Optimal immune systems: A resource allocation and information processing view of immune defense

Andreas Mayer

THÈSE DE DOCTORAT
de l'Université de recherche
Paris Sciences Lettres –
PSL Research University

préparée à
l'École normale supérieure

Optimal immune systems

A resource allocation and information processing view of immune defense

par Andreas Mayer

École doctorale Physique en Ile-de-France
Spécialité: Physique
Soutenue le 23 juin 2017

Composition du Jury :

M Edo Kussell
New York University
Rapporteur

M Olivier Martin
INRA
Rapporteur

Mme Aleksandra Walczak
École normale supérieure
Directrice de thèse

M Thierry Mora
École normale supérieure
Directeur de thèse

Mme Martine Ben Amar
École normale supérieure
Présidente du Jury

M Pierre Boudinot
INRA
Membre du Jury

M Olivier Tenaillon
Paris Diderot - INSERM
Membre du Jury

ENS ÉCOLE NORMALE SUPÉRIEURE | PSL RESEARCH UNIVERSITY PARIS

Seul l'inconnu épouvante les hommes. Mais, pour quiconque l'affronte, il n'est déjà plus l'inconnu.

Terre des hommes — Antoine de Saint-Exupéry

ABSTRACT

Biological organisms have evolved diverse immune mechanisms to defend themselves against pathogens. Here we build mathematical models of immune systems optimally tuned to the statistics of pathogens.

Beyond molecular details, different immune mechanisms differ in how protection is acquired, processed and passed on to subsequent generations – differences that may be essential to long-term survival. To explain the observed diversity of strategies we compare the long-term adaptation of populations as a function of the pathogen dynamics that they experience and of the immune strategy that they adopt. We find that the two key determinants of an optimal immune strategy are the frequency and the characteristic timescale of the pathogens. Depending on these two parameters, we identify distinct modes of immunity, including adaptive, innate, bet-hedging and CRISPR-like immunities, which recapitulate the diversity of natural immune systems. Our results carry over to the general question of evolution in fluctuating environments, for which we provide novel analytical results in temporally correlated environments.

The adaptive immune system provides protection through a broad repertoire of cells specific to different pathogens. To predict statistical features of well-adapted repertoires we analyze which repertoire minimizes cost of infection for a given distribution of pathogens. The theory predicts that the immune system has more receptors for rare antigens than expected from the frequency of encounters; and individuals exposed to the same infections have sparse repertoires that are largely different, but nevertheless exploit cross-reactivity to provide the same coverage of antigens. Our results follow from a tension between the statistics of pathogen detection, which favor a broader receptor distribution, and the effects of cross-reactivity, which tend to concentrate the optimal repertoire onto a few highly abundant clones. These predictions can be tested in high throughput surveys of receptor and pathogen diversity. We then explicitly consider how the adaptive immune system can learn the statistics of the environments from its past infection history in a Bayesian manner. We show that optimal repertoires can be reached by keeping memory of an infection through the selective proliferation of stimulated cells. The Bayesian perspective on repertoire dynamics provides an unifying conceptual framework to explain a number of features of immunological memory and suggests further experiments.

RESUMÉ

Les organismes biologiques ont développé divers mécanismes immunitaires afin de se protéger des pathogènes. Nous développons ici des modèles mathématiques de systèmes immunitaires, adaptés de façon optimale aux statistiques des pathogènes.

Au delà des détails moléculaires, ces mécanismes immunitaires diffèrent dans la manière d'acquérir, de réguler et de transmettre une protection immunitaire; différences qui pourraient s'avérer essentielles pour la survie à long terme. Afin d'expliquer la diversité des stratégies qui sont observées, nous comparons l'adaptation à long terme de populations en fonction de la dynamique des pathogènes à laquelle elles sont confrontées et de la stratégie immunitaire qu'elles adoptent. Nous démontrons que la fréquence et l'échelle de temps caractéristique des pathogènes sont les deux déterminants clés d'une stratégie immunitaire optimale. En fonction de ces deux paramètres, nous identifions des modes d'immunité distincts, comprenant immunités innées, adaptatives, ou ressemblant au système CRISPR, qui récapitulent la diversité de systèmes immunitaires naturels. Nos résultats viennent s'étendre à la question générale de l'évolution dans des environnements variables pour laquelle nous apportons de nouveaux résultats analytiques au sein d'environnements temporairement corrélés.

Le système immunitaire adaptatif assure une protection à partir d'un large répertoire de cellules spécifiques à différents pathogènes. Pour prédire des propriétés statistiques de répertoires adaptés, nous étudions quel répertoire minimise au mieux le risque d'infections pour une distribution de pathogènes donnée. La théorie prédit que les cellules spécifiques contre les antigènes rares sont surreprésentées par rapport à la fréquence de leurs rencontres et que les individus, exposés aux mêmes infections, possèdent des répertoires avec des récepteurs largement différents mais exploitent la réactivité croisée afin de parvenir à la même couverture d'antigènes. Nos résultats sont issus d'une opposition entre les statistiques de détection des pathogènes, qui soutiennent l'idée d'une plus large distribution de récepteurs, et les effets de la réactivité croisée, qui tend à concentrer le répertoire optimal sur un petit nombre de clones. Nos prédictions peuvent être testées à partir des données à haut débit sur la diversité des récepteurs et de pathogènes. Par la suite, nous examinons explicitement comment le système immunitaire adaptatif peut apprendre de manière bayésienne les statistiques de l'environnement à partir de l'historique des infections précédentes. Nous montrons que les répertoires optimaux peuvent être atteints par prolifération sélective des cellules spécifiques. La perspective bayésienne sur la dynamique des répertoires fournit un cadre conceptuel unificateur qui explique un certain nombre de caractéristiques de la mémoire immunitaire et appelle à des expériences complémentaires.

PUBLICATIONS

This thesis is based on a number of published and submitted publications.

A preliminary version of the results presented in Chapter 4 was submitted as a Master thesis to the University of Göttingen. Chapter 4 was previously published in [1], chapter 6 in [2]. Chapter 7 is currently under review and is available as a preprint [3]. The optimization algorithm used within these chapters has beeen released as a stand-alone library described in [4]. Chapter 5 is not yet submitted for publication. For the ease of reading some of the supplementary material from the publications has been incorporated in the introductory chapters.

During my thesis I had the pleasure to collaborate with another PhD student, Jonathan Desponds, on a joint project about population dynamics models of immune repertoires. Some of this work was published in [5]. This work is not directly discussed in this thesis, but some of its ideas are touched upon in the introductory and closing chapters.

[1] Andreas Mayer, Vijay Balasubramanian, Thierry Mora, and Aleksandra M Walczak. "How a well-adapted immune system is organized." In: *Proceedings of the National Academy of Sciences of the United States of America* 112.19 (2015), pp. 5950–5955.

[2] Andreas Mayer, Thierry Mora, Olivier Rivoire, and Aleksandra M Walczak. "Diversity of immune strategies explained by adaptation to pathogen statistics." In: *Proceedings of the National Academy of Sciences* 113.31 (2016), pp. 8630–8635.

[3] Andreas Mayer, Thierry Mora, Olivier Rivoire, and Aleksandra M. Walczak. "Transitions in optimal adaptive strategies for populations in fluctuating environments." In: *arXiv preprint arXiv:1703.09780* (2017), pp. 1–18.

[4] Andreas Mayer. "Noisyopt: A Python library for optimizing noisy functions." In: *Journal of Open Source Software* 2.13 (2017).

[5] Jonathan Desponds, Andreas Mayer, Thierry Mora, and Aleksandra M. Walczak. "Population dynamics of immune repertoires." In: *arXiv preprint arXiv:1703.00226* (2017).

ACKNOWLEDGMENTS

I want to first and foremost thank Aleksandra Walczak and Thierry Mora, which have supervised me for the last four years first for my Master thesis and then during my PhD. It is your excellent ideas and encouragement that have enabled me to do the work presented in this thesis. Thank you for introducing me to the fascinating world of theoretical biophysics and teaching me both intuition and rigour. Thank you for sometimes subtly pushing me back towards thinking about the larger questions when my fondness of numerics, or perfectionism had been taking over for too long. Thank you also for the liberty to follow my own curiosity and to work at my own (sometimes irregular) rhythm. I am also indebted to you for providing me so many opportunities to learn broadly and to encounter the community of biophysicists by attending conferences and schools, and for having assembled such a stimulating group of people in Paris.

Thank you for deeply caring about the members of your group through the good times and the more difficult times. Whenever I had a question or whenever there was a deadline looming no matter the circumstances I had a reply from you within hours. Whenever we talked about science you needed uncannily little prompting to know precisely the ongoing work and to have deep discussions despite your large panel of projects. Whenever we talked about life no lunch break in the Collège des Irlandais was ever to long to get to the bottom of the newest discussion topic or gossip. Finally, even in the moments when we disagreed on something (or thought to do so) we had good natured discussions (I admit my part of responsibility in the length of these discussions). I will keep fond memories of these years in Paris learning from you and working with you and hope that our paths will cross again many times in the future.

I thank Edo Kussell and Olivier Martin for having agreed to serve as referees for this thesis and for their close reading of the manuscript. A big thank you also to Martine Ben Amar, Pierre Boudinot, and Olivier Tenaillon for being part of my defense committee. Furthermore I thank Rémi Monasson for having followed my progress towards this PhD as an advisory committee member.

I then want to thank Vijay Balasubramanian and Olivier Rivoire with which I had the pleasure to collaborate. I still remember the pre-Christmas week three and a half years ago in which I first got to know Vijay when he visited us in Paris. Leading up to his visit I had obtained some first results on optimal repertoires, propelled by the enthusiasm of Vijay and daily, hour-long discussions between all four of us we made enormous progress in interpreting what we had learned so far and spinning further ideas. I count myself lucky to have experienced such excitement and intense scientific discussions early into my work, which has helped me getting through tougher moments later on. Olivier was an invaluable collaborator for the work on adaptive strategies in fluctuating environments. With his deep knowledge of models of informa-

tion processing in evolutionary dynamics he seemingly always had already thought about answers to most of the questions and problems I encountered along the way.

A special thank you goes to my lab colleagues, in order of appearance – Jonathan Desponds, Rhys Adams, Yuval Elhanati, Quentin Marcou, Christophe Gardella, Paulina Szymanska, Maximilian Puelma Touzel, Huy Tran, Jacopo Marchi, Alec Douglas, and Nicholas Rossi. Your enthusiasm for discussing wide-ranging topics in science and beyond have shaped and helped develop my own thinking as a scientist. The last years would have been much less enjoyable without your company and our lunches, coffee breaks, and parties. Merci Jonathan, pour avoir partagé ton humour et ta culture avec moi dans les endroits variés où nos thèses nous ont emmenés en commençant par notre tournée dans l'ouest américain, des pistes de ski des Houches, jusqu'aux bars du Marais. Thank you Yuval, for the countless discussions for example fueled by your encyclopaedic knowledge of Parisian geography and museums, and for our joint running excursions on the coulée verte du sud-parisien and elsewhere. Merci Quentin, pour nos débats de science, politique, musique, et philosophie; pour les fêtes dans la «commune de Villejuif »; et pour toutes ces moments de joie et de souffrance partagés le long de nos thèses.

Pour leur aide organisationnel, leur efficacité, et leur bonne humour en m'aidant naviguer la bureaucratie parfois opaque je tiens à remercier Viviane Sebille et Sandrine Patacchini.

Während meiner Schulzeit wurde mein Interesse an Naturwissenschaften und Mathematik durch einige besondere Lehrer gestärkt. Zwei möchte ich besonders hervorheben: Edith Pfinder, die mich wahlweise als Mathe- und Physiklehrerin fast die ganze Schulzeit begleitet hat; und Bernhard Horlacher, Leiter des Kepler-Seminars für Naturwissenschaften in Stuttgart, an dessen Ideenreichtum und Weltsicht ich oft noch gerne zurück denke. Während meiner Studienzeit hatte ich das Glück die Grundvorlesungen bei sehr engagierten Professoren wie Thomas Pruschke und Kurt Schönhammer zu belegen, denen ich hier stellvertretend für viele andere danken möchte. Mein Dank gilt weiterhin Ulrich Parlitz, der mir die wunderbare Welt der nichtlinearen Dynamik eröffnet hat, und durch seine Co-Betreuung die externe Masterarbeit in Frankreich ermöglicht hat, die schließlich zu dieser Dissertation geführt hat.

Finally, I want to express my deep gratitude towards my family and close friends who even if sometimes from afar have always been with me. Danke, merci, and thank you to all of you!

CONTENTS

1 INTRODUCTION

The availability of increasingly large data sets gathered in a (semi-)automated fashion holds the promise to change many fields. In biology, methods for such facilitated data generation, known as high-throughput methods, have become increasingly powerful and affordable over the last two decades. A prime example is sequencing, whose cost in monetary terms and in terms of effort has dropped massively: the sequencing of a human genome, which just 20 years ago took a large internationally coordinated effort to establish, has now become a routine analysis [1].

These massive data sets coupled to computational algorithms allow one to extract patterns and correlations based on which one can make predictions. This has prompted some to declare the end of theory [2]. There certainly is a need for purely data-driven work in biology but on its own this approach can suffer from issues of interpretability and generalizability. Therefore such work should rather be seen as complementary to hypothesis-driven research and theoretical work.

What the availability of abundant quantitative data calls for in my view is the advent of more quantitative and mathematical theory. The importance of mathematical modelling in this data-rich era stems from its unique ability to connect simple hypotheses to the kind of statistical signatures that can be checked against the large-scale, quantitative data characteristic of modern biological experiments. What has been especially called for in biology is not only modelling aimed at fitting already obtained experimental data but theory that turns conceptual ideas into testable predictions informing future experiments [3]. The paradigm of hypothesis-driven theory development takes inspiration from how theoretical physics has proceeded in its tremendously successful description of non-animate matter [4]. While biological systems operate within the confines of the same physical laws, "more is different" [5] and the kind of theory we are discussing here should aim to provide relevant mescoscopic descriptions of the emergent complexity of biological phenomena.

The work presented within this thesis applies this paradigm to the study of how organisms protect themselves against pathogens. Immunology is a field of research of obvious practical importance to public health and moreover a field where many exciting developments have taken place in the last decade from discoveries of previously unknown immune mechanisms [6, 7, 8] to the opportunities created by new experimental techniques such as high-throughput sequencing [9, 10, 11]. The new questions posed by these discoveries and the older ones now within better experimental reach, have kept me fascinated during the last four years that I have been thinking about immunology.

This thesis aims to contribute to answering two such questions in particular. First, what explains the diversity and convergent evolution of adaptive strategies used by different organisms in their immune defense? Second, what

determines the statistical properties of the repertoire of receptors of the cells of the adaptive immune system? The former question is posed by discoveries of hitherto unknown immune mechanisms such as of an adaptive immune system of non-jawed vertebrates [6] or the CRISPR-Cas system of bacteria [7] that have substantially expanded our view of how pathogen defense is organized. Renewed interest into the latter question has been sparked by new data coming from high-throughput sequencing applied to cells of the adaptive immune system [9, 10, 11].

The central hypothesis around which my work has been structured is that immune systems are well-adapted to the statistical structure of pathogenic environments. As we will see this hypothesis leads to a number of interesting and potentially experimentally testable predictions when applied to concrete questions such as "how should immune repertoires be organized to minimize harm from infection?" or "how should a population of organisms adapt its immune defense over generations?". To answer these questions we build simple mathematical models of immunity and ask what would be the optimal organization or dynamics of immune protection in a given pathogenic environment.

A basic property of immune defense is that it aims to provide protection against stochastic events, – infections. The importance of stochasticity is thus a common thread throughout the different models. This thesis thus heavily relies on concepts and tools from the theory of stochastic processes and statistical physics. A large part of my work has been devoted to model development, and to finding adequate mathematical descriptions that capture the relevant biology. To obtain insights from the models we then determined optimal strategies by either analytical or numerical optimization of parameters within biological constraints.

Eloquent critiques of the use of optimality principles in biology have been made, among which Gould and Lewontin [12] is a classic example. While we investigate how immune systems might operate optimally for a specific task, we certainly do not claim that immune systems are precisely optimized. We rather use the assumption of optimality to have a mathematically well-posed way of asking the question of what general consequences arise from the need to perform a certain function well [13]. The use of optimality theory in biology can be compared to the use of the equilibrium assumption in statistical mechanics [14] or to the use of maximizing principles in classical economic theory [15]. The dynamics of a many-body system – the evolutionary dynamics of species adaptation, the non-equilibrium processes preceding relaxation to equilibrium, or the dynamical interplay of economic players, respectively – is replaced by a static picture by the appeal to an optimizing principle. One should not forget and potentially question the strong assumptions leading there but such an idealization can allow for progress in our understanding. Numerous examples of the fruitfulness of such thinking in other biological contexts from neuroscience [16, 17, 18] to cellular biology [19, 20] hint at the possibility that something interesting is to be learned in this way about immune systems too.

Marvelling at the complexity of biological systems working on much simplified models sometimes feels ultimately wrong-headed, or even futile. Never-

thesless, to gain understanding with our limited intellects needs simplification. In this sense "models in biology [are] accurate descriptions of our pathetic thinking" [21]. This is especially true of the kind conceptual theory presented in this thesis. In the words of the biophysicist Rob Phillips such "theory [...] is about living dangerously by turning our thinking into formal mathematical predictions and confronting that math with experiments that have not yet been done." The future will hopefully see such a confrontation of the predictions arising from thinking about optimal immune systems with experimental data. It is then that we will know whether ultimately anything useful will have been learned. What has motivated me during this thesis is the hope that the answer will be positive. As the reader will hopefully see in the following the approach we have taken at the very least leads to intellectually stimulating and potentially general theoretical questions.

The rest of this dissertation is structured into four parts the first three of which are further subdivided into two chapters each. The chapters in parts ii and iii are self-contained and can be read independently.

Part i provides an introduction to some basic concepts in immunology, evolution, and mathematics relevant to the work presented in this thesis. In Chapter 2, I give a brief introduction to immunology leading up to a discussion of some open questions concerning the adaptive immune repertoires and the evolution of immune systems. In Chapter 3, I introduce some of the mathematical theory used in this work from optimization, via stochastic processes, to some elements of evolutionary theory.

Part ii presents results about the structure and dynamics of immune repertoires in adaptive immunity. In Chapter 4 we derive optimal distributions of immune receptors in a fixed, known antigenic environment and discuss how statistical features of those optimal repertoires might help explain some puzzling experimental results. In Chapter 5 we relax the assumption that the environment is fixed and known. This leads to a Bayesian view of a well-adapting repertoire dynamics, in which the adaptive immune system combines prior beliefs about pathogenic variation with the past pathogen exposure to estimate the current pathogen statistics. This provides a novel perspective on immunological memory that makes a number of experimentally testable predictions.

Part iii presents results about the evolution of diverse immune strategies understood as a question of optimizing population growth in fluctuating environments. In Chapter 6 we show how the diversity of immune strategies across the tree of life might be explained as an adaptation to differing pathogen statistics. In Chapter 7 we present a number of analytical and numerical results on transitions between strategies maximizing long-term growth rate in fluctuating environments. We use some of these results to obtain a deeper understanding of the phase diagram of immune strategies obtained in the previous chapter.

Chapter 8 concludes this thesis by summarizing the major contributions of this thesis and sketching some ideas for future research.

Part I

BACKGROUND, CONCEPTS AND TOOLS

In this part of the thesis I introduce some of the immunological background, theoretical concepts, and mathematical tools that are used throughout the remainder of the thesis. Without trying to be exhaustive I will give a brief exposition of some core ideas.

Side notes are used to explicitely link concepts and tools to where they are used in the rest of the thesis.

2 PRINCIPLES OF IMMUNE DEFENSE

2.1 INTRODUCTION

Biological organisms facing the threat of infections by pathogens have evolved a multi-layered system of protective mechanisms [22, 23]. The ensemble of these immune defense mechanisms serves to protect the organism from disease. Defense mechanisms in humans range from simple mechanical barriers to pathogen entry such as provided by our skin, to the highly specialized cells of the adaptive immune system that patrol our body to detect and combat foreign intruders. The importance of an intact immune system is underlined by the problems arising from its failures in the elderly or in patients with immunodeficiency [22]. For instance, patients afflicted with the human immunodeficiency virus (HIV), which targets cells of the immune system, often suffer from life-threatening opportunistic infections.

As infectious diseases are one of the leading causes of human mortality much effort has gone into studying the immune system. The development of antibiotics, vaccinations, and other therapeutics has dramatically reduced the death toll of infectious diseases. A particular striking success story of immunological research is eradication of pathogens all-together through coordinated world-wide vaccination programs, such as achieved for small pox in the 1970s [22].

Recently cancer immunotherapy has received a lot of attention [24, 25]. This emerging field aims at engineering the immune system to combat cancer. The somatic mutations leading to the growth of cancer cells can also lead cancerous cells presenting different surface molecules. These neoantigens are potential targets for recognition by T cells of the adaptive immune system. To enhance immune responses drugs can be used to block a checkpoint inhibiting prolonged activation of the cells. Additionally the number of T cells specific to the neoantigens can be enhanced by administering a synthetic vaccine or by expanding them ex-vivo.

The clinical advancements in immunology went hand in hand with and were enabled by fundamental research into the mechanisms of immune protection. This research has furthered our understanding of the constituents of the different defense systems and of how they are interconnected. On a conceptual level there are a few recurring themes. An immune defense mechanism needs to distinguish what is harmful from what is not. Especially an immune system should not react to an organisms' "self", i.e. its own healthy cells, proteins, etc. Failures of such self/non-self recognition lead to autoimmune diseases. Furthermore immune defense mechanisms need to deal with the stochasticity of infections and the (co-)evolutionary dynamics of pathogens.

The remainder of this chapter focuses on a selection of aspects of immunology that are particularly relevant to the work presented in the remainder of this thesis.

2.2 THE ADAPTIVE IMMUNE SYSTEM

used in particular in Ch. 4,5

2.2.1 *Life and death of lymphocytes*

The adaptive immune system recognizes pathogens through specific biophysical binding. Lymphocytes are the cells that make up the adaptive immune system. There a two arms of adaptive immunity, relying on T cells and B cells respectively. B cells are important for combating pathogens in extracellular fluids by the production of antibodies (humoral immunity), whereas T cells co-ordinate the immune response of other immune cell types and kill infected cells (cellular immunity). Lymphocytes bind to pathogen-derived peptides – called antigens – with a particular receptor protein – called T cell receptor (TCR) or B cell receptor (BCR). Every lymphocyte has a unique receptor, but there is large diversity of receptors across cells. The diversity allows for protection against diverse and mutating pathogens. The multiset of all the receptors is known as the immune repertoire. Its composition in an individual determines the breadth and speed of the response of the adaptive immune system to pathogens [26]. The immune repertoire composition in terms of the sequences of the receptors is now directly experimentally accessible by sequencing (Sec. 2.2.3). According to clonal selection theory lymphocytes specific to the pathogen proliferate and differentiate into effector cells during an infection [27]. After an infection most of the effector cells die, but some remain in a memory state to provide long-lasting protection against reinfection by the same pathogen.

Both B cells and T cells can be subdivided into subsets with different characteristics and functions as assessed experimentally by the determination of molecular markers expressed on the surface of the cells. An example is the distinction between T cells expressing the CD4 glycoprotein on their surface from those expressing the CD8 glycoprotein. The glycoproteins are co-receptors involved in the interaction of T cells with the cells presenting antigen. The former are known as T helper cells as they have the role of co-ordinating the immune response by the secretion of signaling molecules, whereas the latter are known as Cytotoxic T cells as they kill cells infected by viruses or bacteria [22].

A diverse set of receptors is generated by a genetic recombination mechanism. The recombination machinery creates a genetic sequence coding for the variable region of the receptor protein by a fuzzy joining of genomic templates. This process creates an enormous diversity of possible sequences by the combinatorial diversity of template choice as well as through the random insertions and deletions in the joining of templates. Not all recombinations are equally likely [28] and the same sequence can be generated by different recombination scenarios [29]. Different receptors are thus not generated with the same frequency resulting in a biased repertoire. These biases might explain

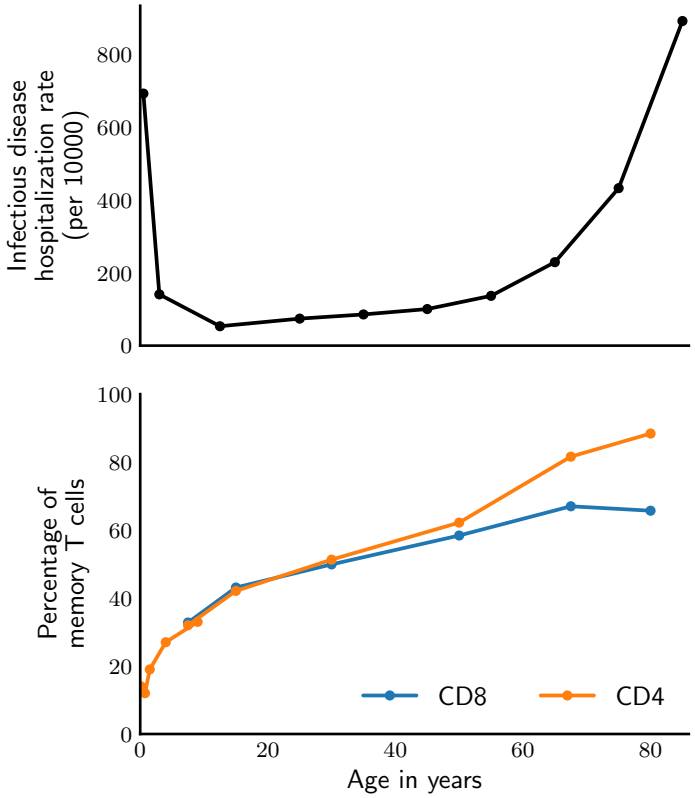

Figure 2.1: Infectious disease hospitalization rates and median percentage of T cells showing memory markers as a function of age. Own work following idea from [35]. Data on infectious disease hospitalization rate from [36]. T cell data from [37] and [38] (early CD4 data).

why, despite the enormous number of potential receptors, sometimes different individuals respond with the same receptor to an antigen [30].

Given the diversity of receptors the immune system has to make sure to eliminate the receptors that recognize benign peptides belonging to its own body. Negative selection of newly generated lymphocytes that interact too strongly with self-molecules is one key step to avoid auto-immunity [22]. There are further peripheral mechanisms including regulatory T cells to suppress unwanted responses.

The composition of the lymphocyte repertoire in the productive repertoire is determined by the population dynamics of lymphocytes. As a lymphocyte keeps its rearranged receptor during proliferation all cells emanating from the same initial cell form a clone with the same specificity to antigen. The survival of lymphocytes in the periphery depends on various stimuli ranging from cytokines (signaling molecules of the immune system) to interaction with self- and foreign-antigens [31, 32]. Lymphocyte survival and homeostasis has been argued to resemble an ecological system of species competing for resources [33, 31]. Antigens have been shown experimentally to be part of these resources [34, 31]

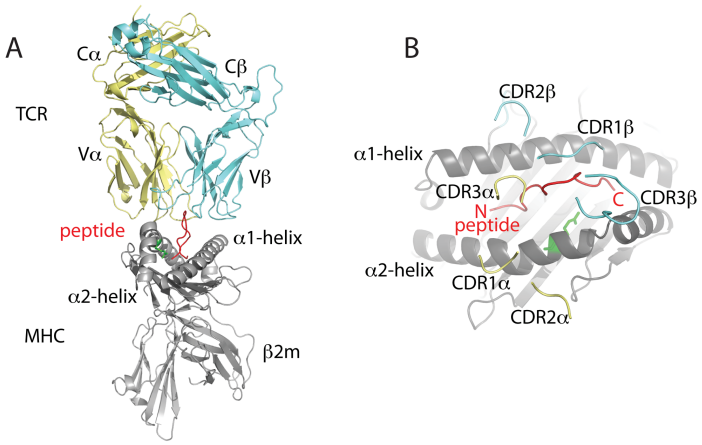

Figure 2.2: (A) Structure of a T cell receptor (TCR) and a major histocompatibility complex (MHC). The MHC binds short peptides in a groove between two α-helices. The TCR normally binds loosely to the MHC and is only activated if its interaction with the peptide provides additional stabilization. (B) Close-up view of the binding site showing the spatial localization of different regions of the TCR relative to the peptide. From: [44] under cc-by License

Immunological memory is a key feature of the adaptive immune system. Let us review some of the details of how memory is built and maintained for CD8+ T cells [39]. Upon primary infection there is a large proliferation that amplifies the number of specific lymphocytes by several orders of magnitude from which a relatively fixed 5-10% has been shown to survive after infection. The naive precursor frequency has been shown to correlate with the immune response magnitude in primary infection [40, 41]. This leads to an about 100-1000 fold increase in antigen-specific T cells [42]. The determinants of the cell fate descision are starting to be elucidated but a precise mechanistic understanding of how a fixed fraction of memory cells is produced remains lacking. Secondary memory responses differ markedly from primary responses. The build up of immunological memory contributes to large decreases in pathogen susceptibility during the first years of life (Fig. 2.1) [35, 43]. Infants are highly susceptible as they do not yet have formed sufficient protective immune memory.

2.2.2 *Molecular basis of pathogen recognition*

T cells recognize peptides presented on so-called major histocompatibility complexes (MHC) (Fig. 2.2A) [45]. The MHC molecule has a groove to which peptides can bind stably. Inside so-called antigen presenting cells MHCs are loaded with antigenic peptides and then shuttled to their surface. To trigger activation the T cell receptor (TCR) must bind to the MHC-peptide complex. The binding of TCR to MHC in general has insufficient affinity to trigger T cell activation. Further stabilization of the binding through interactions between the TCR and the peptide are thus needed. Structural studies of TCRs bound to an MHC-peptide complex have revealed some of the geometry of the binding

[46, 47]. The complementarity determining region 3 (CDR3) is the region of the TCR that is closest to the peptide (Fig. 2.2. Thanks to the genetic recombination mechanism described in the previous section it is highly variable. The matching between the diverse CDR3s and the MHC-bound peptide explains much of the specificity of T cell to diverse pathogens.

B cells do not require antigen presentation as they bind directly to surface markers of pathogens [22]. As in the case of T cells this binding involves a surface receptor, the B cell receptor (BCR), which provides specificity to the recognition. The antibodies produced by B cells are basically soluble versions of the BCR and thus share the same specificity.

For both T and B cells the recognition process is known to be degenerate [48, 30, 49, 50], i.e. receptors are cross-reactive and recognize a number of different antigens. Cross-reactivity predominantly occurs between closely related antigens, although exceptions have also been observed [51, 50]. The conformational flexibility of the TCR-MHC binding site gives a structural explanation for cross-reactivity [50]

Cross-reactivity is currently studied by a number of complementary approaches from structural biology [52] to mutation-experiments [53]. Cross-reactivity is believed to be an essential feature of T cell receptor based pathogen recognition [48, 49]. As the number of possible antigenic peptides by far exceeds the number of cells, full protection can only be achieved if every receptor binds many peptides. Cross-reactivity is thought to be limited by the need to avoid autoimmunity, i.e. receptors that bind too broadly would also bind benign molecules.

2.2.3 *Repertoire sequencing and characteristics*

used in Sec. 8.2.1

Modern high-throughput sequencing has greatly expanded the experimental ability to study immune repertoire composition [9, 54, 55, 11, 26]. With sequencing it is possible to obtain the nucleotide sequences on either the DNA or mRNA level of all T cell or B cell receptors from a biological sample. With unique barcoding – the insertion of random tags before sequencing – reliable counts of the number of times a sequence was in the sample can now additionally be obtained [56, 57].

Repertoire sequencing holds the promise to provide signatures of past pathogen exposure by the enhanced presence of pathogen-specific receptors. Large-scale sequencing efforts have for example helped identify a set of TCRs over-represented in repertoires of individuals chronically infected with cytomegalovirus [58]. It remains however often difficult to find specific signatures in a given individual due to the statistical nature of the adaptive immune response. The precise receptors that respond differ between individuals and only few are shared as "public" responses across individuals [58, 59]. To think about repertoires on a functional-level thus involves thinking about them as statisticial ensembles [26]. It is on such a statistical level that robust signatures of antigen-driven selection might be expected. Broader features than the full CDR3 sequence have been shown to allow to better distinguish betwen pre- and post-immunization repertoires bioinformatically [60, 61].

The recombination process does not always produce productive sequences. If the recombination succeeds on the second try then the first attempt is silenced and a functional cell is created. Due to the silencing there is no selection of cells based on their unproductive sequence. The statistics of unproductive sequences thus provides a window into the recombination process isolated from forces later acting on the repertoire. By learning statistical models of the VDJ-recombination process from such unproductive sequences, a quantative description of the biases in the generation process can be obtained [29]. Statistical models of selection have then been inferred on productive sequences by using the generation model to normalize [62].

An interesting feature recurrently found in sequencing studies is the existence of a very broad clone-size distribution. Such long-tailed distributions of clone-sizes might be explained by a temporally fluctuating fitness of clones [63].

2.3 DIVERSE IMMUNE STRATEGIES ACROSS THE TREE OF LIFE

used in Ch. 6

The question of how to defend against pathogens is present across all of biology. Here I review some of the mechanisms used in different organisms – from bacteria to vertebrates – for immune defense. I also discuss the evolution of the different defense mechanism and discuss how beyond their mechanistic divergences they show diverse adaptive strategies in coping with changing pathogenic environments.

Even bacteria, some of which are pathogens themselves for higher organisms, need to defend themselves from other pathogens. Specifically, they face the threat of getting infected by bacteriophages. Bacteriophages are viruses that infect bacteria. Given the common threat of phage infection bacteria have evolved a range of resistance mechanisms [64, 65]. Bacteriophages infect the bacteria by adsorption which relies on recognition of host-specific surface markers. Infection probability is thus reduced when the surface structure relevant to phage adsorption is either modified or masked. Two other modes of defense are based on cutting the phage DNA: the restriction-modification system and the CRISPR-Cas system. Restriction-modification systems cleave incoming unmethylated phage DNA using restriction enzymes [64]. To protect the host DNA from also being cut another set of proteins methylates it. The CRISPR-Cas system works by incorporating short stretches (30-70 basepairs) of phage DNA, called spacers, into the bacterial genome [7, 66, 67]. These spacers are then expressed as RNA, which serves as a guide to effector proteins cleaving the cognate virus nucleic acids. To distinguish self from non-self the incorporation of spacers from DNA is contingent on a preceding motif absent in the host's CRISPR array. The CRISPR-Cas system provides sequence-specific defense in contrast to the restriction-modification system. The incorporated spacers are heritable so that the offspring are also protected against the same phage. Expression of the acquired spacers is thought to be constitutive [7]. Outside of their role in bacterial immune defense both the restriction-modification system and the more recently discovered CRISPR-Cas system have found important applications in genetic engineering.

In multicellular organism immune defense is more complicated than in unicellular bacteria and often involves specialized cells [68]. In vertebrates immunity is divided into two branches: the innate immune system, which provides an immediate but largely unspecific response to infections, and the evolutionary more recent adaptive immune system, which mounts more targeted responses.

Adaptive immunity was thought to exist only in jawed vertebrates. More recently a similar defense system has been discovered in jawless vertebrates (e.g. lampreys). This newly discovered immune system shares many of the characteristics of the jawed-vertebrate adaptive immune systems such as the expression of somatically-diversified receptors and immunological memory [6]. On a molecular level the machinery used for defense is different but the adaptive strategy of how to deal with changing pathogens is remarkably similar, in an interesting example of convergent evolution.

Innate immunity is evolutionary older than adaptive immunity and also exists in invertebrates. It consists of various defense mechanisms such as macrophages which remove unwanted cells and bacteria by phagocytosis. They are directed against pathogen by pattern recognition receptors (PRRs) which target pathogen-associated molecular patterns. An example of these pattern recognition receptors are Toll-like receptors. They target conserved patterns of pathogens such as lipopolysaccharides, which are found in the outer membrane of Gram-negative bacteria. PRRs are less versatile and less specific then the somatically diversified receptors of the adaptive immune system and thus need to rely on more highly conserved pathogen signatures. Memory which has been thought to be a defining feature of adaptive immunity has been shown recently to also exist in some innate mechanisms [8].

Taking a large view of different immune defense mechanisms, it is interesting to observe that they differ beyond the molecular details in the evolutionary strategies they employ to adapt to changing pathogenic environments [68, 69]. The immune systems differ in how they process information about the environment and how they employ defense. The adaptive immune system is highly adaptable during the lifetime but does not pass on specificities to the next generation. This is in contrast with innate mechanisms which provide a more constitutive defense and where specificities are transmitted through the germline. The CRISPR-Cas system also transmits specific protection but in contrast to the other immune systems acquires protection actively and not by random mutation. The diversity of these adaptive strategies on the one hand and the convergent evolution of some types of strategies on the other hand pose the question of whether there are some common evolutionary determinants explaining the evolution of the adaptive strategies. We provide a proposal in terms of the adaptive value of different strategies in temporally fluctuating pathogenic environments in Chapter 6

2.4 THEORETICAL APPROACHES

The population dynamics of clones of B- and T-cells can be modelled as an ecological process where the clones are treated as different species [70]. One of

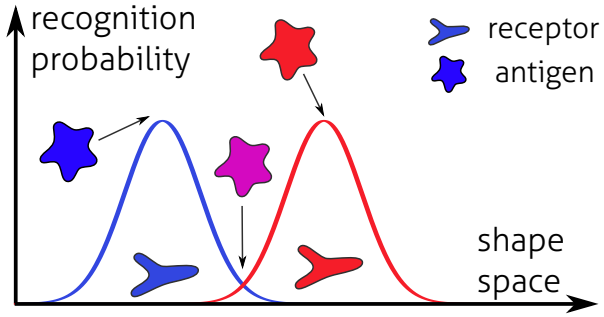

Figure 2.3: Cartoon of how to model cross-reactive binding of receptors to antigens by an effective recognition space, called shape space. Detection probabilty decays with distance (blue and red curve), such that a receptor detects antigens within a certain distance of its position in the space. Close-by receptors have overlapping detection regions and thus detect some of the same antigens (purple antigen).

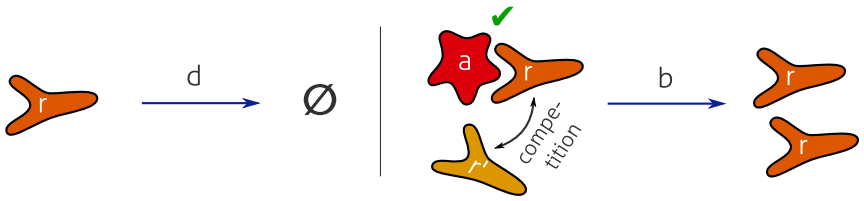

Figure 2.4: Cartoon of a population dynamics of lymphocytes based on competition for antigenic stimulation.

the first such models was introduced by de Boer and Perelson and developed in a series of papers which models competition between clones for antigens using deterministic equations [71, 33, 72, 73]. In more recent works the models have been extended to include the stochasticity arising from the finite number of lymphocytes [74, 75, 76, 63].

In the absence of a precise structural understanding or quantitative data cross-reactivity has mainly been modelled by simple conceptual representations of the receptor-pathogen interaction. In the classical *shape space* picture it is assumed that pathogens and receptors are described by coordinates in a common generally high-dimensional space [77, 51, 78, 30]. The coordinates are an abstract representation of the receptor features that determine antigen recognition. Cross-reactivity is then modeled as a dependence of recognition probability on the distance between receptor and antigen in this space (Fig. 2.3). Short distances in shape space correspond to a good fit between the two molecules, leading to strong recognition, while large distances translate into weak interactions and no or poor recognition.

Let us now more precisely define a population dynamics model of competing lymphocyte clones. The repertoire is described by a set of K clones with abundances $N_r(t)$, $1 \leqslant r \leqslant K$, and the environment by a set of L antigens with concentrations $c_a(t)$, $1 \leqslant a \leqslant L$. The binding probabilities between antigens and clones are encoded in an $K \times L$ interaction matrix f, where $f_{r,a}$ is the probability of receptor r to trigger a response upon encounter with antigen a. The

dynamics of each clone are governed by division b and death rates d (Fig. 2.4)

$$\partial_t N_r = [b(S_r) - d(S_r)]N_r, \tag{2.1}$$

where S_r is a clone-specific antigenic stimulus. It is defined as the sum of the stimuli provided by all antigens

$$S_r = \sum_{a=1}^{L} f_{r,a} A_a(\tilde{N}_a) c_a, \tag{2.2}$$

where A_a is an antigen-specific factor that decreases with increasing coverage of the antigen by the repertoire

$$\tilde{N}_a = \sum_{r=1}^{K} f_{r,a} N_r. \tag{2.3}$$

The factor can be given an interpretation in terms of the availability for binding of antigens given competing receptors [71, 33], which gives rise to a forms of A_a of

$$A_a = 1/(1 + \epsilon \tilde{N}_a), \tag{2.4}$$

where ϵ sets the strength of competition. More generally the model can be seen as an effective model of competition, which might arise through other mechanisms such as faster clearing of antigens for higher coverage. In Chapter 4 we consider a variant of this model with successive infections during which antigen a is present ($c_a = 1$) for a period of time. The infections happen with a probability that is drawn from a probability distribution Q_a. In a mean-field description this reduces to the kind of model we have just described as we will see there.

These models have shown that competition for antigenic stimulation is a sufficient mechanism to maintain a diverse repertoire of lymphocyte receptors [71]. Such competition can furthermore assure a low degree of commonality between the antigens recognized by different receptors [74, 75]. A heavy-tailed distribution of clone sizes can be explained by such models if antigens change over time which gives rise to a fluctuating fitness of clones [63].

Theoretical work in immunology has provided insights into a range of other questions some of which are reviewed in [79, 80]. A major effort of early mathematical work on the immune system has been to understand which mechanisms allow self and non-self to be distinguished [81, 82, 83, 51]. Outside of biology, these studies had a major impact by sparking the field of artificial immune systems, which designs algorithms inspired by the biological tolerance mechanisms [84]. In a series of more recent studies, Kosrmlj and co-workers [85, 86] have used ideas from statistical physics to analyze the effects of thymic selection on the properties of T-cell repertoires. Going beyond the effective descriptions of the recognition process in the population dynamics models, more detailed models have investigated how different trafficking strategies within

lymph nodes impact the efficiency of the adaptive immune system [87]. Finally, the field of ecological immunology has searched an understanding for the evolution of immunity in terms of its benefits and costs [88, 89, 90]. These studies have analyzed how the optimal investment in immune defense depends on trade-offs between immune protection and other traits, on epidemiological factors, and the co-evolutionary dynamics of pathogens.

3

MATHEMATICAL AND COMPUTATIONAL TECHNIQUES

3.1 OPTIMIZATION

Mathematical optimization is an ubiquitous tool across many fields. To give just two examples: in physics many physical laws can be formulated as optimization problems using extremality principles; in computer science machine learning is about finding descriptions of data to minimize some suitably constructed loss function. In this thesis we use optimization techniques to find parameters for which models of biological systems perform best according to some biologically relevant metric. We thus review some definitions, results, and algorithms from the vast literature on optimization problems.

3.1.1 *Classes of optimization problems*

To start let us introduce some basic notation. Consider the problem of minimizing a function $f_0(\mathbf{x})$ over parameters $\mathbf{x}$. Not all parameter values $\mathbf{x}$ are generally attainable – we call those that are feasible. Feasibility is expressed by m inequality and p equality constraints involving functions of the parameters $\mathbf{x}$, $f_i(\mathbf{x})$ and $h_i(\mathbf{x})$:

$$\text{Minimize} \qquad f_0(\mathbf{x}) \tag{3.1}$$

$$\text{subject to} \qquad f_i(\mathbf{x}) \leqslant 0, \quad i = 1, \dots, m \tag{3.2}$$

$$h_i(\mathbf{x}) = 0, \quad i = 1, \dots, p. \tag{3.3}$$

Note that we have followed the commonly used convention of formulating optimization problems as minimization. Maximization can be reduced to this standard form as it is equivalent to minimizing $-f_0(\mathbf{x})$.

It is useful – especially for the choice of a numerical optimization algorithm – to classify optimization problems into classes. *Unconstrained optimization* considers the case where the parameters $\mathbf{x}$ can take arbitrary values, i.e. there are no constraints, $m = p = 0$. The class of *constrained optimization* can further be subdivided by whether the problem only has equality constraints $m = 0$ or also inequality constraints $m > 0$. Note that optimization problems can often be reformulated in equivalent ways, some of which might be more convenient to analyze than others [91]. Equality constraints can for example sometimes be eliminated by redefining the variables of the optimization problem: for example, the constrained problem of optimizing over $(x_1, x_2, \cdots, x_n)$ subject to $\sum x = 1$ can be turned into an unconstrained problem by introducing the variables $(y_1, y_2, \cdots, y_{n-1})$ with $x_i = y_i$ for $i < n$ and $x_n = 1 - \sum_{i=1}^{n-1} y_i$. Further classification is often based on specific problem structure: from least-squares problems arising in data fitting (where f_0 is a sum of squares) to linear programming (where f_0 and the constraints are linear functions).

In the following subsections we review three types of optimization problems of practical relevance to the work presented in this thesis in some more detail. First, we review convex optimization in Sec. 3.1.1.1, which is a broad and particularly important class of optimization problems. Second, we review noisy optimization in Sec. 3.1.1.2, where the goal is to minimize the expectation value of a function based on noisy samples. Third, we review multi-objective optimization in Sec. 3.1.1.3, where there instead of a single objective there are multiple objectives to optimize at the same time.

3.1.1.1 *Convex optimization*

used in Ch. 4

The convexity of an optimization problem is a useful mathematical property. To lay the ground we review some basic definitions and results from convex analysis [91].

A *convex set* is a set which contains any line segment between two of its points, *i.e.*

$$C \text{ convex} \Leftrightarrow \forall \mathbf{x}, \mathbf{y} \in C \, \forall \theta \in [0,1] : \theta \mathbf{x} + (1-\theta)\mathbf{y} \in C. \tag{3.4}$$

A function $f : C \to \mathbf{R}$ defined on a convex domain C is a *convex function* if the line segment between any two points on the graph of the function lies above the graph, *i.e.*

$$f \text{ convex} \Leftrightarrow \forall \mathbf{x}, \mathbf{y} \in C \, \forall \theta \in [0,1] : f(\theta \mathbf{x} + (1-\theta)\mathbf{y}) \leqslant \theta f(\mathbf{x}) + (1-\theta) f(\mathbf{y}). \tag{3.5}$$

In other words a function is convex if its *epigraph* (the set of points above or on its graph) is a convex set. For a differentiable function to be convex it necessarily has to verify the first-order condition

$$\forall \mathbf{x}, \mathbf{y} \in C : f(\mathbf{x}) + \boldsymbol{\nabla} f(\mathbf{x})^{\mathsf{T}} \cdot (\mathbf{y} - \mathbf{x}) \leqslant f(\mathbf{y}) \tag{3.6}$$

as follows directly from the definition of a convex function. Here $\boldsymbol{\nabla}$ denotes taking the gradient, and $\mathbf{x} \cdot \mathbf{y}$ denotes the inner dot product between $\mathbf{x}$ and $\mathbf{y}$. If the function is twice differentiable it also has to verify the second-order condition

$$\forall \mathbf{x}, \mathbf{y} \in C : (\mathbf{y} - \mathbf{x})^{\mathsf{T}} \cdot \boldsymbol{\nabla}^2 f(\mathbf{x}) \cdot (\mathbf{y} - \mathbf{x}) \geqslant 0, \tag{3.7}$$

i.e. its Hessian needs to be positive semi-definite. To establish convexity of a function instead of verifying these conditions it is often more convenient to show that the function can be composed from elementary convex functions by convexity-preserving operations [91].

Convex optimization is the minimization of a convex function over a convex set. It includes many practical important problems such as linear least-squares regression and quadratic programming. Research during the last two decades has provided a relatively complete theory of and efficient algorithms for convex optimization problems [91]. The same are still lacking for more general optimization problems. An important property of convex optimization prob-

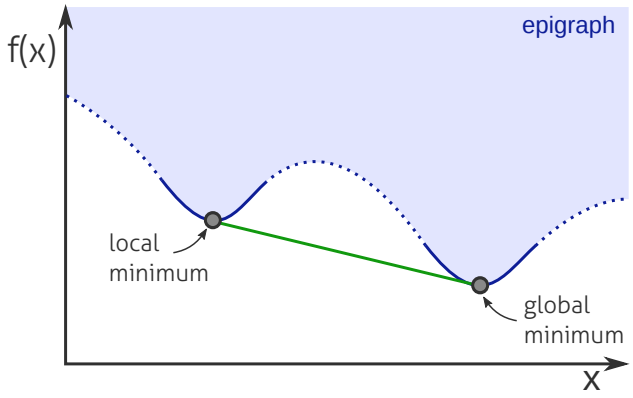

Figure 3.1: A function with multiple distinct minima cannot be convex. In the vicinity of a minimum the function can be approximated by a parabola (solid blue lines). Convexity requires every line connecting two point on the graph of the function to lie completely within the epigraph (light blue shaded region) of the function, *i.e.* above or on the graph. A line connecting two local minima (solid green line) cannot fulfill this requirement.

lems is the absence of distinct local optima: every locally optimal solution of such a problem is also globally optimal (Fig. 3.1).

3.1.1.2 *Noisy optimization*

used in Ch. 6,7

Sometimes a precise evaluation of the objective function is either impossible or exceedingly computationally expensive. Examples are optimization based on a complex simulation, or the evaluation of loss function over large data sets in machine learning. How to optimize a function given noisy evaluations is thus an important problem. Such a noisy optimization problem can be stated as finding the minimum over a bounded domain Ω of a function $f(\mathbf{y})$

$$\min_{\mathbf{x}\in\Omega} f(\mathbf{x}) = \min_{\mathbf{x}\in\Omega} \mathbb{E}[F(\mathbf{x}, \omega)], \tag{3.8}$$

which is not known explicitly, but needs to be approximated by evaluating a function $F(\mathbf{x}, \omega)$ dependent on a random variable ω. The optimization problem Eq. 3.8 is an example of a stochastic programming problem [92].

3.1.1.3 *Multi-objective optimization*

used in Ch. 7

Multi-objective optimization searches for optimal parameters $\mathbf{x}$ within a domain Ω to minimize a vector of objective functions $\mathbf{f}$ [91]:

$$\min_{\mathbf{x}\in\Omega} \mathbf{f}(\mathbf{x}) \tag{3.9}$$

For nontrivial multi-objective optimization problems there is no single solution minimizing all objectives simultaneously. To define what optimality means for a multi-objective problem we can use the minimal requirement that a solution $\mathbf{x}$ should not be worse for all objectives than what is achievable for any other

feasible $\mathbf{x}'$. This criterion is known as Pareto optimality. The family of solutions which are Pareto optimality is called the Pareto frontier.

To go beyond Pareto optimality needs additional information about the relative importance of objectives. Multi-objective optimization problems can be turned into a scalar optimization problem given weights $\boldsymbol{w}$ assigned to different objective functions:

$$\min_{\mathbf{x}\in\Omega} \boldsymbol{w}^{\mathsf{T}}\mathbf{f}(\mathbf{x}) \tag{3.10}$$

Pareto optimality is a necessary condition for optimality in this scalarized problem (if all weights are strictly positive).

3.1.2 *Optimality conditions*

As the solution to an optimization problem is an extremal point a variational principle holds: the optimization criterion is stationary with respect to all infinitesimal changes allowed by the constraints. This yields a set of conditions – called Karush-Kuhn-Tucker conditions – that need to hold at the optimum [93, 91]. They provide a generalization of the method of Lagrange multipliers to problems involving inequality constraints. Analytical solutions to optimization problems can sometimes be obtained by solving the equations defining these conditions. For convex optimization problems these conditions are not just necessary but also sufficient for optimality.

To begin we write down the Lagrangian of the optimization problem defined in Eqs. 3.1-3.3

$$\mathcal{L}(\mathbf{x},\boldsymbol{\lambda},\boldsymbol{\nu}) = f_0(\mathbf{x}) + \sum_{i=1}^{p} \lambda_i h_i(\mathbf{x}) + \sum_{i=1}^{m} \nu_i f_i(\mathbf{x}), \tag{3.11}$$

where $\boldsymbol{\lambda}$ is a p-dimensional vector of Lagrange multipliers enforcing the equality constraints and $\boldsymbol{\nu}$ is a m-dimensional vector of Lagrange multipliers enforcing the inequality constraints. The optimal $\mathbf{x}^\star$ is an extremum of this Lagrangian, so the stationarity condition

$$\boldsymbol{\nabla} f_0(\mathbf{x}^\star) + \sum_{i=1}^{p} \lambda_i^\star \boldsymbol{\nabla} h_i(\mathbf{x}^\star) + \sum_{i=1}^{m} \nu_i \boldsymbol{\nabla} f_i(\mathbf{x}^\star) = 0 \tag{3.12}$$

holds at the optimum for associated $\boldsymbol{\lambda}^\star$ and $\boldsymbol{\nu}^\star$. $\mathbf{x}^\star$ further needs to be feasible, *i.e.* it must not violate any of the constraints

$$f_i(\mathbf{x}) \leqslant 0, \tag{3.13}$$

$$h_i(\mathbf{x}) = 0. \tag{3.14}$$

If an inequality constraint $f_i(\mathbf{x}^\star) > 0$ is satisfied as a strict inequality at the optimum then the inequality does not restrict the allowed infinitesimal changes. Therefore the constraint should not have an influence on the Lagrangian. This so called complementary slackness condition, requires the La-

grange multipliers associated with the inequality constraint to be zero unless the constraint is active,

$$\nu_i^\star f_i(\mathbf{x}^\star) = 0. \tag{3.15}$$

Lastly, even if an inequality constraint is active only changes pointing outwards of the constraint set are prohibited. Therefore the sign of Lagrange multipliers for inequality constraints needs to be non-negative

$$\nu_i^\star \geqslant 0. \tag{3.16}$$

3.1.3 *Numerical techniques*

A plethora of algorithms has been developed to numerically solve optimization problems [94, 93, 91]. I do not review the full breadth of such algorithms here, but introduce some common ideas before discussing two particular algorithms I have used during this thesis.

One strategy to find a minimum is to start somewhere and then go down the hill. The simplest implementation of this idea is gradient descent, which follows the direction of steepest descent. Gradient descent is slow to converge and can suffer from issues of ill-conditioning of the objective function. One can choose better descent directions by using curvature information about the objective function. Newton methods are based on the exact curvature obtained from the Hessian whereas Quasi-Newton methods iteratively approximate the Hessian from gradient information only. Newton methods have better convergence properties then gradient methods and so are generally preferred. For large scale optimization involving many parameters the computational demand of every iteration can be prohibitively high so that one needs to fall back to gradient methods.

To go down the hill you need to know not just where you are but how the height of the hill changes around you. If the objective function has an explicit expression it is best to analytically take its derivatives needed for the methods discussed in the previous section. What if the function is not known explicitly? This is a problem known as black-box optimization. Sometimes the derivatives can be approximated using finite differences, which allows to use the algorithms described above. This is not always possible though, for example the function might not be differentiable everywhere. *Derivative-free optimization* algorithms then find their use [95, 96]. We use and extend such an algorithm to optimize a function based only on noisy function evaluations in Sec. 3.1.3.2.

To terminate iterative optimization algorithms one needs to define a stopping criterion. Stopping criteria are designed to be quantitative measures of convergence. A frequently used measure is that the norm of the gradient falls below some predefined threshold. For convex problems a lower bound for the cost can be established by solving a linear programming problem as follows:

$$f_{lb} = f\left(x^k\right) + \min_{x \in C} \left[\left(x - x^k\right) \cdot \nabla f\left(x^k\right)\right] \leqslant f\left(x^*\right). \tag{3.17}$$

The linear programming problem $\bar{x}^k = \operatorname{argmin}_{x \in C} \nabla f\left(x^k\right)^T \left(x - x^k\right)$ is solved explicitly [93] by

$$\bar{x}^k = e_{i^*},\ i^* = \operatorname*{argmin}_i \left(\nabla f\left(x^k\right)\right)_i, \tag{3.18}$$

in the case where C is the probability simplex and where e_i denotes the ith unit vector. We can use this lower bound to define a stopping criterion for the numerical optimization that ensures convergence up to a tolerance ϵ:

$$\frac{f\left(x^k\right) - f_{lb}}{f_{lb}} < \epsilon \tag{3.19}$$

3.1.3.1 *Projected gradient method*

used in Ch. 4

In constrained optimization simply taking a step in a descent direction might violate the constraints. The idea of the projected gradient method is to project the iterates back onto the constraint set [97]. The projection of a point $\mathbf{y}$ onto a set C is defined as the point within the set closest to $\mathbf{y}$

$$\mathcal{P}(\mathbf{y}) = \operatorname*{argmin}_{\mathbf{x} \in C} \frac{1}{2}\|\mathbf{x} - \mathbf{y}\|^2, \tag{3.20}$$

where $\|\mathbf{x}\|$ denotes the Euclidean norm of $\mathbf{x}$. Projection is a quadratic programming problem, the solution of which generally can be time-consuming. The practical feasibility of optimization methods involving projections relies on the existence of efficient algorithms for solving this problem for some practically important constraints such as for bound or simplex constraints [98].

A projected gradient algorithm iterates according to

$$\mathbf{x}^{k+1} = \mathcal{P}\left(\mathbf{x}^k - s^k \boldsymbol{\nabla} g\left(\mathbf{x}^k\right)\right), \tag{3.21}$$

where $\mathcal{P}$ denotes the projection onto C as defined above, and s^k is the step size taken in the direction of the gradient. The projected gradient method is conceptually simple, but slow to converge, $g\left(\mathbf{x}^k\right) - g\left(\mathbf{x}^\star\right) \simeq \mathcal{O}(1/k)$ [99].

Recently there has been an increasing interest in optimal first-order methods that increase the theoretical bound on the rate of convergence to $\mathcal{O}(1/k^2)$ [99, 97]. These methods only use function and gradient information and retain the simplicity of ordinary first order methods. They achieve accelerated convergence by adding a momentum term to ordinary gradient descent. An accelerated projected gradient method follows

$$\mathbf{y}^{k+1} = \mathbf{x}^k + \omega^k\left(\mathbf{x}^k - \mathbf{x}^{k-1}\right) \tag{3.22}$$

$$\mathbf{x}^{k+1} = \mathcal{P}\left(\mathbf{y}^{k+1} - s^k \boldsymbol{\nabla} g\left(\mathbf{y}^{k+1}\right)\right) \tag{3.23}$$

with a suitably chosen sequence ω^k. Following [97] we use $\omega_k = \frac{k}{k+3}$.

The step size s is determined by backtracking [99]: we iteratively decrease s by multiplication by $\beta < 1$ until $g(z) \leqslant g\left(\mathbf{y}^k\right) + \left(\mathbf{z} - \mathbf{y}^k\right) \cdot \nabla g\left(\mathbf{y}^k\right) + \frac{1}{2s}(\mathbf{z} - \mathbf{y}^k)^2$, where $\mathbf{x} \cdot \mathbf{y}$ denotes the inner dot product between $\mathbf{x}$ and $\mathbf{y}$, and $\mathbf{z} =$

$\mathcal{P}\left(\mathbf{y}^k - s\nabla g\left(\mathbf{y}^k\right)\right)$. In practice we determine s in this way at the first step of the optimization and then keep it fixed based on this initial estimate.

3.1.3.2 *Noisy optimization*

used in Ch. 6,7

Given N samples of the function for independently drawn ω (see Eq. 3.8) we can approximate

$$f(\mathbf{y}) \approx \frac{1}{N}\sum_{i=1} F(\mathbf{y}, \omega_i), \tag{3.24}$$

which will converge by the law of large numbers as $N \to \infty$. By considering this so-called sample average approximation [92] as a deterministic function to be optimized, classical deterministic optimization methods can be used.

Here we adapt a pattern search algorithm to solve a noisy optimization problem with bound constraints. The algorithm combines compass search, a simple pattern search algorithm which allows to easily incorporate bound constraints [95], with the idea of adapting the number of evaluations of F dynamically to control the noise in the approximation [100, 101]. An advantage of the algorithm over alternative methods for noisy optimization such as stochastic approximation [102] is that it allows one to define stopping criteria in terms of parameter convergence instead of relying on more indirect stopping criteria such as decrease conditions.

Let us define the set of search directions considered at each iteration as $\mathcal{D} = \{\pm\boldsymbol{e}_i | i = 1, ..., n\}$, where $\boldsymbol{e}_i$ is the i-th unit vector; and let us further define $P_\Omega(\mathbf{y}) = \operatorname{argmin}_{\mathbf{y}'\in\Omega} |\mathbf{y}' - \mathbf{y}|^2$ as the projection of a point $\mathbf{y}$ onto Ω [91]. The projection onto box constraints considered here is particularly simple and computationally efficient as it just sets coordinate entries outside of the bounds to the bound value. Given an initial guess for the parameter vector $\mathbf{y}^0$, an initial step size Δ^0, and an initial number of times F should be sampled N^0 the algorithm proceeds as follows to find the optimal parameter vector to within a tolerance of Δ_{tol}:

1. **Initialize** the parameter vector $\mathbf{y} \leftarrow \mathbf{y}^0$, step size $\Delta \leftarrow \Delta^0$, and number of samples $N \leftarrow N^0$.
2. **While** ($\Delta \geqslant \Delta_{tol}$) or ($\mathbf{y}$ or N updated during last iteration):
   a) **For each** step $\Delta\mathbf{d}$ along the positive and negative coordinate directions $\mathbf{d} \in \mathcal{D}$:
      i. **If** $f(\mathbf{y}') < f(\mathbf{y})$ (as judged from N samples of F at both points), where $\mathbf{y}' = P_\Omega(\mathbf{y} + \Delta\mathbf{d})$, then update parameter vector:

$$\mathbf{y} \leftarrow \mathbf{y}'$$

      ii. **Else If** new point $\mathbf{y} + \Delta\mathbf{d}$ is feasible, i.e. $\mathbf{y} + \Delta\mathbf{d} \in \Omega$, and if $f(\mathbf{y} + \Delta\mathbf{d}) = f(\mathbf{y})$ can not be ruled out based on N samples of F at both points (criterion below), then either one oversteps the

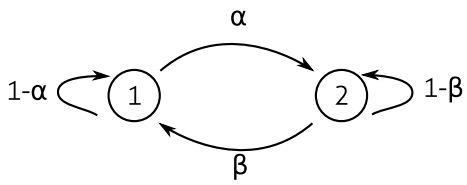

Figure 3.2: Sketch of a two-state Markov chain. The process switches stochastically between states with switching rates that only depend on the current state (Markov property).

minimum or statistical power is insufficient. Therefore first try half-step in the same direction, and if it fails increase sampling:

A. **If** $f(\mathbf{y} + \frac{\Delta}{2}\mathbf{d}) < f(\mathbf{y})$ (as judged from N samples of F at both points), then update the parameter vector and reduce the step size:

$$\mathbf{y} \leftarrow \mathbf{y} + \frac{\Delta}{2}\mathbf{d}$$
$$\Delta \leftarrow \Delta/2$$

B. **Else** increase sampling:

$$N \leftarrow 2N$$

b) **If** no updates during preceding loop then contract pattern size:

$$\Delta \leftarrow \Delta/2$$

For the comparisons between objective function values, we use hypothesis testing on N paired samples of F, i.e. we evaluate $F(\mathbf{y}, \omega_i), F(\mathbf{y}, \omega_i)$ for $\omega_i, i = 1, ..., N$ and calculate pairwise differences. The hypothesis testing uses a confidence level α, which indirectly controls how much the function is sampled. To correct for the multiple tests performed for different directions, we use a Bonferroni correction by using a confidence level $\alpha/(2n)$ for individual tests, where $2n$ is the number of search directions.

3.2 STOCHASTIC PROCESSES

3.2.1 *Markov chains*

used in Ch. 6,7

Markov chains describe stochastic processes that transitions between discrete states without memory (Fig. 3.2). The Markov property that the probability of moving to a new state only depends only on the current state means that a Markov chain is fully defined by its transition matrix $\mathbf{T}$. In discrete time the distribution over states $\mathbf{p}_t$ changes in time according to the Master equation

$$\mathbf{p}_{t+1} = \mathbf{p}_t \mathbf{T}. \tag{3.25}$$

If the chain is irreducible and all of its states are positive recurrent then a unique steady state distribution $\mathbf{x}_s$ exists. As it is unchanged by the dynamics it is a left eigenvector of the transition matrix with eigenvalue 1. The other eigenvalues of $\mathbf{P}$ determine the speed of convergence to the steady-state distribution.

To illustrate these points consider the simple example of a two-state Markov chain (Fig. 3.2). The transition matrix is

$$T = \begin{bmatrix} 1-\alpha & \alpha \\ \beta & 1-\beta \end{bmatrix}. \tag{3.26}$$

Solving for the left eigenvector corresponding to eigenvalue 1 gives the steady-state distribution

$$\boldsymbol{\pi} = \left(\tfrac{\beta}{\alpha+\beta}, \quad \tfrac{\alpha}{\alpha+\beta} \right) \tag{3.27}$$

The second eigenvalue is $\lambda = 1-\alpha-\beta$. It is related to the convergence speed of the distribution against the stationary distribution with the time constant of this approach given by $\tau = -1/\ln(\lambda)$.

3.2.2 *Fokker-Planck and Langevin formalism*

used in Ch. 5

Expanding on the preceding discussion of Markov chains we now consider Markov processes with continuous states and in continuous time. A Markov jump process is defined by the rates $W(x'|x)dx'$ of jumps from state x to $[x', x'+dx')$ in a small time interval dt. The temporal evolution of the probability of being in state x follows the Master equation

$$\partial_t p(x,t) = \int dx' \left[W(x|x')p(x',t) - W(x'|x)p(x,t) \right], \tag{3.28}$$

which can be understood as a continuity equation for probability flux.

The Fokker-Planck equation is an approximation of this Master equation in the limit where $W(x'|x)$ is highly peaked around x [103]. Such a limit of small jumps is relevant in many physical situations. For example, a grain of pollen diffuses in water due to many random interactions with the surrounding water molecules, every one of which only has a very small effect. To derive the Fokker-Planck equation we express the transition probability as a function of the size $r = x - x'$ of the jump

$$W(x|x') = W(x', r). \tag{3.29}$$

The Master equation (Eq. 3.28) in this notation reads

$$\partial_t p(x,t) = \int dr\, W(x-r, r)p(x-r,t) - \int dr\, W(x,-r)p(x,t). \tag{3.30}$$

The shift from x to $x - r$ in the first term is small by assumption so we can Taylor expand up to second order, which yields the Fokker-Planck equation

$$\partial_t p(x,t) = -\partial_x \left[a_1(x)p(x,t)\right] + \frac{1}{2}\partial_y^2 \left[a_2(x)p(x,t)\right], \tag{3.31}$$

where the coefficients are the first and second moment of the jump rates

$$a_n(x) = \int dr W(x,r) r^n. \tag{3.32}$$

The Fokker-Planck equation describes a diffusion process. We write $a_1 = \mu$ for the drift coefficient and $a_2 = D$ for the Diffusion coefficient.

Instead of describing the temporal evolution of the probability density with the Fokker-Planck equation, it is sometimes more intuitive to think about individual trajectories of the stochastic process. Still assuming that the jumps are small we can think of them as a stochastic force added to the deterministic dynamics of the system. This gives rise to the Langevin equation

$$\partial_t x(t) = \mu(x) + \sigma(x)\xi(t), \tag{3.33}$$

where $\xi(t)$ is Gaussian white noise with correlation function $\langle \xi(t)\xi(t')\rangle = \delta(t - t')$. Langevin equations are known in the mathematical literature as stochastic differential equations, where they are defined in terms of Wiener processes. The Langevin equation is equivalent to a Fokker-Planck equation with $\mu(x) = \mu(x)$, and $D(x) = \sigma^2(x)$ in the Ito formalism (see [104] for more details).

Numerically the simplest scheme for simulating trajectories according to a Langevin equation is the Euler-Maruyama algorithm [104]. In Ito formalism given a time step Δt it calculates the next position as

$$x(t + \Delta t) = x(t) + \mu(x(t))\Delta t + \sigma(x(t))\Delta\xi, \tag{3.34}$$

where $\Delta\xi$ is a Gaussian variable with mean 0 and variance Δt. As the Gaussian noise has variations of the order of $\sqrt{\Delta t}$ this algorithm has the lower convergence order of $1/2$ compared to the Euler algorithm for deterministic systems. Higher order schemes have been developed, see e.g. Kloeden and Platen [105] for an in-depth treatment of numerical methods for solving stochastic differential equations.

The formalisms can be extended to multiple dimensions, where the multidimensional Langevin equation is given by

$$\partial_t \mathbf{x}(t) = \boldsymbol{\mu}(\mathbf{x}(t), t) + \boldsymbol{\sigma}(\mathbf{x}(t), t)\boldsymbol{\xi}(t) \tag{3.35}$$

with $\langle \xi_i(t)\xi_j(t')\rangle = \delta_{i,j}\delta(t - t')$. The corresponding Fokker-Planck equation for the probability density $p(\mathbf{x}, t)$ is [104]

$$\frac{\partial p(\mathbf{x},t)}{\partial t} = -\sum_i \frac{\partial}{\partial x_i}[\mu_i(\mathbf{x},t)p(\mathbf{x},t)] + \frac{1}{2}\sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j}[D_{ij}(\mathbf{x},t)p(\mathbf{x},t)], \tag{3.36}$$

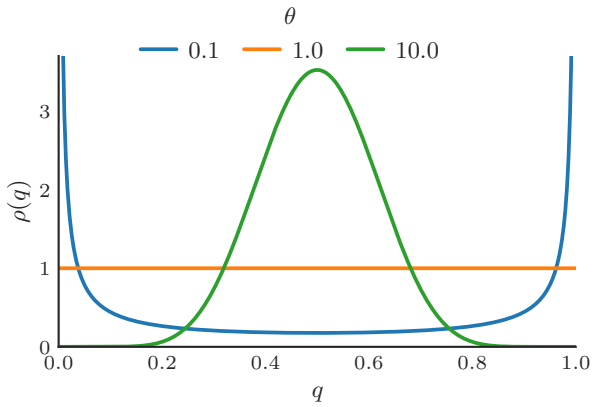

Figure 3.3: Steady-state distribution of allele frequency for neutral Wright-Fisher diffusion with recurrent mutation. The distribution is single-peaked distribution for $\theta > 1$ and bimodal for $\theta < 1$.

with μ defined as before and a diffusion tensor

$$D_{ij}(x,t) = \sum_k \sigma_{ik}(x,t)\sigma_{jk}(x,t). \tag{3.37}$$

3.3 ELEMENTS OF EVOLUTIONARY THEORY

3.3.1 *Genetic drift*

In finite populations the frequency of alleles changes over time even in the absence of selection. Genetic drift can be understood within the Wright-Fisher model, which considers a population of N individuals reproducing at discrete non-overlapping generations. It assumes that offspring sample alleles randomly from the parent generation.

We consider more explicitly the case of two alleles (A and a) in a haploid population. Assume that in the parent generation there are n individuals with allele A. In the next generation the number of offspring with allele A follows a binomial distribution with N draws and success probability $q = n/N$. On average the fraction q of A alleles remains unchanged but there are stochastic changes due to the variance of the fraction of $q(1-q)/N$. These stochastic changes accumulate overtime, which leads to one of the alleles eventually becoming fixed in the population. After fixation the whole population is homogeneous and there is no further change in the absence of new mutations. *used in Ch. 5*

The strength of genetic drift depends on the population size and vanishes in infinitely large populations. In large populations the relative changes in population composition between generations are small. In this limit the dynamics can be approximated by treating q as a continuous variable. Also taking a continous limit in time rescaled by $1/N$ allows applying the reasoning outlined in Sec. 3.2.2 to obtain a Fokker-Planck description,

$$\partial_t \rho = \frac{1}{2}\partial_q^2(\rho q(1-q)), \tag{3.38}$$

of genetic drift for the conditional probability distribution function $\rho(q,t)$. This diffusion approximation is widely applied in population genetics [106, 107]. Note the differing terminology between evolutionary theory and diffusion processes: genetic drift gives rise to a diffusion term, while drift in the diffusion-sense describes deterministic changes.

If there are recurrent mutations between the alleles then fixation is no longer an absorbing boundary condition. The Fokker-Planck equation for such neutral drift with recurrent mutations is

$$\partial_t \rho = -\frac{1}{2}\partial_q(\rho(\alpha(1-q)-\beta q)) + \frac{1}{2}\partial_q^2(\rho q(1-q)), \tag{3.39}$$

given population scaled mutation rates $\alpha = 2Nu$ ($\beta = 2Nv$), where u (v) is the probability of a mutation producing allele A (a) per individual and per generation. Instead of recurrent mutations immigration into the population can also be considered as a source of continued variation. Immigration with rate m from an external population leads to a drift term $m(\bar{q}-q)$, where $\bar{q}$ is the mean of the external population [108]. This reduces to Eq. 3.39 with $\alpha = m\bar{q}$ and $\beta = m(1-q)$.

The stationary solution $\rho_s^{(\alpha,\beta)}(q)$ of Eq. 3.39 is defined by $\partial_t \rho_s^{(\alpha,\beta)}(q) = 0$. Integrating once we obtain

$$\partial_q(\rho_s^{(\alpha,\beta)}(q)q(1-q)) = -\rho_s^{(\alpha,\beta)}(q)(\beta q - \alpha(1-q)), \tag{3.40}$$

which is a separable differential equation with solution

$$\rho_s^{(\alpha,\beta)}(q) = \frac{1}{Zq(1-q)}\exp\left[-\int dq\,\frac{\beta q-\alpha(1-q)}{q(1-q)}\right], \tag{3.41}$$

where Z is set by normalization. Expanding the integrand in partial fractions we obtain

$$\int dq\,\frac{\beta q-\alpha(1-q)}{q(1-q)} = \int dq\,\left[\frac{\beta}{1-q}-\frac{\alpha}{q}\right] = -\beta\ln(1-q)-\alpha\ln(q), \tag{3.42}$$

which leads to

$$\rho_s^{(\alpha,\beta)}(q) = \frac{1}{Z}q^{\alpha-1}(1-q)^{\beta-1}. \tag{3.43}$$

This distribution is known as the Beta distribution and its normalization $Z = B(\alpha,\beta) = \Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha+\beta)$ is the Beta function.

To understand the result consider the case of symmetric mutation rates $\alpha = \beta =: \theta$. For $\theta > 1$ the steady-state distribution is unimodal (Fig. 3.3): mutations dominate genetic drift and q only fluctuates little around the mean frequency for balanced mutations. For $\theta < 1$ the steady-state distribution is bimodal: mutations are more rare and often one of the two alleles is close to fixation with occassional switches.

The model can be extended to more than two alleles. The particular case of parent-independent mutation is notable as it remains analytically solvable. We use $\theta_i = 2Nu_i$ to denote the population-scaled mutation rate towards allele i,

where u_i is the probability of a mutation producing allele i per individual and per generation. The dynamics is described by the Fokker-Planck equation

$$\partial_t \rho = -\frac{1}{2}\sum_{i=1}^{K-1} \partial_{q_i}\left[(\theta_i - |\theta| x_i)\rho\right] + \frac{1}{2}\sum_{i,j=1}^{K-1} \partial_{q_i}\partial_{q_j}\left[q_i(\delta_{i,j} - q_j)\rho\right], \quad (3.44)$$

where $\theta = (\theta_1, ..., \theta_K)$ and $|\theta| = \sum_{i=1}^{K} \theta_i$.

3.3.2 *Evolutionary strategies in fluctuating environments*

used in Ch. 6,7

The question of which evolutionary strategies allow populations to best cope with environmental fluctuations is a long-standing question in evolutionary biology [109, 110, 111, 112, 113]. On a theoretical level, there are connections of this question to information theory [111, 114], and non-equilibrium statistical physics [115, 116].

The dynamics of a population of organisms reproducing at discrete times t in a fluctuating environment can be described by keeping track of the distribution of phenotypes [117] (see Fig. 7.1). At each generation, an individual produces a stochastic number of offspring, whose distribution depends on the phenotype σ of that individual, and the state of the environment x_t. We denote the mean of the number of offspring by $f(\sigma, x_t)$. To adapt its phenotype a population can use different sources of information about the environment [114]: it can actively sense the environmental state, or it can passively adapt by switching stochastically and using the inherited information about phenotypes selected in the past. The value of these sources of information depends on sensing costs and the statistics of the environment [111].

Let $N_t(\sigma|x_{t'<t+1})$ be the mean number of organisms in the population at time t with protection σ, for a given environment history $(x_{t'<t+1})$. The change in population composition from one generation to the next is governed by the reproductive success of individuals in each state σ, modified by stochastic phenotype switches from parents to offspring:

$$N_{t+1}(\sigma|x_{t'<t+1}) = f(\sigma, x_t)\sum_{\sigma'} \pi(\sigma|\sigma', x_t) N_t(\sigma'|x_{t'<t}), \quad (3.45)$$

where $\pi(\sigma|\sigma', x_t)$ is the switching probability from phenotype σ' to phenotype σ. Note that the protection phenotype switching probability might depend on the state x_t of the environment. To simplify notations, we omit the condition on the environment $(\cdot|x_{t'<t})$ when referring to conditional means in the following.

Population growth is an inherently multiplicative process. The total population size at time T is given by the product $N_T = N_0 \prod_{t=0}^{T} Z_T$, where Z_t is the relative increase of the population size at generation t. The success of an adaptive strategy can be measured in term of its long-term growth rate,

$$\Lambda = \lim_{T\to\infty} \frac{1}{T} \ln N_T/N_0 = \lim_{T\to\infty} \frac{1}{T}\sum_{t=0}^{T} \ln(Z_t). \quad (3.46)$$

The strategy with maximal long-term population growth rate outperforms any other strategy in the long run for almost every sequence of environments in populations of infinite size. It thus provides a measure of long-term fitness [117].

A particular class of strategies are so-called bet-hedging strategies which increase long-term growth rate by reducing the variability of Z_t at the expense of its mean [118, 113]. By Jensen's inequality decreasing the variability increases the long-term growth rate in Eq. 3.46 as it is a concave function of Z_t. More concretely, setting $Z_t = \langle Z \rangle + \delta_t$ we can develop $\ln Z_t$ around the mean $\langle \ln Z \rangle$ up to second order in δ_t as $\ln Z_t \approx \ln\langle Z \rangle + \delta_t/\langle Z \rangle - \delta_t^2/\langle Z \rangle^2$. By definition $\langle \delta_t \rangle = 0$ so $\langle \ln Z_t \rangle = \ln\langle Z \rangle - \langle \delta_t^2 \rangle/\langle Z \rangle^2$, which shows that to optimize long-term growth a balance needs to be struck between increasing the arithmetic mean of Z_t and reducing its variance across environmental conditions.

3.4 BAYESIAN INFERENCE AND DECISION THEORY

used in Ch. 5

Assume you observe data $\mathbf{x}$ with a probability $P(\mathbf{x}|\theta)$ (called likelihood) that depends on parameters θ of a statistical model. The task of statistical inference is to deduce something about θ from observations. Bayesian inference is based on a subjective view of probabilities: probabilities describe degrees of belief. This is opposed to the frequentist interpretation of probability as the limit of the relative frequency of an event in a large number of repetitions. Even before having seen any data we already have some beliefs about sensible values of θ which we summarize in a *prior distribution* $P(\theta)$ on θ. To update our belief upon seeing the data we use Bayes theorem as follows:

$$P(\theta|\mathbf{x}) = \frac{P(\mathbf{x}|\theta)P(\theta)}{Z}, \tag{3.47}$$

where $Z = P(\mathbf{x})$ is a normalizing constant The *posterior distribution* $P(\theta|\mathbf{x})$ contains everything there is to know about θ.

As we have seen Bayesian inference is conceptually straightforward. In practice, however, it can sometimes be computationally difficult to apply Eq. 3.47. A common trick to make the calculations tractable is to use special prior distributions. These priors are conjugate with respect to the likelihood in the sense that the posterior stays in the same functional family as the prior. Then the Bayesian update can be done simply in terms of the parameters of the family of functions.

We can extend Bayesian inference to recursive inferences about an only partially observed Markov process [119]. The observation $\mathbf{x}(t)$ at time t is dependent on the current state of the stochastic process $\theta(t)$. Assume that from previous measurements we have a prior distribution $P(\theta, t^-)$ then we update our belief about the state of the stochastic process simply according to Eq. 3.47 to the posterior distribution $P(\theta, t^+)$. Until the next measurement the stochastic process evolves further, which should be reflected in our beliefs about θ.

We can write the probability distribution of θ at time t as an integral of the product of the distribution at an earlier time t and a conditional probability,

$$P(\theta, t) = \int d\theta' \, P(\theta, t|\theta', t')P(\theta', t'). \tag{3.48}$$

Given a stochastic process with a Fokker-Planck operator $\mathcal{A}$ (see Sec. 3.2.2) this leads to the equation,

$$\frac{dP(\theta, t)}{dt} = \mathcal{A}P(\theta, t), \tag{3.49}$$

which needs to be solved starting from $P(\theta, t^+)$ as the initial condition. This represents using the model of the stochastic process to predict the future. The prediction then serves as a prior for the next encounter, which makes the procedure recursive. The procedure we just outlined is variably known as recursive Bayesian estimation of sequential Bayesian filtering [119].

Consider that a decision needs to made with respect to the choice of an action a out of a set $\mathscr{A}$ of admissible actions. The outcome of the action depends on the state of the world encoded by a parameter $\theta \in \Theta$. How should we choose among actions in the presence of statistical information related to their probable outcomes? Answering this question is the aim of statistical decision theory [120, 121]. We define a loss function $L : \Theta \times \mathscr{A} \to \mathbb{R}$ that associates a loss experienced by the decision maker to every action a in every state θ. The statistical information about outcomes is given in the form of the probability distribution $P(\theta)$ of states θ. Bayesian decision theory means choosing the action that minimizes the Bayesian expected loss

The optimist reader might prefer to think in terms of gains instead of losses and equivalently define a utility function as the negative of the loss function.

$$\int_\Theta d\theta \, L(\theta, a)P(\theta). \tag{3.50}$$

A common problem in statistics is to choose a point estimate $\hat{\theta}$ to summarize our knowledge $P(\theta|x)$. We can see this problem as a special case of the theory we have just developed. In this case, the set of actions is the choice of estimates $\hat{\theta}$. For the simple quadratic loss function $L(\theta, \hat{\theta}) = (\hat{\theta} - \theta)^2$, simple algebra shows that the optimal estimate is the posterior mean, $\hat{\theta} = \int d\theta \, \theta P(\theta|x)$. Other loss functions lead to other estimators among which the maximum *a posteriori* estimate $\hat{\theta}_{MAP} = \text{argmax}_\theta \, P(\theta|x)$ deserves special mention due to its wide usage.

3.5 DYNAMICAL SYSTEMS

The temporal evolution of the state $\mathbf{x}$ of a dynamical system in continuous times is described by a system of ordinary differential equations (ODEs),

$$\frac{dx}{dt} = \mathbf{f}(\mathbf{x}, t). \tag{3.51}$$

If $\mathbf{f}$ is not explicitly dependent on t then the dynamical system is called autonomous.

Given an initial condition $\mathbf{x}_0$ an ODE can be integrated to obtain an trajectory of the dynamical system. Initial value problems can sometimes be solved analytically. For example linear systems $\mathbf{f}(\mathbf{x}) = \mathbf{A}\mathbf{x}$ have a formal solution $\mathbf{x}(t) = \mathbf{x}_0 e^{\mathbf{A}t}$ in terms of a matrix exponential. Generally one has to revert to numerical solutions using classical algorithms ranging from the simple Euler algorithm to more refined schemes such as Runge-Kutta integration [122].

To learn something about generic features of the dynamics one can turn to the study of attractors of the dynamics [123, 124]. An attractor governs the dynamics after an initial transient starting from all initial conditions within its basin of attraction. We present two tools related to establishing the properties of the simplest attractors, fixed points, in the following: linear stability analysis and Lyapunov functions.

3.5.1 *Linear stability analysis*

used in Ch. 4

A fixed point $\mathbf{x}^\star$ of an autonomous dynamical system is defined by

$$\mathbf{f}(\mathbf{x}^\star) = 0. \tag{3.52}$$

Local stability of such a fixed point means that any slight perturbation around the fixed point decays with time. A small perturbation $\boldsymbol{\delta}(t) = \mathbf{x}(t) - \mathbf{x}^\star$ follows the linearized dynamical equation

$$\frac{d\boldsymbol{\delta}}{dt} = \mathbf{J}\left(\mathbf{x}^\star\right)\boldsymbol{\delta}, \tag{3.53}$$

where the Jacobian matrix $\mathbf{J}(\mathbf{x})$ is defined as

$$J_{i,j} = \frac{\partial f_i(\mathbf{x})}{\partial x_j}. \tag{3.54}$$

Perturbations along a direction corresponding to a right eigenvector of J grow exponentially with a rate given by the corresponding eigenvalue. Stability requires $\lim_{t\to\infty}\|\delta(t)\| = 0$ for all small perturbations δ, which implies that all real parts of the eigenvalues of the Jacobian at the fixed point need to be negative at a stable fixed point. The Hartmann-Grobman theorem [124] ensures that the higher order nonlinear terms do not change the stability as determined by linearization except for the degenerate case, where one of the eigenvalues of $\mathbf{J}$ is zero.

3.5.2 *Lyapunov functions*

used in Ch. 4

Even if a dynamical system has a single stable fixed point $\mathbf{x}^\star$ it does not necessarily always converge to this point. There might be other attractors such as limit cycles so that not all initial conditions are within the basin of attraction of the fixed point. One way to ensure that the fixed point is eventually reached is to demonstrate the existence of a Lyapunov function. A Lyapunov function is a real valued, positive definite function, that decreases along all trajectories of

the dynamical system. More precisely a continuously differentiable function $V(\mathbf{x})$ is a Lyapunov function, if [124]

1. $V(\mathbf{x}) > V(\mathbf{x}^\star)$, *i.e.* V is positive definite,
2. $\frac{dV(\mathbf{x})}{dt} = \boldsymbol{\nabla} V(\mathbf{x}) \cdot \frac{d\mathbf{x}}{dt} = \boldsymbol{\nabla} V(\mathbf{x}) \cdot \mathbf{f}(\mathbf{x}) < 0$.

In contrast to local stability the existence of such a Lyapunov function shows asymptotic stability of the fixed point.

3.6 MEASURES OF SPATIAL ORDER

used in Ch. 4

In Chapter 4 we analyze optimal distributions of immune receptors in shape space. To analyze spatial order in these distributions we rely on two tools from statistical physics. Namely, we calculate radial distribution functions and a normalized power spectral density similar to the structure factor which we define below. This allows us to get insight into the nature of the tiling patterns, that emerge as optimal distributions.

Let us denote the probability of finding a particle (in our case immune receptor) at position $\mathbf{r}$ in space by $P(\mathbf{r})$. The radial distribution function measures correlations between particle numbers as a function of their distance R in space. It is defined as [125]

$$g(R) = \left\langle P(\mathbf{r}) P(\mathbf{r}') \right\rangle_{\|\mathbf{r}-\mathbf{r}'\|=R}, \tag{3.55}$$

where the average is over all pairs of positions with a distance R. Absence of correlations (as in the case of an ideal gas in physics) leads to a flat radial distribution function equal to one everywhere. If receptor placement is correlated the radial distribution function shows characteristic peaks. The normalized power spectral density of a spatial distribution is defined as

$$S(\mathbf{q}) = \left| \int_{\mathbf{r}} d\mathbf{r}\, P(\mathbf{r}) e^{i\mathbf{q}\cdot\mathbf{r}} \right|^2 \Big/ \int_{\mathbf{r}} d\mathbf{r}\, P(\mathbf{r})^2, \tag{3.56}$$

in terms of the wave vector $\mathbf{q}$. For isotropic distribution one can average over all directions of $\mathbf{q}$ to obtain a quantity that only depends on its modulus $q = \|\mathbf{q}\|$. Large (small) q correspond to short (long) distances in space. The value of $S(q)$ indicates the relative power of fluctuations at the spatial scale $2\pi/q$.

The normalized power spectral density coincides with the *structure factor* familiar in physics [126, 125] if $P(\mathbf{r})$ consists of a sum of Dirac Delta peaks. Let us assume that the distribution $P(\mathbf{r})$ consists of N peaks of the same shape $\phi(\mathbf{x})$ at positions $\mathbf{R}_l$, $P(\mathbf{r}) = \sum_{l=1}^{N} \phi(\mathbf{r} - \mathbf{R}_l)$. The numerator of Eq. 3.56 can then be rewritten as $|\phi(\mathbf{q})|^2 I(\mathbf{q})$ in terms of a quantity $I(\mathbf{q}) = \left| \sum_{l=1}^{N} e^{i\mathbf{q}\cdot\mathbf{R}_l} \right|^2$ that depends only on the peak positions and of a quantity $\phi(\mathbf{q}) = \int_{\mathbf{r}} d\mathbf{r}\, \phi(\mathbf{r}) e^{i\mathbf{q}\mathbf{r}}$ that depends only on the peak shape. A quantity similar to $\phi(\mathbf{q})$ arises also in condensed matter physics where it is called *atomic form factor*. The normalization of $S(\mathbf{q})$ is chosen such that it coincides with the *structure factor* for Dirac delta peaks of uniform heights. For Dirac delta peaks $\phi(\mathbf{q}) = 1$, and

$\int_{\mathbf{r}} d\mathbf{r}\, P(\mathbf{r})^2 = 1/N$, so that $S(\mathbf{q}) = I(q)/N$, which is the definition of the structure factor of a point pattern.

For patterns for which the precise local position of the points is random, the absence of structure at the smallest scales translates into a structure factor that approaches one for large q. The $q \to 0$ limit of $S(q)$ indicates how uniformly space is covered at large scales. In the small q limit the structure factor is equal to the relative variance in the number of peaks contained within large areas [127]. For random point patterns the number of peaks in a given area is Poisson distributed and the variance of the number of peaks grows as the volume of the area. The constant relative variance implies $S(\mathbf{q} \to 0) > 0$. There is an interesting class of patterns for which the variance grows sublinearly with volume, implying vanishing variation at the largest scales $S(\mathbf{q} \to 0) = 0$. Such patterns are called hyperuniform [127, 128]. A trivial example are regular patterns. More interestingly one can also construct disordered hyperuniform patterns that combine local randomness $S(\mathbf{q} \to \infty) = 1$ with large scale uniform space coverage ($S(\mathbf{q} \to 0)$) [127, 129, 130, 128].

3.7 IMPLEMENTATION AND COMPUTATIONAL REPRODUCIBLITY

The numerical work reported in this thesis was performed using the scientific *Python* stack including the *Numpy* [131], *Scipy* [132], and *Matplotlib* [133] packages. Time-sensitive parts of the code were accelerated using *Cython* [134], which is a tool that allows to combine the ease of Python programming with the speed of C where it matters.

Science proceeds as a collective endeavour with studies building onto each other. Scientific truth emerges from cycles of independent replication and refinement of results. In practice, however, it is often already difficult to reproduce published research, that is to obtain the same results using the same methods. There currently is thus a big movement in science towards improving reproducibility [135, 136]. For computational science a minimal requirement for reproducibility is the public availability of the source code. All source code associated with the published papers is thus available online: for ease of use on Github, and for persistent storage on Zenodo. If a problem is common enough it might be worth the effort to go beyond reproducibility and enable reusability. This means to provide a well-documented, and generic solution in the form of a software tool/package. In this spirit I have released the Python package *Noisyopt* to provide implementations of common algorithms to solve noisy optimization problems in the Python programming language [137].

The following table gives the relevant references:

Work	Github Link	Zenodo Link
[138]	`http://github.com/andim/optimmune`	`http://dx.doi.org/10.5281/zenodo.16796`
[139]	`http://github.com/andim/evolimmune`	`http://doi.org/10.5281/zenodo.495494`
[140]	`http://github.com/andim/transitions-paper`	`http://doi.org/10.5281/zenodo.495495`
[137]	`http://github.com/andim/noisyopt`	`http://doi.org/10.5281/zenodo.596660`

Part II

IMMUNE REPERTOIRES

This part of the thesis takes a look at immune repertoires from a perspective of how they allocate resources relative to their knowledge of the pathogenic environment. Specifically we ask:

How is a well-adapted immune system organized? And how should a well-adapting immune system remember?

4

HOW A WELL-ADAPTED IMMUNE SYSTEM IS ORGANIZED

This chapter was previously published in Ref. [138].

4.1 INTRODUCTION

The adaptive immune system protects organisms from a great variety of pathogens by maintaining a population of specialized cells, each specific to particular challenges. Together these cells cover the array of potential threats. To recognize pathogens, the immune system relies on receptor proteins expressed on the surface of its main constituents, the B and T lymphocytes. These receptors interact with antigens (small molecular elements making up pathogens), recognize them through specific binding, and initiate the immune response. Each lymphocyte expresses a unique receptor formed from random combinations encoded in the genome. The receptors later undergo selection through the death and division of the lymphocytes that express them, as well as mutations in the case of B lymphocytes. The diversity of the receptor repertoire determines the range of threats that the adaptive immune system can target.

The detailed *composition* of the immune receptor repertoire, and not just its breadth, is important for conferring effective protection against infections. Broadly speaking, a diverse population of receptors will confer wider immunity, and a larger clonal population of a particular receptor will confer more effective immunity against the pathogens to which it is specific. However, there is a tradeoff between diversity and clone sizes because the number of receptors is limited. By selectively proliferating some receptors at the expense of others, the immune system retains a memory of past infections [27], facilitating subsequent immune responses. Furthermore, while infections increase the populations of receptors with the greatest specificity, they can also lead to a reorganization of the immune repertoire as a whole [60].

How should the repertoire be organized to minimize the cost of infections? We develop a framework for answering this question by abstracting key general features of the adaptive immune system: the receptor repertoire is bounded in size, receptors are "cross-reactive" (each antigen binds many receptors; each receptor binds many antigens), and the cost of an infection increases with time. Given these general assumptions, we consider a simplified landscape of pathogens, where infections are drawn from a fixed distribution. By simplifying the setting in this way, and independently of the detailed dynamics of immune responses, we arrive at broad insights about the composition of immune repertoires that are optimal for their pathogenic environments. Our framework is not meant to give a complete account of immunity. To do so we would need to include several other components of the immune system, such as interaction between its innate and adaptive arms and avoidance of auto-immunity. The

latter problem —the challenge of discriminating self from non-self— has been the focus of many theoretical studies of the immune system [77, 83]. This paper primarily investigates the relation between the adaptive repertoire and the pathogenic environment, but we will also discuss how other components and constraints of the immune system can be incorporated into our model.

The theory predicts, counter-intuitively, that the number of receptors specific to rare pathogens will be amplified relative to the probability of encounter, at the expense of receptors for common infections. We also find that two organisms responding to a pathogen distribution will display unique populations of immune receptors, even though their coverage of pathogens will be similar. How can the immune system achieve these sorts of optima? Surprisingly, we find that simple competition between receptor clones can drive the population to the optimal composition for minimizing the cost of infections.

New high throughput methods are making it possible to survey B-cell and T-cell receptor diversity in fish [10, 141], in mice [28, 60] and humans [54, 142, 29, 11]. As methods are developed to better characterize pathogenic landscapes and receptor cross-reactivity, predictions for the composition of optimal repertoires derived from our framework can be directly compared with experiments. To arrive at our results we ask how the immune system should be organized to perform its function well, rather than starting with the detailed dynamics of its components. We are proposing that the universal features of the adaptive immune system follow simply from general statistical considerations, while the detailed dynamical implementation arises from the historical contingencies of evolution.

4.2 DEFINITION OF THE PROBLEM

To find the optimal repertoire distribution we must consider the nature of antigen-receptor interactions, and a penalty that the immune system pays for not recognizing antigens. This penalty must reflect the facts that recognition should happen within reasonable time, before the pathogen colony can significantly increase its size; the interactions between the immune receptors and antigen are probabilistic; and not all antigens are equally frequent. We assume that, although the immune system cannot predict precisely which antigens it will encounter and when, it incorporates an estimate of the *probabilities* of their occurrences. We also take these probabilities to be constant in time. This is an idealization grounded in a separation of timescales, which assumes the distribution of antigens remains constant on timescales on which the immune system adapts.

The above ideas are the basis for our cost function, which reflects the penalty of non-recognition, for a given repertoire and antigenic environment. In Fig. 4.1 we introduce a quantification of the problem. Given Q_a, the probability that the next infection will be caused by antigen a, we model the immune repertoire by a distribution of receptors P_r, from which lymphocytes with the receptor r are drawn at random.

An antigen a and a receptor r interact with a certain strength set by the binding affinity between the two molecules. This is described by the prob-

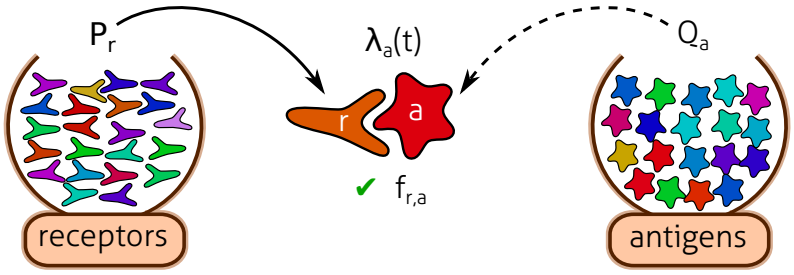

Figure 4.1: Schematic of a statistical model of antigen recognition by the adaptive immune system. After infection, antigen a encounters immune receptor r at random with a rate $\lambda_a(t)$. An encounter leads to a successful recognition with a probability $f_{r,a}$ that reflects the matching between a given antigen–receptor pair.

ability $f_{r,a}$ (which we call the "cross-reactivity function") that an antigen a encountering a receptor r results in a recognition event, leading to the activation of the lymphocyte expressing that receptor. Each encounter of the antigen a with a random receptor has a probability $\tilde{P}_a = \sum_r f_{r,a} P_r$ to lead to recognition and trigger an immune response. Thus, $\tilde{P}_a$ can be viewed as the coverage of antigen a by the repertoire.

Given this coverage, we consider $\bar{F}_a$ – the average harm caused by antigen a. We will show below that, consistent with intuition, $\bar{F}_a$ is a decreasing function of the coverage $\tilde{P}_a$. The overall expected cost is then just the harm averaged over the antigen distribution:

$$\text{Cost}(\{P_r\}) = \langle F \rangle = \sum_a Q_a \bar{F}_a(\tilde{P}_a). \tag{4.1}$$

The need to defend against many antigens at the same time with a limited number of receptors introduces a trade-off. If more receptors recognize an antigen, there are less to protect against other threats.

Lastly, we derive an expression for the average harm $\bar{F}_a$ caused by antigen a. During its time in the periphery, an antigen a will encounter and possibly interact with receptors at a rate $\lambda_a(t)$ which increases with time as the pathogen population grows. Each encounter will occur with a different receptor r drawn from P_r. The mean number of encounters between antigens and receptors after a time t, which we will call *effective time*, is defined as $m_a(t) = \int_0^t d\tau \lambda_a(\tau)$, where $t = 0$ is set by the introduction of the antigen. The time t to the first recognition event, or response time, is random and depends on the coverage $\tilde{P}_a$.

The longer the system fails to detect the antigen, the more likely the infection is to become harmful. We assume that the integrated harm caused by an antigen since the beginning of an infection is an increasing function $F_a(t)$ of the time of first recognition. How exactly F_a grows with time may strongly depend on the type of infection and receptors [79, 143, 40]. The mean harm inflicted to the organism by the attack of an antigen a is then given by this quantity averaged over the distribution of possible response times: $\bar{F}_a = \tilde{P}_a \int_0^{+\infty} dm\, F_a[t_a(m)]\, e^{-m\tilde{P}_a}$, where $t_a(m)$, the inverse function of $m_a(t)$,

is the amount of time it takes for m encounters to occur between the receptors and pathogen a (see App. A.1 for a derivation). The result depends on the cost expressed as a function of the effective time m, $F_a[t_a(m)]$, which we denote $F_a(m)$ to simplify notation. We will consider several specific choices of this effective cost function in Results.

Our aim here is to propose a general framework for thinking about the repertoire. Thus, we do not explicitly model intracellular communication, cell differentiation, activation of co-factors, coordination of different cell types, the interaction with the innate immune system, and the full complexity of the recognition process. The idea is that $F_a(m)$ implicitly summarizes all of these factors in terms of an effective cost.

In general the cost function $F_a(m)$ depends on the antigen a, reflecting the various virulences of different pathogens. To simplify, we can assume that the cost function takes the factorized form: $F_a(m) = \mu_a F(m)$, where μ_a is the pathogen-dependent virulence factor, and $F(m)$ describes how all threats develop with time. The cost will then take the form: $\sum_a \mu_a Q_a \tilde{P}_a \int_0^\infty dm F(m) e^{-m\tilde{P}_a}$. In this expression, the virulence factor μ_a of a pathogen plays the same role as its likelihood Q_a. Some pathogens are rare but very virulent (like anthrax), while others may be common but not very virulent (like the common cold), and an ideal immune system should be able to cope with both. In our model the overall "threat" of a pathogen is expressed as the product of the two, $\mu_a Q_a$. In practice μ_a can be absorbed into the definition of Q_a, and will be omitted in the rest of the paper.

Given such a model of the recognition process, there exists an optimal adaptive immune system, characterized by the choice of the receptor distribution P_r, that minimizes the expected cost in a given antigenic environment Q_a. The optimal repertoire is found by minimizing the expected cost in Eq. 4.1 with respect to P_r, subject to constraints of non-negativity ($P_r \geqslant 0$) and normalization ($\sum_r P_r = 1$). Simple local extremality conditions are sufficient for optimality because our problem can be shown to be convex (see App. A.2). The condition $\sum_r P_r = 1$ is a normalized version of the constraint that the total number of receptors is limited.

4.3 RESULTS

4.3.1 *The optimal repertoire is more uniform than the pathogen distribution*

We can now ask how best to distribute the receptors to minimize the cost (Eq. 4.1) for a given antigenic environment. To begin, we neglect cross-reactivity (later we will see that this is equivalent to looking at the structure of the repertoire at scales larger than the cross-reactivity). In this case antigens and receptors can be associated one by one by a cross-reactivity function $f_{r,a} = 1$ if $r = a$ and 0 otherwise. In this case we can analytically determine the optimal distribution (App. A.4.2) $P_r^* = \max\left[\bar{F}'^{(-1)}(-\lambda/Q_r), 0\right]$, where $\bar{F}'^{(-1)}$ denotes the inverse function of the derivative of $\bar{F}_a = \bar{F}(\tilde{P}_a)$ expressed as a function of $\tilde{P}_a$, and λ is a positive constant fixed by the normalization $\sum_r P_r^* = 1$. Table 4.1 presents results for several representative cost functions.

A simple scenario occurs when the pathogen population grows exponentially in time, as do the cost and the encounter rate—reflecting the proliferative nature of pathogens. In this case the cost grows linearly in the number of encounters, *i.e.* $F(m) = m$ (see App. A.3). Then we find that the optimal fraction of the repertoire taken up by a given receptor is proportional to the square root of the threat (combination of frequency and virulence) of the corresponding antigen $P_r^* \propto \sqrt{Q_r}$. Intuitively, we expect that the optimal repertoire should focus its resources on receptors recognizing the most common or virulent antigens. However this enhanced protection against frequent or virulent antigens comes at the cost of a slower response against the uncommon and harmless antigens, and this bias toward more threatening antigens must remain limited. The square root dependence reflects a particular trade-off between these two opposing constraints, by directing more resources toward more threatening antigens while uniformizing the distribution compared to a linear dependence. Intriguingly, the same square root dependence has been found as an optimal solution for the size of tRNA pools as a function of codon usage [144], and in a model for the screening of suspicious individuals [145].

The extent to which more resources are directed toward more threatening antigens depends on the relative gains and losses of earlier and later recognition events, which are captured in our model by the effective cost function $F(m)$. In general, steeper cost functions imply more flattened distributions of receptors. The cost function $F(m) = m^\alpha$, and its associated optimal distribution $P_r^* \propto Q_r^{1/(1+\alpha)}$, help illustrate this point. Such cost functions can arise when both $m(t)$ and $F(t)$ increase exponentially as a function of time, but with different exponents (see App. A.3). When α is large, the cost of non-recognition increases very quickly with time, calling for an urgent response. Consequently the optimal immune system tends to cover the space uniformly to get all potential threats, even the unlikely ones, under control. Conversely, when α is low, the harm caused by pathogens does not explode with time, permitting the system to recognize the rarer pathogens late, and focus its resources on the common ones.

In some situations, there may be little or even no difference between a late response, or no response at all, because the total harm caused by an infection stabilizes. For example, consider the cost $F(m) = 1 - e^{-\beta m}$ which saturates at large effective times. In this case, the optimal solution (Table 4.1) relates receptor and antigen through a square root as for linear cost, but with a cut-off at low Q_a. This cut-off occurs because there is little benefit to having receptors recognizing rare or harmless antigens, whose recognition is likely to happen late, when differences in recognition times do not matter anymore. This result is consistent with the observation that the immune system may sometimes ignore infections if the harm they cause is too small, as in *e.g.* the case of Simian immunodeficiency virus infection of sooty mangabeys [146].

Real harm may occur only when the effective time m crosses a threshold. This situation can be modeled by taking $F(m) = \Theta(m - m_0) = 0$ for $m < m_0$, and 1 otherwise. In this case the receptor distribution is organized to maximize the chance of detection before m_0. The optimal repertoire for this cost (Table 4.1) has no receptors for the least threatening pathogens (cutoff at

$F(m)$	P^*_r	$A(\tilde{N}_a)$
m^α	$C\,Q_r^{\frac{1}{1+\alpha}}$	$C'(N_{st}/\tilde{N}_a)^{1+\alpha}$
$\ln m$	$C\,Q_r$	$C'(N_{st}/\tilde{N}_a)$
$1-\exp(-\beta m)$	$\max\{C\sqrt{Q_r}-\beta, 0\}$	$C'/(\beta+\tilde{N}_a/N_{st})^2$
$\Theta(m-m_0)$	$\max\{\ln(Q_r)/m_0 - C, 0\}$	$C'\exp(-m_0\tilde{N}_a/N_{st})$

Table 4.1: The cost function $F(m)$ measures the harm caused to an organism by the time that immune receptors have had m encounters with a pathogen. The optimal receptor distribution P^* is determined by minimizing this cost, given a pathogen distribution Q, and a cross-reactivity function $f_{r,a}$ specifying the probability that receptor r binds to antigen a. The second column gives the form of P^* over scales larger than the cross-reactivity. The optimal P^* can be reached as a steady-state resulting from competitive binding between receptors and antigens (see last section of Results) quantified by an "availability function" A. $\tilde{N}_a = \sum_r N_r f_{r,a}$ represents the coverage of antigen a by the repertoire, $N_{st} = \sum_r N_r$ is the total steady state population and $C, C', \beta,$ and m_0 are positive constants.

low probabilities) and a drastically flattened receptor distribution (logarithm of the pathogen distribution).

Is there a cost function for which the receptor distribution is not flattened relative to the pathogen distribution? This occurs in a special case where cost increases very slowly (logarithmically) with effective time. However, in general, cost is minimized by a receptor distribution that is flattened relative to the pathogen distribution.

4.3.2 *Cross-reactivity dramatically reduces diversity in the optimal repertoire*

By allowing receptors to bind to a variety of antigens, cross-reactivity should permit the immune system to reduce the number of receptor types required to cover the whole range of possible threats. We will show that given sufficient cross-reactivity, the optimal immune repertoire concentrates all its resources on a few receptors, which together tile antigenic space.

Following Perelson and Oster [77], we think of receptors and antigens as points in a common high dimensional *shape space*, whose coordinates are associated to unspecified physicochemical properties. For simplicity, we assume that cross-reactivity only depends on the relative position of receptor and antigen in shape space $f_{r,a} = f(r-a)$, where f is a decreasing function of the distance between a and r. Short distances in shape space correspond to a good fit between the two molecules, leading to strong recognition, while large distances translate into weak interactions and poor recognition.

To build intuition, we first consider an analytically solvable example (Fig. 4.2). We describe the space of receptors and antigens by a single continuous number, and assume a Gaussian antigen distribution with variance σ_Q^2, and Gaussian cross-reactivity of width σ, which sets the typical distance within which a receptor and antigen interact. We derive the optimal receptor distributions for costs of the form $F(m) = m^\alpha$ (App. A.4.3.2). For narrow cross-reactivities ($\sigma < \sigma_c = \sigma_Q\sqrt{1+\alpha}$), the optimal receptor distribution is Gaussian with

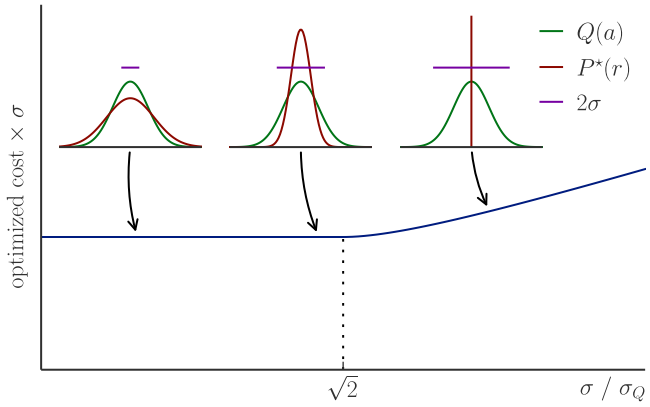

Figure 4.2: The optimal cost and receptor distributions for protecting against a one-dimensional Gaussian antigenic landscape $Q(a)$ of variance σ_Q^2, as a function of the cross-reactivity width σ. As σ increases, the optimal distribution $P^*(r)$ becomes narrower and narrower (left and middle insets), until it concentrates entirely onto a single point, for $\sigma \geqslant \sqrt{2}\sigma_Q$ (right inset). The minimal cost (multiplied by σ for a comparison at constant recognition capability) is constant below the transition point, but increases with σ past it. The cross-reactivity function, which quantifies the affinity between receptor r and antigen a as a function of their distance in shape space, has a Gaussian form: $f(r-a) = \exp[-(r-a)^2/2\sigma^2]$, and the cost function is linear in the effective recognition time, $F(m) = m$.

variance $(1+\alpha)\sigma_Q^2 - \sigma^2$ and the optimal cost is independent of σ. For wide cross-reactivities ($\sigma > \sigma_c$), the receptors are optimally of a single type with reactivity centered on the pathogen distribution, while the optimal normalized cost increases with σ since the receptor is unnecessarily broadly reactive. These results arise from a tension between two opposing tendencies. As in the non cross-reactive case, the need to cover rare pathogens broadens the optimal receptor distribution relative to the pathogen distribution. However, cross-reactivity has the opposite effect, favoring more concentrated distributions.

Does cross-reactivity generically drive the optimal receptor distribution to cluster into peaks? We investigated this question numerically. For concreteness, we consider a linear cost $F(m) = m$, and random pathogen environments in one or two dimensions constructed by drawing each Q_a from a log-normal distribution characterized by a coefficient of variation κ. The shape space is taken to be bounded and discretized, and we use accelerated gradient projection optimization (App. A.5). We find that the optimal repertoire P^* is strongly peaked on a discrete forest of receptors (Fig. 4.3A,B). The width of these peaks decreases as numerical precision is increased, suggesting that in a continuous limit the optimum consists of a weighted sum of Dirac delta functions, *i.e.* distinct, discretely spaced receptors in different amounts (Fig. A.1).

By inspection, the peaks tend to repel each other and to organise into local tiling patterns, as further evidenced by the damped oscillations in the radial distribution function [125] (App. A.6 and Fig. A.2A). This exclusion is a sensible way to distribute resources, as it limits redundant protection against the same pathogens. The spacing between peaks roughly follows the cross-

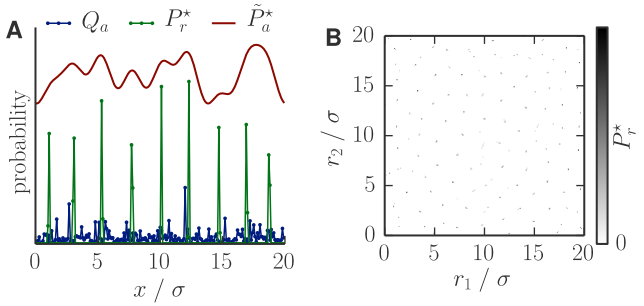

Figure 4.3: Cross-reactivity plays an important role in shaping the optimal repertoire, often leading to highly peaked repertoires. (A)-(B): The optimal receptor distribution P_r^*. for (A) one- and (B) two-dimensional random environments. Despite being peaked, the optimal distribution of receptors covers the antigenic space fairly uniformly, as shown by its coverage by the receptors, $\tilde{P}_a^* = \sum_r f_{r,a} P_r^*$, shown in the one-dimensional case (A). The cross-reactivity and cost functions are the same as in Fig. 4.2. The antigenic landscape Q_a is generated randomly from a log-normal distribution with coefficient of variation $\kappa = 1$.

reactivity scale σ, suggesting that P^* is smooth when viewed at scales larger than σ. Confirming this, $\tilde{P}^*$ (*i.e.* the coverage of the antigenic space by the receptors) smoothly tracks the variations in the antigen distribution Q at a broad scale (Fig. 4.3A). When viewed coarsely in this way, cross-reactivity is irrelevant and P^* tends to the solutions of Table 4.1. Indeed, at these broad scales, the distribution of peaks is uniform, as demonstrated by the very low power in the spectrum of P_r at small wave vectors (App. A.6 and Fig. A.2B), indicating that the number of receptors contained in any given large area of the shape space is very reproducible. This phenomenon of small scale randomness with large-scale regularity is called *disordered hyperuniformity* [127], and arises in jammed packings as evidence of the incompressibility of the material. In biological terms, hyperuniformity means that the distribution of receptor peaks provides a much more uniform coverage of the antigen space than if the peaks were positioned randomly according to a Poisson distribution. For our optimal repertoires small scale fluctuations get smoothed out by cross-reactivity and can be tolerated, while at large scales the fluctuations track variations in the antigenic landscape to provide smooth coverage (see Fig. A.3).

To test the generality of our findings we tested other choices of cross-reactivity functions (App. A.7). We have so far assumed a unique scale σ for cross-reactivity, consistent with recent reports that cross-reactivity is local in antigenic space [147]. However, receptor-antigen recognition can be very specific and sensitive to single mutations [148, 85], or extremely degenerate across very dissimilar antigens [149]. To account for these long-range effects, we also tested fat-tailed cross-reactivity functions. We found that, for a variety of non-Gaussian cross-reactivity functions, long tailed or not, the optimal repertoire remains strongly peaked, although the position, number and strength of the peaks do change (Fig. A.4). Next, to relax the assumption that the cross-reactivity width is uniform across receptors, we tested receptor-dependent cross-reactivities σ_r drawn from a log-normal distribution. While the regu-

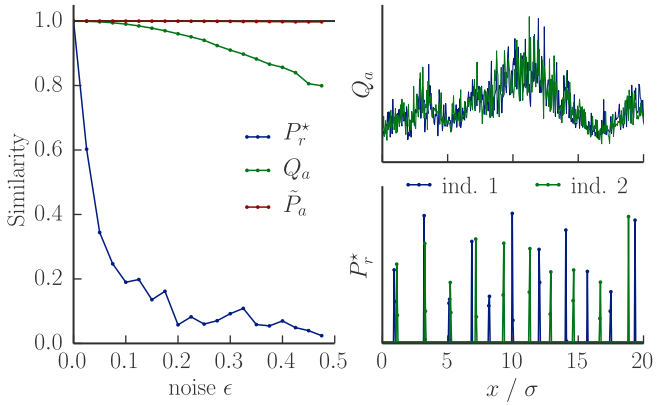

Figure 4.4: Two individuals in the same environment Q_a that see it with slightly different noises have similar coverages of the antigenic space, but achieve it with different receptors. This results in largely non-overlapping repertoires. Shown are the overlaps (normalized to be between 0 and 1) between the experienced pathogen distributions Q_a, the resulting optimal receptor distributions P_r^*, and the corresponding coverages $\tilde{P}_a$, as a function of the noise ϵ with which individuals perceive the environment. The right plots show an example of antigenic environments and optimal receptor distributions for $\epsilon = 0.2$. We calculated the optimal receptor distributions for two individuals 1 and 2 experiencing respective environments Qe^{z_1} and Qe^{z_2}, where Q is a random environment with fluctuations on scales larger than the cross-reactivity σ (power spectrum $\propto 1/(1+(10q\sigma)^2)$) normalized so that its coefficient of variation is 0.5, and z_1, z_2 are Gaussian noises of mean zero and variance ϵ^2. The choice of cost and cross-reactivity functions are the same as in Fig. 4.2.

larity of the local tiling structure is affected by this additionnal heterogeneity (just as we expect in a packing of spheres of variable size) the large-scale hyperuniformity is nonetheless preserved (see Fig. A.5). Next we considered distributions of antigens with correlations across shape space (reflecting *e.g.* phylogenic correlations between pathogens). Again we find peaked optimal receptor distributions (Fig. A.6), similar to those for uncorrelated antigen landscapes. For computational reasons, we restricted our analysis to two dimensional pathogen landscapes, but the analogy with random packing problems that we discussed above allows us to expect that all of these results will hold generally in higher dimensions. Lastly, we incorporated the avoidance of the self by excluding from the optimization all receptors within distance σ of a set of randomly positioned self-antigens (App. A.8 and Fig. A.7). We find that receptors are likely to be found near the boundary of these exclusion zones, but otherwise keep the same general tiling structure.

In summary, the optimal immune repertoire looks random at scales smaller than the cross-reactivity, but has the structure of a disordered tiling at larger scales so that, after accounting for cross-reactivity, the repertoire smoothly covers the pathogen landscape. These findings have an important consequence for different individuals exposed to the same pathogenic environment. Each individual will experience a slightly different spectrum of antigens because of the statistics of encounters and other sources of variability. These slightly different experiences of the same world lead to optimal repertoires with a striking property – the receptor distributions are largely different, even though their coverage of the pathogen landscape is similar after including cross-reactivity (Fig. 4.4). This finding can be compared with surveys of "public" repertoires of immune receptors [30, 60].

4.3.3 *The optimal repertoire can be reached through competition for antigens*

The results presented so far have established how repertoires should be structured to provide optimal protection. Given the complex interdependences between receptors arising from local and global trade-offs, one might think that the globally optimal solution could only be reached via some biologically implausible centralized mechanism distributing resources system-wide. In fact, we will show that the optimal repertoire can be reached through self-organization, via competitive evolution of receptor populations under antigen stimulation.

We consider a model that is similar to that introduced by de Boer, Perelson and collaborators for competitive dynamics of B and T cells [71, 73] (see App. A.9). Its main assumptions are that division of receptor-expressing lymphocytes is driven by antigen stimulation, and that receptors compete for the limited supply of antigens. At each time step, a random antigen α is drawn from the distribution Q_α. Each receptor type r responds to it by expanding or shrinking its population N_r according to its specificity, by an amount $N_r \Delta t[A(\sum_{r'} N_{r'} f_{r',\alpha}) f_{r,\alpha} - d]$, where Δt is the time step: receptors proliferate upon successful recognition of antigens (first term) and die with a constant rate d (second term). In the absence of competition, the proliferation rate should be proportional to $f_{r,\alpha}$, but the antigen α may also bind other receptors,

reducing its availability for receptor r. The *coverage* of antigen a by the repertoire, $\tilde{N}_a = \sum_r N_r f_{r,a}$, quantifies the breadth of the receptor pool competing to bind with a. The availability of antigen a for binding is assumed to be a decreasing function $A(\tilde{N}_a)$ of its coverage. The stimulation of r by a is thus modified to $A(\tilde{N}_a)f_{r,a}$ in the equation for the growth rate. In the limit of fast sampling of antigens, or mean-field limit ($\Delta t \to 0$), this stochastic dynamics is well described by the deterministic differential equations:

$$\frac{dN_r}{dt} = N_r \left[\sum_a Q_a A \left(\sum_{r'} N_{r'} f_{r',a} \right) f_{r,a} - d \right]. \tag{4.2}$$

For a given pathogenic environment, the total steady-state receptor population size N_{st} will be set by the death rate d, which counter-balances growth at steady state. Although in reality the ability of the system to re-organize itself diminishes with age, for simplicity we take all rates to be constant.

The stable fixed points of the mean-field dynamics (Eq. 4.2) realize the optimal repertoires of the previous sections when the availability function A is matched to the cost function $F(m)$ through the relation

$$A\left(\tilde{N}_a\right) = -c'\bar{F}'\left(\tilde{N}_a/N_{st}\right), \tag{4.3}$$

where N_{st} is the total number of receptors $\sum_r N_r$ at steady state. Table 4.1 shows $A(\tilde{N})$ for several cost functions. To understand this result, first note that when binding is *not* cross-reactive the dynamical equations for each receptor are independent, and read: $dN_r/dt = N_r(Q_r A(N_r) - d)$. The availability function now depends only on N_r, meaning that receptors only compete with their own kind — they occupy their own antigenic niche. The steady state size of clone r is thus set by the carrying capacity of that niche, $N_r = A^{(-1)}(d/Q_r)$, or zero if that capacity is negative. With the availability given by Eq. 4.3, this reproduces the optimal repertoire. As seen in Table 4.1, fast growing cost functions correspond to very load-sensitive availability functions. In these cases, rare infections are almost as threatening as frequent ones; therefore the growth of the receptors that are specific to frequent antigens is actively limited to leave room for other receptors. The correspondance of Eq. 4.3 holds when receptor binding is cross-reactive (App. A.10). Cross-reactivity leads to competition amongst receptor types, effectively enforcing an exclusion between similar receptors. This phenomenon, known in ecology as competitive exclusion, is important for lymphocyte dynamics [71], and provides the mechanism by which our dynamical model reproduces the discrete clustering found in the optimal receptor distribution.

To check that the dynamics do converge to the optimum, we simulated numerically the full stochastic dynamics, as well as their mean-field limit (Eq. 4.2), for a random antigenic environment in two dimensions, with $A(\tilde{N}) = 1/(1 + \tilde{N}/N_0)^2$. Fig. 4.5 shows the dynamics of the receptor distribution $P_r(t) = N_r(t)/\sum_{r'} N_{r'}(t)$, as well as its cost relative to the optimal solution, as a function of time. Starting from a random distribution of receptors, the repertoire reorganizes into localized peaks that become increasingly prominent and well-separated with time. Three independent runs of the stochastic simulation all

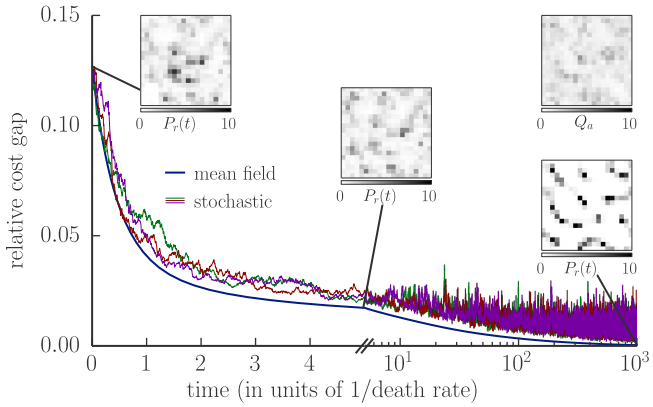

Figure 4.5: The immune repertoire can self-organize to a state that minimizes cost and provides protection against infections via competitive evolution of receptor populations stimulated by antigens. Numerical simulations of the population dynamics, as well as its mean-field limit (Eq. 4.2), show how competition causes a random initial receptor distribution to fragment into a highly peaked pattern (insets representing $P_r(t) = N_r(t)/\sum_{r'} N_{r'}(t)$). The top-right inset represents the antigenic environment Q_a driving the dynamics (generated from a lognormal noise of power spectrum $\propto 1/(1+(5q)^2)$ and coefficient of variation 1). Departure from optimality, as measured by the relative cost gap $[\langle F\rangle(P_r(t)) - \langle F\rangle(P_r^*)]/\langle F\rangle(P_r^*)$, decreases with time and eventually reaches zero in the mean-field limit. The three independent runs of the stochastic dynamics show reproducible results. We use the availability function $A(\tilde{N}) = 1/(1+\tilde{N}/N_0)^2$ with $N_0 = 10^6$, a death rate $d = 0.001$ and a cost function $F(m) = 1 - e^{-\beta m}$ with $\beta = 1/110$. The space size is 10σ. The initial condition was drawn from a lognormal noise of power spectrum $\propto 1/(1+(5q)^2)$, with coefficient of variation 2 and $\sum_r N_r(0) = 1.1 \cdot 10^8$. In the stochastic simulations, the time between antigen presentations is $\Delta t = 0.005 d^{-1}$ (200 infections per cell lifetime).

converge approximately to the global minimum of the cost, with most of the improvement achieved within a few cell lifetimes. Convergence is exact in the mean-field limit, indicating that the steady-state solution is indeed optimal.

In summary, competitive dynamics can allow the immune repertoire to self-organize into a state that confers high protection against infections. In the special case when the availability A is scale invariant, the expected cost is a Lyapunov function of the dynamics (App. A.11), implying that the optimum is reached regardless of the initial condition. Note, however, that the dynamics of Eq. 4.2 is expected to slow down with age, as the plasticity of the adaptive system decreases due to the diminishing number of naive cells [150].

4.4 DISCUSSION

We introduced a general framework for predicting the optimal composition of the immune repertoire to minimize the cost of infections contracted from a given distribution of antigens. This framework can be extended in several ways to be more biologically faithful, *e.g.* by accounting for antigen-dependent infection dynamics, and evolution of the pathogenic landscape. Our predictions can be tested in experiments that study how the environment influences

the composition of immune repertoires, either via high-throughput sequencing surveys of receptor populations [151, 60], or by sequencing receptors specific to given antigens [40]. The comparison between theory and experiment will provide insight into the functional constraints of antigen recognition by the immune system.

There are many situations where living systems must respond to very diverse and often very high dimensional spaces of external influences using strictly limited resources. To sense, internally represent, and then respond to these influences, organisms often employ a large diversity of components, such as cell types or genes [152], each sensitive to a small part of the space. For example, the retina supports a diverse population of ganglion cell types, each sensitive to a different visual feature, that collectively represent the behaviorally salient aspects of visual scenes [153]. Likewise, the mammalian olfactory system contains some ~1000 distinct receptors that each bind widely to odorants, and collectively cover olfactory space [154]. In these cases, the limited repertoire of component types provides a key constraint on information processing. Faced with such constraints, living systems must commit resources wisely, adapting to the structure of the environment, and balancing breadth of coverage against depth of resolution, in light of priorities, costs and constraints [155]. We have shown that these elements also shape the optimal form of the immune repertoire.

Our finding that cross-reactivity causes the optimal repertoire to fragment is related to the concept of *limiting similarity* due to competitive exclusion in ecological settings [156]. In this context, empty regions of phenotypic space result when competition is important on the scale at which resources vary, and continuous coexistence of species only occurs in exceptional cases [157]. In general, niche-space heterogeneity promotes species clustering [158], recalling our finding that any heterogeneous antigen distribution leads to fragmentation of the optimal repertoire. The conceptual connection between the immune repertoire and ecological organization is even clearer in our dynamical model where species compete for an array of resources (the antigens), and grow in relation to their success in securing resources.

Although this study relies on a simple abstraction of the adaptive immune system, we expect that our framework and results will extend to other distributed protection systems where diverse threats are addressed by an array of specific responses. For example, the immune system of bacteria, or CRISPR system [7], for which population dynamics models have already been proposed [159], could be studied within a similar framework to predict the relative abundance of CRISPR spacers and corresponding viruses in a co-evolving population of bacteria and viruses.

Acknowledgements. We thank O. Rivoire and F. Zamponi for helpful discussions. The work was supported by grant ERCStG n. 306312. VB was supported by the Fondation Pierre-Gilles de Gennes, NSF grants PHY-1058202 and EF-0928048. Portions of this work were done at the Aspen Center for Physics, supported by NSF grant PHY-1066293. AM was supported by a DAAD Promos stipend.

5 HOW A WELL-ADAPTING IMMUNE SYSTEM REMEMBERS

This chapter is currently being prepared for publication.

5.1 INTRODUCTION

A defining feature of the adaptive immune system is its ability to learn from pathogen exposure by building immunological memory [22, 81, 160, 51, 79, 35]. The importance of such memory is shown for example by the much higher rate with which newborns contract infectious disease before having developed sufficient immunological memory. Clinically, immunological memory forms the basis of vaccination: exposure to attenuated pathogens or pathogen-derived antigens builds memory which then protects against the disease.

The cells of the adaptive immune system (B and T lymphocytes) detect pathogen by specific binding of receptor proteins expressed on their surface to pathogen markers (antigens). To fight off pathogens immune cells with the required specificity proliferate massively during an infection. After an infection most of these cells die and the number of specific cells drops again, although it remains much higher than before the infection [42]. This increased number of specific cells after an infection is the basis of memory in the adaptive immune system. Despite the relatively short life span of memory cells [161], constant balanced turnover allows the elevated levels of protection to remain relatively stable for decades after the infection in the absence of persistent antigens [162]. Memory helps the organism defend against recurring threats, but needs to be balanced against broad protection from yet unseen threats. Furthermore specific memory might lose its usefulness over time as pathogens evolve away from being detected.

How much benefit can immunological memory provide to the organism? How much memory should be kept to minimize harm from infections? To answer these questions we extend a framework for predicting optimal repertoires for given pathogen statistics [138] by explicitly considering the optimal inference of pathogen frequencies as a sequential Bayesian forecasting problem [119]. We derive the optimal repertoire dynamics in a temporally varying environment and analyze its features. We show how biologically realistic population dynamics can approximate the optimal repertoire dynamics. Furthermore we analyze the conditions under which memory provides a benefit. Comparing the predictions of our theory to signatures observed in real immune systems, strongly argues for a view in which the adaptive immune system is trying to learn a highly sparse distribution of antigens.

Our work provides a novel perspective on immunological memory from a Bayesian viewpoint. Bayesian decision theory has been a useful theoretical framework in other areas of biology from cellular sensing and response [163, 164, 165, 166] to neuroscience [167, 17, 168, 169]. A Bayesian view of adaptive

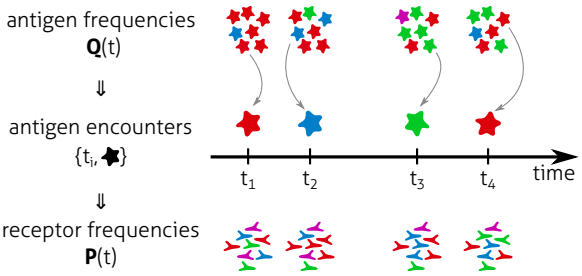

Figure 5.1: Schematic of a model of repertoire dynamics in a changing antigenic environment. An organism encounters antigens at time points $\{t_i\}$ according to a Poisson process with rate λ . The probability of encountering antigen α at time t depends on the antigen environment at that time $Q_\alpha(t)$. Encounters come at a cost that depends on the fraction P_α of the repertoire specific to the antigen. The repertoire of receptor specificities changes over time as the immune systems deals with encountered antigens. In this paper, we ask what is the optimal repertoire dynamics to minimize expected harm from infections.

immunity might similarly be useful to think about the dynamics of immune repertoires. Concretely, we show how this view links the amount of memory production to the variability of pathogenic environments and derive a number of further testable predictions about dynamical features of well-adapting repertoires.

5.2 DEFINITION OF THE PROBLEM

5.2.1 *Pathogen encounters in a changing environment*

In our model the adaptive immune system encounters antigens following a Poisson process with rate λ (Fig. 5.1). There are K different antigens $\alpha = 1, 2..., K$ against which the immune system provides protection through lymphocytes with a cognate receptor. An encounter at time t happens with a probability $Q_\alpha(t)$ with an antigen α, which is equal to its relative frequency in the pathogenic environment at that time point. We denote the vector of antigen frequencies by $\mathbf{Q} = (Q_1, ..., Q_K)$. Here and in the following we denote vectors by bold symbols. How much harm the infection does depends on what fraction $P_\alpha(t)$ of the repertoire is specific to antigen α. The repertoire composition is described by the vector of receptor frequencies $\mathbf{P} = (P_1, ...P_K)$. Given the history of the antigens $\alpha_1, \alpha_2, ..., \alpha_n$ encountered at times $t_1 < t_2 < ... < t_n < t$, we ask how to optimally choose $\mathbf{P}(t)$ to minimize harm from infections. The question of which repertoire dynamics a well-adapting immune system should use, is the question of how the past experience should best be used to prepare for future infections.

We consider that the pathogenic environment follows a stochastic dynamics, which we write generically as a Fokker-Planck equation for the conditional probability distribution $\rho(\mathbf{Q}, t)$

$$\frac{\partial \rho(\mathbf{Q}, t)}{\partial t} = \mathcal{A}\rho(\mathbf{Q}, t), \tag{5.1}$$

where $\mathcal{A}$ is a differential operator pertaining to the dynamics.

For concreteness, we consider the stochastic dynamics of pathogen frequencies in a finite population that changes due to immigration from an external reservoir with scaled rates θ^0 and genetic drift, which using a diffusion-approximation leads to (see Sec. 3.3.1)

$$\begin{aligned}\tau\frac{\partial\rho(\mathbf{Q},t)}{\partial t} = &-\frac{1}{2}\sum_{a=1}^{K-1}\frac{\partial}{\partial Q_a}\left[(\theta_a^0 - |\theta^0|Q_a)\rho(\mathbf{Q},t)\right] \\ &+\frac{1}{2}\sum_{a,b=1}^{K-1}\frac{\partial^2}{\partial Q_a\partial Q_b}\left[Q_a(\delta_{a,b}-Q_b)\rho(\mathbf{Q},t)\right],\end{aligned} \tag{5.2}$$

where τ sets the time scale of the pathogen dynamics, θ^0 is a K-length vector, and $\delta_{a,b}$ is the Kronecker delta, which is 1 if $a = b$ and 0 otherwise. Here and in the following we denote the norm of a vector $\mathbf{x}$ by $|\mathbf{x}| = \sum_i x_i$. This dynamics while clearly simplified reproduces some of the main features of real pathogen environments (Fig. B.2). First, at a fixed point in time different pathogens are present but with large differences in their relative frequencies. Second, the dominant pathogens change over time, such as is the case for e.g. the flu where the dominant strain changes every few years.

5.2.2 *Minimizing the cost of infection*

As the immune system encounters pathogens the immune repertoire – in a statistical language – samples the pathogen distribution. From these samples it can learn something about the pathogen distribution to inform its repertoire choice. For tractability, we consider a limit in which repertoire changes in response to infections are fast compared to the time between infections and, in which costs of repertoire remodelling are secondary to the costs of immune defense. At every point in time the optimal repertoire is then chosen independently based on the current knowledge about the pathogen environment. We have thus split the problem of the optimal repertoire dynamics into two: how to allocate resources, and how to infer pathogen frequencies.

The choice of the repertoire $\mathbf{P}$ is determined by the mean cost of the next infection $c(\mathbf{Q},\mathbf{P})$ given the pathogen frequencies $\mathbf{Q}$. We summarize the knowledge the immune system has at time t about the pathogen frequencies by a belief distribution $B(\mathbf{Q},t)$, which gives the subjective probabilities that the environment is $\mathbf{Q}$. According to Bayesian decision theory the optimal repertoire $\mathbf{P}^\star$ should then minimize the expected cost

$$\mathbf{P}^\star(t) = \underset{\mathbf{P}}{\operatorname{argmin}}\langle c(\mathbf{Q},\mathbf{P})\rangle_{B(\mathbf{Q},t)}, \tag{5.3}$$

where the average is over the belief distribution.

The cost of an infection depends on how well an organism is protected against a particular pathogen: the larger the number of lymphocytes in the repertoire specific to the pathogen, the lower is the infection cost on average. The immune system samples the antigen at random until a cell with the cog-

nate receptor recognizes the pathogen and declenches the immune reaction. The cost of infection depends on the time to recognition. Given a function $F_a(m)$ defining how quickly an infection with pathogen a becomes more harmful with the effective time to its recognition, we can derive the mean cost of an infection, $\bar{F}_a$, which is a decreasing function of the fraction P_a of cells in the repertoire specific to that pathogen (see Chapter 4). To obtain the mean cost of the next infection we need to average over the probabilities $\mathbf{Q}$ of encountering different pathogens,

$$c(\mathbf{Q}, \mathbf{P}) = \sum_a Q_a \bar{F}_a(P_a). \tag{5.4}$$

The expected cost of the next infection for a repertoire $\mathbf{P}$ averaged over the belief distribution thus reduces to

$$\langle c(\mathbf{Q}, \mathbf{P}) \rangle_{B(\mathbf{Q},t)} = c(\hat{\mathbf{Q}}(t), \mathbf{P}), \tag{5.5}$$

where $\hat{\mathbf{Q}}(t) = \langle \mathbf{Q} \rangle_{B(\mathbf{Q},t)}$. Eq. 5.5 reduces the choice of an optimal repertoire given a belief distribution (Eq. 5.3) to the choice of an repertoire in a fixed distribution with the estimated frequencies $\hat{\mathbf{Q}}(t)$. The latter problem has a unique solution as shown in Chapter 4, so knowing $\hat{\mathbf{Q}}(t)$ gives us the optimal repertoire $\mathbf{P}^\star(t)$

$$\mathbf{P}^\star(t) = G(\hat{\mathbf{Q}}(t)), \tag{5.6}$$

where the function G depends on the scaling of costs with effective time $F_a(m)$ (Ch. 4).

5.3 RESULTS

5.3.1 *Optimal repertoire dynamics*

We consider here how to optimally learn the pathogen frequencies from past encounters assuming the immune system knows $\mathcal{A}$, i.e. the dynamics the pathogen frequencies follow.

To infer the pathogen frequencies in a Bayesian manner the immune system starts from a prior belief $B(\mathbf{Q}, 0)$ about the probability of having a particular set of pathogen frequencies $\mathbf{Q}$, which should be equal to the steady-state distribution of the dynamics. The belief is continously updated as time goes on and pathogens are encountered. The problem is a sequential Bayesian prediction problem [119], which can be decomposed into prediction steps that propagate the belief forward in time according to the dynamical model of pathogen dynamics, and into update steps that incorporate the new evidence from pathogen encounters (Fig. 5.2. The prediction steps uses the Fokker-Planck operator of the environmental dynamics $\mathcal{A}$ to project the belief $B(Q, t)$ forwards in time

$$\frac{dB(\mathbf{Q}, t)}{dt} = \mathcal{A}B(\mathbf{Q}, t). \tag{5.7}$$

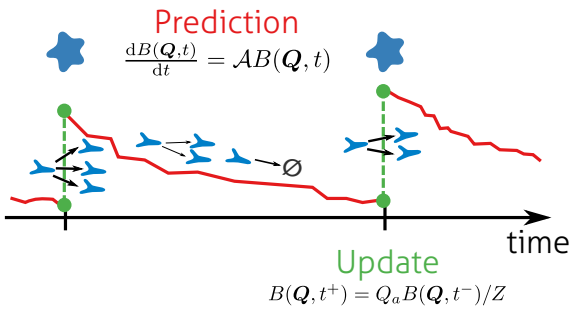

Figure 5.2: Sketch of an optimal repertoire dynamics in a changing antigenic environment. Upon encounters the expected frequency of the pathogen increases and thus the coverage (red line) in the optimal repertoire goes up. This increase in coverage can be biologically interpreted as memory produced by a proliferation of the specific lymphocytes. In the absence encounters the belief relaxes back to steady-state, which corresponds to a decay in memory. The decay of memory can arise from memory cells having a higher death than homeostatic proliferation rate.

The update step combines the Likelihood Q_a of observing the sampled pathogen a given a pathogen distribution Q with the prior belief about the distribution of pathogens $B(\mathbf{Q}, t^-)$ before the eencounter to calculate the belief after the encounter $B(\mathbf{Q}, t^+)$,

$$B(\mathbf{Q}, t^+) = Q_a B(\mathbf{Q}, t^-)/Z \tag{5.8}$$

where $Z = \int d\mathbf{Q}\ Q_a B(\mathbf{Q}, t^-)$.

The update step asks for a higher protection against the pathogen that was just encountered. Such a higher protection can be achieved by proliferating the immune cells specific to the pathogen and keeping some of these proliferated cells around after the infection, which is the basis of immunological memory in the adaptive immune system. The repertoire changes during the prediction step depend on the pathogen dynamics $\mathcal{A}$. The dynamics in Eq. 5.2 is mean-reverting so the protection should decay back to a broad steady-state in the absence of restimulation. This can be achieved by having memory of a finite lifetime, which is a feature of immunological memory. A dynamics close to the optimal Bayesian repertoire dynamics can hence conceivably be implemented using biologically plausible mechanisms (Fig. 5.2).

5.3.2 *Quasi-static limit*

To start we derive the dynamics of the belief in the quasi-static limit, where the sampling is fast compared to the relaxation time scale of the environmental dynamics. In this case, the changes in the prediction phase can be neglected, $\frac{dB(\mathbf{Q},t)}{dt} \approx 0$, and the inference reduces to the iterative application of Eq. 5.8.

The steady-state distribution of the pathogen dynamics (Eq. 5.2) is a Dirichlet distribution with the parameter vector θ^0 [170]. The Dirichlet distribution is a distribution over probability distributions $\mathbf{x}$. It is parametrized by a param-

eter vector θ of the same dimensionality as x, and has the probability density $\rho(\mathbf{x})$

$$\frac{1}{Z(\theta)} \prod_i x_i^{\theta_i - 1}, \tag{5.9}$$

where the normalizing constant is the multivariate Beta function, which can be defined in terms of Gamma functions as $Z(\theta) = \prod_i \Gamma(\theta_i)/\Gamma(|\theta|)$.

Starting from an initial belief $B(\mathbf{Q}, 0)$ about pathogen frequencies equal to the steady-state Dirichlet distribution the belief after the first pathogen encounter is

$$B(\mathbf{Q}, t_1^+) = \prod_a Q_a^{\theta_a^0 + \delta_{a,a_1} - 1}/Z, \tag{5.10}$$

which is a Dirichlet distribution with parameters θ,

$$\theta_a = \theta_a^0 + \delta_{a,a_1}, \tag{5.11}$$

The update of the belief distribution corresponds to updating the parameters of the Dirichlet distribution. This feature of the model makes the updating fully analytically tractable. In the language of Bayesian statistics it stems from the fact that the Dirichlet distribution, which is the steady-state distribution of the pathogen dynamics, is a conjugate prior distribution to the categorical Likelihood function of getting infected by a particular pathogen [171].

By the iterative application of Eq.5.11 we obtain the belief distribution after the n^{th} pathogen encounter as the Dirichlet distribution with parameters

$$\theta(t) = \theta^0 + \mathbf{n}(t), \tag{5.12}$$

where $n_a(t) = \sum_{i=1}^{n} \delta_{a,a_i}$ counts the number of times the immune system has encountered pathogen a up to time t_n. As θ is a sum of counts and the prior parameters, the elements of θ^0 are often called pseudocounts. To turn this update of beliefs into the optimal repertoire dynamics we use Eq. 5.6 and the formula

$$\hat{\mathbf{Q}} = \theta/|\theta| \tag{5.13}$$

for the mean of the Dirichlet distribution.

How much can the immune system lower the cost of protection by remembering the past? Here we answer the question within the quasistatic limit and show that the benefit of memory depends on the sampling rate and the sparsity of the environment (Fig. 5.3).

To do so we numerically simulate the optimal repertoire dynamics for pathogen distributions $\mathbf{Q}$ drawn from a symmetric Dirichlet distribution defined by a parameter vector θ^0 of all θ^0. We consider environments with different degrees of variability by varying θ^0 (Fig. 5.3A-C). For small θ^0 the antigenic environment is effectively sparse with a small number of antigens predominating. We vary θ^0 and the dimensionality of the antigenic space K and calculate

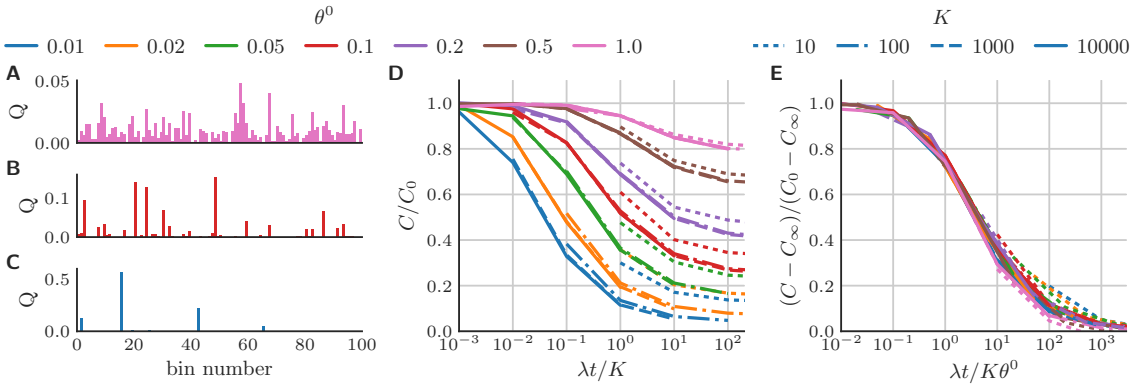

Figure 5.3: Advantage of immunological memory depends on sufficient sampling. (A)-(C): Distributions of K pathogens drawn from symmetric Dirichlet distributions ($\theta^0 = (\theta^0, \theta^0 ..., \theta^0)$), where θ_0 controls the variability in frequencies. (D,E): The cumulative cost of infections C using optimal Bayesian repertoire dynamics depends on time t, pathogen encounter rate λ, and features of pathogen distribution K, θ^0. (D) Cumulative cost normalized by the baseline cost for a uniform repertoire as a function of the number of samples per dimension $\lambda t/K$. (E) Relative cost gap as a function of number of samples per reduced dimension $\lambda t/K\theta^0$. Parameters: Cost is measured as time to recognition, $F_a(m) = m$.

the cumulative cost $C = \sum_i^j \bar{F}_{a_i}(P^\star_{a_i}(t_i^-))$ of all infections up to a time t (sum up to largest j for which $t_j < t$). To reduce variance in this cost we average over independent simulations.

The benefit of immunological memory relative to the baseline infection cost of C_0 for a uniform repertoire increases with the parameter θ^0 (Fig. 5.3D), i.e. memory pays off more if pathogenic frequencies are more variable. For $\theta^0 \ll 1$ memory already provides a benefit even if the number of samples (= pathogen encounter) is small compared to the dimensionality K of the distribution to be learned, this is for $\lambda t/K < 1$ We define the relative cost gap as $(C - C_\infty)/(C_0 - C_\infty)$, where C_0 is the cumulative cost for the uniform repertoire and C_∞ the cumulative cost for the optimal repertoire for precisely known $\mathbf{Q}$. The numerical results suggests that this relative cost gap scales as a function of the sampling of the relevant pathogens $\lambda t/K\theta^0$ for large enough K and short lifetimes (Fig. 5.3E). The more sparse environments are (= the smaller θ^0) the less pathogen encounters are needed for immunological memory to be beneficial.

If the immune system does not know precisely the variability of antigen frequencies, i.e. when it assumes a wrong θ^0, this leads to a suboptimal performance (Fig. 5.4B). Even with a misspecified prior immunological memory still reduces infection costs with better sampling (Fig. 5.4A). In fact, in the limit of high sampling the prior does not matter and the repertoire eventually converges to the correct distribution. The biggest difference is observed at intermediate sampling $\lambda t/K \sim 1$ where a misspecification of θ^0 by a factor of ten leads to an about 30% higher cost (Fig. 5.4C). This demonstrates a remarkable robustness of immunological memory to misspecified priors of the variability of the pathogen distribution.

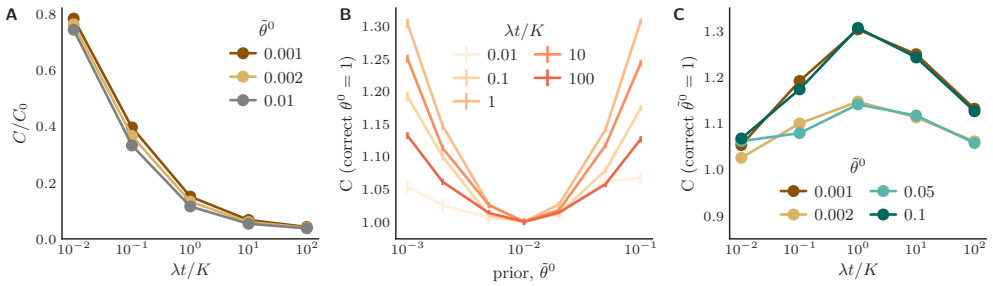

Figure 5.4: Benefit of immunological memory with misspecified prior on variability. (A) Decrease in mean cost of infections with incorrect priors $\tilde{\theta}^0 = \theta^0$ (brown lines) and correct prior (grey line) with longer lifetime. (B) Relative cost as a function of prior for different lifetimes of the organism. (C) Relative cost as a function of lifetime. Parameters: $N = 1000$, $\theta_0 = 0.01$.

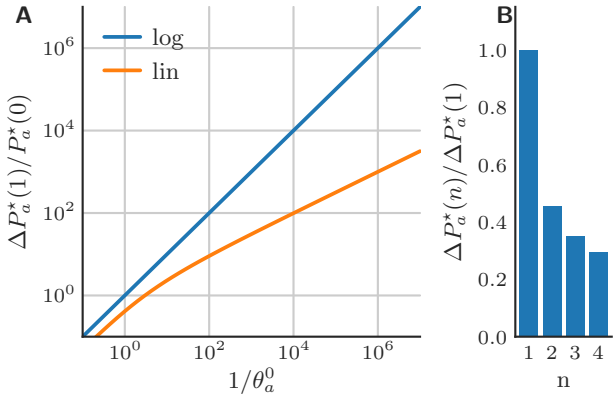

Figure 5.5: Optimal amount of memory production. (A) Change in protection $\Delta P_a^\star(1)$ upon first pathogen encounter relative to prior protection for linear and logarithmic cost functions as a function of pathogen variability. (B) Change in protection $\Delta P_a^\star(n)$ relative to change for first encounter $\Delta P_a^\star(1)$ for linear cost function in the quasistatic limit.

How does the optimal level of protection against a particular pathogen change upon encounters? We analyze this question by showing how the differential change in protection against pathogen a, $\Delta P^\star_a(n)$, upon the n^{th} encounter depends on the variability of the pathogen distribution and the scaling of infection cost with effective time. We find that the more variable pathogen distributions are the more memory should be retained of infections (Fig. 5.5). To gain intuition consider the simple (biologically unrealistic case) that cost increases logarithmically with time to recognition, then the resource allocation is proportional, $\mathbf{P}^\star = \hat{\mathbf{Q}}$ [138]. The change of protection relative to the background level $P^\star_a(0) = \theta^0_a/|\boldsymbol{\theta}^0|$ is

$$\Delta P^\star_a(1)/P^\star_a(0) = 1/\theta^0_a, \tag{5.14}$$

i.e. the memory production is inversely proportional to θ^0_a. The smaller the elements of $\boldsymbol{\theta}^0$ are the quicker the repertoire adjusts.

We can use this relationship to estimate the parameter regime in which the adaptive immune system functions. To explain a large memory production upon infections, we can read off from Fig. 5.5 that the relevant antigens need to be highly sparse in the overall antigen space: assuming a 100-fold increase in protection [42] and a linear cost function leads to $\theta^0 \sim 10^{-4}$.

The theory also makes predictions about how much memory to produce upon secondary infection. There are two effects. First, the n^{th} reinfection with the same pathogen starts from a higher baseline $\theta^0 + n - 1$, which for sublinear G leads to smaller relative updates for reinfections (Fig. 5.5B). Second, the increase of the normalizing factor $|\boldsymbol{\theta}(t)|$ with time, which after the law of large numbers grows as $|\boldsymbol{\theta}^0| + \lambda t$, becomes noticeable after long times. A unit increase of θ_a (Eq. 5.11) leads to a smaller increase in expected pathogen frequency (Eq. 5.13).

We still miss an estimate of K. There is a large number of potential antigens and receptors, but due to cross-reactivity not all of them are independent. We can think of an abstract recognition space as a high-dimensional space, where receptors are points surrounded by their recognition (cross-reactivity) balls [77]. Antigens are then scattered in this space and will be recognized by a receptor if they fall within its balls' radius. We choose an effective number of categories K that lumps together those antigens that are recognized by the same receptors. Given the probability p that a randomly chosen receptor recognizes a given antigen, we can estimate an effective dimension $K \sim 1/p$.

In the literature estimates of the precursor frequencies for common viruses of the order $\sim 10^{-5}$ have been reported [50]. Given that this includes the response to several antigens p should be somewhat smaller so we assume $p \sim 10^{-6}$. The optimal repertoire dynamics with these parameters can reproduce the observed rapid drop in infection susceptibility at young age assuming a biologically reasonable encounter rate λ (Fig. 5.6). Interestingly the optimal dynamics then also roughly reproduces the time-course of the fraction of cells that are memory of previous infections observed experimentally (see Fig. 2.1). In particular the saturation of memory production upon reinfections leads to a slowing down of the memory compartment growth with age.

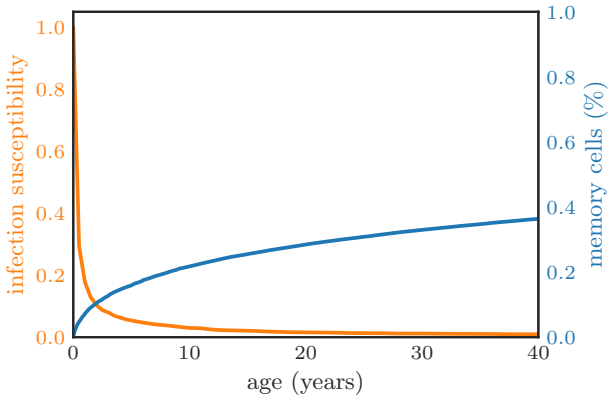

Figure 5.6: Memory cell frequency and infection susceptibility in young age with a repertoire following optimal Bayesian dynamics in a static environment. Parameters where chosen, such that the percentage of the overall repertoire taken up by memory cells and the infection susceptibility follow similar time courses as observed in humans (see Fig. 2.1). Pathogen susceptibility is assumed to be proportional to expected time to first recognition given the repertoire. Parameter: encounter rate $\lambda = 2000$/year, antigen space dimensionality $K = 10^6$, antigen variability $\theta^0 = 10^{-4}$.

5.3.3 *Dynamic case*

Let us now consider the more general case where the prediction phase does matter. As it turns out this problem is still analytically tractable as the eigenfunctions of the backward operator corresponding to Eq. 5.2 are known [172, 108, 173, 174, 175]. For simplicity we consider the case $K = 2$ here. This is without loss of generality as the $K-1$ dimensional dynamics is nested within the K dimensional dynamics: the diffusion between K pathogens can be broken down to the diffusion between one pathogen and the other $K-1$ pathogens as immigration is independent of the population composition.

Specializing Eq. 5.2 to $K = 2$ gives

$$\tau \partial_t \rho = \frac{1}{2}(\beta \partial_q(\rho q) - \alpha \partial_q(\rho(1-q))) + \frac{1}{2}\partial_q^2(\rho q(1-q)), \tag{5.15}$$

where q is the probability of the chosen pathogen. (To map to the symmetric K pathogen case consider $\alpha = \theta^0$, $\beta = (K-1)\theta^0$.) The steady-state distribution specializes from the Dirichlet distribution to the Beta distribution

$$\rho_s(q) = q^{\alpha-1}(1-q)^{\beta-1}/Z, \tag{5.16}$$

where $Z = B(\alpha, \beta) = \Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha+\beta)$ is the Beta function.

Using the known eigenfunction of the backward operator corresponding to Eq. 5.15 [174] its general solution can be written as (see App. B.2)

$$\rho(q,t) = \rho_s(q)\left(1 + \sum_{n=1}^{\infty} d_n(t) F_n(q)\right), \tag{5.17}$$

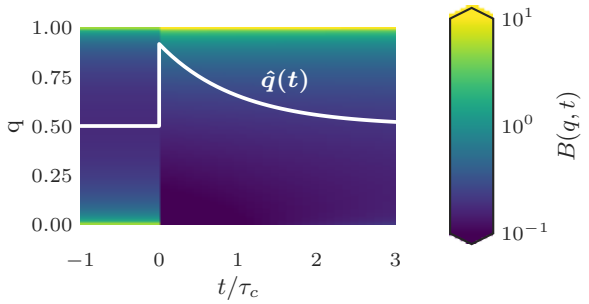

Figure 5.7: Dynamics of the belief distribution in a variable environment. After the pathogen encounter at time 0 it is more likely that the pathogen is currently at high frequency but as time passes the belief distribution relaxes back to its steady state value. The heatmap shows the belief distribution as a function of time, the black line its mean. Parameters: scaled immigration rates $\theta^0 = 0.1$.

in terms of coefficients d_n, with eigenfunctions $F_n(q)$, which are modified Jacobi polynomials (defined in App. B.2). The coefficients follow a simple exponential decay

$$\frac{dd_n(t)}{dt} = -\lambda_n, \tag{5.18}$$

with eigenvalues λ_n of which we give the first two,

$$\lambda_1 = (\alpha+\beta)/2\tau, \tag{5.19}$$
$$\lambda_2 = (1+\alpha+\beta)/\tau. \tag{5.20}$$

The Bayesian procedures (Eq. 5.8,5.7) are analytically tractable in terms of the coefficients d_n (for a derivation see App. B.2). In short, the update uses that the F_n are related by a recursion formula (Eq. B.14), which leads to

$$d_n(t^+) = \chi_n^- d_{n-1}(t^-) + \chi_n^0 d_n(t^-) + \chi_n^+ d_{n+1}(t^-), \tag{5.21}$$

where $\chi_n^-, \chi_n^0, \chi_n^+$ are constants (defined in App. B.2) that depend on which pathogen was encountered. The prediction phase following the i^{th} pathogen encounters changes the coefficients according to Eq. 5.18 starting from the initial condition $d_n(t_i^+)$. In combination with the formula for the expected frequency (App. B.2) this yields an analytical procedure for how to optimally estimate the pathogen frequency,

$$\hat{q}(t) = \alpha/(\alpha+\beta)\,(1 + d_1(t)\,\beta/(\alpha+\beta+1))\,. \tag{5.22}$$

Eq. 5.22 shows that the mean belief decays back to steady-state with simple exponential decay with time constant $\tau_c = 1/\lambda_1 = 2\tau/(\alpha+\beta)$. This leads to a similar relaxation of the optimal repertoire with the same characteristic time scale.

For illustration we apply the formalism to an example with $K = 2$ and thus $\alpha = \beta =: \theta^0$ (Fig. 5.7). Before any encounters the belief distribution is equal

to the steady-state distribution and thus in the eigendecomposition Eq. 5.17 none of the higher terms are present, $d_n = 0$ for $n \geqslant 1$. The first encounter with antigen by Eq. 5.21 leads to a non-zero value for the first eigenfunction, $d_1^+ = \frac{1}{\theta}$. In the following the estimated frequency decays back to steady-state according to $\langle q \rangle(t) = \left(1 + e^{-t/\tau_c}/(2\theta + 1)\right)/2$ as shown by the black line in Fig. 5.7. Following Eq. 5.21 the second encounter with a pathogen leads to belief distribution that also has a non-zero weight for the second eigenfunction, $d_2^+ \neq 0$, as long as d_1 has not completely decayed back to zero. This continues with every encounter, which leads to increasingly higher eigenfunctions arising in the decomposition of the belief distribution.

Let us now consider which mechanisms could allow the adaptive immune system to follow the optimal dynamics in varying environments. Concretely, we ask whether the information encoded in the repertoire composition is sufficient to keep the relevant memory of past infections. We consider the case in which the optimal repertoire $\mathbf{P}^\star$ uniquely determines $\hat{\mathbf{Q}}$. The dynamics of the mean belief generally depends on the full belief distribution and not just on the mean itself. As the repertoire encodes only the mean, we investigate when the equations for the dynamics of the mean can be closed. For the specific pathogen dynamics we consider in this paper the dynamics of the mean during the prediction step is in fact independent of higher-order moments (Eq. 5.22). The update step, however is generally not fully determined by the mean prior belief (Eq. 5.21): the mean frequency depends on the first coefficient d_1 (Eq. 5.22) only, but the recursion for the first coefficient depends also on the second coefficient d_2. In the quasi-static environments discussed earlier d_2 has a fixed relationship to d_1, which explains why the equations were closed. Another limit in which memory production only depends on the mean prior belief is, when pathogen sampling rates are much smaller than the decay constant of the second eigenfunction λ_2 ($\lambda\tau \ll 1 + \alpha + \beta$), so that the second coefficient is approximately zero, $d_2 \approx 0$. Outside of these two limits the immune system can either employ additional memory mechanisms or use an only approximately optimal dynamics.

5.4 DISCUSSION

In this chapter we have developed a Bayesian perspective on immunological memory. We have shown that the optimal repertoire dynamics arising from such a view can be well-approximated by biologically realistic population dynamics of lymphocyte clones. The optimal repertoire dynamics has a number of interesting features with links to observed properties of adaptive immunity. We have shown that the optimal dynamics strongly limits proliferation upon recurrent infections. Mechanisms to achieve such a limitation might explain the observed difficulty to generate big memory responses for booster vaccinations. We also demonstrated the need to forget previous infections over time while going back to a broad steady-state. One way to implement such a dynamics is to have an independently limited pools of cells for memory, which might explain the existence of distinct memory and naive cells in the adaptive immune system.

Our results show that the benefit of immunological memory depends crucially on the variability of the antigen distribution. For more variable distributions perfect adaptation provides more benefit. Furthermore for highly variable distributions memory already provides a benefit in the undersampled limit where the number of encountered pathogens λt is small compared to the dimensionality K of the distribution. On a theoretical level, given the approximate sparsity of the Dirichlet distribution (Fig. 5.3C) for small θ^0 there might be an interesting connection with the literature on compressed sensing [176] that merits further exploration. Biologically, our results imply that memory can provide little benefit against viruses such as HIV which show a broad distribution compared to viruses which are in the low mutation limit and behave as a quasispecies. The large increase in specific protection upon primary infection also hints at a highly sparse pathogen distribution. The relatively small number of important regions of antigenic space might also explain how the immune system can learn relevant features of the antigenic space quickly enough to explain the observed rapid drop in infection susceptibility of children in the first years of their life.

The environmental dynamics is simple enough to be tractable and allows for analytical insight, but clearly the real dynamics is more complicated. An interesting direction for further work is to consider more realistic dynamics. A particularly interesting direction is to explicitly incorporate the mutational dynamics of pathogens in antigenic space.

A similar analysis of the benefits of memory might be used to get a deeper understanding of the evolutionary purpose of memory in innate immunity, which has been a focus of much recent work [8].

Acknowledgements. The work was supported by grant ERCStG n. 306312.

Part III

ADAPTIVE STRATEGIES IN FLUCTUATING ENVIRONMENTS

In this part of the thesis we ask which adaptive strategies optimize long-term growth rate in fluctuating environments. The first chapter investigates with which immune strategy a population can best defend itself against changing pathogens. In the second more technical chapter we show graphically and analytically how the best evolutionary strategy in a fluctuating environment depends on the fitness landscape and environmental statistics.

Note that there are some changes in notations between the two chapters (notably $\pi_{env} \to p$, $\tau_{env} \to t_c$).

6

IMMUNE STRATEGIES AGAINST CHANGING PATHOGENS

This chapter was previously published in Ref. [139].

6.1 INTRODUCTION

Immune systems have evolved to protect organisms against large and unpredictable pathogenic environments. Yet immunity always comes at a cost (metabolic and maintenance costs, auto-immune disorders, etc. [177]), and this cost must be balanced by the benefits that protection confers [178, 179]. Faced with the problem of evolving a suitable defense, different organisms, from archaea to humans, have developed different strategies to identify and target pathogens, which have given rise to a diversity of mechanisms of immunity.

A large effort has been made to elucidate these mechanisms down to their molecular details in a variety of species [22, 180, 181, 182, 64, 7]. Beyond many differences, these studies have revealed many commonalities [6, 183], which hint at a possible general understanding of the trade-offs that shape their design [178, 177]. For instance, independently of the well-known adaptive immune systems of jawed vertebrates, jawless vertebrates (e.g. lampreys) have developed an alternative adaptive system that uses a distinct molecular family of receptors, but both systems function largely in the same way, relying on the generation of a large number of diverse receptors expressed by two types of lymphocytes (B or T-like cells). Likewise, the innate immune systems of invertebrates and vertebrates, share many similarities, relying on the selected expression of germline Toll-like receptors upon infection. Some of the features of vertebrate immunity are even shared with bacteria, who have developed their own targeted immunity based on the CRISPR/Cas system [184, 7], which itself bears strong resemblance with genome protection through interfering RNAs in eukaryotes [185].

Independently of how they evolved and their particular molecular implementation, we may classify these diverse mechanisms into a few broad modes of immunity: heritable but not adaptable within an individual's lifetime, as innate immune systems; heritable and adaptable within a lifetime but with the benefits of adaptation being non heritable, as adaptive immune systems; acquired from the environment and heritable, as the CRISPR/Cas system; and mixed strategies combining several of these elements. These broad distinctions call for general principles to characterize the conditions under which one or another mode of immunity may be expected to evolve [6, 183, 177]. The diversity and variability of threats from the pathogenic environment suggests that different modes of immunity may offer better protection depending on the patterns of occurrence and virulence of pathogens, or the effective population size of the protected population. Here we apply a general theoretical frame-

work for analyzing populations in a varying environment [186] to predict the emergence of the basic forms of observed immunity.

6.2 MODEL

Individuals reproduce in the presence of pathogens, which randomly appear, may persist for several generations and disappear before possibly reappearing at a latter time (Fig. 6.1A). In our framework, a given pathogen has a probability α to appear and a probability β to disappear from one generation to the next (Fig. 6.1B). The pathogenic dynamics is quantified both by the pathogen frequency $\pi_{\text{env}} = \alpha/(\alpha+\beta)$, which is the probability that it is present at any given generation, and by the characteristic timescale $\tau_{\text{env}} = -1/\ln(1-\alpha-\beta)$, which sets how fast pathogens appear and disappear. Although other parametrizations may be considered, this choice for τ_{env} preserves the symmetry between the presence and absence of the pathogen (Appendix C.1).

Pathogens reduce the fitness of the individuals in the population and the immune system is designed to mitigate this effect. An individual's fecundity, defined as its expected number of descendants in the next generation $\bar{\xi}$, depends on the pathogenic environment and its ability to protect itself against it. Each pathogen independently lowers the fecundity of an unprotected individual by a relatively large cost factor $c_{\text{infected}} > 0$ (Fig. 6.1D). This cost is reduced to a lower cost $c_{\text{defense}} < c_{\text{infected}}$ when the individual is protected by its immune system, however this protection comes at a basal cost $c_{\text{constitutive}} < c_{\text{defense}}$ of maintaining the immune defense in absence of the pathogen (Fig. 6.1E).

We explore the choices and tradeoffs underlying various modes of immunity along three axes: adaptability, heritability, and mode of acquisition. The first, adaptability axis concerns how much resources are invested in the protection for the return of an efficient response. This tradeoff imposes a relationship between c_{defense} and $c_{\text{constitutive}}$ (Fig. 6.1E): the more effective the defense (the lower c_{defense}), the higher maintenance cost (the higher $c_{\text{constitutive}}$). For example, having a large number of immune cells specialized against a specific pathogen allows for a quick and efficient response in case of invasion, but this enhanced protection comes at the cost of producing and maintaining these cells in the absence of the pathogen. This strategy is adopted, for example, by much of the innate immune systems of plants and animals [22]. On the contrary, the adaptive immune system keeps a very small specialized pool of lymphocytes for each potential antigen, and makes them proliferate only in case of infection [6]. Having the machinery of adaptive immunity comes at some upfront investment cost, but the huge diversity of adaptive immune repertoires then allows for a response against many pathogens at essentially no marginal constitutive cost. It is this marginal constitutive cost that determines when to use a mode of defense once an organism has the machinery (Appendix C.4). We assume that the cost of defense grows faster at small constitutive costs than at large ones (as reflected in the convexity in Fig. 6.1E). The second, heritability axis is defined by the probability q that the protection is not transmitted to the offspring (Fig. 6.1B). Finally, the third, acquisition axis specifies how individuals may ac-

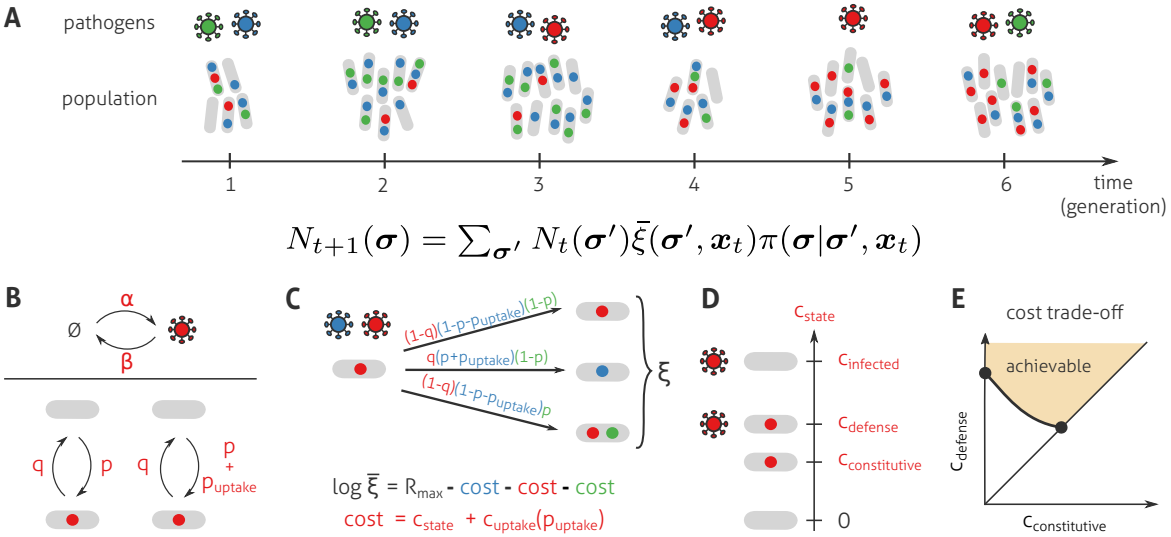

Figure 6.1: A model to explore the incidence of different modes of immunity on the long-term growth of populations. **(A)** A population of organisms, each possibly protected against no, one or several pathogens (no, one or several colored dots) evolves in presence of a pathogenic environment that varies from generation to generation. The mean number of individuals with protection σ at generation t, $N_t(\sigma)$, is given by a recursion equation involving the mean number of offspring $\bar{\xi}(\sigma', x_t)$ for individuals with protection σ' and the probability $\pi(\sigma|\sigma', x_t)$ that each of their offspring inherits a protection σ; both of these quantities may depend on the current pathogenic environment x_t. The long-term growth rate of the population is given by $(1/t)\ln N_t$ at large t, with $N_t = \sum_\sigma N_t(\sigma)$ the total population size. **(B)** Dynamics of appearance and disappearance of pathogens x_t and immune protection σ. A pathogen appears with rate α and disappears with rate β; these rates define the frequency $\pi_{env} = \alpha/(\alpha+\beta)$ and characteristic time $\tau_{env} = -1/\ln(1-\alpha-\beta)$ of the pathogen. Protection against a given pathogen is acquired spontaneously with rate p, and lost from one generation to the next with rate q. Additionally, the presence of the pathogen can increase the rate of acquisition of protection by p_{uptake}, as *e.g.* in the CRISPR/Cas system of prokaryotes. **(C)** The ξ offspring produced by an individual inherits the immune protections of their parent with rules specified in panel B. Each pathogen reduces the mean number of offspring $\bar{\xi}(\sigma, x_t)$ by a cost c_{state} that depends on whether the individual is in state 'infected', 'defense', or 'constitutive' relative to the pathogen, and by a cost $c_{uptake}(p_{uptake})$ that depends on the rate p_{uptake} at which protection is directly induced by the presence of the pathogen. **(D)** An unprotected organism pays a cost of infection $c_{infected}$ if the pathogen is encountered, which is reduced to $c_{defense}$ if it is protected. A protected organism must, however, pay a constitutive cost $c_{constitutive}$ even in the absence of the pathogen, while an unprotected organism pays no cost. **(E)** We assume a trade-off between the constitutive and defense costs: a more efficient defense (lower $c_{defense}$) requires more resources (higher $c_{constitutive}$).

quire the protection without inheriting it from their parent. This acquisition may occur randomly independently of the environment, with probability p, for instance by mutation or phenotypic switching, as is the case for antibiotic resistance in bacteria [187]; or it can be induced by the presence of the pathogen with probability p_{uptake}, as in CRISPR-Cas immunity (Fig. 6.1B) [7]. This mechanism comes at an extra cost $c_{uptake}(p_{uptake})$ due to maintenance and the risks of uptaking foreign genetic material (Fig. 6.1C), in addition to the state-dependent cost c_{state} (Fig. 6.1D). To account for the dangers associated with taking up foreign DNA we assume that c_{uptake} increases superlinearly with p_{uptake}.

6.3 RESULTS

Each choice of the parameters $c_{constitutive}$, q, p and p_{uptake} defines a specific immune strategy. This strategy is optimal if a population that adopts it outgrows in the long run any other population following a different strategy. Our goal is to characterize this optimal strategy, in particular its dependency on the two key properties of the pathogen, its frequency π_{env} and its characteristic time τ_{env}. We achieve this goal by maximizing the long-term growth rate of populations, defined by $(1/t)\ln N(t)$, where $N(t)$ is the total population size at generation t (Fig. 6.1A) [117]. Conveniently, since the fecundity is affected independently by the different pathogens, each pathogen contributes additively to the growth rate and can be studied separately (Fig. 6.1C and Materials and Methods).

Remarkably, we obtain qualitatively different optimal solutions for different values of π_{env}, τ_{env}, with sharp transitions between these strategies as one varies the parameters of the pathogen statistics, allowing us to define distinct immune regimes (Fig. 6.2A). The emergence of these very distinct regimes is not an assumption, but the result of the optimisation itself.

Fig. 6.2B describes these optimal strategies along the three axes of variation outlined earlier. Along the first axis of variation, adaptability, we find that frequent or persistent pathogens are best dealt with by constitutively expressed immunity ($c_{constitutive} = c_{defense}$), and rare and transient pathogens by investing minimally in the defense ($c_{constitutive} = 0$, in blue); between these two extremes, only a limited form of adaptation is required ($c_{constitutive} < c_{defense}$, in green). Along the second axis, heritability, we find that carrying the protection at all times ($q = 0$) is beneficial against fast pathogens but that losing the protection with probability $q > 0$ is more advantageous for slow ones. Finally, along the third axis, acquisition, we verify that there is no need to pay the price of informed acquisition ($p_{uptake} = 0$) whenever protection is systematically inherited ($q = 0$); when it is not the case ($q > 0$), we find that uptake is advantageous for sufficiently infrequent pathogens (yellow and orange regions) but that only for very infrequent pathogens does it becomes the exclusive mode of acquisition of protection ($p = 0$, $p_{uptake} > 0$, in yellow).

Each of these distinct regimes, or phases, is instantiated by natural immune systems (Table 6.1). For transient and rare pathogens (blue phase), the optimal strategy is to inherit a defense with minimal constitutive cost. This strategy is

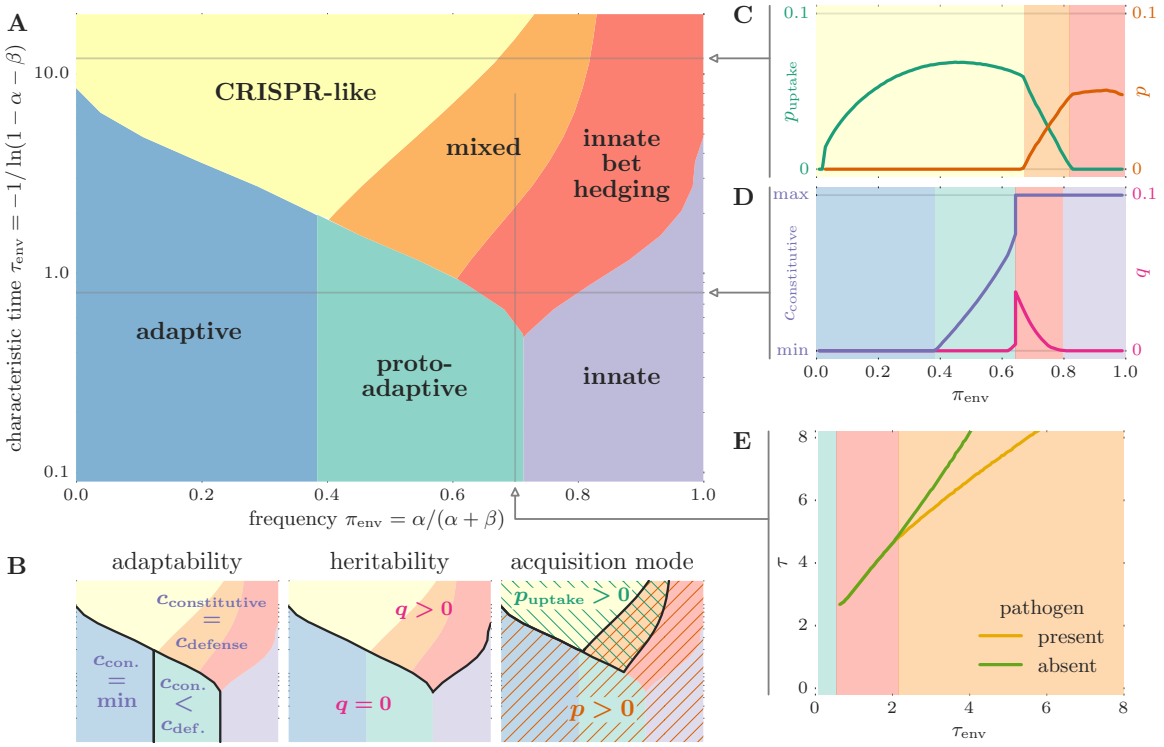

Figure 6.2: Optimal immune strategies as a function of the frequency and characteristic time of pathogens. **(A)** Distinct optimal immune strategies emerge for different statistics of appearance of the pathogens. Each phase is characterized by the value of parameters indicated in panel B and named after a known immune system that has similar characteristics (the name 'adaptive' is after the vertebrate immune system). **(B)** The different phases of immunity are defined by the values of parameters along three main axes: adaptability (constitutive cost $c_{constitutive}$), heritability $(1-q)$ and mode of acquisition (p and p_{uptake}). **(C)** and **(D)** Optimal parameters as a function of π_{env} for $\tau_{env} = 12$ (C) and $\tau_{env} = 0.8$ (D). For slowly varying environments (C), rare pathogens are best targeted by CRISPR-like uptake of protection, while frequent pathogens are best dealt with by spontaneous acquisition of protection, with a crossover in-between where both co-exist. For faster varying environments (D), the constitutive cost invested in the protection goes from negligible to maximal as the pathogen frequency increases. When it is maximal, the best strategy transitions from bet-hedging ($q > 0$) to a full protection of the population ($q = 0$). **(E)** The correlation times of protection in absence of the pathogen, $\tau = -1/\ln(1-p-q)$, and in its presence, $\tau = -1/\ln(1-p-p_{uptake}-q)$, are shown for $\pi_{env} = 0.7$ as a function of τ_{env}. Both increase with the correlation time of the pathogen. In this figure, an infinite population size is assumed and the following parameters are used: $c_{infection} = 3$; $c_{constitutive} = (1.8 - c_{defense})/(c_{defense} - 0.2)$; $c_{uptake}(p_{uptake}) = 0.1 \times p_{uptake} + p_{uptake}^2$ (see Fig. S2 for other choices).

characteristic of the adaptive immune system in vertebrates, where an effective immune response is mounted from a small number of precursor cells, the marginal cost of which is negligible [6]. For transient but frequent pathogens (purple phase), the optimal strategy consists instead in inheriting a maximally efficient protection that makes the individuals effectively insensitive to the presence of the pathogen at the expense, however, of a large constitutive cost. The recognition of pathogen-associated molecular patterns by pattern recognition receptors, as for instance the recognition of lipopolysaccharide by Toll-like receptors, is an example of such an innate strategy [22]. An intermediate phase (in green) separates these two extremes, where adaptation is present with a non-zero constitutive cost. This strategy, which we call proto-adaptive, is represented by certain specialized cells of the innate immune system, such as natural killer cells [188], whose abundance can vary as a function of experienced infections, effectively implementing an adaptive memory within a single generation.

For slow and frequent pathogens (red phase), protection is acquired with probability $p > 0$ and lost with probability $q > 0$ independently of the presence of the pathogen. This bet-hedging strategy is implemented in bacteria that can switch on or off the expression of phage receptors [64]. For slow but infrequent pathogens (yellow phase), a form of bet-hedging is again present, but this time with a non-zero probability to acquire protection only in presence of the pathogen. An example of such a Lamarckian strategy is the CRISPR-Cas immune system in bacteria [7]. Finally, a mixed phase (in orange) is also possible where protection is randomly acquired at a rate that is increased by the presence of the pathogen.

We can gain insight into the transitions between the different phases by considering three analytically solvable simplifications of the model, as detailed in Appendix C.3. In the first of these simplified models, we can calculate the transitions from a purely constitutive to a proto-adaptive to a purely adaptive strategy as the pathogen frequency π_{env} decreases. The second model highlights the transition from a bet-hedging to a deterministic protection, while the third one focuses on the transitions from a purely passive to a purely active acquisition of the protection, with a mixed phase in between.

It is instructive to examine how the parameters of immunity vary within the phases (Fig. 6.2C-D and C.1). As one may expect, the statistical properties of the protection tend to track the pathogen statistics [138]. The more frequent the pathogen, the more prevalent the protection in the population (Fig. 6.2C). Likewise, the characteristic time of the protection, τ, grows with that of the pathogen, τ_{env} (Fig. 6.2E).

The assumptions that we made allow us to treat each pathogen-protection pair independently of each other. However, there are a number of ways in which this assumption may be questioned. We discuss several in Appendix C.4 and we find that these generalizations do not qualitatively affect our conclusions. For example, infections could interact by inflicting more harm together than the sum of the each alone, as *e.g.* HIV in conjunction with other diseases. This case can be incorporated into our approach by considering a modified effective cost including the extra cost of co-infection. Another way in which

strategy	defining characteristics				biological examples
	perfect heritability	acquisition mode		adaptability	
	$q = 0$	$p > 0$	$p_{uptake} > 0$	$c_{defense}$	
innate	yes	yes	no	minimal	innate defense by recognition of pathogen-associated molecular patterns by pattern recognition receptors [22]
protoadaptive	yes	yes	no	intermediate	"trained" innate immunity [8], especially defense by natural killer cells [188]; "systemic acquired resistance" in plants [189]
adaptive	yes	yes	no	maximal	adaptive immune systems of jawed and jawless vertebrates [6]
innate bet hedging	no	yes	no	minimal	mutation of phage receptors by bacteria [64]
CRISPR-like	no	no	yes	minimal	CRISPR-Cas system in bacteria and archea [7]
mixed	no	yes	yes	minimal	concurrent use of CRISPR-Cas system and mutations of surface molecules by bacteria defending against phages [65]

Table 6.1: Optimal strategies found in the phase diagram (Fig. 2), their definition in terms of parameters of our framework, and biological examples.

pathogen-protection pairs could be correlated is through a non-additive cost of protection, if the marginal cost of protection increases or decreases with the number of protections. For example, if having protection against two pathogens is much more costly than twice the cost of having protection against just one, then the optimal strategy may be to hedge bets by keeping a subpopulation protected against one pathogen, and another subpopulation against the other. Lastly, cross-reactivity, the widespread ability of protections to recognize several pathogens, is another departure from independence, which can be partly overcome by grouping together pathogens recognized by a common protection.

6.4 DISCUSSION

The phase portrait of Fig. 2A rationalizes the salient differences between the immune systems of prokaryotes and vertebrates. Bacterial and archeal organisms evolve on timescales that are much closer to those of their pathogens than vertebrates. From the viewpoint of microbes, the pathogenic environment is relatively constant ($\tau_{env} > 1$), while for vertebrates a particular pathogenic strain is unlikely to survive a single generation ($\tau_{env} \ll 1$). Consistent with our results, vertebrates use fully heritable modes of immunity, and do not rely

on bet-hedging. To deal with infrequent and fast evolving pathogens such as viruses, they recourse to adaptive mechanisms by which they can upregulate their protection in case of an invasion. The three predicted strategies – adaptive, proto-adaptive, and innate – correspond to the known modes of immunity in vertebrates [22]. Prokaryotes, on the other hand, almost systematically use bet-hedging strategies. They recourse to both the CRISPR-Cas system of acquired immunity [7], and to innate immunity through *e.g.* restriction endonucleases [64], which correspond to the predicted Lamarckian and innate bet-hedging strategies of the diagram, respectively. These results are robust to changes of parameters, although increasing costs can make bet-hedging beneficial even at short characteristic times (Fig. C.2).

Bacteria and vertebrates also have very different population sizes, which influence their overall survival probability. To evaluate this impact, we ran stochastic simulations competing different strategies for increasing population sizes (see Sec. 6.5 and Fig. C.3). The phase diagram of Fig. 2A is recovered for population sizes as small as a thousand, while for smaller population sizes the boundaries between regimes are smeared. In small populations, adaptive strategies are generally favored over CRISPR-like strategies, and the amount of bet-hedging increases. In fact, for finite populations it is always beneficial to recourse to some degree of bet-hedging in order to react quickly to environmental changes and avoid extinction.

Here, we consider the case of a common environment experienced by all individuals. Having different parts of the population experience different micro-environments that are not persistent over generations does not change the population dynamics on evolutionary time scales (Materials and Methods). These micro-environments can result from differing infection probabilities for different subsets of the population, e.g. arising from spatial niches, or other non-pathogenic factors such as nutrient availability that influence the capability of individuals to cope with pathogens. If there are micro-environments that persist over many generations then our results hold in each micro-environment. An optimal strategy might then exploit the additional predictability stemming from knowing the statistical properties of the micro-environments and use the micro-environment diversity as a means of bet-hedging.

Our results also suggest that plants and some invertebrates, which also have long generation times compared to the variation time of pathogens, should be endowed with adaptive and proto-adaptive immune systems, in addition to innate protection mechanisms [181]. Consistent with this prediction, the innate branch of the plant immune system is able to increase protection in the entire plant following a local infection through "systemic acquired resistance" [189], providing the mechanistic basis of an inducible, proto-adaptive immune system. In addition, virus-derived small interfering RNAs, which accumulate during infections, are portrayed as likely candidate of adaptive immunity in plants and invertebrates – they are induced by the virus, and keep a memory of past infections [190, 191, 192]. Interestingly, small RNA-based immunity has been shown to be inheritable in *C. elegans* [193], an invertebrate with a short generation time of around 4 days, in agreement with our result that CRISPR-like immunity is desirable in this case.

By analyzing the long-term fate of populations under minimal assumptions concerning the rules governing adaptability, heritability and acquisition of immune protections, we have recovered the basic known modes of immunity. Remarkably our results hold even for a single pathogen. The key determinants of optimal immune strategies are found to be the statistical features of pathogen occurrence: its frequency and its characteristic timescale. As an implication, a diverse pathogenic environment, with varying statistics, will favor mixed solutions, consistently with the observation of multiple immune systems within a same organism – such as adaptive and innate immune systems in vertebrates, or CRISPR and innate defense in bacteria. Naturally, the molecular implementation of these general principles differs greatly even between organisms sharing the same type of immunity. Yet an evolutionary perspective that accounts for the costs and benefits of protection is enough to explain the most salient features of immunity. It will be interesting to extend our framework to account for other essential features of immunity, *e.g.* the acquisition of protection by horizontal transfer or the coevolutionary dynamics between pathogens and their hosts. In view of our analysis, it is already less surprising that complex forms of immunity such as the adaptive immune system have evolved separately in jawed and non-jawed vertebrates, with the same general features but different molecular encodings.

6.5 MATERIALS AND METHODS

6.5.1 *Population dynamics*

The pathogenic environment is described by an L-dimensional vector $\mathbf{x}$ (bold symbols refer to vectors), where $x_i = 1$ if pathogen i is present, 0 otherwise. Protection of an organism against these pathogens is also described by an L-dimensional vector $\boldsymbol{\sigma}$, where $\sigma_i = 1$ if the protection (antibody, TCR, CRISPR spacer) against pathogen i is present, and 0 otherwise.

We consider the dynamics of a population of organisms reproducing at discrete times t. At each generation, each individual produces a stochastic number ξ of offspring, whose distribution depends on the state $\boldsymbol{\sigma}$ of that individual, and the environment $\mathbf{x}_t$. We denote its mean by $\bar{\xi}(\boldsymbol{\sigma}, \mathbf{x}_t)$.

Let $N_t(\boldsymbol{\sigma}|\mathbf{x}_{t'<t+1})$ be the mean number of organisms in the population at time t with protection $\boldsymbol{\sigma}$, for a given environment history $(\mathbf{x}_{t'<t+1})$ [117]. The change in population composition from one generation to the next is governed by the reproductive success of individuals in each state $\boldsymbol{\sigma}$, modified by stochastic state switching from parents to offspring:

$$N_{t+1}(\boldsymbol{\sigma}|\mathbf{x}_{t'<t+1}) = \sum_{\boldsymbol{\sigma}'} N_t(\boldsymbol{\sigma}'|\mathbf{x}_{t'<t})\bar{\xi}(\boldsymbol{\sigma}', \mathbf{x}_t)\pi(\boldsymbol{\sigma}|\boldsymbol{\sigma}', \mathbf{x}_t), \tag{6.1}$$

where $\pi(\boldsymbol{\sigma}|\boldsymbol{\sigma}', \mathbf{x}_t)$ is the switching probability from protection state $\boldsymbol{\sigma}'$ to state $\boldsymbol{\sigma}$. Note that the protection state switching probability, which represents to what extent protection is inherited, acquired or lost, generally depends on the

state $\mathbf{x}_t$ of the environment. For ease of notation, we omit in the following the condition on the environment $(\cdot|\mathbf{x}_{t'<t})$ when referring to conditional means.

A similar recursion to Eq. 6.1 can be written for the fraction of the population in each state, $n_t(\boldsymbol{\sigma}) = N_t(\boldsymbol{\sigma})/N_t$, with $N_t = \sum_{\boldsymbol{\sigma}} N_t(\boldsymbol{\sigma})$ the total population size:

$$n_{t+1}(\boldsymbol{\sigma}) = \frac{1}{Z_t} \sum_{\boldsymbol{\sigma}'} n_t(\boldsymbol{\sigma}') \bar{\xi}(\boldsymbol{\sigma}', \mathbf{x}_t) \pi(\boldsymbol{\sigma}|\boldsymbol{\sigma}', \mathbf{x}_t), \tag{6.2}$$

where Z_t is a normalization constant enforcing $\sum_{\boldsymbol{\sigma}} n_t(\boldsymbol{\sigma}) = 1$. The population size is given by $N_t = N_0 \prod_{t'=0}^{t-1} Z_{t'}$, so that the long-term growth rate, $\Lambda = \lim_{T\to\infty} \frac{1}{T} \ln N_T$, is given by:

$$\Lambda = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T} \ln(Z_t). \tag{6.3}$$

The strategy with maximal long-term population growth rate outperforms in the long run any other strategy for almost every sequence of environments in populations of infinite size. This rate thus provides a measure of long-term fitness [117].

We assume that the mutation and inheritance probabilities of different pathogen-protection pairs are independent of each other, *i.e.* that $\pi(\boldsymbol{\sigma}|\boldsymbol{\sigma}', \mathbf{x}_t)$ factorizes over the pathogens,

$$\pi(\boldsymbol{\sigma}|\boldsymbol{\sigma}', \mathbf{x}_t) = \prod_i \pi_i(\sigma_i|\sigma_i', x_{i;t}). \tag{6.4}$$

The entries of $\pi_i(\sigma_i|\sigma_i', x_{i;t})$ are given by Fig. 1B: $\pi_i(1|0, x) = p + x p_{\text{uptake}}$ and $\pi_i(0|1, x) = q$.

In addition, the effects of different pathogen-protection pairs on the growth rate are taken to be additive (Fig. 1C), so that:

$$\begin{aligned} \ln \bar{\xi} = R_{\max} - \sum_{i=1}^{L} \big[& c_{\text{infection},i} (1-\sigma_i) x_i \\ & + c_{\text{constitutive},i} \sigma_i (1-x_i) + c_{\text{defense},i} \sigma_i x_i \\ & + c_{\text{uptake}}(p_{\text{uptake},i}) \big] \end{aligned} \tag{6.5}$$

where $R_{\max}$ is the growth rate in absence of any immune cost. With these assumptions, the distribution $n_t(\boldsymbol{\sigma})$ also factorizes over i:

$$n_t(\boldsymbol{\sigma}) = \prod_{i=1}^{L} (r_i^t \sigma_i + (1-r_i^t)(1-\sigma_i)), \tag{6.6}$$

where r_i^t is the fraction of the population having protection i at time t. Plugging this Ansatz into Eq. 6.2 with Eqs. 6.4 and 6.5 yields the following recursion for r_i^t:

$$\begin{aligned} r_i^{t+1} = [&(1-r_i^t)e^{-c_{\text{infection},i}x_i^t}\left(p_i + p_{\text{uptake},i}x_i^t\right) \\ &+ r_i^t e^{-c_{\text{defense},i}x_i^t - c_{\text{constitutive},i}(1-x_i^t)}(1-q_i)] \\ &/\ [\left(1-r_i^t\right)e^{-c_{\text{infection},i}x_i^t} \\ &+ r_i^t e^{-c_{\text{defense},i}x_i^t - c_{\text{constitutive},i}(1-x_i^t)}]. \end{aligned} \tag{6.7}$$

The recursion depends on the sequence of x_i^t, which is a stochastic binary process switching from 0 to 1 with probability α, and from 1 to 0 with probability β (Fig. 1B). Note that the sequence x_i^t is the same for the whole population (a quenched variable in the statistical mechanics sense). We have $Z_t = e^{R_{\max}} \prod_{i=1}^{L} z_i^t$, with

$$\begin{aligned} z_i^t = e^{-c_{\text{uptake}}(p_{\text{uptake},i})}[&\left(1-r_i^t\right)e^{-c_{\text{infection},i}x_i^t} \\ &+ r_i^t e^{-c_{\text{defense},i}x_i^t - c_{\text{constitutive},i}(1-x_i^t)}] \end{aligned} \tag{6.8}$$

From Eq. 6.3, it then follows that

$$\Lambda = R_{\max} + \sum_{i=1}^{L}\left(\lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\ln z_i^t\right). \tag{6.9}$$

The long-term growth rate is a sum of independent terms for each pathogen-protection pair, which allows us to treat the problem of maximizing long-term growth rate one pathogen at a time.

6.5.2 *Numerical solution*

The cost function of the optimization, Λ, can be approximated by solving the recursion equation describing the relative frequency of organisms with different protection states in the population for a large enough number of generations (we used at least 10^6 generations). Our goal is to optimize Λ over the four parameters $p, q, p_{\text{uptake}}, c_{\text{constitutive}}$ (or over the subset of free parameters for a given strategy) constrained to their domain of definition. For numerical purposes, all four parameters are first mapped onto the unit interval $[0, 1]$. The noise in the evaluation of Λ arising from its approximation from finite time data makes the optimization challenging. Since the process is ergodic, averaging over very long periods is equivalent to repeating the process multiple times. The noise can therefore be reduced by both prolonged simulation or repeated sampling at the expense of a higher computational cost per function evaluation. To find the global optimum of this noisy function under the bound constraints on the parameters we use a two-phase algorithm. In the first phase the DIRECT algorithm [194] provides us with a rough, but global optimization for which we use a relatively low quality approximation. The results of this first phase are then refined by a pattern search algorithm with an adaptive

sampling of the function (described in detail in Appendix C.2) using the parameters $\Delta_{tol} = 0.0005$ (for Fig. S1 $\Delta_{tol} = 0.005$), $\alpha = 0.005$.

To obtain a phase diagram such as the one shown in Fig. 6.2A we first performed a global optimization over all four parameter values as described above, for every environment condition (π_{env}, τ_{env}) (Fig. S1). Based on the optimal parameters found in this step, we defined the features of the emerging phases (Table I). All phases are defined by a subset of the variables lying at a constraint boundary. In order to calculate precise phase boundaries, we find the frequency of pathogens π_{env} at a given characteristic time τ_{env} for which the difference in long-term growth rates between a given pair of strategies vanishes. To obtain the root of the difference function, we use a bisection algorithm. To decrease noise, the difference is calculated across pairs of simulations using the same sequence of pathogens $\{x_t\}$ and the function is sampled adaptively to ascertain statistical significance. The bisection algorithm is run up to a tolerance of 0.025 in π_{env} and then the precise position of the root is interpolated assuming linearity of the difference function within the interval. To prevent e.g. the mixed strategy to reduce to a CRISPR-like strategy, we impose that the parameters which are not set to a fixed value in a particular strategy are not closer than a tolerance 0.005 (0.0005 for q) of the boundary.

6.5.3 *Simulations with finite populations sizes.*

To study the influence of the effects of finite population size we perform direct agent-based simulations of a population of adapting individuals with strategies evolving on a slow timescale. The population has a finite size N that remains fixed over the course of the simulation. At every generation the parents of the N individuals are drawn from the individuals making up the previous generation with probabilities proportional to the mean number of offspring $\bar{\xi}$ of these individuals. The offspring's state $\boldsymbol{\sigma}$ is determined from the state of its parent $\boldsymbol{\sigma}'$ according to the switching rates $\pi(\boldsymbol{\sigma}|\boldsymbol{\sigma}', x_t)$ defined previously. Along with the state $\boldsymbol{\sigma}$, the switching rates themselves, $\pi(\boldsymbol{\sigma}|\boldsymbol{\sigma}', \mathbf{x})$, as well as the degree of adaptability, $c_{constitutive}$ – in other words, the parameters defining the immune strategy – are also transmitted to the offspring. They also change from parent to offspring, although at a much slower rate than the state to preserve a clear separation of timescales between short-term and long-term adaptations. In this setup, selection acts on the strategies. After an equilibration phase, we collect statistics on the strategies adopted by individuals in the population. To get rid of the effect of deleterious mutations that do not eventually fix in the populations the mutation rate and size were scaled down exponentially with time. As population size is finite deleterious mutations can fix in the population, which means that even in the limit of zero mutation rate there remains a spread in the distribution of strategies. Hence we do not only represent the median as a measure of the central tendency of a parameter, but also the interquartile range as a measure of its spread.

Acknowledgements. The work was supported by grant ERCStG n. 306312.

7

TRANSITIONS IN OPTIMAL ADAPTIVE STRATEGIES

This chapter was previously published in Ref. [140].

7.1 INTRODUCTION

Nothing is as constant as change. This age-old adage applies to biological populations, which may respond by evolving mechanisms to mitigate the consequences of environmental fluctuations [109, 110, 111, 112, 113]. This adaptation can be implemented at different levels. At an individual level, the simplest strategy consists in adopting a generalist phenotype that does reasonably well across environments. At a population-level, another strategy is to constantly generate a phenotypically diverse mixture of individuals, each specialized to a different environmental condition. Which strategy provides the largest selective advantage in the long run depends on the nature of environmental fluctuations and on the fitness costs and trade-offs limiting the range of accessible phenotypes. For instance, although tracking the environment to adopt a phenotype specialized to each current condition may seem optimal, this strategy is often precluded by the costs of constantly monitoring environmental changes and of frequently switching between phenotypes.

Which strategies to deal with environmental fluctuations may be selected is a long-standing question in evolutionary biology. Interest in this question has recently been rekindled by novel laboratory experiments with populations growing in controlled fluctuating environments [195, 196, 197], new theoretical developments providing links to ideas from information theory and stochastic thermodynamics [115, 116, 114], and its relevance to understanding non-genetic modes of inheritance [198, 186] and how biological populations might respond to climate change [112, 199].

Here, we study this question in a model of population growth in a randomly fluctuating environment. The model considers a large population of organisms characterized by their phenotype and replicating at discrete generations. An optimal adaptive strategy is defined by the choice of phenotypes and switching rates between them that ensures the largest long-term population growth rate. We analyze how this optimal strategy depends on the environmental statistics and the replication rates. The analysis reveals transitions between qualitatively different strategies: non-switching or single-phenotype strategies, where all of the population is of the same phenotype; and switching or bet-hedging strategies, where the population diversifies. Further transitions arise between strategies where the population adopts a phenotype specialized in a single environment, and strategies relying on a generalist phenotype.

Our work extends the growing literature investigating transitions between optimal adaptive strategies [200, 201, 186, 202, 203] and generalizes some of our previous results on the adaptation of immune strategies to pathogen statis-

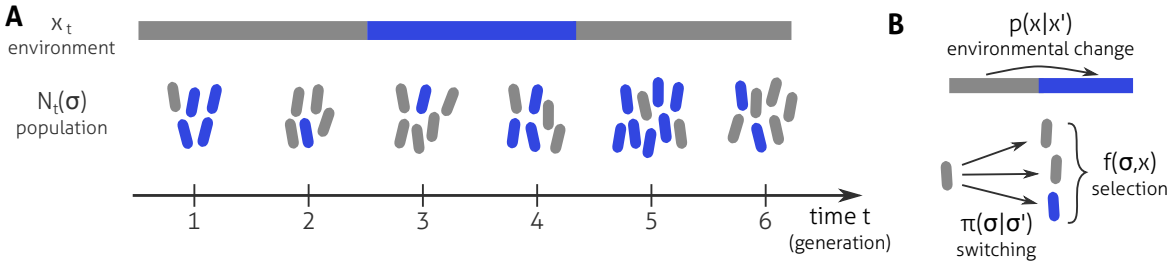

Figure 7.1: Model of population growth in a fluctuating environment. (A) A population composed of individuals of different phenotypes σ grows in a changing environment x_t. Between each discrete generation, the phenotype of each individual may switch. (B) The environment follows a stochastic dynamics described by a Markov chain with transition rates $p(x|x')$. The population composition changes between generations due to the effects of selection (an individual with phenotype σ in environment x produces in average $f(\sigma, x)$ offspring) and phenotype switching (an individual with phenotype σ' has probability $\pi(\sigma|\sigma')$ to have an offspring with phenotype σ).

tics [139]. In particular, we derive exact expressions for the transitions between different modes of immunity in memoryless environments when the strategy includes an adjustable investment into immunity. We also calculate analytically the transitions between switching and non-switching strategies between two phenotypes in temporally correlated environments. After briefly introducing the mathematical framework (Sec. 7.2), we present a graphical method for studying transitions in optimal adaptive strategies in temporally uncorrelated environments (Sec. 7.3), and apply it to the case of an immune system with adjustable investment (Sec. 7.4). We then turn to the case of temporally correlated environments and provide new analytical and numerical results on transitions in this more general setting (Sec. 7.5).

7.2 POPULATION GROWTH IN FLUCTUATING ENVIRONMENTS

We are interested in describing the evolution of a possibly phenotypically heterogeneous biological population (of cells, organisms, etc.) in a fluctuating environment. We describe the population at generation t by the number $N_t(\sigma)$ of individuals with a given phenotype σ. Phenotypes differ by their replication rate $f(\sigma, x)$, which give the mean number of offspring produced by an individual of phenotype σ in environmental condition x (see Fig. 7.1). The environment is described as a discrete Markov chain with a transition matrix $p(x|x')$, which we assume to be stationary and ergodic. The population changes under the influence of the selective pressures generated by the differences in replication rates between phenotypes, and through phenotype switches described by a transition matrix $\pi(\sigma|\sigma')$. In the limit of infinitely large population size, the population composition follows the recursion [117]

$$N_{t+1}(\sigma) = f(\sigma, x_t) \sum_{\sigma'} \pi(\sigma|\sigma') N_t(\sigma'). \tag{7.1}$$

An alternative formulation in which multiplication precedes mutation can be defined by mapping $N_t(\sigma) \to N_t(\sigma)/f(\sigma, x_{t-1})$ [139]. The two formalisms are equivalent as long as the switching rates do not depend on the environmental state. Eq. 7.1 can also be written in a compact matrix notation as

$$\mathbf{N}_{t+1} = \mathbf{A}^{(x_t)}\mathbf{N}_t, \quad \text{with} \quad A^{(x_t)}_{\sigma,\sigma'} = f(\sigma, x_t)\pi(\sigma|\sigma'). \tag{7.2}$$

Here and in the following, we write vectors and matrices in bold notation.

The different modalities by which populations might cope with fluctuating environmental conditions correspond to different properties of the switching matrix $\pi(\sigma|\sigma')$. For *non-switching strategies*, the whole population has the same phenotype $\tilde{\sigma}$ and the switching matrix consists in a row of ones, $\pi(\sigma|\sigma') = 1$ if $\sigma = \tilde{\sigma}$ and 0 otherwise. If the chosen phenotype is a better all-rounder doing intermediately well across environments, this corresponds to an individual-level *generalist strategy*. For *switching strategies*, we may distinguish those with and without memory. In a switching strategy without memory, the probability of switching to a phenotype does not depend on the parental phenotype, $\pi(\sigma|\sigma') = \pi(\sigma)$. Such strategies implement population-level bet-hedging, i.e., diversification of the population into phenotypes that may each be specialized to one of the environmental conditions to come. Switching with memory, where $\pi(\sigma|\sigma')$ does depend on σ', provides the basic ingredients, variation and heritability, to enable adaptive tracking of the environment through Darwinian evolution. In the limit where switching is very rare, $\pi(\sigma|\sigma') \ll \pi(\sigma'|\sigma')$ for $\sigma \neq \sigma'$, the phenotypic dynamics is equivalent to the strong-selection weak-mutation limit of population genetics [186]. The model thus integrates in a common mathematical framework a range of different modes of response to environmental variations.

Over long evolutionary time scales, selection might act on the adaptive mechanisms to adjust them to the statistics of environmental fluctuations. Explicit models of the evolution of the switching rates $\pi(\sigma|\sigma')$ show that variation in switching rates can indeed be selected upon [201, 204, 186]. Transgenerational feedback reinforcing the production of successfull phenotypes provides an alternative mechanism to learn a good strategy [205]. Which adaptive strategy do we expect to evolve in the long run? Here, we focus on the optimal strategy representing the optimal possible end-product of this evolution. In our model, the optimal switching rates maximize long-term growth rate, defined as

$$\Lambda = \lim_{T\to\infty} \frac{1}{T} \ln N_T/N_0, \tag{7.3}$$

where $N_T = \sum_\sigma N_T(\sigma)$ is the total population size. To understand why this is the relevant measure of evolutionary success in the long run, consider a population with two subpopulations following different strategies. Then in the long run the population following the strategy with highest long-term growth rate almost surely outnumbers the one following the other strategy for almost every sequence of environments [206]. The question of which adaptive strategy

$\pi^*(\sigma|\sigma')$ has the largest selective advantage is thus recast as the problem of maximizing the long-term growth rate over possible strategies:

$$\pi^*(\sigma|\sigma') = \underset{\pi(\sigma|\sigma')}{\operatorname{argmax}} \Lambda, \tag{7.4}$$

for given replication rates $f(\sigma, x)$ and given environmental dynamics $p(x|x')$. This is the problem that we address in this paper.

7.3 WHEN AND HOW TO BE A GENERALIST IN UNCORRELATED ENVIRONMENTS

7.3.1 *Extended fitness set and Pareto optimality*

The simplest environmental fluctuations to consider are memoryless fluctuations, where the state of the environment is independent of its state in the previous generation, $p(x|x') = p(x)$. In this case, no gain can be expected from keeping a memory of past phenotypic states, and the optimal adaptive strategy is also memoryless, $\pi(\sigma|\sigma') = \pi(\sigma)$. Since the population composition is constant over generations, the number of offspring depends only on the state of the environment and Eq. 7.1 reduces to a recursion for the total population size $N_t = \sum_\sigma N_t(\sigma)$:

$$N_{t+1} = N_t f(x_t), \tag{7.5}$$

where

$$f(x) = \sum_\sigma f(\sigma, x)\pi(\sigma) \tag{7.6}$$

is the average population fitness. Graphically, it is convenient to represent each possible phenotype σ as a point in the space of environmental conditions x (where each environment x defines a dimension), with coordinates given by the replication rates $f(\sigma, x)$ (orange dots in Fig. 7.2A). The set $D_f = \{\sum_\sigma f(\sigma, x)\pi(\sigma) | \sum_\sigma \pi(\sigma) = 1, \pi(\sigma) \geqslant 0\}$ of achievable $f(x)$ when switching rates $\pi(\sigma)$ are varied then corresponds to the convex hull of these points (orange area in Fig. 7.2B). In the ecological literature, this set of achievable strategies is known as the extended fitness set and was introduced by Levins [109].

The recursion for the total population size (7.5) is solved by $N_T = N_0 \prod_t f(x_t)$. Taking logarithms, we have $\ln N_T/N_0 = \sum_{t=1}^{T} \ln f(x_t)$ and we can apply the law of large numbers to write the long-term growth rate (7.3) as a weighted average of log-fitnesses:

$$\Lambda = \sum_x p(x) \ln f(x), \tag{7.7}$$

with weights given by the frequency of each environment.

Finding the optimal strategy $\pi^*(\sigma)$ that maximizes $\Lambda = \sum_x p(x) \ln \sum_\sigma f(\sigma, x)\pi(\sigma)$ over the domain allowed by the rules of probabilities is a convex optimization

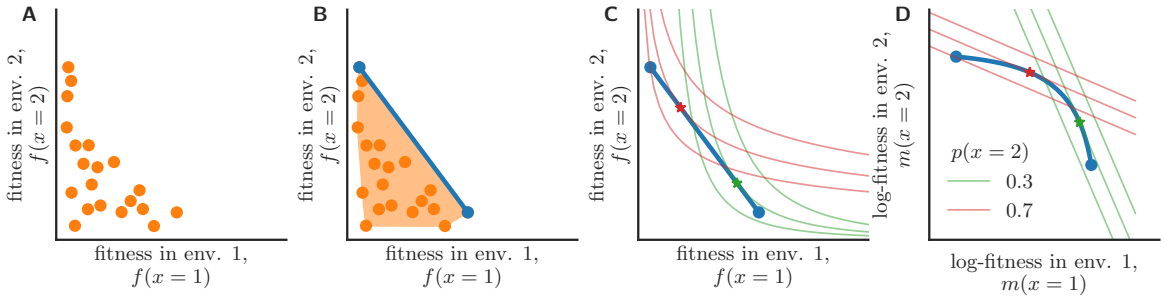

Figure 7.2: Illustration of the steps of a graphical method of finding the best adaptation strategy in uncorrelated environments. (A) Fitness values of phenotypes across environments (orange dots). (B) Fitness values achievable by switching strategies (orange area) are those inside the convex hull of the fitness values of the different phenotypes. A necessary condition for optimality is to lie on the Pareto frontier (blue line). (C, D) The optimal strategy has the fitnesses (red/green star) at which the isolines of the long-term growth rate for given environmental frequencies (red lines for $p(2) = 0.7$, green lines for $p(2) = 0.3$) are tangential to the Pareto frontier. (C) In fitness space the isolines are curved. (D) To determine the optimal strategy it is more convenient to work in log-fitness space, where the isolines are straight lines.

problem whose solution is well known [118, 109, 207, 206, 117, 200]. It is useful to rephrase the problem as the optimization of $\Lambda = \sum_x p(x) \ln f(x)$ over the fitnesses $\mathbf{f}$ constrained to belong to the extended fitness set D_f introduced above. One can go further and equivalently optimize $\Lambda = \sum_x p(x) m(x)$ over the log-fitnesses $m(x) = \ln f(x)$ contrained to belong to $\ln(D_f)$. Going from $\boldsymbol{\pi}$ to $\mathbf{f}$ to $\mathbf{m}$ simplifies the expression of the objective function Λ but makes the domain of optimization more complex.

Eq. (7.7) shows that the long-term growth rate is an increasing function of each environment fitness $f(x)$. Increasing fitness in one environment is always desirable if this can be done without impairing fitness in any other environment. Thus, any optimal solution must lie on the set of fitnesses $\mathbf{f}$ for which no improvement can be made in one environment without impairing performance in another, called the Pareto frontier. Usually, no phenotype provides the best fitness for all environments due to trade-offs between performance under different conditions. Thus the Pareto frontier is generally not a single point but a line when the environment alternates between two conditions (blue line in Fig. 7.2B), and a hyper-surface of dimension $n-1$ when the environment alternates between n conditions. To find the overall optimum along the Pareto front requires to consider the explicit way in which performances for different objectives combine into a scalar measure, which is determined in our case by the frequency of the different environments (7.7).

7.3.2 *Graphical method for finding the optimal strategy*

The various views of the optimization problem discussed in the previous subsection imply a graphical method to determine the optimal strategy. For sim-

plicity, we illustrate it by considering switching between only two environments (Fig. 7.2). Starting from the graphical representation of the Pareto front for the set of achievable fitnesses (Fig. 7.2B), we need to find the point of this frontier with the highest growth rate: this is done graphically by representing the growth rate isolines $\Lambda[f(1), f(2)] = K$ (red and green lines in Fig. 7.2C where the two colors corresponds to different environmental statistics) given by (7.7):

$$f(2) = \frac{e^{K/p(2)}}{f(1)^{p(1)/p(2)}}. \tag{7.8}$$

By plotting the isolines for different K we can find the isoline for the largest K that still intersects with the Pareto frontier, called supporting line. The intersection point defines the optimal adaptive strategy the population should adopt (red and green stars in Fig. 7.2C). This construction was first proposed by Levins [109].

Here, we propose to go one step further and work in log-fitness space to circumvent the difficulty of handling curved isolines. In log-fitness space, the isolines are linear and normal to the vector $\mathbf{p}$:

$$p(1)m(1) + p(2)m(2) = K. \tag{7.9}$$

If the Pareto front has a tangent of slope $-p(1)/p(2)$, the tangent point thus defines the optimal strategy for the environment $\mathbf{p}$ (Fig. 7.2D). More generally, the supporting isoline corresponding to the optimal growth rate shares at least one point with the Pareto frontier but is otherwise entirely above that frontier.

The graphical method generalizes to d environments by studying the extended fitness set in a space of d dimension, according to the following procedure. First, represent the phenotypes' fitnesses as points in the space of different environments, each environment defining a dimension (orange dots in Fig. 7.2A). Second, construct the convex hull of these points to find the fitnesses achievable by switching strategies D_f (orange area in Fig. 7.2B), and find the Pareto-optimal frontier of that set (blue line in Fig. 7.2B). Third, plot this Pareto surface in log-fitness space (blue line in Fig. 7.2D). Finally, find the hyperplane normal to $\mathbf{p}$ that is a supporting hyperplane of the Pareto frontier (red and green lines in Fig. 7.2D), and read off the optimal strategy as the intersection point between that hyperplane and the Pareto frontier (red and green stars in Fig. 7.2D).

When the Pareto frontier is contained in a hyperplane, fitnesses can be rescaled onto the unit simplex, $\sum_x f(x) = 1$, with no loss of generality [200]. In this case the optimal strategy is given in terms of the rescaled fitnesses as $\mathbf{f}^* = \mathbf{p}$, making the graphical construction even simpler (Fig. D.1 and App. D.1).

7.3.3 *Transitions between switching, non-switching, and generalist strategies*

The graphical method provides a visual approach to classify the different possible adaptive strategies. For the sake of simplicity, we start again with the

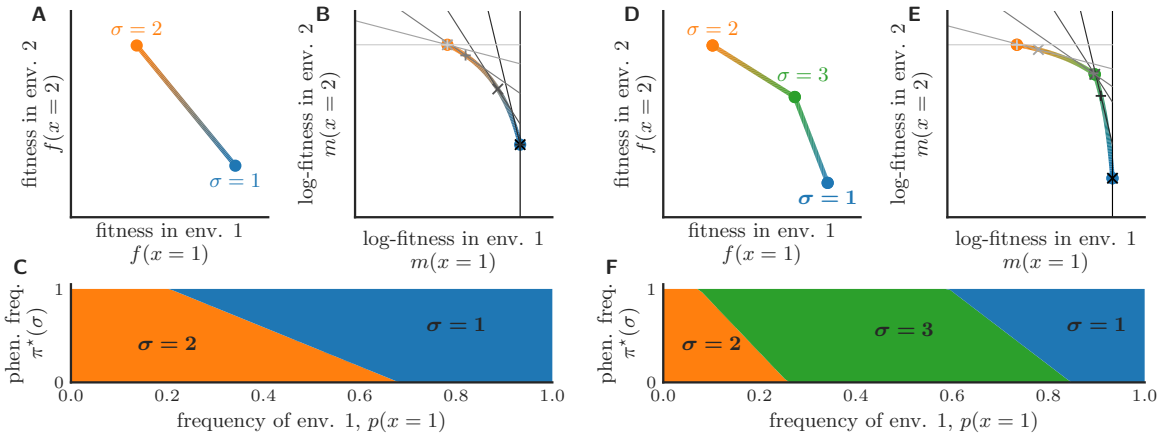

Figure 7.3: Transitions of the optimal strategy as a function of environmental frequencies without (A-C) and with (D-F) a generalist phenotype. (A,D) Pareto frontier of achievable fitness vectors by phenotypes (dots) and their mixtures (lines). (B,E) In log-fitness space a tangent construction (grey lines) yields the optimal strategy (grey crosses) for different environments (from dark to light grey for $p_1 = 1 \to p_1 = 0$ in 0.2 steps). (C,F) Transitions between switching and non-switching strategies as a function of the probability of encountering environment 1. Parameters: (A-C) $f(\sigma = 1) = (1, 0.3), f(\sigma = 2) = (0.4, 1)$, (D-F) $f(\sigma = 1) = (1, 0.2), f(\sigma = 2) = (0.3, 1.0), f(\sigma = 3) = (0.8, 0.7)$.

case of a two-state environment and first assume that only two phenotypes are accessible: a blue phenotype ($\sigma = 1$) best suited to environment 1 and an orange phenotype ($\sigma = 2$) best suited to environment 2 (Fig. 7.3A-C). In this case, the Pareto front is a segment joining the two phenotypes. In log-fitness space, this segment is curved and concave, implying that $\partial m(2)/\partial m(1)$ is a decreasing function of $m(1)$. Different environmental statistics are characterized by the frequencies $p(1)$ and $p(2) = 1 - p(1)$ of the two environmental states. The value of $p(1)$ sets the slope $-p(1)/p(2)$ of the isolines of growth rate that we should consider Eq. 7.9.

Depending on the value of $p(1)$, different cases arise. First, if $p(1)$ is too high or too low, there is no tangent to the Pareto front of slope $-p(1)/p(2)$ and the support point lies at one of the two extremities of the Pareto front. In these cases, the optimal strategy (crosses in Fig. 7.3B) is to adopt a constant phenotype – the phenotype optimal for the most frequent environmental state. When $p(1)$ takes an intermediate value, the isoline is tangent to the Pareto frontier at an intermediate support point, indicating an optimal strategy involving switching between the two possible phenotypes. As a function of the frequency of encountering different environments, there are thus two transitions, from non-switching to switching and to non-switching again. This succession of optimal strategies is read off as a function of the environmental frequency from the Pareto line (Fig. 7.3C).

One can make the problem more interesting by adding a third "generalist" phenotype, which does relatively well across both environments (Fig. 7.3D-F, green dot). This generalist creates a kink in the Pareto frontier, meaning that it will be optimal as a constant phenotype for a certain range of environmen-

tal conditions. Thus, depending on the frequencies of the two environmental states, the optimal strategy consists either of having a constant specialized phenotypes (blue or orange) when one environment is much more frequent than the other, a constant generalist phenotype (green) when the two environments have similar frequencies, or switching between a specialized phenotype and the generalist phenotype in intermediate situations (Fig. 7.3F). The transition from specialist to generalist was studied in a similar model in [200], but in the slightly different context of a continuous choice of strategies.

These conclusions generalize to an arbitrary number d of environmental states. It follows from the graphical construction that for a given statistics of the environment, the number of discrete phenotypes between which the population may switch in optimal strategies is at most equal to the number of different environmental conditions, d: the subset of the extended fitness set corresponding to this switching is the polytope of dimension $d-1$ whose vertices are these d phenotypes (a segment for $d=2$, a triangle for $d=3$). This observation may be viewed as extending to changing environments the principle of competitive exclusion stating that a single niche cannot support more than one species.

We complement the graphical analysis by analytical results in the simplest case of two environments and two phenotypes illustrated by Fig. 7.4A-C. Since only the relative fitnesses in each environment is relevant for the dynamics, we set without restriction of generality the replication rate of each phenotype in its preferred environment to 1. The other phenotype has a selective disadvantage, with replication rate $w_x < 1$:

$$f(\sigma, x) = \begin{cases} 1 & \sigma = x, \\ w_x & \sigma \neq x. \end{cases} \tag{7.10}$$

The parameter w_x can be interpreted as the degree of specialization: $w_x = 1$ means no specialization, while $w_x = 0$ means extreme specialization. Since $p(1) = 1 - p(2)$ and $\pi(1) = 1 - \pi(2)$, there are just two free parameters $p_2 \equiv p(2)$ and $\pi_2 \equiv \pi(2)$. In these variables the long-term growth rate (7.7) is written as

$$\begin{aligned} \Lambda = {} & p_2 \log[(1-\pi_2)w_2 + \pi_2] \\ & + (1-p_2)\log[1-\pi_2+\pi_2 w_1)]. \end{aligned} \tag{7.11}$$

To find the optimal fraction of the population with phenotype 2, $\pi_2^\star$, Eq. (7.11) is to be maximized over $\pi_2 \in [0,1]$. The optimization yields

$$\pi_2^\star = \begin{cases} 0 & \text{if} \quad p_2 \leqslant p_2^{lb}, \\ \frac{p_2 - p_2^{lb}}{p_2^{ub} - p_2^{lb}} & \text{if} \quad p_2^{lb} < p_2 < p_2^{ub}, \\ 1 & \text{if} \quad p_2 \geqslant p_2^{ub} \end{cases} \tag{7.12}$$

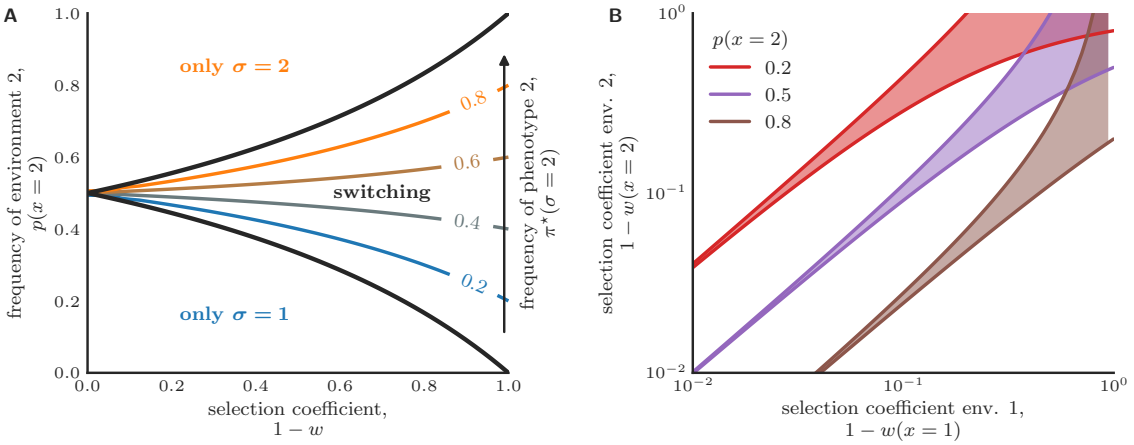

Figure 7.4: Transitions between switching and non-switching strategies depend on environmental selectivity and environmental frequencies. In a temporally uncorrelated environment changing randomly between two states, 1 and 2, a population of organisms is adapted optimally by either being in a single phenotypic state or by having a mixture of phenotypes (bet-hedging) depending on the statistics of the environment and the degree to which the phenotypes are specialized. In environment $x = 1$ (2) phenotype 2 (1) has replication rate w_x relative to the other phenotype. (A) Transitions as a function of specialization level and environmental frequency in the symmetric case, $w_1 = w_2 = w$. The black lines mark the transition from single-phenotype to bet-hedging strategies: above the upper (lower) line the entire population optimally has phenotype 2 (1), between the two lines phenotypic diversification provides an advantage. The optimal fraction of phenotype 2 in the bet-hedging region is shown by the colored lines. (B) Regions of selection factors in which bet-hedging is the preferred strategy (shaded areas) for environments with different frequencies of being in state 2. Either strong selection or a precise mapping between the relative selection factors and the relative environmental frequencies are needed to make bet-hedging optimal.

with lower and upper bounds

$$p_2^{lb} = \frac{w_2}{1+(1-w_2)w_1/(1-w_1)} \tag{7.13}$$

$$p_2^{ub} = \frac{1}{1+(1-w_2)w_1/(1-w_1)} \tag{7.14}$$

on the environmental frequencies for which diversification is optimal. The calculation recapitulates the conclusions from the graphical method (Fig. 7.4C). In the limit were selection is very stringent, $w_1 \to 0$ and $w_2 \to 0$, the transitions disappear, $p_2^{lb} \to 0$ and $p_2^{ub} \to 1$ and the optimal strategy reduces to proportional betting,

$$\pi_2^\star = p_2. \tag{7.15}$$

In the context of biological bet-hedging, this result was already noted by Cohen [118]; it was also derived earlier in the context of gambling by Kelly [206].

The range of environmental frequencies for which bet-hedging is favored over the non-switching strategies depends strongly on the selectivity of the environments (Fig. 7.4A). Consider for simplicity the symmetric case $w_1 = w_2 = w$, then non-switching strategies are favored for $|p_2 - 1/2| \geqslant (1-w)/[2(1+w)]$. In the limit, $w \to 0$, the strategy tends to proportional bet-hedging as discussed earlier. The larger w, the smaller the region of environmental frequencies for which switching strategies are optimal. As there is smaller variability in fitness across generations for the same phenotype, switching is less needed to hedge against environmental fluctuations.

Instead of considering transitions in optimal strategies as environmental frequencies are varied, we can also consider transitions as selection pressures are varied at fixed environmental frequencies (Fig. 7.4B). As selection pressures are decreased there are transitions to a non-switching strategy (white areas in Fig. 7.4B). The optimality of bet-hedging (shaded areas in Fig. 7.4B) for weak selection pressures depends on a precise matching between the asymmetry in selection pressures and environmental frequencies. This conclusion generalizes the results of [201], which considered numerically asymmetric fitness landscapes, $w_1 \neq w_2$, but only with a symmetric environment, $p_1 = p_2 = 1/2$.

7.4 TRANSITIONS BETWEEN OPTIMAL IMMUNE STRATEGIES

Fitnesses achievable by single phenotypes (orange dots in Fig. 7.2) can fill a set delimited by a continuous line, called *trade-off function*, which is the Pareto frontier of non-switching strategies. It is common to consider such a continuous set of phenotypes with all possible switching strategies between them [109, 200]. The Pareto frontier of switching strategies defined in the previous section then delimits the convex hull of that continuous set. The two Pareto frontiers (of switching and non-switching strategies) coincide if the trade-off function is concave, i.e. if the set of achievable phenotypes is convex; in that case non-switching strategies are optimal everywhere. Otherwise, similar transitions as in the previous section will arise [200]. In some biological situations

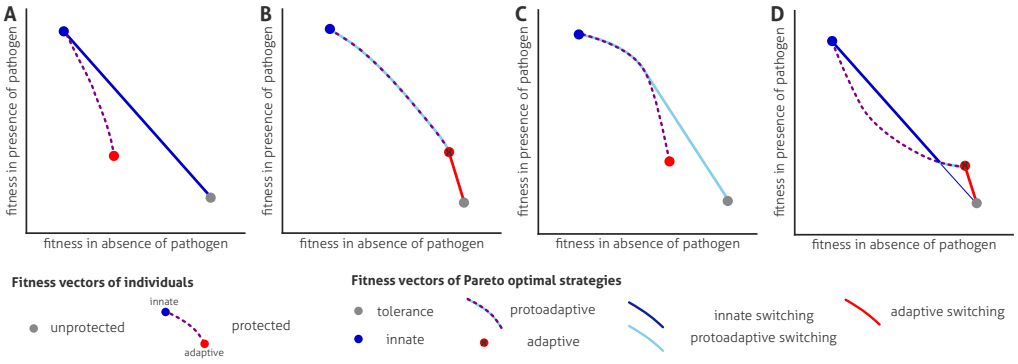

Figure 7.5: Strength of trade-offs between constitutive and defense cost of protection determine adaptation strategy in a fluctuating pathogenic environment. In the model, unprotected individuals have a fixed fitness profile (grey dot). Protection comes in various degrees of adaptability (dashed purple line) between maximal (blue dot) and minimal (red dot) level of constitutive investement in defense. Switching strategies are possible where only parts of the population are protected. They have fitnesses that are a linear combination of the fitness of unprotected and protected indviduals for a given level of adaptability. The optimal strategy needs to lie along the Pareto frontier of the possible fitnesses. The strategies that lie on the Pareto surface allow reading off the succession of optimal strategies as the probability of encountering the pathogen is decreased. (A) Strong trade-offs lead to switching strategies being better then adaptable protection. (B) For shallow trade-offs the Pareto frontier is achieved by adaptable defenses. (C) A combination of shallow and steep trade-offs can lead to only some degree of adaptability being used. (D) A concave trade-off function can lead to first order transitions in strategy and potential co-existence of locally optimal solutions.

however, only some combinations of phenotypes along a trade-off function may be accessible at the same time, meaning that one cannot switch between all phenotypes on the trade-off line. Such a constraint on switching rates can induce discontinuous transitions, or cause the co-existence of multiple locally optimal solutions, as we now illustrate in a simple model of evolution of immunity.

Our illustrative example is a model that we proposed to explain the diversity of immune strategies observed across the tree of life [139]. The purpose is to show how different strategies are associated with different statistics of pathogen dynamics. In its simplest form, the model has two environmental states, presence ($x = 1$) or absence ($x = 0$) of a pathogen. In a given strategy, it has two accessible phenotypes, protected ($\sigma = 1$) or unprotected ($\sigma = 0$). Strategies are represented by $\mathbf{f} = (f(x = 0), f(x = 1))$ as before.

The unprotected phenotype is fixed in fitness space: $\mathbf{f} = (f_{base}, f_{inf})$ (grey dot in Fig. 7.5), where $f_{inf} < f_{base}$ is the reduced fitness in infected unprotected individuals. By constrast, the protected phenotype lies on a trade-off function: $\mathbf{f} = (f_{con}, f_{def}(f_{con}))$, with $f_{con} \in [f_{con}^{min}, f_{con}^{max}]$ (dashed purple line delimited by red and blue dots in Fig. 7.5). $f_{con} < f_{base}$ represents the reduced fitness due to the investment into the protection, while $f_{def} > f_{inf}$ is the fitness of protected individuals in presence of the pathogen.

The choice of f_{con} along the trade-off function sets the investment into the protection, and is part of the strategy: once this strategy is fixed, it is possible to switch between protected and unprotected phenotypes, but not between different points of the trade-off function. This constraint can be justified biologically by the high cost of plasticity that such switches would incur.

The function $f_{def}(f_{con})$ encodes the trade-off between the efficiency of the protection and its cost. By analogy with immune mechanisms in vertebrates, we interpret it in terms of *adaptivity* of the response within the lifetime of the organism, with higher adaptivity enabling lower cost at the expense of lower protective efficiency [139]. We therefore refer to the maximally protective and costly strategy with $f_{con} = f_{con}^{max}$ as *innate immunity* and to the minimally protective and costly strategy with $f_{con} = f_{con}^{min}$ as *adaptive immunity*. Intermediate strategies with $f_{con}^{min} < f_{con} < f_{con}^{max}$ are referred to as *protoadaptive*.

Within this model, the equation for long-term growth rate in an uncorrelated environment (7.7) becomes

$$\begin{aligned}\Lambda =& p \ln[\pi f_{def} + (1-\pi) f_{inf}] \\ &+ (1-p) \ln[\pi f_{con} + (1-\pi) f_{base}],\end{aligned} \tag{7.16}$$

where $p \equiv p(x = 1)$ is the probability of the presence of the pathogen and $\pi \equiv \pi(x = 1)$ the probability of being protected. Here, the problem is not only to find the optimal switching probability $\pi^\star$, but also to find the optimal protection adaptability, $f_{con}^\star$. To summarize, the problem is as follows: for a given p, f_{inf}, f_{base} and $f_{def}(f_{con})$, find $f_{con}^\star$ and $\pi^\star$ that maximize long-term growth rate in Eq. 7.16.

We are particularly interested in transitions between $f_{con}^\star, \pi^\star$ taking intermediate or extremal values within their respective ranges. Given that each of these two variables can either reach its lower or upper bound or take an intermediate value, nine different cases may arise. However, since the level of adaptability of the response is inconsequential if none of the population is protected ($\pi^\star = 0$), only seven qualitatively different immune defense strategies are relevant: tolerance ($\pi^\star = 0$, grey dot in Fig. 7.5), innate ($\pi^\star = 1$, $f_{con} = f_{con}^{max}$, blue dot in Fig. 7.5), adaptive ($\pi^\star = 1$, $f_{con} = f_{con}^{min}$, red crossed dot in Fig. 7.5), protoadaptive ($\pi^\star = 1$, $f_{con}^{min} < f_{con} < f_{con}^{max}$, light blue line with purple dashes in Fig. 7.5), innate switching ($0 < \pi^\star < 1$, $f_{con} = f_{con}^{max}$, blue line in Fig. 7.5), adaptive switching ($0 < \pi^\star < 1$, $f_{con} = f_{con}^{min}$, red line in Fig. 7.5), and protoadaptive switching ($0 < \pi^\star < 1$, $f_{con}^{min} < f_{con} < f_{con}^{max}$, light-blue line in Fig. 7.5).

Which of these strategies is optimal in a given environment? And what is the nature of the transitions between strategies as the frequency of encountering the pathogen is varied? Here, we apply the graphical method to answer these questions and show how the answers depend critically on the shape of the trade-off function. Our conclusions, summarized in Fig. 7.5, are supported by analytical results derived in Appendix D.2.

The simplest case is when adaptability comes at an excessive cost, as depicted in Fig. 7.5A: an innate switching strategy is then always preferable to an adaptive strategy. In this case, as the probability of encountering the pathogen increases, the optimal strategy transitions from tolerance (grey dot in

Fig. 7.5A) to an innate defense strategy (blue dot in Fig. 7.5A) via an innate switching (blue line in Fig. 7.5A). When adaptability of the defense does not impair its effectiveness as severely, as in Fig. 7.5B, two new transitions occur. As the probability of encountering the pathogen increases, the optimal strategy now transitions from tolerance to, successively, adaptive switching, adaptive, protoadaptive and finally innate defense strategy. In other cases, a switching protoadaptive defense strategy may also be optimal, as in the case of the trade-off function of Fig. 7.5C. In this case, as the probability of encountering the pathogen increases, the optimal strategy transitions from tolerance to, successively, protoadaptive switching, protoadaptive and finally innate defense strategy. Finally, we may consider a case where the trade-off line is not convex as in Fig. 7.5D. The Pareto frontier is then not necessarily concave, and we might have first order transitions between strategies. For the trade-off shape shown in Fig. 7.5D, there is a transition from protoadaptive switching (blue line with purple dashes) directly to innate switching (blue line), with a discontinuity in the level of adaptability of the response.

7.5 WHEN AND HOW TO USE MEMORY IN TEMPORALLY CORRELATED ENVIRONMENTS

In temporally correlated environments, the past phenotypes of an individual carry information about the next environmental state. The optimal switching strategy may thus involve memory, i.e. it may be advantageous for $\pi(\sigma|\sigma')$ to depend on σ'. Stochastic switching with memory serves an additional purpose relative to the memoryless switching strategies considered so far: in addition to providing a bet-hedging mechanism against the uncertainty of the environment, it provides the variation and heritability needed for tracking the environmental state. Here, we extend the previous analysis to characterize the conditions under which temporal correlations in environmental fluctuations favor switching strategies with memory over non-switching strategies. The graphical method does not extend to correlated environments but we show that the transitions between switching and non-switching strategies can be characterized analytically.

7.5.1 *Insights from the adiabatic limit*

It is instructive to start with long correlation times, when the duration of each environmental state is much longer than the time that it takes for the population to reach its steady state composition. In this adiabatic limit, the model is analytically solvable [111, 117]. We present the solution for the case where switching takes place between a number of different phenotypes, with each phenotype σ being best in one environment x, which we denote by the same symbol $\sigma = x$ (other cases can in fact always be reduced to this one [117]). A calculation based on a series of eigendecompositions of the growth matrix in

different environments (see Appendix D.3 for derivation) leads to an expression of the long-term growth rate as [117]

$$\begin{aligned}\Lambda = &\sum_x p(x) \ln f(x,x) \\ &+ \sum_{x,x'} p(x|x') p(x') \ln[\pi(x|x')\Gamma(x,x')],\end{aligned} \tag{7.17}$$

which involves the overlap $\Gamma(x,x')$ between steady-state population compositions in environments x, x', given by

$$\Gamma(x,x') = \frac{f(x,x')}{f(x',x') - f(x,x')} + \frac{f(x,x)}{f(x,x) - f(x',x)}, \tag{7.18}$$

if the environment changes, $x \neq x'$, and 1 otherwise.

Optimizing Eq. 7.17 over $\pi(x|x')$ subject to the normalization constraint leads to $\pi^\star(x|x') = p(x|x')$. Within the adiabatic limit, the optimal strategy is therefore always to diversify, with switching rates equal to the environmental switching rates. This generalizes the result that proportional betting is optimal in the limit of strong selection, Eq. 7.15, to the case where reaching steady state takes longer but environmental switches are rarer. In contrast to the results in the previous section, switching is always favored in the adiabatic limit, even when selection is weak.

We can use the expression of Eq. 7.17 to ask how much each phenotype σ should be specialized to its environement $x = \sigma$. Being more specialized means higher fitnesses of the adapted phenotypes, $f(x,x)$, at the expense of lower fitnesses for the maladapted phenotypes, $f(x,x' \neq x)$, assuming a trade-off between the two. More specialized phenotypes have lower relative replication rate $w(x,x') = f(x,x')/f(x',x')$ [$w(x,x')$ reduces to $w_{x'}$ of Eq. 7.10 in the case of two environmental states]. Rewriting

$$\Gamma(x,x') = \frac{w(x,x')}{1 - w(x,x')} + \frac{1}{1 - w(x',x)}, \tag{7.19}$$

we see that specialization also implies lower overlaps $\Gamma(x,x')$, and thus lower values for the second term in the long-term growth rate Eq. 7.17. On the other hand, the first term in Eq. 7.17 grows with $f(x,x)$, i.e. with higher specialization. Because of these contradictory terms, the optimal strategy along the trade-off between $f(x,x)$ and $f(x,x')$ will depend on the details of trade-off function and of the environmental statistics. However, as environment fluctuations become slower, $p(x|x' \neq x) \to 0$, the second term in Eq. 7.17 vanishes for $x \neq x'$, letting the first term dominate. In that limit, highly specialized phenotypes become more and more advantageous. This observation is again in contrast with the results of the preceding section (Fig. 7.3D-F), which have shown that generalists are optimal under certain environmental conditions.

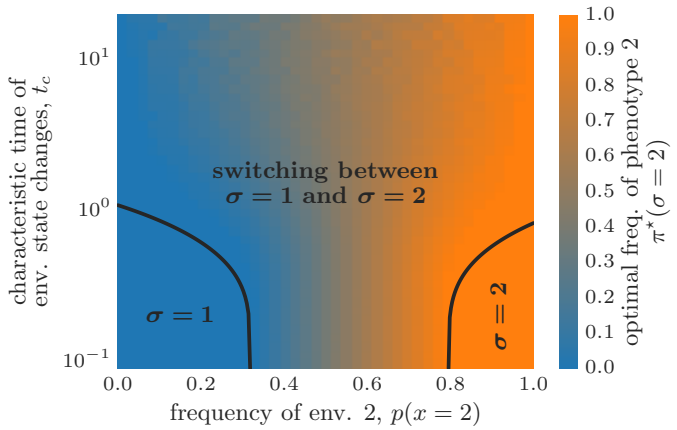

Figure 7.6: Switching strategies are favored over a larger range of conditions if environmental states are temporally autocorrelated. Here we generalize the results of Fig. 7.3A-C about transitions between switching and non-switching strategies by considering the influence of environmental correlation. The numerically obtained optimal switching rate $\pi^\star(\sigma = 2)$ is plotted as a function of t_c, the characteristic time scale of environmental changes, and $p(x = 2)$, the fraction of the time the environment is in state 2. The range of environmental frequencies in which there is switching ($0 < \pi^\star < 1$) increases with temporal correlations. As a comparison we also show the analytical transition lines obtained in Sec. 7.5.3, Eqs. (7.31)-(7.32).

7.5.2 *Connecting the limit of uncorrelated and adiabatically switching environments numerically*

So far we have considered two opposite limits: temporally uncorrelated environments in Sec. 7.3 and 7.4, and temporally correlated environment with long correlation times in Sec. 7.5.1. These two limits give very different answers to the questions of whether bet-hedging is desirable, or whether generalist phenotypes can be optimal. To study the intermediate regime between these two extremes, we first start by presenting the results of a numerical study, based on the recursion equation Eq. 7.1. We apply the numerical approach described in Ref. [139]. In short, we approximate the long-term growth rate numerically by simulating for a large number of generations, and then use a derivative-free global optimization algorithm to roughly find the global optimum. In practice, we focus on two-state environments, which we characterize by their characteristic time scale, t_c, defined by $e^{-1/t_c} = 1 - p(1|2) - p(2|1)$ and the probability of being in state 2, $p(x = 2)$. The numerical results show how the two limits are connected for the case without (Fig. 7.6) and with a generalist phenotype (Fig. 7.7). In temporally correlated strategies, phenotype frequencies vary with the environmental history. To represent strategies in a simple way that generalizes the case of memoryless strategies, we define $\pi(\sigma)$ as the steady-state frequency of phenotype σ in a lineage, $\sum_{\sigma'} \pi(\sigma|\sigma')\pi(\sigma') = \pi(\sigma)$. Consistent with results in the adiabatic limit, for large t_c switching strategies dominate across the range of environmental frequencies (Figs. 7.6 and 7.7): $\forall\sigma, \pi^\star(\sigma) < 1$. In the case where there is an intermediate, generalist phenotype ($\sigma = 3$) the switching takes place primarily between the specialist types: $\pi^\star(\sigma = 3) \ll 1$

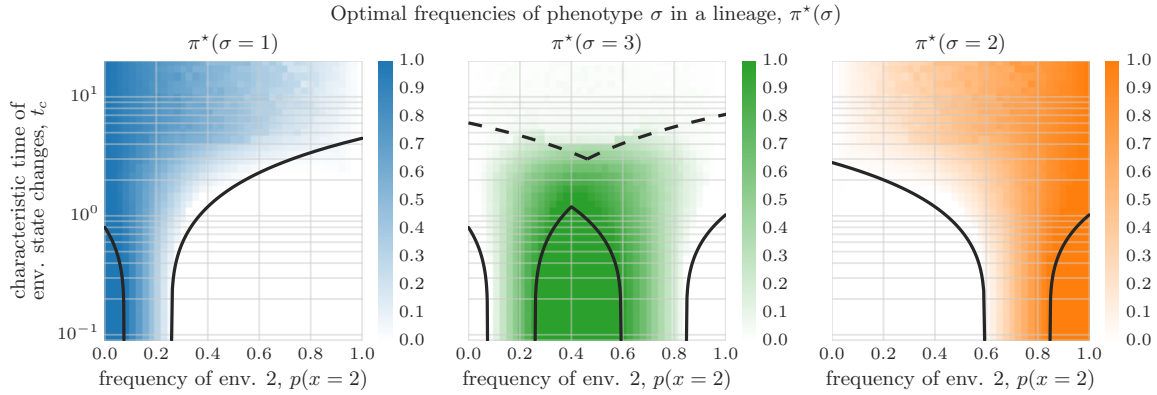

Figure 7.7: Switching between specialists is the preferred adaptation strategy in highly correlated environments even if a generalist phenotype is optimal in uncorrelated environments. Here we generalize the results of Fig. 7.3D-F about transitions between switching, specialist, and single-phenotype generalist strategies by considering the influence of environmental correlation. The numerically determined optimal frequencies of different phenotypes $\pi^\star(\sigma)$ in a lineage are plotted as a function of t_c, the characteristic time scale of environmental changes, and $p(x = 2)$, the fraction of the time the environment is in state 2. As a comparison we also show the analytical transition lines between single-phenotype and switching strategies obtained in Sec. 7.5.3, Eqs. (7.31)-(7.32) (solid lines) and the approximate transition line above which switching takes place between the two specialist phenotypes as obtained in Sec. 7.5.1 and App. D.3, Eq. (D.46) (dashed lines).

for large t_c (Fig. 7.7), consistent with the argument that specialized phenotypes are optimal in the adiabatic limit (Sec. 7.5.1). The transition to a regime where non-switching strategies are optimal happens when the temporal correlations of the environment are of the order of the generation time, $t_c \sim 1$. In this regime, all three phenotypes (two specialists and one generalist) may co-exist in the optimal strategy, $\forall \sigma, \pi^*(\sigma) > 0$. Recall that such mixtures involving more phenotypes than distinct environments are suboptimal in memoryless environments, $t_c = 0$, as deduced from the graphical construction (see Sec. 7.3.3).

7.5.3 *An analytical result for intermediate timescales*

We present here an approach to derive analytically the boundaries between optimal switching and non-switching strategies in correlated environments. The approach is based on an expansion at small switching rates of the Master equation of the joint environmental and population switching process near the transition boundary.

For notational convenience we assume replication precedes mutation and rewrite the recursion equation for the fraction of the population in each state $n_t(\sigma) = N_t(\sigma)/N_t$ with $N_t = \sum_\sigma N_t(\sigma)$,

$$n_{t+1}(\sigma) = \frac{1}{Z_t} \sum_{\sigma'} \pi(\sigma|\sigma') f(\sigma', x_t) n_t(\sigma'), \tag{7.20}$$

where Z_t is a normalization constant enforcing $\sum_\sigma n_t(\sigma) = 1$. Since $N_T = N_0 \prod_{t=0}^{T} Z_t$ the long-term growth rate given by Eq. 7.3 becomes

$$\Lambda = \lim_{T\to\infty} \frac{1}{T} \sum_{t=0}^{T} \ln Z_t. \tag{7.21}$$

For simplicity, we consider a two-state model again. We introduce the simplified notations $\pi(1|2) = \pi_{12}$, $\pi(2|1) = \pi_{21}$ for the type switching rates, $p(1|2) = p_{12}$, $p(2|1) = p_{21}$ for the environment switching rates, and denote $n_t(2) = n_t$, and redefine x_t to be 1 if the environment is in state 2 and 0 otherwise. We use the same convention as in Eq. 7.10, $f(1,1) = f(2,2) = 1$, $w(2,1) = w_1$ and $w(1,2) = w_2$. This allows us to rewrite the recursion equation as

$$n_{t+1} = \frac{1}{Z_t}\left(n_t(1-\pi_{21})w_1^{1-x_t} + (1-n_t)\pi_{12}w_2^{x_t}\right) \tag{7.22}$$

with

$$Z_t = n_t w_1^{1-x_t} + (1-n_t)w_2^{x_t}. \tag{7.23}$$

To analyze the transition from an optimal strategy where all individuals have phenotype 1 to a strategy with some switching to the other phenotype, we need to know whether a small π_{12} is better than $\pi_{12} = 0$ – if that is the case, then switching is advantageous. We thus consider $\pi_{12} \ll 1$ and $n_t \ll 1$. The recursion Eq. 7.22 becomes at leading order in n_t and π_{12}:

$$n_{t+1} = \pi_{12} + (1-\pi_{21})w_1^{1-x_t}w_2^{-x_t}n_t. \tag{7.24}$$

$\ln Z_t$ can also be expanded:

$$\begin{aligned} \ln Z_t &= x_t \ln w_2 - n_t + n_t w_1^{1-x_t}w_2^{-x_t} \\ &= x_t \ln w_2 + n_t(w_1 - 1) + x_t n_t(w_2^{-1} - w_1). \end{aligned} \tag{7.25}$$

Over long times the joint environmental-population process is ergodic. The long-term growth rate is thus given as $\langle \ln Z \rangle$, where $\langle . \rangle$ indicates an average over the steady state distribution of x, n. No switching ($n = 0$) gives a long-term growth rate of $\langle x \rangle \ln w_2$ Thus the difference in long-term growth rate between stochastic switching and the single-phenotype strategy is

$$\Delta\Lambda = \langle n \rangle (w_2^{-1} - w_1)\left(\frac{\langle xn \rangle}{\langle n \rangle} - \frac{1-w_1}{w_2^{-1} - w_1}\right), \tag{7.26}$$

which shows that stochastic switching is advantageous if

$$\frac{\langle xn \rangle}{\langle n \rangle} > \frac{1-w_1}{w_2^{-1} - w_1}. \tag{7.27}$$

We can identify the right-hand side of this equation with the lower bound environmental frequency p_2^{lb} in uncorrelated environments defined in Eq. 7.13.

When there is no memory, n and x are uncorrelated, the left-hand side reduces to $\langle x \rangle = p(x = 2)$ and we recover the result of Eq. 7.12. If there is memory, then n and x are positively correlated through the effects of selection on the population composition, which increases the fraction on the left-hand side. This leads us to a first important conclusion: switching is favored over non-switching strategies under a wider range of environmental parameters in the presence of temporal autocorrelation.

We go further and calculate analytically the left-hand side of Eq. 7.27 at the transition. Some algebra shows that $\rho_{1,t} = \langle (1 - x_t) n_t \rangle$ and $\rho_{2,t} = \langle x_t n_t \rangle$ satisfy the recursion

$$\begin{aligned}
\rho_{1,t+1} =& p_{21}[\pi_{12} p_2 + (1 - \pi_{21}) w_2^{-1} \rho_{2,t}] \\
&+ (1 - p_{12})[\pi_{12}(1 - p_2) + (1 - \pi_{21}) w_2 \rho_{1,t}], \qquad (7.28) \\
\rho_{2,t+1} =& (1 - p_{21})[\pi_{12} p_2 + (1 - \pi_{21}) w_2^{-1} \rho_{2,t}] \\
&+ p_{12}[\pi_{12}(1 - p_2) + (1 - \pi_{21}) w_2 \rho_{1,t}], \qquad (7.29)
\end{aligned}$$

where we use the short-hand notation $p_2 = p_{21}/(p_{12} + p_{21}) = \langle x \rangle$ for the average fraction of generations the environment is in state 2. Therefore at steady state, $\rho_{\sigma,t} = \rho_{\sigma,t+1} = \rho_\sigma$, we have

$$\begin{aligned}
\frac{\langle xn \rangle}{\langle n \rangle} &= \frac{\rho_2}{\rho_2 + \rho_1} \\
&= \frac{p_2 \left[1 - (1 - \pi_{21}) e^{-1/t_c} w_1\right]}{1 - (1 - \pi_{21}) e^{-1/t_c} w_1 \left[(1 - p_2) w_1 w_2 + p_2\right]},
\end{aligned} \qquad (7.30)$$

where we recall that $e^{-1/t_c} = 1 - p(1|2) - p(2|1)$ quantifies memory in the environment. The expression in (7.30) is a decreasing function of π_{21} so its maximum is achieved in the limit of π_{21} going to zero. Setting $\pi_{21} = 0$ in Eq. 7.30 and plugging the result into Eq. 7.27, we obtain the condition needed for optimal switching to outperform always being in state $\sigma = 1$:

$$p_2 > \frac{(1 - w_1)(1 - e^{-1/t_c} w_2^{-1})}{(w_2^{-1} - w_1)(1 - e^{-1/t_c})}. \qquad (7.31)$$

The second transition, between optimal switching and always being in state $\sigma = 2$, is given by the replacements $w_1 \to w_2$, $w_2 \to w_1$, $p_2 \to 1 - p_2$, yielding the condition:

$$p_2 < \frac{(w_1^{-1} - 1)(1 - e^{-1/t_c} w_2)}{(w_1^{-1} - w_2)(1 - e^{-1/t_c})}. \qquad (7.32)$$

These transitions reduce to (7.12) in the limit of no environment memory, $t_c = 0$. The transition curves reach $p_2 = 0$ and $p_2 = 1$ at $t_c = -1/\ln(w_2)$ and $t_c = -1/\ln(w_1)$, respectively. The resulting phase diagram is shown in Fig. 7.6 along with a numerical optimization, which confirms the results.

The analytical results show that temporal correlations in the environment favor the evolution of stochastic switching. We can compare to the case of uncorrelated environments considered in Fig. 7.4. While switching is only op-

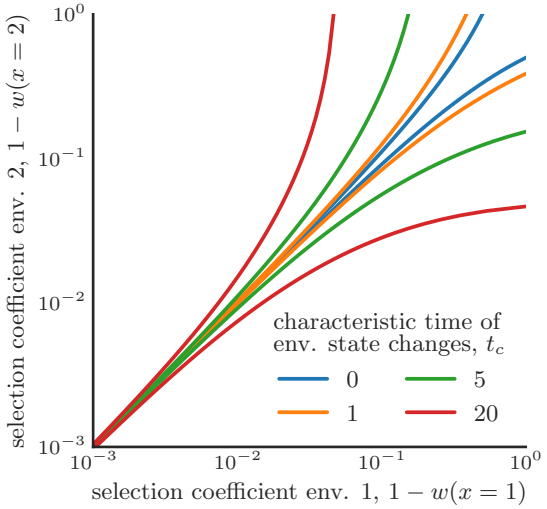

Figure 7.8: Environmental correlations increase the range of fitness landscapes for which switching strategies are optimal. Region where switching is optimal (in between colored lines) as a function of environmental correlation time. Two state environment as in Fig. 7.4 with symmetric environmental frequencies, $p_2 = p_1 = 0.5$. Selection coefficient $s(x)$ quantifies how much the best adapted phenotype to environment x outperforms the suboptimal phenotype for that environment.

timal in uncorrelated environments if selection is strong in both environments (blue line in Fig. 7.8), temporally-correlated environments make it optimal for smaller or asymmetric selection (e.g. red line in Fig. 7.8). We may interpret this broadening of the range where switching is optimal by noting that, in correlated environments, switching does not just contribute to bed-hedging but also to adaptively tracking the state of the environment.

7.5.4 *Continuous time limit*

Lastly, we discuss the continuous time limit of Eq. 7.1 where our results take a simple form. The limit is obtained by rescaling the switching rates, growth rates, and times by δt, $p(x|x') \to p(x|x')\delta t$ for $x \neq x'$, $\pi(\sigma|\sigma') \to \pi(\sigma|\sigma')\delta t$ for $\sigma \neq \sigma'$, and $\ln[f(\sigma, x)] \to m(\sigma, x)\delta t$, $t \to t/\delta t$, $t_c \to t_c/\delta t$, $\ln w_x \to \ln w_x/\delta_t$, $\Lambda \to \Lambda/\delta t$ and sending $\delta t \to 0$, which yields

$$\frac{d\mathbf{N}}{dt} = \mathbf{A}^{x(t)}\mathbf{N}(t), \tag{7.33}$$

where $A^{x(t)}_{\sigma,\sigma'} = m(\sigma, x(t))\delta_{\sigma,\sigma'} + \pi(\sigma|\sigma')$.

We can take the limit of the results obtained in Sec. 7.5.3 to see how temporal autocorrelation influences the results in this case. From Eq. 7.31 we obtain

$$p_2 > \frac{1 + t_c \ln w_2}{1 + \ln w_2 / \ln w_1}, \tag{7.34}$$

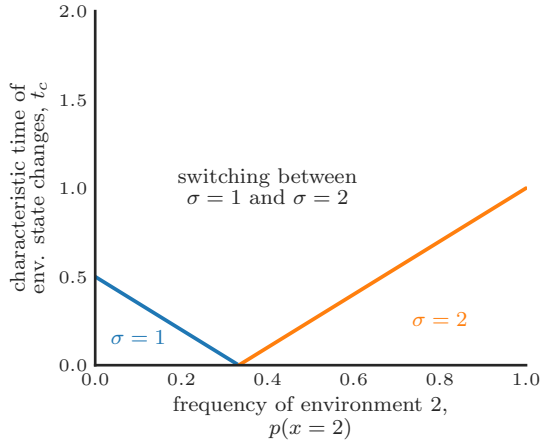

Figure 7.9: Phase diagram in the continuous time limit. t_c is the environment correlation time, and p is the fraction of the time the environment is in state 2. On the left of the blue line, the optimal solution is for the population to have single phenotype $\sigma = 1$. On the right of the red line, the optimal solution is to have single phenotype $\sigma = 2$. In between the optimal solution is to switch between both phenotypes. The blue transition line reaches $p = 0$ at $t_c = -1/\ln w_1$, while the red transition reaches $p = 1$ at $t_c = -1/\ln w_2$. The two transitions meet at $p = \ln w_1/\ln(w_1 w_2)$ (dashed line). Parameters: $\ln w_2 = -2$ and $\ln w_1 = -1$.

and from Eq. 7.32

$$p_2 < \frac{1 - t_c \ln w_2}{1 + \ln w_2/\ln w_1}. \tag{7.35}$$

The formulas are simpler and notably linear in the correlation time t_c. The range of environmental frequencies for which stochastic switching is optimal thus grows linearly with the environmental correlation time scale t_c, as $-2t_c \ln w_1 \ln w_2/\ln(w_1 w_2)$.

The point $p_2 = 0$ is reached by the first transition Eq. 7.34 from the non-switching to switching regime for $t_c = -1/\ln w_2$, and the point $p_2 = 1$ reached by the second transition for $t_c = -1/\ln w_1$ Eq. 7.35, as in the discrete time case. In the limit of no environmental memory, $t_c = 0$, the two transitions are at the same point $\ln w_1/\ln(w_1 w_2)$: this means that bet-hedging is never advantageous and the transition is from one single-phenotype strategy to the other. This is in contrast with the solution in discrete time, where there always is a window in which bet-hedging is favored, regardless of the environmental memory. Since in any finite time interval, the environment cycles through all its states, the population effectively only sees the mean environment. The long-term growth rate in continuous time is thus given by

$$\Lambda = \sum_x p(x) \sum_\sigma f(\sigma, x)\pi(\sigma) \tag{7.36}$$

which is a linear function in $\pi(\sigma)$. Λ is optimized by putting all weight on the phenotype σ with largest average fitness $\sum_x p(x)f(\sigma,x)$. Thus no switching strategies can be optimal and the optimal strategy always consists of a single phenotype.

7.6 DISCUSSION

Our results provide a unified view of transitions between optimal adaptive strategies in randomly fluctuating environments. By revisiting the fitness set representation of Levins [109], valid for temporally uncorrelated environments, we presented a graphical method, supplemented by analytical calculations, to determine the transitions between bet-hedging and single-phenotype strategies, as well as between specialist and generalist phenotypes (Fig. 7.3), generalizing previous results [118, 109, 207, 206, 117, 200]. Extending the method to phenotypes constrained by a trade-off function, we constructed graphically and calculated analytically the transitions between optimal strategies of diversification and adaptability in a simple model of evolution of immunity [139] (Fig. 7.5).

As noticed in previous studies, temporal correlations in the environmental conditions influences the choice of optimal adaptation strategies [208, 199, 203]. The intermediate timescale regime, where the environmental correlation time is of the same order as the generation time, has been notoriously difficult to handle analytically. Here, we presented an analytical approach to show how temporal correlations in environments can be exploited by switching strategies that keep some memory of previous phenotypes. Our results show that temporal correlations broaden the range of selective pressures for which a switching strategy is better than a single-phenotype one. Everything else being equal, switching strategies are thus more favorable in correlated environments than in uncorrelated environments. To our knowledge, only one other analytical approach is available to analyze optimal strategies in correlated environments [203].

The results are independent of mechanisms, which may take different forms. For instance, one mechanism to achieve a generalist phenotype is through *plasticity*, i.e., a generalist phenotype may partly or totally be induced by the environmental condition. In our approach, however, only the value of the replication rate $f(\sigma,x)$ in environmental condition x given the inherited type σ matters, not the process by which it is achieved. Only when the induced phenotype may be transmitted to the next generation, as for instance with the Lamarckian CRISPR-like strategy of [139], does the distinction between inherited and induced phenotype, and therefore the concept of plasticity, become relevant.

Possible extensions of our results include the influence of non-random environmental changes, such as periodic environments [209, 204, 199, 203], constraints on relative switching rates [201, 209, 202], active sensing mechanisms [111] and heritable plasticity [186, 139], or finite population size effects [210]. Some of these factors are known to lead to transitions between adaptive strategies, e.g. the variability of environmental durations [203], or cause the transi-

tions to become discontinuous, e.g. when switching rates are constrained to be independent of phenotype [201, 209, 202].

Acknowledgements. We thank B. Xue and O. Carja for helpful discussions. The work was supported by grant ERCStG n. 306312.

8 CONCLUSIONS

8.1 MAIN CONTRIBUTIONS OF THIS THESIS

As outlined in Ch. 1 this thesis has explored consequences of the hypothesis that immune systems are well-adapted to the statistical structure of pathogenic environments. This perspective has led us to a number of interesting results about how optimal immune systems allocate resources and process information with regard to the pathogenic environment.

The composition of immune repertoires determines the breadth and speed of protection provided by the adaptive immune system. We have developed a statistical model of pathogen recognition by the adaptive immune system (Ch. 4,5). By analyzing optimal resource allocation within this model we provide a novel theoretical framework for understanding the organization and dynamics of immune repertoires.

Our work on optimal resource allocation in a fixed pathogen environment (Ch. 4) makes concrete testable prediction about statistical features of well-adapted repertoires that can inform analysis of high-throughput repertoire sequencing experiments. We have shown that optimal repertoires are biased towards common pathogens but less than proportionally to the relative frequency. This speeds up defense by providing efficient protection against common threats while leaving sufficient resources for combatting rare threats that could become harmful if left unchecked. Considering resource allocation with cross-reactivity of receptors we demonstrated that optimal repertoires focus resources on a few clones that collectively tile the space of pathogens. The precise position of these peaks is ill-constrained so that a small difference in the experienced antigen distribution can lead to large changes of the repertoire. Repertoires that differ a lot on the sequence level can provide comparable protection once cross-reactivity is factored in. This cautions against measures of immune protection based solely on the idea of "public" repertoires of specific cells shared across individuals [30]. Our results argue for a view in which protection should be defined on a broader level that takes into account cross-reactivity. There are some encouraging hints coming from bioinformatic analyses to sustain this view. For example, using the frequency of short stretches of CDR3 amino acid sequences provides better discrimination of pre- and post-immunization repertoires than using full sequence identity [60, 61].

We have also shown that optimal repertoires in a fixed environment can be reached via a self-organized dynamics based on antigen-driven proliferation and competition between clones cross-reactive to the same antigens (Ch. 4). This result has prompted us to investigate in more generality optimal repertoire dynamics in changing pathogenic environments. We have derived the optimal dynamics under the assumption of additive costs of infections in a Bayesian framework (Ch. 5). In this framework the repertoire optimally infers

the statistics of the pathogen environment from its infection history. The theory predicts features of immunological memory in the adaptive immune system and provides a way to quantify the benefit of memory as a function of the relevant parameters. Our results link the amount of memory production upon infection to the variability of the pathogenic environment, and the decay of memory to the time scale of pathogens. We argue that the observed usefulness of immunological memory implies that the relevant features the immune system is trying to learn are highly sparse in the space of possible antigens.

Biological organisms have evolved a wide range of immune mechanisms to defend themselves against pathogens. Beyond molecular details, these mechanisms differ in how protection is acquired, processed and passed on to subsequent generations – differences that may be essential to long-term survival. We have extended a model for evolution in fluctuating environments [117] to include features specific to immunity (Ch. 6). We analyzed which adaptive strategies maximize long-term growth rates in the model to get insights into the evolution of diverse immune strategies. We find that the two key determinants of an optimal immune strategy are the frequency and the characteristic timescale of the pathogens. Depending on these two parameters, our framework identifies distinct modes of immunity, including adaptive, innate, bet-hedging and CRISPR-like immunities, which recapitulate the diversity of natural immune systems. The resulting phase diagram of optimal immune strategies might be used to guide future research into so-far unidentified or underappreciated immune mechanisms.

Prompted by our results on the distinct optimal immune strategies we have shown more generally the existence of transitions between adaptive strategies in fluctuating environments as environmental frequencies and fitness landscapes are varied (Ch. 7). In particular, we provide novel analytical results for transitions between switching and non-switching strategies in temporally correlated environments. Our results show that temporal correlations between environmental conditions favor evolutionary tracking of the environmental state by specialist phenotypes over more generalist phenotypes.

8.2 IDEAS FOR FUTURE RESEARCH

8.2.1 *Detour: Analysis of repertoire data*

A strong motivation of developing theoretical models of biological systems is to inspire novel ways of analyzing data. Theory can help make progress even without immediate direct tests if it helps sharpen intuition and provides inspiration of novel ways of analyzing data. By turning verbal hypotheses into precise mathematical models one is forced to make assumptions explicit and develop tools to analyze these mathematical models. Method development is often easier on the synthetic data coming from models, where there is full control over the process that generates data. The analysis method might then be adapted to the analysis of real data. To illustrate this assertion we present some preliminary results of an analysis of immune repertoire sequencing data based on ideas emanating from the work presented in Part ii.

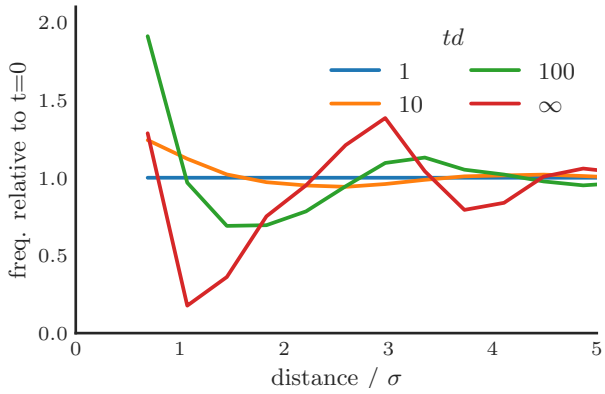

Figure 8.1: Radial distribution functions relative to the initial distribution for different times t in a simulation of the competitive mean-field dynamics introduced in Chapter 4 ($t = \infty$ from optimization). Parameters: $A(x) = 1/(1 + x^2)$, death rates $d = 0.001$, discretization 0.5σ, two dimensional recognition space of length 20σ with periodic boundary conditions. Pathogen (initial repertoire) frequencies correlated in shape space were generated as in Fig. 4.4 with coefficient of variation 1.0 (0.2) and correlation length 5σ (5σ).

Recall that in the adaptive immune system the composition of the repertoire of lymphocyte receptors determines immune protection against diverse pathogens. The repertoire is shaped by the past population dynamics of lymphocytes. There is an interesting possibility to study lymphocyte dynamics indirectly based on single time point sequencing data by looking for signatures of a particular dynamics in the repertoire statistics. We have shown that competitive proliferation upon antigenic stimulation is a candidate dynamical principle that allows a repertoire to self-organize into a state of good antigenic space coverage (see Part ii). Such a dynamics results in characteristic signatures in the distribution of receptor sequences in recognition space. For large times compared to the lifetime of lymphocytes competitive exclusion leads to a dip in the radial distribution at a distance of the order of the typical cross-reactivity width (Fig. 8.1). For small times the dominant effect is convergent selection by the co-variation of stimulation by antigens rather than competition, leading to an increased correlation in abundances for small distances. Such an effect was previously observed e.g. in the context of microbiota species abundances [211]. Here, we take a first step towards constructing an analogous measure for real data to determine whether any similar signatures of either co-stimulation or competition are present.

We use human immune repertoire sequencing data from Pogorelyy et al. [212]. In the study the hypervariable region of T cell receptors from peripheral blood cells of healthy human adults were sequenced using the Illumina HiSeq platform. In the experimental protocol randomly generated identifiers are inserted as barcodes before sequencing, to allow reliable determination of relative frequencies of different clones. The data were preprocessed as described in Pogorelyy et al. [212]. The analysis was done on productive α-chain sequences without separation into T cell subsets (CD4/CD8 or naive/memory).

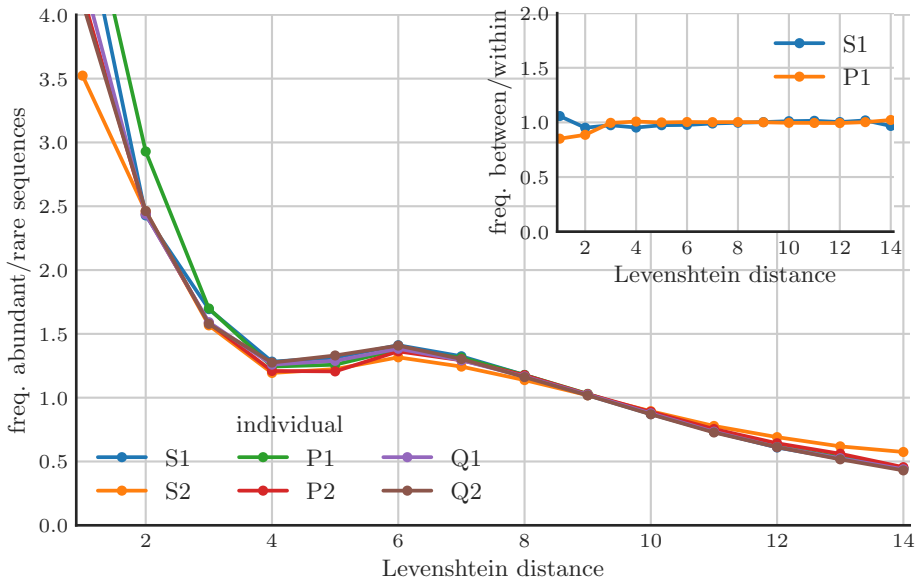

Figure 8.2: Normalized histogram of Levenshtein distances of randomly drawn pairs of sequences of abundant vs. rare sequences in real data. Inset shows normalized histogram of distances between abundant sequences between two individuals vs. within the individuals. Abundant sequences were defined as having at least three counts. Individuals starting with the same letter are twins.

To test for an antigen-dependent competition effect we must define a proxy for the functional distance of receptors. The sequence-to-function map is complicated but we can try and use a simple metric capturing some of the relevant features of cross-reactive binding. As a first step we here use the Levenshtein distance between the CDR3 amino acid sequences. The Levenshtein distance is defined as the minimal number of single-character edits (deletions, insertions, and substitutions) required to change one sequence into another.

To overcome the issue of normalization we compare the correlations found between abundant sequences relative to those in rarer sequences for which we expect to see less effects of competition. The analysis shows an increase in close-by abundant as compared to rare sequences and there is a hint of a small dip in normalized frequency at a distance of around 4 edits (Fig. 8.2). Abundant sequences are overall more similar to each other as shown by the decreased frequency of the largest distances. No such difference is observed when comparing distances of abundant sequences between individuals relative to within individuals (Fig. 8.2 inset). The data is compatible with a picture of antigen-driven proliferation based not on the unique infection history but on universal antigenic signatures shared accross individuals of e.g self or microbiotal origin. Convergent selection seems to be pre-dominant, although the small dip might be a possible signature of some degree of competition. Note that this is a preliminary analysis and that further checks are needed to rule out an influence of confounding effects and to establish the robustness of the analysis procedure. Possible confounding effects include the possible influ-

ence of differing origins of abundant and rare sequences, sequencing errors, and the preprocessing pipeline. If robust and strong signatures of competition based on antigen-binding indeed exist there is an exciting possibility of learning something about what determines cross-reactivity by comparing different proxies for functional distance.

8.2.2 *Some other ideas*

Our thinking about repertoires in this thesis has been informed by a view of repertoires as statistical ensembles [26]. This view emphasizes that protection is driven not by a core repertoire of receptors shared among most individuals but emerges from the collection of specificities in a particular individual. In this view it is natural to search for the effects of antigen-driven selection on the level of statistical signatures such as proposed in the previous section. Such statistical signatures might provide robust markers that transcend the stochasticity of individual responses to antigens. There is an interesting expansion of this view to the question of self/non-self recognition. Can we think of the self and pathogen proteome as statistical ensembles and analyze self/non-self discrimination as the statistical distinguishability of the ensembles? To answer this question we need to take into the account the rules of what fragments of proteins can serve as antigens. Then using known information about self and foreign proteomes it should be possible to map the two distributions in antigenic space.

This idea has emerged from discussions with Quentin Marcou.

On the road towards a direct confrontation of the theoretical predictions about immune repertoires (Part ii) with experiments a few major hurdles need to be cleared. Specifically, in repertoire sequencing studies the knowledge of the pathogen history often is lacking or incomplete, and despite much progress our knowledge of the determinants of cross-reactive binding remains incomplete. I discuss some current developments and future ideas that might help surmount these particular hurdles below.

Some partial information about the pathogen history is starting to become available. This knowledge pertains to particular perturbations to the repertoire such as arising from immunization [60] or chronic infection with cytomegalovirus [58]. Germ-free mice have long been studied in immunology, but going forward an intermediate model system which is neither germ-free nor has the full complexity of pathogenic history in uncontrolled settings would be very useful. Studying the long-term immune repertoire dynamics under controlled pathogen exposure might provide rich data to compare with our theoretical predictions. Alternatively – although currently this remains far-fetched – frequent longitudinal sequencing might in principle allow the reconstruction of the infection history under natural conditions.

To work towards a functional view of the immune protection afforded by different repertoires requires a better link between the simple shape space picture and the underlying sequence space. Inspiration about how to make progress on understanding cross-reactivity might be found in the approaches used in other biological systems in which specific but cross-reactive binding underlies biological function (Fig. 8.3): the nose differentiates odors by binding

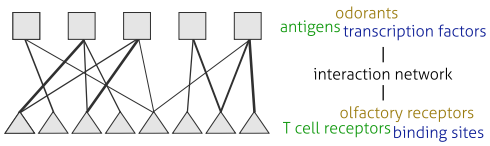

Figure 8.3: Specific recognition between molecules from two classes is necessary to achieve various biological functions. There is no unique specificity in any of these systems but some degree of cross-reactivity. Due to their shared features inspiration for how to better describe cross-reactive binding between immune receptors and antigens might be found in these other systems.

of odorants by olfactory receptors, cells regulate gene expression by binding of transcription factors to sites on the DNA, and so on. Across these biological systems interaction patterns are significantly correlated, i.e. two elements from the same class are likely to share more interaction partners than expected from independent connections. The question of how to find relevant mesoscopic descriptions of the interaction that reproduce these correlations is thus shared across systems.

A particular promising avenue in immunology is to use data from new single cell methods that enable sequencing lymphocyte receptors with known affinity to particular epitopes [213] or to pair sequencing with affinity measurements [214]. By inferring statistical physics based models of cross-reactivity landscapes from this data deeper insights into the molecular determinants of specific binding and the structure of the cross-reactivity network can be obtained.

As a complementary avenue, progress may be made by connecting the geometry of the cross-reactivity network to macroscopic experimental observables. Concretely, this suggests extending the kind of analysis proposed in Sec. 8.2.1 to find signatures of the ecological dynamics of lymphocytes under different assumptions on the geometry of cross-reactivity space in macroscopic observables such as clone size distributions and cross-correlations in abundances of different receptors. Finally, insights might be obtained from an analysis of the constraints functional requirements put on the patterns of cross-reactivity. Cross-reactivity needs to both enable specificity (to avoid recognition of self) and the ability to cover a large space of possible pathogens with a limited number of cells. Which structures of cross-reactivity networks are well-suited to fulfill these dual requirements may be analyzed by exploiting analogies to high-dimensional packing problems. Furthermore one might analyze how the features of cross-reactivity influence the efficiency of the (co-)evolution of receptors relative to pathogenic antigens.

Our results on transitions in adaptive strategies in fluctuating environment may be tested by means of experimental evolution. Varying selection strengths in experimental evolution studies like those performed by Beaumont et al. [196] could provide a way to test the conditional optimality of switching strategies. In the last years many studies have challenged old dogmas about the evolution of immunity, from the discovery of an adaptive immune system

in jawless vertebrates to the memory characteristics of many innate defense mechanisms. Our phase diagram of optimal immune strategies (Ch. 6) could help guide the discovery of further overlooked immune defense mechanisms. To give an example, vertebrates have on average longer lifespans than other organisms and thus according to our theory should be particularly prone to use adaptive immunity. However, there are also non-vertebrate organisms with long lifespans: Some molluscs for example, which are invertebrates, can live for several centuries [215]. Based on our theoretical work we might hypothesize that they have evolved alternative adaptive or proto-adaptive immune defense mechanisms. While certainly much contingency is to be expected in the evolution of a system so complex as immune defense, maybe thinking in terms of adaptive strategies will provide a fruitful conceptual framework for such exploration.

Part IV

APPENDIX

A

HOW A WELL-ADAPTED IMMUNE SYSTEM IS ORGANIZED

A.1 PROBABILITY DISTRIBUTION OF THE TIME OF FIRST RECOGNITION

In order to calculate the cost of not-recognizing an antigen a, we need to find the distribution of times when a successful encounter takes place. The probability of having the first recognition of antigen a by receptor r in the time between t and $t + dt$ reads:

$$H_a(t)dt = \lambda_a(t)dt \cdot \sum_r P_r f_{r,a} \times \lim_{N\to\infty} \prod_{i=1}^{N} \left(1 - \lambda_a(t_i)\frac{t}{N}\sum_r P_r f_{r,a}\right),$$

where the first term is the probability of having an encounter between t and $t + dt$, the second the probability of this encounter being successful, and the third the probability of there not being any prior recognition events. For the calculation of the last term we have decomposed the time leading up to t into N intervals of length t/N. Taking the $N \to \infty$ limit yields:

$$H_a(t) = \lambda_a(t)\tilde{P}_a e^{-\int_0^t dt'\lambda_a(t')\tilde{P}_a}, \tag{A.1}$$

where we have used the short-hand notation $\tilde{P}_a = \sum_r P_r f_{r,a}$ for the probability that a randomly chosen receptor recognizes antigen a.

A.2 CONVEXITY OF THE EXPECTED COST

In this Appendix we show that the cost function $\langle F \rangle$ is a convex function of its argument $\{P_r\}$ (the receptor distribution). We start by introducing an alternative expression of $\bar{F}_a$, obtained by integration by parts:

$$\bar{F}_a = \int_0^\infty dm F'_a(m)e^{-m\tilde{P}_a} + F(0). \tag{A.2}$$

We calculate the derivatives of this average cost with respect to $\tilde{P}_a$:

$$\frac{d\bar{F}_a}{d\tilde{P}_a} = -\int_0^\infty dm\, m F'_a(m)e^{-m\tilde{P}_a} \tag{A.3}$$

$$\frac{d^2\bar{F}_a}{d\tilde{P}_a^2} = \int_0^\infty dm\, m^2 F'_a(m)e^{-m\tilde{P}_a} \tag{A.4}$$

Since by assumption $F'_a(m)$ is positive, the second derivative of $\bar{F}_a$ with respect to $\tilde{P}_a$ is positive. This establishes the convexity of $\bar{F}_a$ as a function of $\tilde{P}_a$. Since

$\langle F \rangle = \sum_a Q_a \bar{F}_a$ (with $Q_a \geqslant 0$), it is a convex function of $\{\tilde{P}_a\}$. Therefore it is also a convex function of $\{P_r\}$, as $\{P_r\}$ and $\{\tilde{P}_a\}$ are linearly related.

A.3 BIOLOGICAL MOTIVATION OF POWER-LAW COST FUNCTIONS

In the main text we have developed a general framework for discussing the antigen-receptor recognition process. To fully specify the model we need to choose an effective cost function $F_a(m) = F_a(t_a(m))$. In the main text we derive optimal receptor distributions for a number of effective cost functions, including power-law functions $F(m) = m^\alpha$. Here we sketch plausible scenarios motivating that choice.

Consider an organism being infected with a antigen a. As long as there is no immune reaction, the antigens divide inside its host and thus increase its population size. If the initial population size is small it is reasonable to assume exponential growth.

The more antigens there are at the time of the immune reaction the more damage they can potentially do. Likewise, the more antigens, the higher the rate of encounters. These two quantities are also expected to grow exponentially in time:

$$F_a(t) = F_a(0)e^{\nu_a t}, \tag{A.5}$$

$$\lambda_a(t) = \lambda_a(0)e^{\nu'_a t} \tag{A.6}$$

The two exponents may be different in general, because the number of pathogenic agents that cause the harm may grow differently than the number of antigens that can be recognized by the immune system. This difference could for example stem from the fact that both the pathogen's antigenic exposure and its virulence are cooperative effects, and thus scale as a power of the number of invading individuals. Using $m_a(t) = \lambda_a(0)(e^{\nu'_a t} - 1)/\nu'_a$, and eliminating time $t \approx \ln[m_a/\lambda_a(0)]/\nu'_a$ (for t large compared to $1/\nu'_a$), we rewrite the effective cost function in terms of the number of encounters:

$$F_a(m) = F_a(0)\left(\frac{m}{\lambda_a(0)}\right)^{\frac{\nu_a}{\nu'_a}} \propto m^\alpha, \tag{A.7}$$

with $\alpha = \nu_a/\nu'_a$.

A.4 ANALYTICAL OPTIMIZATION

A.4.1 *Optimality conditions*

In the following we give optimality conditions for the optimization problem defined in the main text, which will be used for the following analytical determination of optimal receptor distributions. These conditions, called Karush-Kuhn-Tucker conditions [91], are derived from a generalization of the method of Lagrange multipliers to inequality as well as equality constraints.

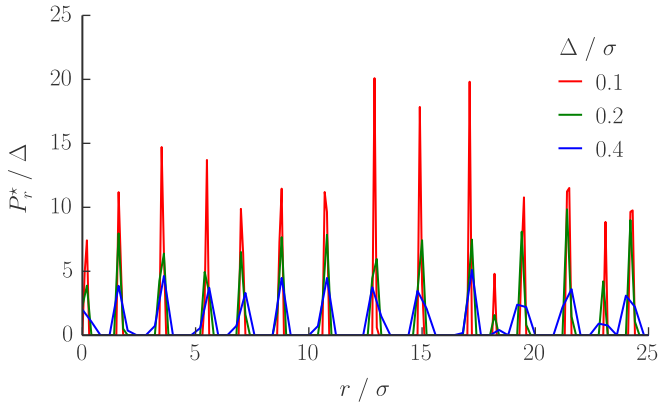

Figure A.1: Solving the optimization problem with a finer and finer discretization step suggests that the peaks found in the optimal receptor distributions converge to true Dirac delta functions. Starting from a problem with a discretization step of $\Delta = 0.1\sigma$, we construct coarse-grained versions of it by downsampling the antigen distribution two and four fold, yielding $\Delta = 0.2\sigma$ and 0.4σ respectively. The resulting coarse-grained optimization problems are then solved, and the optimal distributions P_r^*/Δ represented (after appropriate normalization by the step size). The random antigen distribution is log-normal with coefficient of variation $\kappa = 0.25$.

The Lagrangian for the optimization problem is

$$\mathcal{L}(P, \lambda, \nu) = \langle F \rangle (P) + \lambda \left(\sum_r P_r - 1 \right) - \sum_r \nu_r P_r, \tag{A.8}$$

with

$$\langle F \rangle = \sum_a Q_a \bar{F}_a. \tag{A.9}$$

λ is a Lagrange multiplier enforcing the normalization constraint and ν_r are Lagrange multipliers enforcing the non-negativity constraint. The optimal P^* is an extremum of this Lagrangian. Therefore the stationarity conditions:

$$\left. \frac{\partial \langle F \rangle}{\partial P_r} \right|_{P^*} + \lambda^* - \nu_r^* = 0, \tag{A.10}$$

with

$$\frac{\partial \langle F \rangle}{\partial P_r} = \sum_a Q_a \bar{F}'_a(\tilde{P}_a) f_{r,a}, \tag{A.11}$$

must hold for some value of λ^* and ν_r^* that enforce the constraints. The inequality constraint $P_r \geqslant 0$ further requires that:

$$\nu_r^* \geqslant 0 \tag{A.12}$$

$$\nu_r^* P_r^* = 0, \tag{A.13}$$

where the second condition is known as the complementary slackness condition. It requires the Lagrange multipliers associated with the non-negativity to be zero unless the constraint is active, *i.e.* unless the corresponding receptor probability is zero.

The three conditions may be reformulated as:

$$\left.\frac{\partial \langle F \rangle}{\partial P_r}\right|_{P^*} + \lambda^* \geqslant 0 \tag{A.14}$$

$$\left(\left.\frac{\partial \langle F \rangle}{\partial P_r}\right|_{P^*} + \lambda^*\right) P_r = 0 \tag{A.15}$$

For all receptors that are present in the optimal repertoire ($P_r^* > 0$) these conditions imply

$$\left.\frac{\partial \langle F \rangle}{\partial P_r}\right|_{P^*} = -\lambda^*. \tag{A.16}$$

If a receptor is not present in the optimal repertoire ($P_r^* = 0$) then the less stringent condition holds:

$$\left.\frac{\partial \langle F \rangle}{\partial P_r}\right|_{P^*} \geqslant -\lambda^*. \tag{A.17}$$

We note here that $\partial \langle F \rangle / \partial P_r \leqslant 0$ (because more receptors always yield a lower cost), so that $\lambda^* \geqslant 0$.

These two conditions can be explained as follows: if a repertoire is optimal, all changes allowed by the constraints will lead to a higher cost, *i.e.* moving receptors from one type to another will not yield an improvement. All partial derivatives of the cost with respect to the receptor probabilities should thus be equal to the same value (Eq. A.16). If there are already no receptors of a certain type, *i.e.* $P_r = 0$, we get a less stringent condition. We can no longer remove receptors away from this type r, but only add some to it, at the expense of other receptor types. The increase in cost due to the depletion of these other types should be higher than the gain of moving them to type r. The partial derivatives of the cost with respect to the receptors that are not present in the repertoire must thus be larger than the partial derivatives of the present receptors, which are given by $-\lambda^*$ (Eq. A.17).

A.4.2 *Solution for uniquely specific receptors*

We now solve Eqs. A.16 and A.17 for a repertoire of uniquely specific receptors (no cross-reactivity). Eq. A.11 becomes

$$\frac{\partial \langle F \rangle}{\partial P_r} = Q_r \bar{F}_r'(P_r), \tag{A.18}$$

$F(m)$	$\bar{F}(\tilde{P}_a)$	$h(x)$
m^α	$\Gamma(1+\alpha)/\tilde{P}_a^\alpha$	$(-x/(\alpha\Gamma(1+\alpha)))^{\frac{1}{1+\alpha}}$
$\ln m$	$\gamma - \ln \tilde{P}_a$	$-1/x$
$1-\exp(-\beta m)$	$\beta/\left(\beta+\tilde{P}_a\right)$	$\sqrt{-\beta/x}-\beta$
$\Theta(m-m_0)$	$\exp\left(-m_0\tilde{P}_a\right)$	$-\ln(-x/m_0)/m_0$

Table A.1: Intermediate results in the derivation of the optimal solution. The first column shows several choices of the effective cost function, $F(m)$. For these cost functions the second column shows the average cost of a pathogenic attack, $\bar{F}(\tilde{P}_a)$, and the third column shows the inverse of its derivative, $h = \left(\bar{F}'\right)^{-1}$. Γ is the Gamma function, γ is Euler's constant, β and m_0 are positive constants.

where we have used the fact that in the absence of cross-reactivity $\tilde{P}_a = P_a$. If all optimal receptor probabilities are positive then we can insert this relationship into Eq. A.16 to obtain

$$Q_r \bar{F}_r'\left(P_r^*\right) = -\lambda^*. \tag{A.19}$$

and thus:

$$P_r^* = h_r\left(-\lambda^*/Q_r\right), \tag{A.20}$$

where $h_r = \bar{F}_r'^{(-1)}$ denotes the inverse function of $\bar{F}_r'$. Since that function $\bar{F}_r'$ is always negative, h_r must take a negative argument.

For some cost functions, solving this equation may yield some negative receptor probabilities. In these cases some of the non-negativity constraints need to be active. Setting $P_r = 0$ when Eq. A.20 is negative yields the correct optimal distribution under the non-negativity constraint. We verify that for these r, Eq. A.17 is satisfied by $P_r = 0$, because:

$$Q_r \bar{F}_r'\left(P_r = 0\right) \geqslant Q_r \bar{F}_r'[h_r\left(-\lambda^*/Q_r\right)] = -\lambda^*, \tag{A.21}$$

where we have used the fact that $\bar{F}_r'$ is an increasing function of its argument (due to the positivity of its derivative, *cf.* Eq. A.4), and $h_r\left(-\lambda^*/Q_r\right) \leqslant 0$

In summary, the solution to the optimization problem is

$$P_r^* = \max\{h_r\left(-\lambda^*/Q_r\right), 0\}, \tag{A.22}$$

where the value of λ^* is fixed by the normalization condition $\sum_r P_r = 1$.

In Table A.1 we give the explicit expressions of $\bar{F}_a$ and h_a, for the particular choices of the cost function $F(m)$ considered in the main text.

A.4.3 *Solution for cross-reactive receptors*

The previous results can be generalized to cross-reactive receptors in a continuous space, using Fourier transforms. This generalization will lead up to the

results presented in the *Cross-reactivity dramatically limits optimal repertoire diversity* section of the main text, and notably the Gaussian case discussed therein.

A.4.3.1 *Deconvoluting the optimality conditions in Fourier space*

We consider a continuous receptor-antigen space and we assume a translation invariant cross-reactivity function $f_{r,a} = f(r-a)$. We write the optimality condition Eq. A.16

$$\int da\, Q(a)\bar{F}'\left[\tilde{P}^*(a)\right] f(r-a) = -\lambda^*, \tag{A.23}$$

where in continuous space the coverage is defined as:

$$\tilde{P}(a) = \int dr\, P(r) f(r-a). \tag{A.24}$$

We notice that both expressions involve integrals, which are convolutions with the cross-reactivity kernel. Since the convolution of a constant is also a constant, a solution of

$$Q(a)\bar{F}'\left(\tilde{P}^*(a)\right) = -\lambda', \quad \text{with } \lambda' > 0, \tag{A.25}$$

is also a solution of Eq. A.23. As in the case of uniquely specific receptors, we can solve this equation for $\tilde{P}^*(a)$:

$$\tilde{P}^*(a) = h\left[-\lambda'/Q(a)\right], \tag{A.26}$$

where $h = \bar{F}'^{(-1)}$ as in A.20. If there was no cross-reactivity, there would be no difference between P and $\tilde{P}$, and we would be done. Here we need to perform a deconvolution to obtain the optimal receptor distribution P from the optimal coverage $\tilde{P}$. We do so in Fourier space, where the convolution turns into a product. Deconvolution is therefore much simpler in Fourier space as it corresponds to a division

$$\mathcal{F}[\tilde{P}] = \mathcal{F}[P]\mathcal{F}[f] \quad \Leftrightarrow \quad \mathcal{F}[P] = \mathcal{F}[\tilde{P}]/\mathcal{F}[f], \tag{A.27}$$

where we have defined the Fourier transform of a function $g(x)$ as $\mathcal{F}[g](k) = \int_{-\infty}^{\infty} dx g(x) e^{ikx}$. To calculate the optimal receptor distribution we insert Eq. A.26 into Eq. A.27 and perform an inverse Fourier transform $\mathcal{F}^{-1}[\tilde{g}](x) = (1/2\pi)\int_{-\infty}^{\infty} dk \tilde{g}(k) e^{-ikx}$ to obtain

$$P^* = \mathcal{F}^{-1}\left[\mathcal{F}[h\left(-\lambda'/Q\right)]/\mathcal{F}[f]\right]. \tag{A.28}$$

This result is only valid as long as the above quantity is positive and normalizable, as we shall see below.

A.4.3.2 *The Gaussian case*

In this section we apply the general results of the previous section to a concrete example. In order to find the optimal receptor distribution analytically we use

Eq. A.28, we assume the antigen distribution and cross-reactivity function are Gaussian

$$Q(a) = \frac{1}{\sqrt{2\pi\sigma_Q^2}} \exp\left(-a^2/2\sigma_Q^2\right), \tag{A.29}$$

$$f(r-a) = \exp\left[-(r-a)^2/2\sigma^2\right], \tag{A.30}$$

and we take

$$F(m) = m^\alpha. \tag{A.31}$$

Inserting h from Tab. A.1 into Eq. A.28 allows us to write

$$P^* \propto \mathcal{F}^{-1}\left[\mathcal{F}[Q^{\frac{1}{1+\alpha}}]/\mathcal{F}[f]\right] \tag{A.32}$$

as an equivalent equation determining the optimal repertoire. We can calculate the modified antigen distribution as

$$Q(a)^{\frac{1}{1+\alpha}} \propto \exp\left(-\frac{a^2}{2(1+\alpha)\sigma_Q^2}\right). \tag{A.33}$$

The Fourier transform of a Gaussian function of variance σ^2 is a Gaussian function of variance $1/\sigma^2$ [216]. Therefore we have

$$\mathcal{F}[Q^{\frac{1}{1+\alpha}}](q) \propto \exp\left[-(1+\alpha)\sigma_Q^2 q^2/2\right], \tag{A.34}$$

$$\mathcal{F}[f](q) \propto \exp\left[-\sigma^2 a^2/2\right], \tag{A.35}$$

from which

$$\mathcal{F}[Q^{\frac{1}{1+\alpha}}]/\mathcal{F}[f] \propto \exp\left\{-[(1+\alpha)\sigma_Q^2 - \sigma^2]q^2/2\right\} \tag{A.36}$$

follows. Taking the inverse Fourier transform and normalizing, we obtain

$$P^*(r) = \frac{1}{\sqrt{2\pi[(1+\alpha)\sigma_Q^2 - \sigma^2]}} \exp\left(-\frac{r^2}{2[(1+\alpha)\sigma_Q^2 - \sigma^2]}\right). \tag{A.37}$$

Normalization is only possible for $\sigma < \sigma_Q\sqrt{1+\alpha} \equiv \sigma_c$. In the limit $\sigma \to \sigma_c$ the Gaussian converges to a Dirac delta function. Intuition suggests that a Dirac delta function centered on the peak position should remain optimal for further increases in σ. To prove this assertion we note that a Dirac delta function is zero everywhere, except in one point. Since all but one receptor probabilities are at the boundary defined by the non-negativity constraints,

we only need to check Eq. A.17. We compute the left-hand side of Eq. A.23 as a function of r

$$\int dp\ Q(a)\bar{F}'[\tilde{P}^*(a)]f(r-a) \propto -\exp\left\{\frac{-r^2[\sigma^2-(1+\alpha)\sigma_Q^2]}{2\sigma^2(\sigma^2-\alpha\sigma_Q^2)}\right\}, \tag{A.38}$$

and note that it has a minimum for $r = 0$. This shows that the partial derivatives of the expected cost at $r \neq 0$ are greater than at $r = 0$, implying that Eq. A.17 holds.

The cost of the optimal repertoires as a function of the cross-reactivity width σ is given by

$$\langle F\rangle(P^*) = \left(\frac{\sigma_Q}{\sigma}\right)^{\alpha} \begin{cases} (1+\alpha)^{\frac{1+\alpha}{2}} & \text{if } \sigma < \sigma_c, \\ \frac{(\sigma/\sigma_Q)^{\alpha}}{\sqrt{1-\alpha(\sigma_Q/\sigma)^2}} & \text{otherwise.} \end{cases} \tag{A.39}$$

Both expressions give the same cost at the transition $\sigma = \sigma_c$. After multiplying by $(\sigma/\sigma_Q)^{\alpha}$ to compare at constant recognition capability $\int f = \sqrt{2\pi}\sigma$, this expression is constant for $\sigma < \sigma_c$, and grows for $\sigma > \sigma_c$.

A.4.3.3 *General argument for peakedness*

A simple argument can help understand why cross-reactivity generically leads to peaked optimal solutions. The convolution with a kernel is a smoothening operation, represented by a low-pass filter in the Fourier domain. The optimal solution in the absence of the non-negativity constraints requires that $\tilde{P}_a = h(Q_a)$. As $\tilde{P}_a$ is the low-passed filtered version of P_r, the high-frequency components of $h(Q_a)$ will be magnified by the deconvolution. These high-frequency wiggles can lead to negative values of $\mathcal{F}^{-1}[h(Q_a)]$, which are not allowed, leading to set many values of $P(r)$ to zero. This effects results in a peaked solution. Because the size of the cross-reactivity kernel is inversely proportional to the cutoff frequency in the Fourier domain, we expect the spacing of the peaks to be related to the size of the cross-reactivity kernel.

A.5 NUMERICAL OPTIMIZATION

We numerically minimize the cost function subject to the normalization and non-negativity constraints by using a fast projected gradient algorithm as described in Sec. 3.1.3. For all reported numerical results we have chosen $\epsilon = 10^{-8}$. To minimize finite size effects in the simulations we have used periodic boundary conditions in the receptor/antigen space.

The discretization steps used in the figures are listed below:

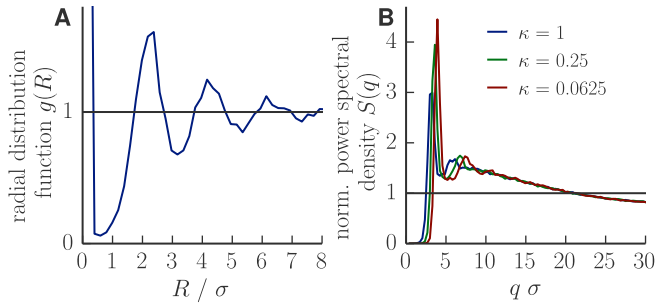

Figure A.2: Radial distribution function and normalized power spectral density of the optimal receptor distribution P_r^* for random environments in two dimensions. (A) The radial distribution function of P_r^* shows an exclusion zone around each peak, followed by oscillations characteristic of a local tiling pattern. (B) Normalized power spectral density $S(q)$ of P_r^* for different values of the parameter κ quantifying the heterogeneity of the antigenic landscape. The high suppression of fluctuations at large scales (small q) indicates that the pattern has very little fluctuations in the number of receptors used to cover large surface areas.

Step	Figure
0.5σ	5
0.1σ	3, A.3, A.4, A.6
0.05σ	4

A.6 TILING PROPERTIES: RADIAL DISTRIBUTION FUNCTION AND POWER SPECTRAL DENSITY OF THE RECEPTOR DISTRIBUTION

In this appendix we analyse the tiling structure of the peaks in the optimal distribution P_r^* found in Fig. 4.3 using the methods described in Sec. 3.6. Fig. A.2A presents the radial distribution function $g(R)$ for P^* in two dimensions. The initial drop at small r indicates that peaks in P^* are rarely close – *i.e.*, peaks in the optimal repertoire tend to repel each other. This exclusion, which operates over the range of strong cross-reactivity, is a sensible way to distribute resources, as it limits redundant protection against the same pathogens. The damped oscillation of the peaks of $g(R)$ confirm that the receptors in P^* are organized into a disordered tiling pattern. A similar radial distribution function is seen in high density random packings of hard spheres where the spheres must cover as much space as possible but exclude each other. In both cases, the tiling ensures uniform coverage of space at large scales.

To quantify the regularity of the tiling, we calculate the normalized power spectral density $S(\mathbf{q})$ of the pattern. Fig. A.2B shows $S(q)$ averaged over many realizations of the antigen landscape, and over all directions of q so that it only depends on its modulus $|q|$. $S(q)$ approaches 1 for large q, showing that the precise local positions of the peaks are random. (The small departure from 1 is attributable to numerical discretization introducing a non-Dirac shape of the peaks.) $S(q)$ is very low for small q, indicating that the number of receptors contained in any given large area of the shape space is very reproducible,

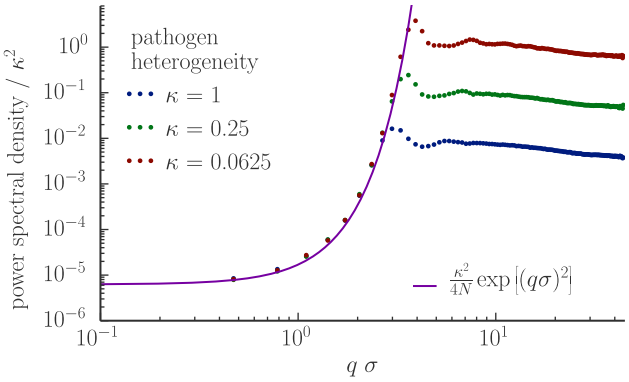

Figure A.3: Power spectral density normalized by the squared antigenic environment heterogeneity index κ: $|\sum_r P_r e^{iqr}|^2/\kappa^2$. The data collapse for different κ shows that the fluctuations at large scale are entirely attributable to those of the antigenic environment, and scale with them. At these large scales, the power spectrum of the receptor distribution is approximately given by: $\exp[(q\sigma)^2]\kappa^2/4$. The exponential term stems from the inverse of the Fourier transform of f (see Eq. A.28). In other words, the coverage of the antigenic space exactly tracks the distribution of antigens, with no additional fluctuations due to the random positioning of peaks (which would be present if this positioning was Poisson distributed). This property is called disordered hyperuniformity in the physics of jammed materials [127, 129, 130]. Parameters are the same as in Fig. 4.3.

providing uniform coverage. This behavior might be compared to disordered hyperuniform patterns which arise e.g. in the context of jammed materials [127, 129, 130]. For our optimal repertoires small scale fluctuations (large q) get smoothed out by cross-reactivity and can be tolerated, while at large scales the fluctuations track variations in the antigenic landscape to provide smooth coverage (see Fig. A.3).

A.7 NON-GAUSSIAN, LONG-TAILED, AND NON-UNIFORM CROSS-REACTIVITY FUNCTIONS

To assess the impact of different assumptions about the nature of cross-reactivity on the results we performed a number of simulations with different kernel functions.

First, we investigated the family of kernel functions defined by $f(r-a) = \exp[-(|r-a|/\eta)^\gamma]$ (Fig. A.4 left). By changing the parameter γ we can go from an exponential ($\gamma = 1$) via a Gaussian $\gamma = 2$ to a top-hat kernel ($\gamma \to \infty$). Up to $\gamma = 2$ all such kernels have positive Fourier transforms, whereas for $\gamma > 2$ the Fourier transforms also take negative values. The positive definiteness has been shown to be an important property in a related problem in ecology [157]. Second, we also investigated how long-tailed kernel functions change the optimal repertoire by considering the functional form $f(r-a) = 1/(1 + (|r-a|/\eta)^2)$ (Fig. A.4 right).

Dropping one further assumption we investigated the influence of varying the width of the cross-reactivity function between receptors (Fig. A.5). The

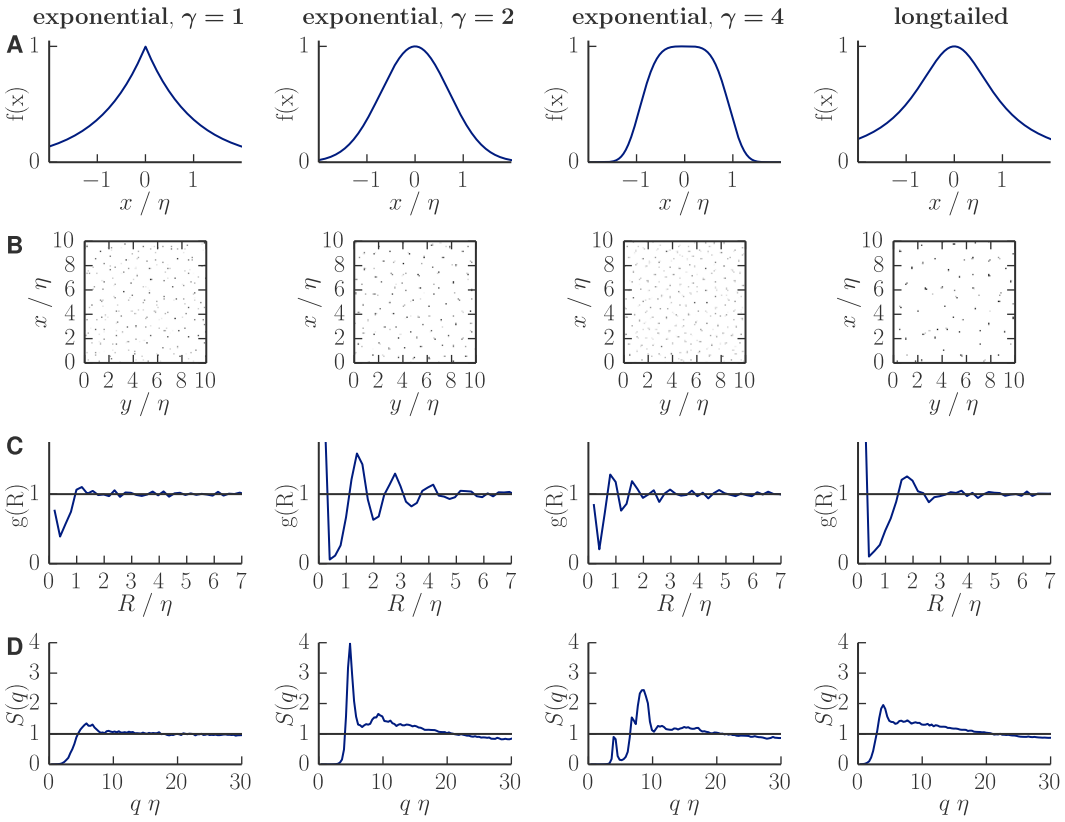

Figure A.4: Influence of the choice of the cross-reactivity kernel $f(a - r)$ on the optimization problem. Regardless of the kernel choice the optimal repertoire is peaked for non-uniform antigen distributions. The details of distribution depend on the cross-reactivity kernel. (A): Kernel functions used to describe cross-reactivity. First three column show exponential kernels of the form $f(r - a) = \exp[-(|r - a|/\eta)^{\gamma}]$ with different values of the parameter γ. Last column shows a long-tailed kernel of the form $f(r - a) = 1/(1 + (|r - a|/\eta)^2)$. (B): Examples of optimal receptor distributions in two dimensions, for antigenic environments generated as in Fig. 4.3B (with coefficient of variation $\kappa = 0.25$). (C) Radial distribution function of the optimal distribution. (D) Structure factor of the optimal distribution. The results in both (C) and (D) are averaged over 10 independent runs. A linear effective cost function $F(m) = m$ is assumed throughout.

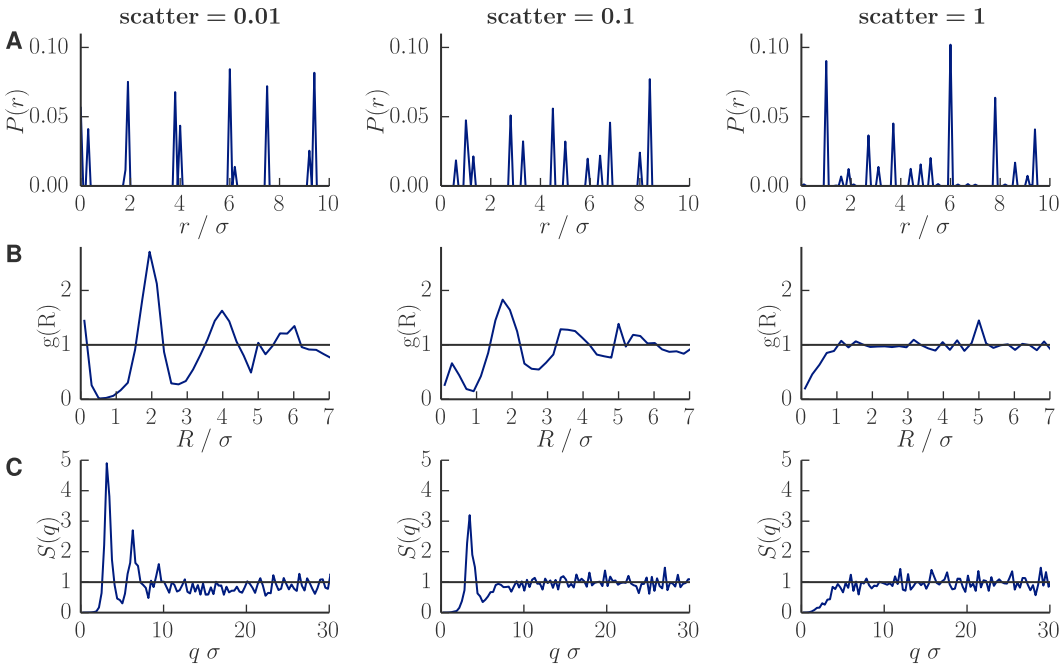

Figure A.5: Influence of varying amounts of non-uniformity in the widths of the cross-reactivity kernel on the optimization problem. (A): Examples of optimal receptor distributions in one dimensions, for antigenic environments generated as in Fig. 4.3B (with coefficient of variation $\kappa = 0.25$). A linear effective cost function $F(m) = m$ was assumed. (B) Radial distribution function of the optimal distribution. (C) Structure factor of the optimal distribution. The results in both (B) and (C) are averaged over 30 independent runs.

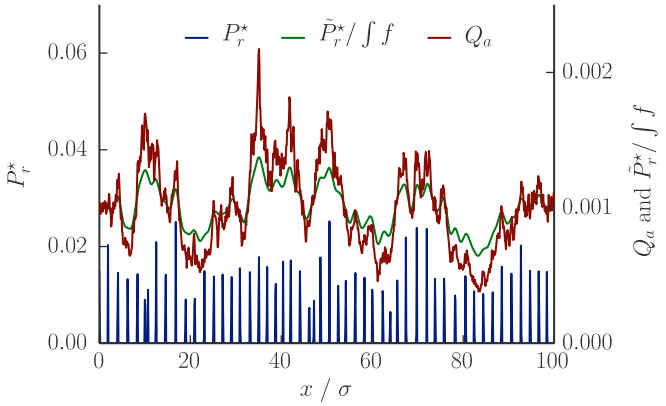

Figure A.6: Adding correlations to the antigen distribution does not change the peakedness of optimal receptor distributions. The result of the optimization is shown for a random antigen landscape with correlations. The antigen distribution is generated by Fourier filtering. First we generate an uncorrelated, normally distributed random series. This series is then filtered to obtain a power spectrum $\propto 1/(1 + (10q\sigma)^2)$. Finally, the filtered series is exponentiated to ensure the non-negativity of the generated values.

width of the cross-reactivity was drawn randomly from a log-normal distribution with different coefficients of variation (corresponding to different amounts of scatter in the width). Biologically, the overall stimulatory capacity of receptors is constrained, we rescaled the cross-reactivity so that all receptors had same overall stimulatory potency.

A.8 EXCLUDING STRONGLY SELF-BINDING RECEPTORS

The presence of self-antigens that should not be recognized puts constraints on which receptors the repertoire might contain. As a first step to understand how such an requirement interacts with the tradeoff considered in this paper we analyzed a simple model: a number of self-antigens are picked at random positions. The repertoire is not allowed to have receptors that are to highly reactive to any of the self-antigens. In practice this is assured by adding a constraint to the optimization that none of the receptors in the repertoire can have a distance smaller than σ to any self-antigen. Introducing this constraint changes the optimal repertoire, but key features such as the fragmentation of the repertoire and the tiling are conserved (Fig. A.7).

A.9 MODEL FOR RECEPTOR DYNAMICS

Here we describe our model for competitive receptor dynamics. We then show how, in a mean-field limit where antigen encounters are very frequent, this model reduces to a system of differential equations for the population dynamics.

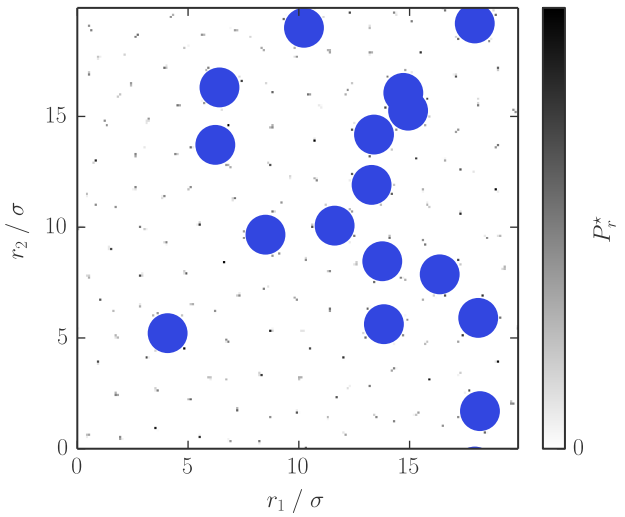

Figure A.7: Effect of exclusion zones around self-antigens on the optimization problem. No receptors are allowed in exclusion regions around self-antigens (shaded in blue). Parameters are the same as in Fig. 4.3B.

At every step we update the number of receptors according to

$$\Delta N_r = \Delta t \cdot N_r \left[A \left(\sum_{r'} N_{r'} f_{r',a} \right) f_{r,a} - d \right] \tag{A.40}$$

where the antigen a is drawn randomly with probability Q_a and Δt is a parameter determining how much the repertoire changes per step.

In the limit where Δt is small the dynamics cycle through different encountered pathogens so fast that they effectively become the following dynamics:

$$\frac{dN_r}{dt} = N_r \left[\sum_a Q_a A \left(\sum_{r'} N_{r'} f_{r',a} \right) f_{r,a} - d \right] \tag{A.41}$$

This dynamics is of mean-field type, *i.e* it neglects the effect of the stochasticity in the encounter of pathogens.

A.10 THE STABLE FIXED POINT OF THE MEAN-FIELD POPULATION DYNAMICS MINIMIZES THE COST FUNCTION

In this section we show that the stable fixed point $\{N_r^*\}$ of the population dynamics, Eq. A.41, gives a probability distribution $P_r = N_r/N_{tot}$ (with $N_{tot} = \sum_r N_r$) that minimizes the cost $\langle F \rangle$. For this correspondence to be exact, the availability function of the dynamics and the effective cost function of the optimization must be related by:

$$A(\tilde{N}_a) = -c' \bar{F}'(\tilde{N}_a / N_{st}), \tag{A.42}$$

where $\tilde{N}_a = \sum N_r f_{r,a}$, and N_{st} is the total number of receptors N_{tot} at the fixed point.

A fixed point is characterized by $dN_r/dt = 0$. If $N_r > 0$, this translates into

$$\sum_a Q_a A\left(\sum_{r'} N_{r'} f_{r',a}\right) f_{r,a} - d = 0. \tag{A.43}$$

Using the correspondence between availability and cost function given by Eq. A.42 we rewrite this condition as

$$\sum_a Q_a \bar{F}'\left(\tilde{P}_a\right) f_{r,a} = -c'd, \tag{A.44}$$

which is equivalent to the optimality condition Eq. A.16, with the identification $\lambda^* = c'd$.

For $N_r = 0$ we need to work a bit harder to show that the optimality condition at the boundary Eq. A.17 is satisfied. Here the key assumption establishing the minimization of the cost function is the stability of the fixed point. A fixed point is stable if the real parts of the Jacobian's eigenvalues are all negative. The Jacobian reads:

$$\begin{aligned} J_{r,r'} =& \delta_{r,r'}\left(\sum_a Q_a A\left(\sum_{r''} N_{r''} f_{r'',a}\right) f_{r,a} - d\right) \\ &+ N_r \sum_a Q_a A'\left(\sum_{r''} N_{r''} f_{r'',a}\right) f_{r,a} f_{r',a}. \end{aligned} \tag{A.45}$$

We remark that for $N_r = 0$ the r^{th} row of the Jacobian is non-zero only on the diagonal. That value on the diagonal is an eigenvalue of the Jacobian and must be negative:

$$\sum_a Q_a A\left(\sum_{r'} N_{r'} f_{r',a}\right) f_{r,a} - d < 0, \tag{A.46}$$

Again we replace $A\left(\sum_r N_r f_{r,a}\right)$ by $-\bar{F}'_a\left(\tilde{P}_a\right)$ according to Eq. A.42 to obtain

$$\sum_a Q_a \bar{F}'\left(\tilde{P}_a\right) f_{r,a} > -c'd, \tag{A.47}$$

which is equivalent to the optimality condition at the boundary Eq. A.17, provided that $\lambda^* = c'd$.

A.11 COST FUNCTION AS A LYAPUNOV FUNCTION OF THE MEAN-FIELD DYNAMICS

Here we show rigorously that, when the availability function is scale invariant, as in the case for the simple cost function $F(m) = m^\alpha$, the dynamics must converge towards a fixed point. This fixed point is unique and corresponds to the optimal of the cost $\langle F \rangle$, as we have shown in the previous section.

$A(x)$ is scale invariant if there exists a function ν such that $A(\gamma x) = \nu(\gamma)A(x)$. In this case we will see that the changes of relative frequencies P_r in the repertoire over time only depend on the total number of receptors through a prefactor. Below we derive the equations governing this dynamics and will then prove that this dynamics is assured to converge to a stable fixed point. We do so by showing that the dynamics admits the expected cost $\langle F \rangle$ as a Lyapunov function, *i.e.*, a function that continually decreases under the dynamics.

For ease of notation we rewrite Eq. A.41 as:

$$\frac{dN_r}{dt} = N_r[\pi_r(N) - d], \tag{A.48}$$

where N is a short-hand for $\{N_r\}$, and $\pi_r = \sum_a Q_a A\left(\sum_{r'} N_{r'} f_{r',a}\right) f_{r,a}$ is the growth rate of receptor type r. The relative frequencies $P_r = N_r/N_{tot}$ evolve according to:

$$\frac{dP_r}{dt} = \frac{1}{N_{tot}}\frac{dN_r}{dt} - \frac{N_r}{N_{tot}^2}\frac{dN_{tot}}{dt} \tag{A.49}$$

$$= P_r\left[\pi_r(N) - \sum_{r'} P_{r'}\pi_{r'}(N)\right]. \tag{A.50}$$

If A is scale invariant, so is π_r and $\pi_r(N) = \pi_r(N_{tot}P) = \nu(N_{tot})\pi_r(P)$. Then the equations further simplify to

$$\frac{dP_r}{dt} = \nu(N_{tot})P_r\left[\pi_r(P) - \sum_{r'} P_{r'}\pi_{r'}(P)\right], \tag{A.51}$$

$$= \nu(N_{tot})P_r\left(\pi_r - \bar{\pi}\right), \tag{A.52}$$

where $\bar{\pi} = \sum_r P_r\pi_r$.

We can now write how the expected cost $\langle F \rangle$ evolves in time:

$$\frac{d\langle F\rangle}{dt} = \sum_r \frac{\partial\langle F\rangle}{\partial P_r}\frac{dP_r}{dt} \tag{A.53}$$

$$= \nu(N_{tot})\sum_r P_r\left[\sum_a Q_a \bar{F}'_a(\tilde{P}_a) f_{r,a}\right](\pi_r - \bar{\pi}) \tag{A.54}$$

$$= -\frac{\nu(N_{tot})}{c'}\sum_r P_r\left[\sum_a Q_a A(N_{st}\tilde{P}_a) f_{r,a}\right](\pi_r - \bar{\pi}) \tag{A.55}$$

$$= -\frac{\nu(N_{tot})\nu(N_{st})}{c'}\sum_r P_r\pi_r\left(\pi_r - \bar{\pi}\right) \tag{A.56}$$

$$= -\frac{\nu(N_{tot})\nu(N_{st})}{c'}\sum_r P_r\left(\pi_r - \bar{\pi}\right)^2 \leqslant 0. \tag{A.57}$$

This proves that the cost always decreases with time, *i.e.* is a Lyapunov function of the dynamics. Therefore the dynamics will reach a stable fixed point at steady state, which is guaranteed to be the global minimum of the expected cost $\langle F \rangle$.

B

HOW A WELL-ADAPTING IMMUNE SYSTEM REMEMBERS

B.1 NUMERICAL METHODS

To simulate trajectories according to a neutral symmetric Wright-Fisher model in the diffusion-limit we use the Langevin equation [217]

$$\frac{dQ_a}{dt} = \frac{1}{2\tau}(\theta_0 - K\theta_0 Q_a)\,dt + \sum_{b=1}^{K} \sqrt{Q_a/\tau}\left(\delta_{ab} - \sqrt{Q_a Q_b}\right)\eta_b(t) \quad \text{(B.1)}$$

with Gaussian white noise $\langle \eta_a(t)\eta_b(t')\rangle = \delta_{a,b}\delta(t-t')$. We simulate the stochastic differential equation using the Euler-Maruyama algorithm [104]. To ensure the solution stays within the probability simplex we project the results of the iterations back onto the simplex at every iteration.

B.2 SOLVING THE DIFFUSION PATHOGEN DYNAMICS

To simplify notations we absorb τ by rescaling time by $1/\tau$. To obtain a solution to Eq. 5.15 we write $\rho(q,t)$ as the product of the steady-state distribution with a time-varying function $f(q,t)$ [104]

$$\rho(q,t) = \rho_s(q) f(q,t). \quad \text{(B.2)}$$

By substitution into Eq. 5.15 $f(q,t)$ can be shown to obey the backward equation

$$\partial_t f(q,t) = \frac{1}{2}(-\beta q + \alpha(1-q))\partial_q f(q,t) + \frac{1}{2}q(1-q)\partial_q^2 f(q,t). \quad \text{(B.3)}$$

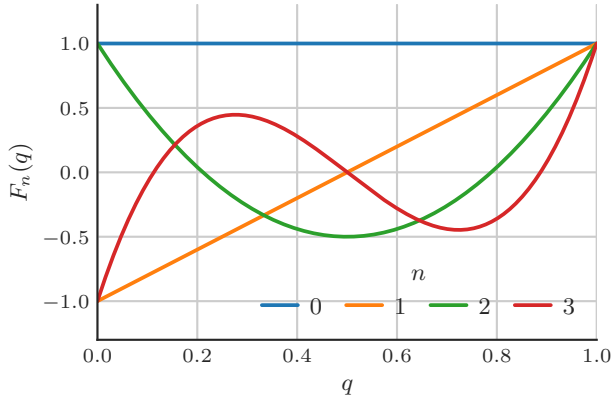

Figure B.1: First few modified Jacobi polynomials $F_n(q)$, which are the eigenfunctions pathogen dynamics propagator.

The backward equation is linear in $f(q,t)$ so its solution can be written as a superposition of eigenfunctions,

$$f(q,t) = \sum_{n=0}^{\infty} d_n(t) F_n(q) \tag{B.4}$$

with $d_n(t) = d_n(0)e^{-\lambda_n t}$. This leads to the equation

$$\frac{1}{2}q(1-q)\frac{d^2F_n}{dq^2}(q) + \frac{1}{2}(-\beta q + \alpha(1-q))\frac{dF_n}{dq}(q) = -\lambda_n F_n(q), \tag{B.5}$$

which up to a rescaling is a Jacobi differential equation. Its eigenfunctions are thus the modified Jacobi Polynomials (Fig. B.1)

$$F_n(q) = P_n^{(\beta-1,\alpha-1)}(2q-1), \tag{B.6}$$

where $P_n^{(a,b)}(x)$ is the n-th Jacobi polynomial and its eigenvalues are

$$\lambda_n = \frac{1}{2}n(n+\alpha+\beta-1). \tag{B.7}$$

As the backward equation is linear the general solution is a superposition of the eigenfunctions

For the following discussion we recapitulate some of the properties of these polynomials. The first three polynomials are

$$F_0(q) = 1, \tag{B.8}$$

$$F_1(q) = \alpha(q-1) + \beta q, \tag{B.9}$$

$$F_2(q) = \frac{1}{2}(\alpha(1+\alpha) - 2(1+\alpha)(1+\alpha+\beta)q + (1+\alpha+\beta)(2+\alpha+\beta)q^2). \tag{B.10}$$

The polynomials form an orthogonal system with respect to the weight function $q^{\alpha-1}(1-q)^{\beta-1}$, i.e.

$$\int_0^1 dq F_n(q)F_m(q)q^{\alpha-1}(1-q)^{\beta-1} = \delta_{n,m}\Delta_n(\alpha,\beta), \tag{B.11}$$

where the normalization coefficients $\Delta_n(\alpha,\beta)$ are given by

$$\Delta_n(\alpha,\beta) = \frac{\Gamma(n+\alpha)\Gamma(n+\beta)}{(2n+\alpha+\beta-1)\Gamma(n+\alpha+\beta-1)\Gamma(n+1)}. \tag{B.12}$$

Combining Eq. B.8 with Eq. B.11 we obtain

$$\int_0^1 dq F_n(q)\rho_s(q) = \delta_{n,0}, \tag{B.13}$$

where we have used the fact that the steady state distribution is a multiple of the weight function. Normalization of the probability distribution implies $d_0 = 1$ always.

For $n \geqslant 1$ the polynomials are related by the recursion formula [174]

$$qF_n(q) = c_n^- F_{n-1}(q) + c_n^0 F_n(q) + c_n^+ F_{n+1}(q). \tag{B.14}$$

with the coefficients

$$c_n^- = \frac{(n+\alpha-1)(n+\beta-1)}{(2n+\alpha+\beta-1)(2n+\alpha+\beta-2)}, \tag{B.15}$$

$$c_n^0 = \frac{1}{2} - \frac{\beta^2-\alpha^2-2(\beta-\alpha)}{2(2n+\alpha+\beta)(2n+\alpha+\beta-2)}, \tag{B.16}$$

$$c_n^+ = \frac{(n+1)(n+\alpha+\beta-1)}{(2n+\alpha+\beta)(2n+\alpha+\beta-1)}, \tag{B.17}$$

for $n \geqslant 1$ while for $n = 0$

$$qF_0(q) = \frac{\alpha}{\alpha+\beta}F_0(q) + \frac{1}{\alpha+\beta}F_1(q). \tag{B.18}$$

To calculate first moments we combine the recursion formula Eq. B.14 with the equation for the zeroth moments to obtain

$$\langle q \rangle_{\rho_s F_n} = \int_0^1 dq F_n(q)\rho_s(q)q = \begin{cases} c_0^0 = \frac{\alpha}{\alpha+\beta} & \text{for } n = 0, \\ c_1^- = \frac{\alpha\beta}{(\alpha+\beta+1)(\alpha+\beta)} & \text{for } n = 1 \\ 0 & \text{for } n \geqslant 2 \end{cases} \tag{B.19}$$

The properties are useful for analyzing recursive Bayesian estimation. We can write the prior belief as $B(q,t^-) = \rho_s(q)\sum_n d_n^- F_n(q)$ and ask how it should change if we encounter pathogen 1 at time t. (In the following we suppress the explicit notation for the dependence of these function on α and β.) Applying Bayes rule the update equation is

$$B(q,t^+) = \frac{qB(q,t^-)}{Z} = \frac{\rho_s(q)}{Z}\sum_n d_n^- qF_n(q), \tag{B.20}$$

which with the recursion formula Eq. B.14 yields

$$B(q,t^+) = \frac{\rho_s(q)}{Z}\{[c_0^0 + d_1^- c_1^-]F_0(q) \tag{B.21}$$

$$+ \sum_{n=1}^{\infty}[d_{n-1}^- c_{n-1}^+ + d_n^- c_n^0 + d_{n+1}^- c_{n+1}^-]F_n(q)\}. \tag{B.22}$$

We can write the numerator as a multiplication of $\mathbf{d}$ with the triadiagonal matrix which has c_n^0 along the diagonal, c_{n+1}^- above the diagonal, and c_{n-1}^+ below. Normalization requires $d_0^+ = 1$ so $Z = c_0^0 + d_1^- c_1^-$. Finally, we define the update coefficients for d_n as

$$\chi_n^x = c_n^x / Z \tag{B.23}$$

for x in $-, 0, +$. If pathogen 2 instead of pathogen 1 is encountered one can obtain a similar recursion equation.

The prediction equation is

$$B(q,t) = \rho_s(q)\{1 + \sum_{n=1}^{\infty} d_n(0)e^{-\lambda_n t}F_n(q)\} \tag{B.24}$$

from which with Eq. B.19 it follows that

$$\langle q\rangle(t) = \frac{\alpha}{\alpha+\beta}\left[1 + d_1(0)e^{-\lambda_1 t}\frac{\beta}{\alpha+\beta+1}\right]. \tag{B.25}$$

The prior evidence determines the value of $d_1(0)$, but the decay towards the long-term average prediction always follows a simple exponential decay with a time scale $\tau = 1/\lambda_1 = 2/(\alpha+\beta)$

B.3 SUPPLEMENTARY FIGURES

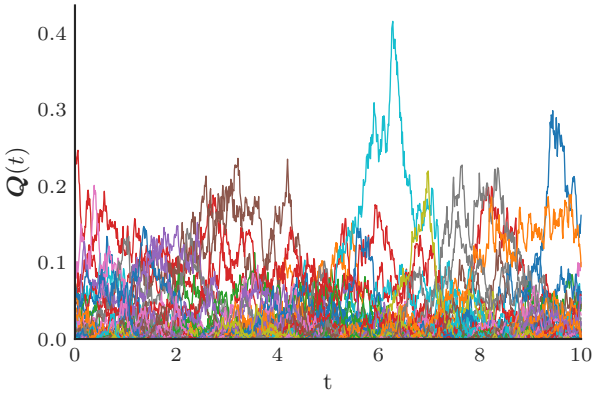

Figure B.2: Antigenic environmental dynamics according to Eq. 5.2. Lines show the frequencies Q_α of different antigens over time. Simulations were based on a Langevin description of the dynamics as described in Sec. B.1. Parameters: $K = 500$, $\theta_0 = 0.02$, $\tau = 10$.

C

IMMUNE STRATEGIES AGAINST CHANGING PATHOGENS

C.1 PARAMETRIZING A TWO-STATE MARKOV CHAIN

The parametrization that we use is based on the average frequency with which the chain is in one of its state and on a characteristic time scale of state changes. Concretely, the first parameter is the fraction of generations during which a pathogen is present $\pi_{env} = \langle x \rangle = \alpha/(\alpha+\beta)$, and the second is the autocorrelation time of the chain, defined from its auto-correlation function: $\langle x_t x_{t'} \rangle - \langle x \rangle^2 = \pi_{env}(1-\pi_{env})(1-\alpha-\beta)^{|t-t'|} \equiv \pi_{env}(1-\pi_{env})\exp(-|t-t'|/\tau_{env})$, or $\tau_{env} \equiv -1/\ln(1-\alpha-\beta)$. We have chosen the autocorrelation time over alternative time scales such as the persistence time of the pathogen $\tilde{\tau}_{env} = -1/\ln(1-\beta)$, because the autocorrelation time is symmetric in the switching rates: long stretches of continuous pathogen absence play a role in the choice of a strategy as long stretches of continuous pathogen presence. Our choice of characteristic time provides a measure for the degree of predictability of the next state given the current state. For instance, for a characteristic time of zero, $1-\alpha = \beta$, no information on the next state can be gained from knowing the current state. The parametrization has the additional property that all combinations $\pi_{env} \in [0,1]$ and $\tau_{env} \in [0,\infty)$ correspond to valid values of $\alpha \in [0,1], \beta \in [0,1]$, which is not the case for all combinations of $(\pi_{env}, \tilde{\tau}_{env})$, for example.

C.2 PATTERN-SEARCH BASED OPTIMIZATION FOR PROBLEMS WITH NOISY FUNCTION EVALUATIONS

The numerical optimization of the long-term growth rate considered in this work falls into the class of noisy optimization problems described in Sec. 3.1.1.2. The long-term growth rate can generally not be calculated explicitly but is approximated numerically from long but finite simulations of the population dynamics. These simulations depend on the history of pathogen presence and absence x_t, which is a random variable. We solve this optimization problem using the algorithm described in Sec. 3.1.3.2 The notations used there can be mapped onto the problem of optimizing long-term population growth rates through $f \leftrightarrow -\Lambda, y \leftrightarrow (p, q, p_{uptake}, c_{constitutive}), \omega \leftrightarrow x_t$. For the optimization problem considered in this paper, this algorithm works reliably and efficiently enough to allow for the many optimizations needed for a phase diagram such as shown in Fig. 2.

C.3 ANALYTICAL INSIGHT INTO THE TRANSITIONS BETWEEN STRATEGIES

By analytically solving three simplified problems, we provide additional insights into the choice of strategy. For brevity of notation, we set $c_{def} = c_{defense}$, $c_{con} = c_{constitutive}$, and $c_{inf} = c_{infected}$.

C.3.1 *When to regulate the response*

For pathogens changing with a small characteristic time scale, there is a transition from adaptive to protoadaptive to innate strategies for standard parameters (Fig. 2) as a function of π_{env}. For all three strategies considered here, the complete population is always protected, $q = 0$, and there is no active acquisition, $p_{uptake} = 0$. The equation for the instantaneous growth rate at generation t (Eq. 6.8 of the main text) thus simplifies to:

$$z^t = \begin{cases} e^{-c_{def}} & \text{if } x^t = 1 \\ e^{-c_{con}} & \text{if } x^t = 0 \end{cases} \tag{C.1}$$

where growth only depends on the absence or presence of pathogen during the current generation regardless of what happened at previous generations. The optimal long-term growth rate can then be calculated analytically by weighting the instantaneous growth rates in the presence and absence of pathogen by the frequency of the two environmental states

$$\Lambda = -\pi_{env} c_{def}(c_{con}) - (1 - \pi_{env}) c_{con}. \tag{C.2}$$

This expression for the long-term growth rate directly gives us some insight into how the frequency of pathogen affects how much the response should be regulated. The more frequent the pathogen, the more often the defense is actually used and thus the less it should be regulated. By maximizing Λ over $c_{con} \in [0, c_{con}^{max}]$ for a given trade-off function $c_{def}(c_{con})$, we obtain analytical expressions for the phase boundaries. One finds the following conditions for local optimality of the three strategies:

$\pi_{env} < \pi_{env}^{(ap)}$	$\pi_{env}^{(ap)} \leqslant \pi_{env} \leqslant \pi_{env}^{(po)}$	$\pi_{env}^{(po)} < \pi_{env}$
$c_{con} = 0$	$0 \leqslant c_{con} \leqslant c_{con}^{max}$	$c_{con} = c_{con}^{max}$

with

$$\pi_{env}^{(ap)} = \left(1 - \left.\frac{dc_{def}}{dc_{con}}\right|_{c_{con}=0}\right)^{-1}, \tag{C.3}$$

$$\pi_{env}^{(po)} = \left(1 - \left.\frac{dc_{def}}{dc_{con}}\right|_{c_{con}=c_{con}^{max}}\right)^{-1}. \tag{C.4}$$

As we assume a convex trade-off shape, we have $\pi_{env}^{(ap)} < \pi_{env}^{(pi)}$, which implies a succession of adaptive, protoadaptive and innate strategies for increasing

π_{env} as seen in Fig. 6.2. If instead the trade-off function $c_{def}(c_{con})$ is concave, then the protoadaptive phase vanishes.

C.3.2 *When to hedge your bets*

For the very frequent pathogens, the optimal strategy is to have protection at all times, whereas for less frequent pathogens some bet-hedging is often favored (Fig. 6.2 and C.2). To understand the transition from bet hedging innate to deterministic innate immunity, we compare the long-term growth rates of populations using these strategies. For simplicity, we restrict the analysis to strategies with no heritability, $p = 1 - q$, and no regulation, $c_{def} = c_{con}$. The fraction of protected individuals is constant across generations and the long-term growth rate can be calculated analytically as

$$\Lambda = \pi_{env} \ln[(1-p)e^{-c_{inf}} + pe^{-c_{con}}] + (1-\pi_{env}) \ln[1-p+pe^{-c_{con}}]. \quad (C.5)$$

Optimizing the long-term growth rate over the fraction of protected organisms p yields

$\pi_{env} < \pi_{env}^{(0i)}$	$\pi_{env}^{(0i)} \leqslant \pi_{env} \leqslant \pi_{env}^{(io)}$	$\pi_{env}^{(io)} < \pi_{env}$
$p = 0$	$0 \leqslant p \leqslant 1$	$p = 1$

with

$$\pi_{env}^{(0i)} = \frac{e^{c_{con}} - 1}{e^{c_{inf}} - 1}, \quad (C.6)$$

$$\pi_{env}^{(io)} = \frac{1 - e^{-c_{con}}}{1 - e^{-c_{inf}}}. \quad (C.7)$$

This shows the existence of three regimes. For rare pathogens tolerance is optimal (as we are only looking at unregulated strategies), for frequent pathogens it is best to always protect, while in-between bet-hedging is favored. The existence of these different phases is a known result in the bet-hedging literature when both phenotypes can survive in both environmental states [200], as is the case here. The assumption $p = 1 - q$ makes the derivation of this result exact when the environment itself is memoryless, $\alpha = 1 - \beta$. In the presence of temporal correlations in pathogen occurrence, we expect bet-hedging strategies to be favored for a larger range of pathogen frequencies, as they can exploit the predictability of the environment.

C.3.3 *When to acquire actively*

For pathogens with large temporal correlations, the optimal strategy changes from an active, to a mixed, to a passive mode of acquisition (Fig. 6.2). To understand these transitions, we again turn to an analytical solvable limit. As these strategies are favored in the presence of temporal correlations, the limit of temporally uncorrelated strategies $p = 1 - q$ considered in the previous section is not the most pertinent. We turn instead to another analytical solvable limit, in which growth rate differences are very large compared to the generation time,

$c_{con} \gg 1, c_{inf} - c_{def} \gg 1$. In this limit, the fraction of protected individuals is Markovian as all parents of individuals in the current generation were in the favored state of the last environment (all maladapted individuals die). We note that similar results can be obtained in the limit of large environmental correlation times τ_{env} without assuming completely specialized phenotypes [111]. The long-term growth rate can therefore be expressed analytically based on the probabilities Q_{ij} of observing an environmental state i followed by state j ($Q_{00} = (1-\pi_{env})(1-\alpha), Q_{01} = (1-\pi_{env})\alpha, Q_{10} = \pi_{env}\beta, Q_{11} = \pi_{env}(1-\beta)$) as

$$\begin{aligned}\Lambda =& Q_{00}\ln(1-p) + Q_{10}\ln q + Q_{01}\ln((p+p_{uptake})e^{c_{def}}) \\ &+ Q_{11}\ln((1-q)e^{c_{def}}) - c_{uptake}(p_{uptake}).\end{aligned} \tag{C.8}$$

By comparing the terms in which p and p_{uptake} appear in this expression, the strengths and weaknesses of the two acquisition modes become evident. Passive acquisition has a diversification cost due to unnecessary switching into state 1 in the absence of pathogen ($Q_{00}\ln(1-p)$). Active acquisition does not have this penalty, but is more difficult to implement and comes with an extra cost $c_{uptake}(p_{uptake})$ dependent on its uptake rate. As the probability Q_{00} is high for rare and temporally correlated pathogens, the relative cost of random acquisition is especially high for these pathogens, where most of the time mutations conferring gain of protection are deleterious. Optimizing the expression of the long-term growth rate over $p, p_{uptake} \in [0,1]$ we find the following optimality conditions:

$\pi_{env} < \pi_{env}^{(cm)}$	$\pi_{env}^{(cm)} \leqslant \pi_{env} \leqslant \pi_{env}^{(mi)}$	$\pi_{env}^{(mi)} < \pi_{env}$
$p = 0, p_{uptake} > 0$	$p > 0, p_{uptake} > 0$	$p > 0, p_{uptake} = 0$

with

$$Q_{00}(\pi_{env}^{(cm)}) = \frac{Q_{01}(\pi_{env}^{(cm)})}{g^{-1}(Q_{01}(\pi_{env}^{(cm)}))}, \quad \text{with } g(p_{uptake}) = p_{uptake}\frac{dc_{uptake}}{dp_{uptake}}, \tag{C.9}$$

$$\pi_{env}^{(mi)} = 1 - \left.\frac{dc_{uptake}}{dp_{uptake}}\right|_{p_{uptake}=0}. \tag{C.10}$$

Thus, in this limit, a CRISPR-like strategy is favored for rare pathogens, an innate bet-hedging strategy for frequent pathogens, and mixed strategies in-between, in agreement with the numerical results reported in Fig. 6.2 of the main text.

C.4 NON-INDEPENDENT PATHOGEN-PROTECTION PAIRS

The factorization of the recursion relation defining the population dynamics allows us to treat the problem one pathogen at a time. This makes the problem mathematically tractable and the results easily interpretable. Different protection-pathogen pairs can only be treated independently, however, if

a number of assumptions are met: the costs must be additive, one protection protects against only one pathogen and *vice versa*, and the dynamics of different pathogens are independent. A full treatment of the general, non-factorized problem is outside the scope of this work, but in the following we discuss how relaxing some of these assumptions affects the optimal strategy. Specifically, we consider simple cases with only two pathogen-protection pairs to build intuition of where we expect qualitative changes in optimal strategies, and where and how we can relate back to the results for the factorizing case.

C.4.1 *Non-additive cost of infection*

If the cost of an infection is amplified by co-infections by other pathogens, then we expect the optimal strategies to be similar to the ones emerging for a single pathogen, but with an higher effective cost of infection $c_{\text{infection}}$ (for the influence of a higher cost of infection on the phase diagram see Fig. C.2F). The effective cost should take into account the extra cost incurred by the presence of a co-infection weighted its probability of occurrence.

To test this intuition, we consider a simple case with two pathogens, where we impose $c_{\text{constitute}} = c_{\text{defense}}$ and $p_{\text{uptake}} = 0$. A completely unprotected organism pays a cost $c_{\text{infection}}$ if it gets infected by one pathogen, and a cost $2c_{\text{infection}} + \nu$ if it gets infected by both. Solving the problem numerically shows that the optimal fraction of protection against the two pathogens increases with ν (Fig. C.4) as expected. The Pearson correlation coefficient between being protected against one or the other pathogen remains small even for ν of the order of $c_{\text{infection}}$, meaning that the optimal strategy remains close to the independent case.

C.4.2 *Non-additive cost of protection*

As with the case of non-additive cost of infection, we expect non-additive costs of protection to result in a modified effective cost of protection (for the influence of changing the cost of protection see Fig. C.2G). However, for a non-independent cost of protection, an optimal immune strategy might differ significantly from the factorizing case. In particular, the optimal strategy may regulate the total number of protections at a given time to either exploit the economies of scale (if protection against many pathogens is relatively cheaper) or avoid an overburdening cost (if protection against many pathogens at the same time is relatively more costly).

Some of the immune strategies that require a lot of machinery to function, such as vertebrate adaptive immunity or CRISPR-Cas immunity, might come at the expense of a large fixed investment cost, c_{system}, in addition to their state dependent costs. This non-additive cost can be viewed as shared equally between all pathogen-protection pairs concerned by the adaptive strategy. It does not break the independence between them, but rather adds an offset cost c_{system}/L, where L is the number of pathogen-protection pairs, which will shift the transition at which adaptive immunity becomes favorable.

C.4.3 *Cross-reactive protection*

In most biological defense systems, there is some degree of cross-reactivity, i.e. defense against several pathogens can be achieved with the same protection. This feature can be incorporated in our framework by introducing a more complicated form of the dependency of the number of offspring on the protection state σ. We expect the optimal strategy to exploit cross-reactivity by having dissimilar protections that collectively tile the space of possible pathogens [138]. Then, the dynamics of pathogens can be effectively reduced to the presence of absence of any of the pathogens within the scope of a given protection.

To validate this intuition, we consider a single protection that is efficient against two pathogens of frequencies $\pi_{\text{env},1}$ and $\pi_{\text{env},2}$. Assume that the cost of defense is the same whether we defend against one or both pathogens, as summarized by the costs in the table below:

$\sigma\backslash(x_1,x_2)$	$(0,0)$	$(1,0)$	$(1,0)$	$(1,1)$
0	0	c_{inf}	c_{inf}	$2c_{\text{inf}}$
1	c_{con}	c_{def}	c_{def}	c_{def}

where (x_1,x_2) indicates which of the two pathogens are present. If the protection strategy is memoryless ($p = 1-q$), then the long term growth rate is

$$\begin{aligned}\Lambda =&(1-\pi_{\text{env},1})(1-\pi_{\text{env},2})\ln r_{00} + \pi_{\text{env},1}(1-\pi_{\text{env},2})\ln r_{10}\\ &+(1-\pi_{\text{env},1})\pi_{\text{env},2}\ln r_{01} + \pi_{\text{env},1}\pi_{\text{env},2}\ln r_{11},\end{aligned} \tag{C.11}$$

where $r_{x_1x_2}$ is the average growth rate in environment (x_1,x_2): $r_{00} = pe^{-c_{\text{con}}} + 1-p$, $r_{01} = r_{10} = pe^{-c_{\text{def}}} + (1-p)e^{-c_{\text{inf}}}$, $r_{11} = pe^{-c_{\text{def}}} + (1-p)e^{-2c_{\text{inf}}}$. The long term growth rate can be alternatively expressed as

$$\Lambda = (1-\pi_{\text{env,eff}})\ln r_{00} + \pi_{\text{env,eff}}\ln r_{10} + \pi_{\text{env},1}\pi_{\text{env},2}\ln\frac{r_{11}}{r_{01}}, \tag{C.12}$$

with $\pi_{\text{env,eff}} = \pi_{\text{env},1} + \pi_{\text{env},2} - \pi_{\text{env},1}\pi_{\text{env},2}$. The last term in this expression is small for either infrequent pathogens ($\pi_{\text{env},1}\pi_{\text{env},2} \ll \pi_{\text{env,eff}}$) or if a large fraction of the population is protected ($1-p \ll 1$ hence $r_{01} \approx r_{11}$). Neglecting this second-order term, we are left with the expression corresponding to a single pathogen with effective frequency $\pi_{\text{env,eff}}$, in agreement with our expectation.

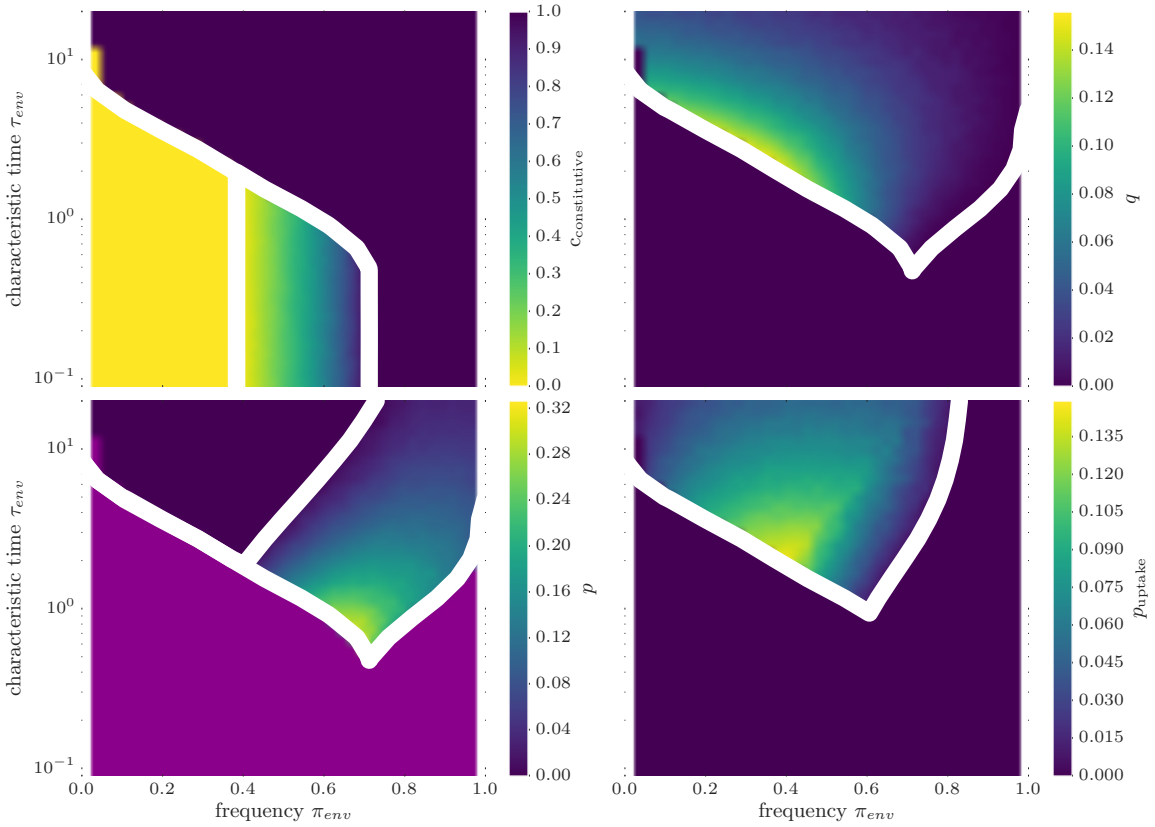

Figure C.1: Optimal parameters from a global optimization of long-term growth rate. Regions where a parameter is unconstrained at the optimum are shown in purple. Phase boundaries pertaining to the shown parameter in white. A maximum number of 10000 function evaluations is used for the first phase of the optimization. The second phase of the optimization is terminated at a tolerance in the parameter values of 0.005. The same model parameters as in Fig. 6.2 are used.

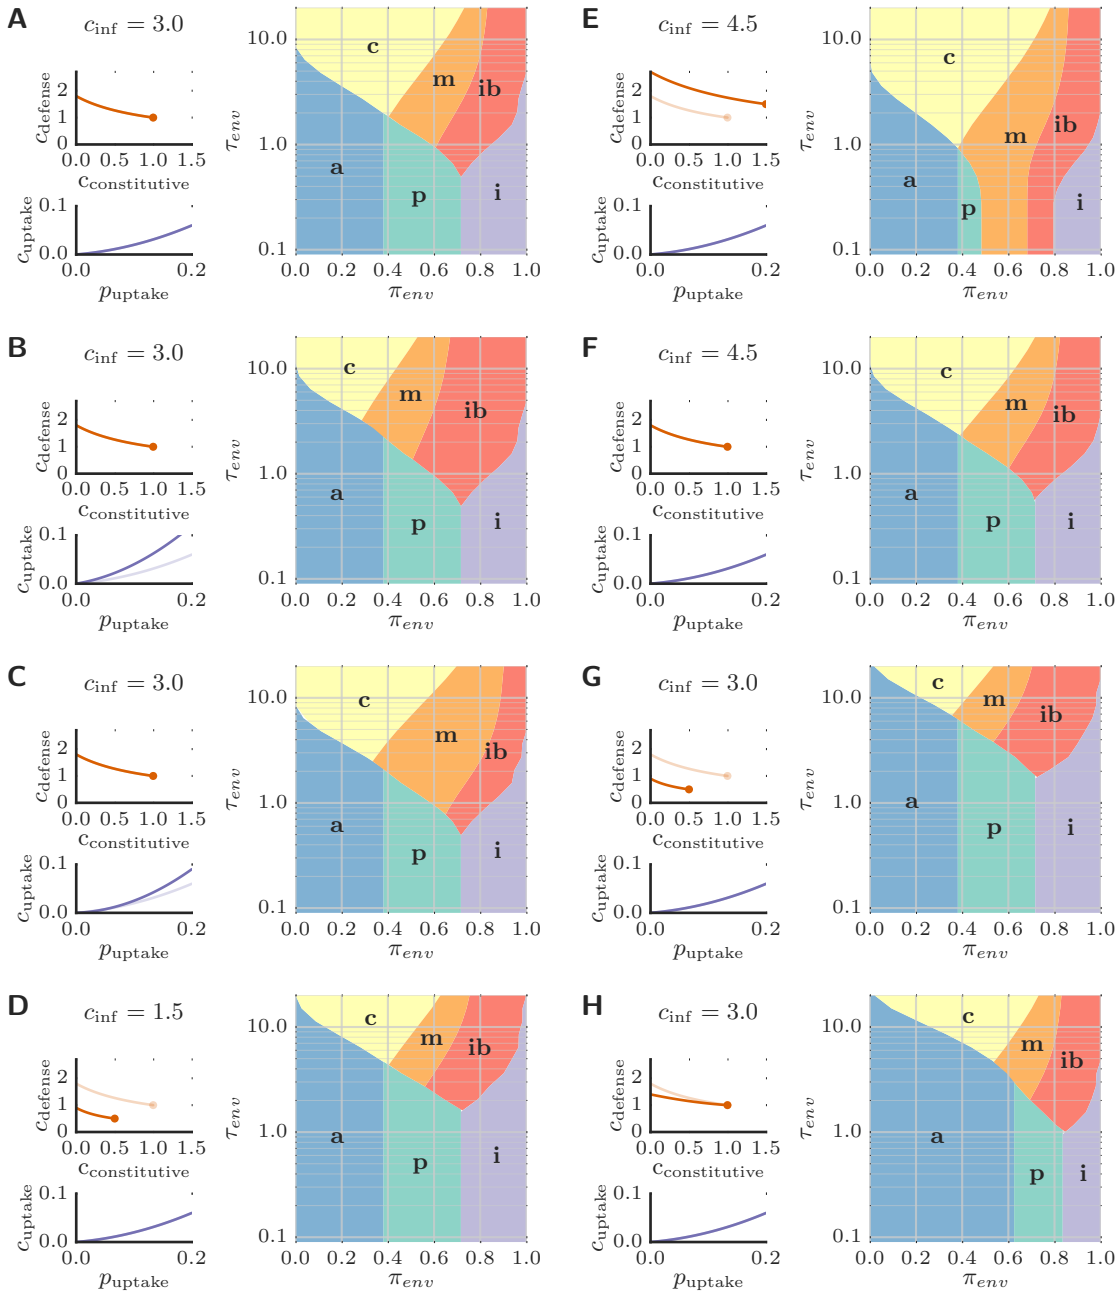

Figure C.2: Influence of parameter choice on the phase diagram presented in Fig. 6.2. For every panel the parameter choices are shown on the left and the phase boundaries between **a**daptive, **p**roto-adaptive, **i**nnate, **i**nnate **b**et hedging, **m**ixed and **C**RISPR-like strategies are shown on the right. As a reference, lines in lighter color show trade-off and uptake cost for parameter set used in Fig. 2. **(A)** Phase diagram for parameters used in Fig. 2. **(B)** More expensive active acquisition (c_{uptake} multiplied by a factor of two). **(C)** Different functional form for cost of active acqusition: $c_{uptake} = 0.05 \times p_{uptake} + 2 \times p_{uptake}^2$. **(D)** More permissive state-dependent costs (costs multiplied by a factor of 0.5). **(E)** Less permissive state-dependent costs (costs multiplied by a factor of 1.5). **(F)** Higher cost of infection. **(G)** Higher cost of immune protection. **(H)** Different functional form for cost trade-off, $c_{defense} = 1.4 - 0.6 \times c_{constitutive} + 0.2 \times c_{constitutive}^2$

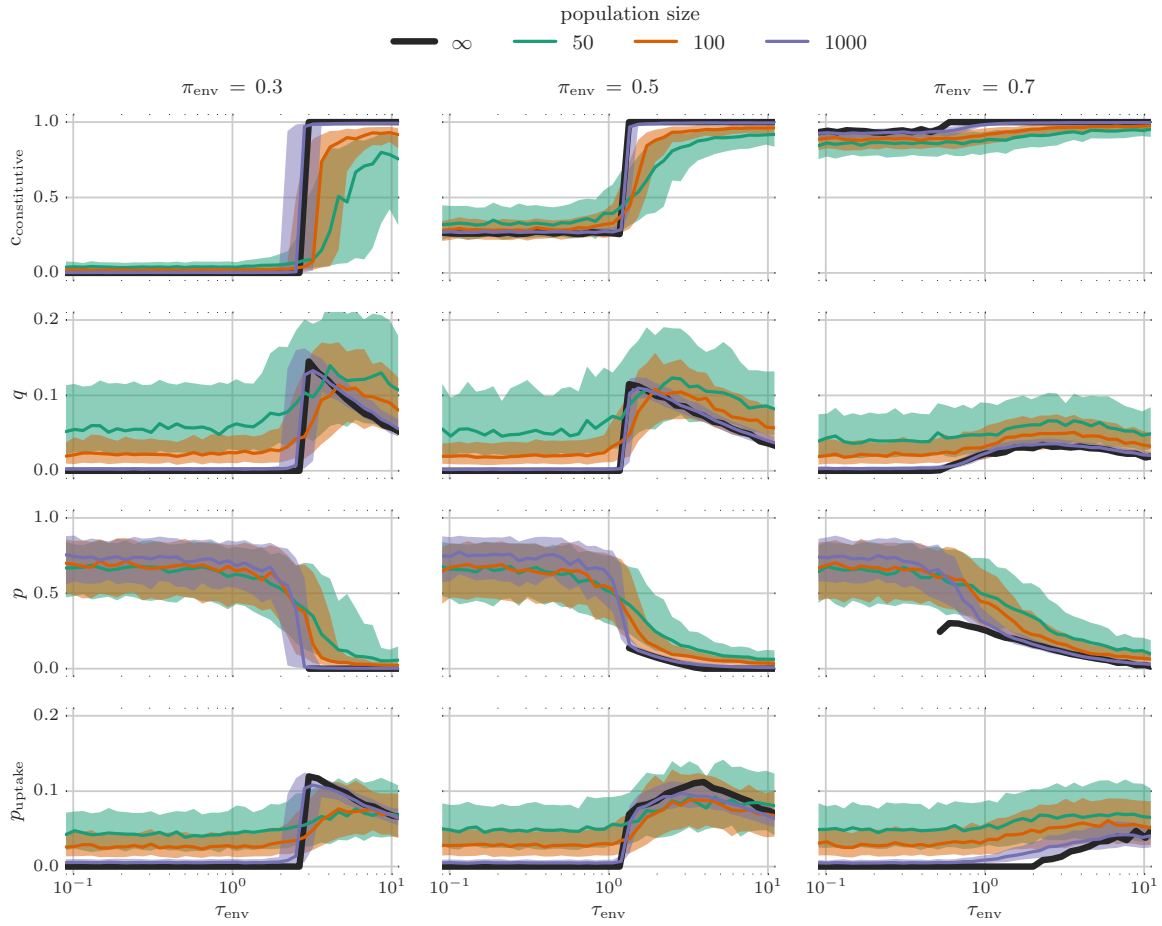

Figure C.3: Influence of finite population size on optimal immune strategies from an agent-based simulation with evolving strategy parameters (switching rates and degree of adaptability) as described in the text. For the infinite population, p is only shown for $q > 0$, because for $q = 0$ the value of p is not constrained other than being positive. Subplots show the median (solid line) and interquartile range (shaded area) of the strategy parameters at the end of a simulation of 100000 generations length. Both are calculated from 500 independent simulations. In each simulation, the strategy parameters evolve from a random initial distribution via mutation and selection. Mutations take place with a rate $0.01\exp(-t/10000)$ per generation and are normally distributed with mean zero and standard deviation $0.25\exp(-t/10000)$. The bound constraints on the parameters were enforced by setting the strategy parameters to the boundary value if outside after a mutation. Costs of different immune states as in Fig. 6.2.

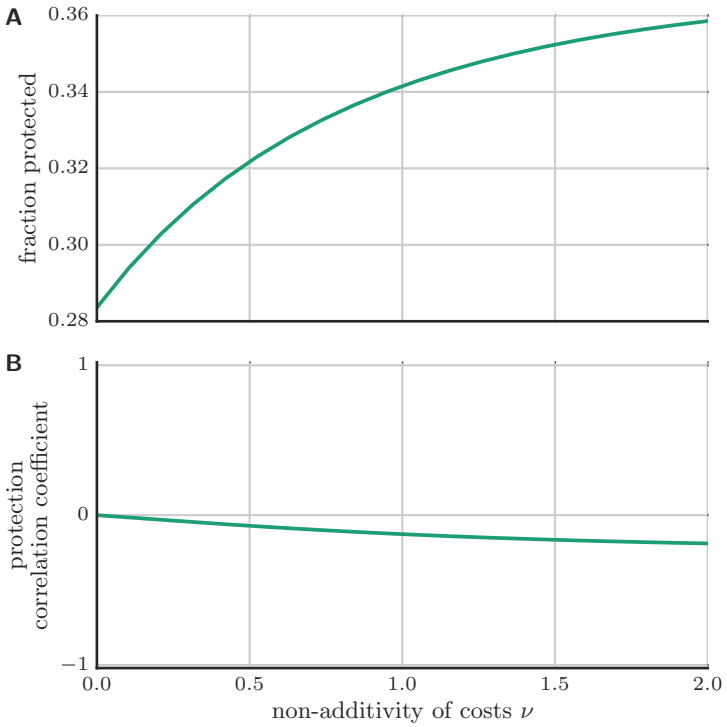

Figure C.4: Optimal protection strategy against two equally frequent pathogens $\pi_{env,1} = \pi_{env,2} = 0.4$ as a function of the degree of non-additivity of the cost of infection ν. (**A**) Fraction of population protected against a particular pathogen. (**B**) Pearson correlation coefficient between the protection states against the two pathogens. As costs are non-additive, the problem no longer factorizes and the optimal strategy no longer chooses protections against different pathogens independently. However, here the optimal strategy treats each pathogen almost indendently, as measured by the low correlation coefficient. With an increasing cost of co-infection, more protection is needed, in agreement with our intuition that co-infection leads to higher effective costs. Parameters: $c_{infection} = 2$, $c_{defense} = c_{constitutive} = 1$, optimization of the distribution over protection states respecting the probability simplex constraints using an accelerated projected gradient algorithm as described in [138].

D

TRANSITIONS IN OPTIMAL ADAPTIVE STRATEGIES

D.1 OPTIMAL STRATEGIES BY MAPPING TO UNIT SIMPLEX

If the sum of fitnesses of a phenotype over environments $f(\sigma) = \sum_x f(\sigma, x)$ is constant for all phenotypes, then any mixture will also have a constant sum of fitnesses. The normalization constraint on $\boldsymbol{\pi}$ then translates into an equivalent constraint on $\mathbf{f}$. The solution of the optimization problem in its fitness form is then particularly simple [207, 200]. Therefore, where possible, it is worthwhile to map the optimization problem to this simpler case by a rescaling of fitnesses in different environments. Here we show how to perform the rescaling and the conditions under which it is possible. Fig. D.1 illustrates such a mapping in a simple case with two environmental states.

The optimization problem is invariant with respect to additions of terms that are constant with respect to the variables over which one optimizes. Specifically, we can add the term $\sum_x p(x) \log c(x)$ to Eq. (7.7) with all positive $c(x)$, which is constant with respect to $\boldsymbol{\pi}$. This gives us a new optimization problem with the objective function $\tilde{\Lambda} = \sum_x p(x) \log[f(x)c(x)] = \sum_x p(x) \log \tilde{f}(x)$, in terms of the rescaled fitnesses $\tilde{f}(x) = \sum_\sigma \pi(\sigma)\tilde{f}(\sigma, x)$ and $\tilde{f}(\sigma, x) = f(\sigma, x)c(x)$. The equivalence of these problems shows that a rescaling of the axes of fitness space does not change the optimal adaptation strategy.

We can now try and use this rescaling to make the sum of scaled fitnesses a constant, which we chose to be 1 without restriction of generality. This means we aim to chose $c(x)$, such that $\sum_x \tilde{f}(\sigma, x) = \sum_x f(\sigma, x)c(x) = 1$ holds for all σ. In matrix-vector notation we can represent these conditions as the systems of equation

$$\mathbf{F}\mathbf{c} = \mathbf{1}, \tag{D.1}$$

where $\mathbf{1} = (1, 1, ..., 1)^T$ is the vector of all ones and $\mathbf{F}$ the matrix of phenotype fitness profiles with entries $F_{\sigma,x} = f(\sigma, x)$. Eq. (D.1) requires that the scalar products of $\mathbf{c}$ with the row vectors of $\mathbf{f}$ (the phenotypes fitness profiles) are equal to 1 for all rows. The vector $\mathbf{c}$ thus is a normal vector to the hyperplane spanned by the fitness profiles. The mapping is therefore only possible if a hyperplane passing through all fitness profiles exists. The intercept d_x of the hyperplane with the x axis is given by $\mathbf{c}(d_x \boldsymbol{e}_x) = 1 \Leftrightarrow d_x = 1/c_x$, where $\boldsymbol{e}_x$ is the x-th unit vector. Eq. (D.1) thus specificies that we should rescale fitnesses by dividing through these intercepts to achieve our goal of mapping the problem to the unit simplex. The positivity of the scaling constants $c(x)$ puts further requirements on $\mathbf{F}$ for the mapping to work: geometrically, all intercepts need to be positive, or algebraically, the inverse of the fitness matrix needs to have all positive row sums. In the case where $\mathbf{F}$ is an invertible matrix fitnesses should be rescaled using $\mathbf{c} = \mathbf{F}^{-1}\mathbf{1}$. If the scaling is possible then in the scaled

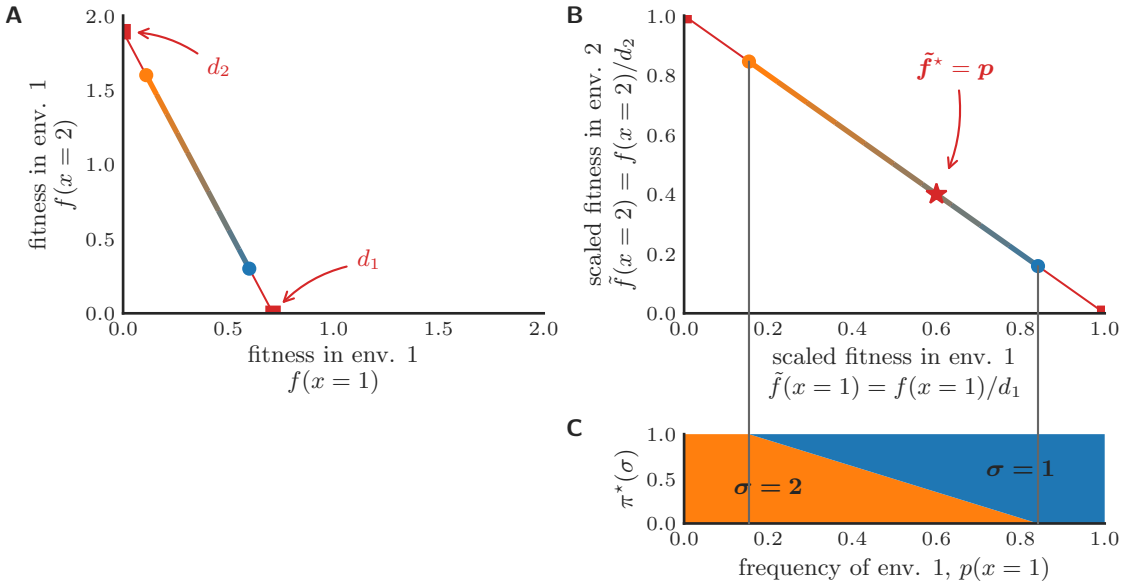

Figure D.1: Mapping of the problem to the unit simplex helps optimizing long-term growth rate graphically. To determine the best strategy using two phenotypes (blue/orange dots) and their mixtures (colored line) we rescale the original fitnesses (A) such that the sum of fitnesses is constant (B). To do so fitnesses are rescaled by dividing through the intercepts (red squares) of the line passing through the two points with the axes (red line). In the scaled fitnesses the optimal strategy has fitness vector $\tilde{\mathbf{f}}^\star = \mathbf{p}$ (red star), which can be be mapped back to the original problem by reverting the rescaling. Where the so-determined fitnesses lie between the fitnesses of the two phenotypes the optimal strategy switches between both phenotypes with frequencies relative to how far the optimum is from the two phenotypes. If the optimal rescaled fitness lies outside the achievable range of fitnesses using the closest phenotype is optimal. (C) Optimal mixture of the two phenotypes as a function of the frequency of environmental state 1.

variables the normalization constraint on $\boldsymbol{\pi}$ leads to a normalization constraint on $\tilde{\mathbf{f}}$.

We can derive the optimal fitness profile using the Lagrange formalism. The Lagrangian of the optimization problem is

$$\mathcal{L} = \sum_x p(x) \ln \tilde{f}(x) - \lambda \left(\sum_x \tilde{f}(x) - 1 \right), \tag{D.2}$$

where we have assumed that the optimum is in the interior of $D_{\tilde{f}}$, i.e. none of the non-negativity constraints on elements $\boldsymbol{\pi}$ are active. Taking the derivative with respect to $\tilde{f}(x)$ and setting it to zero yields $\tilde{f}^\star(x) = p(x)/\lambda$. As we have rescaled fitnesses such that the sum of fitnesses scale to 1 the Lagrange multiplier is $\lambda = 1$. In the rescaled variables the optimal strategy thus allocates fitness to each environment proportional to its frequency:

$$\tilde{\mathbf{f}}^\star = \mathbf{p}. \tag{D.3}$$

From the optimum in the rescaled variables the optimum in the original variables can be obtained by reversing the scaling $f_i^\star = \tilde{f}_i^\star / c_i$.

Due to the non-negativity constraints on $\pi(\sigma)$, which we have neglected so far in the discussion, only a subset of the unit simplex is accessible if there are no phenotypes that are not completely specialized to the different environments. Where the unconstrained solution lies outside the feasible region a value on the boundary of the fitness set is constrained optimum instead. The fitness allocation among the remaining unconstrained directions still is proportional to the frequency of the respective environments.

D.2 ANALYTICAL RESULTS ON OPTIMAL IMMUNE STRATEGIES IN UNCORRELATED ENVIRONMENTS

D.2.1 *Optimization problem*

The cost function of the optimization problem is the long-term population growth rate, which depends on the environmental statistics and the chosen strategy. The long-term growth rate in uncorrelated environments for a given $p, f_{base}, f_{inf}, f_{def}(f_{con})$ is given by Eq. 7.16, which we recall here

$$\begin{aligned} \Lambda(\pi, f_{con}) = {} & p \ln[\pi f_{def} + (1-\pi) f_{inf}] \\ & + (1-p) \ln[\pi f_{con} + (1-\pi) f_{base}]. \end{aligned} \tag{D.4}$$

To find the optimal strategy we need to solve the following optimization problem

$$\begin{aligned} \underset{\pi, f_{con}}{\text{maximize}} \quad & \Lambda(\pi, f_{con}) \\ \text{subject to} \quad & 0 \leqslant \pi \leqslant 1 \\ & f_{con}^{min} \leqslant f_{con} \leqslant f_{con}^{max} \end{aligned} \quad . \tag{D.5}$$

The optimization consists in finding the (global) maximum of a two-variable objective function subject to bound constraints on both variables. In the following derivations we make use of the ordering of the costs in the non-trivial case $f_{base} \geqslant f_{con}, f_{def} > f_{inf}$ and of the Pareto condition on the trade-off line $f'_{def}(f_{con}) < 0$.

This problem can be solved numerically, but as is shown in the following a lot of information is available from a purely analytical treatment of the optimization problem. The Karush-Kuhn-Tucker conditions give necessary conditions for local optimality of a point $\pi^\star, f^\star_{con}$. For bound constrained problems these conditions boil down to the statement that the partial derivative of the objective function with respect to either variable needs to be [91]: zero if the variable is in the interior of its feasible interval, negative if the variable is at the lower end of its feasible domain, and positive if the variable is at the upper end of its feasible domain. Expressed in equations the necessary conditions for $\pi^\star, f^\star_{con}$ to be locally optimal is that

$$\partial_\pi \Lambda(\pi^\star, f^\star_{con}) \begin{cases} \leqslant 0, \text{ if } \pi^\star = 0 \\ \geqslant 0, \text{ if } \pi^\star = 1 \\ = 0, \text{ otherwise} \end{cases} \tag{D.6}$$

and that

$$\partial_{f_{con}} \Lambda(\pi^\star, f^\star_{con}) \begin{cases} \leqslant 0, \ f^\star_{con} = 0 \\ \geqslant 0, \ f^\star_{con} = 1 \\ = 0, \text{ otherwise} \end{cases} . \tag{D.7}$$

The conditions provide only necessary but not sufficient conditions for local optimality. A condition ensuring sufficiency is that the Hessian at the optimum constricted to the feasible directions is negative definite.

D.2.2 *Derivatives of the cost function*

The optimality conditions derived in the previous subsection involve the derivatives of the cost function, which can be obtained using simple algebra and which we give below. The derivative of the cost function with respect to π is given by

$$\partial_\pi \Lambda = \frac{p(f_{def} - f_{inf})}{f_{inf}(1-\pi) + f_{def}\pi} - \frac{(1-p)(f_{base} - f_{con})}{f_{base}(1-\pi) + f_{con}\pi} \tag{D.8}$$

$$\partial_{f_{con}} \Lambda = \pi \left[\frac{p f'_{def}}{f_{inf}(1-\pi) + f_{def}\pi} + \frac{(1-p)}{f_{base}(1-\pi) + f_{con}\pi} \right] . \tag{D.9}$$

For sufficiency we also need to consider the second derivatives of the cost function:

$$\partial_\pi^2 \Lambda = -\frac{p(f_{def} - f_{inf})^2}{(f_{inf}(1-\pi) + f_{def}\pi)^2} - \frac{(1-p)(f_{base} - f_{con})^2}{(f_{base}(1-\pi) + f_{con}\pi)^2} \tag{D.10}$$

$$\partial_{f_{con}}^2 \Lambda = -\pi \left[\frac{p(\pi f_{def}'^2 - (f_{inf}(1-\pi) + \pi f_{def}) f_{def}'')}{(f_{inf}(1-\pi) + f_{def}\pi)^2} + \frac{(1-p)\pi}{(f_{base}(1-\pi) + f_{con}\pi)^2} \right]. \tag{D.11}$$

The second derivative with respect to π is always negative which shows that the long-term growth rate is a concave function of π. For a fixed value of f_{con} the optimization thus corresponds to a maximization of a concave function and always yields a unique optimum. The second derivative of the long-term growth rate with respect to f_{con} is also always negative, if $f_{def}'' \leqslant 0$. This condition on the trade-off function is fulfilled if individuals might bet hedge in their degree of specialization in environment 1. Otherwise the second derivative might be positive for some p and there can thus exist more than one local optimal in the full optimization problem.

A sufficient condition for having a local maximum is the negative definiteness of the Hessian. As one of its diagonal elements is always negative this is equivalent to showing that the determinant of the Hessian is positive. The determinant of the Hessian at an interior stationary point can be calculated to be

$$\det \nabla^2 \Lambda(\pi^\star, f_{con}^\star) = -f_{def}''(f_{con}^\star) \frac{(f_{base} - f_{con}^\star)^2 (f_{def}^\star - f_{inf})\pi}{(f_{base} f_{def}^\star - f_{inf} f_{con}^\star)(f_{base}(1-\pi) + f_{con}^\star \pi)(f_{inf}(1-\pi) + \pi f_{def}^\star)}. \tag{D.12}$$

It follows that for $f_{con}^\star$ to be optimal for an intermediate $\pi^\star$ the trade-off curve needs to be locally concave $f_{def}''(f_{con}^\star) < 0$.

D.2.3 *Regions of local optimality for different phases*

The optimality conditions Eqs. (D.6) and (D.7) allow for three different cases for $\pi^\star$ and $f_{con}^\star$ each. This makes for a total of $3 \times 3 = 9$ different combinations. For the case $\pi^\star = 0$ the growth rate does not depend on $f_{con}^\star$, so there exists up to seven distinct phases. Under which conditions are these strategies locally optimal? In the following we analytically derive the interval of p for which these strategies are optimal.

D.2.3.1 *Tolerance ($\pi^\star = 0$, arbitrary $f_{con}^\star$)*

The condition of local optimality is $\partial_\pi \Lambda(0, f_{con}^\star) \leqslant 0$ (see Eq. (D.6)), which needs to hold for all feasible $f_{con}^\star$. This translates to the condition $p \leqslant \frac{f_{inf}(f_{base} - f_{con}^\star)}{f_{base} f_{def}^\star - f_{inf} f_{con}^\star}$.

The condition needs to hold for the $f^{\star}_{con}$ giving the strictest bound. Tolerance thus is optimal for

$$p \leqslant \min_{f_{con}} \frac{f_{inf}(f_{base} - f_{con})}{f_{base}f_{def} - f_{inf}f_{con}} =: p^{(0)}, \tag{D.13}$$

i.e. for the rarest pathogens. If the adaptive strategy comes without constitutive cost ($f^{min}_{con} = 0$), then the tolerance phase disappears ($p^{(0)} = 0$). Where the phase exists it is followed by one of the bet hedging strategies.

D.2.3.2 *Innate* ($\pi^{\star} = 1, f^{\star}_{con} = f^{min}_{con}$)

From Eq. D.6 the condition of local optimality is $\partial_{\pi}\Lambda(1, f^{min}_{con}) \geqslant 0$. This translates to the condition

$$p \geqslant \frac{(f_{base} - f^{min}_{con})f^{max}_{def}}{f_{base}f^{max}_{def} - f_{inf}f^{min}_{con}} =: p^{(i\tilde{i})}. \tag{D.14}$$

The second optimality condition Eq. D.7 is $\partial_{f_{con}}\Lambda(1, f^{min}_{con}) \geqslant 0$, leading to

$$p \geqslant \frac{f^{max}_{def}}{f^{max}_{def} - f^{min}_{con}f'_{def}(f^{min}_{con})} =: p^{(ip)}. \tag{D.15}$$

Both conditions need to hold at the same time for local optimality so an innate strategy is optimal for

$$p \geqslant \max(p^{(ip)}, p^{(i\tilde{i})}), \tag{D.16}$$

i.e. for the most frequent pathogens. Depending on which of the two conditions is more stringent it is followed either by a innate bet hedging strategy or a protoadaptive phase.

D.2.3.3 *Adaptive* ($\pi^{\star} = 1, f^{\star}_{con} = f^{max}_{con}$)

Eq. (D.6) leads to

$$p \geqslant \frac{(f_{base} - f^{max}_{con})f^{min}_{def}}{f_{base}f^{min}_{def} - f_{inf}f^{max}_{con}} =: p^{(a\tilde{a})} \tag{D.17}$$

and Eq. (D.7) to

$$p \leqslant \frac{f^{min}_{def}}{f^{min}_{def} - f^{max}_{con}f'_{def}(f^{max}_{con})} =: p^{(ap)}. \tag{D.18}$$

Taken together an adaptive strategy is optimal for

$$p^{(a\tilde{a})} \leqslant p \leqslant p^{(ap)}. \tag{D.19}$$

D.2.3.4 *Protoadaptive ($\pi^\star = 1$, intermediate $f^\star_{con}$)*

Eq. (D.6) leads to $p \geqslant \frac{(f_{base}-f^\star_{con})f^\star_{def}}{f_{base}f^\star_{def}-f_{inf}f^\star_{con}}$ and Eq. (D.7) to $p = \frac{f^\star_{def}}{f^\star_{def}-f^\star_{con}f'_{def}(f^\star_{con})}$. The two conditions together lead to

$$f'_{def}(f^\star_{con}) \geqslant -\frac{f_{def}(f^\star_{con}) - f_{inf}}{f_{base} - f^\star_{con}}, \tag{D.20}$$

i.e. the derivative of the trade-off function needs to be more shallow then the derivative of costs of a mixture with the current type. As we have an intermediate level of regulation we need to check the second derivative to assure the extremum is a maximum. As shown in the main text this leads to the condition $\frac{d^2 \ln f_{def}}{d(\ln f_{con})^2} < 0$. If the trade-off function is assumed to be fulfill both conditions everywhere and to be smooth then by the intermediate value theorem there is a $f^\star_{con}$, which is optimal for a p in the region

$$p^{(ap)} \leqslant p \leqslant p^{(ip)} \tag{D.21}$$

D.2.3.5 *innate switching (intermediate $\pi^\star$, $f^\star_{con} = f^{min}_{con}$)*

Eq. (D.6) leads to

$$p^{(0\tilde{i})} \leqslant p \leqslant p^{(i\tilde{i})} \tag{D.22}$$

with

$$p^{(0\tilde{i})} = \frac{f_{inf}(f_{base} - f^{min}_{con})}{f_{base}f^{max}_{def} - f_{inf}f^{min}_{con}} \tag{D.23}$$

and Eq. (D.7) to

$$f'_{def}(f^{min}_{con}) \leqslant -\frac{f^{max}_{def} - f_{inf}}{f_{base} - f^{min}_{con}} \tag{D.24}$$

i.e. the derivative of the trade-off shape needs to be steeper then the line joining the unprotected state.

D.2.3.6 *adaptive switching (intermediate $\pi^\star$, $f^\star_{con} = f^{max}_{con}$)*

Eq. (D.6) leads to

$$p^{(0\tilde{a})} \leqslant p \leqslant p^{(a\tilde{a})} \tag{D.25}$$

with

$$p^{(0\tilde{a})} = \frac{f_{inf}(f_{base} - f^{max}_{con})}{f_{base}f^{min}_{def} - f_{inf}f^{max}_{con}} \tag{D.26}$$

and Eq. (D.7) to

$$f'_{def}(f^{min}_{con}) \leqslant -\frac{f^{min}_{def} - f_{inf}}{f_{base} - f^{max}_{con}}. \tag{D.27}$$

D.2.3.7 *protoadaptive switching (intermediate $\pi^\star, f^\star_{con}$)*

Eq. (D.6) leads to

$$p^{(0\tilde{p})} \leqslant p \leqslant p^{(p\tilde{p})} \tag{D.28}$$

with

$$p^{(0\tilde{p})} = \frac{f_{inf}(f_{base} - f^\star_{con})}{f_{base}f^\star_{def} - f_{inf}f^\star_{con}} \tag{D.29}$$

and Eq. (D.7) to

$$f'_{def}(f^\star_{con}) = -\frac{f_{def}(f^\star_{con}) - f_{inf}}{f_{base} - f^\star_{con}}. \tag{D.30}$$

The derivative needs to be equal to the slope of the line connecting the fitness profile to the non-protected type. The sufficiency condition $\det H(\pi^\star, f^\star_{con}) > 0$ leads to

$$f''_{def}(f^\star_{con}) < 0. \tag{D.31}$$

D.3 DERIVATION OF LONG-TERM GROWTH RATE IN THE ADIABATIC LIMIT

The study of the adiabatic limit in which the durations of environmental periods are large relative to the time scales of population composition change goes back to [111]. Mathematically the long-term growth rate can be approximated by an eigenvalue perturbation approach. In the following we give a derivation following the notations of [117].

The transfer matrix connecting the population composition at successive time points is $\langle\sigma'|A^{(x)}|\sigma\rangle = f(\sigma', x)\pi(\sigma'|\sigma)$ (in bra-ket notation), which one can decompose as $A^{(x)} = A^{(x)}_0 + A^{(x)}_1$ with

$$\langle\sigma'|A^{(x)}_0|\sigma\rangle = \begin{cases} f(\sigma, x) \text{ if } \sigma' = \sigma \\ 0 \text{ otherwise.} \end{cases} \tag{D.32}$$

and

$$\langle\sigma'|A^{(x)}_1|\sigma\rangle = \begin{cases} -f(\sigma, x)(1 - \pi(\sigma|\sigma)) \text{ if } \sigma' = \sigma \\ f(\sigma', x)\pi(\sigma'|\sigma) \text{ otherwise.} \end{cases} \tag{D.33}$$

Using this decomposition we treat $A^{(x)}_1$ as a perturbation to $A^{(x)}_0$ to approximately solve the eigenvalue problem of $A^{(x)}$. As A_0 is diagonal its eigenvalues are $\lambda_{0,\sigma} = f(\sigma, x)$ with corresponding eigenvectors $|\sigma\rangle$, which have all but the σ-th element set to zero. Applying the formulas for the eigenvalues and eigenvectors of the perturbed problem we obtain

$$\lambda_\sigma = f(\sigma, x)\pi(\sigma|\sigma) \tag{D.34}$$

and the corresponding right eigenvectors

$$|\psi_\sigma^{(x)}\rangle = |\sigma\rangle + \sum_{\sigma'\neq\sigma} \frac{\langle\sigma'|A_1^{(x)}|\sigma\rangle}{f(\sigma,x) - f(\sigma',x)}|\sigma'\rangle \tag{D.35}$$

$$= |\sigma\rangle + \sum_{\sigma'\neq\sigma} \frac{f(\sigma',x)\pi(\sigma'|\sigma)}{f(\sigma,x) - f(\sigma',x)}|\sigma'\rangle. \tag{D.36}$$

In order to calculate overlaps we also need to calculate left eigenvectors. The left eigenvectors of $A_0^{(x)}$ are equal to its right eigenvectors as its a diagonal matrix. The left eigenvectors of the perturbed problem are

$$\langle\psi_\sigma^{(x)}| = \langle\sigma| + \sum_{\sigma'\neq\sigma} \frac{\langle\sigma'|(A_1^{(x)})^T|\sigma\rangle}{f(\sigma,x) - f(\sigma',x)}\langle\sigma'| \tag{D.37}$$

$$= \langle\sigma| + \sum_{\sigma'\neq\sigma} \frac{\langle\sigma|A_1^{(x)}|\sigma'\rangle}{f(\sigma,x) - f(\sigma',x)}\langle\sigma'| \tag{D.38}$$

$$= \langle\sigma| + \sum_{\sigma'\neq\sigma} \frac{f(\sigma,x)\pi(\sigma|\sigma')}{f(\sigma,x) - f(\sigma',x)}\langle\sigma'|. \tag{D.39}$$

Let us assume that for every environment x there is a type $\sigma = x$, which provides optimal growth. The overlap between the largest eigenvectors in environments x and x' is given by

$$Q(x,x') := \langle\psi_x^{(x)}|\psi_{x'}^{(x')}\rangle = \pi(x|x')\Gamma(x,x') \tag{D.40}$$

with

$$\Gamma(x,x') = \frac{f(x,x')}{f(x',x') - f(x,x')} + \frac{f(x,x)}{f(x,x) - f(x',x)} \tag{D.41}$$

In the adiabatic limit the long-term growth rate is given by

$$\Lambda = \sum_x p(x)\ln\lambda_x - \sum_{x,x';x\neq x'} p(x'|x)p(x)\ln\frac{1}{Q(x,x')} \tag{D.42}$$

$$= \sum_x p(x)\ln f(x,x) + \sum_{x,x'} p(x'|x)p(x)\ln[\pi(x|x')\Gamma(x,x')], \tag{D.43}$$

where we have defined $\Gamma(x,x) = 1$.

We can write out the sums in the case of an environment switching between two states as

$$\begin{aligned}\Lambda =& p(1)\ln f(1,1) + p(2)\ln f(2,2)\\ &+ p(1|2)p(2)\ln[\pi(1|2)\Gamma(1,2)]\\ &+ p(2|1)p(1)\ln[\pi(2|1)\Gamma(2,1)].\end{aligned} \tag{D.44}$$

To compare the best switching strategies using phenotypes of fitness f or $\tilde{f}$ we calculate the long-term growth rate difference

$$\begin{aligned}\Delta\Lambda =&(1-p_2)\ln\frac{f(1,1)}{\tilde{f}(1,1)}+p_2\frac{f(2,2)}{\tilde{f}(2,2)}\\&+(1-e^{-1/t_c})p_2(1-p_2)\ln\frac{\Gamma(1,2)\Gamma(2,1)}{\tilde{\Gamma}(1,2)\tilde{\Gamma}(2,1)},\end{aligned}\tag{D.45}$$

where we have used short-hand notations for the environmental switching frequencies as introduced in the text. Setting $\Delta\Lambda = 0$ we can solve for the transition line between the two sets of phenotypes,

$$e^{-1/t_c}=1-\frac{(1-p_2)\ln\frac{\tilde{f}(1,1)}{f(1,1)}+p_2\ln\frac{\tilde{f}(2,2)}{f(2,2)}}{(1-p_2)p_2\ln\frac{\Gamma(1,2)\Gamma(2,1)}{\tilde{\Gamma}(1,2)\tilde{\Gamma}(2,1)}}.\tag{D.46}$$

Such an analysis can be applied to the case where a generalist phenotype is on the Pareto frontier to find when switching only uses specialists (Fig. 7.7). To do so we compare the growth rate of switching between the specialist phenotypes $\sigma = 1$ and $\sigma = 2$ to the growth rates of switching between one of the specialists and the generalist $\sigma = 3$. (D.46) then gives an approximate result for the time scale of environmental correlations above which switching only involves specialists.

BIBLIOGRAPHY

[1] Amalio Telenti et al. "Deep Sequencing of 10,000 Human Genomes." In: *Proceedings of the National Academy of Sciences* 113.42 (2016), pp. 11901–11906.

[2] Chris Anderson. "The End of Theory: The Data Deluge Makes the Scientific Method Obsolete." In: *Wired Magazine* 16.07 (2008), pp. 1–2.

[3] Rob Phillips. "Theory in Biology: Figure 1 or Figure 7?" In: *Trends in Cell Biology* 25.12 (2015), pp. 723–729.

[4] William Bialek. "Perspectives on theory at the interface of physics and biology." In: *arXiv preprint arXiv:1512.08954* (2015).

[5] Philip W. Anderson. "More is Different." In: *Science* 177.4047 (1972), pp. 393–396.

[6] Thomas Boehm. "Design principles of adaptive immune systems." In: *Nature Reviews Immunology* 11.5 (2011), pp. 307–317.

[7] Luciano A Marraffini and Erik J Sontheimer. "CRISPR interference: RNA-directed adaptive immunity in bacteria and archaea." In: *Nature reviews. Genetics* 11.3 (2010), pp. 181–90.

[8] Mihai G Netea, Leo A B Joosten, Eicke Latz, Kingston H G Mills, Gioacchino Natoli, Hendrik G Stunnenberg, Luke A J O'Neill, and Ramnik J Xavier. "Trained immunity: A program of innate immune memory in health and disease." In: *Science* 352.6284 (2016), aaf1098.

[9] Pierre Boudinot, Maria Encarnita Marriotti-Ferrandiz, Louis D. Pasquier, Abdenour Benmansour, P. A. Cazenave, and Adrien Six. "New perspectives for large-scale repertoire analysis of immune receptors." In: *Molecular Immunology* 45.9 (2008), pp. 2437–2445.

[10] Joshua A J A Weinstein, Ning N Jiang, Richard A R A White, Daniel S D S Fisher, and Stephen R S R Quake. "High-throughput sequencing of the zebrafish antibody repertoire." In: *Science* 324.5928 (2009), pp. 807–810.

[11] Adrien Six, Maria Encarnita Mariotti-Ferrandiz, Wahiba Chaara, Susana Magadan, Hang Phuong Pham, Marie Paule Lefranc, Thierry Mora, Véronique Thomas-Vaslin, Aleksandra M. Walczak, and Pierre Boudinot. "The past, present, and future of immune repertoire biology - the rise of next-generation repertoire analysis." In: *Frontiers in Immunology* 4.NOV (2013).

[12] Stephen Jay Gould and Richard C Lewontin. "The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist programme." In: *Proceedings of the Royal Society of London. Series B. Biological Sciences* 205.1161 (1979), pp. 581–598.

[13] Geoffrey A Parker and J. Maynard Smith. "Optimality theory in evolutionary biology." In: *Nature* 348.6296 (1990), pp. 27–33.

[14] William Bialek. *Biophysics: searching for principles*. Princeton University Press, 2012.

[15] Paul Krugman. "What economists can learn from evolutionary theorists." In: *A talk given to the European Association for Evolutionary Political Economy* (1996).

[16] J. H. van Hateren. "A theory of maximizing sensory information." In: *Biological Cybernetics* 68.1 (1992), pp. 23–29.

[17] Jeffrey M. Beck, Wei Ji Ma, Roozbeh Kiani, Tim Hanks, Anne K. Churchland, Jamie Roitman, Michael N. Shadlen, Peter E. Latham, and Alexandre Pouget. "Probabilistic Population Codes for Bayesian Decision Making." In: *Neuron* 60.6 (2008), pp. 1142–1152.

[18] Alexander P. Nikitin, Nigel G. Stocks, Robert P. Morse, and Mark D. McDonnell. "Neural population coding is optimized by discrete tuning curves." In: *Physical Review Letters* 103.13 (2009).

[19] Erez Dekel and Uri Alon. "Optimality and evolutionary tuning of the expression level of a protein." In: *Nature* 436.7050 (2005), pp. 588–92.

[20] Alfonso Pérez-Escudero et al. "Structure of deviations from optimality in biological systems." In: *Proceedings of the National Academy of Sciences* 106.48 (2009), pp. 20544–20549.

[21] Jeremy Gunawardena. "Models in biology: 'accurate descriptions of our pathetic thinking'." In: *BMC biology* 12 (2014), p. 29.

[22] Kenneth Murphy, Paul Travers, and Mark Walport. *Janeway's Immunobiology*. 7th. Vol. 2. Garland Science, 2001.

[23] Ruslan Medzhitov. "Recognition of microorganisms and activation of the immune response." In: *Nature* 449.7164 (2007), pp. 819–826.

[24] T. N. Schumacher and R. D. Schreiber. "Neoantigens in cancer immunotherapy." In: *Science* 348.6230 (2015), pp. 69–74.

[25] Ning Jiang. "Immune engineering: from systems immunology to engineering immunity." In: *Current Opinion in Biomedical Engineering* (2017).

[26] Curtis G. Callan, Thierry Mora, and Aleksandra M. Walczak. "Repertoire sequencing and the statistical ensemble approach to adaptive immunity." In: *Current Opinion in Systems Biology* (2016).

[27] F M Burnet. "A modification of Jerne's theory of antibody production using the concept of clonal selection." In: *The Australian Journal of Science* 20.3 (1957), pp. 67–69.

[28] Wilfred Ndifon, Hilah Gal, Eric Shifrut, Rina Aharoni, Nissan Yissachar, Nir Waysbort, Shlomit Reich-Zeliger, Ruth Arnon, and Nir Friedman. "Chromatin conformation governs T-cell receptor J gene segment usage." In: *Proceedings of the National Academy of Sciences* 109.39 (2012), pp. 15865–15870.

[29] Anand Murugan, Thierry Mora, Aleksandra M Walczak, and Curtis G Callan. "Statistical inference of the generation probability of T-cell receptors from sequence repertoires." In: *Proceedings of the National Academy of Sciences* 109 (2012), pp. 16161–16166.

[30] Vanessa Venturi, David a Price, Daniel C Douek, and Miles P Davenport. "The molecular basis for public T-cell responses?" In: *Nature Reviews Immunology* 8.3 (2008), pp. 231–238.

[31] Antonio A Freitas and Benedita Rocha. "Population biology of lymphocytes: the flight for survival." In: *Annual review of immunology* 18.1 (2000), pp. 83–111.

[32] Charles D. Surh and Jonathan Sprent. "Homeostasis of Naive and Memory T Cells." In: *Immunity* 29.6 (2008), pp. 848–862.

[33] Rob J De Boer and Alan S Perelson. "Towards a general function describing T cell proliferation." In: *J Theor Biol* 175 (1995), pp. 567–576.

[34] A R Mclean, M M Rosado, F Agenes, R Vasconcellos, A A Freitas, and Robert M May. "Resource competition as a mechanism for B cell homeostasis." In: *Immunology* 94 (1997), pp. 5792–5797.

[35] Donna L Farber, Naomi a Yudanin, and Nicholas P Restifo. "Human memory T cells: generation, compartmentalization and homeostasis." In: *Nature Reviews Immunology* 14.1 (2014), pp. 24–35.

[36] Krista L Yorita Christensen, Robert C Holman, Claudia a Steiner, James J Sejvar, Barbara J Stoll, and Lawrence B Schonberger. "Infectious disease hospitalizations in the United States." In: *Clinical infectious diseases* 49 (2009), pp. 1025–1035.

[37] Pasquine Saule, Jacques Trauet, Virginie Dutriez, Véronique Lekeux, Jean Paul Dessaint, and Myriam Labalette. "Accumulation of memory T cells from childhood to old age: Central and effector memory cells in CD4+ versus effector memory and terminally differentiated memory cells in CD8+ compartment." In: *Mechanisms of Ageing and Development* 127.3 (2006), pp. 274–281.

[38] W T Shearer et al. "Lymphocyte subsets in healthy children from birth through 18 years of age: The pediatric AIDS clinical trials group P1009 study." In: *Journal Of Allergy And Clinical Immunology* 112.5 (2003), pp. 973–980.

[39] John T. Harty and Vladimir P. Badovinac. "Shaping and reshaping CD8+ T-cell memory." In: *Nature Reviews Immunology* 8.2 (2008), pp. 107–19.

[40] James J. Moon, H. Hamlet Chu, Marion Pepper, Stephen J. McSorley, Stephen C. Jameson, Ross M. Kedl, and Marc K. Jenkins. "Naive CD4+ T Cell Frequency Varies for Different Epitopes and Predicts Repertoire Diversity and Response Magnitude." In: *Immunity* 27.2 (2007), pp. 203–213.

[41] Marc K Jenkins and James J Moon. "The role of naive {T} cell precursor frequency and recruitment in dictating immune response magnitude." In: *The Journal of Immunology* 188.9 (2012), pp. 4135–4140.

[42] Veit R. Buchholz, Ton N M Schumacher, and Dirk H. Busch. "T Cell Fate at the Single-Cell Level." In: *Annual review of immunology* 34.1 (2016), pp. 65–92.

[43] A Katharina Simon, Georg A Hollander, and Andrew McMichael. "Evolution of the immune system in humans from infancy to old age." In: *Proc Biol Sci* 282.1821 (2015), pii: 20143085.

[44] Cyril F. Reboul, Grischa R. Meyer, Benjamin T. Porebski, Natalie A. Borg, and Ashley M. Buckle. "Epitope flexibility and dynamic footprint revealed by molecular dynamics of a pMHC-TCR complex." In: *PLoS Computational Biology* 8.3 (2012).

[45] Ananda W Goldrath and Michael J Bevan. "Selecting and maintaining a diverse T-cell repertoire." In: *Nature* 402.6759 (1999), pp. 255–262.

[46] K Christopher Garcia, Massimo Degano, Robyn L Stanfield, Anders Brunmark, Michael R Jackson, P A Peterson, Luc Teyton, and Ian A Wilson. "An ab T cell receptor structure at 2.5?and its orientation in the TCR-MHC complex." In: *Science* 274 (1996), pp. 209–219.

[47] Markus G Rudolph, Robyn L Stanfield, and Ian A Wilson. "How TCRs bind MHCs, peptides, and coreceptors." In: *Annual Review Of Immunology* 24 (2006), pp. 419–466.

[48] Don Mason. "A very high level of crossreactivity is an essential feature of the T- cell receptor." In: *Immunology Today* 19.9 (1998), pp. 395–404.

[49] Andrew K Sewell. "Why must T cells be cross-reactive?" In: *Nature Reviews Immunology* 12.9 (2012), pp. 669–677.

[50] Veronika I. Zarnitsyna, Brian D. Evavold, Louis N. Schoettle, Joseph N. Blattman, and Rustom Antia. "Estimating the diversity, completeness, and cross-reactivity of the T cell repertoire." In: *Frontiers in Immunology* 4.DEC (2013).

[51] Alan S Perelson and Gérard Weisbuch. "Immunology for physicists." In: *Reviews of Modern Physics* 69.4 (1997), pp. 1219–1268.

[52] InYoung Song, Anna Gil, Rabinarayan Mishra, Dario Ghersi, Liisa K Selin, and Lawrence J Stern. "Broad TCR repertoire and diverse structural solutions for recognition of an immunodominant CD8+ T cell epitope." In: *Nature Structural & Molecular Biology* February (2017), pp. 1–15.

[53] Rhys M. Adams, Justin B. Kinney, Thierry Mora, and Aleksandra M. Walczak. "Measuring the sequence-affinity landscape of antibodies with massively parallel titration curves." In: *eLife* 5 (2016), e23156.

[54] Harlan S. Robins, Paulo V. Campregher, Santosh K. Srivastava, Abigail Wacher, Cameron J. Turtle, Orsalem Kahsai, Stanley R. Riddell, Edus H. Warren, and Christopher S. Carlson. "Comprehensive assessment of T-cell receptor β-chain diversity in αβ T cells." In: *Blood* 114.19 (2009), pp. 4099–4107.

[55] Jennifer Benichou, Rotem Ben-Hamo, Yoram Louzoun, and Sol Efroni. "Rep-Seq: Uncovering the immunological repertoire through next-generation sequencing." In: *Immunology* 135.3 (2012), pp. 183–191.

[56] Mikhail Shugay et al. "Towards error-free profiling of immune repertoires." In: *Nature methods* 11.6 (2014), pp. 653–5.

[57] James M. Heather, Mazlina Ismail, Theres Oakes, and Benny Chain. "High-throughput sequencing of the T-cell receptor repertoire: pitfalls and opportunities." In: *Briefings in Bioinformatics* November 2016 (2017), bbw138.

[58] Ryan O Emerson et al. "Immunosequencing identifies signatures of cytomegalovirus exposure history and HLA-mediated effects on the T cell repertoire." In: *Nature Genetics* April (2017), pp. 1–10.

[59] Guobing Chen, Xinbo Yang, Annette Ko, Xiaoping Sun, Mingming Gao, Yongqing Zhang, Alvin Shi, Roy A. Mariuzza, and Nan-ping Weng. "Sequence and Structural Analyses Reveal Distinct and Highly Diverse Human CD8 + TCR Repertoires to Immunodominant Viral Antigens." In: *Cell Reports* 19.3 (2017), pp. 569–583.

[60] Niclas Thomas, Katharine Best, Mattia Cinelli, Shlomit Reich-Zeliger, Hilah Gal, Eric Shifrut, Asaf Madi, Nir Friedman, John Shawe-Taylor, and Benny Chain. "Tracking global changes induced in the CD4 T cell receptor repertoire by immunization with a complex antigen using short stretches of CDR3 protein sequence." In: *Bioinformatics* 30.22 (2014), pp. 3181–3188.

[61] Yuxin Sun, Katharine Best, Mattia Cinelli, James M. Heather, Shlomit Reich-Zeliger, Eric Shifrut, Nir Friedman, John Shawe-Taylor, and Benny Chain. "Specificity, Privacy, and Degeneracy in the CD4 T Cell Receptor Repertoire Following Immunization." In: *Frontiers in Immunology* 8.April (2017), pp. 1–12.

[62] Yuval Elhanati, Anand Murugan, Curtis G Callan, Thierry Mora, and Aleksandra M Walczak. "Quantifying selection in immune receptor repertoires." In: *Proceedings of the National Academy of Sciences of the United States of America* 111.27 (2014), pp. 9875–80.

[63] Jonathan Desponds, Thierry Mora, and Aleksandra M. Walczak. "Fluctuating fitness shapes the clone size distribution of immune repertoires." In: *Proceedings of the National Academy of Sciences* 113.2 (2016), pp. 274–279.

[64] Simon J Labrie, Julie E Samson, and Sylvain Moineau. "Bacteriophage resistance mechanisms." In: *Nature Reviews Microbiology* 8.5 (2010), pp. 317–327.

[65] Edze R. Westra, Stineke Van houte, Sam Oyesiku-Blakemore, Ben Makin, Jenny M. Broniewski, Alex Best, Joseph Bondy-Denomy, Alan Davidson, Mike Boots, and Angus Buckling. "Parasite exposure drives selective evolution of constitutive versus inducible defense." In: *Current Biology* 25.8 (2015), pp. 1043–1049.

[66] Serena Bradde, Marija Vucelja, Tiberiu Tesileanu, and Vijay Balasubramanian. "Dynamics of adaptive immunity against phage in bacterial populations." In: *Proceedings of the National Academy of Sciences* I.1 (2015), pp. 1–6.

[67] Parthana Mohanraju, Kira S Makarova, Bernd Zetsche, Feng Zhang, Eugene V Koonin, and John Van der Oost. "Diverse evolutionary roots and mechanistic variations of the CRISPR-Cas systems." In: *Science* 353.6299 (2016), aad5147.

[68] Nadia Danilova. "The evolution of immune mechanisms." In: *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution* 306.6 (2006), pp. 496–520.

[69] Jacob Rimer, Irun R. Cohen, and Nir Friedman. "Do all creatures possess an acquired immune system of some sort?" In: *BioEssays* 36.3 (2014), pp. 273–281.

[70] Jonathan Desponds, Andreas Mayer, Thierry Mora, and Aleksandra M. Walczak. "Population dynamics of immune repertoires." In: *arXiv preprint arXiv:1703.00226* (2017).

[71] Rob J De Boer and Alan S Perelson. "T cell repertoires and competitive exclusion." In: *Journal of theoretical biology* 169.4 (1994), pp. 375–390.

[72] Rob J. De Boer and Alan S. Perelson. "Competitive control of the self-renewing T cell repertoire." In: *International Immunology* 9.5 (1997), pp. 779–790.

[73] Rob J De Boer, Antod A Nio Freitas, and Alan S Perelsona. "Resource Competition Determines Selection of B Cell Repertoires." In: *J. theor. Biol* 212 (2001), pp. 333–343.

[74] Emily R. Stirk, Carmen Molina-París, and Hugo A. van den Berg. "Stochastic niche structure and diversity maintenance in the T cell repertoire." In: *Journal of Theoretical Biology* 255.2 (2008), pp. 237–249.

[75] Emily R. Stirk, Grant Lythe, Hugo A. van den Berg, and Carmen Molina-Paris. "Stochastic competitive exclusion in the maintenance of the naive T cell repertoire." In: *Journal of Theoretical Biology* 265.3 (2010), pp. 396–410.

[76] Grant Lythe, Robin E. Callard, Rollo L. Hoare, and Carmen Molina-París. "How many TCR clonotypes does a body maintain?" In: *Journal of Theoretical Biology* 389 (2016), pp. 214–224.

[77] Alan S Perelson and George F Oster. "Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination." In: *Journal of theoretical biology* 81.4 (1979), pp. 645–670.

[78] Derek J Smith, Stephanie Forrest, Ron R Hightower, and Alan S Perelson. "Deriving shape space parameters from immunological data." In: *Journal of theoretical biology* 189.2 (1997), pp. 141–50.

[79] Rustom Antia, Vitaly V Ganusov, and Rafi Ahmed. "The role of models in understanding CD8+ T-cell memory." In: *Nature Reviews Immunology* 5.2 (2005), pp. 101–111.

[80] Arup K. Chakraborty. "A Perspective on the Role of Computational Models in Immunology." In: *Annual Review Of Immunology* 35 (2017), pp. 403–39.

[81] J. Doyne Farmer, Norman H. Packard, and Alan S. Perelson. "The immune system, adaptation, and machine learning." In: *Physica D: Nonlinear Phenomena* 22.1-3 (1986), pp. 187–204.

[82] Stephanie Forrest, Brenda Javornik, Robert E Smith, and Alan S Perelson. "Using genetic algorithms to explore pattern recognition in the immune system." In: *Evolutionary computation* 1.3 (1993), pp. 191–211.

[83] Stephanie Forrest, A.S. Perelson, Lawrence Allen, and R. Cherukuri. "Self-nonself discrimination in a computer." In: *Proceedings of 1994 IEEE Computer Society Symposium on Research in Security and Privacy* (1994), pp. 202–212.

[84] Steven A Hofmeyr and Stephanie Forrest. "Architecture for an Artificial Immune System." In: *Evolutionary Computation* 8.4 (2000), pp. 443–473.

[85] Andrej Kosmrlj, Abhishek K Jha, Eric S Huseby, Mehran Kardar, and Arup K Chakraborty. "How the thymus designs antigen-specific and self-tolerant T cell receptor sequences." In: *Proceedings of the National Academy of Sciences of the United States of America* 105.43 (2008), pp. 16671–6.

[86] Andrej Košmrlj, Mehran Kardar, and Arup K. Chakraborty. "Statistical Physics of T-Cell Development and Pathogen Specificity." In: *Annual Review of Condensed Matter Physics* 4.1 (2013), pp. 339–360.

[87] Minyi Lee, Judith N. Mandl, Ronald N. Germain, and Andrew J. Yates. "The race for the prize: T-cell trafficking strategies for optimal surveillance." In: *Blood* 120.7 (2012), pp. 1432–1438.

[88] M van Baalen. "Coevolution of recovery ability and virulence." In: *Proceedings of the Royal Society of London Series B-Biological Sciences* 265.1393 (1998), pp. 317–325.

[89] Martin R. Miller, Andrew White, and Michael Boots. "Host life span and the evolution of resistance characteristics." In: *Evolution* 61.1 (2007), pp. 2–14.

[90] Ben M. Sadd and Paul Schmid-Hempel. "Perspective: principles of ecological immunology." In: *Evolutionary Applications* 2.1 (2009), pp. 113–121.

[91] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge, UK: Cambridge University Press, 2004.

[92] Alexander Shapiro and Andy Philpott. "A tutorial on stochastic programming." In: *Preprint available at isye.gatech.edu* (2007), pp. 1–35.

[93] Dimitri P Bertsekas. *Nonlinear programming*. Athena Scientific, Belmont, MA, USA, 1999.

[94] S J Wright and J Nocedal. *Numerical optimization*. Vol. 2. Springer, 1999.

[95] Tamara G Kolda, Robert M Lewis, and Virginia Torczon. "Optimization by Direct Search: New Perspectives on Some Classical and Modern Methods." In: *SIAM Review* 45.3 (2003), pp. 385–482.

[96] A R Conn, K Scheinberg, and L N Vicente. *Introduction to derivative-free optimization*. 2009.

[97] Neal Parikh and Stephen Boyd. "Proximal Algorithms." In: *Foundations and Trends in Optimization* 1.3 (2013), pp. 123–231.

[98] John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. "Efficient projections onto the L1 -ball for learning in high dimensions." In: *Proceedings of the 25th international conference on Machine learning - ICML* (2008), pp. 272–279.

[99] Amir Beck and Marc Teboulle. "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems." In: *SIAM Journal on Imaging Sciences* 2.1 (2009), pp. 183–202.

[100] Edward J Anderson and Michael C. Ferris. "A Direct Search Algorithm for Optimization with Noisy Function Evaluations." In: *SIAM Journal on Optimization* 11.3 (2001), pp. 837–857.

[101] Geng Deng and Michael C Ferris. "Variable-number sample-path optimization." In: *Mathematical Programming* 117.1-2 (2009), pp. 81–109.

[102] James C Spall. "Implementation of the Simulteaneous Perturbation Algorithm for Stochastic Optimization." In: *IEEE Transactions on Aerospace and Electronic Systems* 34.3 (1998), pp. 817–823.

[103] NG van Kampen. *Stochastic Processes in Physics and Chemistry*. 2007.

[104] Crispin W Gardiner et al. *Handbook of stochastic methods*. Vol. 4. Springer Berlin, 1985.

[105] Peter E. Kloeden and Eckhard Platen. *Numerical Solution of Stochastic Differential Equations*. Vol. 23. 1. Berlin Heidelberg: Springer-Verlag, 1992, pp. 1–7.

[106] Kimura. "Diffusion models in population genetics." In: *Journal of Applied Probability* 1.2 (1964), pp. 177–232.

[107] Warren J. Ewens. *Mathematical Population genetics*. 2nd. New York: Springer, 2004.

[108] G. J. Baxter, R. A. Blythe, and A. J. McKane. "Exact solution of the multi-allelic diffusion model." In: *Mathematical Biosciences* 209.1 (2007), pp. 124–170.

[109] Richard Levins. *Evolution in changing environments: some theoretical explorations*. 1968, p. 120.

[110] J Seger and H J Brockmann. "What is bet-hedging?" In: *Oxford Surveys in Evolutionary Biology* 4 (1987), pp. 182–211.

[111] Edo Kussell and Stanislas Leibler. "Phenotypic diversity, population growth, and information in fluctuating environments." In: *Science* 309.5743 (2005), pp. 2075–2078.

[112] Luis Miguel Chevin, Russell Lande, and Georgina M. Mace. "Adaptation, plasticity, and extinction in a changing environment: Towards a predictive theory." In: *PLoS Biology* 8.4 (2010).

[113] Andrew M Simons. "Modes of response to environmental change and the elusive empirical evidence for bet hedging." In: *Proceedings of the Royal Society of London B: Biological Sciences* 278.1712 (2011), pp. 1601–1609.

[114] Olivier Rivoire. "Informations in Models of Evolutionary Dynamics." In: *Journal of Statistical Physics* 162.5 (2016), pp. 1324–1352.

[115] Dror A. Vinkler, Haim H. Permuter, and Neri Merhav. "Analogy between gambling and measurement-based work extraction." In: *Journal of Statistical Mechanics: Theory and Experiment* (2016), p. 043403.

[116] Tetsuya J. Kobayashi and Yuki Sughiyama. "Fluctuation Relations of Fitness and Information in Population Dynamics." In: *Physical Review Letters* 115.23 (2015), pp. 1–5.

[117] Olivier Rivoire and Stanislas Leibler. "The value of information for populations in varying environments." In: *Journal of Statistical Physics* 142.6 (2011), pp. 1124–1166.

[118] Dan Cohen. "Optimizing reproduction in a randomly varying environment." In: *Journal of Theoretical Biology* 12.1 (1966), pp. 119–129.

[119] Z H E Chen. "Bayesian Filtering: From Kalman Filters to Particle Filters, and Beyond." In: *Statistics* 182.1 (2003), pp. 1–69.

[120] David J C MacKay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003.

[121] James O. Berger. *Statistical Decision Theory and Bayesian Analysis*. 2nd. Springer-Verlag, 1985.

[122] William H Press, Saul a Teukolsky, William T Vetterling, and Brian P Flannery. *Numerical recipes 3rd edition: The art of scientific computing*. Vol. 1. Cambridge University Press, 2007, p. 1262.

[123] Walter Greiner. *Classical mechanics: Systems of particles and hamiltonian dynamics*. 2010, pp. 1–579.

[124] Steven H. Strogatz. *Nonlinear Dynamics and Chaos*. 1994.

[125] Paul M Chaikin and Tom C Lubensky. *Principles of condensed matter physics*. Vol. 1. Cambridge University Press, Cambridge, UK, 2000.

[126] Jean-Pierre Hansen and Ian R I.R McDonald. *Theory of simple liquids*. November. Elsevier, 1990.

[127] Salvatore Torquato and Frank H Stillinger. "Local density fluctuations, hyperuniformity, and order metrics." In: *Physical Review E* 68 (2003), p. 041113.

[128] Salvatore Torquato. "Hyperuniformity and its generalizations." In: *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 94.2 (2016), pp. 1–21.

[129] Aleksandar Donev, Frank H Stillinger, and Salvatore Torquato. "Unexpected density fluctuations in jammed disordered sphere packings." In: *Physical Review Letters* 95.9 (2005), p. 90604.

[130] Ludovic Berthier, Pinaki Chaudhuri, Corentin Coulais, Olivier Dauchot, and Peter Sollich. "Suppressed compressibility at large scale in jammed packings of size-disperse spheres." In: *Physical Review Letters* 106.12 (2011), pp. 1–4.

[131] Stéfan Van Der Walt, S. Chris Colbert, and Gaël Varoquaux. "The NumPy array: A structure for efficient numerical computation." In: *Computing in Science and Engineering* 13.2 (2011), pp. 22–30.

[132] Travis E. Oliphant. "Python for scientific computing." In: *Computing in Science and Engineering* 9.3 (2007), pp. 10–20.

[133] J D Hunter. "Matplotlib: A 2D graphics environment." In: *Computing In Science & Engineering* 9.3 (2007), pp. 90–95.

[134] S Behnel, R Bradshaw, C Citro, L Dalcin, D S Seljebotn, and K Smith. "Cython: The Best of Both Worlds." In: *Computing in Science Engineering* 13.2 (2011), pp. 31–39.

[135] Roger D. Peng. "Reproducible Research in Computational Science." In: *Science* 334.6060 (2011), p. 1226.

[136] Geir Kjetil Sandve, Anton Nekrutenko, James Taylor, and Eivind Hovig. "Ten Simple Rules for Reproducible Computational Research." In: *PLoS Computational Biology* 9.10 (2013), pp. 1–4.

[137] Andreas Mayer. "Noisyopt: A Python library for optimizing noisy functions." In: *Journal of Open Source Software* 2.13 (2017).

[138] Andreas Mayer, Vijay Balasubramanian, Thierry Mora, and Aleksandra M Walczak. "How a well-adapted immune system is organized." In: *Proceedings of the National Academy of Sciences of the United States of America* 112.19 (2015), pp. 5950–5955.

[139] Andreas Mayer, Thierry Mora, Olivier Rivoire, and Aleksandra M Walczak. "Diversity of immune strategies explained by adaptation to pathogen statistics." In: *Proceedings of the National Academy of Sciences* 113.31 (2016), pp. 8630–8635.

[140] Andreas Mayer, Thierry Mora, Olivier Rivoire, and Aleksandra M. Walczak. "Transitions in optimal adaptive strategies for populations in fluctuating environments." In: *arXiv preprint arXiv:1703.09780* (2017), pp. 1–18.

[141] Thierry Mora, Aleksandra M Walczak, William Bialek, and Curtis G Callan. "Maximum entropy models for antibody diversity." In: *Proceedings of the National Academy of Sciences* 107.12 (2010), pp. 5405–5410.

[142] Kevin Larimore, Michael W McCormick, Harlan S Robins, and Philip D Greenberg. "Shaping of Human Germline IgH Repertoires Revealed by Deep Sequencing." In: *The Journal of Immunology* 189.6 (2012), pp. 3221–30.

[143] Roland R Regoes, Daniel L Barber, Rafi Ahmed, and Rustom Antia. "Estimation of the rate of killing by cytotoxic T lymphocytes in vivo." In: *Proceedings of the National Academy of Sciences* 104.5 (2007), pp. 1599–1603.

[144] Xuhua Xia. "How optimized is the translational machinery in Escherischia coli, Salmonella typhimirium and Saccharomyces cerevisiae?" In: *Genetics* 149 (1998), pp. 37–44.

[145] William H Press. "Strong profiling is not mathematically optimal for discovering rare malfeasors." In: *Proceedings of the National Academy of Sciences of the United States of America* 106.6 (2009), pp. 1716–1719.

[146] Guido Silvestri, Donald L. Sodora, Richard A. Koup, Mirko Paiardini, Shawn P. O'Neil, Harold M. McClure, Silvija I. Staprans, and Mark B. Feinberg. "Nonpathogenic SIV infection of sooty mangabeys is characterized by limited bystander immunopathology despite chronic high-level viremia." In: *Immunity* 18.3 (2003), pp. 441–452.

[147] Michael E. Birnbaum et al. "Deconstructing the peptide-MHC specificity of T cell recognition." In: *Cell* 157.5 (2014), pp. 1073–1087.

[148] Eric S. Huseby, Janice White, Frances Crawford, Tibor Vass, Dean Becker, Clemencia Pinilla, Philippa Marrack, and John W. Kappler. "How the {T} cell repertoire becomes peptide and {MHC} specific." In: *Cell* 122.2 (2005), pp. 247–260.

[149] V Bhardwaj, V Kumar, H M Geysen, and E E Sercarz. "Degenerate recognition of a dissimilar antigenic peptide by myelin basic protein-reactive T cells. Implications for thymic education and autoimmunity." In: *Journal of immunology* 151.9 (1993), pp. 5000–10.

[150] Richard A Miller. "The Aging Immune System: Primer and Prospectus." In: *Science* 273.5271 (1996), pp. 70–74.

[151] Christopher Vollmers, Rene V Sit, Joshua A Weinstein, Cornelia L Dekker, and Stephen R Quake. "Genetic measurement of memory B-cell recall using antibody repertoire sequencing." In: *Proceedings of the National Academy of Sciences of the United States of America* 110.33 (2013), pp. 13463–8.

[152] Gašper Tkačik, Aleksandra M. Walczak, and William Bialek. "Optimizing information flow in small genetic networks." In: *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 80.3 (2009), pp. 1–18.

[153] Tim Gollisch and Markus Meister. "Eye Smarter than Scientists Believed: Neural Computations in Circuits of the Retina." In: *Neuron* 65.2 (2010), pp. 150–164.

[154] Linda Buck and Richard Axel. "A novel multigene family may encode odorant receptors: A molecular basis for odor recognition." In: *Cell* 65.1 (1991), pp. 175–187.

[155] Vijay Balasubramanian and Peter Sterling. "Receptive fields and functional architecture in the retina." In: *The Journal of physiology* 587.Pt 12 (2009), pp. 2753–2767.

[156] Robert Macarthur and Richard Levins. "The Limiting Similarity, Convergence, and Divergence of Coexisting Species." In: *The American Naturalist* 101.921 (1967), p. 377.

[157] Simone Pigolotti, Cristóbal López, and Emilio Hernández-García. "Species clustering in competitive {L}otka-{V}olterra models." In: *Physical Review Letters* 98.25 (2007), p. 258101.

[158] Péter Szabó and Géza Meszéna. "Limiting similarity revisited." In: *Oikos* 112.3 (2006), pp. 612–619.

[159] Jiankui He and Michael W. Deem. "Heterogeneous diversity of spacers within CRISPR (Clustered Regularly Interspaced Short Palindromic Repeats)." In: *Physical Review Letters* 105.12 (2010).

[160] R Ahmed and D Gray. "Immunological memory and protective immunity: understanding their relation." In: *Science* 272.5258 (1996), pp. 54–60.

[161] Rustom Antia, Sergei S Pilyugin, and Rafi Ahmed. "Models of immune memory: On the role of cross-reactive stimulation, competition, and homeostasis in maintaining immune memory." In: *Immunology* 95.December (1998), pp. 14926–14931.

[162] Federica Sallusto, Antonio Lanzavecchia, Koichi Araki, and Rafi Ahmed. "From vaccines to memory and back." In: *Immunity* 33.4 (2010), pp. 451–463.

[163] Theodore J Perkins and Peter S Swain. "Strategies for cellular decision-making." In: *Molecular systems biology* 5.326 (2009), p. 326.

[164] Tetsuya J. Kobayashi. "Implementation of dynamic bayesian decision making by intracellular kinetics." In: *Physical Review Letters* 104.22 (2010), pp. 1–4.

[165] Eric D Siggia and Massimo Vergassola. "Decisions on the fly in cellular sensory systems." In: *Proceedings of the National Academy of Sciences* (2013), pp. 3704–3712.

[166] David A. Sivak and Matt Thomson. "Environmental Statistics and Optimal Regulation." In: *PLoS Computational Biology* 10.9 (2014).

[167] Michael R DeWeese and Anthony M Zador. "Asymmetric Dynamics in Optimal Variance Adaptation." In: *Neural Computation* 10.5 (1998), pp. 1179–1202.

[168] Sophie Deneve. "Bayesian spiking neurons I: inference." In: *Neural computation* 20.1 (2008), pp. 91–117.

[169] Barry Wark, Adrienne Fairhall, and Fred Rieke. "Timescales of Inference in Visual Adaptation." In: *Neuron* 61.5 (2009), pp. 750–761.

[170] Alison Etheridge. *Some Mathematical Models from Population Genetics*. Lecture No. Springer-Verlag, 2012, p. 119.

[171] Andrew Gelman, John B. Carlin, Hal S. Stern, and Donald B. Rubin. *Bayesian data analysis*. Boca Raton: CRC Press, 2004.

[172] R.C. Griffiths. "A Transition Density Expansion for a Multi-Allele Diffusion Model." In: *Advances in Applied Probability* 11.2 (1979), pp. 310–325.

[173] Robert C Griffiths and Dario Spanò. "Diffusion processes and coalescent trees." In: *arXiv preprint arXiv:1003.4650* (2010).

[174] Yun S. Song and Matthias Steinrücken. "A simple method for finding explicit analytic transition densities of diffusion processes with general diploid selection." In: *Genetics* 190.3 (2012), pp. 1117–1129.

[175] Matthias Steinrücken, Y. X Rachel Wang, and Yun S. Song. "An explicit transition density expansion for a multi-allelic Wright-Fisher diffusion with general diploid selection." In: *Theoretical Population Biology* 83.1 (2013), pp. 1–14.

[176] Lenka Zdeborová and Florent Krzakala. "Statistical physics of inference: Thresholds and algorithms." In: *arXiv preprint arXiv:1511.02476* 8732.i (2015), pp. 1–62.

[177] Paul Schmid-Hempel. "Evolutionary ecology of insect immune defenses." In: *Annu Rev Entomol* 50 (2005), pp. 529–551.

[178] Robert L Lochmiller and Charlotte Deerenberg. "Trade-Offs in Evolutionary Immunology: Just What Is the Cost of Immunity?" In: *Oikos* 88.1 (2000), pp. 87–98.

[179] Kurt a McKean, Christopher P Yourth, Brian P Lazzaro, and Andrew G Clark. "The evolutionary costs of immunological maintenance and deployment." In: *BMC Evolutionary Biology* 8.8 (2008), p. 76.

[180] Bergljót Magnadóttir. "Innate immunity of fish (overview)." In: *Fish & Shellfish Immunology*. Vol. 20. 2. Elsevier, 2006, pp. 137–151.

[181] Jonathan D.G. Jones and Jeffery L Dangl. "The plant immune system." In: *Nature* 444.7117 (2006), pp. 323–329.

[182] Michael R. Strand. "The insect cellular immune response." In: *Insect Science* 15.1 (2008), pp. 1–14.

[183] Tom J. Little, Dan Hultmark, and Andrew F Read. "Invertebrate immunity and the limits of mechanistic immunology." In: *Nature immunology* 6.7 (2005), pp. 651–654.

[184] Philippe Horvath and Rodolphe Barrangou. "CRISPR/Cas, the immune system of bacteria and archaea." In: *Science* 327.5962 (2010), pp. 167–170.

[185] Colin D. Malone and Gregory J. Hannon. "Small RNAs as Guardians of the Genome." In: *Cell* 136.4 (2009), pp. 656–668.

[186] Olivier Rivoire and Stanislas Leibler. "A model for the generation and transmission of variations in evolution." In: *Proceedings of the National Academy of Sciences* 111.19 (2014), E1940–9.

[187] M Gniadkowski. "Evolution of extended-spectrum beta-lactamases by mutation." In: *Clinical microbiology and infection* 14 Suppl 1 (2008), pp. 11–32.

[188] Eric Vivier, David H Raulet, Alessandro Moretta, Michael A Caligiuri, Laurence Zitvogel, Lewis L Lanier, Wayne M Yokoyama, and Sophie Ugolini. "Innate or Adaptive Immunity? The Example of Natural Killer Cells." In: *Science* 331.6013 (2011), pp. 44–49.

[189] Steven H Spoel and Xinnian Dong. "How do plants achieve immunity? Defence without specialized immune cells." In: *Nature reviews. Immunology* 12.2 (2012), pp. 89–100.

[190] Olivier Voinnet. "RNA silencing as a plant immune system against viruses." In: *Trends in Genetics* 17.8 (2001), pp. 449–459.

[191] P M Waterhouse, M B Wang, and T Lough. "Gene silencing as an adaptive defence against viruses." In: *Nature* 411.6839 (2001), pp. 834–42.

[192] Roghiyh Aliyari, Qingfa Wu, Hong Wei Li, Xiao Hong Wang, Feng Li, Lance D. Green, Cliff S. Han, Wan Xiang Li, and Shou Wei Ding. "Mechanism of Induction and Suppression of Antiviral Immunity Directed by Virus-Derived Small RNAs in Drosophila." In: *Cell Host and Microbe* 4.4 (2008), pp. 387–397.

[193] Oded Rechavi, Gregory Minevich, and Oliver Hobert. "Transgenerational inheritance of an acquired small RNA-based antiviral response in C. elegans." In: *Cell* 147.6 (2011), pp. 1248–1256.

[194] D R Jones, CD D Perttunen, and BE E Stuckman. "Lipschitzian Optimization Without the Lipschitz Constant." In: *Journal of Optimization Theory and Application* 79.1 (1993), pp. 157–181.

[195] Murat Acar, Jerome T Mettetal, and Alexander van Oudenaarden. "Stochastic switching as a survival strategy in fluctuating environments." In: *Nature genetics* 40.4 (2008), pp. 471–475.

[196] Hubertus J E Beaumont, Jenna Gallie, Christian Kost, Gayle C Ferguson, and Paul B Rainey. "Experimental evolution of bet hedging." In: *Nature* 462.7269 (2009), pp. 90–93.

[197] Guillaume Lambert and Edo Kussell. "Quantifying Selective Pressures Driving Bacterial Evolution Using Lineage Analysis." In: *Physical Review X* 5.1 (2015), p. 011016.

[198] M Lachmann and E Jablonka. "The inheritance of phenotypes: an adaptation to fluctuating environments." In: *Journal of theoretical biology* 181.1 (1996), pp. 1–9.

[199] Carlos A. Botero, Franz J. Weissing, Jonathan Wright, and Dustin R. Rubenstein. "Evolutionary tipping points in the capacity to adapt to environmental change." In: *Proceedings of the National Academy of Sciences* 112.1 (2015), pp. 184–189.

[200] M.C. Donaldson-Matasci, Michael Lachmann, and C.T. Bergstrom. "Phenotypic diversity as an adaptation to environmental uncertainty." In: *Evolutionary Ecology Research* 10.4 (2008), pp. 493–515.

[201] Marcel Salathé, Jeremy Van Cleve, and Marcus W. Feldman. "Evolution of stochastic switching rates in asymmetric fitness landscapes." In: *Genetics* 182.4 (2009), pp. 1159–1164.

[202] Pintu Patra and Stefan Klumpp. "Emergence of phenotype switching through continuous and discontinuous evolutionary transitions. TL - 12." In: *Physical biology* 12.4 (2015), p. 46004.

[203] Antun Skanata and Edo Kussell. "Evolutionary Phase Transitions in Random Environments." In: *Physical Review Letters* 117.3 (2016), p. 038104.

[204] Oana Carja, Uri Liberman, and Marcus W Feldman. "Evolution in changing environments: modifiers of mutation, recombination, and migration." In: *Proceedings of the National Academy of Sciences of the United States of America* 111.50 (2014), pp. 17935–40.

[205] Bingkan Xue and Stanislas Leibler. "Evolutionary learning of adaptation to varying environments through a transgenerational feedback." In: *Proceedings of the National Academy of Sciences of the United States of America* 113.40 (2016), pp. 11266–11271.

[206] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Hoboken, NJ: Wiley, 2005, pp. 1–748.

[207] Yoh Iwasa and Patsy Haccou. "Optimal mixed strategies in stochastic environments." In: *Theoretical Population Biology* 47 (1995), pp. 212–243.

[208] William C. Ratcliff, Peter Hawthorne, and Eric Libby. "Courting disaster: How diversification rate affects fitness under risk." In: *Evolution* 69.1 (2015), pp. 126–135.

[209] Bernadett Gaál, Jonathan W. Pitchford, and A. Jamie Wood. "Exact results for the evolution of stochastic switching in variable asymmetric environments." In: *Genetics* 184.4 (2010), pp. 1113–1119.

[210] BingKan Xue and Stanislas Leibler. "Bet-hedging against demographic variations." In: *arXiv preprint arXiv:1701.00523* (2017).

[211] Charles K. Fisher, Thierry Mora, and Aleksandra M. Walczak. "Variable habitat conditions drive species covariation in the human microbiota." In: *PLoS Computational Biology* 13.4 (2017), e1005435.

[212] Mikhail V Pogorelyy et al. "Persisting fetal clonotypes influence the structure and overlap of adult human T cell receptor repertoires." In: *PLoS Computational Biology* 13.7 (2017), e1005572.

[213] George C Wang, Pradyot Dash, Jonathan A McCullers, Peter C Doherty, and Paul G Thomas. "T cell receptor αβ diversity inversely correlates with pathogen-specific antibody levels in human cytomegalovirus infection." In: *Science translational medicine* 4.128 (2012), 128ra42.

[214] Shu-Qi Zhang et al. "Direct measurement of T cell receptor affinity and sequence from naïve antiviral T cells." In: *Science Translational Medicine* 8.341 (2016), 341ra77–341ra77.

[215] Bernd R. Schöne, Jens Fiebig, Miriam Pfeiffer, Renald Gleß, Jonathan Hickson, Andrew L A Johnson, Wolfgang Dreyer, and Wolfgang Oschmann. "Climate records from a bivalved Methuselah (Arctica islandica, Mollusca; Iceland)." In: *Palaeogeography, Palaeoclimatology, Palaeoecology* 228.1-2 (2005), pp. 130–148.

[216] Ilía Nikolaevich Bronshtein, Konstantin A Semendyayev, Gerhard Musiol, and Heiner Muehlig. *Handbook of mathematics*. Vol. 3. Springer, Berlin Heidelberg, Germany, 2007.

[217] Soumik Pal. "Wright-fisher diffusion with negative mutation rates." In: *Annals of Probability* 41.2 (2013), pp. 503–526.

Abstract

Biological organisms have evolved diverse immune mechanisms to defend themselves against pathogens. Here we build mathematical models of immune systems optimally tuned to the statistics of pathogens.

Beyond molecular details, different immune mechanisms differ in how protection is acquired, processed and passed on to subsequent generations – differences that may be essential to long-term survival. To explain the observed diversity of strategies we compare the long-term adaptation of populations as a function of the pathogen dynamics that they experience and of the immune strategy that they adopt. We find that the two key determinants of an optimal immune strategy are the frequency and the characteristic timescale of the pathogens. Depending on these two parameters, we identify distinct modes of immunity, including adaptive, innate, bet-hedging and CRISPR-like immunities, which recapitulate the diversity of natural immune systems. Our results carry over to the general question of evolution in fluctuating environments, for which we provide novel analytical results in temporally correlated environments.

The adaptive immune system provides protection through a broad repertoire of cells specific to different pathogens. To predict statistical features of well-adapted repertoires we analyze which repertoire minimizes cost of infection for a given distribution of pathogens. The theory predicts that the immune system has more receptors for rare antigens than expected from the frequency of encounters; and individuals exposed to the same infections have sparse repertoires that are largely different, but nevertheless exploit cross-reactivity to provide the same coverage of antigens. Our results follow from a tension between the statistics of pathogen detection, which favor a broader receptor distribution, and the effects of cross-reactivity, which tend to concentrate the optimal repertoire onto a few highly abundant clones. These predictions can be tested in high throughput surveys of receptor and pathogen diversity. We then explicitly consider how the adaptive immune system can learn the statistics of the environments from its past infection history in a Bayesian manner. We show that optimal repertoires can be reached by keeping memory of an infection through the selective proliferation of stimulated cells. The Bayesian perspective on repertoire dynamics provides an unifying conceptual framework to explain a number of features of immunological memory and suggests further experiments.

Keywords

biophysics, immunology, statistical physics, evolution

Résumé

Les organismes biologiques ont développé divers mécanismes immunitaires afin de se protéger des pathogènes. Nous développons ici des modèles mathématiques de systèmes immunitaires, adaptés de façon optimale aux statistiques des pathogènes.

Au delà des détails moléculaires, ces mécanismes immunitaires diffèrent dans la manière d'acquérir, de réguler et de transmettre une protection immunitaire; différences qui pourraient s'avérer essentielles pour la survie à long terme. Afin d'expliquer la diversité des stratégies qui sont observées, nous comparons l'adaptation à long terme de populations en fonction de la dynamique des pathogènes à laquelle elles sont confrontées et de la stratégie immunitaire qu'elles adoptent. Nous démontrons que la fréquence et l'échelle de temps caractéristique des pathogènes sont les deux déterminants clés d'une stratégie immunitaire optimale. En fonction de ces deux paramètres, nous identifions des modes d'immunité distincts, comprenant immunités innées, adaptatives, ou ressemblant au système CRISPR, qui récapitulent la diversité de systèmes immunitaires naturels. Nos résultats viennent s'étendre à la question générale de l'évolution dans des environnements variables pour laquelle nous apportons de nouveaux résultats analytiques au sein d'environnements temporairement corrélés.

Le système immunitaire adaptatif assure une protection à partir d'un large répertoire de cellules spécifiques à différents pathogènes. Pour prédire des propriétés statistiques de répertoires adaptés, nous étudions quel répertoire minimise au mieux le risque d'infections pour une distribution de pathogènes donnée. La théorie prédit que les cellules spécifiques contre les antigènes rares sont surreprésentées par rapport à la fréquence de leurs rencontres et que les individus, exposés aux mêmes infections, possèdent des répertoires avec des récepteurs largement différents mais exploitent la réactivité croisée afin de parvenir à la même couverture d'antigènes. Nos résultats sont issus d'une opposition entre les statistiques de détection des pathogènes, qui soutiennent l'idée d'une plus large distribution de récepteurs, et les effets de la réactivité croisée, qui tend à concentrer le répertoire optimal sur un petit nombre de clones. Nos prédictions peuvent être testées à partir des données à haut débit sur la diversité des récepteurs et de pathogènes. Par la suite, nous examinons explicitement comment le système immunitaire adaptatif peut apprendre de manière bayésienne les statistiques de l'environnement à partir de l'historique des infections précédentes. Nous montrons que les répertoires optimaux peuvent être atteints par prolifération sélective des cellules spécifiques. La perspective bayésienne sur la dynamique des répertoires fournit un cadre conceptuel unificateur qui explique un certain nombre de caractéristiques de la mémoire immunitaire et appelle à des expériences complémentaires.

Mots Clés

biophysique, immunologie, physique statistique, évolution