N^h)}{F_N(v_N^*)}$	41
3-7	Frequency plots of R^h and R^f for $n = 2$	41
3-8	Frequency plots of r^h , r^f , and r^n for $n = 2$	42
3-9	Frequency plot of $\frac{F_N^h(v_N^h)}{F_N(v_N^*)}$ for $n = 2$	42
4-1	R-square values for reward-to-go approximations at the 1st DP iteration.	47
4-2	R-square values for reward-to-go approximations at the 2nd DP iteration.	52
4-3	R-square values for reward-to-go approximations at the 3rd DP iteration.	53
4-4	R-square values for reward-to-go approximations at the 4th DP iteration.	54
4-5	R-square values for reward-to-go approximations at the 5th DP iteration.	55
4-6	R-square values for reward-to-go approximations at the 6th DP iteration.	56
4-7	R-square values for reward-to-go approximations at the 7th DP iteration.	57
4-8	R-square values for reward-to-go approximations at the 8th DP iteration.	58
4-9	R-square values for reward-to-go approximations at the 9th DP iteration.	59
5-1	Reward ratios of the heuristic, follow, and no-vaccine policy in each case study.	65

5-2	Performance of the heuristic, follow, and no-vaccine policy under different values of κ_v . The line shown in each plot is a 45 degree line. If a data point lies above the line, this means that the performance of the y -axis case is better than that of the x -axis case, and vice-versa.	66
5-3	Performance of the heuristic, follow, and no-vaccine policy under different values of α . If a data point lies above the line, this means that the performance of the y -axis case is better than that of the x -axis case, and vice-versa.	67
5-4	Performance of the heuristic, follow, and no-vaccine policy under different values of θ . If a data point lies above the line, this means that the performance of the y -axis case is better than that of the x -axis case, and vice-versa. . . .	69
5-5	Performance of the heuristic, follow, and no-vaccine policy under different values of p_e . If a data point lies above the line, this means that the performance of the y -axis case is better than that of the x -axis case, and vice-versa.	70
5-6	Performance of the heuristic, follow, and no-vaccine policy under different values of $\ d\ $. If a data point lies above the line, this means that the performance of the y -axis case is better than that of the x -axis case, and vice-versa.	71
5-7	Performance of the heuristic, follow, and no-vaccine policy under different values of c_v . If a data point lies above the line, this means that the performance of the y -axis case is better than that of the x -axis case, and vice-versa.	72

Chapter 1

Introduction

1.1 Health and Economic Impact of Influenza

Influenza is an acute respiratory disease that has caused significant morbidity and mortality throughout human history. Influenza pandemic (a worldwide epidemic) arose at irregular intervals. The most devastating pandemic of the 20th century occurred in 1918-19 when 500,000 people died in the U.S., and 20 million died worldwide [1]. Two other pandemics occurred in 1957-58 (70,000 deaths in the U.S.) and 1968-69 (34,000 deaths in the U.S.) [1]. These two pandemics together killed more than 1.5 million people and caused an estimated \$32 billion in economic damages worldwide due to productivity losses and medical expenses [2]. Between global pandemics, local influenza epidemics occur on an annual basis. According to the Centers for Disease Control and Prevention (CDC), influenza epidemic affects 25 to 50 million people and is linked to 35,000 deaths and more than 114,000 hospitalizations each year in the United States. Influenza and pneumonia (often a complication of influenza) together is the sixth leading cause of death in the United States (65,000 deaths in 2000). More than 85 percent of these deaths are among people 65 years of age or older. In 1993, it was estimated that Medicare reimbursement for excess hospitalizations during influenza epidemics ranged from \$750 million to \$1 billion [5]. In view of the expensive hospitalization costs related to preventable influenza complications, Congress authorized an \$80 million per year Medicare flu shot entitlement in 1993.

1.2 Influenza Viruses

Influenza viruses are negative-strand RNA viruses with segmented genomes. Influenza in human beings are caused by the A, B, and C serotype of influenza viruses. Influenza due to the C virus has seldom been reported; this is probably due to the widespread of antibody to the C serotype among adults. In contrast, the A and B virus are both responsible for the outbreak of annual epidemics.

Influenza A and B viruses share the same basic biological structure. A schematic drawing of this viral structure is shown in Figure 1-1 [3]. From the viewpoint of vaccine design and

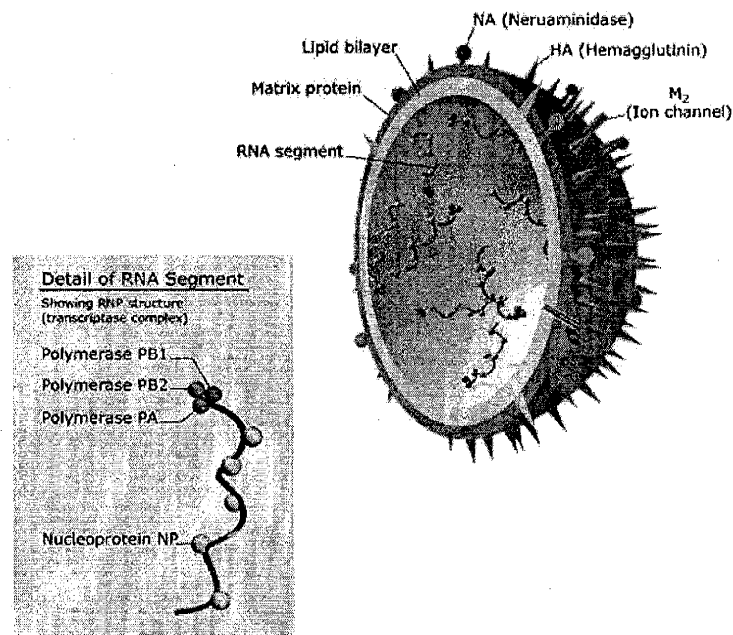


Figure 1-1: The basic structure of influenza A and B.

development, the two surface proteins, hemagglutinin (HA) and neuraminidase (NA), are of paramount importance. HA allows the virus to enter host cells and begin infection; NA destroys host cell's membrane for release of newly formed viral particles from the surface of the host cell. These two proteins are the primary targets of the humoral immune response (immunity mediated by circulating antibodies) in humans. Influenza A serotype are further divided into subtypes according to the major differences in HA and NA (e.g. A (H1N1), A (H3N2)). Unlike influenza A, influenza B has no subtypes. Within each A subtype and the B serotype, new strains arise frequently due to the high mutation rate of HA and NA. (Here,

“strains” mean mutated variants within an A subtype or the B serotype. In general, strains are functionally identical, but the HA and NA structures of nonidentical strains are slightly different.) These mutations often result in new strains against which human beings have little or no immunity, thus necessitate frequent updates of vaccine strains to fight against these new strains. Small mutations are known as *antigenic drifts*, which are caused by point mutations in the HA and NA genes. Both Influenza A and B virus undergo antigenic drifts which constitute yearly epidemics. Large mutations are known as *antigenic shifts*, which result from gene reassortment of viral RNA genomes among different A subtypes; antigenic shifts give rise to unpredictable pandemics.

1.3 Influenza Vaccination

During the last decade, standard influenza vaccines (commonly known as flu shots) contain inactivated antigens against the three most epidemiologically important flu viruses: A (H1N1), A (H3N2), and B. Vaccine potency requirements have been focused on HA instead of NA for two main reasons [16]. First of all, the degree of structural difference in HA among influenza strains can be measured conveniently by methods such as hemagglutinin-inhibition (HI) assay and single radial immunodiffusion (SRD) test; NA, on the other hand, is less well characterized due to the lack of reliable tests. Second of all, HA has been confirmed to be capable of stimulating a protective immune response to influenza; in contrast, the immunogenicity of NA has yet to be established definitely.

In making annual vaccine strain recommendations to countries around the globe, the World Health Organization (WHO) chooses a vaccine strain whose HA structure matches closely with that of the upcoming dominant epidemic strain. In the rest of this thesis, we define “follow” policy as the strategy that uses the expected epidemic strain as the vaccine strain. The actual WHO vaccine strain selection strategy is a slight variant of the follow policy: The vaccine strain is updated if there is a 4-fold difference (measured by HI assays) between the previous vaccine strain and the expected epidemic strain. That is, vaccine recommendation is updated if there is a large enough difference between the previous vaccine strain and the expected epidemic strain. The purpose of this modification is to facilitate the vaccine manufacturing process. Therefore, the follow policy is an idealized version of the current WHO vaccine strain selection strategy. It is obvious that an accurate

epidemic strain prediction is pivotal to the effectiveness of the follow policy. To predict the epidemic strains that will prevail in future flu seasons, a worldwide surveillance network has been set up by the WHO to monitor the evolution of influenza viruses.

1.4 Motivation and Objectives

Data from the National Center for Health Statistics show that 15.1%, 32.2%, and 63.0% of the 18-49, 50-64, and >65 age group in the U.S. received influenza vaccination in 2001 [4]. The CDC estimated that vaccine efficacy is 70-90% in healthy adults and 30-70% in elderly people when there is a good match between the vaccine and epidemic strain. In addition to physique and age, influenza vaccine efficacy varies among individuals for another major reason: vaccine efficacy depends on the individual's prior exposure to influenza. It has been shown that influenza vaccination works effectively in first-time vaccinees [11]. However, vaccine efficacy in repeat vaccinees has not been established definitely. Different studies on the effect of prior immunization on vaccine efficacy have generated paradoxical conclusions. Beyer *et al.* found that annually repeated vaccination did not reduce vaccine efficacy [6]. Gross *et al.*, on the other hand, showed that adults who were previously vaccinated had significantly lower serologic response to influenza vaccine than first-time vaccinees [12]. The widely cited "Hoskins" and "Keitel" studies have also reached different conclusions on the efficacy of repeated vaccination: Hoskins *et al.* concluded that repeated vaccination was not effective [13], while Keitel *et al.* concluded the contrary [14].

To resolve the seemingly contradictory observations of vaccine efficacy in repeated vaccinations, Smith *et al.* proposed an explanation – the *antigenic distance hypothesis* [21]. This hypothesis states that variations in repeat vaccine efficacy are due to antigenic differences between the vaccine strain and all the strains to which the vaccinee has previously been exposed. To test this hypothesis, Smith *et al.* constructed a computer model which estimated vaccine efficacy in repeat vaccinees relative to vaccine efficacy in first-time vaccinees. The computed efficacy had good correlation with the observed efficacy in the Hoskins and Keitel studies ($r = 0.87$) [21]. An independent study by Lapedes and Farber [15] also supports the validity of the hypothesis. Thus, the antigenic distance hypothesis provides a simple and credible explanation for the observed discrepancies in the efficacy of repeated vaccinations.

While most efforts in influenza vaccination have been focused on predicting the upcoming

epidemic strains [7, 19], the potential of devising more effective vaccine strain selection strategies (other than the follow policy) has not been actively explored. Moreover, over the last decade, numerous medical researchers have questioned the cost-effectiveness of influenza vaccines and advocated the “no-vaccine” policy. A review of the cost-effectiveness of current influenza vaccine selection practice calls for a performance evaluation of the follow policy, the no-vaccine policy, and possibly other more effective policies. In this thesis, we will use the antigenic distance hypothesis as a basis to analyze the seasonal influenza vaccination process. The seasonal vaccine strain selection problem is formulated as an N -period discrete-time continuous-state-space stochastic dynamic program. The objective of this thesis is to obtain insights on how vaccine strains should be chosen for individuals with different histories of influenza exposure. The remainder of this thesis takes the following form. In Chapter 2, we give an overview on (1) the nature of anti-flu immune response, (2) the concept of “shape space” which allows a mathematical characterization of influenza strains, and (3) the theoretical influenza vaccination model (a dynamic program). In Chapter 3, we analyze the first iteration of the dynamic program (which is equivalent to a single-period problem) and develop a simple way to visualize the vaccine selection problem. In Chapter 4, we use the results in Chapter 3 to devise a efficient algorithm for solving the high-dimensional stochastic dynamic program approximately (to a reasonable accuracy). In Chapter 5, we review the numerical results of our model and discuss their implications on vaccine selection policies. We will also discuss the practicability of different policies and the limitations of our model.

Chapter 2

Formulation of the Model

2.1 The Role of Antibodies in Anti-Flu Immune Response

Antibodies are our major line of defense against influenza infection. Antibodies are proteins that circulate inside the bloodstream and their primary function is to bind to and remove foreign substances that invade into the body. Antibodies bind to the surface of foreign substances via target sites known as *antigens*. By binding to the antigens on a foreign substance, antibodies can prevent the foreign substance from binding to cells and causing infections. Moreover, formation of antibody-antigen complexes facilitate the destruction of foreign substances by complements and macrophages as well as the uptake of foreign substances by antigen-presenting cells (which is necessary for triggering T-cell responses).

In order to recognize the wide range of foreign substances that can invade into the body, the antibody repertoire comprises antibody molecules with a huge collection of structural specificities. The range of antibody specificities is large enough such that given any foreign substance, at least one type of antibody can recognize and bind to this substance. The most important antigen found on influenza viruses is the surface protein hemagglutinin (HA). The goal of vaccination is to stimulate the production of antibodies that can bind to the HA of the upcoming epidemic strain. Production of these antibodies require activation of vaccine-specific CD4 T cells in the lymphoid tissues.

Some of the antibodies generated against a given flu strain (call it Strain X) can also bind to a different strain (call it Strain Y) if the two strains have similar HA structures. Such binding of anti-Strain-X antibodies to Strain Y are termed *cross-reactions*. Because of cross-reactions, the antibodies generated against Strain X provide a certain level of protec-

tion against subsequent invasion of Strain Y. The level of protection is determined by the degree of cross-reactions between anti-Strain-X antibodies and Strain Y, which is in turn determined by the structural similarity of the two strains. For convenience in later discussions, we say that Strain X and Y cross-react or that Strain X and Y are cross-reactive when there are cross-reactions between anti-Strain-X antibodies and Strain Y. The dynamics of cross-reactions will be illustrated in more detail in the next section.

In our vaccination model, we are concerned with the protection mediated by circulating antibodies. Thus, we do not model explicitly cell-mediated immunity that happens after the infection of cells.

2.2 Shape Space and the Antigenic Distance Hypothesis

“Shape space” was initially proposed by Perelson and Oster [17] as a mathematical characterization for antigen-antibody binding. Shape space has played a major role in theoretical immunology since its inception [20, 18, 8, 9, 22]. The basic idea is that each antigen and antibody can be represented as a unique point in a Euclidean space called the shape space; the vector representation of an antigen or antibody is called its shape vector. The *antigenic distance* between an antigen and an antibody is defined as the Euclidean distance between their shape vectors. The binding strength of an antigen-antibody pair is inversely proportional to their antigenic distance. The antigenic distance between two antigens (antibodies) is also defined as the distance between their shape vectors. The structural similarity between two antigens (antibodies) is inversely proportional to their antigenic distance. We say that an antibody and an antigen are complementary to each other if they have the same shape vector. Based on the above definitions, it is clear that binding strength is maximum when an antibody and antigen are complementary to each other. For convenience, we call an antigen (antibody) with shape vector a “antigen a ” (“antibody a ”).

Lapedes and Farber [15] showed that HA of influenza A (H3N2) strains could be robustly characterized using Euclidean shape space with 2-5 dimensions. In view of this and the central role of HA in anti-influenza immune response, we identify each influenza strain with its HA shape vector. That is, if a strain has HA with shape vector a , then we call this strain Strain a . In this context, every influenza A virus subtype will have its own shape space; the influenza B virus will also have its own shape space.

As mentioned in the Chapter 1, the antigenic distance hypothesis was proposed by Smith *et al.* to explain the observed differences in vaccine efficacy among people with different histories of influenza exposures. The hypothesis is based on the dynamics of cross-reactions, which is best illustrated using the 2-dimensional shape space examples in Figure 2-1. We shall explain Figure 2-1 in detail in the following paragraph.

Let us first look at Figure 2-1(a) where an individual is being vaccinated with strain v_1 . Upon challenge by the vaccine strain v_1 (\bullet), antibodies (\times) whose shape vectors are close enough to v_1 are generated. In Smith *et al.*, “close enough” refers to antibodies that lie within the *ball of stimulation* (BOS) of v_1 . Antibodies that lie inside the BOS of v_1 bind to and help eliminate v_1 . The affinity between an antibody and an antigen increases as their antigenic distance decreases. Reciprocally, an antibody a (\star) can bind to antigens that lie within its BOS. In Figure 2-1(b), we assume that the individual is attacked by an epidemic strain e_1 after vaccination. The potency of this epidemic attack is attenuated if the vaccine and epidemic strain are cross-reactive. Two strains are cross-reactive if their BOS intersect. Antibodies that lie within this intersection can bind to both strains. We call these antibodies *cross-reactive antibodies* of the two strains. In this diagram, the cross-reactive antibodies of strain v_1 and e_1 (generated earlier by vaccination with strain v_1) bind to and remove the epidemic viral particles (of strain e_1), thus help preventing infection. Because of cross-reactions, vaccination can be effective even if the vaccine and epidemic strain are not identical, as is the case shown here. Vaccine efficacy of v_1 on e_1 increases as the number of cross-reactive antibodies of the two strains increases. Some of the epidemic viral particles escape neutralization and trigger production of anti- e_1 antibodies (those that lie within the BOS of e_1). Next, let us suppose this individual receives vaccination with strain v_2 in the following year. Vaccine efficacy of v_2 depends on the level of cross-reactions between v_2 and all previously exposed strains (i.e., v_1 and e_1 in this example). In Figure 2-1(c), such cross-reactions are intense. Therefore, before the production of anti- v_2 antibodies can be triggered, the v_2 vaccine particles are eliminated by cross-reactive antibodies generated earlier by v_1 and e_1 . In this case, efficacy of repeated vaccination with strain v_2 is low. In Figure 2-1(d), the level of cross-reactions is low and the vaccine v_2 is successful in triggering significant production of anti- v_2 antibodies. In this case, efficacy of repeated vaccination is high. Note that the degree of cross-reactions depends on the relative positions of v_1 , e_1 , and v_2 in the shape space as well as the radius of the BOS.

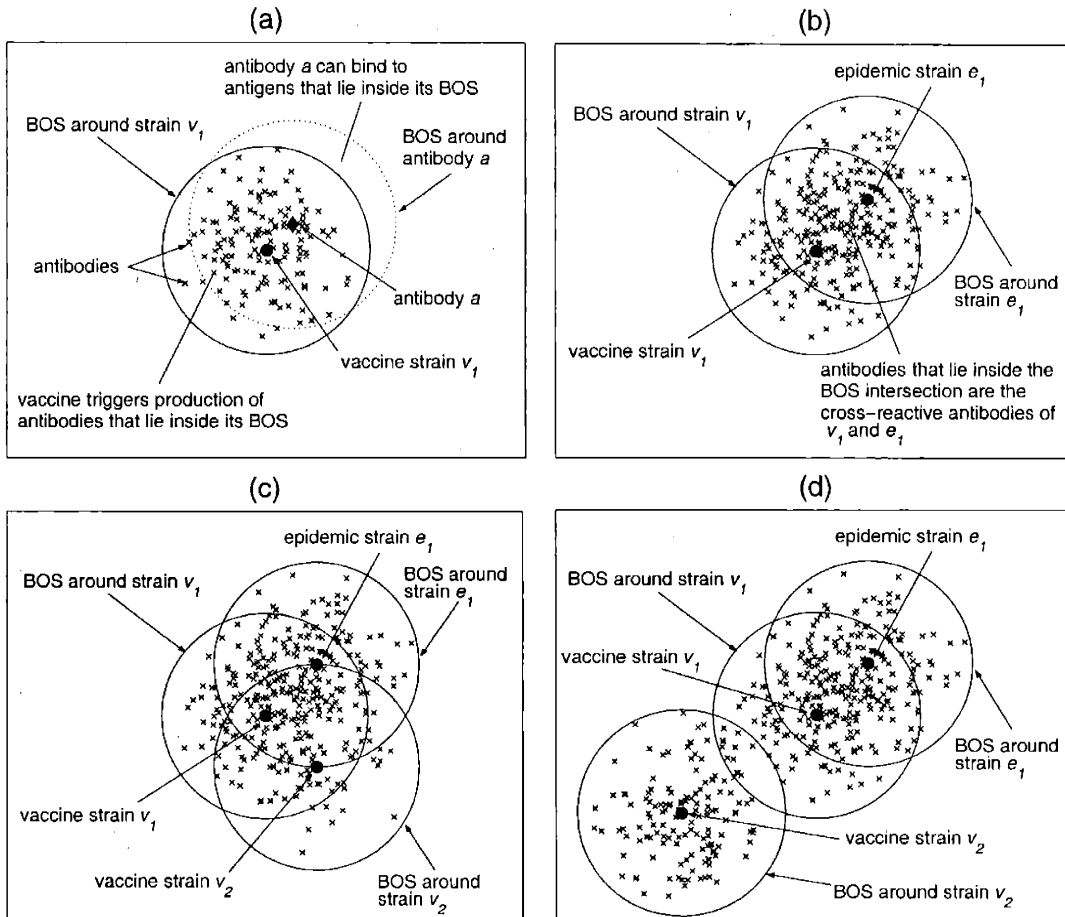


Figure 2-1: The dynamics of cross-reactions.

In summary, the main theme of the antigenic distance hypothesis is as follows: all flu strains (including both vaccines and epidemic viruses) that were seen by the individual before the administration of a vaccine can potentially reduce the efficacy of this vaccine. In other words, the observed low vaccine efficacy in some cases of repeated vaccinations can be attributed to the strong cross-reactions between the vaccine strain and all previously exposed strains.

2.3 Influenza Vaccine Strain Selection as a Discrete-time Control Problem

Having presented the relevant immunological and shape space concepts, we now give a brief description of the influenza vaccine strain selection problem. Our goal is to build a decision framework for choosing an efficacious vaccine strain for an Influenza A subtype or

the Influenza B serotype.

We consider the case of an individual. Let $t = 1, 2, \dots$ index time in years. The shape space is an n -dimensional Euclidean space \mathfrak{R}^n . For mnemonic purpose, we denote the shape space by S instead of \mathfrak{R}^n . Each $s \in S$ corresponds to a physical influenza strain, which can be a vaccine and/or an epidemic strain. The *immune state* of an individual at the beginning of period t is characterized by $x_t = \{x_t(s) \mid s \in S\}$, where $x_t(s)$ is a collective measure of the number of antibodies and B cells that are complementary to strain s at time t . Hence, the function x_t summarizes the individual's humoral immunocompetence against different influenza strains. For convenience in later discussions, we shall simply call $x_t(s)$ the number of antibodies against strain s henceforth. The immune state x_t is constructed from the individual's *immunization history* h_t ; the construction of x_t will be discussed shortly. The immunization history h_t is defined as the sequence of vaccine and epidemic strains to which the individual was exposed between period 1 and period t .

The sequence of events is as follows. At the beginning of each period t , we are given the individual's immunization history h_t and the probability density function of the upcoming epidemic strain e_t . (Since there is a time interval of at least 7 months between vaccine strain decision and the flu season, we model the shape vectors of future epidemic strains as random variables.) Given this information, we choose and administer a vaccine strain $v_t \in S$ for this individual. Vaccination triggers an immune response which prompts the immune state to heighten from x_t to y_t (via antibodies production and immunological memory). That is, for each $s \in S$, the corresponding number of antibodies is increased from $x_t(s)$ to $y_t(s)$. The epidemic strain e_t is then realized and the individual is exposed to the epidemic strain with probability p_e , which is independent of time. A reward $m_t(h_t, v_t, e_t)$ is then incurred. The reward function m_t measures how well the vaccine strain v_t protects this individual (who has immunization history h_t) from the epidemic strain e_t if he is exposed to the epidemic attack. After the realization of e_t , period t ends and period $t + 1$ begins. As a new period begins, the immune state y_t evolves into x_{t+1} due to the immune response triggered by e_t and the natural decay of antibodies. The same cycle then repeats itself over time.

In the next few pages, we will spell out the specific details of our model. Before moving on, we remark that the human immune system is an extremely complex system for modelling. In this thesis, we take a theoretical approach and model the relevant immunodynamics on a macroscopic level. Our goal is to build a model that strikes a reasonable

balance between precision (is the model realistic?) and mathematical tractability (is the model analyzable?). In an attempt to build a realistic model that can capture the immunodynamics described in the previous paragraph, we ask the following questions: (1) How does the epidemic strain sequence $\{e_t\}$ evolve? (2) How does the immune state change after an immune response? (3) How do we measure the efficacy of a vaccine strain? We shall address these questions one by one in the following subsections.

2.3.1 How does the epidemic strain sequence $\{e_t\}$ evolve?

Regarding the first question, we make the following assumption:

Assumption 2.1. *The random epidemic sequence $\{e_t\}$ evolves according to a Markov process that is independent of the vaccine strain selection policy. In particular, $\{e_t\}$ follows a random walk: $e_{t+1} = e_t + u_t$ where $u_t \sim N(d, \sigma^2 \mathbf{I})$.*

The vector d represents the mean *antigenic drift* of the flu virus. As discussed in the last chapter, antigenic drift is a result of random mutations of the flu virus. Our assumption that the epidemic evolution is independent of vaccination policies is supported by Fitch *et al.* [10] and Plotkin *et al.* [19]; both papers have shown that in retrospect, vaccination has not significantly altered the evolutionary course of the influenza virus in the Northern Hemisphere. The assumption of normal distribution and independence of the components of u_t are for tractability purposes. These assumptions do not reflect established experimental results. However, they are not contradictory to the preliminary results observed in the flu shape space map. As more data from the flu shape space map become available, we can refine these assumptions accordingly.

2.3.2 How does the immune state change after an immune response?

This question requires more elaboration. Since the immune state is a measure of the abundance of antibodies and B cells, we restrict $x(s)$ to be nonnegative for all $s \in S$ for every immune state x . Upon challenge by a strain a (which can be a vaccine or epidemic strain), cross-reactions occur between strain a and the existing antibodies in the immune state x . We denote the level of cross-reactions by $c(x, a)$, which we call the *clearance* of a by x . We

define $c(x, a)$ as

$$c(x, a) = \min \left\{ \underbrace{\frac{1}{R} \int_S w(s-a)x(s)ds}_{\text{weight}}, 1 \right\}. \quad (2.1)$$

Let us give a physical interpretation of (2.1). The function $w : S \rightarrow [0, 1]$ is a weight function. Given a strain with antigen a and an antibody s , $w(s-a)$ measures the affinity of the antigen-antibody pair as a function of their antigenic distance $\|s-a\|$. In this paper, we shall let $w(s-a) = e^{-\eta\|s-a\|^2}$ where $\eta > 0$ is some scaling factor. This functional form is consistent with the results in Lapedes and Farber (see Figure 5 in [15]). Note that affinity decreases as antigenic distance increases. Moreover, the weight function is spherically symmetric around 0; this means that same antigenic distances result in the same level of affinity. The constant R in (2.1) is a normalization factor. We defer our discussion on the choice of R until we have described all the dynamics of the immune state. Having explained the notation in (2.1), we can now see that the normalized integral in (2.1) is a weighted number of antibodies in the immune state x , where more weight is given to antibodies that are antigenically closer to a . The clearance is bounded above by 1, meaning that clearance cannot be increased further after the weighted average has risen beyond a certain level (which is controlled by the yet to be explained normalization factor R). This upper bound guarantees that the immune state is always nonnegative. For convenience in later discussions, we denote the weighted average in (2.1) by $\bar{c}(x, a)$:

$$\bar{c}(x, a) = \frac{1}{R} \int_S x(s)w(s-a)ds. \quad (2.2)$$

Now suppose we start with an immune state x_t at time t and we administer a vaccine strain v_t . The immune state then changes from x_t to y_t as follows:

$$y_t(s) = x_t(s) + \underbrace{\kappa_v(1 - c(x_t, v_t))}_{\text{effective dose}} \underbrace{w(s - v_t)}_{\text{weight}}, \quad s \in S. \quad (2.3)$$

(2.3) means that the use of vaccine strain v_t increases the number of antibodies by $\kappa_v(1 - c(x_t, v_t))w(v_t - s)$ for each $s \in S$. We call $\kappa_v(1 - c(x_t, v_t))$ the *effective dose* of the vaccine v_t . The constant κ_v represents the maximum effective dose of a vaccine. The term $1 - c(x_t, v_t)$ admits the following interpretation. After vaccine particles are injected into the body, they are filtered by the existing antibodies that circulate inside the bloodstream. Therefore,

we assume that if a vaccine v_t is administered on an immune state x_t , a fraction $c(x_t, v_t)$ (see (2.1)) of vaccine particles are neutralized by the existing antibodies in x_t . Thus, only $1 - c(x_t, v_t)$ of the total administered vaccine particles are able to trigger immunological memory and antibodies production in the lymphoid tissues. The function $w : S \rightarrow \mathfrak{R}$ in (2.3) is the same as the weight function that we used in (2.1). Therefore, the increase in number is larger for antibodies that are antigenically closer to the vaccine strain v_t . We remark that (2.3) is consistent with the computer model developed by Smith *et al.* [21].

After vaccination, the epidemic strain e_t is realized. With probability p_e , the individual is exposed to the epidemic attack, which prompts the immune state to evolve in a way that is similar to (2.3). Moreover, as a new period begins, we assume that the immune state is discounted by a factor of $\alpha < 1$ due to the natural decay of antibodies. Therefore, y_t evolves into x_{t+1} as follows:

$$x_{t+1}(s) = \alpha \left[y_t(s) + I_{e_t} \kappa_e (1 - c(y_t, e_t)) w(s - e_t) \right], \quad s \in S \quad (2.4)$$

where

$$I_{e_t} = \begin{cases} 1 & \text{if exposed to epidemic strain } e_t \text{ in period } t \text{ (w.p. } p_e); \\ 0 & \text{otherwise (w.p. } 1 - p_e). \end{cases} \quad (2.5)$$

We assume $\kappa_e > \kappa_v$, which means that epidemic attacks can trigger stronger immune responses than vaccinations. This is because vaccines are usually non-replicating viral particles while epidemic viruses are live viruses.

Let us now go back to define the normalization factor R in (2.1). First, we introduce the following definition:

Definition 2.1. *If the immunization history is empty, then the immune state is 0. That is, $x_t = 0$ if $h_t = \emptyset$. We call this the naive immune state.*

We let

$$R = \frac{\kappa_v}{\theta} \int_S w(s - a) w(s - a) ds = \frac{\kappa_v}{\theta} \int_S e^{-\eta \|s - a\|^2} e^{-\eta \|s - a\|^2} ds = \frac{\kappa_v}{\theta} \sqrt{\frac{\pi}{2\eta}}, \quad (2.6)$$

where $0 < \theta \leq 1$ is called the *clearance constant* (note that the above calculation is independent of the choice of a , which can be any finite vector). This definition of R admits the following interpretation. Suppose we have a naive immune state at time t , i.e. $x_t = 0$. Consider

the case where this naive state is boosted by a vaccine strain a and then challenged by an epidemic attack of the same strain a immediately after the vaccination. After vaccination, the immune state x_t is boosted to $y_t(s) = \kappa_v w(s - a)$, $s \in S$, according to the dynamics that we have just described. (2.6) then implies that the weighted average in (2.1) at the time of the ensuing epidemic challenge is $\bar{c}(y_t, a) = \frac{1}{R} \int_S y_t(s) w(s - a) ds = \frac{\kappa_v}{R} \int_S w(s - a) w(s - a) ds = \theta$. Hence, the clearance of the epidemic strain a by y_t is θ , meaning that the vaccine is successful in eliminating $\theta \times 100\%$ of the epidemic viral particles for a naive immune state when the vaccine and epidemic strain match exactly.

We close our discussion on the construction of immune state with a few shorthand notations. We let $D_{v_t} = \kappa_v(1 - c(x_t, v_t))$ in (2.3) and $D_{e_t} = \kappa_e(1 - c(y_t, e_t))$ in (2.4). For $i < t$, we call $\alpha^{t-i} D_{v_i}$ the *residual dose* of v_i at time t ($\alpha^{t-i} D_{e_i}$ is defined analogously). Moreover, we define a function p such that given an element $a \in h_t$, $p(a)$ outputs the time index of a ; for instance, $p(v_{t-1}) = t - 1$. Under such notation, (2.4) and (2.3) can be expressed compactly as

$$x_t(s) = \sum_{a \in h_t} \alpha^{t-p(a)} D_a w(s - a), \quad s \in S, \quad (2.7)$$

$$y_t(s) = \sum_{a \in \{h_t, v_t\}} \alpha^{t-p(a)} D_a w(s - a), \quad s \in S. \quad (2.8)$$

From these equations, we see that the constructed immune state is additive in the sense that an immune state at a given time is the sum of the effect of all previously exposed strains (discounted by the decay factor α). Note that the contribution of each strain in the immunization history depends on its cross-reactions with all previously exposed strains via its effective dose. Also, note that the immune state at any two different times are related additively. For instance, x_t and x_{t+1} are related by the equation

$$x_{t+1}(s) = \alpha \left[x_t(s) + D_{v_t} w(s - v_t) + I_{e_t} D_{e_t} w(s - e_t) \right], \quad s \in S. \quad (2.9)$$

Relationships between immune state at different times such as (2.9) will be invoked frequently throughout the rest of this thesis.

2.3.3 How do we measure the efficacy of a vaccine strain?

Recall that $c(y_t, e_t)$ is the fraction of epidemic viral particles that are cleared by the post-vaccine immune state at time t if the individual is exposed to the epidemic. Hence, we use $c(y_t, e_t)$ as the benchmark for measuring the efficacy of our vaccine choice. To simplify our analysis, however, we use $\bar{c}(y_t, e_t)$ as the measure instead (recall that $c(y_t, e_t) = \min\{\bar{c}(y_t, e_t), 1\}$). That is, we define the reward at time t to be $\bar{c}(y_t, e_t)$. For an N -period problem, the overall objective function is

$$E_{I_{e_1}, \dots, I_{e_N}, e_1, \dots, e_N} \left[\sum_{t=1}^N \beta_t \bar{c}_t(y_t, e_t) \mid e_0 \right] \quad (2.10)$$

where $\beta_t \leq 1$ is the probability that the individual will be alive in period t given that he is alive in period 1 (of course, we have $\beta_t \geq \beta_{t+1}$ for all t).

2.4 Dynamic Programming Formulation

Having described in detail the formulation of our influenza vaccine strain selection model, we can now write down our dynamic program. At each period t , the components of our dynamic program are as follows:

State: Immunization history h_t and the expected upcoming epidemic strain \bar{e}_t

Decision variable: Vaccine strain v_t

Random disturbance: Upcoming epidemic strain e_t

Reward: Approximate clearance of epidemic viral particles by the post-vaccine

immune state, $\bar{c}(y_t, e_t)$, weighted by the probability of survival at time t , β_t

Note that the origin of the shape space is arbitrary and we are only interested in the relative positions of the strains in the shape space. This means that we have the freedom of defining the origin of the shape space to be \bar{e}_t at each iteration of the dynamic program (DP). Hence, without loss of generality, we can reformulate the DP such that the state at time t comprises only the immunization history h_t . In summary, for a time horizon of N periods, the objective is to find the optimal vaccine strain selection policies $\mu_t^*(h_t)$, $t = 1, \dots, N$, to maximize (2.10) subject to (2.1), (2.2), (2.5), (2.6), (2.7), and (2.8). Note that the number of state variables increases as time increases because we accumulate more and more strains in the immunization history as we move forward in time. The expansion of the state vector

results in a high-dimensional continuous state space for the early iterations of the dynamic program. In Chapter 4, we show how we can overcome the curse of dimensionality in this problem and solve this high-dimensional DP.

Chapter 3

Analysis of the 1-Period Problem

3.1 Some Useful Formulas

We begin our analysis with some formulas that will be used repeatedly in the rest of this paper. Substituting (2.7) and (2.8) into (2.2), we have

$$\begin{aligned}\bar{c}(x_t, v_t) &= \frac{1}{R} \int_S \sum_{a \in h_t} \alpha^{t-p(a)} D_a w(s-a) w(s-v_t) ds = \sum_{a \in h_t} \alpha^{t-p(a)} D_a g(a-v_t), \\ \bar{c}(y_t, e_t) &= \frac{1}{R} \int_S \sum_{a \in \{h_t, v_t\}} \alpha^{t-p(a)} D_a w(s-a) w(s-e_t) ds = \sum_{a \in \{h_t, v_t\}} \alpha^{t-p(a)} D_a g(a-e_t),\end{aligned}\tag{3.1}$$

where $g(a-b) = \frac{1}{R} \int_S w(s-a) w(s-b) ds$. Since $w(s-a) = e^{-\eta \|s-a\|^2}$ and $R = \frac{\kappa_v}{\theta} \sqrt{\frac{\pi}{2\eta}}$ (see (2.6)), it can be easily verified that $g(a-b) = \frac{\theta}{\kappa_v} e^{-\eta \|a-b\|^2/2}$. Note that (3.1) implies

$$\bar{c}(y_t, e_t) = \bar{c}(x_t, e_t) + D_{v_t} g(e_t - v_t)\tag{3.2}$$

and other similar additive relations which we shall use frequently throughout the thesis. Substituting (3.1) into the definitions of the effective doses, we have

$$\begin{aligned}D_{v_t} &= \kappa_v (1 - c(x_t, v_t)) = \kappa_v \left(1 - \sum_{a \in h_t} \alpha^{t-p(a)} D_a g(a-v_t) \right)^+, \\ D_{e_t} &= \kappa_e (1 - c(y_t, e_t)) = \kappa_e \left(1 - \sum_{a \in \{h_t, v_t\}} \alpha^{t-p(a)} D_a g(a-e_t) \right)^+.\end{aligned}\tag{3.3}$$

By Assumption 2.1, the conditional pdf of e_t is

$$f(e_t | e_{t-1}) = \left(\frac{1}{\sqrt{2\pi}\sigma} \right)^n e^{-\|e_t - \bar{e}_t\|^2 / 2\sigma^2}, \quad s \in S, \quad (3.4)$$

where $\bar{e}_t = e_{t-1} + d$. Therefore, if s is independent of e_{t+k} for all $k \geq 0$, we have

$$E_{e_{t+k}} [g(s - e_{t+k}) | e_{t-1}] = \frac{\theta}{\kappa_v} B_{k+1}^{n/2} e^{-\eta B_{k+1} \|s - \bar{e}_t - kd\|^2 / 2} \quad (3.5)$$

where $B_k = \frac{1}{1+k\eta\sigma^2}$. Combining (3.1) and (3.5), we see that

$$E_{e_{t+k}} [\bar{c}(y_t, e_{t+k}) | e_{t-1}] = \sum_{a \in \{h_t, v_t\}} \alpha^{t-p(a)+k} D_a E_{e_{t+k}} [g(a - e_{t+k}) | e_{t-1}] \quad (3.6)$$

admits a closed form expression.

3.2 Reformulation of the 1-Period Problem

We now begin the analysis of the first DP iteration, which is the same as a 1-period problem (except for the difference in β_t). The objective function is simply the expected reward at time N :

$$\beta_N E_{e_N} [\bar{c}(y_N, e_N) | e_{N-1}] = \beta_N \underbrace{E_{e_N} [\bar{c}(x_N, e_N) | e_{N-1}]}_{\text{residual immunity in } x_N} + \beta_N \underbrace{D_{v_N} E_{e_N} [g(v_N - e_N) | e_{N-1}]}_{\text{immunity boosted by } v_N}. \quad (3.7)$$

The equality holds because of (3.2). Since the first term in (3.7) is independent of v_N , it suffices to maximize just the second term. Moreover, since we are maximizing the function, we are only interested in choices of v_N that give positive D_{v_N} (we can always achieve $D_{v_N} = 0$ by not using any vaccine). Using (3.3) and (3.5), we can easily verify that the 1-period problem is the same as $\max_{v_N} F_N(v_N)^+$ where

$$F_N(v_N) = \left(1 - \frac{\theta}{\kappa_v} \sum_{a \in h_N} \alpha^{N-p(a)} D_a e^{-\eta \|a - v_N\|^2 / 2} \right) e^{-\eta B_1 \|\bar{e}_N - v_N\|^2 / 2}. \quad (3.8)$$

Note that an optimal solution to $\max_{v_N} F_N(v_N)$ is also an optimal solution to $\max_{v_N} F_N(v_N)^+$. We choose to work with the former formulation because the objective function is continuous. Let v_N^* be the optimal solution to this maximization problem. Clearly, if $h_N = \emptyset$

(i.e., no immunization history), then $v_N^* = \bar{e}_N$ and the 1-period problem is trivial. So let us assume $h_N \neq \emptyset$ for the rest of our analysis. Also, we assume that v_N^* is unique and \bar{e}_N is not a stationary point of F_N (which can be achieved by slightly perturbing the data h_N if necessary). This means that $\nabla F_N(\bar{e}_N) \neq 0$. We are interested in obtaining analytical insights on the behavior of v_N^* . Unfortunately, v_N^* does not admit a closed-form solution and its properties is not easily analyzable. Our goal in this section is therefore to find a close approximation to v_N^* that is analytically tractable. Let us begin by visualizing the geometry of this problem. The term inside the parentheses in (3.8) shows that there are penalties for placing v_N near the points in h_N : a penalty is radially distributed around each point in h_N . These penalties arise from the cross-reactions of the vaccine strain v_N and each strain in the immunization history h_N . On the other hand, there is a reward for placing v_N near \bar{e}_N : a reward is radially distributed around \bar{e}_N . This reward stems from the expected cross-reactions between the vaccine strain and the upcoming epidemic strain. Hence, v_N^* is the point that strikes an optimal balance between the penalties and the reward. We proceed as follows.

We begin with a simple example where h_N has only one element, say v_{N-1} . That is, $h_N = v_{N-1}$. From the symmetry of the problem shown in Figure 3-1, we immediately see




Figure 3-1: Symmetry in the example where $h_N = v_{N-1}$.

that $v_N^* = \bar{e}_N - \rho^* \frac{v_{N-1} - \bar{e}_N}{\|v_{N-1} - \bar{e}_N\|}$ for some $\rho^* > 0$. Therefore, we only need to know ρ^* in order to locate v_N^* . By expressing v_N in the form of $\bar{e}_N - \rho \frac{v_{N-1} - \bar{e}_N}{\|v_{N-1} - \bar{e}_N\|}$ in (3.8), we can transform the multi-dimensional problem $\max_{v_N} F_N(v_N)$ into a one-dimensional problem

$$\max_{\rho > 0} \left(1 - \frac{\theta}{\kappa_v} \alpha D_{v_{N-1}} e^{-\eta(\rho + \|v_{N-1} - \bar{e}_N\|)^2/2} \right) e^{-\eta B_1 \rho^2/2}. \quad (3.9)$$

However, this optimization problem does not have a closed-form solution neither. Nevertheless, as we will show later in our analysis, we can extract some important relationships between the optimal solution and the parameters of the problem.

Let us now consider the case where h_N has an arbitrary number of points. Our next step is to approximate the problem $\max_{v_N} F_N(v_N)$ with a one-dimensional problem that resembles (3.9). In the previous example, we assume $h_N = v_{N-1}$ and then deduce the direction in which v_N^* lies (relative to \bar{e}_N). Let us now reverse the scenario and assume that we know the optimal solution v_N^* . In particular, we know the optimal direction $\frac{v_N^* - \bar{e}_N}{\|v_N^* - \bar{e}_N\|}$, denoted by \bar{v}_N^* . Then by symmetry, we see that it is possible to construct a 1-period problem that has (i) the same optimal solution v_N^* , and (ii) an immunization history h_N that has only one point, denoted by $r_{r_N}^* = \bar{e}_N - \gamma_N^* \bar{v}_N^*$ for some $\gamma_N^* > 0$. That is, there exist $D_{r_N}^*$ and r_N^* (equivalently, γ_N^*) such that

$$v_N^* = \arg \max_{v_N} \left(1 - \frac{\theta}{\kappa_v} D_{r_N}^* e^{-\eta \|r_N^* - v_N\|^2/2} \right) e^{-\eta B_1 \|v_N - \bar{e}_N\|^2/2}, \quad (3.10)$$

If we can determine $D_{r_N}^*$ and r_N^* without knowing v_N^* , then we can transform $\max_{v_N} F_N(v_N)$ into a one-dimensional problem that is much more amenable to analysis. Note that (3.10) implies that in our search for v_N^* , all relevant information in h_N are compressed into $D_{r_N}^*$ and r_N^* . That is, $D_{r_N}^*$ and r_N^* contains necessary and sufficient information for finding v_N^* . We call the transformation of h_N to $D_{r_N}^*$ and r_N^* the *history reduction*. Recall that our goal in this chapter is to find a good approximation to v_N^* that is easily analyzable. In view of the 1-dimensional problem transformation under history reduction, we shall approximate v_N^* by obtaining a good approximation to $D_{r_N}^*$ and r_N^* . Before we do this, however, let us first examine the behavior of v_N^* under the setting of the transformed problem in (3.10).

3.3 Properties of the Optimal Solution to the Transformed 1-D Problem

Analogous to (3.9), the problem in (3.10) can be reduced to a 1-dimensional problem:

$$\max_{\rho > 0} \left(1 - \frac{\theta}{\kappa_v} D_{r_N}^* e^{-\eta(\rho + \gamma_N^*)^2/2} \right) e^{-\eta B_1 \rho^2/2}. \quad (3.11)$$

Let ρ^* be the optimal solution to (3.11). As mentioned before, ρ^* does not admit a closed-form solution. Our next step is to analyze how ρ^* varies as a function of the parameters in (3.11). In particular, we are interested in eliciting the dependence of ρ^* on γ_N^* , which is the distance between r_N^* and \bar{e}_N , while holding other parameters constant. For this purpose, we nondimensionalize the objective function by letting $q = \sqrt{\frac{\eta}{2}}\rho$, $\tau = \sqrt{\frac{\eta}{2}}\gamma_N^*$, and $C = \frac{\theta}{\kappa_v}D_{r_N}^*$. The nondimensionalized objective function is therefore

$$H(q) = (1 - Ce^{-(q+\tau)^2})e^{-B_1q^2} \quad (3.12)$$

where $C > 0$ and $0 < B_1 \leq 1$. For any $\tau > 0$, H is bounded and has only one local maximum in $q > 0$. Hence, the global maximum of H can be found by solving the equation

$$\frac{dH}{dq} = 2e^{-B_1q^2-(q+\tau)^2} \underbrace{\left(C(B_1q + q + \tau) - B_1qe^{(q+\tau)^2} \right)}_{f(q,\tau)} = 0 \quad (3.13)$$

in $q > 0$. Let ϕ be the unique function that satisfies $f(\tau, \phi(\tau)) = 0$ for all $\tau > 0$. The optimal solution v_N^* can now be expressed as

$$v_N^* = \bar{e}_N - \sqrt{\frac{2}{\eta}}\phi\left(\sqrt{\frac{\eta}{2}}\gamma_N^*\right) \frac{r_N^* - \bar{e}_N}{\|r_N^* - \bar{e}_N\|}. \quad (3.14)$$

Theorem 3.1 and 3.2 below give us some important information on how ϕ varies with τ , C , and B_1 . For ease of presentation, we assume $C \neq \frac{B_1}{1+B_1}$ and include $\tau = 0$ in the domain of ϕ , albeit $\tau > 0$ in our problem. The proofs of Theorem 3.1 and 3.2 can be found in the Appendix.

Theorem 3.1. $\phi \in C^\infty([0, \infty))$ and $\lim_{\tau \rightarrow \infty} \phi(\tau) = 0$. The (C, B_1) parameter plane can be partitioned into 4 regions as follows based on the general behavior of ϕ therein:

Region 1: $C < \frac{B_1}{1+B_1}$

Region 2: $\frac{B_1}{1+B_1} < C < \frac{B_1}{1+B_1}e^{\frac{1}{2(1+B_1)}}$

Region 3: $\frac{B_1}{1+B_1}e^{\frac{1}{2(1+B_1)}} \leq C < \frac{B_1}{1+B_1}e^{\frac{3}{2(1+2B_1)}}$

Region 4: $C \geq \frac{B_1}{1+B_1}e^{\frac{3}{2(1+2B_1)}}$


The behavior of ϕ in each region is summarized in Table 3.1. Furthermore, $\tau_{max} \leq \sqrt{\frac{1}{2}}$ and $\tau_{inflect} \leq \sqrt{\frac{3}{2}}$ for all $C > 0$ and $0 < B_1 \leq 1$ that satisfy $C \neq \frac{B_1}{1+B_1}$.

Theorem 3.2. Suppose we vary C while holding B_1 or $\frac{C}{B_1}$ fixed. Let ϕ_C be the function ϕ

Region	General behavior of ϕ
1	$\phi(0) = 0$; $\phi(\tau)$ is strictly concave and strictly increasing on $[0, \tau_{max})$; $\phi(\tau)$ attains a global maximum at $\tau = \tau_{max}$; $\phi(\tau)$ is strictly concave and strictly decreasing on $(\tau_{max}, \tau_{inflect})$; $\phi(\tau)$ has a point of inflection at $\tau = \tau_{inflect}$; $\phi(\tau)$ is strictly convex and strictly decreasing on $(\tau_{inflect}, \infty)$.
2	Same as Region 1 except $\phi(0) > 0$.
3	$\phi(0) > 0$; $\frac{d\phi}{d\tau}(0) = 0$ if $C = \frac{B_1}{1+B_1} e^{\frac{1}{2(1+B_1)}}$, $\frac{d\phi}{d\tau}(0) < 0$ otherwise; $\phi(\tau)$ is strictly concave and strictly decreasing on $(0, \tau_{inflect})$; $\phi(\tau)$ has a point of inflection at $\tau = \tau_{inflect}$; $\phi(\tau)$ is strictly convex and strictly decreasing on $(\tau_{inflect}, \infty)$.
4	$\phi(0) > 0$; $\frac{d^2\phi}{d\tau^2}(0) = 0$ if $C = \frac{B_1}{1+B_1} e^{\frac{3}{2(1+2B_1)}}$, $\frac{d^2\phi}{d\tau^2}(0) > 0$ otherwise; $\phi(\tau)$ is strictly convex and strictly decreasing for all $\tau > 0$.

Table 3.1: General behavior of ϕ in the four different parameter regions in Theorem 3.1.

parameterized by C . If $\hat{C} > 0$ and $\tilde{C} > 0$ satisfy $\hat{C} < \tilde{C}$, then $\phi_{\hat{C}}(\tau) \leq \phi_{\tilde{C}}(\tau)$ for all $\tau \geq 0$.



3.4 Construction of the Approximate Optimal Solution

We will complete our analysis of the 1-period problem by constructing a good and robust approximate optimal solution. As mentioned before, we will do this by finding a good approximation to D_{TN}^* and r_N^* . By comparing $F_N(v_N)$ in (3.8) and the objective function in (3.10), we see that D_{TN} and r_N will be good approximations to D_{TN}^* and r_N^* , respectively,

if

$$\sum_{a \in h_N} \alpha^{N-p(a)} D_a e^{-\eta \|a - v_N\|^2/2} \approx D_{r_N} e^{-\eta \|r_N - v_N\|^2/2} \quad \text{for } v_N \text{ near } v_N^*. \quad (3.15)$$

Consider expanding the left-hand-side of (3.15) as follows:

$$\begin{aligned} & \sum_{a \in h_N} \alpha^{N-p(a)} D_a e^{-\eta \|a - v_N\|^2/2} \\ &= e^{-\eta \|\bar{e}_N - v_N\|^2/2} \sum_{a \in h_N} \alpha^{N-p(a)} D_a e^{-\eta \|a - \bar{e}_N\|^2/2} e^{\eta (v_N - \bar{e}_N)^T (a - \bar{e}_N)} \\ &= e^{-\eta \|\bar{e}_N - v_N\|^2/2} \sum_{a \in h_N} \alpha^{N-p(a)} D_a e^{-\eta \|a - \bar{e}_N\|^2/2} \sum_{k=0}^{\infty} \frac{1}{k!} \left(\eta (v_N - \bar{e}_N)^T (a - \bar{e}_N) \right)^k. \end{aligned} \quad (3.16)$$

The second equality is obtained by writing $a - v_N$ as $(a - \bar{e}_N) - (v_N - \bar{e}_N)$. Next, we let

$$r_N = \bar{e}_N - \gamma_N \frac{G_N}{\|G_N\|} \quad (3.17)$$

where γ_N and G_N are to be determined. We then substitute (3.17) into the right-hand-side of (3.15) and expand the result as follows:

$$\begin{aligned} & D_{r_N} e^{-\eta \|r_N - v_N\|^2/2} \\ &= D_{r_N} e^{-\eta \|\bar{e}_N - \gamma_N \frac{G_N}{\|G_N\|} - v_N\|^2/2} \\ &= e^{-\eta \|\bar{e}_N - v_N\|^2/2} D_{r_N} e^{-\eta \gamma_N^2/2} e^{-\eta \gamma_N (v_N - \bar{e}_N)^T \frac{G_N}{\|G_N\|}} \\ &= e^{-\eta \|\bar{e}_N - v_N\|^2/2} D_{r_N} e^{-\eta \gamma_N^2/2} \sum_{k=0}^{\infty} \frac{1}{k!} \left(-\eta \gamma_N (v_N - \bar{e}_N)^T \frac{G_N}{\|G_N\|} \right)^k. \end{aligned} \quad (3.18)$$

We can now match (3.16) and (3.18) (i.e., the left- and right-hand-side of (3.15)) up to the second order by choosing

$$G_N = \sum_{a \in h_N} \alpha^{N-p(a)} (\bar{e}_N - a) D_a e^{-\eta \|a - \bar{e}_N\|^2/2}, \quad (3.19)$$

$$\gamma_N = \|G_N\| \left(\sum_{a \in h_N} \alpha^{N-p(a)} D_a e^{-\eta \|a - \bar{e}_N\|^2/2} \right)^{-1}, \quad (3.20)$$

$$D_{r_N} = e^{\eta \gamma_N^2/2} \frac{\|G_N\|}{\gamma_N}. \quad (3.21)$$

Since we assume $h_N \neq \emptyset$, the term inside the parentheses in (3.20) is nonzero and γ_N is well-defined. Interestingly, we notice from (3.17) and (3.21) that r_N and D_{r_N} are both functions of only γ_N and G_N . From (3.17), we see that γ_N is the distance between r_N and \bar{e}_N and $-\frac{G_N}{\|G_N\|}$ is the unit vector that points to r_N from \bar{e}_N . Moreover, it can be easily verified that $\frac{\nabla F_N(\bar{e}_N)}{\|\nabla F_N(\bar{e}_N)\|} = \frac{G_N}{\|G_N\|}$, so $G_N \neq 0$ (otherwise \bar{e}_N is a stationary point of F_N which violates our assumption) and r_N is well-defined.

We now have r_N and D_{r_N} as approximations to r_N^* and $D_{r_N}^*$. By substituting these approximations into the transformed problem (3.10), we are in essence approximating F_N by

$$F_N^h(v_N) = \left(1 - \frac{\theta}{\kappa_v} D_{r_N} e^{-\eta \|r_N - v_N\|^2/2}\right) e^{-\eta B_1 \|v_N - \bar{e}_N\|^2/2}. \quad (3.22)$$

Let v_N^h be the optimal solution to $\max_{v_N} F_N^h(v_N)$; we will use v_N^h as an approximation to v_N^* . As discussed before, we have

$$v_N^h = \bar{e}_N - \sqrt{\frac{2}{\eta}} \phi\left(\sqrt{\frac{\eta}{2}} \gamma_N\right) \frac{r_N - \bar{e}_N}{\|r_N - \bar{e}_N\|} = \bar{e}_N + \sqrt{\frac{2}{\eta}} \phi\left(\sqrt{\frac{\eta}{2}} \gamma_N\right) \frac{G_N}{\|G_N\|}. \quad (3.23)$$

Notice that (3.23) means that we are approximating the optimal direction \bar{v}_N^* with $\frac{G_N}{\|G_N\|}$. Since $\frac{\nabla F_N(\bar{e}_N)}{\|\nabla F_N(\bar{e}_N)\|} = \frac{G_N}{\|G_N\|}$, we are essentially approximating the optimal direction \bar{v}_N^* with the steepest ascent direction of F_N at \bar{e}_N . We use the superscript “h” in v_N^h because we will use v_N^h as a time-independent heuristic policy in our dynamic program in the next chapter. The performance of v_N^h in the 1-period problem will be demonstrated shortly; the performance of the heuristic policy in the multi-period problem will be tested in the next chapter when we solve the DP numerically.

3.5 Testing the Performance of Different Policies in the 1-Period Problem

In this section, we test the performance of the heuristic policy, the follow policy (denoted by $v_N^f = \bar{e}_N$), and the no-vaccine policy (denoted by $v_N^n = \infty$) in the 1-period problem. Before we do so, let us first summarize our results so far. Our goal in this chapter is to obtain an analytically tractable approximation to the optimal solution of the 1-period problem. We showed that a general 1-period problem with an arbitrary immunization history h_N can be reformulated as a special 1-period problem where h_N has only one representative strain

r_N^* whose effective dose is $D_{r_N}^*$ (see (3.10)). We then extract some important properties of the optimal solution to the transformed 1-dimensional 1-period problem (see (3.11)) in Theorem 3.1 and 3.2. Finally, we obtained an approximation to v_N^* , denoted by v_N^h (see (3.23)), by approximating r_N^* and $D_{r_N}^*$ with r_N and D_{r_N} . We now close this chapter by testing the performance of v_N^h , v_N^f , and v_N^n and how well $F_N^h(v_N^h)$ (see (3.22)) approximates $F_N(v_N^*)^+$. To this purpose, we construct a simulation as follows:

1. Without loss of generality, set $\kappa_e = 1$ and $\eta = 5$. $\eta = 5$ is convenient because the value of the weight function $w(s)$ drops to negligible level when $\|s\| \geq 1$. See Figure 3-2.




Figure 3-2: Plot of the weight function for $\eta = 5$. The weight drops to negligible level when $\|s\| \geq 1$.

2. Since we do not have concrete data to help us set parameter values, we generate parameters randomly according to Table 3.2. We set $\sigma = c_v \frac{\|d\|}{\sqrt{n}}$ where c_v measures the degree of variation in the epidemic evolution (which is a random walk).
3. Compute

$$R^h = \frac{F_N(v_N^h)^+}{F_N(v_N^*)^+} \quad \text{and} \quad R^f = \frac{F_N(v_N^f)^+}{F_N(v_N^*)^+} \quad (3.24)$$

which measures the performance of the heuristic and follow policy. Note that $R^n = F_N(v_N^n)^+/F_N(v_N^*)^+ = 0$ for the no-vaccine policy. Also, we compute $\frac{F_N^h(v_N^h)}{F_N(v_N^*)^+}$ which measures how well $\max_{v_N} F_N^h(v_N)$ (the transformed problem under history reduction) approximates $\max_{v_N} F_N(v_N)^+$ (the original 1-period problem). This will lead us to devise an efficient algorithm to approximate the DP in the next chapter.

Parameter	Distribution	Interpretation
n (Integer)	$U(2, 5)$	Dimension of shape space is between 2 and 5
M (Integer)	$U(1, 10)$	1 to 10 years of immunization history
κ_v	$U(0.3, 0.6)$	Vaccines are 30 to 60% as potent as epidemic attacks
α	$U(0.3, 0.6)$	Immune state decay rate is between 0.3 and 0.6.
θ	$U(0.4, 0.8)$	Vaccine removes 40 to 80% of epidemic viral particles in a naive immune state if vaccine strain matches epidemic strain
p_e	$U(0.3, 0.6)$	Probability of epidemic exposure is between 0.3 and 0.6
$\ d\ $	$U(0, 0.5)$	Magnitude of mean antigenic drift
c_v	$U(0.1, 1)$	$\sigma = c_v \frac{\ d\ }{\sqrt{n}}$

Table 3.2: Randomly generated parameter values for testing the accuracy of v_N^h in the 1-period problem.

The simulation is repeated 800,000 times to test the performance of the three policies against that of the optimal policy. The frequency plots of R^h and R^f are shown in Figure 3-3. From the plot of R^h , we see that the heuristic policy is an excellent approximation to the optimal policy for all the randomly generated combinations of parameter values in the simulations. In this sense, the heuristic policy is a robust approximation to the optimal policy. The frequency plot of R^f shows that the follow policy is less robust, yet still performs quite well in many cases.




Figure 3-3: Frequency plots of R^h and R^f .

Note that the performance measures that we just used (i.e., R^h and R^f) compare only the immunity boosted by v_N . Recall that F_N is the part of the 1-period objective function that depends on v_N (see (3.7)). In practice, however, we are most interested in the overall immunocompetence of the individual and we should therefore include the residual immunity term (see (3.7)) in the performance measures. To this end, let us compare the performance

measures

$$r^h = \frac{E_{e_N}[\bar{c}(y_N^h, e_N) | e_{N-1}]}{E_{e_N}[\bar{c}(y_N^*, e_N) | e_{N-1}]}, \quad r^f = \frac{E_{e_N}[\bar{c}(y_N^f, e_N) | e_{N-1}]}{E_{e_N}[\bar{c}(y_N^*, e_N) | e_{N-1}]}, \quad r^n = \frac{E_{e_N}[\bar{c}(y_N^n, e_N) | e_{N-1}]}{E_{e_N}[\bar{c}(y_N^*, e_N) | e_{N-1}]}, \quad (3.25)$$

where y_N^* , y_N^h , y_N^f , and y_N^n are the post-vaccine immune states under the optimal, heuristic, follow, and no-vaccine policy (note that $y_N^n = x_N$). The frequency plots of r^h , r^f , and r^n are shown in Figure 3-4. Also, we are interested in knowing how the performance of policies change as a function of residual immunity. Let R^{res} be the ratio of residual immunity to vaccine-boosted immunity under the optimal policy. That is, let $R^{res} = \frac{E_{e_N}[\bar{c}(x_N, e_N) | e_{N-1}]}{\theta B_1^{n/2} F_N(v_N^*)^+}$. In Figure 3-5, we plot R^h and R^f against R^{res} . With the effect of

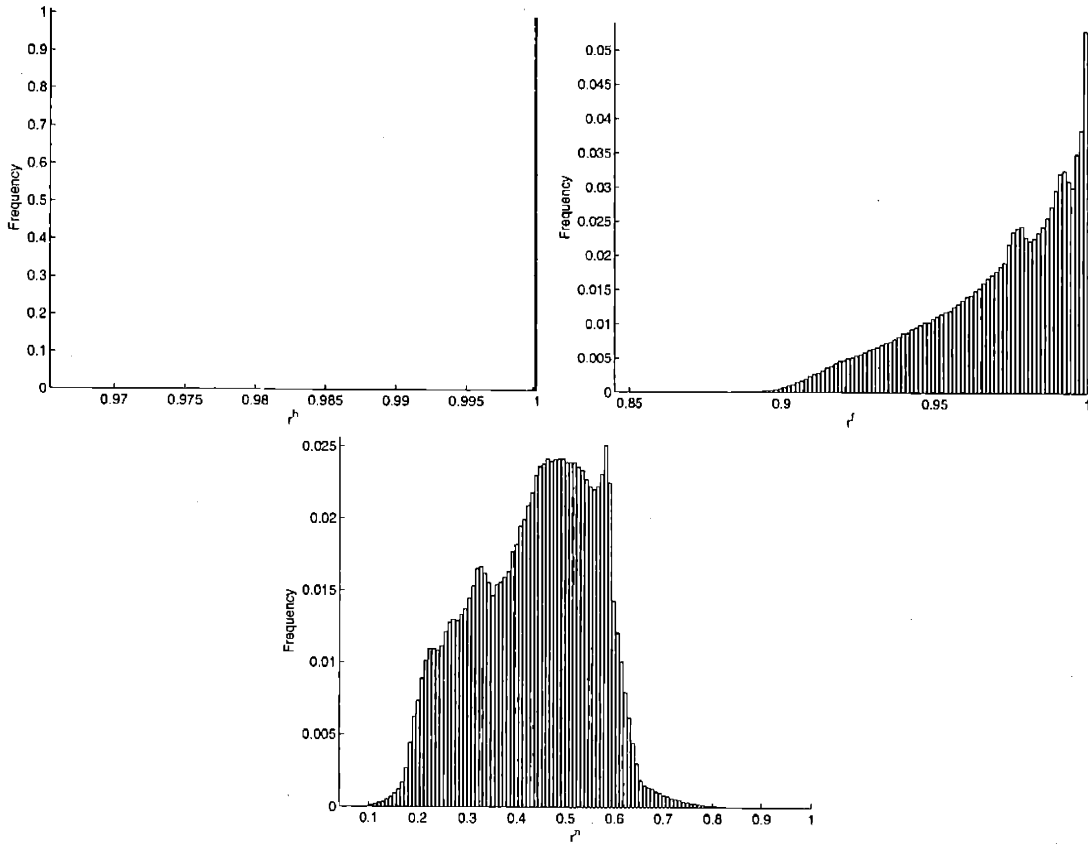


Figure 3-4: Frequency plots of r^h , r^f , and r^n .

residual immunity now taken into account, let us reexamine the performance of the three policies. From the frequency plot of r^h in Figure 3-4, we see that the performance of the heuristic policy is still robust. Indeed, as shown in the plot of R^h in Figure 3-5, its performance is very insensitive to the level of residual immunity. Next, from the frequency




Figure 3-5: Scatterplot of R^h and R^f versus the level of residual immunity.

plot of r^f , we see that the follow policy is performing very well after the residual immunity term is taken into account. This is because in cases where the follow policy provides only relatively low level of boosted immunity (i.e., low R^f), the attenuation of vaccine efficacy implies a high level of residual immunity. Thus, the low vaccine efficacy is compensated by a high level of residual immunity. This view is confirmed by the plot of R^f versus R^{res} in Figure 3-5, which shows that vaccine efficacy under the follow policy deteriorates as the level of residual immunity increases. Although the performance of the follow policy is far less spectacular than that of the heuristic, it is fair to say that the follow policy is still a good and robust suboptimal policy. Finally, the frequency plot of r^n shows that the performance of the no-vaccine policy varies widely depending on the parameter values and the immunization history. In this sense, the no-vaccine policy is neither good nor robust.

Next, we turn our attention to the frequency plot of the ratio $\frac{F_N^h(v_N^h)}{F_N(v_N^*)^+}$ in Figure 3-6. The plot shows that $F_N^h(v_N^h)$ is an excellent approximation to $F_N(v_N^*)^+$. This fact will help us identify an efficient way to approximate the value function of the first DP iteration in the next chapter. When we solve the DP numerically in the next chapter, we will focus on the case $n = 2$ (i.e., the dimension of the shape space is $n = 2$). In Figure 3-7 to 3-9, we show the frequency plots of R^h , R^f , r^h , r^f , r^n , and $\frac{F_N^h(v_N^h)}{F_N(v_N^*)^+}$ for simulation trials that have $n = 2$. Comparing the results in Figure 3-7 to 3-9 with that in Figure 3-3 to 3-6, we see that the performance of the three policies are not greatly affected by the value of n . However, comparing the frequency plot of r^f in Figure 3-4 to that in Figure 3-8, we can deduce that the follow policy performs worse when n is larger. This is because when the dimension of the shape space is higher, it is easier for the optimal vaccine strain to avoid cross-reactions




Figure 3-6: Frequency plot of $\frac{F_N^h(v_N^h)}{F_N^*(v_N^*)}$.

with past strains.




Figure 3-7: Frequency plots of R^h and R^f for $n = 2$.




Figure 3-8: Frequency plots of r^h , r^f , and r^n for $n = 2$.




Figure 3-9: Frequency plot of $\frac{F_N^h(v_N^h)}{F_N(v_N^h)}$ for $n = 2$.

Chapter 4

Numerical Methods for DP Approximation

4.1 Parameter Values for DP

In the previous chapter, we analyzed the 1-period problem, which is the same as the first iteration of our dynamic program. In addition, we have shown that our heuristic policy is an excellent and robust near-optimal policy in the 1-period problem. Our goal in this chapter is to solve an N -period problem numerically and then compare the performance of the heuristic policy, the follow policy, the no-vaccine policy to that of the optimal policy. We restrict our attention to a time horizon of $N = 10$ years and a shape space dimension of $n = 2$. Specifically, we consider the case of a 65-year-old individual who will be receiving flu vaccines for the next $N = 10$ years. According to the United States Life Tables in 2001, the values of β_t , $t = 1, \dots, N$, for this individual are shown in Table 4.1. (Recall that β_t is the probability that the individual will be alive in period t .) As mentioned before, we do not have concrete data that allow us to choose specific values for the rest of our model parameters. For this reason, we will solve the DP for the parameter values shown in Table 4.2. We will run the DP for all 216 combinations of parameter values given in

t	1	2	3	4	5	6	7	8	9	10
β_t	0.984	0.967	0.949	0.929	0.908	0.885	0.862	0.837	0.810	0.782

Table 4.1: Values of β_t , $t = 1, \dots, N$, obtained from the U.S. Life Table, 2001.

Parameter	Values	Interpretation
κ_v	{0.3, 0.6}	Vaccines are 30 or 60% as potent as epidemic attacks.
α	{0.3, 0.6}	Immune state decay rate is 0.3 or 0.6.
θ	{0.4, 0.6, 0.8}	Clearance constant is 0.4, 0.6, or 0.8
p_e	{0.3, 0.6}	Probability of epidemic exposure is between 0.3 or 0.6.
$\ d\ $	{0.1, 0.4, 0.8}	The magnitude of mean antigenic drift is 0.1, 0.4 or 0.8
c_v	{0.1, 0.5, 1}	$\sigma = c_v \frac{\ d\ }{\sqrt{n}}$

Table 4.2: DP parameter values

Table 4.2, which we shall index as Case 1 to Case 216. The combination of parameter values in each case can be found in Table A.1 in the Appendix.

4.2 Breaking the Curse of Dimensionality via Features Extraction

As mentioned at the end of Chapter 2, this DP is difficult to solve because the dimension of the state vector increases as time increases. In the first DP iteration, there are $N - 1 = 9$ years of immunization history. Therefore, h_N has $2(N - 1)n = 36$ real variables (one year of history is consisted of one vaccine strain and one epidemic strain; each strain has a shape vector of dimension $n = 2$). Since a closed form expression of the value function is not available, we resort to finding a closed form approximation to the value function. However, learning the value function in a 36-dimensional continuous state space is a daunting task. The computational requirement for solving this DP (approximately) is thus prohibitive under the usual straightforward DP implementation. In what follows, we develop an efficient algorithm to get around the curse of dimensionality in our DP. Our approach is to first identify a small number of “features” that can effectively capture the information contained in the state vector. We then approximate the value function by a low order polynomial of these features. This approach allows us to learn the value function from a reduced feature space whose dimension is much smaller than that of the original state space. We proceed as follows.

4.3 Reward-To-Go Approximation at the 1st DP Iteration

Let μ_t^h , μ_t^f , and μ_t^n denote the heuristic, follow, and no-vaccine policy at period t . Recall that $\mu_t^h(h_t) = v_t^h$ and $\mu_t^f(h_t) = \bar{e}_t$, where v_t^h is given in (3.23) (with N replaced by t). Since all three of these policies are time-independent, we shall drop their time subscripts henceforth. Let μ_N^* be the optimal policy at time N . Let J_N^* , J_N^h , J_N^f , and J_N^n be the the optimal, heuristic, follow, and no-vaccine reward-to-go at time N . Our first step is to extract the desirable features for approximating J_N^* , J_N^h , and J_N^f (J_N^n has a closed form solution so no approximation is necessary). To do this, we need to obtain information on how J_N^* , J_N^h , and J_N^f depend on the state h_N . By (3.7), the reward-to-go function at time N under a given policy $\mu \in \{\mu_N^*, \mu^h, \mu^f\}$, denoted by J_N , is

$$J_N(h_N) = \beta_N E_{e_N}[\bar{c}(x_N, e_N) \mid e_{N-1}] + \beta_N \theta B_1^{n/2} F_N(\mu(h_N))^+. \quad (4.1)$$

As shown in the last chapter, the first term on the right-hand-side is the residual immunity in the immune state x_N and is independent of the policy μ . So let us look at how $F_N(\mu(h_N))^+$ depends on the state h_N . Let $\delta_N = \frac{\|G_N\|}{\gamma_N} = \sum_{a \in h_N} \alpha^{t-p(a)} D_a e^{-\eta \|a - \bar{e}_N\|^2/2}$, where G_N and γ_N are given in (3.19) and (3.20). For the follow policy, we have

$$F_N(\mu^f(h_N))^+ = F_N(\bar{e}_N)^+ = \left(1 - \underbrace{\frac{\theta}{\kappa_v} \sum_{a \in h_t} \alpha^{t-p(a)} D_a e^{-\eta \|a - \bar{e}_N\|^2/2}}_{\delta_N}\right)^+. \quad (4.2)$$

Thus, we see that $F_N(\mu^f(h_N))^+$ depends on the state h_N via only δ_N . Next, from Figure 3-9, we see that $F_N^h(\mu^h(h_N))$ is an excellent approximation to $F_N(\mu_N^*(h_N))^+$. Substituting (3.23) into (3.22), we obtain

$$\begin{aligned} F_N^h(\mu^h(h_N)) &= \left(1 - \frac{\theta}{\kappa_v} D_{\tau_N} e^{-(\phi(\sqrt{\frac{\eta}{2}}\gamma_N) + \sqrt{\frac{\eta}{2}}\gamma_N)^2}\right) e^{-B_1 \phi(\sqrt{\frac{\eta}{2}}\gamma_N)^2} \\ &= \left(1 - \frac{\theta}{\kappa_v} \frac{\|G_N\|}{\delta_N} e^{\frac{\eta}{2} \frac{\|G_N\|^2}{\delta_N^2}} e^{-(\phi(\sqrt{\frac{\eta}{2}} \frac{\|G_N\|}{\delta_N}) + \sqrt{\frac{\eta}{2}} \frac{\|G_N\|}{\delta_N})^2}\right) e^{-B_1 \phi(\sqrt{\frac{\eta}{2}} \frac{\|G_N\|}{\delta_N})^2}. \end{aligned} \quad (4.3)$$

Thus, we see that $F_N^h(\mu^h(h_N))$ depends on the state h_N via only δ_N and $\|G_N\|$. Therefore, (4.2) and (4.3) suggest that in approximating the second term of the reward-to-go in (4.1) under all three policies, we may use δ_N and $\|G_N\|$ as features to summarize the information contained in the state h_N . In particular, for each policy $\mu \in \{\mu_N^*, \mu^h, \mu^f\}$, we consider

approximating $\beta_N \theta B_1^{n/2} F_N(\mu(h_N))^+$ with a low order polynomial function of δ_N and $\|G_N\|$:

$$\beta_N \theta B_1^{n/2} F_N(\mu(h_N))^+ \approx C_{N,1} + C_{N,2} \delta_N + C_{N,3} \delta_N^2 + C_{N,4} \|G_N\|^2. \quad (4.4)$$

The coefficients will be obtained by running a regression on randomly generated state samples and their corresponding reward-to-go values. The functional form in (4.4) is chosen because (i) it yields a good approximation as we will demonstrate shortly, and (ii) it gives us a closed form approximation to $E_{e_{N-1}}[F_N(\cdot)^+ | e_{N-2}]$ which we will need in the next iteration of the DP. (A $\|G_N\|$ term is not included in the polynomial because we cannot get closed form expression for the expected value of $\|G_N\|$.) Let $C_{N,i}^*$, $C_{N,i}^h$, and $C_{N,i}^f$, $i = 1, \dots, 4$, be the regression coefficients for the optimal, heuristic, and follow reward-to-go, respectively. The next step is to test the performance of our approximation in (4.4) using simulations. We compute the regression coefficients $C_{N,i}^*$, $C_{N,i}^h$, and $C_{N,i}^f$ in (4.4) from a training set that comprises 2,500 randomly generated state samples. We then test the performance of the regression with an independent test set that comprises 2,500 randomly generated state samples. The R-square values of the training and test set for each policy are plotted in Figure 4-1. It can be seen from Figure 4-1 that our chosen functional form in (4.4) works very well in approximating the reward-to-go functions at time N for all three policies. Let us denote the resulting approximations to J_N^* , J_N^h , and J_N^f by \tilde{J}_N^* , \tilde{J}_N^h , and \tilde{J}_N^f , respectively. Note that if $F_N(\mu^f(h_N))^+ > 0$ for the follow policy, then (4.2) implies that the approximation in (4.4) is exact with $C_{N,1}^f = \beta_N \theta B_1^{n/2}$, $C_{N,2}^f = -\beta_N \theta B_1^{n/2} \frac{\theta}{\kappa_v}$, and $C_{N,3}^f = C_{N,4}^f = 0$.

4.4 Reward-To-Go Approximation at the 2nd DP Iteration

We now move on to the second iteration of the DP. For conciseness, we shall show the calculations for the optimal reward-to-go function only; the calculations for the heuristic, follow, and no-vaccine reward-to-go are analogous. The optimal reward-to-go function at time $N - 1$ is

$$J_{N-1}^*(h_{N-1}) = \max_{v_{N-1}} \left\{ E_{I_{e_{N-1}}, e_{N-1}} \left[\beta_{N-1} \bar{c}(y_{N-1}, e_{N-1}) + J_N^*(h_N) \mid e_{N-2} \right] \right\}. \quad (4.5)$$

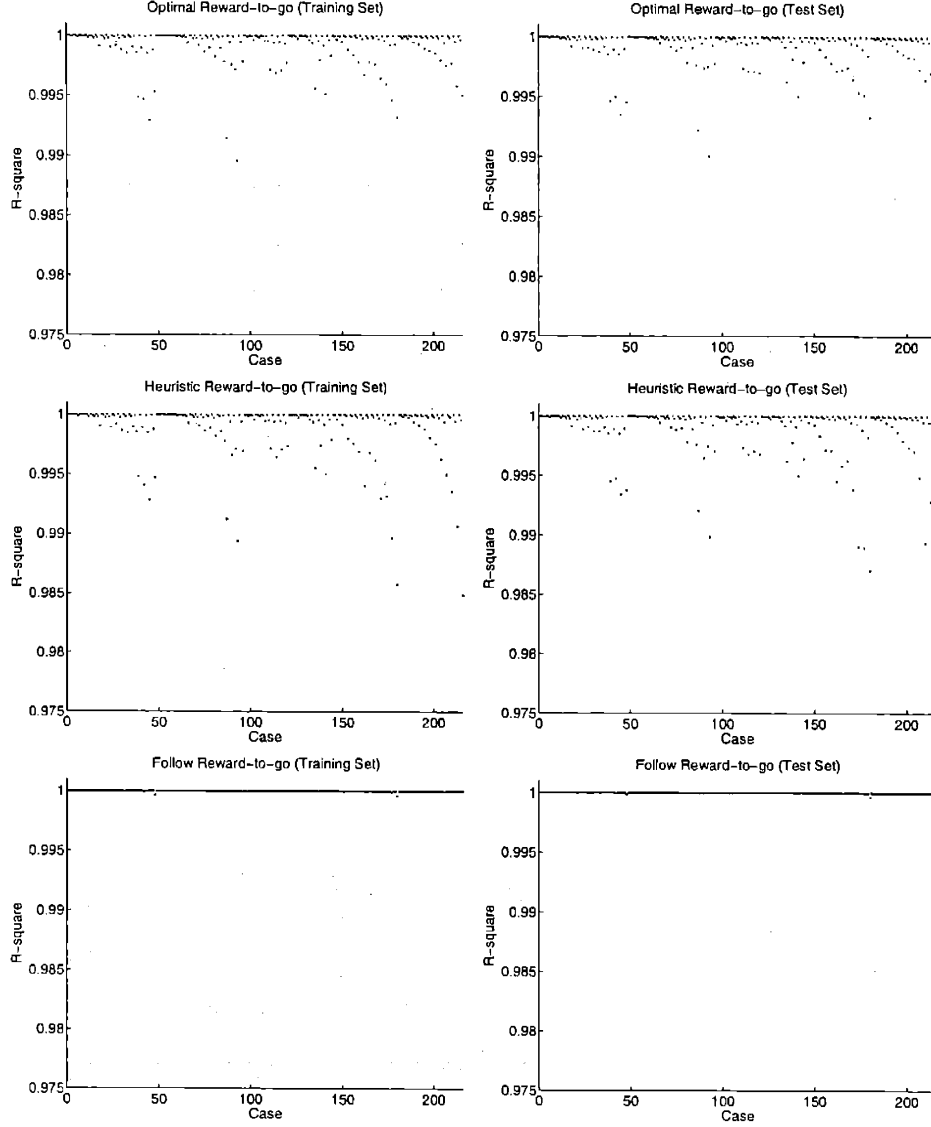


Figure 4-1: R-square values for reward-to-go approximations at the 1st DP iteration.

The next step is thus to evaluate $E_{I_{e_{N-1}}, e_{N-1}} [J_N^* | e_{N-2}]$. Using our reward-to-go approximation in (4.4), we have

$$\begin{aligned}
& E_{I_{e_{N-1}}, e_{N-1}} [J_N^*(h_N) | e_{N-2}] \\
& \approx E_{I_{e_{N-1}}, e_{N-1}} \left[\beta_N E_{e_N} [\bar{c}(x_N, e_N) | e_{N-1}] + C_{N,1}^* + C_{N,2}^* \delta_N + C_{N,3}^* \delta_N^2 + C_{N,4}^* \|G_N\|^2 | e_{N-2} \right] \\
& = \beta_N E_{e_N} [\alpha \bar{c}(x_{N-1}, e_N) | e_{N-2}] + \beta_N \alpha D_{v_{N-1}} E_{e_N} [g(v_{N-1} - e_N) | e_{N-2}] \\
& \quad + p_e \beta_N \alpha E_{e_{N-1}} [D_{e_{N-1}} | e_{N-2}] E_{e_{N-1}, e_N} [g(e_{N-1} - e_N) | e_{N-2}] \\
& \quad + E_{I_{e_{N-1}}, e_{N-1}} [C_{N,1}^* + C_{N,2}^* \delta_N + C_{N,3}^* \delta_N^2 + C_{N,4}^* \|G_N\|^2 | e_{N-2}]. \tag{4.6}
\end{aligned}$$

The equality follows since

$$\bar{c}(x_N, e_N) = \alpha \bar{c}(x_{N-1}, e_N) + \alpha D_{v_{N-1}} g(v_{N-1} - e_N) + I_{e_{N-1}} \alpha D_{e_{N-1}} g(e_{N-1} - e_N). \quad (4.7)$$

Substituting (4.6) into (4.5), we have

$$\begin{aligned} & J_{N-1}^*(h_{N-1}) \\ & \approx \max_{v_{N-1}} \left\{ \beta_{N-1} E_{e_{N-1}} \left[\bar{c}(x_{N-1}, e_{N-1}) \mid e_{N-2} \right] + \beta_N E_{e_N} \left[\bar{c}(x_{N-1}, e_N) \mid e_{N-2} \right] \right. \\ & \quad + \beta_{N-1} D_{v_{N-1}} E_{e_{N-1}} \left[g(v_{N-1} - e_{N-1}) \mid e_{N-2} \right] + \beta_N \alpha D_{v_{N-1}} E_{e_N} \left[g(v_{N-1} - e_N) \mid e_{N-2} \right] \\ & \quad + p_e \beta_N \alpha E_{e_{N-1}} \left[D_{e_{N-1}} \mid e_{N-2} \right] E_{e_{N-1}, e_N} \left[g(e_{N-1} - e_N) \mid e_{N-2} \right] \\ & \quad \left. + E_{I_{e_{N-1}, e_{N-1}}} \left[C_{N,1}^* + C_{N,2}^* \delta_N + C_{N,3}^* \delta_N^2 + C_{N,4}^* \|G_N\|^2 \mid e_{N-2} \right] \right\}. \end{aligned} \quad (4.8)$$

Note that the first two terms of this objective function are independent of v_{N-1} (these are the residual immunity terms). Let us denote the sum of the last four terms of the approximate objective function in (4.8) by $F_{N-1}^*(v_{N-1})$. Therefore, we are approximating the second iteration of the DP with the problem $\max_{v_{N-1}} F_{N-1}^*(v_{N-1})$. In the Appendix, we compute closed formed approximations to $E_{e_{N-1}} \left[D_{e_{N-1}} \mid e_{N-2} \right]$, $E_{I_{e_{N-1}, e_{N-1}}} \left[\delta_N \mid e_{N-2} \right]$, $E_{I_{e_{N-1}, e_{N-1}}} \left[\delta_N^2 \mid e_{N-2} \right]$, and $E_{I_{e_{N-1}, e_{N-1}}} \left[\|G_N\|^2 \mid e_{N-2} \right]$. This means that we can obtain a closed form approximation to F_{N-1}^* , which greatly facilitates the optimization process. Let $\tilde{\mu}_{N-1}^*$ be the optimal policy to our approximate problem in (4.8). Since we have shown that \tilde{J}_N^* is a good approximation to J_N^* (see Figure 3-9), the approximate optimal policy $\tilde{\mu}_{N-1}^*$ so obtained should be very close to the optimal policy μ_{N-1}^* .

As mentioned before, we will apply the same approximation techniques that we used for the optimal reward-to-go to the heuristic, follow, and no-vaccine reward-to-go. Our next step is to choose the features for approximating $F_{N-1}^*(\tilde{\mu}_{N-1}^*(h_{N-1}))$, $F_{N-1}^h(\mu^h(h_{N-1}))$, $F_{N-1}^f(\mu^f(h_{N-1}))$, and $F_{N-1}^n(\mu^n(h_{N-1}))$. In view of our success in using δ_N and $\|G_N\|$ as features at time N , let us consider again approximating these functions with the functional form

$$C_{N-1,1} + C_{N-1,2} \delta_{N-1} + C_{N-1,3} \delta_{N-1}^2 + C_{N-1,4} \|G_{N-1}\|^2. \quad (4.9)$$

As before, we test the performance of our approximation in (4.9) using simulations. The coefficients $C_{N,i}^*$, $C_{N,i}^h$, $C_{N,i}^f$, and $C_{N,i}^n$ in (4.8) are obtained from the regressions at time N .

The coefficients $C_{N-1,i}^*$, $C_{N-1,i}^h$, $C_{N-1,i}^f$, and $C_{N-1,i}^n$ are then computed from a training set that comprises 2,500 randomly generated state samples at time $N - 1$; the performance of the regression is then tested with an independent test set that comprises 2,500 randomly generated state samples. The R-square values of the training and test set for the approximate optimal, heuristic, and follow reward-to-go approximations are plotted in Figure 4-2. It can be seen from Figure 4-2 that the chosen functional form in (4.9) performs very well in approximating the reward-to-go functions at time $N - 1$. Let us denote the resulting approximate reward-to-go functions by \tilde{J}_{N-1}^* , \tilde{J}_{N-1}^h , \tilde{J}_{N-1}^f , and \tilde{J}_{N-1}^n , respectively.

4.5 DP Approximation Algorithm

The results in the last two sections show that our scheme of using δ_t and $\|G_t\|$ as features is very successful in approximating the reward-to-go functions at time N and $N - 1$ for all four policies. Let us now consider extending our approximation scheme to all subsequent DP iterations. The algorithm is as follows:

1. Evaluate the closed form approximate objective function at time t . The approximate objective function is

$$E_{I_{e_t}, e_t} \left[\beta_t \bar{c}(y_t, e_t) + \tilde{J}_{t+1}^*(h_{t+1}) \mid e_{t-1} \right] \quad (4.10)$$

where

$$\begin{aligned}
& E_{I_{e_t}, e_t} \left[\tilde{J}_{t+1}^*(h_{t+1}) \mid e_{t-1} \right] \\
&= E_{I_{e_t}, e_t} \left[\sum_{k=1}^{N-t} \beta_{t+k} \alpha^{k-1} E_{e_{t+k}} \left[\bar{c}(x_{t+1}, e_{t+k}) \mid e_t \right] \mid e_{t-1} \right] \\
&\quad + E_{I_{e_t}, e_t} \left[C_{t+1,1}^* + C_{t+1,2}^* \delta_{t+1} + C_{t+1,3}^* \delta_{t+1}^2 + C_{t+1,4}^* \|G_{t+1}\|^2 \mid e_{t-1} \right] \\
&= \sum_{k=1}^{N-t} \beta_{t+k} \sum_{a \in h_t} \alpha^{t+k-p(a)} D_a E_{e_{t+k}} \left[g(a - e_{t+k}) \mid e_{t-1} \right] \\
&\quad + \sum_{k=1}^{N-t} \beta_{t+k} \alpha^k D_{v_t} E_{e_{t+k}} \left[g(v_t - e_{t+k}) \mid e_{t-1} \right] \\
&\quad + p_e \sum_{k=1}^{N-t} \beta_{t+k} \alpha^k E_{e_t} \left[D_{e_t} \mid e_{t-1} \right] E_{e_t, e_{t+k}} \left[g(e_t - e_{t+k}) \mid e_{t-1} \right] \\
&\quad + E_{I_{e_t}, e_t} \left[C_{t+1,1}^* + C_{t+1,2}^* \delta_{t+1} + C_{t+1,3}^* \delta_{t+1}^2 + C_{t+1,4}^* \|G_{t+1}\|^2 \mid e_{t-1} \right] \tag{4.11}
\end{aligned}$$

- Let $F_t^*(v_t)$ be the part of the objective function that depends on v_t . The approximate DP problem at time t is thus $\max_{v_t} F_t^*(v_t)$. Let $\tilde{\mu}_t^*$ be the optimal policy to this approximate problem. Generate K (K even) state samples randomly and solve the problem $\max_{v_t} F_t^*(v_t)$ for each sample. Let Y_i , $i = 1, \dots, K$, be the optimal values of these samples. Also, for each state sample, compute the value of the features

$$\delta_t = \sum_{a \in h_t} \alpha^{t-p(a)} D_a e^{-\eta \|a - \bar{e}_t\|^2 / 2}, \tag{4.12}$$

$$\|G_t\| = \left\| \sum_{a \in h_t} \alpha^{t-p(a)} (\bar{e}_t - a) D_a e^{-\eta \|a - \bar{e}_t\|^2 / 2} \right\|. \tag{4.13}$$

Denote these computed values by $\delta_{t,i}$ and $\|G_{t,i}\|$, $i = 1, \dots, K$.

- Approximate $F_t^*(\tilde{\mu}_t^*(h_t))$ with

$$C_{t,1}^* + C_{t,2}^* \delta_t + C_{t,3}^* \delta_t^2 + C_{t,4}^* \|G_t\|^2 \tag{4.14}$$

by computing the regression coefficients $C_{t,i}^*$, $i = 1, \dots, 4$, from $\delta_{t,i}$, $\|G_{t,i}\|$, and Y_i , $i = 1, \dots, K/2$.

- Test the performance of the approximation in (4.14) with a test set that comprises

$\delta_{t,i}$, $\|G_{t,i}\|$, and Y_i , $i = K/2 + 1, \dots, K$. Compute the R-square value for this test set.

5. If the R-square value of the test set is satisfactory, denote the resulting approximate reward-to-go by \tilde{J}_t^* and go to the next iteration of the DP.

The R-square values for each case in each DP iteration are plotted in Figure 4-3 to Figure 4-9. From these regression results, we see that our DP approximation algorithm works very well in most cases. In the next chapter, we will examine the results of our computational study and discuss their implications on the practicability and cost-effectiveness of the four vaccine strain selection policies.

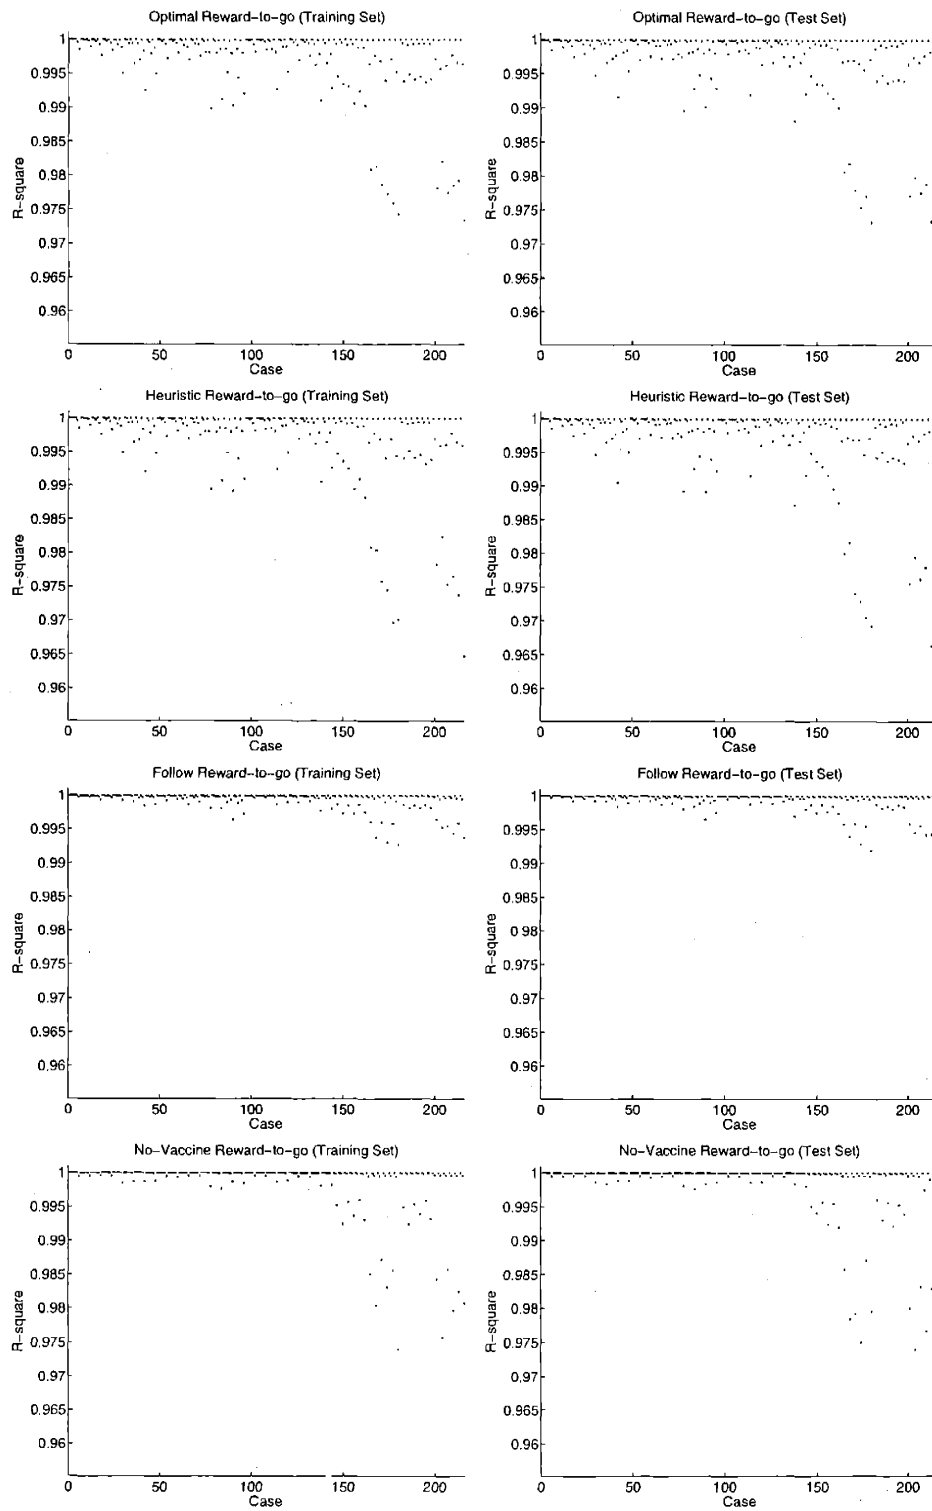


Figure 4-2: R-square values for reward-to-go approximations at the 2nd DP iteration.

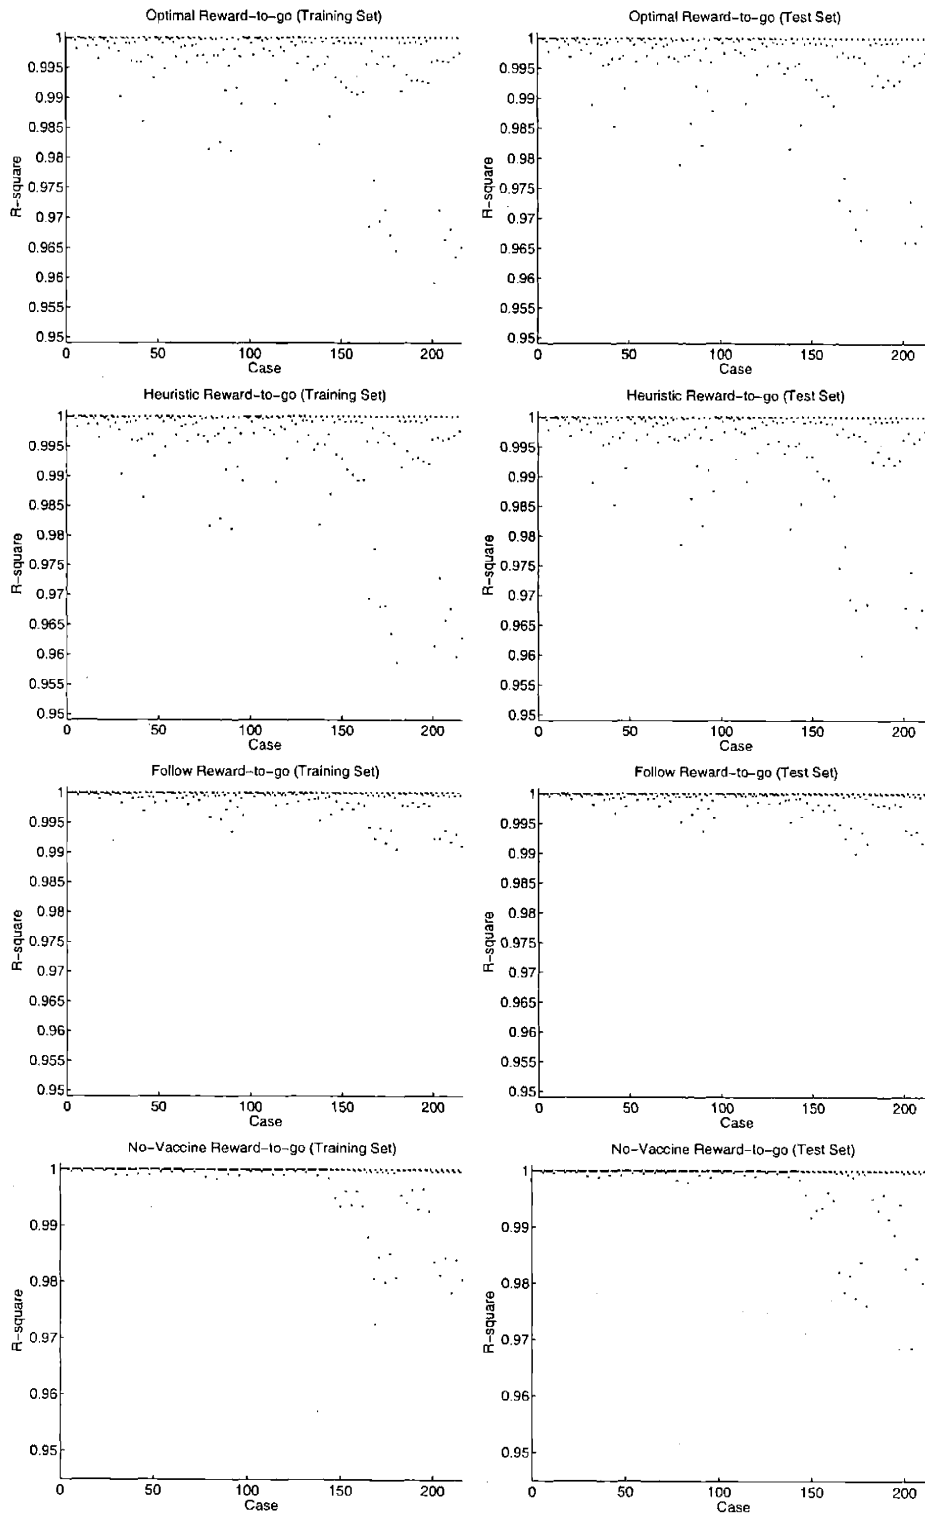


Figure 4-3: R-square values for reward-to-go approximations at the 3rd DP iteration.

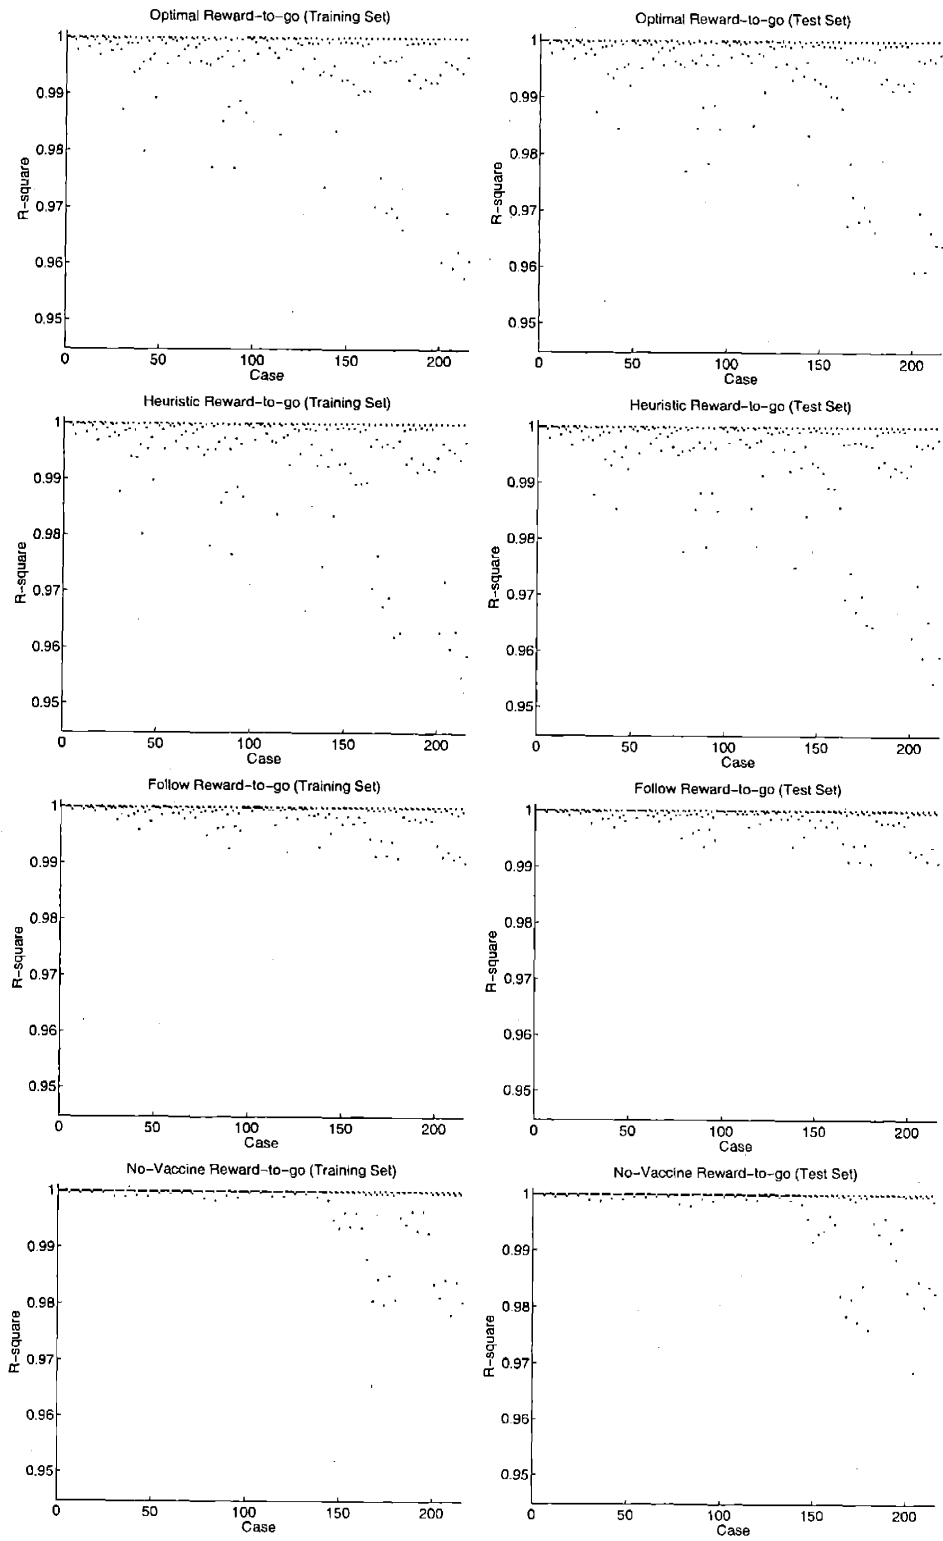


Figure 4-4: R-square values for reward-to-go approximations at the 4th DP iteration.

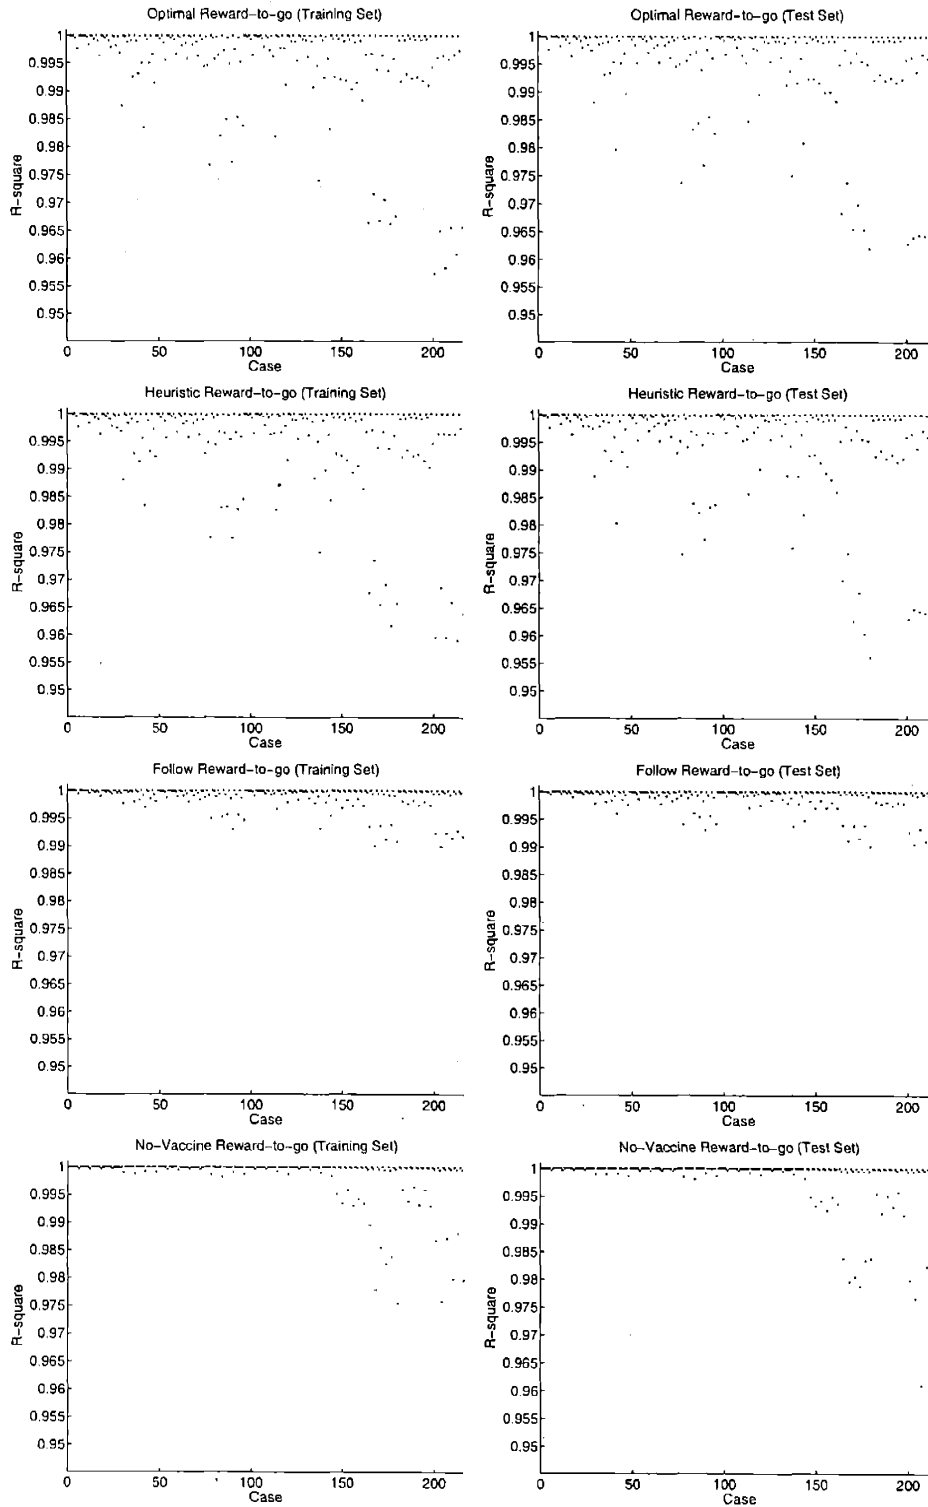


Figure 4-5: R-square values for reward-to-go approximations at the 5th DP iteration.

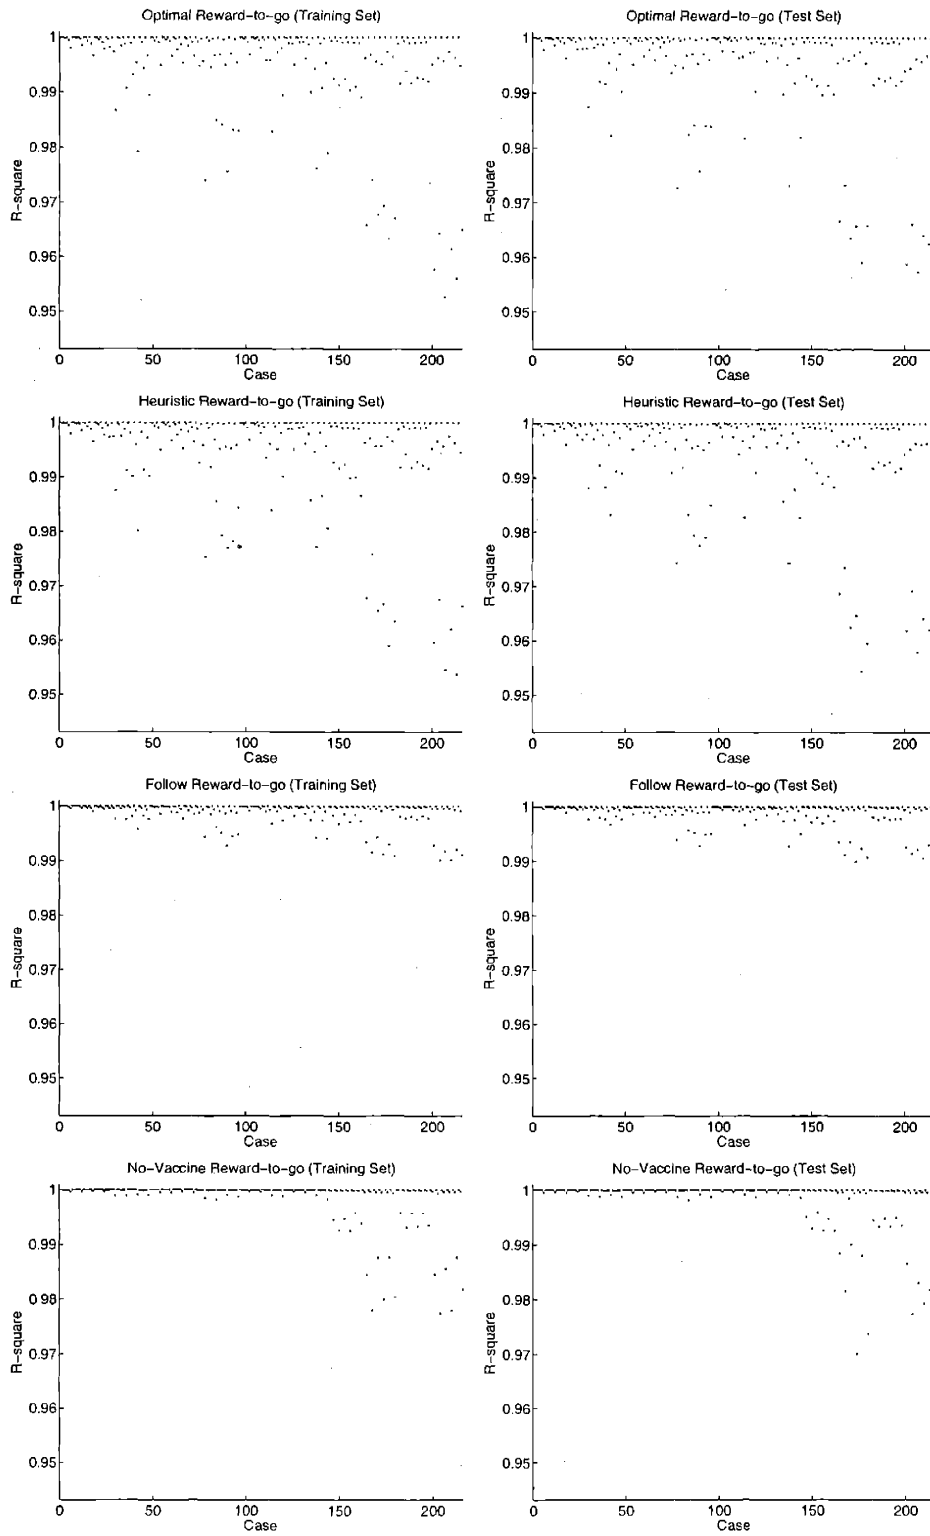


Figure 4-6: R-square values for reward-to-go approximations at the 6th DP iteration.

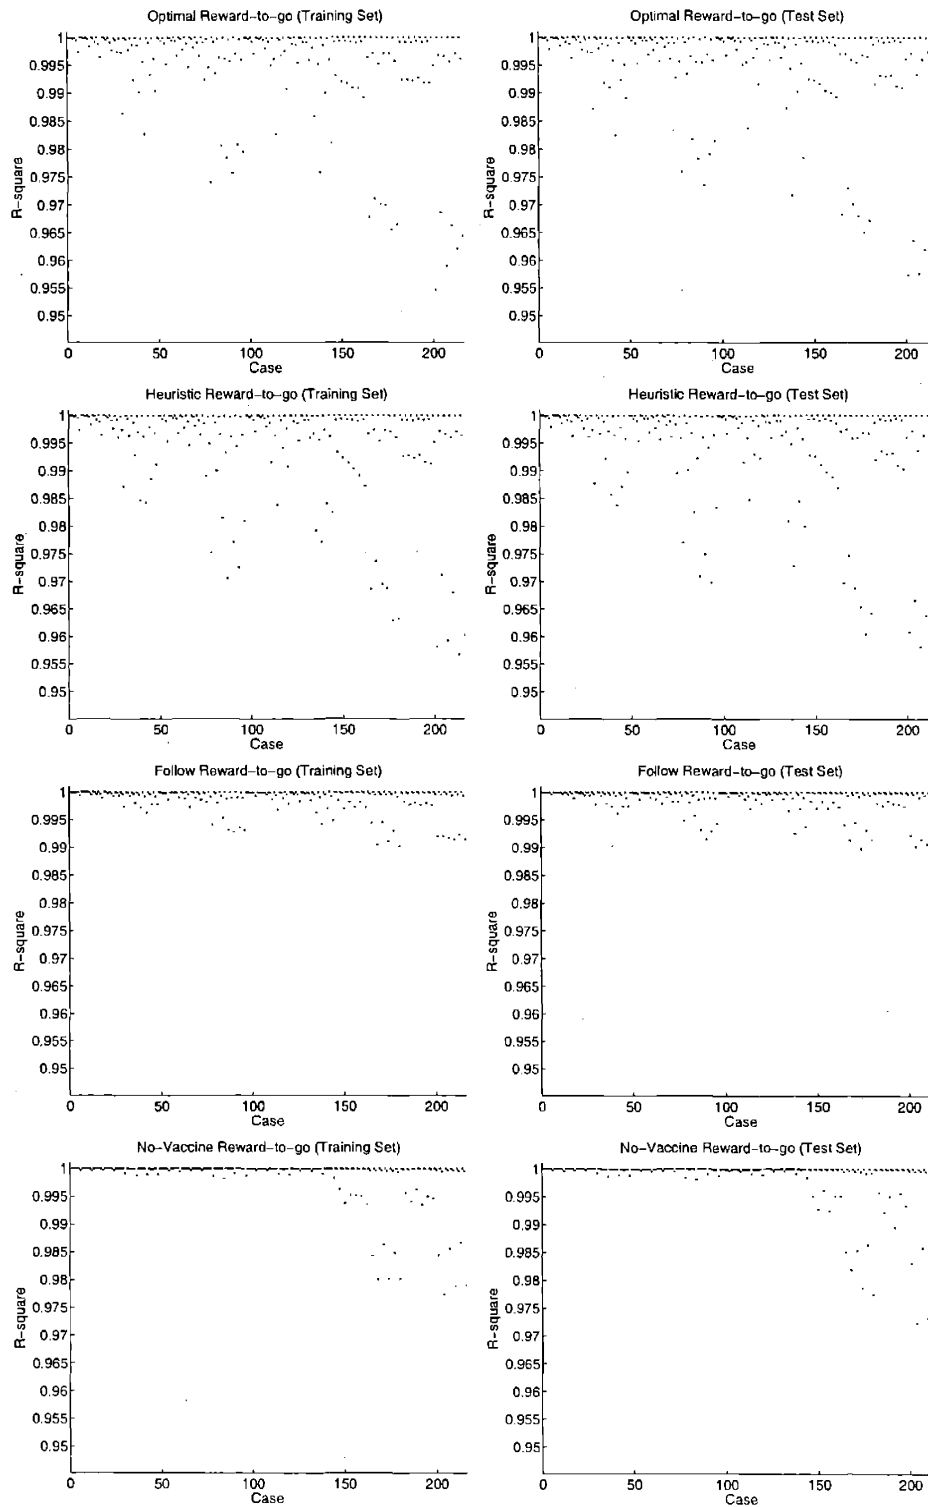


Figure 4-7: R-square values for reward-to-go approximations at the 7th DP iteration.

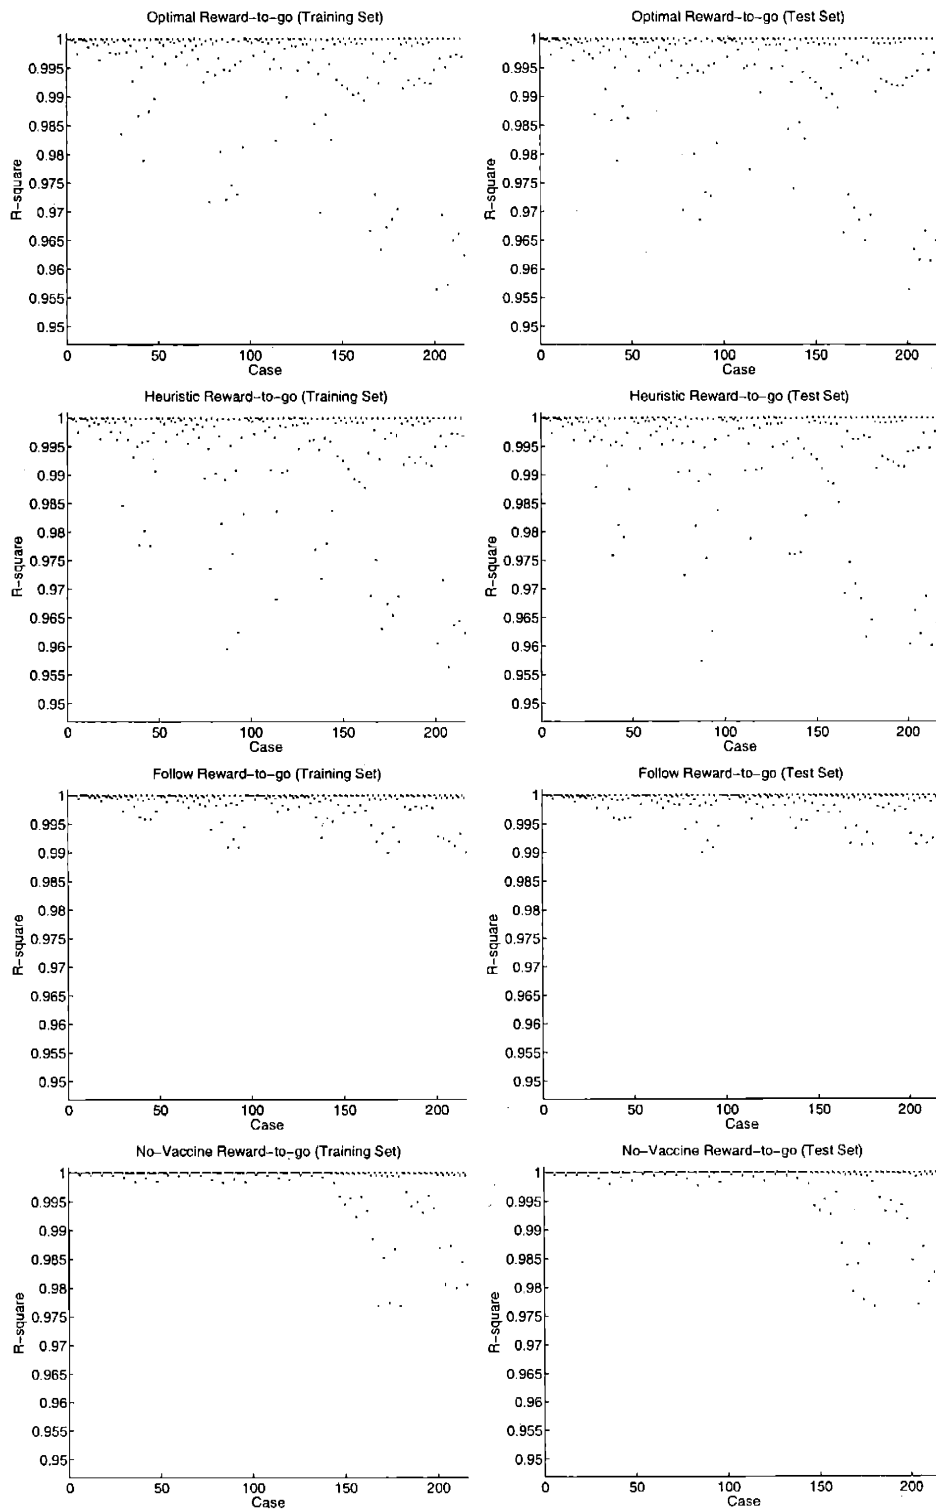


Figure 4-8: R-square values for reward-to-go approximations at the 8th DP iteration.

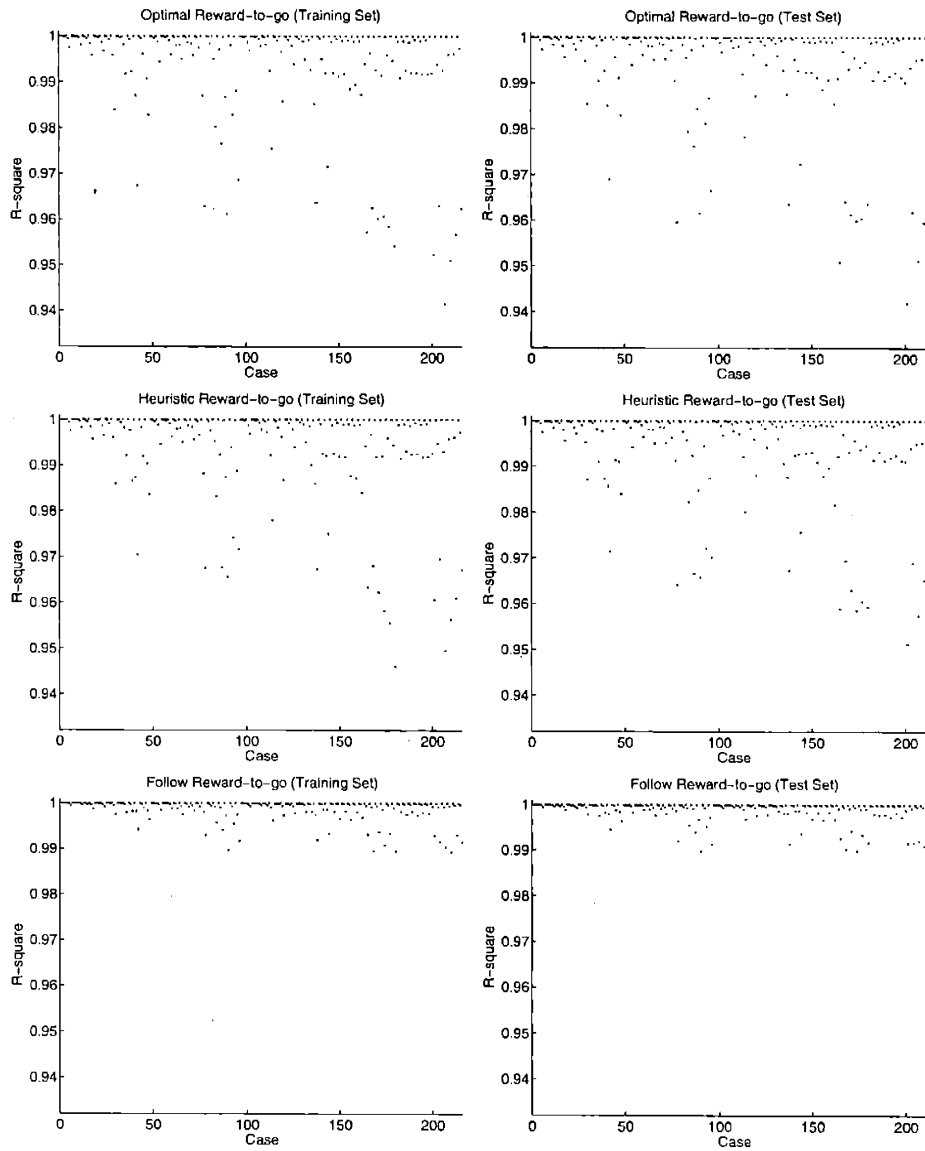


Figure 4-9: R-square values for reward-to-go approximations at the 9th DP iteration.

Chapter 5

Discussion

In the last chapter, we solved the DP to reasonable accuracy for 216 different combinations of parameter values (see Table A.1 in the Appendix). In this chapter, we shall look at the results of our computational study and discuss their implications on vaccine strain selection policies.

5.1 Performance Evaluation of Different Policies

Following the notation that we introduced in the last chapter, \bar{J}_1^* , \bar{J}_1^h , \bar{J}_1^f , and \bar{J}_1^n are the approximate DP objective function value (which we shall simply call “rewards” to facilitate our discussions) under the optimal, heuristic, follow, and no-vaccine policy, respectively. Let

$$r_h = \frac{\bar{J}_1^h}{\bar{J}_1^*}, \quad r_f = \frac{\bar{J}_1^f}{\bar{J}_1^*}, \quad r_n = \frac{\bar{J}_1^n}{\bar{J}_1^*}. \quad (5.1)$$

We shall use r_h , r_f , and r_n as the performance measures for the three policies. The plots in Figure 5-1 show the values of r_h , r_f , and r_n for each of our 216 computational case studies.

Let us first look at the performance of the heuristic policy. From the plot of the heuristic-to-optimal reward ratios (i.e., r_h), we see that the heuristic policy is very close to optimal in almost all cases ($r_h > 0.95$ except in one case). Recall that the time-independent heuristic policy is a near-optimal policy to the 1-period problem. So the high values of r_h means that when choosing a vaccine strain, we can achieve near-optimal performance by using a 1-period optimal policy (instead of having to solve a DP with a long time-horizon). This also

implies that when predicting future epidemic evolution, we need to be accurate on only the most immediate upcoming epidemic strain in order to identify an effective vaccine strain. The close-to-optimal performance of the heuristic policy also indicates that our approximate history reduction is an accurate and simple way of visualizing the vaccine selection problem. To understand why the 1-period optimal policy works so well in the multi-period problem, let us look into the objective function of a 2-period problem for insights:

$$\begin{aligned}
& E_{I_{e_t}, I_{e_{t+1}}, e_t, e_{t+1}} \left[\beta_t \bar{c}(y_t, e_t) + \beta_{t+1} \bar{c}(y_{t+1}^*, e_{t+1}) \mid e_{t-1} \right] \\
= & \beta_t \underbrace{E_{e_t} \left[\bar{c}(x_t, e_t) \mid e_{t-1} \right]}_{\text{residual immunity against } e_t \text{ in } x_t} + \beta_{t+1} \underbrace{E_{e_{t+1}} \left[\bar{c}(x_t, e_{t+1}) \mid e_{t-1} \right]}_{\text{residual immunity against } e_{t+1} \text{ in } x_t} \\
& + \beta_t \underbrace{D_{v_t} E_{e_t} \left[g(v_t - e_t) \mid e_{t-1} \right]}_{v_t - e_t \text{ cross-reaction}} + \beta_{t+1} \underbrace{\alpha D_{v_t} E_{e_{t+1}} \left[g(v_t - e_{t+1}) \mid e_{t-1} \right]}_{v_t - e_{t+1} \text{ cross-reaction}} \\
& + p_e \beta_{t+1} \underbrace{\alpha E_{e_t} \left[D_{e_t} \mid e_{t-1} \right] E_{e_t, e_{t+1}} \left[g(e_t - e_{t+1}) \mid e_{t-1} \right]}_{e_t - e_{t+1} \text{ cross-reaction}} \\
& + \beta_{t+1} \underbrace{E_{I_{e_t}, e_t} \left[E_{e_{t+1}} \left[D_{\mu^*(h_{t+1})} g(\mu^*(h_{t+1}) - e_{t+1}) \mid e_t \right] \mid e_{t-1} \right]}_{\text{(optimal) } v_{t+1} - e_{t+1} \text{ cross-reaction}} \tag{5.2}
\end{aligned}$$

The first two residual immunity terms are unaffected by the choice of v_t . The high performance of the 1-period optimal policy in the multi-period problem suggests that the $v_t - e_t$ cross-reaction term (which is simply the objective function of the 1-period problem) is the dominant term in the 2-period objective function in terms of its dependence on v_t . This is equivalent to saying that the sum of the last three terms is a weak function of v_t . These terms corresponds to the effect of v_t on the clearance of e_{t+1} . We can decompose this effect into the following components:

1. The most obvious interaction between v_t and e_{t+1} is, of course, the $v_t - e_{t+1}$ cross-reaction.
2. If the individual is exposed to e_t at time t , the cross-reaction of v_t and e_t lowers the effective dose of e_t (which is the objective of the 1-period problem). Then at time $t + 1$, e_t can cross-react with v_{t+1} and e_{t+1} . Lower effective dose of e_t therefore implies weaker $e_t - v_{t+1}$ and $e_t - e_{t+1}$ cross-reaction. Weaker $e_t - v_{t+1}$ cross-reaction, on the other hand, leads to higher effective dose of v_{t+1} , which in turn implies stronger

$v_{t+1}-e_{t+1}$ cross-reaction. Therefore, if the individual is exposed to e_t , then the v_t-e_t cross-reaction results in (i) weaker e_t-e_{t+1} cross-reaction and (ii) stronger $v_{t+1}-e_{t+1}$ cross-reaction.

3. v_t itself can cross-react with v_{t+1} , thus lowering the effective dose of v_{t+1} and resulting in weaker $v_{t+1}-e_{t+1}$ cross-reaction.

We can argue qualitatively why the net effect of the above components are immaterial when compared to the v_t-e_t cross-reaction term. Relative to the v_t-e_t cross-reaction term in (5.2), each of the above components is a weaker function of v_t because (1) the immunity provided by v_t is discounted by α at time $t + 1$ and (2) the reward at time $t + 1$ is discounted by $\beta_{t+1} \leq \beta_t$. Moreover, the v_t-e_{t+1} cross-reaction term in Component 1 has less “room for optimization” than the v_t-e_t cross-reaction term because the variance of e_{t+1} prediction is larger. Besides being weaker functions of v_t , these components also have counterbalancing effects, as we explain as follows. Recall from our analysis in Chapter 3 that the follow policy is a good and robust policy for the 1-period problem. This implies that the optimal v_{t+1} must be close to \bar{e}_{t+1} . Thus, if v_t is close to \bar{e}_{t+1} , then v_t is likely to be somewhat close to v_{t+1} too. It then follows that Component 1 and 3 have opposite effects: The stronger the v_t-e_{t+1} cross-reaction in Component 1, the weaker the $v_{t+1}-e_{t+1}$ cross-reaction in Component 3 (because the v_t-v_{t+1} cross-reaction lowers the effective dose of v_{t+1}), and vice versa. Therefore, the contribution of v_t in Component 1 and 3 tend to cancel out. Next, within Component 2, the e_t-e_{t+1} cross-reaction effect in (i) and the $v_{t+1}-e_{t+1}$ cross-reaction effect in (ii) are obviously opposing forces. In addition, the contribution of Component 2 is further discounted by p_e , which is the probability that the individual is exposed to e_t . In short, all these factors suggest that the net effect of v_t in all these components is much weaker than the effect of v_t in the v_t-e_t cross-reaction term, hence the good performance of the 1-period optimal policy in the 2-period problem. For a multi-period problem with a longer time-horizon, the residual effects of v_t in later periods are further discounted by α , $\{\beta_{t+k}\}$, and the larger variance and drift of the epidemic evolution. Despite the qualitative nature of the above arguments, their implications are consistent with the quantitative results in Figure 5-1.

Next, we turn our attention to the performance of the follow policy. The plot of the follow-to-optimal reward ratios (i.e., r_f) shows that the follow policy performs quite well in

all cases. This outcome is consistent with the performance results of the heuristic policy. As discussed in the last paragraph, we can achieve good performance in the multi-period problem by using the 1-period optimal solution at every time period. Since the follow policy is a good and robust 1-period policy, it should come as no surprise that the follow policy also performs quite well under the multi-period setting.

Finally, we look at the performance of the no-vaccine policy. From the plot of the no-vaccine-to-optimal reward ratios (i.e., r_n), we see that the performance of the no-vaccine policy varies widely depending on the parameter values. The value of the no-vaccine-to-optimal reward ratio can reach as low as 0.1 and as high as 0.78. Under the no-vaccine policy, we rely entirely on past epidemic exposures to provide immunity against future epidemic attacks. Therefore, we expect the performance of the no-vaccine policy to be relatively good in the following scenarios (or a combination of them):

1. The decay constant α is high, i.e. antibodies decay at a slow rate. Unlike immunity provided by vaccines, which are administered shortly before the flu seasons, immunity triggered by prior epidemic exposures suffers a longer time delay. Larger values of α will therefore increase the performance of the no-vaccine policy.
2. The norm of d is small, i.e. the magnitude of antigenic drift is small. If the drift is large, cross-reactions between successive epidemic strains are weaker, thus lowering the performance of the no-vaccine policy.
3. The value of p_e is large, i.e. the probability of epidemic exposure is high. Since the no-vaccine policy relies on the protection given by prior epidemic exposures, its performance increases as the probability of exposure increases.

In the next section, we provide numerical results to support these hypotheses.

5.2 Sensitivity of Policies Performance to Parameter Values

Our next focus is to see how individual parameter affects the performance of different policies. To this end, we look into how sensitive r_h , r_f , and r_n are to the value of each individual parameter (while holding other parameters constant) in our case studies. We begin with κ_v , which is the potency of vaccines (relative to epidemic viruses). Recall from Table 4.2 that we ran the DP for $\kappa_v \in \{0.3, 0.6\}$. Let us divide the 216 case studies into

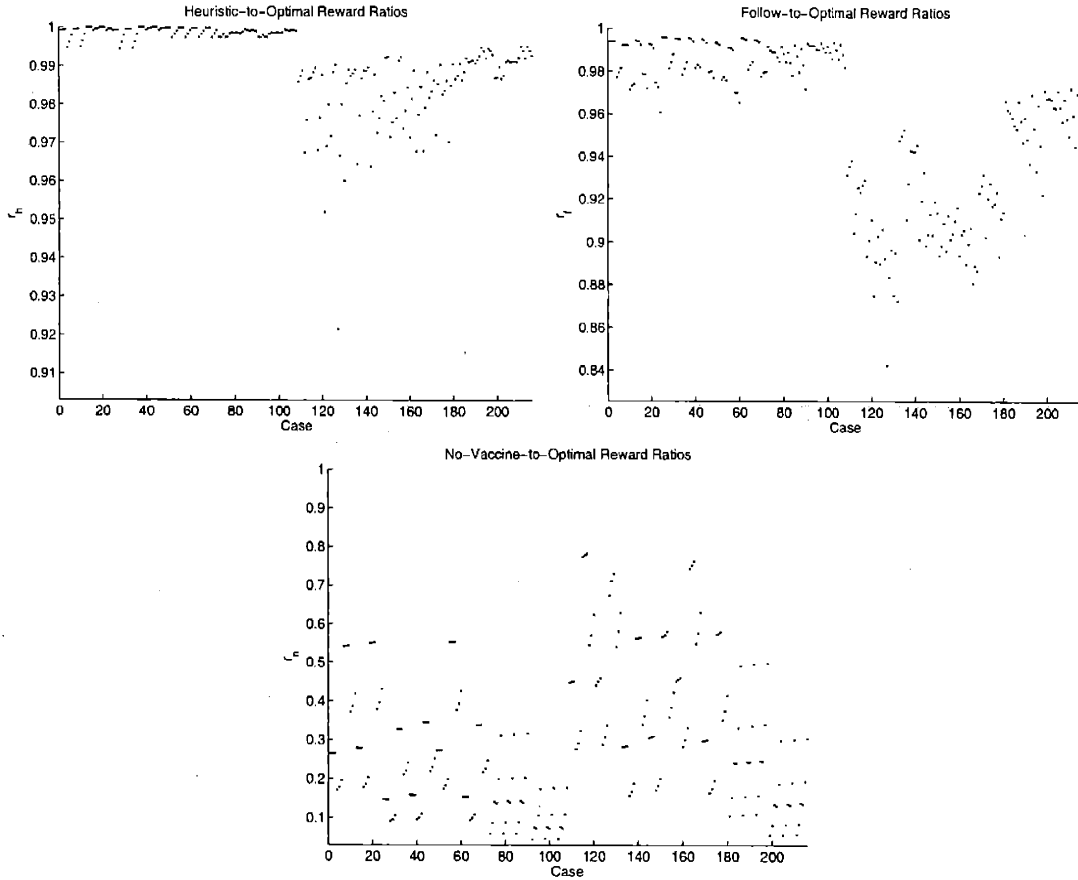


Figure 5-1: Reward ratios of the heuristic, follow, and no-vaccine policy in each case study.

two groups according to the value of κ_v in each case. We then pair each member of one group with a member of the other if the two have the same parameter values except κ_v . In Figure 5-2, we plot r_h , r_f , and r_n from one group against r_h , r_f , and r_n from the other group (each data point in the plots corresponds to a pair). From Figure 5-2, it appears that both the heuristic and follow policy performs better when κ_v is larger. This suggests that we can boost the performance of these policies by increasing κ_v via, for instance, a higher dosage of vaccine injection or the use of attenuated live viruses. On the other hand, we see that the no-vaccine policy becomes relatively less effective (relative to the optimal policy) as κ_v increases. This is to be expected since a higher vaccine potency is of no utility if there are no vaccinations.

In Figure 5-3 to Figure 5-7, we repeat the same sensitivity analysis for α , θ , p_e , $\|d\|$, and c_v . The plots for α in Figure 5-3 show that the performance of the heuristic and follow policy deteriorate significantly as α increases. This implies that the 1-period optimal policy

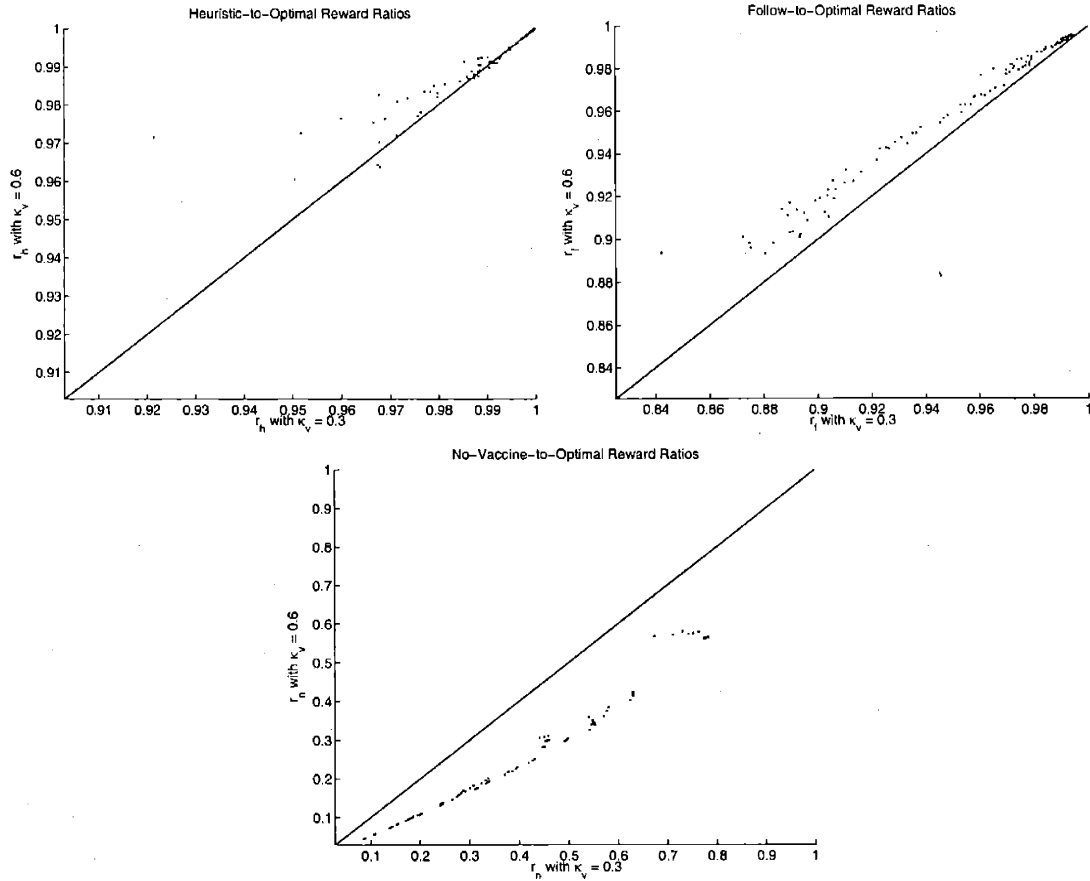


Figure 5-2: Performance of the heuristic, follow, and no-vaccine policy under different values of κ_v . The line shown in each plot is a 45 degree line. If a data point lies above the line, this means that the performance of the y -axis case is better than that of the x -axis case, and vice-versa.

becomes less effective in the multi-period problem if the antibodies decay at a slower rate. This is because with a slower decay rate, the immunity provided by a vaccine lasts longer and therefore exerts more interference on vaccine efficacy and epidemic viral clearance in future periods. For instance, the last three components in (5.2) become stronger functions of v_t which weaken the dominance of the v_t-e_t cross-reaction term in the 2-period objective. As a result, we expect the optimal vaccine policy at different times to be more strongly coupled when the value of α is large. This explains the weaker performance of the 1-period optimal policy in the multi-period problem when α is large. Since the heuristic and follow policy are similar to the 1-period optimal policy, it follows that their performance also degrade for the same reason as α gets large. Unlike the heuristic and follow policy, the no-vaccine policy performs better when α is large. As mentioned in the last section, we expect

that a slower decay rate will favor the performance of the no-vaccine policy because we rely entirely on epidemic exposures in past periods to provide protection against upcoming epidemic attacks. The results here support our hypothesis.

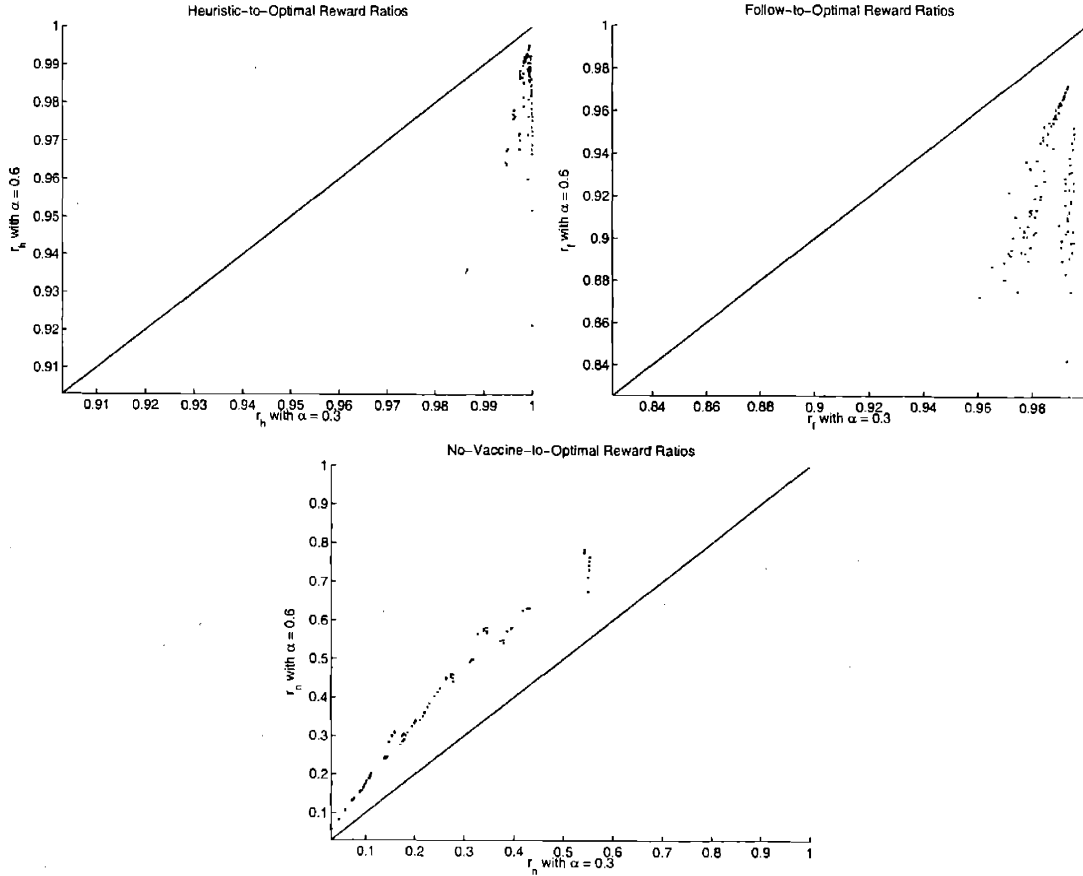


Figure 5-3: Performance of the heuristic, follow, and no-vaccine policy under different values of α . If a data point lies above the line, this means that the performance of the y -axis case is better than that of the x -axis case, and vice-versa.

Next, we look at the effect of θ , which is the clearance constant. Recall from (2.6) that $\theta \times 100$ is the percentage of epidemic viral particles removed by a vaccine for a naive immune state if the vaccine strain matches exactly the epidemic strain. In Figure 5-4, the plots of r_h reveal that the performance of the heuristic policy exhibits no monotonic dependence on the value of θ . From the plots of r_f , we see that the follow policy appears to perform better with a smaller θ . Finally, the plots of r_n show that the performance of the no-vaccine policy is rather insensitive to the value of θ .

Next, we turn our attention to the effect of p_e , which is the probability of epidemic exposure. From the plot of r_h in Figure 5-5, we see that the value of p_e has little impact

on the performance of the heuristic policy. The follow policy, on the other hand, performs better when p_e is smaller, as shown in the plot of r_f . An explanation for these results is that unlike the heuristic policy, the follow policy is not designed to avoid cross-reactions with past strains in the immunization history. With a larger p_e , the number of epidemic strains in the immunization history that can cross-react with the “follow” vaccine strain tends to be larger. As a result, the performance of the follow policy is worse when the probability of exposure is higher. The plot of r_n shows that the no-vaccine policy is more effective when p_e is larger, which agrees with our intuitions in Section 5.1.

Our next focus is $\|d\|$, which is the magnitude of the mean antigenic drift. We expect that the heuristic and follow policy to perform better as $\|d\|$ increases since the multi-period problem can be decomposed into a series of 1-period problems for $\|d\|$ large. On the other hand, the no-vaccine policy will perform worse for large $\|d\|$ since cross-reactions between successive epidemic strains are weaker. The results in Figure 5-6 confirm these views. Moreover, it is interesting to note that the coupling between the optimal decision at different times are still quite strong for $\|d\| = 0.4$ although $w(\|d\|) \approx 0.5$ (see Figure 3-2).

Finally, we look at the effect of c_v . Recall that we set $\sigma = c_v \frac{\|d\|}{\sqrt{n}}$. That is, the larger the value of c_v , the larger the variance in the epidemic evolution. The plot of r_h in Figure 5-7 shows that the heuristic policy performs better as the variance gets larger in all cases. A possible explanation for this is that when the variance is large, the prediction for epidemic strain e_{t+1} is much less accurate than the prediction of the epidemic strain e_t at time t (due to the Markov property of the epidemic evolution). Therefore, the larger variance leads the optimal policy to focus on hitting e_t instead of chasing after e_{t+1} . In other words, a larger variance favors the 1-period optimal policy, hence the increased performance of the heuristic policy. The performance results of the follow policy are consistent with this hypothesis with a few exceptional cases. These exceptions may be attributed to the fact that the follow policy is a less robust 1-period suboptimal policy than the heuristic policy. Finally, a larger variance also favors the performance of the no-vaccine policy. This is because with a larger variance, the epidemic strain prediction is less accurate, thus lowering the rewards of doing vaccinations.

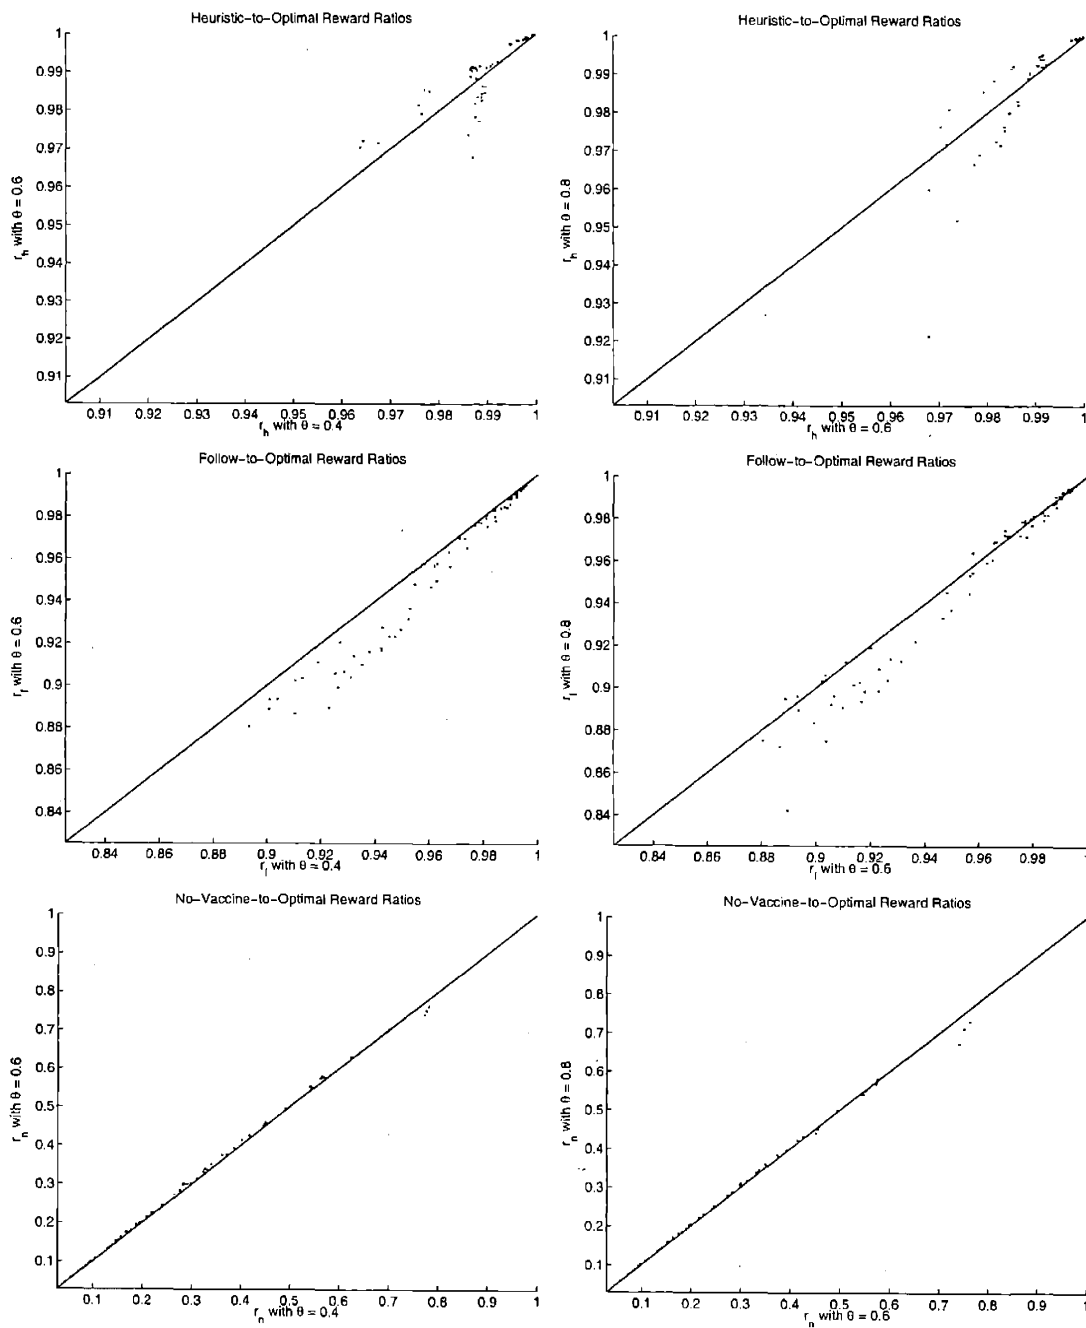


Figure 5-4: Performance of the heuristic, follow, and no-vaccine policy under different values of θ . If a data point lies above the line, this means that the performance of the y -axis case is better than that of the x -axis case, and vice-versa.

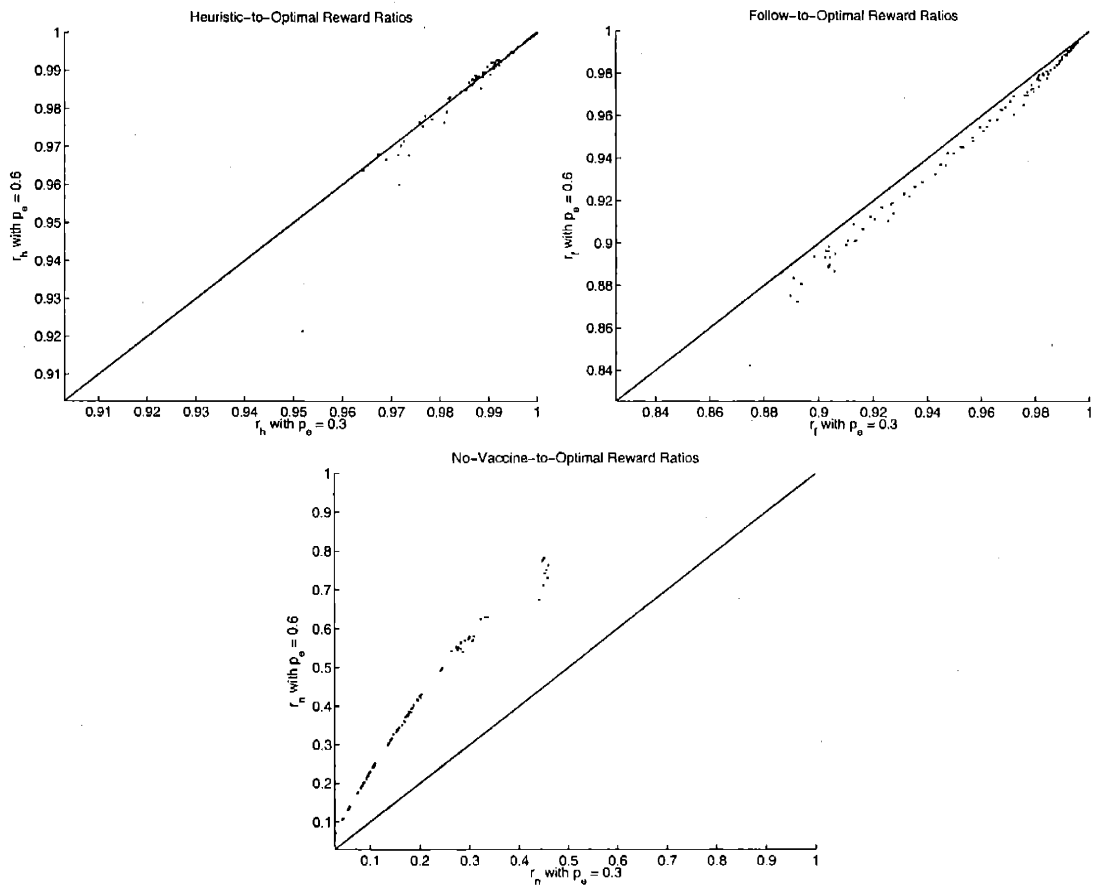


Figure 5-5: Performance of the heuristic, follow, and no-vaccine policy under different values of p_e . If a data point lies above the line, this means that the performance of the y -axis case is better than that of the x -axis case, and vice-versa.

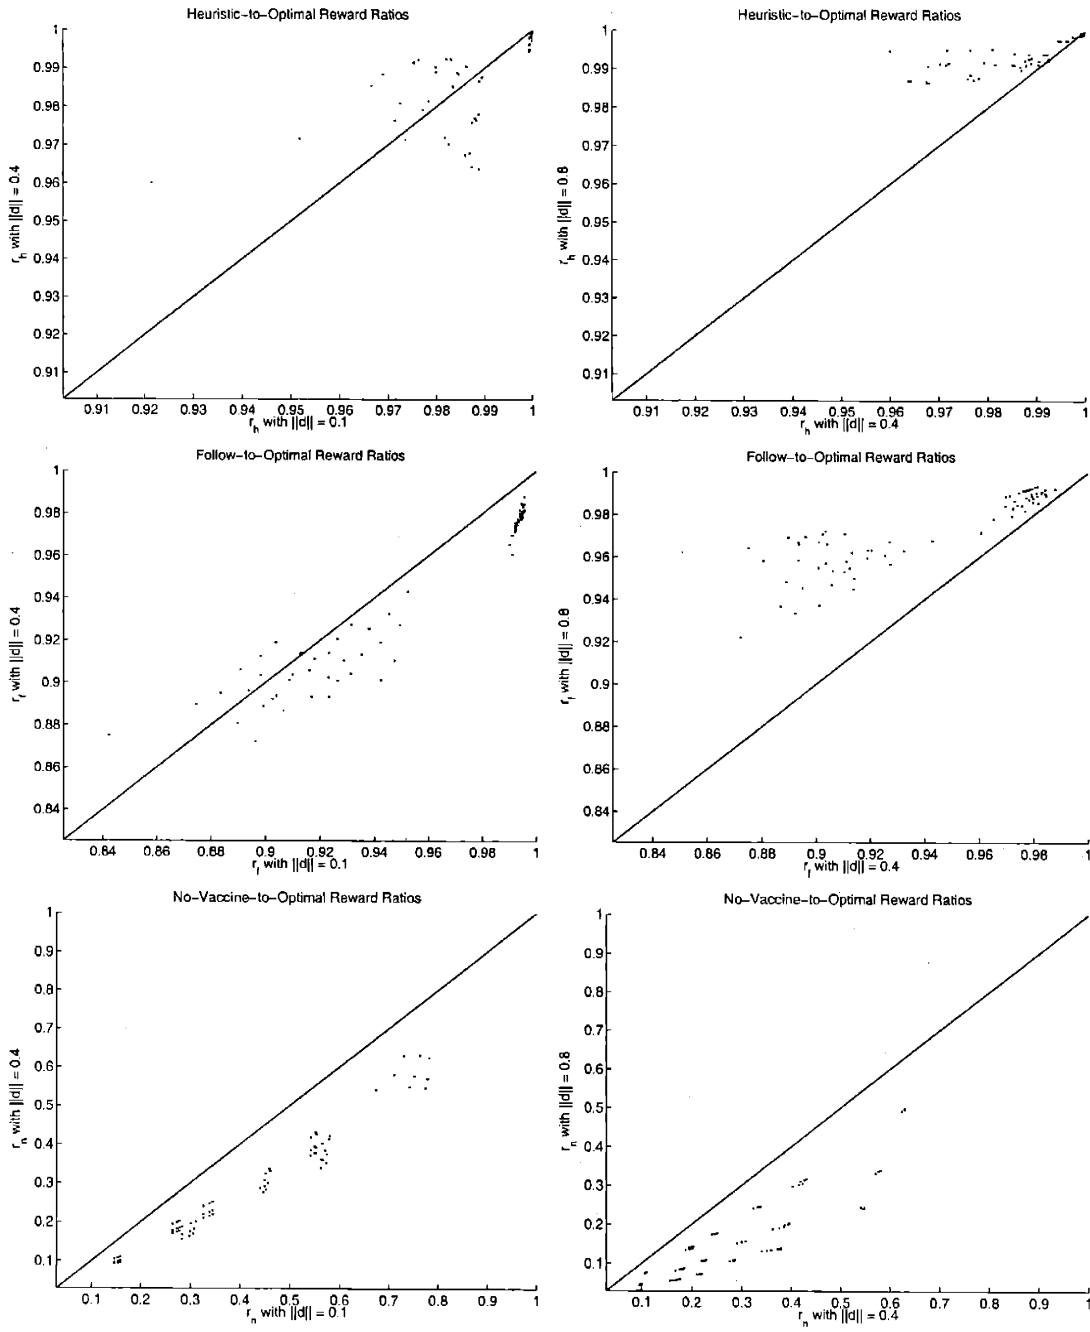


Figure 5-6: Performance of the heuristic, follow, and no-vaccine policy under different values of $\|d\|$. If a data point lies above the line, this means that the performance of the y -axis case is better than that of the x -axis case, and vice-versa.

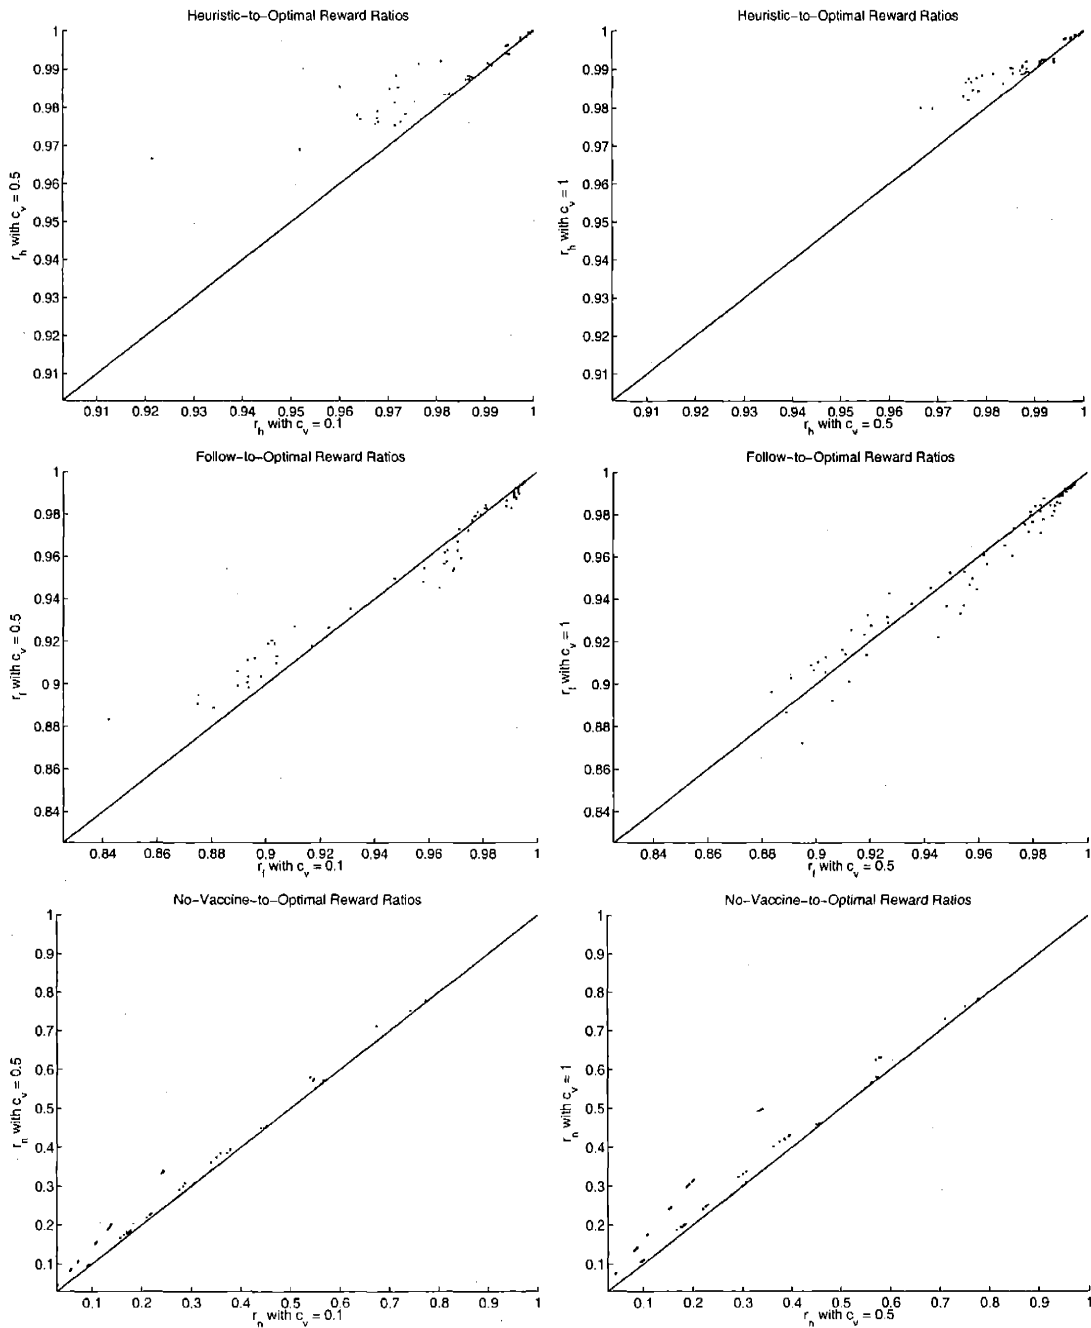


Figure 5-7: Performance of the heuristic, follow, and no-vaccine policy under different values of c_v . If a data point lies above the line, this means that the performance of the y -axis case is better than that of the x -axis case, and vice-versa.

5.3 Limitations of the Model and Recommendations on Vaccine Strain Selection

Before we summarize our results, let us first remark on some of the limitations of our model. Our model is a theoretical model that built upon the antigenic distance hypothesis, preliminary results from the flu shape space map, and some intuitions on the dynamics of the immune system. We have assumed that the dynamics of anti-flu humoral immunity are governed by the immune state transition equations. Given the complexity of the immune system, it has yet to be established that these equations provide a satisfactory description of the relevant immunodynamics. The continuous nature of the flu shape space (which is simply an n -dimensional Euclidean space) implies that the number of flu strains are uncountable, while in reality, the number of possible physical flu strains is almost certainly finite. We have also restricted the epidemic evolution to follow a random walk with a constant drift. While there is evidence from the flu shape space map that the epidemic evolution is a drifted process, there is no evidence to support (or contradict) that distances between two successive epidemic strains can be modelled as independent normal random variables. Finally, the reward in our model is a theoretical quantity (which is the clearance of epidemic viral particles by the post-vaccine immune state). In practice, however, we are most interested in the actual post-vaccine serological response against epidemic attacks. Although we expect the two to be monotonically increasing functions of each other, it would be ideal to construct a more clinically-based objective function by establishing a link between our theoretical model and measurable anti-flu humoral immunocompetence. In view of the above limitations, the validity of our model has yet to be tested by future work from immunology and the flu shape space map. Until then, the results presented in this work should be interpreted with caution.

Despite all the above-mentioned shortcomings, our model is the first analytical model to look at the feedback-control nature of the annual influenza vaccine selection problem. Let us now summarize our findings and draw some conclusions on the vaccine strain selection problem within the context of our model. We have shown that the heuristic policy, which is a robust near-optimal policy for the 1-period problem, is also an excellent and robust policy for the multi-period problem. This means that solving a 1-period problem will give us a vaccine strain that is almost as effective as that provided by the multi-period problem. The

follow policy, although less spectacular than the heuristic policy, performs quite well also in the multi-period problem. The performance of these two policies are rather insensitive to the values of the model parameters except α , which is the antibodies decay constant. In Section 5.2, we showed that when α is increased from 0.3 to 0.6, the value of r_h and r_f can decrease by as much as 7.85% and 15.1%, respectively. Therefore, in terms of both performance and robustness, the heuristic policy is the winner.

However, several factors deter the implementation of the heuristic policy and favor the follow policy over the heuristic policy. First of all, the heuristic policy requires detailed and accurate knowledge of anti-flu immunodynamics (e.g. values of α , θ , κ_v , etc.). Such information has yet to be determined precisely. In contrast, the follow policy requires only the prediction of epidemic strains, which is also one of the ingredients of the heuristic policy. Second of all, the heuristic policy prescribes different vaccine strains to people with different immunization histories. With a population that comprises a diverse set of immunization histories, the heuristic policy requires (1) tracking of personal immunization history (e.g., via blood tests), (2) production of a potentially wide collection of vaccines, and (3) efficient distribution of vaccines to regions with different vaccine demands (because of different composition of immunization histories in different regions). Therefore, implementation of the heuristic policy requires a whole new network of clinical, manufacturing, and logistical coordinations. The follow policy, on the other hand, is a one-size-fit-all policy that is exempt from such complementary setup. As a result, the cost of influenza vaccination will be much higher under the heuristic policy. Therefore, in terms of practicability and cost, the follow policy is the winner.

We have also shown that the performance of the no-vaccine policy varies widely depending on the parameter values of the model. In particular, the value of r_n can reach as low as 0.09 and as high as 0.78. Several factors favor the no-vaccine policy, which include (1) a slow antibodies decay rate, (2) a small antigenic drift, (3) a high probability of epidemic exposure, (4) low vaccine potency. Since we do not have concrete data on the parameter values of our model, we do not recommend the use of the no-vaccine policy because of its lack of robustness.

In summary, although the heuristic policy performs better and is more robust than the follow policy, the follow policy is much easier to implement and less expensive. From both the scientific and practical point of view, it would be interesting to see how much (clinically

measurable) improvements the heuristic policy can provide and whether its implementation cost is justified by the improvements. Before such evidence can be established, we believe that the follow policy, which is the idealized version of current WHO practice, is a cost-effective influenza vaccine strain selection policy.

Appendix A

Proofs and Formulas

A.1 Proof of Theorem 3.1

To ease our exposition, we rephrase the theorem into the following parts:

1. $\phi \in C^\infty([0, \infty))$ if $C \neq \frac{B_1}{1+B_1}$.
2. Let τ_0 be the smallest nonnegative value of τ such that $\frac{d\phi}{d\tau}(\tau) \leq 0$. Then τ_0 exists, $\tau_0 \leq \sqrt{\frac{1}{2}}$ and ϕ is strictly decreasing for all $\tau > \tau_0$. In particular, τ_0 is the global maximum of ϕ .
3. Let τ_1 be the smallest nonnegative value of τ such that $\frac{d\phi}{d\tau}(\tau) \leq 0$ and $\frac{d^2\phi}{d\tau^2}(\tau) \geq 0$. Then τ_1 exists, $\tau_1 \leq \sqrt{\frac{3}{2}}$ and ϕ is strictly convex for all $\tau > \tau_1$. Furthermore, $\tau_1 \geq \tau_0$.
4. $\phi(0) > 0$ if and only if $C > \frac{B_1}{1+B_1}$; $\frac{d\phi}{d\tau}(0) < 0$ if and only if $C > \frac{B_1}{1+B_1} e^{\frac{1}{2(1+B_1)}}$; $\frac{d^2\phi}{d\tau^2}(0) > 0$ if and only if $C > \frac{B_1}{1+B_1} e^{\frac{3}{2(1+2B_1)}}$;

For a given $\tau > 0$, $\phi(\tau)$ is the unique positive value of q that satisfies the equation

$$f(\tau, q) = C(B_1q + q + \tau) - B_1qe^{(q+\tau)^2} = 0. \quad (\text{A.1})$$

Multiplying (A.1) by $e^{-(q+\tau)^2}$ and taking the limit $\tau \rightarrow \infty$, we see that $\lim_{\tau \rightarrow \infty} \phi(\tau) = 0$. Moreover, we define $\phi(0) = \text{Re} \sqrt{\ln\left(\frac{C}{B_1} + C\right)}$. Note that $\phi(\tau) = 0$ only if $\tau = 0$.

Let us prove Part 1. Suppose $C \neq \frac{B_1}{1+B_1}$. Straightforward differentiation gives $\nabla_q f(\tau, q) = CB_1 + C - B_1e^{(q+\tau)^2}(1 + 2q^2 + 2\tau q)$. By rearranging terms in the equation $f(\tau, \phi(\tau)) = 0$, we obtain $B_1e^{(\phi(\tau)+\tau)^2} = C\left(B_1 + 1 + \frac{\tau}{\phi(\tau)}\right)$ for all $\tau > 0$. Substituting this result into

$\nabla_q f(\tau, \phi(\tau))$, we get $\nabla_q f(\tau, \phi(\tau)) = -C \frac{\tau}{\phi(\tau)} - 2C\phi(\tau)(\phi(\tau) + \tau) \left(B_1 + 1 + \frac{\tau}{\phi(\tau)} \right) < 0$ for all $\tau > 0$. At $\tau = 0$, we have $\phi(0) = 0$ if $C < \frac{B_1}{1+B_1}$ and $\phi(0) = \sqrt{\ln\left(\frac{C}{B_1} + C\right)}$ if $C > \frac{B_1}{1+B_1}$. In the first case, $\nabla_q f(0, \phi(0)) = CB_1 + C - B_1 < 0$. In the second case, $\nabla_q f(0, \phi(0)) = -2C(B_1 + 1) \ln\left(\frac{C}{B_1} + C\right) < 0$. Thus, we conclude that $\nabla_q f(\tau, \phi(\tau)) < 0$ for all $\tau \geq 0$. Since $f \in C^\infty$, the Implicit Function Theorem guarantees that $\phi \in C^\infty([0, \infty))$.

Let us now prove Part 2. Differentiating $f(\tau, \phi(\tau)) = 0$ with respect to τ , we obtain

$$\frac{d\phi}{d\tau}(\tau) = -\frac{2\phi(\tau)(\phi(\tau) + \tau) - \frac{C}{B_1}e^{-(\phi(\tau)+\tau)^2}}{2\phi(\tau)(\phi(\tau) + \tau) - \frac{C}{B_1}e^{-(\phi(\tau)+\tau)^2} + (1 - Ce^{-(\phi(\tau)+\tau)^2})}. \quad (\text{A.2})$$

By rearranging the terms in the equation $f(\tau, \phi(\tau)) = 0$, we can see that $Ce^{-(\phi(\tau)+\tau)^2} = \frac{B_1\phi(\tau)}{B_1\phi(\tau)+\phi(\tau)+\tau}$ whenever $\phi(\tau) > 0$. Substituting this relation into (A.2), we get

$$\frac{d\phi}{d\tau}(\tau) = \begin{cases} \frac{\frac{C}{B_1}}{1 - \frac{C}{B_1} - C} & \text{at } \tau = 0 \text{ if } \phi(0) = 0, \\ -\frac{2\phi(\tau)(\phi(\tau)+\tau)(B_1\phi(\tau)+\phi(\tau)+\tau) - \phi(\tau)}{2\phi(\tau)(\phi(\tau)+\tau)(B_1\phi(\tau)+\phi(\tau)+\tau) + \tau} = -\frac{p(\tau, \phi(\tau))}{l(\tau, \phi(\tau))} & \text{otherwise.} \end{cases} \quad (\text{A.3})$$

We claim that there exists at most one $\tau \in [0, \infty)$ such that $\frac{d\phi}{d\tau}(\tau) = 0$. Clearly, $\frac{d\phi}{d\tau}(\tau) = 0$ if and only if $p(\tau, \phi(\tau)) = 0$. By solving the equation $p(\tau, q) = 0$ for q , we find that $q = \psi(\tau) = \frac{-\tau(2+B_1) + \sqrt{\tau^2 B_1^2 + 2B_1 + 2}}{2(1+B_1)}$ is the only root of $p(\tau, q) = 0$ that can be positive. Moreover, it can be easily verified that $\frac{d\psi}{d\tau}(\tau) < 0$ for all $\tau \geq 0$, so ψ is a one-to-one mapping. It then follows that there can be at most one $\tau \in [0, \infty)$ such that $\frac{d\phi}{d\tau}(\tau) = 0$, thus verifying our claim. Note that if $\frac{d\phi}{d\tau}(\tau) = 0$ at some $\tau \geq 0$, then $\psi(\tau) = \phi(\tau)$. Since $\phi(\tau)$ is nonnegative for all $\tau \geq 0$ and $\lim_{\tau \rightarrow \infty} \phi(\tau) = 0$, we must have $\lim_{\tau \rightarrow \infty} \frac{d\phi}{d\tau}(\tau) = 0$ with $\frac{d\phi}{d\tau}(\tau) < 0$ for all $\tau > \tilde{\tau}$ for some large enough $\tilde{\tau}$. Combining this with the fact that $\frac{d\phi}{d\tau}(\tau) = 0$ for at most one $\tau \in [0, \infty)$, we conclude that (i) $\frac{d\phi}{d\tau}(\tau) < 0$ for all $\tau > 0$ if $\frac{d\phi}{d\tau}(0) \leq 0$ and (ii) $\phi(\tau)$ attains a unique global maximum at $\tau_{max} > 0$ (with $\frac{d\phi}{d\tau}(\tau_{max}) = 0$) if $\frac{d\phi}{d\tau}(0) > 0$. Thus, τ_0 as defined in the theorem is

$$\tau_0 = \begin{cases} 0 & \text{if } \frac{d\phi}{d\tau}(0) \leq 0; \\ \tau_{max} & \text{if } \frac{d\phi}{d\tau}(0) > 0. \end{cases}$$

It remains to show that $\tau_0 \leq \sqrt{\frac{1}{2}}$. Clearly, it suffices to show that $\tau_{max} \leq \sqrt{\frac{1}{2}}$. At τ_{max} ,

we have $\phi(\tau_{max}) = \psi(\tau_{max}) \geq 0$. Straightforward calculations show that $\psi(\frac{1}{\sqrt{2}}) = 0$. Since $\frac{d\psi}{d\tau}(\tau) < 0$ for all $\tau \geq 0$, it follows that $\psi(\tau) \geq 0$ only if $\tau \leq \sqrt{\frac{1}{2}}$. Thus, $\tau_{max} \leq \sqrt{\frac{1}{2}}$. This completes the proof of Part 2.

We now move on to Part 3. We claim that there exists at most one $\tau \in (0, \infty)$ such that $\frac{d^2\phi}{d\tau^2}(\tau) = 0$. Straightforward differentiation gives

$$\frac{d^2\phi}{d\tau^2}(\tau) = \begin{cases} 0 & \text{at } \tau = 0 \text{ if } \phi(0) = 0, \\ \frac{2\phi(\tau)(\phi(\tau)+\tau)^2(B_1\phi(\tau)+\phi(\tau)+\tau)h(\tau,\phi(\tau))}{(2B_1\phi(\tau)^3+2\phi(\tau)^3+4\phi(\tau)^2\tau+2\tau B_1\phi(\tau)^2+2\phi(\tau)\tau^2+\tau)^3} & \text{otherwise,} \end{cases}$$

where $h(\tau, q) = 2\tau^2 + 4\tau B_1 q + 4q\tau + 2q^2 - 3 + 4B_1 q^2$. Clearly, $\frac{d^2\phi}{d\tau^2}(\tau) = 0$ at $\tau > 0$ if and only if $h(\tau, \phi(\tau)) = 0$. Solving $h(\tau, q) = 0$ for q shows that $q = \lambda(\tau) = \frac{-2\tau(B_1+1) + \sqrt{4B_1^2\tau^2 + 12B_1 + 6}}{2(1+2B_1)}$ is the only root of $h(\tau, q) = 0$ that can be positive. Moreover, it can be easily verified that $\frac{d\lambda}{d\tau}(\tau) < 0$ for all $\tau > 0$, so λ is a one-to-one mapping. It then follows that there exists at most one $\tau \in (0, \infty)$ such that $\frac{d^2\phi}{d\tau^2}(\tau) = 0$, thus verifying our claim. Note that if $\frac{d^2\phi}{d\tau^2}(\tau) = 0$ at some $\tau > 0$, then $\lambda(\tau) = \phi(\tau)$. Since $\lim_{\tau \rightarrow \infty} \frac{d\phi}{d\tau}(\tau) = 0$ with $\frac{d\phi}{d\tau}(\tau) < 0$ for all $\tau > \tilde{\tau}$ for some large enough $\tilde{\tau}$ (from Part 2), we must have $\frac{d^2\phi}{d\tau^2}(\tau) > 0$ for all $\tau > \hat{\tau}$ for some large enough $\hat{\tau}$. Moreover, if ϕ has a global maximum at τ_{max} , then $\frac{d^2\phi}{d\tau^2}(\tau_{max}) < 0$ and ϕ must have a unique point of inflection at some $\tau_{inflect} > \tau_{max}$ (this follows from our verified claim). Note that $\frac{d\phi}{d\tau}(\tau_{inflect}) < 0$. Combining all these observations with Part 2, we conclude that (i) $\frac{d\phi}{d\tau}(\tau) < 0$ and $\frac{d^2\phi}{d\tau^2}(\tau) > 0$ for all $\tau > 0$ if $\frac{d\phi}{d\tau}(0) < 0$ and $\frac{d^2\phi}{d\tau^2}(0) \geq 0$, and (ii) ϕ has a unique point of inflection at $\tau_{inflect} > 0$ (with $\frac{d\phi}{d\tau}(\tau_{inflect}) < 0$ and $\frac{d^2\phi}{d\tau^2}(\tau_{inflect}) = 0$) otherwise. Thus, τ_1 as defined in the theorem is

$$\tau_1 = \begin{cases} 0 & \text{if } \frac{d\phi}{d\tau}(0) < 0 \text{ and } \frac{d^2\phi}{d\tau^2}(0) \geq 0; \\ \tau_{inflect} & \text{otherwise.} \end{cases}$$

Let us now show $\tau_1 \leq \sqrt{\frac{3}{2}}$. It suffices to show that $\tau_{inflect} \leq \sqrt{\frac{3}{2}}$. At $\tau_{inflect}$, we have $\lambda(\tau_{inflect}) = \phi(\tau_{inflect})$. Simple calculations show that $\lambda(\sqrt{\frac{3}{2}}) = 0$. Since $\frac{d\lambda}{d\tau}(\tau) < 0$ for all $\tau \geq 0$, it follows that $\lambda(\tau) \geq 0$ only if $\tau \leq \sqrt{\frac{3}{2}}$, which implies $\tau_{inflect} \leq \sqrt{\frac{3}{2}}$. It remains to show that $\tau_1 \geq \tau_0$. This is trivially true if $\tau_0 = 0$. If $\tau_0 = \tau_{max}$, then $\tau_1 = \tau_{inflect} > \tau_{max}$ as we have already shown. Thus, the proof of Part 3 is complete.

Finally, we turn to the proof of Part 4. Clearly, $\phi(0) > 0$ if and only if $\ln\left(\frac{C}{B_1} + C\right) > 0$,

which is equivalent to $C > \frac{B_1}{1+B_1}$. Next, we notice that $\frac{d\phi}{d\tau}(0) < 0$ only if $\phi(0) > 0$. Straightforward calculations give $\frac{d\phi}{d\tau}(0) = \frac{1-2(1+B_1)\ln\left(\frac{C}{B_1}+C\right)}{(1+B_1)\ln\left(\frac{C}{B_1}+C\right)}$ if $\phi(0) > 0$, so $\frac{d\phi}{d\tau}(0) < 0$ is the same as $C > \frac{B_1}{1+B_1}e^{\frac{1}{2(1+B_1)}}$. Similarly, $\frac{d^2\phi}{d\tau^2}(0) > 0$ only if $\phi(0) > 0$, and $\frac{d^2\phi}{d\tau^2}(0) = \frac{\sqrt{\ln\left(\frac{C}{B_1}+C\right)}\left(3-2(1+2B_1)\ln\left(\frac{C}{B_1}+C\right)\right)}{4(1+B_1)^2\ln\left(\frac{C}{B_1}+C\right)}$ if $\phi(0) > 0$. Hence, $\frac{d^2\phi}{d\tau^2}(0) > 0$ is the same as $C > \frac{B_1}{1+B_1}e^{\frac{3}{2(1+2B_1)}}$. \square

A.2 Proof of Theorem 3.2

From the Part 2 of the proof of Theorem 3.1, we see that ϕ is a solution to the ODE system specified by the initial condition $\phi(0) = \text{Re}\sqrt{\ln\left(\frac{C}{B_1}+C\right)}$ and the differential equation (A.3). Moreover, ϕ is the unique solution to this ODE system since the right-hand-side of the differential equation is continuously differentiable.

Suppose we vary C while holding B_1 fixed, $\hat{C} > 0$, $\tilde{C} > 0$ and $\hat{C} < \tilde{C}$. Clearly, $\phi_{\hat{C}}(0) \leq \phi_{\tilde{C}}(0)$. If $\phi_{\hat{C}}(0) = \phi_{\tilde{C}}(0) = 0$, then we must have $\hat{C} < \frac{B_1}{1+B_1}$ and $\tilde{C} < \frac{B_1}{1+B_1}$. It then follows that $0 < \frac{d\phi_{\hat{C}}}{d\tau}(0) = \frac{\frac{1}{B_1}}{\frac{1}{\hat{C}} - \frac{1}{B_1} - 1} < \frac{\frac{1}{B_1}}{\frac{1}{\tilde{C}} - \frac{1}{B_1} - 1} = \frac{d\phi_{\tilde{C}}}{d\tau}(0)$. Therefore, for small enough $\epsilon > 0$, $0 < \phi_{\hat{C}}(\tau) < \phi_{\tilde{C}}(\tau)$ for all $\tau \in (0, \epsilon]$. At $\tau = \epsilon$, both $\phi_{\hat{C}}$ and $\phi_{\tilde{C}}$ are governed by the same differential equation (A.3) since B_1 is fixed. Thus, $\phi_{\hat{C}}(\tau) \leq \phi_{\tilde{C}}(\tau)$ for all $\tau > \epsilon$, since otherwise there exist $\tilde{\tau} > \epsilon$ and $\delta > 0$ such that $\phi_{\hat{C}}$ intersects $\phi_{\tilde{C}}$ at $\tau = \tilde{\tau}$ with $\phi_{\hat{C}}(\tau) > \phi_{\tilde{C}}(\tau)$ for $\tau \in (\tilde{\tau}, \tilde{\tau} + \delta)$, which contradicts the uniqueness of the solution to the ODE system with initial condition $\phi(\tilde{\tau}) = \phi_{\tilde{C}}(\tilde{\tau}) = \phi_{\hat{C}}(\tilde{\tau})$. If $\phi_{\hat{C}}(0) < \phi_{\tilde{C}}(0)$, the same arguments are again valid. Hence, $\phi_{\hat{C}}(\tau) \leq \phi_{\tilde{C}}(\tau)$ for all $\tau \geq 0$.

Now suppose we vary C while holding $\frac{C}{B_1}$ fixed, $\hat{C} > 0$, $\tilde{C} > 0$ and $\hat{C} < \tilde{C}$. The arguments below (A.2) show that (A.3) can be written as

$$\frac{d\phi}{d\tau}(\tau) = \begin{cases} \frac{\frac{C}{B_1}}{1 - \frac{C}{B_1} - C} & \text{at } \tau = 0 \text{ if } \phi(0) = 0; \\ -\frac{2\phi(\tau)^2(\phi(\tau)+\tau) - \frac{C}{B_1}\phi(\tau)e^{-(\phi(\tau)+\tau)^2}}{2\phi(\tau)^2(\phi(\tau)+\tau) + \frac{C}{B_1}\tau e^{-(\phi(\tau)+\tau)^2}} & \text{otherwise.} \end{cases} \quad (\text{A.4})$$

From (A.4), we see that the if $\phi(0) > 0$, the differential equation is invariant for all $\tau \geq 0$ as long as $\frac{C}{B_1}$ is fixed. Thus, we can repeat the arguments in the previous paragraph to assert that $\phi_{\hat{C}}(\tau) \leq \phi_{\tilde{C}}(\tau)$ for all $\tau \geq 0$. \square

A.3 Closed Form Formulas for Reward-To-Go Approximations

Using the approximation algorithm given in Chapter 4, the approximate objective function at time t is

$$E_{I_{e_t}, e_t} \left[\bar{c}(y_t, e_t) + \tilde{J}_{t+1}^*(h_{t+1}) \mid e_{t-1} \right] \quad (\text{A.5})$$

where

$$\begin{aligned} \tilde{J}_{t+1}^*(h_{t+1}) &= \sum_{k=1}^{N-t} E_{e_{t+k}} \left[\beta_{t+k} \alpha^{k-1} \bar{c}(x_{t+1}, e_{t+k}) \mid e_t \right] \\ &\quad + C_{t+1,1}^* + C_{t+1,2}^* \delta_{t+1} + C_{t+1,3}^* \delta_{t+1}^2 + C_{t+1,4}^* \|G_{t+1}\|^2. \end{aligned} \quad (\text{A.6})$$

Our goal is to get a closed form expression for the approximate objective function in (A.5). We will be done if we can evaluate the expected value of \tilde{J}_{t+1}^* in closed form. Let us begin by expanding \tilde{J}_{t+1}^* as follows:

$$\begin{aligned} &E_{I_{e_t}, e_t} \left[\tilde{J}_{t+1}^*(h_{t+1}, e_t) \mid e_{t-1} \right] \\ &= E_{I_{e_t}, e_t} \left[\sum_{k=1}^{N-t} E_{e_{t+k}} \left[\beta_{t+k} \alpha^{k-1} \bar{c}(x_{t+1}, e_{t+k}) \mid e_t \right] \right. \\ &\quad \left. + C_{t+1,1}^* + C_{t+1,2}^* \delta_{t+1} + C_{t+1,3}^* \delta_{t+1}^2 + C_{t+1,4}^* \|G_{t+1}\|^2 \mid e_{t-1} \right] \\ &= \sum_{k=1}^{N-t} \beta_{t+k} \sum_{a \in h_t} \alpha^{t+k-p(a)} D_a E_{e_{t+k}} \left[g(a - e_{t+k}) \mid e_{t-1} \right] \\ &\quad + \sum_{k=1}^{N-t} \beta_{t+k} \alpha^k D_{v_t} E_{e_{t+k}} \left[g(v_t - e_{t+k}) \mid e_{t-1} \right] \\ &\quad + p_e \sum_{k=1}^{N-t} \beta_{t+k} \alpha^k E_{e_t} \left[D_{e_t} \mid e_{t-1} \right] E_{e_t, e_{t+k}} \left[g(e_t - e_{t+k}) \mid e_{t-1} \right] \\ &\quad + E_{I_{e_t}, e_t} \left[C_{t+1,1}^* + C_{t+1,2}^* \delta_{t+1} + C_{t+1,3}^* \delta_{t+1}^2 + C_{t+1,4}^* \|G_{t+1}\|^2 \mid e_{t-1} \right]. \end{aligned} \quad (\text{A.7})$$

Simple calculations give $E_{e_t, e_{t+k}} [g(e_t - e_{t+k}) | e_{t-1}] = \frac{\theta}{\kappa_v} B_k^{n/2} e^{-\eta B_k \|kd\|^2/2}$. So our next step is to compute $E_{I_{e_t}, e_t} [\delta_{t+1} | e_{t-1}]$, $E_{I_{e_t}, e_t} [\delta_{t+1}^2 | e_{t-1}]$, and $E_{I_{e_t}, e_t} [\|G_{t+1}\|^2 | e_{t-1}]$:

$$\begin{aligned}
& E_{I_{e_t}, e_t} [\delta_{t+1} | e_{t-1}] \\
&= \frac{\kappa_v}{\theta} E_{I_{e_t}, e_t} \left[\sum_{a \in \{h_t, v_t\}} \alpha^{t+1-p(a)} D_a g(a - \bar{e}_{t+1}) + \alpha I_{e_t} D_{e_t} g(e_t - \bar{e}_{t+1}) | e_{t-1} \right] \\
&= \frac{\kappa_v}{\theta} \left(\sum_{a \in h_t} \alpha^{t+1-p(a)} D_a E_{e_t} [g(a - e_t - d) | e_{t-1}] \right. \\
&\quad \left. + \alpha D_{v_t} E_{e_t} [g(v_t - e_t - d) | e_{t-1}] + \alpha p_e E_{e_t} [D_{e_t} | e_{t-1}] g(d) \right). \tag{A.8}
\end{aligned}$$

$$\begin{aligned}
& E_{I_{e_t}, e_t} [\delta_{t+1}^2 | e_{t-1}] \\
&= \left(\frac{\kappa_v}{\theta} \right)^2 \left(\sum_{a, b \in h_t} \alpha^{2(t+1)-p(a)-p(b)} D_a D_b E_{e_t} [g(a - e_t - d) g(b - e_t - d) | e_{t-1}] \right. \\
&\quad + 2 \sum_{a \in h_t} \alpha^{t+2-p(a)} D_a D_{v_t} E_{e_t} [g(a - e_t - d) g(v_t - e_t - d) | e_{t-1}] \\
&\quad + \alpha^2 D_{v_t}^2 E_{e_t} [g(v_t - e_t - d)^2 | e_{t-1}] + 2p_e \sum_{a \in h_t} \alpha^{t+2-p(a)} D_a g(d) E_{e_t} [g(a - e_t - d) D_{e_t} | e_{t-1}] \\
&\quad \left. + 2p_e \alpha^2 D_{v_t} g(d) E_{e_t} [g(v_t - e_t - d) D_{e_t} | e_{t-1}] + p_e \alpha^2 g(d)^2 E_{e_t} [D_{e_t}^2 | e_{t-1}] \right). \tag{A.9}
\end{aligned}$$

$$\begin{aligned}
& E_{I_{e_t}, e_t} [\|G_{t+1}\|^2 | e_{t-1}] \\
&= \left(\frac{\kappa_v}{\theta} \right)^2 \left(\sum_{a, b \in h_t} \alpha^{2(t+1)-p(a)-p(b)} D_a D_b E_{e_t} [(e_t + d - a)^T (e_t + d - b) g(a - e_t - d) g(b - e_t - d) | e_{t-1}] \right. \\
&\quad + 2 \sum_{a \in h_t} \alpha^{t+2-p(a)} D_a D_{v_t} E_{e_t} [(e_t + d - a)^T (e_t + d - v_t) g(a - e_t - d) g(v_t - e_t - d) | e_{t-1}] \\
&\quad + \alpha^2 D_{v_t}^2 E_{e_t} [\|e_t + d - v_t\|^2 g(v_t - e_t - d)^2 | e_{t-1}] \\
&\quad + 2p_e \sum_{a \in h_t} \alpha^{t+2-p(a)} D_a g(d) E_{e_t} [d^T (e_t + d - a) g(a - e_t - d) D_{e_t} | e_{t-1}] \\
&\quad + 2p_e \alpha^2 D_{v_t} g(d) E_{e_t} [d^T (e_t + d - v_t) g(v_t - e_t - d) D_{e_t} | e_{t-1}] \\
&\quad \left. + p_e \alpha^2 g(d)^2 \|d\|^2 E_{e_t} [D_{e_t}^2 | e_{t-1}] \right). \tag{A.10}
\end{aligned}$$

We now compute $E_{e_t} [D_{e_t} | e_{t-1}]$ from (A.8) and $E_{e_t} [D_{e_t}^2 | e_{t-1}]$ from (A.10). We adopt the following approximations:

$$D_{e_t} \approx \kappa_e \left(1 - \bar{c}(y_t, e_t)\right). \quad (\text{A.11})$$

That is, we ignore the minimum operator in the definition of D_{e_t} . Under this approximation, we have

$$\begin{aligned} E_{e_t} [D_{e_t}^2 | e_{t-1}] &= E_{e_t} \left[\kappa_e^2 \left(1 - \bar{c}(y_t, e_t)\right)^2 | e_{t-1} \right] \\ &= \kappa_e^2 \left(E_{e_t} \left[1 - 2\bar{c}(y_t, e_t) | e_{t-1} \right] \right. \\ &\quad + \sum_{a, b \in h_t} \alpha^{2t-p(a)-p(b)} D_a D_b E_{e_t} \left[g(a - e_t) g(b - e_t) | e_{t-1} \right] \\ &\quad + 2 \sum_{a \in h_t} \alpha^{t-p(a)} D_a D_{v_t} E_{e_t} \left[g(a - e_t) g(v_t - e_t) | e_{t-1} \right] \\ &\quad \left. + D_{v_t}^2 E_{e_t} \left[g(v_t - e_t)^2 | e_{t-1} \right] \right) \end{aligned} \quad (\text{A.12})$$

Next, we evaluate $E_{e_t} [g(a - e_t - d) D_{e_t} | e_{t-1}]$ from (A.9):

$$\begin{aligned} &E_{e_t} [g(a - e_t - d) D_{e_t} | e_{t-1}] \\ &= E_{e_t} [g(a - e_t - d) \kappa_e (1 - \bar{c}(y_t, e_t)) | e_{t-1}] \\ &= \kappa_e \left(E_{e_t} [g(a - e_t - d) | e_{t-1}] - \sum_{b \in h_t} \alpha^{t-p(b)} D_b E_{e_t} [g(a - e_t - d) g(b - e_t) | e_{t-1}] \right. \\ &\quad \left. - D_{v_t} E_{e_t} [g(a - e_t - d) g(v_t - e_t) | e_{t-1}] \right). \end{aligned} \quad (\text{A.13})$$

Now $E_{e_t} [d^T (e_t + d - a) g(a - e_t - d) D_{e_t} | e_{t-1}]$ from (A.10):

$$\begin{aligned} &E_{e_t} [d^T (e_t + d - a) g(a - e_t - d) D_{e_t} | e_{t-1}] \\ &= E_{e_t} [d^T (e_t + d - a) g(a - e_t - d) \kappa_e (1 - \bar{c}(y_t, e_t)) | e_{t-1}] \\ &= \kappa_e \left(E_{e_t} [d^T (e_t + d - a) g(a - e_t - d) | e_{t-1}] \right. \\ &\quad - \sum_{b \in h_t} \alpha^{t-p(b)} D_b E_{e_t} [d^T (e_t + d - a) g(a - e_t - d) g(b - e_t) | e_{t-1}] \\ &\quad \left. - D_{v_t} E_{e_t} [d^T (e_t + d - a) g(a - e_t - d) g(v_t - e_t) | e_{t-1}] \right). \end{aligned} \quad (\text{A.14})$$

Finally, a few other terms that appear in the above equations:

$$\begin{aligned}
& E_{e_t} \left[g(a - e_t)g(b - e_t) \mid e_{t-1} \right] \\
&= \frac{\theta^2}{\kappa_v^2} \int_S e^{-\eta\|a-e_t\|^2/2} e^{-\eta\|b-e_t\|^2/2} \left(\frac{1}{\sqrt{2\pi\sigma}} \right)^n e^{-\|e_t-\bar{e}_t\|^2/2\sigma^2} ds \\
&= \frac{\theta^2}{\kappa_v^2} B_2^{n/2} e^{-\eta B_2(\eta\sigma^2\|a-b\|^2 + \|a-\bar{e}_t\|^2 + \|b-\bar{e}_t\|^2)/2}. \tag{A.15}
\end{aligned}$$

$$\begin{aligned}
& E_{e_t} \left[(e_t - a)^T (e_t - b)g(a - e_t)g(b - e_t) \mid e_{t-1} \right] \\
&= \frac{\theta^2}{\kappa_v^2} \int_S (e_t - a)^T (e_t - b) e^{-\eta\|a-e_t\|^2/2} e^{-\eta\|b-e_t\|^2/2} \left(\frac{1}{\sqrt{2\pi\sigma}} \right)^n e^{-\|e_t-\bar{e}_t\|^2/2\sigma^2} ds \\
&= \frac{\theta^2}{\kappa_v^2} \left[B_2(\bar{e}_t - (\eta\sigma^2 + 1)a + \eta\sigma^2 b)^T (\bar{e}_t - (\eta\sigma^2 + 1)b + \eta\sigma^2 a) + n\sigma^2 \right] \\
&\quad B_2^{n/2+1} e^{-\eta B_2(\eta\sigma^2\|a-b\|^2 + \|a-\bar{e}_t\|^2 + \|b-\bar{e}_t\|^2)/2}. \tag{A.16}
\end{aligned}$$

$$E_{e_t} \left[d^T (e_t - a)g(a - e_t) \mid e_{t-1} \right] = \frac{\theta}{\kappa_v} d^T (\bar{e}_t - a) B_1^{n/2+1} e^{-\eta B_1\|\bar{e}_t-a\|^2/2}. \tag{A.17}$$

$$\begin{aligned}
& E_{e_t} \left[d^T (e_t - a)g(a - e_t)g(b - e_t) \mid e_{t-1} \right] \\
&= \frac{\theta^2}{\kappa_v^2} d^T (\bar{e}_t - (\eta\sigma^2 + 1)a + \eta\sigma^2 b) B_2^{n/2+1} e^{-\eta B_2(\eta\sigma^2\|a-b\|^2 + \|a-\bar{e}_t\|^2 + \|b-\bar{e}_t\|^2)/2}. \tag{A.18}
\end{aligned}$$

A.4 Parameter Values in Case Studies

Case	κ_v	α	θ	p_e	$\ d\ $	c_v	Case	κ_v	α	θ	p_e	$\ d\ $	c_v
1	0.3	0.3	0.4	0.3	0.1	0.1	37	0.3	0.6	0.8	0.3	0.1	0.1
2	0.3	0.3	0.4	0.3	0.1	0.5	38	0.3	0.6	0.8	0.3	0.1	0.5
3	0.3	0.3	0.4	0.3	0.1	1.0	39	0.3	0.6	0.8	0.3	0.1	1.0
4	0.3	0.3	0.4	0.3	0.4	0.1	40	0.3	0.6	0.8	0.3	0.4	0.1
5	0.3	0.3	0.4	0.3	0.4	0.5	41	0.3	0.6	0.8	0.3	0.4	0.5
6	0.3	0.3	0.4	0.3	0.4	1.0	42	0.3	0.6	0.8	0.3	0.4	1.0
7	0.3	0.3	0.4	0.6	0.1	0.1	43	0.3	0.6	0.8	0.6	0.1	0.1
8	0.3	0.3	0.4	0.6	0.1	0.5	44	0.3	0.6	0.8	0.6	0.1	0.5
9	0.3	0.3	0.4	0.6	0.1	1.0	45	0.3	0.6	0.8	0.6	0.1	1.0
10	0.3	0.3	0.4	0.6	0.4	0.1	46	0.3	0.6	0.8	0.6	0.4	0.1
11	0.3	0.3	0.4	0.6	0.4	0.5	47	0.3	0.6	0.8	0.6	0.4	0.5
12	0.3	0.3	0.4	0.6	0.4	1.0	48	0.3	0.6	0.8	0.6	0.4	1.0
13	0.3	0.3	0.8	0.3	0.1	0.1	49	0.6	0.3	0.4	0.3	0.1	0.1
14	0.3	0.3	0.8	0.3	0.1	0.5	50	0.6	0.3	0.4	0.3	0.1	0.5
15	0.3	0.3	0.8	0.3	0.1	1.0	51	0.6	0.3	0.4	0.3	0.1	1.0
16	0.3	0.3	0.8	0.3	0.4	0.1	52	0.6	0.3	0.4	0.3	0.4	0.1
17	0.3	0.3	0.8	0.3	0.4	0.5	53	0.6	0.3	0.4	0.3	0.4	0.5
18	0.3	0.3	0.8	0.3	0.4	1.0	54	0.6	0.3	0.4	0.3	0.4	1.0
19	0.3	0.3	0.8	0.6	0.1	0.1	55	0.6	0.3	0.4	0.6	0.1	0.1
20	0.3	0.3	0.8	0.6	0.1	0.5	56	0.6	0.3	0.4	0.6	0.1	0.5
21	0.3	0.3	0.8	0.6	0.1	1.0	57	0.6	0.3	0.4	0.6	0.1	1.0
22	0.3	0.3	0.8	0.6	0.4	0.1	58	0.6	0.3	0.4	0.6	0.4	0.1
23	0.3	0.3	0.8	0.6	0.4	0.5	59	0.6	0.3	0.4	0.6	0.4	0.5
24	0.3	0.3	0.8	0.6	0.4	1.0	60	0.6	0.3	0.4	0.6	0.4	1.0
25	0.3	0.6	0.4	0.3	0.1	0.1	61	0.6	0.3	0.8	0.3	0.1	0.1
26	0.3	0.6	0.4	0.3	0.1	0.5	62	0.6	0.3	0.8	0.3	0.1	0.5
27	0.3	0.6	0.4	0.3	0.1	1.0	63	0.6	0.3	0.8	0.3	0.1	1.0
28	0.3	0.6	0.4	0.3	0.4	0.1	64	0.6	0.3	0.8	0.3	0.4	0.1
29	0.3	0.6	0.4	0.3	0.4	0.5	65	0.6	0.3	0.8	0.3	0.4	0.5
30	0.3	0.6	0.4	0.3	0.4	1.0	66	0.6	0.3	0.8	0.3	0.4	1.0
31	0.3	0.6	0.4	0.6	0.1	0.1	67	0.6	0.3	0.8	0.6	0.1	0.1
32	0.3	0.6	0.4	0.6	0.1	0.5	68	0.6	0.3	0.8	0.6	0.1	0.5
33	0.3	0.6	0.4	0.6	0.1	1.0	69	0.6	0.3	0.8	0.6	0.1	1.0
34	0.3	0.6	0.4	0.6	0.4	0.1	70	0.6	0.3	0.8	0.6	0.4	0.1
35	0.3	0.6	0.4	0.6	0.4	0.5	71	0.6	0.3	0.8	0.6	0.4	0.5
36	0.3	0.6	0.4	0.6	0.4	1.0	72	0.6	0.3	0.8	0.6	0.4	1.0

Table A.1: Parameter values in DP cases

Case	κ_v	α	θ	p_e	$\ d\ $	c_v	Case	κ_v	α	θ	p_e	$\ d\ $	c_v
73	0.6	0.6	0.4	0.3	0.1	0.1	109	0.3	0.6	0.6	0.3	0.1	0.1
74	0.6	0.6	0.4	0.3	0.1	0.5	110	0.3	0.6	0.6	0.3	0.1	0.5
75	0.6	0.6	0.4	0.3	0.1	1.0	111	0.3	0.6	0.6	0.3	0.1	1.0
76	0.6	0.6	0.4	0.3	0.4	0.1	112	0.3	0.6	0.6	0.3	0.4	0.1
77	0.6	0.6	0.4	0.3	0.4	0.5	113	0.3	0.6	0.6	0.3	0.4	0.5
78	0.6	0.6	0.4	0.3	0.4	1.0	114	0.3	0.6	0.6	0.3	0.4	1.0
79	0.6	0.6	0.4	0.6	0.1	0.1	115	0.3	0.6	0.6	0.6	0.1	0.1
80	0.6	0.6	0.4	0.6	0.1	0.5	116	0.3	0.6	0.6	0.6	0.1	0.5
81	0.6	0.6	0.4	0.6	0.1	1.0	117	0.3	0.6	0.6	0.6	0.1	1.0
82	0.6	0.6	0.4	0.6	0.4	0.1	118	0.3	0.6	0.6	0.6	0.4	0.1
83	0.6	0.6	0.4	0.6	0.4	0.5	119	0.3	0.6	0.6	0.6	0.4	0.5
84	0.6	0.6	0.4	0.6	0.4	1.0	120	0.3	0.6	0.6	0.6	0.4	1.0
85	0.6	0.6	0.8	0.3	0.1	0.1	121	0.6	0.3	0.6	0.3	0.1	0.1
86	0.6	0.6	0.8	0.3	0.1	0.5	122	0.6	0.3	0.6	0.3	0.1	0.5
87	0.6	0.6	0.8	0.3	0.1	1.0	123	0.6	0.3	0.6	0.3	0.1	1.0
88	0.6	0.6	0.8	0.3	0.4	0.1	124	0.6	0.3	0.6	0.3	0.4	0.1
89	0.6	0.6	0.8	0.3	0.4	0.5	125	0.6	0.3	0.6	0.3	0.4	0.5
90	0.6	0.6	0.8	0.3	0.4	1.0	126	0.6	0.3	0.6	0.3	0.4	1.0
91	0.6	0.6	0.8	0.6	0.1	0.1	127	0.6	0.3	0.6	0.6	0.1	0.1
92	0.6	0.6	0.8	0.6	0.1	0.5	128	0.6	0.3	0.6	0.6	0.1	0.5
93	0.6	0.6	0.8	0.6	0.1	1.0	129	0.6	0.3	0.6	0.6	0.1	1.0
94	0.6	0.6	0.8	0.6	0.4	0.1	130	0.6	0.3	0.6	0.6	0.4	0.1
95	0.6	0.6	0.8	0.6	0.4	0.5	131	0.6	0.3	0.6	0.6	0.4	0.5
96	0.6	0.6	0.8	0.6	0.4	1.0	132	0.6	0.3	0.6	0.6	0.4	1.0
97	0.3	0.3	0.6	0.3	0.1	0.1	133	0.6	0.6	0.6	0.3	0.1	0.1
98	0.3	0.3	0.6	0.3	0.1	0.5	134	0.6	0.6	0.6	0.3	0.1	0.5
99	0.3	0.3	0.6	0.3	0.1	1.0	135	0.6	0.6	0.6	0.3	0.1	1.0
100	0.3	0.3	0.6	0.3	0.4	0.1	136	0.6	0.6	0.6	0.3	0.4	0.1
101	0.3	0.3	0.6	0.3	0.4	0.5	137	0.6	0.6	0.6	0.3	0.4	0.5
102	0.3	0.3	0.6	0.3	0.4	1.0	138	0.6	0.6	0.6	0.3	0.4	1.0
103	0.3	0.3	0.6	0.6	0.1	0.1	139	0.6	0.6	0.6	0.6	0.1	0.1
104	0.3	0.3	0.6	0.6	0.1	0.5	140	0.6	0.6	0.6	0.6	0.1	0.5
105	0.3	0.3	0.6	0.6	0.1	1.0	141	0.6	0.6	0.6	0.6	0.1	1.0
106	0.3	0.3	0.6	0.6	0.4	0.1	142	0.6	0.6	0.6	0.6	0.4	0.1
107	0.3	0.3	0.6	0.6	0.4	0.5	143	0.6	0.6	0.6	0.6	0.4	0.5
108	0.3	0.3	0.6	0.6	0.4	1.0	144	0.6	0.6	0.6	0.6	0.4	1.0

Case	κ_v	α	θ	p_e	$\ d\ $	c_v	Case	κ_v	α	θ	p_e	$\ d\ $	c_v
145	0.3	0.3	0.4	0.3	0.8	0.1	181	0.6	0.3	0.4	0.3	0.8	0.1
146	0.3	0.3	0.4	0.3	0.8	0.5	182	0.6	0.3	0.4	0.3	0.8	0.5
147	0.3	0.3	0.4	0.3	0.8	1.0	183	0.6	0.3	0.4	0.3	0.8	1.0
148	0.3	0.3	0.4	0.6	0.8	0.1	184	0.6	0.3	0.4	0.6	0.8	0.1
149	0.3	0.3	0.4	0.6	0.8	0.5	185	0.6	0.3	0.4	0.6	0.8	0.5
150	0.3	0.3	0.4	0.6	0.8	1.0	186	0.6	0.3	0.4	0.6	0.8	1.0
151	0.3	0.3	0.6	0.3	0.8	0.1	187	0.6	0.3	0.6	0.3	0.8	0.1
152	0.3	0.3	0.6	0.3	0.8	0.5	188	0.6	0.3	0.6	0.3	0.8	0.5
153	0.3	0.3	0.6	0.3	0.8	1.0	189	0.6	0.3	0.6	0.3	0.8	1.0
154	0.3	0.3	0.6	0.6	0.8	0.1	190	0.6	0.3	0.6	0.6	0.8	0.1
155	0.3	0.3	0.6	0.6	0.8	0.5	191	0.6	0.3	0.6	0.6	0.8	0.5
156	0.3	0.3	0.6	0.6	0.8	1.0	192	0.6	0.3	0.6	0.6	0.8	1.0
157	0.3	0.3	0.8	0.3	0.8	0.1	193	0.6	0.3	0.8	0.3	0.8	0.1
158	0.3	0.3	0.8	0.3	0.8	0.5	194	0.6	0.3	0.8	0.3	0.8	0.5
159	0.3	0.3	0.8	0.3	0.8	1.0	195	0.6	0.3	0.8	0.3	0.8	1.0
160	0.3	0.3	0.8	0.6	0.8	0.1	196	0.6	0.3	0.8	0.6	0.8	0.1
161	0.3	0.3	0.8	0.6	0.8	0.5	197	0.6	0.3	0.8	0.6	0.8	0.5
162	0.3	0.3	0.8	0.6	0.8	1.0	198	0.6	0.3	0.8	0.6	0.8	1.0
163	0.3	0.6	0.4	0.3	0.8	0.1	199	0.6	0.6	0.4	0.3	0.8	0.1
164	0.3	0.6	0.4	0.3	0.8	0.5	200	0.6	0.6	0.4	0.3	0.8	0.5
165	0.3	0.6	0.4	0.3	0.8	1.0	201	0.6	0.6	0.4	0.3	0.8	1.0
166	0.3	0.6	0.4	0.6	0.8	0.1	202	0.6	0.6	0.4	0.6	0.8	0.1
167	0.3	0.6	0.4	0.6	0.8	0.5	203	0.6	0.6	0.4	0.6	0.8	0.5
168	0.3	0.6	0.4	0.6	0.8	1.0	204	0.6	0.6	0.4	0.6	0.8	1.0
169	0.3	0.6	0.6	0.3	0.8	0.1	205	0.6	0.6	0.6	0.3	0.8	0.1
170	0.3	0.6	0.6	0.3	0.8	0.5	206	0.6	0.6	0.6	0.3	0.8	0.5
171	0.3	0.6	0.6	0.3	0.8	1.0	207	0.6	0.6	0.6	0.3	0.8	1.0
172	0.3	0.6	0.6	0.6	0.8	0.1	208	0.6	0.6	0.6	0.6	0.8	0.1
173	0.3	0.6	0.6	0.6	0.8	0.5	209	0.6	0.6	0.6	0.6	0.8	0.5
174	0.3	0.6	0.6	0.6	0.8	1.0	210	0.6	0.6	0.6	0.6	0.8	1.0
175	0.3	0.6	0.8	0.3	0.8	0.1	211	0.6	0.6	0.8	0.3	0.8	0.1
176	0.3	0.6	0.8	0.3	0.8	0.5	212	0.6	0.6	0.8	0.3	0.8	0.5
177	0.3	0.6	0.8	0.3	0.8	1.0	213	0.6	0.6	0.8	0.3	0.8	1.0
178	0.3	0.6	0.8	0.6	0.8	0.1	214	0.6	0.6	0.8	0.6	0.8	0.1
179	0.3	0.6	0.8	0.6	0.8	0.5	215	0.6	0.6	0.8	0.6	0.8	0.5
180	0.3	0.6	0.8	0.6	0.8	1.0	216	0.6	0.6	0.8	0.6	0.8	1.0

Bibliography

- [1] <http://www.popularpower.com/applications/influenza-vaccination.html>.
- [2] <http://www.who.int/inf-fs/en/fact211.html>.
- [3] <http://www.niaid.nih.gov/newsroom/focuson/flu00/flupic.htm>.
- [4] http://www.cdc.gov/nchs/about/major/nhis/released200207/figures04_1-4_3.htm.
- [5] J.D. Babish A.M. McBean and J.L. Warren. The impact and cost of influenza in the elderly. *Arch Intern Med*, 153:2105–2111, 1993.
- [6] W. Beyer, A. Iris, A.M. Palache, R. Westendorp, and A. Osterhaus. Protection against influenza after annually repeated vaccination. *Arch Intern Med*, 159:182–188, 1999.
- [7] R.M. Bush, C.A. Bender, K. Subbarao, N.J. Cox, and W.M. Fitch. Predicting the evolution of human influenza A. *Science*, 289:1921–1925, December 1999.
- [8] R. DeBoer, P. Hodgeweg, and A.S. Perelson. *Theoretical and Experimental Insights into Immunology*, pages 223–247. Springer-Verlag, 1992.
- [9] R. DeBoer, L. Segel, , and A. Perelson. Pattern formation in one and two dimensional shape space models of the immune system. *J Theor Biol*, 155:295–333, 1992.
- [10] W.M. Fitch, R.M. Bush, C.A. Bender, and N.J. Cox. Long term trends in the evolution H(3) HA1 human influenza type A. *PNAS*, 94:7712–7718, July 1997.
- [11] P.A. Gross, D. Herrington, H.S. Sachs, J. Laufer, and R.A. Lavandowski. The efficacy of influenza vaccine in elderly persons. *Ann Intern Med*, 123:518–527, 1995.

- [12] P.A. Gross, S.J. Sperber, A. Donabedian, S. Dran, G. Morchel, P. Cataruozolo, and G. Munk. Paradoxical response to a novel influenza virus vaccine strain: the effect of prior immunization. *Vaccine*, 17:2284–2289, 1999.
- [13] T.W. Hoskins, J.R. Davis, A.J. Smith, C.L. Miller, and A. Allchin. Assessment of inactivated influenza-a vaccine after three outbreaks of influenza a at christ’s hospital. *Lancet*, 1:33–35, 1979.
- [14] W.A. Keitel, T.R. Cate, R.B. Couch, L.L. Huggins, and K.R. Hess. Efficacy of repeated annual immunization with inactivated influenza virus vaccines over a five year period. *Vaccine*, 15(10):1114–1122, 1997.
- [15] A. Lapedes and R. Farber. The geometry of shape space: Application to influenza. *J Theor Biol*, 212:57–69, 2001.
- [16] K.G. Nicholson, R.G. Webster, and A.J. Hay. *Textbook of Influenza*, chapter 25. Blackwell Science, 1998.
- [17] A. Perelson and G. Oster. Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination. *J Theor Biol*, 81:645–667, 1979.
- [18] A.S. Perelson. Immune network theory. *Immun Rev*, 110:5–35, 1989.
- [19] J.B. Plotkin, J. Dushoff, and S. Levin. Hemagglutinin sequence clusters and the antigenic evolution of influenza A virus. *PNAS*, 99(9):6263–6268, April 2002.
- [20] L.A. Segel and A.S. Perelson. *Theoretical Immunology, Part Two*, pages 321–343. Addison-Wesley, RedWood City, CA, 1988.
- [21] D.J. Smith, S. Forrest, D.H. Ackley, and A.S. Perelson. Variable efficacy of repeated annual influenza vaccination. *PNAS*, 96(24):14001–14006, 1999.
- [22] D.J. Smith, S. Forrest, R.R. HighTower, and A.S. Perelson. Deriving shape space parameters from immunological data. *J Theor Biol*, 189:141–150, 1997.