

Rare event simulation for probabilistic models of T-cell activation

Der Technischen Fakultät der Universität Bielefeld
vorgelegt zur Erlangung des akademischen Grades
Doktor der Naturwissenschaften

Dipl.-Inform. Florian Lipsmeier
Bielefeld, Germany, July 21, 2010

Biomathematics & Theoretical Bioinformatics Group
Faculty of Technology
Bielefeld University



Supervisors:

Prof. Dr. Ellen Baake

Prof. Dr. Sven Rahmann

Acknowledgements

Undertaking and finishing a PhD project certainly requires great effort in many respects. Without help, support, and encouragement from several persons, I would never have been able to finish this work.

First of all, I would like to thank my supervisor Prof. Ellen Baake for many interesting and fruitful discussions, for keeping up with my sometimes confusing explanations and for allowing me much freedom in my research.

Thanks to all the members and former members of the Biomathematics and theoretical Bioinformatics group, especially for their patience with my many test talks during our group seminars and their helpful criticism.

Finally, I'd like to thank my family for encouraging me to pursue my studies and especially my companion in life Vera Drees for supporting me with your love and understanding.

I am grateful for the financial support I received from the NRW Graduate School for Bioinformatics and Genome Research.

Bielefeld, March 2010

Florian Lipsmeier

CONTENTS

1	Introduction	1
2	Biological background	4
2.1	The immune system at a glimpse	4
2.1.1	Innate immune system	5
2.1.2	The adaptive immune system	7
2.2	T-Lymphopoiesis	12
2.2.1	The T-cell receptor	14
2.2.2	Positive selection	16
2.2.3	Negative selection	18
2.2.4	T-Lymphopoiesis in numbers	21
2.3	T-cell activation	23
2.3.1	Introduction	24
2.3.2	TCR binding	25
2.3.3	TCR triggering	26
2.3.4	Models of T-cell activation	27
3	BRB model of T-cell activation	31
3.1	Additional remarks	34
3.2	T-cell activation in numbers	35
4	Mathematical background and computational methodology	37
4.1	Large deviation probabilities	38
4.2	Simulating rare event probabilities	39
5	Analysis and extension of the BRB model of T-cell activation	44
5.1	Rare event simulation: The T-cell model	44
5.1.1	Large deviations for independent but not identically distributed random variables	45
5.1.2	Tilting of transformed random variables	46
5.1.3	The algorithm	47
5.2	Results	50
5.2.1	Performance of the simulation method	51
5.2.2	Analysis of the T-cell model	56
5.3	Negative Selection	63
5.3.1	BRB model with negative selection	64
5.3.2	Simulation method	68

5.3.3	Results	70
5.3.4	Discussion	78
6	A discrete T-cell activation model	82
6.1	The model	83
6.2	Simulation approach	86
6.3	Results	94
6.4	Discussion	111
7	A model for T-cell migration in the thymic medulla	115
7.1	Simulation method	118
7.2	Results	119
7.3	Discussion	125
8	Summary	127
8.1	Outlook	129
	Bibliography	131

NOMENCLATURE

$\bar{\tau}$	mean of the exponentially distributed mean binding time between an antigen and a TCR
η	relative error; standard deviation divided by the mean
\mathcal{T}	random variable for the mean binding time between an antigen and a TCR
τ	binding time between an antigen and a TCR
ϑ	tilting parameter
g	density of W
$G(z^{(f)})$	random variable for the total stimulus induced by all the antigens on an APC to a random T-cell, as a function of the copy number of the foreign antigen
g_{act}	activation threshold; value that has to be reached by $G(z^{(f)})$ in order for a T-cell to get activated
g_{thy}	thymic activation threshold; if this value is reached during negative selection, the T-cell is induced to die
$m^{(c)}$	number of constitutive antigen types presented on an APC
$m^{(v)}$	number of variable antigen types presented on an APC
W	stimulation rate induced by a single antigen to a random T-cell
w	realisation of W
$z^{(c)}$	number of copies of an individual constitutive antigen type presented on an APC
$z^{(f)}$	number of copies of the foreign antigen type presented on an APC
$z^{(v)}$	number of copies of an individual variable antigen type presented on an APC
z_{s}	number of copies of a self antigen type presented on an APC in the new T-cell activation model
AIRE	autoimmune regulator; responsible for pGE regulation

APC	antigen presenting cell
BCR	B-cell receptor
DC	dendritic cell
IL	interleukin; group of different second messengers
IS	importance sampling
LD	large deviation
MHC	major histocompatibility complex
pGE	promiscuous gene expression
pMHC	peptide-MHC; complex of an antigen and an MHC presented on an APC
SS	simple sampling
T_C1, T_C2, T_C17	different types of cytotoxic T-cells
T_H1, T_H2	different types of T helper cells
TCR	T-cell receptor
TEC	thymic epithelial cell
TRA	tissue restricted antigen

CHAPTER 1

INTRODUCTION

When it comes to T-cells in immunobiology a summary of the standard school book descriptions basically looks like as follows. There are cells of the innate immune system that directly recognise and attack a pathogen and there are the cells of the adaptive immune system that are able to learn to recognise and attack all the pathogens that circumvent the innate immune system. One important family of cells of the adaptive immune system are the so-called T-cells. These cells have a special receptor that helps them to detect molecules on cell-surfaces called antigens. When a pathogen enters the body, it leaves a trail of antigens which are recognised by one of the T-cells and this T-cell clears the body of the pathogen.

So, where is the problem one might think. This seems like a pretty straightforward description of a mechanism that is completely understood. However, as we will show, quite the opposite is true. It is possible to give some general, oversimplifying explanations but most of the rest is still very unclear. With this thesis we provide insights and new ideas to illuminate the T-cell's ability to detect pathogens. Let us make clear why this is so important and justifies intense research.

If we observe ourselves and our environment by means of a microscope, we see that we constantly interact in many different ways with many different microorganisms at any given time. This can happen in a symbiotic way as it is with the bacteria in our colon that support digestion or with the bacteria on our skin. Often, however, it is a hostile interaction, that is bacteria, fungi or viruses try to invade our body to use it for their survival and for their reproduction. Thereby they can make us sick or even kill us, eventually.

Most people are living a healthy life, at least over long periods of their lifetime. We rarely get ill and if so, we normally get better quite fast. This is a quite amazing observation, though none of us thinks a lot about it. It is so amazing because our body is under constant attack from a huge variety of pathogenic microbes, like bacteria and viruses. A fact hardly acknowledged by anyone, as we only get to know an attack if it is successful. The counterpart to these microbes in our body is the so called immune system. This highly complex system of many different cells enables the recognition and killing of pathogens in most of the cases.

One central part of this immune system are the so-called T-cells. Since decades much research is devoted not only to T-cells but of course also to all other cells of our immune system. This led to new drugs and therapies in order to prevent diseases. Despite all these efforts much is still unknown.

T-cells belong to the part of the immune system that is termed adaptive because

they have the ability to recognise unknown pathogens and in case of a second meeting with the same pathogen they are able to react much faster. T-cells can on the one hand support other cell types in the direct attack of a pathogen or on the other hand detect and destroy cells of our own body which are infiltrated by pathogens in order to reproduce. They are even able to detect mutated (cancer) cells to a certain degree and destroy them. Hence, it is highly advantageous to know all details regarding these cells.

Surprisingly, there is still a central aspect with regard to T-cells that can only be explained insufficiently. T-cells are activated via the recognition of short amino acid sequences that are displayed on cell surfaces in the body. These sequences are residues from the degradation cycle of the cells and are called antigens. On this molecular level there is no definite characteristic that identifies an antigen as foreign, that is as coming from a pathogen, or as self, that is coming from our own cells. The T-cell constantly 'scans' for antigens and mostly of course they encounter self antigens. A reaction to self antigens would lead to autoimmune reactions. A missing reaction to foreign antigens on the other hand can lead to severe diseases and death.

In order to explain the mystery of the recognition of a foreign antigen against a background of many different self antigens, experiments only are not sufficient. Therefore, there are efforts to use the experimental knowledge as a basis for the development of mathematical models that explain foreign-self discrimination.

At this intersection of immunology, mathematics and computer science this thesis is situated. We use a combination of mathematical modeling on the basis of biological hypothesis and an efficient simulation method to explore the mechanism of T-cell activation with emphasis on its foreign-self discrimination capability. We introduce some of the already existing models in the next chapter and then concentrate on one special model of T-cell activation developed by van den Berg, Rand and Burroughs (BRB model of T-cell activation) in 2001 [205]. In contrast to the other models, this model describes T-cell activation probabilistically. We will explain why this is necessary and beneficial. Previous work could show that this model is capable of explaining foreign-self discrimination [205, 232]. This model is our starting point for the exploration of T-cell activation and foreign-antigen discrimination. From here we go on to extensions of the model and the development of a related model that captures the biological reality in a better way.

An integral part of this thesis is the development of a general and efficient simulation method for a certain type of stochastic models. We use this method for the analysis of the BRB model and its extensions. The development of a new simulation method is necessary because, as we pointed out, T-cells mainly meet self antigens and only very rarely foreign antigens. Hence, the recognition and activation of a T-cell has to be a rare event. But in order to investigate such rare events in detail, sophisticated methods are needed.

The other important part of this thesis is the development and analysis of a new T-cell activation model based on our results from the analysis of the BRB model and additional recent experimental findings. In this model we include a special 'educational' mechanism during T-cell development termed 'negative selection'. Plainly speaking, this mechanism helps to sort out T-cells that are too self-reactive. We succeed in combining negative

selection with the central ideas from the BRB model and show its major influence on foreign-self discrimination.

Because of its importance we devote an additional chapter to a first modeling approach of T-cell migration in the thymus, the place where negative selection occurs. This opens up a new direction of research that should on the one hand clarify the scope of negative selection with regard to the foreign-self discrimination ability of T-cells. On the other hand it should be the first step to the better understanding of negative selection in order to find ways to actually manipulate this process to prevent autoimmune diseases and enhance the effectiveness of T-cells against certain pathogens.

In summary, our present work has three important cornerstones. We introduce all necessary biological details with regard to T-cells that not only suffice as a solid background for this thesis but also as a starting point for future work. Furthermore, we develop a powerful simulation method whose application area goes beyond the models we explore here. Most importantly, we deliver new hypotheses and parameter estimates for T-cell activation, foreign-self discrimination and negative selection via the analysis of the models presented in this thesis. These results can be one step in the direction towards new experimental research in order to finally really explain T-cell activation and therewith foreign-self discrimination.

The thesis is composed as follows:

Chapter 2 – 4 deal with the development of the basic biological and mathematical/informatics knowledge that is needed. In Chapter 2 we therefore introduce the immune system in general and concentrate then on T-cells. Afterwards we describe in Chapter 3 the BRB model of T-cell activation and finally in chapter 4 the basics for our simulation method.

Chapter 5 and 6 represent the core of this thesis. In chapter 5 we develop the simulation method and prove its efficiency. Then we use it to analyse the BRB model. Furthermore we develop and analyse extensions of the model as a consequence of our first results. The sixth chapter is devoted to a new model of T-cell activation which we develop on the basis of our ideas gained in chapter 5. We describe the development of the necessary modified simulation method and test the model with different parameter values in order to explain foreign-self discrimination by T-cells.

Chapter 7 deals with the development of a new model of T-cell migration that should help to clarify the negative selection process. The model we develop in this chapter does not describe T-cell activation but as this mechanism has a vital influence on foreign-self discrimination of T-cells it fits well in the context of this thesis.

Finally we summarise our results in chapter 8 and give an outlook on the implications of our work on future research.

BIOLOGICAL BACKGROUND

In this chapter we elaborate on the immunobiological background of this thesis. In this respect, it is our aim not only to introduce the necessary facts, but we incorporate them in a kind of review. Thereby we want to draw a picture of the newest relevant findings with regard to T-cells and present them in a way that leads to new approaches and ideas also beyond this thesis. We therefore decided to introduce special sections which summarise important data which we extrapolated from different publications as an asset for further model development.

This chapter consists of three parts. At first we give a short, very general introduction on the immune system in order to motivate the role of T-cells in this framework. The second part is concerned with the T-cell development process. This is necessary because it already conveys important information on the establishment of foreign-self discrimination of the T-cell repertoire. Finally, we explore the mechanism of T-cell activation. Besides experimental and theoretical basics, we also introduce different models of T-cell activation that already exist.

2.1 THE IMMUNE SYSTEM AT A GLIMPSE

We and all kinds of (jawed) vertebrates share a similar complex defence system, called the immune system, which protects its host against all sorts of pathogenic microbes. Under the roof of the immune system there is a multi-faceted collection of cells, molecules and their interactions, which enable a specific and non-specific recognition and elimination of a variety of pathogens. Generally, we distinguish between two different parts of the immune system. There is the innate immune system, comprising all innate, that is non-specific, immune responses and there is the adaptive immune system, which comprises all specific immune responses. However, we have to point out that these two parts are entwined with each other in many different ways. Adaptive responses develop as innate responses occur. The innate immune system can therefore be seen as a first response unit, that reacts directly on a pathogen encounter. Furthermore, the adaptive immune system really is adaptive, that is, it allows for the development of a memory and thereby amplifies the reaction efficacy to new encounters with the same pathogen. Just in contrast, responses of the innate immune system do not vary, no matter how often the same pathogen is encountered.

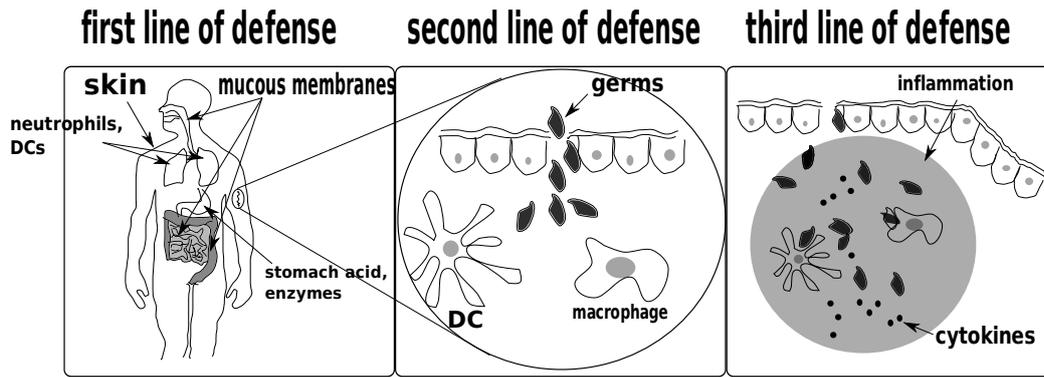


Figure 2.1: The three lines of defence of the innate immune system. The first line tries to prevent the penetration of the body. It consists of a mixture of passive (e.g. skin, stomach acid, membranes) and active (e.g. special types of dendritic cells (DCs), neutrophils, macrophages) elements, where the former ones are just physical barriers and the latter ones are situated at/in this barriers and attack the possible intruder. The second line consists of different types of dendritic cells, neutrophils or macrophages which are situated inside the body but near the surface. Finally, the third line of defence produces an inflammatory environment via cytokines that facilitates DC and macrophage movement and their ability to attack germs.

2.1.1 Innate immune system

We are surrounded by myriads of microorganisms, some of which try to invade our body. To prevent this in the first place, we have different innate lines of defence (see Fig. 2.1) [152, 142]. The most obvious one is the skin and its inner-body equivalent, the mucous membranes (inner-body surfaces which are exposed to the environment). Both form a physical barrier through so called tight junctions, which are firm cell-cell connections. They are furthermore covered with different epithelial cells. Those on our skin form a dry, protecting layer which is hardly penetrable. This is why most of the pathogens try to overcome the mucous membranes instead. Their epithelial cells utilise different mechanisms to prevent this intrusion. They produce mucus to trap microbes, propel them away using cilia and they produce special enzymes and anti-bacterial peptides to eliminate pathogens. Additionally, they facilitate the colonisation of their surface with friendly microorganisms which compete against foreign microbes and can synthesise anti-bacterial substances.

If this first barrier is penetrated, for example by an injury of our skin, the second barrier comes into play. It comprises special cell types which are capable of endo- or phagocytosis, that is the internalisation and degradation of either macromolecules (endocytosis) or whole microorganisms (phagocytosis) [105, 200]. The main cell types responsible for these actions are the dendritic cells (DCs), macrophages and the neutrophils. They identify targets via pattern recognition receptors on their cell surface that bind to surface molecules of invading microorganisms [132, 101].

Before proceeding with the innate immune response it is at this point necessary to have a closer look on receptors and receptor binding, as this is crucial not only here, but also for the adaptive immune system and in particular for this thesis. In order to 'see' their en-

environment all cell types rely on different types of receptors. These are special molecules which are embedded into the membrane of a cell. Mostly they are trans-membrane molecules that detect molecules with their extracellular part and start signalling with their intracellular part. Depending on aspects as different as three-dimensional conformation, charging, amino acid composition and so on, different molecules (normally called ligands, in the special case of immunology we speak of antigens) bind with different affinities to a receptor. If a molecule with a certain affinity binds to a receptor, the receptor is activated and starts its task, which normally involves the start of some type of intracellular signalling cascade. There are quite unspecific receptors which react to many ligands with low binding affinities, as well as very specialised receptors, which only react to one or two well-fitting ligands with very high binding affinities. In general ligands are peptides, hormones, toxins or drugs. In the special case of foreign microbe recognition by T-cells, they are normally just peptide chains (strings of amino acids) which result from a preprocessing of the native (three-dimensional) antigen and are presented by a special class of cells, called antigen presenting cells. As there are 22 different amino acids which can be concatenated to peptide chains in all combinations, the number of possible antigens is very high, which will come into play later. Germ detection by other cell types of the immune system is not so restricted. Their receptors are able to detect antigens in their native, three-dimensional form.

Receptors used in the innate immune response are germline encoded and therefore cannot be changed. That is, every receptor type is hardcoded by a specific gene. This is why there is only a limited number of receptors available in a host. To overcome the disadvantage of this restricted repertoire in the face of the enormous variety of pathogens, these receptors are quite unspecific (also termed crossreactive) concerning the targets they bind (e.g. CD14 is a receptor which binds to all kinds of bacterial lipopolysaccharides). They detect special patterns that are pathogen associated [132]. Thereby they can detect whole families of pathogens that share structural elements which are detected by such a pattern recognition receptor.

A third barrier of the innate immune system is formed by inflammations, which come about if tissue is damaged or if pathogens are recognised [142]. Inflammatory reactions facilitate the movement of effector molecules and cells to the affected area in order to support the killing of the pathogens. Furthermore, the infection is restricted through the healing of the tissue and the building of physical barriers to prevent a further movement of the pathogens.

It is obvious that, although fast and highly effective in killing pathogens, the innate immune system has its Achilles heel in the restricted number of possible receptors for the recognition of pathogenic organisms. In the course of evolution microbes have developed numerous ways to prevent recognition, e.g. by the development of a thick polysaccharide capsule. Even if they are recognised some microbes are capable of (mis-)using the degradation cycle of macrophages for their own purpose and grow within these macrophages. The answer to these threats was the development of a second detection system, which had to be more specific in its recognition mechanism and not restricted to a limited number of receptor types.

2.1.2 *The adaptive immune system*

It is an inherent feature of evolution that through mutation and selection it produces either new variants of a species or even new species. The rate of change of a species is dependent on the length of its reproduction cycle, as during reproduction the important changes occur via genetic mutation and recombination. Most of the microorganisms have very short reproduction cycles, especially compared to vertebrates. Hence, they can change faster and thereby avoid detection. It follows that to prevent bypassing of the host's immune system, there were two possible ways in evolution to go. The first one would have been the expansion of the repertoire of germline encoded receptors. This happened probably during the evolution of invertebrates, which often have a much bigger repertoire of these receptors than vertebrates [151]. The second way has been the development of a new detection mechanism, which is more flexible and specific in its pathogen detection capability. This came about during the evolution of (jawed) vertebrates. Although no one can say for sure whether this way was the optimal one (in terms of selective pressure during evolution), there exists a convincing hypothesis. In vertebrates there is a huge variety of microorganisms, living in symbiosis with the host. In order to guarantee the safety of these symbiotic arrangements, the vertebrates had to deplete their innate receptor repertoire of all receptors which could recognise antigens from these 'friendly' microorganisms. This would have made the host also more vulnerable to other pathogenic microorganisms [151, 163]. The development of a second, adaptive, immune system with a receptor repertoire that is more specific was therefore necessary. This section highlights the most important aspects of this adaptive immune system, with special emphasis on T-cells.

The adaptive immune system has a complex task. It should defend its host against all pathogenic microbes that circumvent the innate immune system and thereby deal with all occurring new mutants of such a microbe, keep it then in a kind of memory to act faster on the occasion of a second infection attempt and of course it should not attack the host itself or microorganisms living in symbiosis with its host. These are exactly the very characteristics of the adaptive immune system [142, 152]. There are two key players involved in adaptive immunity, the so-called lymphocytes and the antigen-presenting cells (APC). Both can enter lymphoid organs and tissues and otherwise circulate around the body by means of the vascular and lymphoid system.

Antigen presenting cells are, as the name tells, highly specialised cells, which can internalise and present antigens. The most important group of APCs are the dendritic cells (see Figure 2.2). Here, we can see one of the several connections between the innate and the adaptive immune system. DCs also play a role in the former one because they internalise pathogenic material, as already mentioned. However, they have developed a very specialised signalling and receptor apparatus that makes them the ideal interaction partner for one group of lymphocytes, so-called T-cells. Another group of APCs are special lymphocytes called B-cells, which besides being APCs serve other important purposes, which is explained in the next paragraph. An APC produces so called major histocompatibility complex molecules (MHC) and acquires antigens. An MHC needs the antigen to form a stable macromolecule, which is then expressed on the surface of an

APC as an antigen (or peptide)-MHC (pMHC) complex. An antigen can only bind to some of the several types of MHC molecules. Currently, for humans 3371 different MHC alleles are known which belong to two different classes (2351 for MHC class I, 1020 for MHC class II) (see www.ebi.ac.uk/ipd for updated numbers), from which only very few are expressed in every single individual [160]. In fact all cells in the body express MHC I molecules on their surface together with fragments out of their interior, but APCs are especially equipped to produce them in great numbers together with many more antigens.

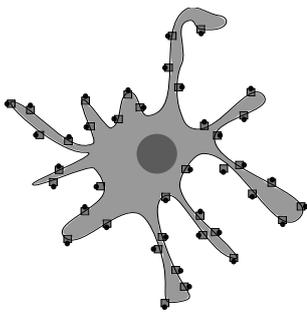


Figure 2.2: A dendritic cell equipped with antigens

The pMHCs on the APCs are scanned by the T-cells. Antigen acquisition is primarily done by the already mentioned mechanisms of phago/endocytosis. The internalised cell fragments are degraded and the resulting very small peptide fragments form the antigens. Figure 2.3 illuminates this, for the present thesis, very important mechanism. The different aspects of it are explained in detail in different sections of this chapter.

At this point it is important to note that an APC has only very limited capabilities of pathogen detection through pattern recognition receptors. In principle it internalises all sorts of cell material from its surroundings. This implies that most of them are parts of dead cells from the host itself. Hence, most of the antigen produced fall into the category of so-called self antigens, that is antigens of the host itself. On the molecular level there exists no distinction between a foreign and a self antigen, that is they are just strings of amino acids with no special marker for pathogenic material. As long as no infection occurs we have to assume that there are even no pathogenic antigens presented by an APC. However, these APC also lack an activation signal that enables them to stimulate T-cells. This signal is supplied via the pattern recognition receptors if they detect pathogens. A signal of these receptors does not only activate the APC, but also leads to an enhanced incorporation of (presumably) foreign antigen [150, 132]. However, it should be clear that even APCs that encounter pathogens will mostly present self antigens because they are flooded with them constantly. This is a crucial observation with regard to the topic of this thesis. The recognition of the limited amount of foreign antigens against a background of many self antigens. A reaction to self antigens would lead to autoimmune reactions and a missing reaction to foreign antigens would let the pathogen invade the body. Furthermore, APCs can co-determine the reaction of the T-cells to the pMHC by means of other types of molecules such as co-receptors or cytokines, which will be elucidated later.

Lymphocytes constitute 20-40% of the white blood cell population. Their two major subsets are the already mentioned T- and B-cells.

As this thesis deals with T-cells, we only briefly outline the function of B-cells. They are responsible for the humoral immune response, that is they induce a reaction on substances in the extracellular fluids [142]. B-cells are generated in the bone marrow, where they also mature. During maturation they develop a special receptor, the B-cell receptor (BCR), which is unique for every B-cell. This is achieved by the genetic rearrangement

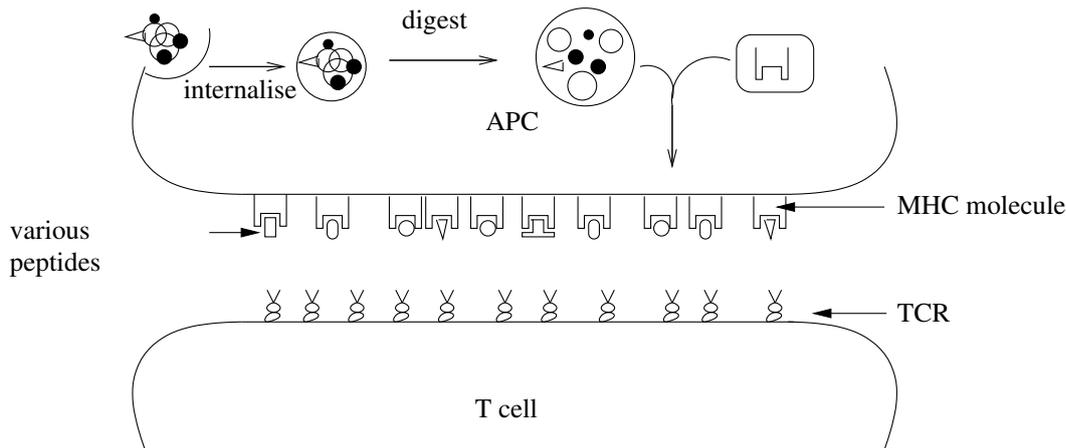


Figure 2.3: A T-cell and an antigen-presenting cell (based on Fig. 1 of [206]). An *APC* absorbs molecules and particles from its vicinity and breaks them down. The emerging fragments, so-called peptides (short sequences of amino acids), serve as antigens. They are bound to so-called MHC molecules (still within the cell), and the resulting complexes, each composed of an MHC molecule and a peptide, are displayed on the surface of the cell (the MHC molecules serve as “carriers” or “anchors” to the cell surface). Since most of the molecules in the vicinity of an APC are “self” molecules, every APC displays a large variety of different types of self antigens and, possibly, one (or a small number of) foreign types. The various antigen types occur in various copy numbers. Each *T-cell* is characterised by a specific type of T-cell receptor (TCR), which is displayed in many *identical* copies on the surface of the particular T-cell. When a T-cell meets an APC, the contact between them is established by a temporary bond between the cells, in which the TCRs and the MHC-peptide complexes interact with each other, which results in stimuli to the T-cell body. If the added stimulation rate is above a given threshold, the T-cell is activated to reproduce, and the resulting clones of T-cells will initiate an immune reaction against the intruder.

of the genes responsible for the expression of this receptor (VDJ recombination) [126]. A very similar mechanism is also involved in T-cell development, therefore we skip any details. With the BCR a B-cell can bind to free (soluble) antigens in their native form. To prevent autoimmune reactions to self antigen, all B-cells have to undergo a selection process during maturation. They meet a huge amount of self antigens in a special environment and only survive if they do not react [2, 144]. On the encounter of its cognate (= perfectly fitting = agonistic) antigen outside the bone marrow, the B-cell is activated and can become either a plasma B-cell, which secretes antibodies (exact copies of its BCR), which bind to every cognate antigen and thereby mark the associated microbe for termination, or it can become a memory B-cell, which lives for a long time and can react much faster to a second encounter with the same antigen (see Figure 2.4). So B-cells fulfill all the requirements mentioned above for the adaptive immune system.

T-cells

T-cells belong to the group of white blood cells (lymphocytes). They develop in the thymus and form an integral part of the adaptive immune system, with plenty of different tasks. Mainly they detect and attack pathogens which bypass the detection via B cells

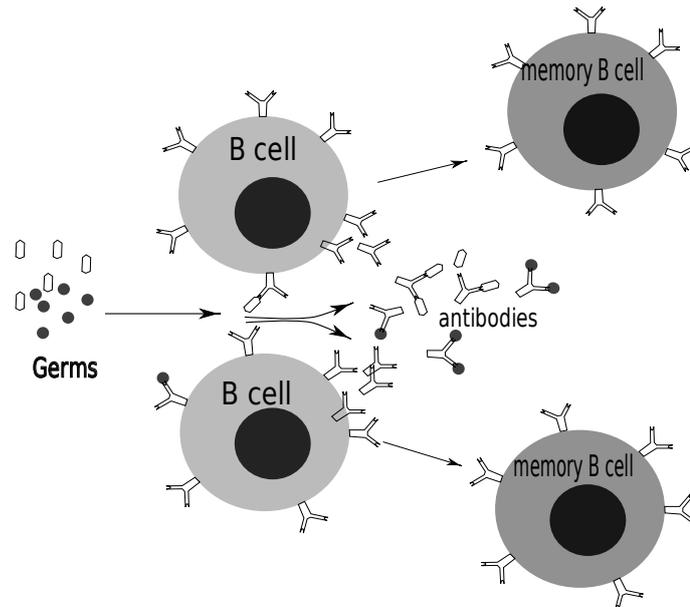


Figure 2.4: B lymphocytes are responsible for the humoral immunity. That is, they detect soluble antigens (small germs or toxins) in the blood via their special B-cell receptor (BCR). This BCR binds to antigens in their native, 3 dimensional shape (in contrast to the T-cell receptor as we will explain later). After a successful detection, the B-cell can release antibodies, which are exact replicas of their BCR. These antibodies bind to all cognate molecules and surface proteins and thereby mark them for other cells like macrophages for termination by phago/endocytosis. Some of the activated B-cells become memory B-cells that survive in a resting-state and can react much faster on a second encounter with the same antigen.

and their antibodies. Moreover, they support B-cells with the effect of new cognate antibody production. One subpopulation of T-cells, called regulatory T-cells (Treg), even acts to stop immune responses in order to prevent autoimmune reactions [142].

A T-cell carries several copies of a specific unique receptor, called T-cell receptor (TCR), on its surface. Upon leaving the thymus, T-cells migrate through blood vessels and especially the lymph nodes. If they encounter an APC, they sort of 'scan' it via their TCR, that is the copies of their TCR bind to the pMHC molecules on the surface of the APC. Every TCR can only bind to one (or a very restricted set) of the several MHC types, which is the first step in a stable binding of the TCR-pMHC complex. The second step is the binding to the presented antigen. If this binding is stable enough, that is the binding duration of such a complex exceeds a certain threshold, the TCR is triggered, signals this to the T-cell and eventually the T-cell is activated.

Generally, naive T-cells, that is T-cells which have not been activated, belong to either of two main types: $CD8^+ CD4^-$ (cytotoxic) and $CD4^+ CD8^-$ (helper) T-cells (the + in $CD4^+$ just means that this receptor is expressed on the cell surface. A - means the opposite. In the following we omit the + and the whole - term). These two molecules are important co-receptors besides the regular T-cell receptor. CD8 helps to stabilise the binding of its T-cell with a MHC molecule of class I, whereas CD4 helps with the binding to MHC II molecules. Tregs are a special group of CD4 T-cells with additional

CD25 molecules. Less than 10% of the T-cells have neither of the two molecules. These special T-cells do not bind to pMHC complexes, but have different other molecules which enhance the binding to for example glycolipid antigens [37].

T-cell differentiation

Without going into the details of T-cell activation, which forms the integral part of this thesis and will be described in greater detail later on (see Section 2.3), let us assume that a T-cell is activated by a pMHC complex on an APC. Both, activated CD4 and CD8 T-cells start a rapid proliferation process in order to generate many clones. Furthermore, the activation event leads to a differentiation of the T-cell clones into various possible T-cell types, depending on the co-stimulatory signals and cytokines being present during the activation process [141]. Until now, 4 main subpopulations of helper T-cells and 3 main subpopulations of cytotoxic T-cells were identified. In the presence of interleukin 12 (IL-12) naive CD4 T-cells become T_H1 cells, which support cell-mediated immune responses by secreting the second messengers IL-2, $IFN\gamma$ and $TNF-\alpha$ and thereby communicate to other cells. This leads, for example, to an improved killing efficacy of macrophages and an improved proliferation of cytotoxic T-cells as well as production of antibodies. They also support T_H1 differentiation through a positive feedback loop [191, 180, 122].

Activation of naive CD4 T-cells via IL-4 stimulated dendritic cells lead to a differentiation to T_H2 cells [51], which support humoral, antihelminthic (against worms) and allergic immune responses [227, 44]. By the production of IL-4, IL-5, IL-6, IL-10 and IL-13 and GATA-3, they enhance B-cell proliferation and antibody secretion. Additionally, they inhibit T_H1 cell differentiation while simultaneously promoting T_H2 differentiation [223].

A third very recently found subpopulation of helper T-cells are the T_H17 cells, which are developed in the presence of $TGF-\beta$ and IL-6. They owe their name to the cytokine IL-17, which they produce beside IL-17a and IL-22. Due to the broad distribution of IL-17 and IL-22 receptors, they thereby induce a massive tissue reaction. Consequently, they promote tissue inflammation (especially during autoimmune diseases). A second very important task of T_H17 cells is the clearance of extracellular pathogens during infections [188, 109, 147].

A special group of CD4 T-cells are the regulatory T-cells. For these cells two distinct ways of development were identified [38]. A larger part of the Treg cells are CD4 CD25 natural Tregs, which are indispensable for the maintenance of immunological tolerance in the host, that is they prevent autoimmune reactions [171, 170, 172, 149]. These Tregs are thymus-derived, similar to the naive CD4 and CD8 T-cells. In contrast to these, the second group, so-called induced Tregs, are generated during CD4 T-cell differentiation upon activation and the simultaneous presence of IL-2 and TGF-Beta. This is important for achieving oral tolerance to allergens and other antigens. It has been found that differentiation of these Tregs and T_H17 cells is reciprocally regulated [226]. All these descriptions are summarised in Figure 2.5.

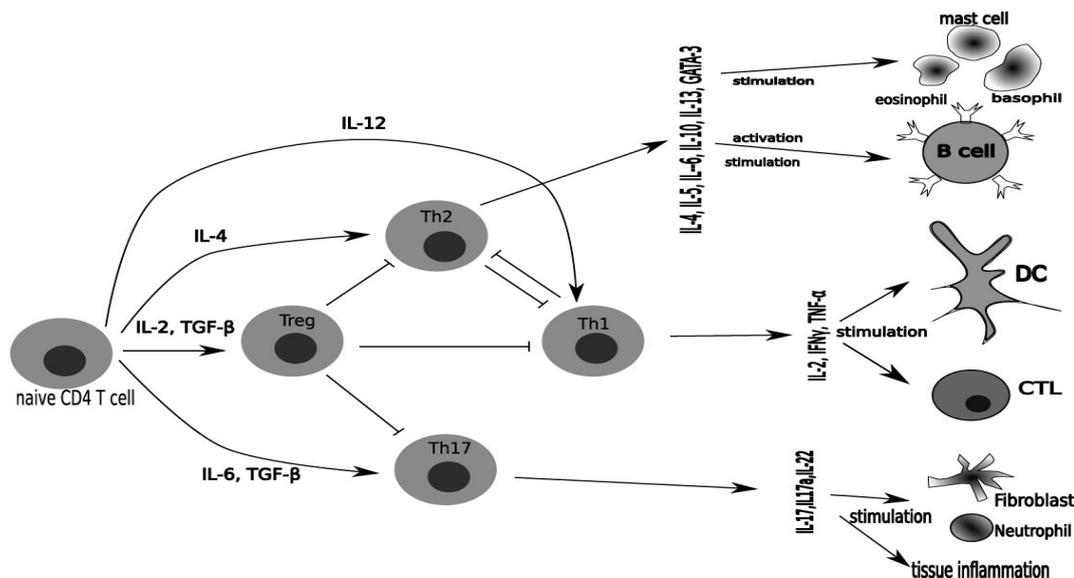


Figure 2.5: CD4 T-cell differentiation; During activation, the T-cell receives different stimulation and inhibition signals in form of different interleukin molecules. These signals influence the differentiation of the T-cell into different types of activated T-cells, which then also fulfill different tasks in the immune reaction. Consequently a naive CD4 T-cell can ultimately influence the immune reaction in all different parts of the immune system, that is humoral, cellular and innate immunity.

In a similar manner CD8 T-cells differentiate into the subpopulations T_C1 , T_C2 , and T_C17 , whereby these play a much more active role in comparison to T helper cells. [223, 120, 137]. After activation they scan all cells in their surrounding and as these cells also present MHC I molecules with antigens (but only in small numbers), they might encounter their activating antigen again if the pathogen is reproducing inside the cell. If so, they attack this cell and thereby prevent this reproduction.

All in all, it is important to note that the whole process of T-cell activation and differentiation is much more than a binary on-off decision upon antigen encounter. On the contrary, the final outcome is very much dependent on a well-orchestrated mixture of different second messengers that bind to other receptors on the T-cell. However, we are only interested in the very first step, namely the activation signal induced by antigens and we assume that this process follows the same rules for all types of T-cells.

2.2 T-LYMPHOPOIESIS

After the short general introduction to the immune system and T-cells, we now focus on all the different aspects regarding T-cells, starting in this section with their developmental process (T-Lymphopoiesis). We will explain why this can already help to understand many aspects of foreign-self discrimination.

T-cells as well as the other lymphocytes (B-cells, natural killer (NK) cells), but also a portion of dendritic cells, develop from the same origin, so called hematopoietic stem cells (HSCs) [19]. These cells are situated in the bone marrow and are the source of all types of blood cells. Through a complicated signalling network, HSCs differentiate

into different progenitor populations, ultimately resulting in all lineages of blood cells. For most of these cell types, the maturation process takes place in the marrow, with one important exception, the T-cells. For their development a whole organ, the thymus, stands by [41].

At the top of the T-cell development hierarchy in the thymus are the so called thymus-resident T-cell progenitors. These cells originate from the bone marrow and are periodically imported into the thymus [66]. The continual settling of new progenitor cells is necessary due to the fact that these cells have only a limited self-renewing capacity. The wave-like behaviour of the progenitor cell influx into the thymus accompanied by a wave of intrathymic DC formation is assumed to optimise T-cell selection, the ultimate stage of T-cell development in the thymus [60]. T-cell selection helps to sort out inactive or autoaggressive T-cells and is described in detail later.

During T-Lymphopoiesis three main stages are identified. They are called double negative (DN), double positive (DP) and single positive (SP), depending on the expression of none of the co-receptors CD4 and CD8, the expression of both of them and ultimately the expression of only one of the two co-receptors [183]. This differentiation leads to two different T-cell subpopulations, whose functions have already been characterised in Section 2.1.2. Each of the three stages occurs in different areas of the thymus consisting of different microenvironments [5]. Actually, it is important to note that the development of thymocytes is crucially dependent on the interaction with thymic epithelial cells (TECs) in the different microenvironments, but on the other hand TECs need the interaction with thymocytes to develop the appropriate microenvironments [210, 209, 21]. Thymocyte development and migration through the thymus seems to be governed by the sequential expression of different chemokine receptors and the release of chemokines via TECs in individual microenvironments [192].

The first two stages occur in the cortex of the thymus. During the DN stage the lymphocytes migrate to the outer cortex called subcapsular zone. Meanwhile they promote the development of cortical TECs (cTECs) from TEC progenitors and start the development of the already mentioned T-cell receptor (TCR). This receptor consists of two different protein chains. The successful assembly of the first protein chain and the formation of a pre-TCR complex on the cell surface marks the transfer of a T-cell from the DN to the DP stage. In this second stage T-cells fully develop the TCR and express it in low levels on its surface besides the two co-receptors CD4 and CD8. They now undergo the process of positive selection, during which most of the T-cells die [63, 73]. After surviving positive selection, T-cells switch to the SP stage and migrate to the medulla. Here they encounter a negative selection mechanism, which can beside killing a T-cell, also turn it into a Treg cell. The few surviving T-cells finally enter the periphery in order to defend the host. In the following we highlight some of the important aspects of T-cell development, including these selection mechanisms, to clarify this on the first sight rather complicated developmental process.

As mentioned briefly in the introduction, T-cell activation is dependent on the TCR, which will be highlighted in Section 2.3. It is therefore necessary to have a closer look at this T-cell receptor. Indeed, the TCR of a T-cell will be its defining element throughout this thesis.

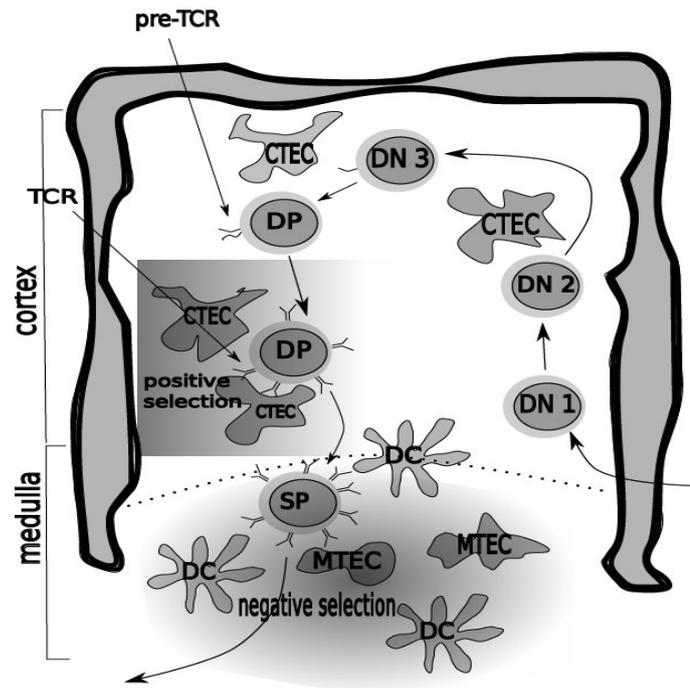


Figure 2.6: T-lymphopoiesis. T-cell precursors enter the thymus and start to interact with the thymic microenvironment. In the 3 double negative stages (DN1, DN2, DN3) they express no co-receptor and start to develop the β -chain of the TCR. The successful assembly of a β and pre- α -chain marks the transfer to the double positive (DP) stage where both co-receptors, CD4 and CD8, are expressed. The T-cells undergo positive selection during which different α -chains of the TCR are generated until the T-cell receives signals from the cTECs or dies by neglect. Ultimately, the T-cells move from the thymic cortex to the thymic medulla and by downregulation of one of the co-receptors enter the single positive stage (SP). Then the T-cells meet with mTECs and DCs which can induce death by apoptosis. Finally, after circulating around for some days in the medulla, the surviving T-cells leave the thymus.

During the DN stage lymphocytes start to generate the so called T-cell receptor (TCR). Generally, two different classes of TCRs are identified, the $\alpha\beta$ and the $\gamma\delta$ receptor. T-cell progenitors become committed to one of these two different T-cell lineages [20].

2.2.1 The T-cell receptor

A T-cell receptor is a heterodimer that consists mainly of two transmembrane glycoprotein chains, α and β or γ and δ . Each of the protein chains is anchored through the cell membrane and consists of two extracellular domains, a constant one and a variable one. It is this variable domain that makes the TCR so special and allows for the detection of many different antigens. The domain is called variable, because in the process of its generation the expression of the associated genes undergoes a procedure called V(D)J recombination [4]. This procedure rearranges variable (V), diversity (D) and joining (J)

gene segments, such that every variable domain and thereby every TCR becomes almost unique. This is true not only for every single individual, but also, to a certain degree, in between different individuals. Recently it has been shown that there are certain TCRs which are shared by groups of individuals. This is due to “convergent recombination”, that is their probability to be expressed is higher because they can be built up from several genetic combinations (degenerated genetic code, alternative splicing possibilities), rather than because of a bias in recombination [213]. Furthermore, the full TCR complex consists of several other protein chains, serving as co-receptors and having a direct influence on the downstream signalling events inside the T-cell. Figure 2.7 shows a cartoon version of the full TCR complex.

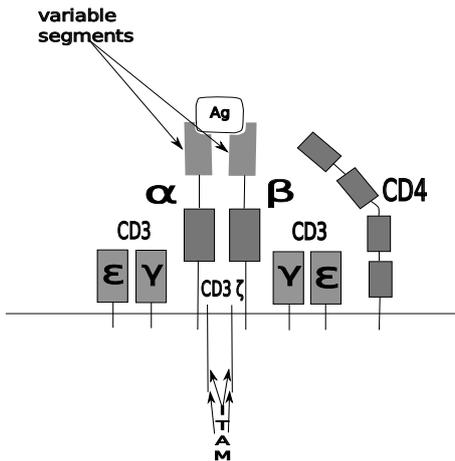


Figure 2.7: TCR complex; the α and β chain are the actual TCR, whereas the other protein chains play the role of co-receptors. ITAMs (immunoreceptor tyrosine-based activation motif) are segments of the CD3 ζ chains that are phosphorylated during TCR-antigen binding by receptor-associated kinases like LCK. This phosphorylation is the starting point of T-cell signalling and is also influenced by the co-receptors.

antigens which are expressed by stressed or (near)-apoptotic cells, due to for example tumor development or infection by viruses [195], just in contrast to $\alpha\beta$ T-cells which are strongly dependent on foreign antigens and have all the different functions already described previously in Section 2.1.2. Therefore, $\alpha\beta$ and $\gamma\delta$ T-cells are two very distinct T-cell lineages which have not very much in common and need to be examined separately. In the following we omit the $\gamma\delta$ T-cells and concentrate on $\alpha\beta$ T-cells, the integral part of the adaptive immune system.

As mentioned earlier, the checkpoint for a T-cell to turn from DN to DP stage is the development of a pre-TCR complex. This checkpoint is called β -selection, as this pre-TCR complex contains the successfully assembled β chain, together with a pre- α chain. The latter chain is then also assembled and the T-cell has a fully functional TCR. From this point on, the survival of a T-cell is crucially dependent on its receptor. A T cell

The generation of the TCR follows a strict order in which the genes for the different protein chains are recombined sequentially ($\delta > \gamma > \beta > \alpha$) [59]. Interestingly, $\gamma\delta$ T-cells can generate more unique receptors than $\alpha\beta$ T-cells and B-cells combined, but their repertoire is quite restricted only to a specific subset of these receptors [27]. The role of $\gamma\delta$ T-cells in the immune system is not well understood until now. They play a role in innate as well as adaptive immune responses and have different functions depending on the tissue they reside in. Indeed, there seems even to be a particular association between certain $\gamma\delta$ TCR repertoires and certain tissues [84]. Receptor binding and activation is not MHC restricted and they are assumed to bind to so-called phospho-antigens, that is non-peptidic, phosphorylated compounds. It is assumed that they, for example, influence DC functions and thereby their immune responses to infectious pathogens and/or are a complement and regulator of $\alpha\beta$ T-cells [29]. $\gamma\delta$ T-cell activation seems to rely mainly on self

regularly needs certain signals in order to survive. These signals can only be induced by the binding of its TCR to peptide-MHC complexes on the cell surface of antigen presenting cells. This is true for the thymus, where the TECs and DCs play the role of APCs, as well as for the periphery where there are further types of APCs. As a result of this dependence, T-cells start to scan cTECs and DCs in the cortex for pMHCs and if their receptors are not able to bind to any of these pMHC complexes, the T-cells die (death by neglect). This process is called positive selection.

2.2.2 Positive selection

T-cell receptors as such are just macromolecules with the capability to bind to cell surface proteins. In this respect they are not different from for example antibodies. This also implies that initially there should be no bias for a T-cell to bind to pMHC molecules on specialised APCs, but given the broad spectrum of possible surface proteins in an individual, it is very probable that every freshly produced TCR could bind to several of these surface proteins [69]. On the other hand T-cells need to specifically bind with their TCRs to pMHC complexes on APCs in order to receive surviving signals and ultimately detect foreign antigens. Consequently, not every T-cell produced in the thymus should be able to fulfill these task, but every T-cell circulating in the periphery has the ability to bind to one type of pMHC complex and thereby to receive surviving signals. Furthermore, it is obvious that there cannot be a common binding site for all types of MHC molecules that can be recognised by every TCR, because then every T-cell in the periphery would recognise all MHC types. The specificity of the TCR suggests that the binding site of the TCR to the MHC is built either completely randomly during V(D)J recombination or the genetic combinations are restricted such that a TCR can bind to only one MHC type. In either case, many T-cells are produced which cannot bind to one of the very few MHC types present in an individual. It is obvious that these T-cells are superfluous and further maturation of them would be a waste of energy for the thymus. Therefore it is necessary to sort out these useless cells, which is fortunately quite easy. The thymus itself has to take no active role in it, but acts as a passive bystander. It is this passivity, which eventually leads to the death of anergic T-cells. Up to this point developing T-cells have constantly received signals from TECs to ensure their survival and development. With a fully functional TCR developed, T-cells rely on a new surviving signal, the activation signal of the TCR, which is induced by a sufficiently long-lasting binding of the TCR to a pMHC on an APC. T-cells which do not receive such a signal die (death by neglect). This process is termed positive selection [184]. During the whole phase between successful β -selection and either successful positive selection or cell death, the T-cell is able to change the TCR α -chain via rearrangements several times [153] and thereby enhance its chance to get a positive selection signal. T-cells in the DP stage survive about 3-4 days, which limits the number of possible rearrangements. It is assumed that this restricted time window is necessary to regulate the TCR α repertoire [79].

A question yet to be answered is the influence of specific antigens in positive selection.

Antigens stabilise the binding of the TCR with a pMHC complex. Therefore it is obvious that, even if the threshold (in terms of TCR-pMHC half-life) for TCR triggering (and thereby positive selection) is low, antigens play a role in it. Several experiments have shown that, at least under certain conditions, there is no need for a diverse self antigen repertoire in the cortex in order to induce positive selection [95, 96]. In fact, only a single antigen is needed to ensure positive selection of quite a diverse T-cell repertoire, which hints to a very low threshold for TCR triggering. Recently it could be shown that cTECs have a unique complex for proteolysis (the degradation of proteins), termed the thymoproteasome [139]. It belongs to the same family of proteasomes that are responsible for the generation of antigenic peptides presented by MHC I molecules in the periphery. The thymoproteasome fulfills the very same task in the cTECs. In contrast to the other proteasomes, it has a weak chymotryptic activity (chymotrypsin is the enzyme responsible for the proteolysis), such that only a unique restricted repertoire of MHC I associated antigens can be generated [193]. It is further hypothesized that for MHC II associated antigens a similar protein cleavage process occurs, involving cathepsin L or S, which is highly expressed in cTECs but not in mTECs [12, 90]. Consequently, positive selection seems to depend on the recognition of MHC in conjunction with a special restricted set of antigens, which are generated by the so called 'modest' cTEC protein degradation [193, 140].

For a long time, positive selection has been seen to be the process that shapes the T-cell repertoire such that every T-cell binds uniquely to a specific MHC. But research into the reasons for organ rejection after transplantation has already shown that T-cells of the host were able to bind to foreign pMHC molecules of the organ, which implies a certain crossreactivity of a TCR to several MHC molecules. In several experiments the T-cell repertoire before positive selection was tested for its crossreactivity and it could be shown that many of them are crossreactive even to the few MHC molecules of the host itself [230, 69]. The same is true for the repertoire after positive selection. In the beginning there is actually even no restriction on either of the two MHC classes [65]. These observations are clarified by the investigation of TCR-pMHC crystal structures. Although the part of the receptor which binds to MHC molecules is generated randomly, there seem to be restrictions such that this randomness is almost limited to a specific pool of sequences. On the other hand an MHC molecule has several surface residues to which a TCR can bind, so called 'codons' [69]. Thus, there is no shared structural binding epitope for all MHC molecules, but a collection. Both, this collection of codons and the pool of TCR binding sequences may have co-evolved, such that many TCRs are in fact quite crossreactive with respect to MHC molecules. All these observations imply that the role of positive selection in rejecting T-cells may be overrated [65].

Its role in the CD4 CD8 lineage decision is still not clear. The strength of the signal induced by pMHC complexes may play a role in it [82, 177]. Here the restricted repertoire of selecting antigens may play a role. There appears to be a negative/positive feedback from TCR-pMHC I/II complexes which influences this decision [65].

Nevertheless it is a fact that T-cells in the periphery do not bind to different MHC molecules, hence, are restricted to one class of MHC. Furthermore, T-cells also rarely bind and react to pMHC complexes where the peptide is a self antigen. Therefore, there

has to be another process which shapes the T-cell repertoire that leaves the thymus. This process is the last checkpoint a T-cell has to pass before entering the periphery and is termed negative selection.

2.2.3 Negative selection

During the DP stage, the TCR is not fully up-regulated. Hence, the TCR sensitivity is biased toward low-affinity interactions. Upon positive selection T-cells reach the SP stage. They are now committed to either the CD4 or CD8 lineage and express a fully functional signalling apparatus, such that in principle they could be released into the periphery in order to meet with APCs and look for pMHC molecules with high affinities to their TCRs in order to get activated [45]. Unfortunately, most of these presented peptides are self antigens, that is antigens derived from proteins of the host itself. A reaction to such antigens would lead to an autoimmune reaction and eventually the death of the host. This is obviously not the usual case in reality. Furthermore, it remains to be clarified why T-cells are restricted to one type of MHC and what the advantage of this restriction is. All of this can be explained by a close look at the last selection step in the thymus, negative selection. In short words, during this process all T-cells which react too strongly to self pMHC are killed and how this is achieved will be the topic of this section.

In the SP stage T-cells migrate to the medulla. During this migration and later in the medulla they constantly meet APCs in form of DCs, macrophages, cTECs and mTECs [192, 129]. These cells present self-derived pMHCs and many of the T-cells encounter their cognate antigen while scanning these cells, which leads to strong signals induced through the triggered TCRs to their T-cells. In the periphery this would lead to a T-cell activation, but in the controlled environment of the thymus the opposite is happening. A T-cell which is triggered too much induces its own apoptosis program. Although, in principle, in the cortex T-cell triggering (for full activation, in contrast to 'just' receiving survival signals) and deletion are also possible, it is far more unlikely [87, 116, 129]. Besides TCR-pMHC binding duration, T-cell activation is also dependent on extrinsic factors such as co-stimulatory molecules. Only in the medulla these are expressed in sufficient amounts (comparable to the periphery) by mTECs and DCs, such that strong TCR triggering is much more likely.

It is this process which also imposes the strong MHC restriction on the peripheral T-cell repertoire [221, 95, 110]. All T-cells that have the ability to bind to several of the few different MHC types present in a host have a huge disadvantage during negative selection. As already mentioned before, different MHC molecules can present different antigens. Hence, if a T-cell can bind to several MHCs the probability to encounter a cognate antigen is elevated significantly and it is very improbable that this T-cell survives.

The need for MHC restriction and diversity

At this point it is appropriate to think about the reasons behind the MHC restriction and the reason why there are different MHC types available. MHC restriction leads to a T-cell reaction in a controlled environment. Without the restriction to one MHC type, the probability of activation by self antigens would be quite high, especially if there would be no restriction to MHC at all, which would be thinkable, as TCRs can bind to several other molecules. On the other side it is advantageous to have more than one MHC molecule type in a host, in order to counter evading strategies of pathogens. By means of mutations and recombinations these change their peptides in the course of evolution, which could lead to antigens that cannot be connected to the given MHC molecule and thereby detection via T-cells would be impossible. Thus, a diverse repertoire of MHC types is necessary. However, as mentioned before, for a single individual this diversity is very restricted, whereas over the whole population it is very high [160]. There is the paradox of high inter-individual diversity and low intra-individual diversity [221]. This can be explained best out of the perspective of evolution and selective pressure. If a single individual has too many different MHC types many more T-cells would be depleted from the repertoire during negative selection, because many more self antigens could be presented. Moreover, the risk of an autoreactive T cell to escape negative selection rises, as it is impossible to present all self antigens to a T-cell. On the other hand, this puts a single individual at risk to be defenceless to a mutated pathogen. However, it is very unlikely that the same pathogen can circumvent the immune system of other individuals of the same populations, because they have many distinct MHC molecules. Establishing a diverse repertoire of MHCs in an individual is therefore restricted by the risk of autoimmunity, while it is enhanced over the total population to ensure the survival of the other individuals [221].

How to obtain self antigens

Up to this point we have been generally concerned with self antigens that are presented via APCs in the thymus, but we ignored a crucial question. Where do these self antigens actually come from? In the periphery, APCs constantly collect cells, cell fragments, proteins etc., digest them and thereby produce antigens. It is thinkable that this is also true for the thymus. However, we have to think about the thymus as a very special organ, with very special tasks, that are different from all other tissues in the body. Therefore, the proteins and cells involved might also be different from those in the rest of the body and thus also the resulting antigens. This could be true for regular antigens, resulting from molecules involved in the normal cell cycle of every cell, and is definitely true for tissue restricted antigens (TRAs), degradation products of molecules which are only present in a certain tissue. Furthermore the body changes while developing and new kinds of cells and proteins occur. If the T-cell repertoire is not prepared for these, autoimmune reactions are provoked.

The thymus is often described as an autarkic regime, which releases T-cells, NK cells and DCs, but lets nothing inside from the periphery. This is not true at all. In fact

dendritic cells constantly migrate from the periphery to the thymus, loaded with pMHCs [94]. Recently it could be shown that this is true for a specific group of DCs, termed circulatory DCs, whereas there are also groups of DCs that are thymus residents [155]. Until now it is not clear how both types of DCs are involved in mediating central tolerance. For thymic DCs it could be shown *in vitro* that they mediate Treg development [217]. Circulatory DCs seem to play a role in both negative selection and Treg induction [155]. It is clear that this constant influx of pMHC can mediate autoaggressive T-cell deletion at least for the regular antigens, although it is not clear how the thymus manages to prevent DCs with foreign pMHC to enter the thymus, which would be devastating for immune protection. Very recently it could also be shown that there are also certain T-cells that migrate back to the thymus [80, 23]. Speculations on their role in the thymus are even more diverse, from the maintenance of certain thymic microenvironments, over import of self pMHCs up to deletion of autoreactive T-cells or the conversion to Tregs.

While circulatory DCs might carry enough pMHCs to mediate tolerance to regular antigens, it is improbable that they present enough tissue restricted antigens, let alone antigens from proteins involved in later stages of host development. Therefore, there is the need for other tolerance mechanisms. Until recently it was proposed that peripheral tolerance (that is tolerance mechanisms outside the thymus) is required to keep T-cells from reacting to these TRAs, although it was hard to explain how this is established. This is not necessary anymore with the recent discovery of a mechanism termed promiscuous gene expression (pGE) [107, 108]. This mechanism allows mTECs to express antigens from all organs and tissues and even from developmentally and temporally regulated genes [54, 75, 181]. With this discovery several questions arose.

There were two competing hypothesis in which way these TRAs are expressed. Either randomly, that is an mTEC expresses antigens from different tissues at the same time, or in a tissue emulating pattern, that is an mTEC plays the role of a cell from a specific tissue and only expresses antigens from this tissue. The latter hypothesis suggests that in different compartments of the medulla different tissues are emulated and T-cells learn there to be tolerant to the specific antigens. This hypothesis is rather intuitive as this should eliminate T-cells quite efficiently and the mechanisms for gene expression are copies of the actual mechanisms in the tissue cells. However, the reality looks different. TRAs are expressed randomly. There is no indication that antigens from specific tissues are expressed more often together [116]. Instead, mTECs express random TRAs colocalised in chromosomal clusters [75, 102].

A second question deals with the regulation of pGE. The central and until now the only known molecule in this context is the autoimmune regulator (AIRE) [6], although some co-regulators like the interferon pathway begin to appear [72]. The discovery of TRAs which can be expressed independently of AIRE implies $LT\beta R$ to regulate their expression [174, 54]. However, the outstanding role of AIRE in pGE regulation could be shown in several experiments with AIRE deficient or altered mice, that lead to different autoimmune diseases [6, 83, 159]. Interestingly, pGE is highly conserved between mouse and human [75], which underlines its importance in the course of evolution and is an explanation for many similar autoimmune diseases in both vertebrates. It is evident that AIRE regulates TRA expression directly and via epigenetic mechanisms, but its exact

function(s) are not identified and several models exist [199].

To understand pGE it is also of interest to have a look at the mTECs themselves. They share the same progenitor cells as the cTECs, but little is known about the first lineage decisions to becoming mTECs [18]. Medullary thymic epithelial cells can be divided into different subsets depending on the expression of the surface markers CD80, CD40, MHC class II and AIRE. It has recently been shown that mTEC differentiation follows the so called 'terminal differentiation model', that is mTECs develop from CD80^{low} AIRE⁻ (few CD80 molecules, no AIRE) to CD80^{high} AIRE⁺ (many CD80 molecules, AIRE present) and meanwhile expand the repertoire of genes they express [116]. It is suggested that there exists a unique mTEC lineage up to the mature mTECs [81]. The fully mature mTECs with the widest ranges of gene expression only survive about 2 weeks, which could be due to an overload of the gene expression machinery [114]. The mTECs in the medulla form separated areas called microdomains, which consist of one to three clonal islets [199, 18]. These islets again contain varying numbers of mTECs of all developmental stages.

One problem with mTECs is their limited capability of inducing T-cell apoptosis and their short lifetime in which they can present TRAs [68]. DCs on the other hand are highly capable of T-cell apoptosis induction and many of them are present in the medulla. It was shown that some of these DCs derive TRAs via so-called cross-presentation from mTECs. How this acquisition is achieved is under discussion, one mechanism might just be the collection of cell material from mTECs that died by apoptosis [76, 77]. This could help to magnify the effectivity of negative selection to TRAs. Furthermore, a special set of DCs together with a special set of mTECs is implied to be the mediator of Treg development in the medulla [217].

In a nutshell, pGE provides for an expression of tissue restricted antigens, whereby every antigen is expressed by only some of the mTECs, randomly in spatially distinct regions of the medulla. Only a certain number of TRAs is expressed simultaneously by neighboring mTECs [199]. Somehow there has to be a mechanism to ensure the effectiveness of pGE in order to prevent autoimmune diseases, especially in light of the fact that in some situations the loss of expression of only one particular TRA can have devastating effects [57]. This mechanism is implied to be stochastic and has to involve epigenetic regulators that are responsible for chromosomal remodelling [55, 215]. It remains to be shown how it really works.

Mature T-cells which survive negative selection eventually migrate to the periphery. The exact mechanism behind this is not identified, but it appears to happen in an ordered fashion, that is the oldest T-cells leave first [129].

2.2.4 *T-Lymphopoiesis in numbers*

After the qualitative description of T-cell development, we here describe it briefly quantitatively.

On average about 10 – 100 hematopoietic precursors enter the thymus per day [116]. But, as already mentioned, this happens in a cyclic manner, such that one week after

leaving the bone marrow these cells enter the thymus and start seeding [60]. It then follows a stage of proliferative expansion and further differentiation over 10 – 14 days [154]. In total, T-cell production has a periodicity of about 3 – 5 weeks [66]. Upon T-cell lineage commitment a T-cell divides about 20 times, mostly in the 2 week long double negative stage [116, 16]. All in all the thymus produces about $5 \cdot 10^7$ T-cells daily, but only $1 - 2 \cdot 10^6$ mature T-cells are released.

Upon β selection a T-cell has to be positively selected. The T-cell constantly tries to bind to pMHC molecules on DCs and cTECs, meanwhile editing the α chain of its receptor via sequential recombination rounds. This process lasts about 3 – 5 days, resulting in the death of the T-cell, if it is not positively selected. The theoretical repertoire size of unique TCRs which can thereby be created is $> 10^{15}$ for mice and $> 10^{18}$ for humans [46, 213]. About 3% of all T-cells survive thymic selection (10% survive positive selection from which again only 35% survive negative selection, other estimates are up to 60% survival rate during negative selection), which reduces the number of possible unique TCRs in the periphery to about 10^{13} for mice and 10^{16} for humans [175, 128, 56]. However, the estimated number of TCRs in the periphery is much lower (10^8 mice, 10^{12} humans) and the number of unique TCRs is only 10^6 for mice and 10^7 for humans [8, 30, 145]. This observation and the fact that there is a pool of TCRs, called public TCRs, which are present in many humans, lead to the theory of convergent recombination. It could be shown that certain TCRs can in principle be generated with a much higher probability [214, 213]. The overall effect of convergent recombination is until now not quantified and therefore it is not clear if it suffices to explain the quite small size of the unique TCR repertoire.

On the other hand we have the vast amount of possible antigens. This number is theoretically estimated to be in the same range as the original number for the TCRs, that is 10^{17} or higher [127]. But in every individual this number is much lower. Here, the restriction is imposed by the MHC molecules, as they cannot bind to all antigens. On the contrary, it is estimated theoretically and experimentally that every type of MHC can 'only' bind to about 10^9 different antigens [67]. This also underlines the need for several different MHC molecules in order to expand the space of presentable antigens. However, this number is still higher than the number of unique TCRs in an individual. Consequentially, a T-cell has to be cross-reactive (also termed poly-specific). Experimental measurements imply that a TCR can bind to about 10^6 antigen types [67]. Generally, these types share similarities on the molecular level, that is a T-cell cannot bind to different totally distinct antigens. This suggests that negative selection creates 'holes' in the space of detectable antigens, which could be shown by an experiment where antigens resulting from HIV share similarities with self antigens and thereby evade detection [67].

It can be shown that prior to entering the selecting stromal environment T-cells move with an average speed of $3 - 8 \mu\text{m}/\text{min}$ and top speed of $30 \mu\text{m}/\text{min}$, following random trajectories [16]. This changes during (positive) selection. The T-cells scan APCs and bind to them. Two different contact types with cTECs were identified, a short one lasting 13 – 36 minutes and a long one lasting over 6 hours. Upon positive selection T-cells move more rapidly and in a more directed fashion in direction of the medulla. In the medulla

they constantly scan mTECs and DCs. New findings indicate that this stage takes only 4 – 5 days, instead of the 12 – 16 days previously assumed [130, 220]. We can further assume a mean scanning time of 3min, which is the mean binding time between a T-cell and a DC during in a lymph node [85, 231]. If we put these numbers together, we can assume that a given T-cell has about 2400 APC meetings in the medulla. Afterwards T-cells migrate to the periphery following the rule that the oldest mature T cells leave first [130, 129].

The important site for negative selection is the medulla and it is important to have a look at its structure. There are several microdomains with 1 to 3 so called clonal islets which merge to large domains [199]. All in all one can identify about 300 areas adding up to 900 different clones [18]. In the early developmental stage each clone consists of 5 – 110 cells with an average of 40, with mTECs in different developmental stages. However, the number of mTECs is estimated to be about 300000, which leads to an average of 166 mTECs per clonal islet [114]. 70% of all mature mTECs are AIRE⁺, 30% AIRE⁻ [55]. The average turnover time of a mature mTEC is 2 weeks and the total cell cycle time between 12 to 30 hours [18]. As already mentioned mTECs present an array of TRAs, which increases with their maturation. These epithelial cells express about $1 - 5 \cdot 10^4$ MHC molecules and many of them are also equipped with regular antigens [138, 52]. It is estimated that at least a total of 1200 – 3000 TRA genes are overexpressed in mTECs or 5 – 10% of all human (mouse) transcripts [75]. The half-life of a pMHC on a mTEC is about 20h and one type of TRA is presented by 1 – 3% of all mTEC [55]. The comparison of mRNA levels of some antigens in mTECs and mammary epithelial cells showed much lower levels in mTECs, 167 (53) fold for Csna (Csnb), which will probably be true for all types of TRAs [55]. This indicates of course that the number of copies of a TRA presented by an mTEC is quite low, compared to the periphery.

2.3 T-CELL ACTIVATION

Adaptive immune responses are crucially dependent on one central event, the activation of a T-cell upon an encounter with an antigen bearing cell and the subsequent recognition of an antigen in form of a pMHC complex. This recognition event has to be both, highly specific and highly sensitive.

However, despite decades of research, a growing body of observations and different proposed activation models, the overall explanation for the whole mechanism is still lacking, or as stated recently: 'In the back of our minds, we all know how antigen receptors trigger the activation of T- and B-lymphocytes. It seems relatively well worked out. The binding of an antigen to the T-cell receptor results in the phosphorylation of the cytoplasmic domains of these receptors, and off the cells go. Yet how exactly does that happen?'[197]

The problem here is the word 'exactly'. T-cell activation can be viewed on different levels of abstraction. One could build a model incorporating all the different molecules and their interaction on the T-cell and APC surfaces, incorporate signalling cascades and so on. This would lead to a model which really tries to explain everything exactly.

But, at the moment this is hardly possible as the significance and role of many of the molecules and dynamics are quite speculative. Moreover, it is quite likely that in such a big model much useless information would be included, which could even lead to wrong results. For example, T-cell activation in general is often thought to require two simultaneous signals [1, 146]. One coming from the various co-stimulatory molecules like CD80, CD86 and CD40 on the APC and their interaction with CD28 and CD40L on the T-cell [39, 104] and the other one coming from the interaction of pMHC and TCR. However, only a modulation of the latter effects the T-cell response, whereas the former shows only minor variability between different T-cells and APCs [103, 1, 157, 146]. It follows that for a model that describes foreign-self discrimination the co-stimulatory interactions are of minor importance and can be omitted.

Generally it should be possible to explain important facets of T-cell activation without going into the deepest molecular details but rely on some important, yet perhaps unexplained observations. Here, we try to illuminate the foreign-self discrimination capability of the peripheral T-cell repertoire. Next we therefore introduce T-cell activation briefly, highlight the important processes involved, before we present some possible models. This will prepare the ground for the formulation of the T-cell activation model that is the basis of this thesis and is introduced in the chapter thereafter.

2.3.1 Introduction

Models that describe the discrimination of foreign antigens from self antigens by T-cells have ultimately to fulfill the requirement of high specificity and high sensitivity. Normally, there is a trade-off between both of them, which makes T-cell activations very unique [70].

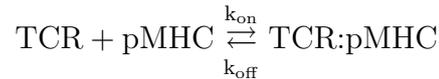
Imagine a TCR and its agonistic antigen, the specificity of a T-cell is then regularly described in terms of the number of different antigens detected by the TCR, where these different antigens result from amino acid substitutions of the original, agonistic antigen. Already one amino acid substitution can prevent T cell activation [179], although usually a TCR is at least a little bit unspecific (also termed crossreactive) [127, 67]. Sensitivity of a TCR on the other hand is measured as the number of cognate antigens necessary to invoke T-cell activation. One early marker of T-cell activation is the flux of Ca^{2+} . It was measured that already one cognate antigen induces a Ca^{2+} flux. Its concentration increases with the number of engaged cognate antigens up to 10 cognate antigens, which seems sufficient to form the immunological synapse [97, 156]. It was also observed that only 3 cognate pMHCs were sufficient for cytotoxic T-cells to switch into killing model. There are about $10^5 - 10^6$ pMHCs on the APC surface, which are about 3000 different antigen types in various copy numbers from the previously mentioned vast space of possible antigens [93, 127, 187, 70]. With an average DC surface of $500\mu\text{m}^2$ and a DC-T-cell area of $50\mu\text{m}^2$ it follows that a T-cell seems to discriminate only a few cognate antigens from a background of at least 10^4 irrelevant antigens [176, 85].

Next we present the relevant biological background together with different models that try to explain this extraordinary discrimination capability. As we will see, many of

these models inherently assume that one TCR-pMHC interaction is sufficient to invoke T-cell activation. We explain the difficulties with this assumption and infer that there have to be additional mechanisms in place.

2.3.2 TCR binding

The central element in T-cell activation is the interaction of the TCRs with the pMHC molecules on the surface of an APC. In contrast to affinity-matured antibodies, the TCR-pMHC interaction has a low affinity in the range of $1 - 100\mu\text{M}$ [47, 208]. In its simplest form it is described by the reaction:



where k_{on} is the association rate which is often quite slow. In fact it is slower than a binding which is dependent on random collisions of TCR and pMHC molecules [47, 208]. The dissociation rate k_{off} on the other hand is high. This rate is often used to calculate the half-life of the interaction $t_{1/2} = \ln 2/k_{\text{off}}$ which is easier to measure and to deal with. There is a good correlation between $t_{1/2}$ and T-cell activity, which makes it an important parameter for model building [47, 208]. But there are exceptions that also point to some influence of the association rate, which would make the binding constant $K_D = k_{\text{off}}/k_{\text{on}}$ the important model parameter [164, 64, 196, 189, 62]. Another study showed that a combination of heat capacity changes ΔC_p and $t_{1/2}$ can overcome interactions that do not follow the $t_{1/2}$ rule [112]. The problem with many of these studies is the use of different peptides as antigens, such that other factors like the influence of MHC binding or measurement inaccuracies with very fast dissociation rates come into play [189]. A strong argument for the $t_{1/2}$ rule is supplied by Cole et al [42]. They show that the TCR binding affinity is governed by the MHC class, that is there are different on-rates of the two different MHC classes. The off-rates for the TCRs on the other hand fall into the same narrow regime regardless of the MHC class. For a list of different TCRs and pMHCs with their reaction parameters and their activation capability see [189]. Unfortunately all this measurements are made at 25°C instead of 37°C and are therefore not directly applicable in models.

For this thesis, the three-dimensional conformations of a TCR and a pMHC and their actual binding behaviour are not relevant. For a recent review of the topic see [168]. Here, it suffices to assume that the association of TCR and pMHC follows a 2-step binding process as for example proposed in [225]. In a first step the TCR binds to the MHC and in a second step to the antigen. Whereas the first step leads to a very transient TCR-MHC complex, the second binding step seemingly suffices to stabilise the binding of the TCR-pMHC complex long enough to invoke some sort of signal which could lead to T-cell activation. More sophisticated models of TCR-pMHC binding would also have to incorporate the influence of co-receptor. For example, Wooldridge et al could show that CD8 stabilises the TCR-pMHC complex which has influence on the activation of cytotoxic T-cells [224]. However, it is generally assumed that T cell activation is similar

for both cytotoxic and helper T-cells, the effect of CD8 might just be to modulate the activation behaviour as proposed and modeled in [206].

2.3.3 TCR triggering

The question is, how does the binding of the TCR to the pMHC invoke a signal to the T-cell. The consensus seems to be that through the binding to the pMHC molecule a cascade of phosphorylation events starts, which ultimately sends an activation signal or as it is often put, the TCR is triggered.

The non-variable part of the TCR complex consists of different signal transduction subunits and these contain so called immunoreceptor tyrosine-based activation motifs (ITAMs). Upon TCR engagement by pMHC, it has been detected that these ITAMs are phosphorylated with the help of Src family tyrosine kinases (SFK) Lck and Fyn [219, 211]. It is still unclear how these early events are started, which is why most of the TCR triggering models try to explain this early activation events, but have quite different approaches.

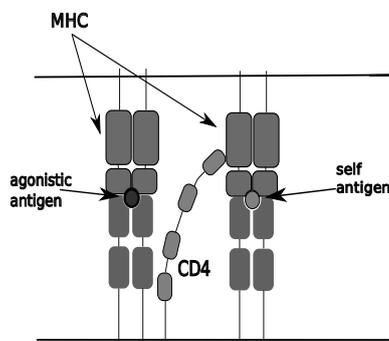


Figure 2.8: Pseudodimer model; One TCR binds to an agonistic antigen and its co-receptor CD4 binds to an MHC molecule of another pMHC-TCR complex where the antigen is self.

There are models which explain TCR triggering through aggregation of TCRs. Models like the co-receptor heterodimerization model or the pseudodimer model explain this aggregation by different mechanisms of dimerization [198, 3, 111]. For example, in the latter model there are two TCRs involved. One binds to an agonistic pMHC, the other one to a self pMHC and both are interconnected via a co-receptor, as can be seen in Figure 2.8. Such a TCR dimer together with both pMHCs would then result in T-cell activation. One major problem with these models is that they do not take into account the ability to react to very low numbers of agonistic pMHC or they instead make use of special endogenous pMHCs, where it is quite improbable that a fitting agonistic and endogenous pMHC are localised next to each other. For a short critical review see [208].

A second type of models uses conformational changes of the TCR induced by the binding of the TCR to the pMHC. These changes could either free certain sites to make them available for phosphorylation or bring together TCR subunits, such that ITAMs are phosphorylated [208]. The advantage of this type of model is that it can explain TCR triggering if only one pMHC is present. Generally, these models have to take into account that the TCR and pMHC structures are quite versatile and flexible, but the triggering mechanism has to work for all of them in the same way [207]. Moreover, it has been shown that there is no conformation change leading to signal transduction within the TCR $\alpha\beta$ heterodimer upon binding [168]. These findings constrain the possibilities how conformational change could work and many older models therefore fail [208]. However, new evidence points to a new mechanism of conformational change [7, 123, 40, 212].

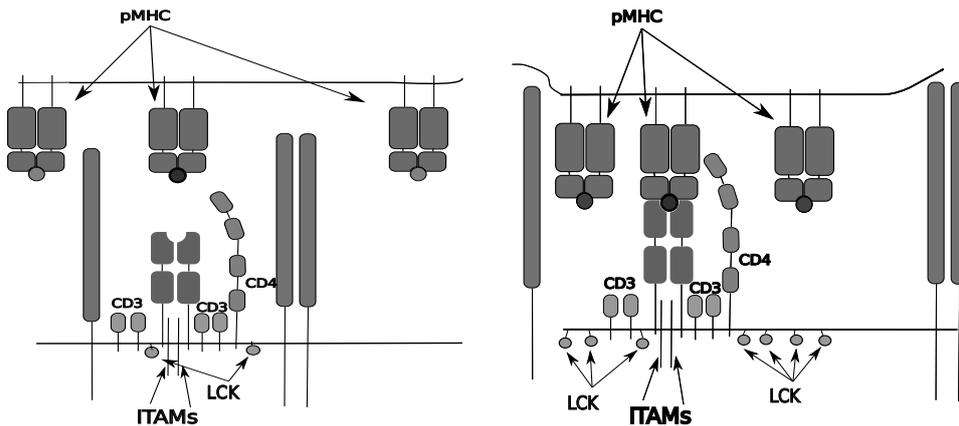


Figure 2.9: Kinetic segregation model: In the state without pMHC binding, phosphorylation by LCK molecules at the TCR complex is in equilibrium with dephosphorylation events (left). By binding of the TCR to a pMHC, large molecules are excluded from the binding domain. More LCK molecules get the possibility for phosphorylation that cannot be compensated by dephosphorylation. If the TCR-pMHC is stable enough the phosphorylation lasts long enough to activate the T-cell.

Here it is assumed that the constant movement of the T-cells and DCs induce a pulling force on the bonded TCR and thereby its conformation is changed.

A third type of model is based on the segregation or redistribution of the TCR complex, either by kinetic segregation [50, 26, 40] or by lipid rafts [91]. In these models special spatial domains are built after TCR-binding which are favourable for ITAM phosphorylation. For an illustration see Figure 2.9.

As the exact mechanism of TCR triggering is still highly controversial it might be more helpful at this point to assume that TCRs are triggered, most probably in dependence on the half-life of the TCR-pMHC complex. Thereby it is possible to make more general propositions about T-cell activation.

2.3.4 Models of T-cell activation

In order to explain the specificity of T-cell activation a kinetic proofreading system was proposed [131, 157]. One interpretation can be seen in Figure 2.10. In this model there is an obligatory chain of signalling intermediates resulting in a final complex. During TCR engagement this chain is run through and if the final complex is reached the TCR is triggered and the T-cell is activated. However, if the TCR disengages before chain completion it is reset to the start complex. It can be shown that this model explains specificity, given that each signalling step occurs at the same rate [34]. Unfortunately this comes with a big loss of sensitivity. Alternatively, this kind of specificity can also be explained by the use of models with different feedback pathways [185, 36].

The problem with all the models mentioned in this and the previous section is that they assume that T-cell activation is more or less the same as TCR triggering. If a TCR is triggered (or only very few) the T-cell is activated. Of course, we mentioned earlier that for a T-cell it is possible to detect one cognate antigen, but this does not

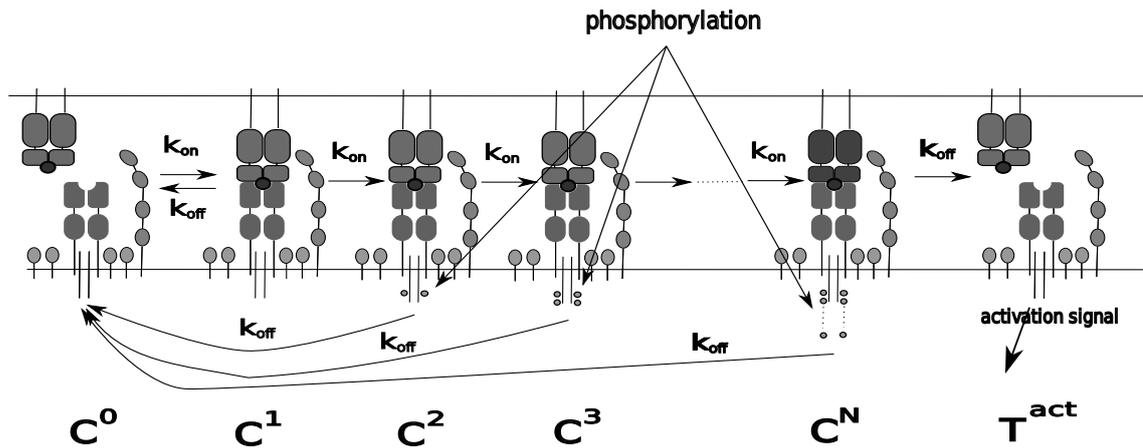


Figure 2.10: Kinetic proofreading: A TCR binds to a pMHC complex. Phosphorylation events start. The longer the binding lasts, the more phosphorylation is achieved. If there is enough phosphorylation and the TCR unbinds, the T-cell gets an activation signal. Otherwise, if there was not enough phosphorylation, the TCR unbinds and is immediately dephosphorylated.

imply that T-cell activation depends on one such engagement alone. On the contrary, there is a big problem with such an assumption. If we solely assume that a T-cell is activated when a certain binding time between a TCR and a pMHC is exceeded then we do not take into account the fact, that the dissociation of a TCR-pMHC complex is stochastic. A half-life of for example 10 seconds means that in 50% of all bindings of such a TCR-pMHC complex, this complex dissociates before 10 seconds (see Figure 2.11). This rather says that, although negative selection might sort out all TCRs with long half-lives with self antigens, this might not affect the activation probabilities very much, because it can be compensated by an increase of engagements between TCRs and self antigens that have lower half-lives. Thereby the overall probability that at least one antigen binds long enough just by chance is elevated. This is obviously no problem given the enormous amount of self antigens on an APC. Therefore, models that explain foreign-self discrimination have also to deal with triggering signals induced by self antigens.

One way to deal with such problems are cooperative models, where TCRs use second-messengers or physical contact to communicate with each other [173, 71, 185]. In these models short binding times desensitise a TCR and TCRs in its neighbourhood, whereas longer binding times work in the opposite way.

It is not hard to imagine that a pMHC triggers a TCR, dissociates from it and afterwards this pMHC can associate to another TCR and perhaps also trigger this one. Valitutti et al could show this serial engagement of many TCRs by few pMHCs and developed the serial-engagement model [203, 202], which then was analysed and modified via experiments and mathematical tools, see for example [10, 100, 182, 222, 201]. Hence, this sequentially binding of multiple TCRs amplifies the signal that one pMHC can induce to the T-cell and only a few agonistic antigens are needed to activate the T-cell. If we additionally assume that a TCR is triggered if a minimum binding time between pMHC and TCR is exceeded, then it follows that there has to be an optimal

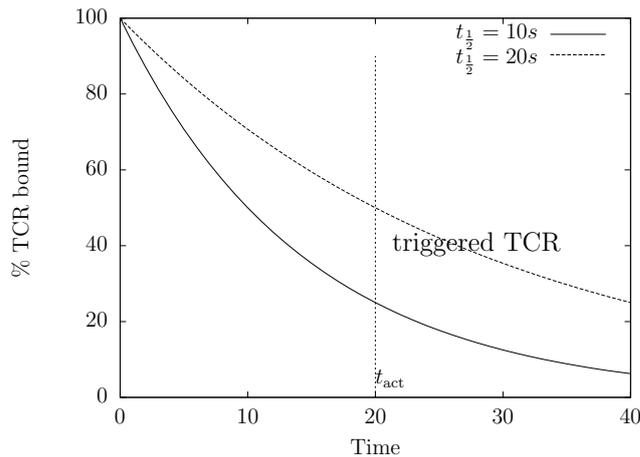


Figure 2.11: Stochastic dissociation: Given a hypothetical triggering threshold of 20s, a TCR is triggered in 50% of all cases if its binding to an antigen has a half-life of 20s. If the half-life is 10s, it is triggered in 25% of all binding events. This just exemplifies that, given enough binding events with one type of antigen, there will always be one event where the binding time exceeds any given triggering threshold. It is therefore insufficient to just use such a triggering threshold for T-cell activation.

half-life which cannot be too short or too high. With a too short half-life a pMHC can engage many TCRs but only trigger them with a very low probability. On the other side, with a too long half-life a pMHC triggers the TCR with high probability but only engages very few TCRs because it stays associated to all of them very long.

These considerations result in the optimal half-life(/dwell-time) hypothesis [103, 74, 28, 146, 43, 92, 106, 190, 167]. Although from a theoretical point of view it is quite compelling, there is plenty experimental evidence for and against this hypothesis. There are several experiments with mutated or newly engineered TCRs that have very slow dissociation rates but lead to T-cell activation, see for example [89, 118, 218, 189]. On the other hand there are examples where long half-lives lead to an impaired TCR-pMHC interaction, see for example [103, 146, 190]. Many of these studies have been performed in vitro and Carreno et al. could show that there are differences between in vitro and in vivo results, where their in vivo results point to the correctness of the hypothesis [146]. Furthermore it is important to note that by increasing the cognate antigen density on an APC long half-life pMHCs can also induce T-cell activation even if the optimal dwell-time hypothesis is correct [74]. One question is therefore if in studies that seem to disprove the hypothesis the antigen density was just too high to see the effect. This could be for example due to an overall higher dosage of antigens or if antigens were used that bind better to MHC molecules, such that they do not dissociate so fast. One further problem is that the half-life is normally measured at 25°C, and there are some indications that there are alterations in the half-life hierarchy in comparison to 37°C measurements [146].

If we consider the optimal dwell-time hypothesis as true, T-cell activation becomes a problem of statistical recognition. The question is, how the signal induced by a set of pMHCs including self and foreign antigens differs from the signal induced by a set of self antigens and how a T-cell can discriminate between these signals. Obviously, the pMHCs

on an APC trigger different amounts of TCRs and there is no difference between these triggered TCRs. Thus, we do not look on a specific signalling pattern where the T-cell can tell how many TCRs were triggered by each antigen. There is only one accumulated signal and the T-cell has to tell if there is a cognate antigen induced signal hidden in this signal, which should lead to T-cell activation. In the following we introduce the T-cell activation model which builds upon these findings and tries to explain the foreign-self discrimination capability of the T-cells.

BRB MODEL OF T-CELL ACTIVATION

In this chapter, we motivate and introduce the model of T-cell recognition as first proposed by BRB in 2001 [205] and further developed by Zint, Baake and den Hollander [232].

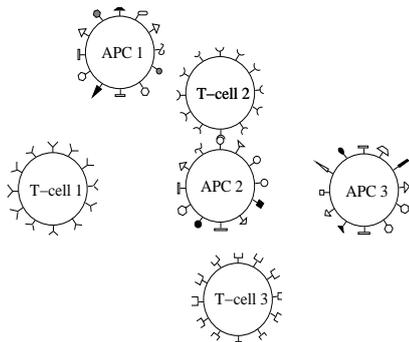


Figure 3.1: Caricature of T-cells and APCs (from [232]). Note that every T-cell has many copies of one particular receptor type, but different T-cells have different receptor types. In contrast, every APC carries a mixture of antigen types, which may appear in various copy numbers.

When T-cells and APCs meet, the T-cell receptors bind to the various antigens presented by the APC [49]. For every single receptor-antigen pair, there is an association-dissociation reaction, the rate constants for which depend on the “match” of the molecular structures of receptor and antigen. Assuming that association is much faster than dissociation and that there is an abundance of receptors (so that the antigens are mostly in the bound state), one can describe the reaction in terms of the dissociation rates only.

Every time a receptor unbinds from an antigen, it sends a signal to the T-cell, provided the association has lasted for at least one time unit (i.e., we rescale time so that the unit of time is this minimal association time required). The duration of a binding of a given receptor-antigen pair follows the $\text{Exp}(1/\tau)$ distribution, i.e. the exponential distribution with mean τ , where τ is the inverse dissociation rate of the pair in question. The rate of stimuli induced by the interaction of our antigen with the receptors in its vicinity is then given by

$$w(\tau) = \frac{1}{\tau} \exp\left(-\frac{1}{\tau}\right) \quad (3.1)$$

(i.e., the dissociation rate times the probability that the association has lasted long enough). As shown in Fig. 3.2, the function w first increases and then decreases with τ with a maximum at $\tau = 1$, which reflects the fact that, for $\tau < 1$, the bindings tend not

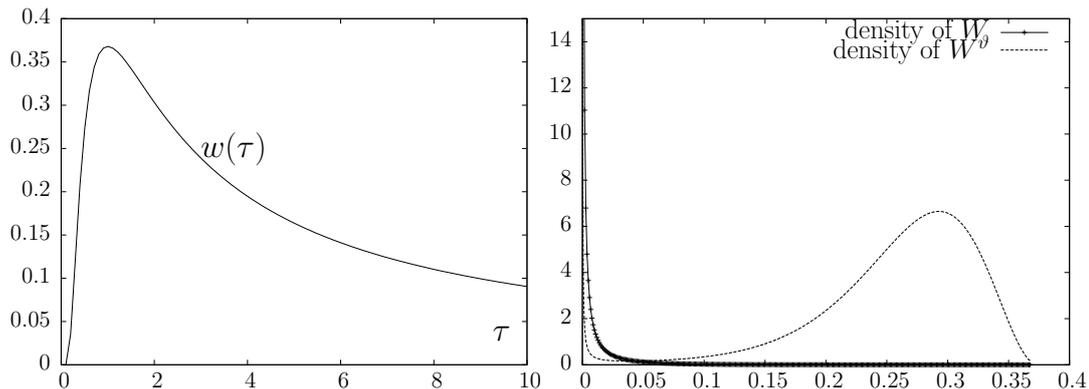


Figure 3.2: Left: The function w . Right: the densities of $W = w(\mathcal{T})$ and W^ϑ with tilting parameter $\vartheta = 46$ (The densities have poles at $w(0) = 0$ and $w(1) = 0.3679$ (due to the vanishing derivative of w at $\tau = 0$ and $\tau = 1$), but the right poles are invisible because they support very little probability mass. In fact, for $\varepsilon = 0.01$, one has $\mathbb{P}(0 \leq W \leq \varepsilon) = 0.98$ and $\mathbb{P}(w(1) - \varepsilon \leq W \leq w(1)) = 2.17 \cdot 10^{-9}$, whereas $\mathbb{P}(0 \leq W^\vartheta \leq \varepsilon) = 0.137138$ and $\mathbb{P}(w(1) - \varepsilon \leq W^\vartheta \leq w(1)) = 0.0050$.

to last long enough, whereas for $\tau > 1$, they tend to last so long that only few stimuli are expected per time unit.

The T-cell sums up the signals induced by the different antigens on the APC, and if the total stimulation rate reaches a certain threshold value, the cell initiates an immune response. This model relies on several previously mentioned hypotheses, namely as kinetic proofreading [131, 158, 121, 86], serial triggering [203, 202, 182, 22, 201, 61], counting of stimulated TCR's [216, 166], and the optimal dwell-time(/half-life) hypothesis.

Due to the huge amount of different receptor and antigen types, it is impossible (and unnecessary) to prescribe the binding durations for all pairs of receptor and antigen types individually. Therefore, BRB chose a probabilistic approach to describe the meeting of APCs and T-cells. A randomly chosen T-cell (that is, a randomly chosen type of receptor) encounters a randomly chosen APC (that is, a random mixture of antigens). The mean binding time that governs the binding of this random receptor to the j th type of antigen is taken to be a random variable denoted by \mathcal{T}_j . The \mathcal{T}_j are independent and identically distributed (i.i.d.) and are assumed to follow the $\text{Exp}(1/\bar{\tau})$ distribution, i.e., the exponential distribution with mean $\bar{\tau}$, where $\bar{\tau}$ is a free parameter. Note that there are two exponential distributions (and two levels of averaging) involved here. First, the duration of an *individual binding* between a type- j antigen and a random receptor is $\text{Exp}(1/\mathcal{T}_j)$ distributed (see the discussion of Eq. (3.1)). Second, \mathcal{T}_j , the *mean duration* of such a binding (where the receptor is chosen once and the times are averaged over repeated bindings with a j antigen) is itself an exponential random variable, with realisation τ_j . Finally, its mean, $\mathbb{E}(\mathcal{T}_j) = \bar{\tau}$, is the mean binding time of a j -antigen (and, due to the i.i.d. assumption, of any antigen) when averaged over all encounters with the various receptor types. The exponential distribution of the individual binding time is an immediate consequence of the (first-order) unbinding kinetics. In contrast, the corresponding assumption for the \mathcal{T}_j is made for simplicity; the concept is compatible with various other distributions as well, see [205] and [232]. The i.i.d. assumption, however, is

crucial, since it implies, in particular, that there is no difference between self and foreign antigens here; i.e., no a priori distinction is built into the model. The total stimulation a T-cell receives is the sum over all stimulus rates $W_j = w(\mathcal{T}_j)$ that emerge from antigens of the j 'th type (we will denote the density of W as g). It is further assumed that there is at most one type of foreign antigen in $z^{(f)}$ copies on an APC, whose signal must be discriminated against the signals of a huge amount of self antigens. The self antigens are here divided into two distinct classes, c and v , that are present in different copy numbers $z^{(c)}$ and $z^{(v)}$. An APC displays $m^{(c)}$ and $m^{(v)}$ different types of class c and v . The indices c and v stand for “constitutive” and for “variable”, respectively; but for the purpose of this article, only the abundances are relevant, in particular, $z^{(c)} > z^{(v)}$ and $m^{(c)} < m^{(v)}$. Over the whole APC the total number of antigens is then $m^{(c)}z^{(c)} + m^{(v)}z^{(v)} =: M$ if no foreign antigen is present. If $z^{(f)}$ foreign molecules are also present, the self molecules are assumed to be proportionally displaced (via the factor $q := (M - z^{(f)})/M$), so that the total number of antigens remains unchanged at

$$z^{(f)} + m^{(c)}qz^{(c)} + m^{(v)}qz^{(v)} = M. \quad (3.2)$$

The total stimulation rate in a random encounter of T-cell and APC can then be described as a function of $z^{(f)}$:

$$G(z^{(f)}) := \left(\sum_{j=1}^{m^{(c)}} qz^{(c)}W_j \right) + \left(\sum_{j=m^{(c)+1}^{m^{(c)}+m^{(v)}} qz^{(v)}W_j \right) + z^{(f)}W_{m^{(c)}+m^{(v)}+1}, \quad (3.3)$$

i.e., a weighted sum of i.i.d. random variables. Alternatively, we consider the extension of the model proposed by Zint et al. [232], which, instead of the deterministic copy numbers $z^{(c)}, z^{(v)}$, uses random variables $Z_j^{(c)}, Z_j^{(v)}$ distributed according to binomial distributions with $\mathbb{E}(Z_j^{(c)}) = z^{(c)}$, $\mathbb{E}(Z_j^{(v)}) = z^{(v)}$ (so the expected number of antigens per APC is still M). The model then reads

$$G(z^{(f)}) := \left(\sum_{j=1}^{m^{(c)}} qZ_j^{(c)}W_j \right) + \left(\sum_{j=m^{(c)+1}^{m^{(c)}+m^{(v)}} qZ_j^{(v)}W_j \right) + z^{(f)}W_{m^{(c)}+m^{(v)}+1}. \quad (3.4)$$

In line with [205, 232], we numerically specify the model parameters as follows: $\bar{\tau} = 0.04$; $m^{(c)} = 50$, $m^{(v)} = 1500$, $z^{(c)} = 500$, $z^{(v)} = 50$ (and hence $M = 10^5$). The distributions in the extended model are the binomials $\text{Bin}(\zeta^{(c)}, p)$ and $\text{Bin}(\zeta^{(v)}, p)$ for $Z_j^{(c)}$ and $Z_j^{(v)}$ respectively, where $\zeta^{(c)} = 1000$, $\zeta^{(v)} = 100$, and $p = 0.5$.

The relevant quantity for us is now the probability

$$\mathbb{P}(G(z^{(f)}) \geq g_{\text{act}}) \quad (3.5)$$

that the stimulation rate reaches or surpasses a threshold g_{act} . To achieve a good foreign-self discrimination, there must be a large difference in probability between the stimulation rate in the case with self antigens only ($z^{(f)} = 0$), and the stimulation rate with the foreign antigen present, i.e.,

$$1 \gg \mathbb{P}(G(z^{(f)}) \geq g_{\text{act}}) \gg \mathbb{P}(G(0) \geq g_{\text{act}}) \geq 0 \quad (3.6)$$

for realistic values of $z^{(f)}$. Note that both events must be rare events – otherwise, the immune system would “fire” all the time. Thus g_{act} must be much larger than $\mathbb{E}(G(z^{(f)}))$ (which, due to (3.2) and the identical distribution of the W_j , is independent of $z^{(f)}$). Evaluating these small probabilities is a challenge. So far, two routes have been used: analytic (asymptotic) theory based on large deviations (LD) and straightforward simulation (so-called simple sampling). Both have their shortcomings: The LD approach is only exact in the limit of infinitely many antigen types (and the available error estimates are usually too crude to be useful); the simulation strategy, on the other hand, is so time-consuming that it becomes simply impossible to obtain sample sizes large enough for a detailed analysis, in particular for large values of g_{act} . Therefore, an importance sampling approach is required.

The BRB model in its just described version is of course still quite abstract, as it for example ignores the fact that even very low numbers of foreign antigens can activate a T-cell and also the selection mechanisms are not included. There already exist extended versions of the model [205, 232], but we will postpone their description to the analysis of the model. There they will emerge as a consequence of the analysis as well as new developed extensions and different models.

3.1 ADDITIONAL REMARKS

In the beginning of our explanations regarding T-cell activation we already made clear that there are still many open questions and as pointed out there are quite a few different possible models under consideration. Before concluding this chapter and thereby the biological introduction we present some recent experimental observations that have not been captured in T-cell activation models, yet.

In the search of a TCR triggering mechanism new results point again to the involvement of receptor deformation [7, 123, 40, 212]. Ma et al. describe and include these results in their receptor deformation model [123]. Although the triggering kinetics as described in the BRB model would be quite different from a T-cell activation model that builds upon the receptor deformation assumptions, there will most probably be commonalities. Especially, there will be some kind of serial triggering. Through the induction of force on the TCR-pMHC complex, they dissociate even faster and consequently a pMHC molecule can bind even more TCRs. On the other side we do not see an analogue to the optimal dwell-time hypothesis as the force induced by the T-cell and APC movements is probably so high that the dissociation times between different TCRs and pMHCs should be quite similar. This would point to a mechanism where we get a saturation in signalling intensity induced by a pMHC with increasing half-life. However, this could be disturbed if we further assume that a very long half-life leads to a too strong force induced to the TCR such that the receptor deformation is too strong to induce TCR triggering. Thereby some kind of optimal dwell-time would be restored and the models should be quite similar.

It is also of note that there are indications that T-cells that have a low affinity to certain self antigens evade negative selection and do not react in the periphery if they

see these antigens in endogenous numbers. However, if the amount of antigen is increased they are activated [228]. It seems even to be the case that T-cells are activated more often than previously thought. Even weak pMHCs can activate a T-cell [229]. The difference between T-cells activated by weak pMHC and by strong pMHC is then in their behaviour afterwards. These observations imply that the activation threshold as set in the BRB model has not to be too high and foreign-self discrimination has not to be perfect. It is important to note that it certainly does not contradict the assumptions of the BRB model. Instead it is thinkable to make use of the fact that we count triggered TCRs and for example introduce additional thresholds.

The last observation we want to mention could as well be seen in this context as it can stand for its own. Henrickson et al. show that T-cell activation in the lymph node is (often) not an instant process but follows different stages [85, 133]. First, T-cells undergo brief serial contacts with DCs for several hours, followed by a phase of stable T-cell-DC interactions. Finally the T-cells return to short T-cell-DC interactions and start to proliferate. Although these results are questioned [33, 32], they are quite interesting as they point to a signal integration mechanism that goes beyond the meeting of a T-cell with only one DC. The real point of T-cell activation would here be the time when a T-cell switches from phase one to phase two. How long this takes can be modulated by the cognate antigen density on the DCs, which could be the reason why Celli et al. do not see a first phase in their experiments [33, 32]. A first computational simulation supports these observations [231]. Again, this result does not contradict the BRB model but opens new possibilities for foreign-self discrimination. There could be an intra-DC activation threshold and an inter-DC activation threshold, that is a T-cell has to be activated several times (perhaps in a time window or in dependence of the amount of DCs met) and is then finally really activated. If we want to combine the last two mentioned observations this could also be interpreted in another way. A T-cell that is activated more often by DCs takes longer to turn to stage two and a T-cell that is only activated very rarely turns fast to stage two. This would be consistent with the observation that weak pMHCs can activate a T-cell but this T-cell does not expand very long but leaves the thymus very fast. It is also possible to interpret this the other way around. High-potency pMHC lead to a faster T-cell migration stop and proliferation whereas there is no stop signal for weak bindings but proliferation is possible for intermediate binding strength [178]

It is evident that with every experiment we get new information on T-cell activation and there is still much unknown. However, these recent results do not contradict the BRB model but leave room for extensions of this model. This feels more than sufficient to revisit the analysis of the BRB model and develop methods that enable a deeper analysis and can be used for extended and follow-up models. This is the motivation for the research that will be described in the main chapters of this thesis.

3.2 T-CELL ACTIVATION IN NUMBERS

To become biologically relevant, models of T-cell activation need to be based on experimentally derived data. A quantification of all different aspects of T-cell activation is

beyond the scope of this thesis. We therefore provide only some data and refer to different relevant publications. We make use of a top down approach, starting with APCs and T-cells in general ending with signalling events at the TCR-pMHC complex.

It is safe to assume that every APC is unique with respect to its antigen profile, that is the collection of pMHC complexes it presents. This is due to the fact that these pMHC complexes are not part of the APC but are a result of the constant internalisation of cell material in the APC's surrounding. For T-cells the situation is slightly different. We already mentioned in 2.2.4 the number of unique T-cells. However, through different events in the periphery a T-cell exists in a low number of copies rather than being really unique (for numbers see Section 2.2.4). On the other side, the number of possible antigens in an individual is estimated to be in the range of 10^{13} .

The next lower level are the single APCs and T-cells. For the former the number of antigens per cell is estimated to be in the range of 300000 with about 3000 unique antigen types in several copies [205, 52, 128]. The number of TCRs per T-cell is in the range of 30000 [93, 127, 187, 70, 43, 189]. The surface area of a DC as the most important APC is estimated to be $500\mu\text{m}^2$ [176, 85] to $1800 - 2400\mu\text{m}^2$, where typically more than two-thirds of this volume is deployed as dendrites [134]. A T-cell on the other hand is much smaller with an estimated surface area between $19 - 40\mu\text{m}^2$ [78] and $150\mu\text{m}^2$ [134]. The size of the contact region between an APC and a T-cell can vary greatly and was experimentally observed to be in the range of 1 to $> 70\mu\text{m}^2$ with a mean of $8\mu\text{m}^2$ [134].

On the level of the TCR-pMHC interaction we have to deal with the association and dissociation rates and of course the crossreactivity of a single TCR. The association and dissociation rates vary greatly and there exists many measurements for different TCRs and antigens together with their capability to induce a T-cell reaction, see for example [189, 146, 28, 48, 103]. The crossreactivity, that is the number of different antigens a TCR can react to, is also under discussion and is estimated to be in the range of 100 to 10^6 , where the lower bound seems to be more reasonable [127, 145, 67].

TCR triggering is dependent on intracellular phosphorylation events, which can be captured by the kinetic proof-reading model. Coombs et al., for example, assume 6 proof-reading steps with a rate of 0.25s^{-1} .

Many of the here mentioned numbers are still under experimental investigation and others are still missing. However, these numbers should be a good point to start the development of T-cell activation models.

MATHEMATICAL BACKGROUND AND COMPUTATIONAL METHODOLOGY

This thesis has two main focuses. The one side is the development of a deeper understanding of the mechanism of T-cell activation and therewith foreign-self discrimination. The other side is the development of computational methods that allow for an efficient analysis of probabilistic T-cell activation models that already exist, like the BRB model, or will emerge in the course of this thesis. In this chapter we introduce the relevant mathematical and computational background and establish a first theoretical result that will allow us in the next chapters to develop and modify such an efficient simulation method.

In probabilistic models of T-cell activation the main task is the estimation of the probability of T-cell activation given different parameter values. Hence, the general problem we now consider is to estimate the probability $P(A)$ of a (rare) event A under a probability measure P . The straightforward approach, known as simple sampling, uses the estimate

$$(\widehat{P(A)})_N := \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{S^{(i)} \in A\} = \frac{1}{N} \text{card}\{1 \leq i \leq N \mid S^{(i)} \in A\}, \quad (4.1)$$

where the $\{S^{(i)}\}_{1 \leq i \leq N}$ are independent and identically distributed (i.i.d.) random variables with distribution P , $\mathbb{1}\{.\}$ denotes the indicator function, and N is the sample size; we will throughout use \widehat{v} for an estimate of a quantity v . $(\widehat{P(A)})_N$ is obviously an unbiased and consistent estimate, but, for small $P(A)$, the convergence to $P(A)$ is slow, and large samples are required to get reliable estimates.

Various simulation methods are available that deal with this problem and yield a better rate of convergence under the right circumstances (see the monograph by Bucklew [25] for an overview). Most of them achieve this improvement by reducing the variance of the estimator. We will concentrate here on the most wide-spread class of methods, namely importance sampling. As is well known, one introduces a new sampling distribution Q here under which A is more likely to happen, produces samples from this distribution and gets back to the original distribution by reweighting. In general, finding a good importance sampling distribution that reduces the variance as much as possible is an art, and much of the literature revolves around this. Some “general purpose” and many ad hoc strategies exist, but usually, importance sampling distributions are best tailored by exploiting the structure of the specific problem at hand. However, if the problem can be embedded into a sequence of problems for which a so-called large deviation

principle is valid, a unified theory is available that identifies the most efficient simulation distribution. This technique of “large deviation simulation” was introduced by Sadowski and Bucklew [169], laid down in the monograph by Bucklew [25], and further developed by Dieker and Mandjes [58]. It rests on the well-established theory of large deviations, as summarised, for example, in the books by Dembo and Zeitouni [53] or den Hollander [88]. Let us recapitulate the basic background.

4.1 LARGE DEVIATION PROBABILITIES

Consider a sequence $\{S_n\}$ of random variables on the probability space $(\mathbb{R}^d, \mathcal{B}, \mathbb{P})$, where \mathcal{B} is the Borel σ -algebra of \mathbb{R}^d . Let $\{P_n\}$ be the family of probability measures induced by $\{S_n\}$, i.e., $P_n(B) = \mathbb{P}(S_n \in B)$ for $B \in \mathcal{B}$. We assume throughout that $\{S_n\}$ satisfies a large deviation principle (LDP) according to the following definition [53, 58]:

Definition 4.1 (Large deviation principle). A family of probability measures $\{P_n\}$ on $(\mathbb{R}^d, \mathcal{B})$ satisfies the large deviation principle (LDP) with rate function I if $I : \mathbb{R}^d \rightarrow [0, \infty]$ is lower semicontinuous and, for all $B \in \mathcal{B}$,

$$-\inf_{x \in B^\circ} I(x) \leq \liminf_{n \rightarrow \infty} \frac{1}{n} \log P_n(B) \leq \limsup_{n \rightarrow \infty} \frac{1}{n} \log P_n(B) \leq -\inf_{x \in \bar{B}} I(x), \quad (4.2)$$

where $B^\circ := \text{int}(B)$ and $\bar{B} := \text{clos}(B)$ denote the interior and the closure of B , respectively. I is said to be a good rate function if it has compact level sets in that $I^{-1}([0, c]) = \{x \in \mathbb{R}^d : I(x) \leq c\}$ is compact for all $c \in \mathbb{R}^d$. \square

A set B is called an *I-continuity set* if

$$\inf_{x \in B^\circ} I(x) = \inf_{x \in B} I(x) = \inf_{x \in \bar{B}} I(x). \quad (4.3)$$

If B is such a set, the LDP means that $P_n(B)$ decays exponentially for large n , with decay coefficient $\inf_{x \in B} I(x)$. A point b is called a *minimum rate point* of B if $\inf_{x \in B} I(x) = I(b)$.

Large deviation principles are well known for many families of random variables, like empirical means of i.i.d. random variables or empirical measures of Markov chains. For the application we have in mind, which involves sums of independent, but not identically distributed random variables, we need the fairly general setting of the Gärtner-Ellis theorem, which we recapitulate here (cf. [53, Thm. 2.3.6] and [88, Ch. V]). Let $\varphi_n(\vartheta) := \mathbb{E}_{P_n}(e^{\langle \vartheta, S_n \rangle})$, $\vartheta \in \mathbb{R}^d$, be the moment-generating function of S_n , where $\langle \cdot, \cdot \rangle$ denotes the scalar product and $\mathbb{E}_\mu(\cdot)$ denotes the expectation of a random variable with respect to the probability measure μ .

Theorem 4.2 (Gärtner-Ellis). Assume that

$$(G1) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \log \varphi_n(n\vartheta) =: \Lambda(\vartheta) \in [-\infty, \infty] \text{ exists,}$$

$$(G2) \quad 0 \in \text{int}(\mathcal{D}_\Lambda), \quad \text{where } \mathcal{D}_\Lambda := \{\vartheta \in \mathbb{R}^d : \Lambda(\vartheta) < \infty\} \text{ is the effective domain of } \Lambda,$$

(G3) Λ is lower semi-continuous on \mathbb{R}^d ,

(G4) Λ is differentiable on $\text{int}(\mathcal{D}_\Lambda)$,

(G5) Either $\mathcal{D}_\Lambda = \mathbb{R}^d$ or Λ is steep at its boundary $\partial\mathcal{D}_\Lambda$, i.e., $\lim_{\text{int}(\mathcal{D}_\Lambda) \ni \vartheta \rightarrow \partial\mathcal{D}_\Lambda} |\nabla\Lambda(\vartheta)| = \infty$.

Then, $\{P_n\}$ satisfies the LDP on \mathbb{R}^d with good rate function I , where I is the Legendre transform of Λ , i.e.,

$$I(x) = \sup_{\vartheta \in \mathbb{R}^d} [\langle x, \vartheta \rangle - \Lambda(\vartheta)], \quad x \in \mathbb{R}^d. \quad (4.4)$$

□

The function Λ in (G1) is convex. If there is a solution ϑ^* of

$$\nabla\Lambda(\vartheta) = x, \quad (4.5)$$

one has

$$I(x) = \langle \vartheta^*, x \rangle - \Lambda(\vartheta^*). \quad (4.6)$$

If Λ is strictly convex in all directions, ϑ^* is unique. See Fig. 5.1 for a one-dimensional example (the T-cell application, in fact).

4.2 SIMULATING RARE EVENT PROBABILITIES

Let now $A \in \mathcal{B}$ be a *rare event* in the sense that $0 < \inf_{x \in A} I(x) < \infty$. Here, the first inequality implies that A becomes exponentially unlikely as $n \rightarrow \infty$, whereas the second inequality serves to exclude nongeneric cases (in particular cases where the event is impossible). An important notion for the rare event simulation of $P_n(A)$ is that of a *dominating point* [25, p. 83]: A point a is a *dominating point* of the set A if it is the unique point such that

- a) $a \in \partial A$,
- b) \exists a unique solution ϑ^* of $\nabla\Lambda(\vartheta) = a$, and
- c) $A \subset \{x \in \mathbb{R}^d : \langle \vartheta^*, x - a \rangle \geq 0\}$.

A dominating point, if it exists, is always a unique minimum rate point (see [25, p. 83]). Convexity of A implies existence of a dominating point (cf. [58]).

Following [58] we now turn to the problem of simulating $P_n(A) = \mathbb{E}_{P_n}(\mathbb{1}\{S_n \in A\})$. The naive simple-sampling estimate obtained from N i.i.d. copies $S_n^{(i)}$ ($1 \leq i \leq N$), drawn from P_n , is, as in (4.1), given by

$$\widehat{(P_n(A))}_N := \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{S_n^{(i)} \in A\}. \quad (4.7)$$

It is clearly unbiased and converges (almost surely) to $P_n(A)$ in the limit $N \rightarrow \infty$, but, as clearly, it is inefficient since it requires that N increase exponentially with n to yield a meaningful estimate. Instead of $\{S_n\}$, one therefore considers an alternative family of

random variables, $\{T_n\}$ with distribution family $\{Q_n\}$, again on $(\mathbb{R}^d, \mathcal{B})$, under which A occurs more frequently. Assuming that P_n and Q_n are absolutely continuous with respect to each other, one can use the identity

$$P_n(A) = \mathbb{E}_{P_n}(\mathbb{1}\{S_n \in A\}) = \mathbb{E}_{Q_n}\left(\mathbb{1}\{T_n \in A\} \frac{dP_n}{dQ_n}(T_n)\right), \quad (4.8)$$

where dP_n/dQ_n is the Radon-Nikodym derivative of P_n with respect to Q_n . The resulting importance sampling estimate then relies on i.i.d. samples $T_n^{(i)}$ from $\{Q_n\}$ and reads

$$\widehat{(P_{Q_n}(A))}_N := \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{T_n^{(i)} \in A\} \frac{dP_n}{dQ_n}(T_n^{(i)}), \quad (4.9)$$

where $(dP_n/dQ_n)(\cdot)$ acts as a reweighting factor from the sampling distribution to the original one. It is reasonable to assume that (dP_n/dQ_n) is continuous to avoid the usual problems with L^1 -functions; this is no restriction for our targeted application.

An adequate optimality concept in this context is that of *asymptotic efficiency*. According to [58], it is based on the *relative error* $\eta_N(Q_n, A)$ defined via its square

$$\eta_N^2(Q_n, A) := \frac{\mathbb{V}_{Q_n}(\widehat{(P_{Q_n}(A))}_N)}{(P_n(A))^2} \quad (4.10)$$

(where $\mathbb{V}_\mu(\cdot)$ denotes the variance of a random variable with respect to the probability measure μ). The relative error is proportional to the width of the confidence interval relative to the (expected) estimate itself. Asymptotic efficiency is then defined as follows.

Definition 4.3 (Asymptotic efficiency). An importance sampling family $\{Q_n\}$ is called *asymptotically efficient* for the rare event A if

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log N_{Q_n}^* = 0, \quad (4.11)$$

where $N_{Q_n}^* := \inf\{N \in \mathbb{N} : \eta_N(Q_n, A) \leq \eta_{\max}\}$ for some given maximal relative error η_{\max} , $0 < \eta_{\max} < \infty$.

In words, asymptotic efficiency means that the number of samples required to keep the relative error below a prescribed bound η_{\max} increases only subexponentially (rather than exponentially as with simple sampling). The concrete choice of η_{\max} is actually irrelevant, see Lemma 1 in [58].

An obvious idea from large deviation theory would be to use, as sampling distributions, the family of measures $\{P_n^\vartheta\}$ that are exponentially tilted with parameter ϑ , that is,

$$\frac{dP_n^\vartheta}{dP_n}(x) = \frac{e^{n\langle \vartheta, x \rangle}}{\varphi_n(n\vartheta)}, \quad x \in \mathbb{R}^d; \quad (4.12)$$

P_n^ϑ then takes the role of Q_n . The task remains to find “the right” ϑ , i.e., a (or the) tilting parameter that makes $\{P_n^\vartheta\}$ asymptotically efficient (if at all possible). Necessary and sufficient conditions for this are given in [58, Assumption 1 and Corollary 1] and are summarised below, in a form adapted to the present context.

Theorem 4.4 (Dieker-Mandjes 2005). Assume that, for some given ϑ^* ,

(V1) $\{P_n\}$ satisfies an LDP with good rate function I ,

(V2) $\limsup_{n \rightarrow \infty} \frac{1}{n} \log \varphi_n(\gamma n \vartheta^*) < \infty$ for some $\gamma > 1$, and, likewise, with ϑ^* replaced by $-\vartheta^*$,

(V3) The rare event A is both an I -continuity set and an $(I + \langle \vartheta^*, \cdot \rangle)$ -continuity set.

Then, the tilted measure $\{P_n^{\vartheta^*}\}$ is asymptotically efficient for simulating A if and only if

$$\inf_{x \in \mathbb{R}^d} [I(x) - \langle \vartheta^*, x \rangle] + \inf_{x \in \bar{A}} [I(x) + \langle \vartheta^*, x \rangle] = 2 \inf_{x \in A^\circ} I(x). \quad (4.13)$$

We use assumption (V2) here to replace the weaker but less easy to verify condition (2) in Assumption 1 of [58], in line with the paragraph below (2) in [58], or [53, Thm. 4.3.1]. Note also that (V2) holds automatically if $\varphi_n(n\vartheta)$ exists for all ϑ – but this is not mandatory here, since only a given ϑ^* is considered.

The proof of Theorem 4.4 is given in [58] and need not be recapitulated here; but we would like to comment briefly on what happens in the central condition (4.13). Replacing Q_n by $P_n^{\vartheta^*}$ in (4.10) and (4.9), we can rewrite η_N^2 as

$$\begin{aligned} \eta_N^2(P_n^{\vartheta^*}, A) &= \frac{\mathbb{V}_{P_n^{\vartheta^*}}(\widehat{P_{P_n^{\vartheta^*}}(A)})_N}{(P_n(A))^2} = \frac{1}{N} \frac{\mathbb{V}_{P_n^{\vartheta^*}}(\widehat{P_{P_n^{\vartheta^*}}(A)})_1}{(P_n(A))^2} \\ &= \frac{1}{N} \frac{1}{(P_n(A))^2} \left[\int_A \left(\frac{dP_n}{dP_n^{\vartheta^*}} \right)^2 dP_n^{\vartheta^*} - (P_n(A))^2 \right]. \end{aligned} \quad (4.14)$$

Obviously (by (V1) and (V3)), $2 \inf_{x \in A^\circ} I(x)$ (i.e., the right-hand side of (4.13)) is the exponential decay rate of $(P_n(A))^2$. Inspection of the proof of Theorem 4.4 reveals that the left-hand side of (4.13) is the exponential decay rate of $\int_A \left(\frac{dP_n}{dP_n^{\vartheta^*}} \right)^2 dP_n^{\vartheta^*}$. It is clear from (4.14) that, for asymptotic efficiency to hold, $\int_A \left(\frac{dP_n}{dP_n^{\vartheta^*}} \right)^2 dP_n^{\vartheta^*}$ must tend to zero at least as fast as $(P_n(A))^2$. But it cannot decrease faster, since $\mathbb{V}_{P_n^{\vartheta^*}}(\widehat{P_{P_n^{\vartheta^*}}(A)})_1$ is nonnegative, so that $\int_A \left(\frac{dP_n}{dQ_n} \right)^2 dQ_n \geq (P_n(A))^2$ for arbitrary Q_n . Hence, the exponential decay rates must be exactly equal, as stated by (4.13). (A closely related argument is given in [25, Ch. 5.2].)

Theorem 4.4 is widely applicable. It holds in many standard situations, in particular in many of those that arise in applications.

Proposition 4.5. Let $\{P_n\}$ be a family of probability measures that satisfy the conditions of the Gärtner-Ellis theorem, with (good) rate function I . Let A be a rare event with dominating point a , let ϑ^* be the unique solution of $\nabla \Lambda(\vartheta) = a$, and assume (V2) and (V3). Then $\{P_n^{\vartheta^*}\}$ is the unique tilted family that is asymptotically efficient for simulating $P_n(A)$.

Proof. The proof is a simple application of Thm. 4.4. (V1) follows from the Gärtner-Ellis theorem; we only need to verify condition (4.13). For the first infimum in (4.13), one obtains

$$\inf_{x \in \mathbb{R}^d} [I(x) - \langle \vartheta^*, x \rangle] = -\Lambda(\vartheta^*) = I(a) - \langle \vartheta^*, a \rangle. \quad (4.15)$$

Here, the first step follows from the convex duality lemma (compare [53, Lemma 4.5.8]), which is applicable since Λ is lower semicontinuous by (G3), and convex and $> -\infty$ everywhere (this follows from (G1) and (G2) by [88, Lemma V.4]). The second step is due to part b) of the dominating point property of a , together with Eq. (4.6).

As to the second infimum in (4.13), note that, by the dominating point property, a minimises both I and $\langle \vartheta^*, \cdot \rangle$ on A , which, together with (V3), gives

$$\inf_{x \in \bar{A}} [I(x) + \langle \vartheta^*, x \rangle] = I(a) + \langle \vartheta^*, a \rangle. \quad (4.16)$$

Eqs. (4.15) and (4.16) together give (4.13) because $\inf_{x \in A^\circ} I(x) = \inf_{x \in \partial A} I(x) = I(a)$. \square

Remark 4.6. Note that an efficiency result closely related to Proposition 4.5 has previously been given by Bucklew [25, Thm. 5.2.1], but this is based on the variance rather than the relative error; and it is only a sufficient condition.

Note that our assumption of a dominating point greatly simplifies the situation. Theorem 2 also allows to cope with situations without a dominating point – but this is not needed below.

The theory presented so far is general enough to provide asymptotically efficient simulation distributions for a wide range of problems. Nevertheless, its concrete applications have, so far, been somewhat restricted, in three ways:

1. Concrete applications require the availability of an explicit large deviation result; in particular, the $\Lambda(\vartheta)$ appearing in the Gärtner-Ellis theorem, and the corresponding tilting parameter ϑ^* , must be known in a more explicit form, or must at least be easily accessible numerically. So far, popular examples include sums of i.i.d. random variables, sums of a functional of a Markov chain, level crossing problems, or queueing problems (see [25] for an overview). For example, for sums of i.i.d. random variables, the Gärtner-Ellis theorem reduces to Cramér's theorem.
2. A simulation method must be available to sample from the corresponding tilted distribution. In the case of sums of independent random variables, tilting of the sum with parameter ϑ is simply achieved by tilting each summand with ϑ . If the individual terms have uniform, Bernoulli, exponential, or Gaussian distributions, their tilted variants are known explicitly (for example, the exponential distribution with parameter λ , $\text{Exp}(\lambda)$, turns into $\text{Exp}(\lambda - \vartheta)$ under tilting with ϑ), and are easily simulated via transformation of random variables from $\text{Uni}_{[0,1]}$, the uniform distribution on the unit interval.

3. So far, concrete applications have mainly come from within stochastic processes as such (like queueing theory), or information theory (like models of digital communication systems).

Rare event simulation is very successful for the above examples – but the range of applications is still small as compared to the generality of the theory; for example, “LDP sampling” is hardly used in physics and biology. Our immunological problem thus provides a nice opportunity to extend the range of applications.

ANALYSIS AND EXTENSION OF THE BRB MODEL OF T-CELL ACTIVATION

In the previous chapter we introduced the BRB T-cell activation model and the theoretical background for the development of efficient simulation methods. In this chapter we use the BRB model as a starting point for the development of ideas on how foreign-self discrimination by T-cells is achieved. This model was already analysed to a certain extent with the help of analytical and numerical calculations and the capability for foreign-self discrimination could be shown [205, 232]. In order to analyse it even deeper and more thoroughly we develop a new simulation technique and use these previous results in order to confirm our first results from our analysis. Building upon new insights we get by our analysis we go one step further and extend the basic BRB model. These extensions are also analysed by a modified version of our newly developed simulation method.

5.1 RARE EVENT SIMULATION: THE T-CELL MODEL

At first, we have to develop an efficient simulation method in order to estimate the activation probabilities in the basic BRB model. Recall that simulating the T-cell model means sampling the random variables $G(z^{(f)})$ of (3.3) and estimating the corresponding tail probabilities $\mathbb{P}(G(z^{(f)}) \geq g_{act})$. Inspection of Eq. (3.3) reveals two difficulties, which correspond to 1. and 2. in the previous Section:

1. $G(z^{(f)})$ is a weighted sum of i.i.d. random variables, to which the standard results for sums of i.i.d. random variables (in particular, Cramér’s theorem) are not applicable. We therefore need an extension to weighted sums – or, better, to general sums of independent, but not identically distributed random variables, which include weighted sums as a simple special case. This is straightforward and will be the subject of Sect. 5.1.1. In particular, it will be seen that, like in the i.i.d. case, every term in the sum must be tilted with the same parameter, but now this “global” tilting factor is a function of all the individual distributions involved.
2. Simulating the random variables $W_j = w(\mathcal{T}_j)$ is straightforward via simple sampling: draw $\text{Exp}(1/\bar{\tau})$ distributed random numbers τ_j (as realisations of \mathcal{T}_j) and apply the transformation (3.1). However, simulating the corresponding tilted variables is a difficult task, for two reasons. First of all, there is no clue how to sample from the tilted distribution via transformation of one of the “basic” distributions (like $\text{Uni}_{[0,1]}$, or $\text{Exp}(\lambda)$) for which efficient random number generation is possible.

Although such a transformation might exist in principle, there is no systematic way of finding it. One reason for this is that tilting acts at the level of the densities, but even the original (untilted) density of $W = w(\mathcal{T})$ is not available explicitly. (With W and \mathcal{T} (without indices) we mean any representative of the family.) This is because its calculation requires the inverse functions and derivatives of the two branches (increasing and decreasing) of the function w , but these are unavailable analytically.

In lack of a transformation method, one might consider to determine the tilted density numerically, integrate it (again numerically) and discretise and tabulate the resulting distribution function. However, this is, again, forbidding for our particular function w : Due to the vanishing derivatives at $\mathcal{T} = 0$ and $\mathcal{T} = 1$, the transformation formula for densities yields singularities in the density of W at these values, with a sizeable fraction of the probability mass concentrated very close to 0 (see Fig. 3.2). This renders numerical calculations unreliable. To circumvent these problems, we will, in Sect. 5.1.2, present a sampling method for the tilted random variable W^ϑ that is based on tilting \mathcal{T} rather than W itself.

5.1.1 Large deviations for independent but not identically distributed random variables

We consider K independent families of i.i.d. \mathbb{R}^d -valued random variables, $\{Y_\ell^{(1)}\}, \dots, \{Y_\ell^{(K)}\}$ (i.e., the distribution within any given family $\{Y_\ell^{(k)}\}$, $1 \leq k \leq K$, is fixed, but the distributions may vary across families). Assume that $\Lambda^{(k)}(\vartheta) := \log \mathbb{E}(e^{\langle \vartheta, Y_1^{(k)} \rangle})$, the log moment-generating function of $Y_1^{(k)}$, is finite for all $\vartheta \in \mathbb{R}^d$ and $1 \leq k \leq K$ (here, $\mathbb{E}(\cdot)$ refers to the probability measure induced by the random variable involved). Let $n^{(1)}, \dots, n^{(K)}$ be positive integers, $n := \sum_{k=1}^K n^{(k)}$,

$$V_n := \sum_{\ell=1}^{n^{(1)}} Y_\ell^{(1)} + \dots + \sum_{\ell=1}^{n^{(K)}} Y_\ell^{(K)}, \quad (5.1)$$

and P_n be the probability measure induced by $S_n = V_n/n$. In the limit $n \rightarrow \infty$, subject to $n^{(k)}/n \rightarrow \gamma^{(k)}$ for all $1 \leq k \leq K$, the limiting log-moment generating function of $\{S_n\}$ becomes

$$\Lambda(\vartheta) = \lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{E}(e^{\langle \vartheta, V_n \rangle}) = \lim_{n \rightarrow \infty} \sum_{k=1}^K \frac{n^{(k)}}{n} \Lambda^{(k)}(\vartheta) = \sum_{k=1}^K \gamma^{(k)} \Lambda^{(k)}(\vartheta), \quad (5.2)$$

where the second step is due to independence. Since, by assumption, $\Lambda^{(k)}(\vartheta) < \infty$ for all $\vartheta \in \mathbb{R}^d$ and $1 \leq k \leq K$, the $\Lambda^{(k)}$ are differentiable on all of \mathbb{R}^d (see [53, Lemma 2.2.31]); in fact, they are even $C^\infty(\mathbb{R}^d)$ [53, Exercise 2.2.24]. Thus, Λ is $C^\infty(\mathbb{R}^d)$ as well.

By (5.2), we have (G1). Again due to $\Lambda^{(k)}(\vartheta) < \infty$, (G2) and (G5) are automatically satisfied. Furthermore, the differentiability of Λ entails (G3) and (G4). We have therefore shown

Lemma 5.1. Under the assumptions of this paragraph, $\{P_n\}$ satisfies the Gärtner-Ellis theorem, with rate function I given by Eq. (4.4). \square

Such $\{P_n\}$ are therefore candidates for efficient simulation according to Prop. 4.5. The tilting factor ϑ^* may not be accessible analytically, but can be evaluated numerically from (4.5). Due to independence, tilting of S_n with $n\vartheta^*$ (that is, tilting of V_n with ϑ^*) is equivalent to tilting each $Y_\ell^{(k)}$ with ϑ^* .

5.1.2 Tilting of transformed random variables

Unlike the W_j , the $\text{Exp}(1/\bar{\tau})$ -distributed random variables \mathcal{T}_j are tilted easily (tilting with ϑ simply gives $\text{Exp}(1/\bar{\tau} - \vartheta)$). One is therefore tempted to tilt the \mathcal{T}_j rather than the W_j , or, in other words, to interchange the order of tilting and transformation. The following Theorem states the key idea.

Theorem 5.2. Let X be an \mathbb{R}^d -valued random variable with probability measure μ , and let $Y := h \circ X$ (or $Y = h(X)$ by slight abuse of notation), where $h : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is μ -measurable. Then Y has probability measure $\nu = \mu \circ h^{-1}$, where h^{-1} denotes the preimage of y . Assume now that $\mathbb{E}_\mu(e^{\langle \vartheta, h(X) \rangle})$ exists, let \tilde{X}^ϑ be an \mathbb{R}^d -valued random variable with probability measure $\tilde{\mu}^\vartheta$ related to μ via

$$\frac{d\tilde{\mu}^\vartheta}{d\mu}(x) = \frac{e^{\langle \vartheta, h(x) \rangle}}{\mathbb{E}_\mu(e^{\langle \vartheta, h(X) \rangle})} \quad (5.3)$$

(so that $\tilde{\mu}^\vartheta \ll \mu$), and let $\tilde{Y}^\vartheta = h(\tilde{X}^\vartheta)$. Then, the measures $\tilde{\nu}^\vartheta$ (of \tilde{Y}^ϑ) and ν^ϑ (for the tilted version of ν , belonging to Y^ϑ) are equal, where $\nu^\vartheta \ll \nu$ with Radon-Nikodym density

$$\frac{d\nu^\vartheta}{d\nu}(y) = \frac{e^{\langle \vartheta, y \rangle}}{\mathbb{E}_\nu(e^{\langle \vartheta, Y \rangle})}. \quad (5.4)$$

Proof. Note first that $e^{\langle \vartheta, y \rangle}$ is clearly μ -measurable, and

$$\mathbb{E}_\nu(e^{\langle \vartheta, Y \rangle}) = \int_{\mathbb{R}^d} e^{\langle \vartheta, y \rangle} d\nu(y) = \int_{\mathbb{R}^d} e^{\langle \vartheta, h(x) \rangle} d\mu(x) = \mathbb{E}_\mu(e^{\langle \vartheta, h(X) \rangle}), \quad (5.5)$$

which exists by assumption, so ν^ϑ is well-defined. We now have to show that $\tilde{\nu}^\vartheta(B) = \nu^\vartheta(B)$ for arbitrary Borel sets B . Observing that $\tilde{\nu}^\vartheta = \tilde{\mu}^\vartheta \circ h^{-1}$ and employing the formulas for transformation of measures [17, (13.7)] and change of variable [17, Thm. 16.13], together with (5.3), one indeed obtains

$$\begin{aligned} \tilde{\nu}^\vartheta(B) &= \tilde{\mu}^\vartheta(h^{-1}(B)) = \int_{h^{-1}(B)} \frac{d\tilde{\mu}^\vartheta}{d\mu}(x) d\mu(x) \\ &= \frac{1}{\mathbb{E}_\mu(e^{\langle \vartheta, h(X) \rangle})} \int_{h^{-1}(B)} e^{\langle \vartheta, h(x) \rangle} d\mu(x) = \frac{1}{\mathbb{E}_\nu(e^{\langle \vartheta, Y \rangle})} \int_B e^{\langle \vartheta, y \rangle} d\nu(y) \\ &= \int_B \frac{d\nu^\vartheta}{d\nu}(y) d\nu(y) = \nu^\vartheta(B), \end{aligned} \quad (5.6)$$

which proves the claim. \square \square

In words, Theorem 5.2 is nothing but the simple observation that, to obtain the tilted version of $Y = h(X)$, one can reweight the measure μ of X with the factors $e^{\langle \vartheta, h(x) \rangle}$, rather than reweighting the measure ν of Y with $e^{\langle \vartheta, y \rangle}$. It should be clear, however, that the measure $\tilde{\mu}^\vartheta$ differs from the “usual” tilted version of μ , which would involve tilting factors $e^{\langle \vartheta, x \rangle}$ rather than $e^{\langle \vartheta, h(x) \rangle}$; for this reason, we use the notation $\tilde{\mu}^\vartheta$ rather than μ^ϑ . Nevertheless, this simple observation is the key to simulation if μ (and $\tilde{\mu}^\vartheta$) are readily accessible at least numerically, but ν (and ν^ϑ) are not.

This is exactly our situation, with \tilde{T}^ϑ , αW^ϑ and αw ($\alpha \in \{qz^{(c)}, qz^{(v)}, z^{(f)}\}$), respectively, taking the roles of \tilde{X}^ϑ , Y^ϑ and h (we will use f , \tilde{f}^ϑ , g and g^ϑ for the corresponding densities of \mathcal{T} , \tilde{T}^ϑ , αW , and $(\alpha W)^\vartheta$). Still, reweighting of the exponential density of \mathcal{T} with $e^{\vartheta \alpha w(\tau)}$ does not yield an explicit closed-form density (let alone an exponential one), and no direct simulation method is available for the corresponding random variables. However, the reweighted densities are easily accessible numerically, in contrast to those of W and its tilted variant, W^ϑ . The way to go is therefore to calculate and integrate \tilde{f}^ϑ numerically and discretise and tabulate the resulting distribution function \tilde{F}^ϑ . Samples of \tilde{T}^ϑ may then be drawn according to this table (i.e., by formally “looking up” the solution of $\tilde{F}(\tilde{T}^\vartheta) = U$ for $U \sim \text{Uni}_{[0,1]}$), and $\alpha W^\vartheta = \alpha w(\tilde{T}^\vartheta)$ is then readily evaluated. The only difficulty left is the time required for searching the table. But this is a practical matter and will be dealt with in the next paragraph.

5.1.3 The algorithm

Taking together our theoretical results, we can now spell out the specific importance sampling algorithm for the simulation of the T-cell model of Sect. 3. If not stated otherwise, we will refer to the basic model (3.3). Recall that it describes the stimulation rate $G(z^{(f)})$ and we wish to evaluate the probability $\mathbb{P}(G(z^{(f)}) \geq g_{\text{act}})$.

To apply LD sampling, let us embed the model into a sequence of models with increasing total number $n = n^{(c)} + n^{(v)} + n^{(f)}$ of antigen types, where $n^{(c)}$, $n^{(v)}$, and $n^{(f)}$ are the numbers of constitutive, variable and foreign antigen types. Let

$$G_n(z^{(f)}) = \left(\sum_{j=1}^{n^{(c)}} q_n z^{(c)} W_j \right) + \left(\sum_{j=n^{(c)}+1}^{n^{(c)}+n^{(v)}} q_n z^{(v)} W_j \right) + \left(\sum_{j=n^{(c)}+n^{(v)}+1}^{n^{(c)}+n^{(v)}+n^{(f)}} z^{(f)} W_{n^{(c)}+n^{(v)}+j} \right), \quad (5.7)$$

where

$$q_n = \frac{n^{(c)} z^{(c)} + n^{(v)} z^{(v)} - n^{(f)} z^{(f)}}{n^{(c)} z^{(c)} + n^{(v)} z^{(v)}} \quad (5.8)$$

(where $z^{(c)}$, $z^{(v)}$, and $z^{(f)}$ are independent of n). Clearly, $G_n(z^{(f)})$ coincides with $G(z^{(f)})$ of (3.3) if $n^{(c)} = m^{(c)}$, $n^{(v)} = m^{(v)}$, and $n^{(f)} = m^{(f)}$, where $m^{(f)} = 0$ or $m^{(f)} = 1$ depending on whether $z^{(f)} = 0$ or $z^{(f)} > 0$; then, $n = m = m^{(c)} + m^{(v)} + m^{(f)}$. We have to consider $\mathbb{P}(G_n(z^{(f)})/n > g_{\text{act}}/m)$ (this reflects the fact that g_{act} must scale with system size). The sequences $\{G_n(z^{(f)})\}$ and $\{G_n(z^{(f)})/n\}$ take the roles of $\{V_n\}$ and $\{S_n\}$, respectively, in Secs. 4.1 and 5.1.1, with P_n the law of $G_n(z^{(f)})/n$; and we consider $A = [g_{\text{act}}/m, \infty)$ with $\mathbb{E}(G_m(z^{(f)})/m) < g_{\text{act}}/m < Mw(1)/m$ (the latter is

the maximum value of $G_m(z^{(f)})/m$ since $w(\tau)$ has its maximum at $\tau = 1$). The limit $n \rightarrow \infty$ is then taken so that $\lim_{n \rightarrow \infty} n^{(c)}/n = m^{(c)}/m$, $\lim_{n \rightarrow \infty} n^{(v)}/n = m^{(v)}/m$, as well as $\lim_{n \rightarrow \infty} n^{(f)}/n = m^{(f)}/m$, that is, the relative amounts of constitutive, variable, and foreign antigens approach those fixed in the original model, (3.3). (Note that, in [232], a different limit was employed, namely, $n \rightarrow \infty$ with $\lim_{n \rightarrow \infty} n^{(c)}/n^{(v)} = C_1 \in (0, \infty)$ and $\lim_{n \rightarrow \infty} n^{(f)}/n = 0$; this is appropriate for exact asymptotics, but not for simulation, because the asymptotic tilting factor to be used in the latter then does not “feel” the foreign antigens.)

Lemma 5.3. Let f be the density of $\text{Exp}(1/\bar{\tau})$ (i.e., $f(\tau) = e^{-\tau/\bar{\tau}}/\bar{\tau}$), and

$$\psi(t) := \mathbb{E}(e^{tW}) = \int_0^\infty \exp(tw(\tau))f(\tau)d\tau = \frac{1}{\bar{\tau}} \int_0^\infty \exp\left(t\frac{\exp(-1/\tau)}{\tau} - \frac{\tau}{\bar{\tau}}\right) d\tau \quad (5.9)$$

be the moment-generating function of W_1 . Under the assumptions of Sect. 5.1.3, the unique solution ϑ^* of

$$\begin{aligned} \frac{g_{\text{act}}}{m} &= \frac{m^{(c)}}{m} qz^{(c)} \left[\frac{d}{dt} \log \psi(t) \right] \Big|_{t=qz^{(c)}\vartheta} + \frac{m^{(v)}}{m} qz^{(v)} \left[\frac{d}{dt} \log \psi(t) \right] \Big|_{t=qz^{(v)}\vartheta} \\ &+ \frac{1}{m} z^{(f)} \left[\frac{d}{dt} \log \psi(t) \right] \Big|_{t=z^{(f)}\vartheta} \end{aligned} \quad (5.10)$$

is the unique asymptotically efficient tilting parameter for LD simulation of $P_n(A)$.

Proof. Clearly, P_n satisfies the assumptions of Sect. 5.1.1. Note, in particular, that $\psi(t) < \infty$ for all $t \in \mathbb{R}$ since W is bounded above and below, and so

$$\Lambda(\vartheta) = \lim_{n \rightarrow \infty} \log \mathbb{E}(e^{\vartheta G_n(z^{(f)})/n}) = \frac{m^{(c)}}{m} \log \psi(qz^{(c)}\vartheta) + \frac{m^{(v)}}{m} \log \psi(qz^{(v)}\vartheta) + \frac{1}{m} \log \psi(z^{(f)}\vartheta) < \infty \quad (5.11)$$

for all ϑ ; hence, the Gärtner-Ellis theorem holds by Lemma 5.1. To verify the remaining assumptions of Prop. 4.5, recall from Sec. 5.1.1 that $\Lambda(\vartheta)$ is differentiable (with continuous derivative) on all of \mathbb{R} . The bounds on g_{act}/m lead to

$$\Lambda'(0) = \frac{\mathbb{E}(G(z^{(f)}))}{m} < \frac{g_{\text{act}}}{m} < \frac{Mw(1)}{m} = \lim_{\vartheta \rightarrow \infty} \Lambda'(\vartheta). \quad (5.12)$$

Λ is strictly convex (since $(d^2/dt^2) \log \psi(t)$ is the variance of W^t , the tilted version of W (cf. [9, Prop. XII.1.1]), which is positive since W and hence W^t is nondegenerate). Eq. (5.12) thus entails that $\Lambda'(\vartheta) = g_{\text{act}}/m$ has a unique solution ϑ^* , which is positive (and clearly satisfies (V2)). As a consequence, g_{act}/m is a dominating point of A , which is a rare event since $0 < I(g_{\text{act}}/m) < \infty$ (by $\Lambda(0) = 0$ together with (5.12) and (4.6); cf. Fig. 5.1, left). Finally, A is a continuity set of both I and $I + \langle \vartheta^*, \cdot \rangle$ simply because I and $\langle \vartheta^*, \cdot \rangle$ are continuous at g_{act}/m , and $A = \overline{A^\circ}$. Realising that the right-hand side of (5.10) equals $\Lambda'(\vartheta)$ (see also Eq. (20) in [232]), one obtains the claim from Prop. 4.5. \square \square

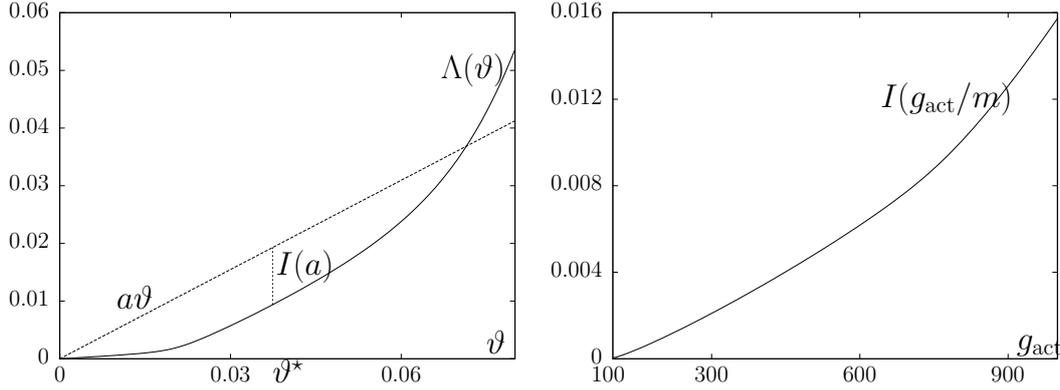


Figure 5.1: The cumulant-generating function Λ (left) and the rate function I (right) for the T-cell model (3.3). The slope of the straight line in the left panel is $a = g_{\text{act}}/m$, where $g_{\text{act}} = 800$ and $m = 1551$. At ϑ^* , $\alpha\vartheta - \Lambda(\vartheta)$ assumes its maximum, $I(a)$ (cf. (4.4)–(4.6)).

The solution of (5.10) is readily calculated numerically. The function Λ , and the resulting rate function I , are shown in Fig. 5.1.

As described in Sect. 5.1.2, we now tilt the density f of the \mathcal{T}_j with ϑ^* according to Eq. (5.3). This yields three different densities $\tilde{f}_\alpha^{\vartheta^*}$, depending on the weighting factors $\alpha \in \{qz^{(c)}, qz^{(v)}, z^{(f)}\}$, namely

$$\tilde{f}_\alpha^{\vartheta^*}(\tau) = \frac{\exp(\alpha\vartheta^*w(\tau))f(\tau)}{\psi(\alpha\vartheta^*)} = \frac{\frac{1}{\tau} \exp\left(\alpha\vartheta^* \frac{\exp(-1/\tau)}{\tau} - \frac{\tau}{\tau}\right)}{\psi(\alpha\vartheta^*)}. \quad (5.13)$$

As discussed in Sect. 5.1.2, this is not the density of any known standard distribution (let alone an exponential one), and simulating from it requires numerical integration (which is well-behaved since the $\tilde{f}_\alpha^{\vartheta^*}$ are numerically well-behaved), and discretisation and tabulation of the resulting distribution functions $\tilde{F}_\alpha^{\vartheta^*}$, followed by “looking up” the solution $\tilde{\tau}^{\vartheta^*}$ of $\tilde{F}_\alpha^{\vartheta^*}(\tilde{\mathcal{T}}^{\vartheta^*}) = U$ for $U \sim \text{Uni}_{[0,1]}$, to finally yield αW^{ϑ^*} via $\alpha W^{\vartheta^*} = \alpha w(\tilde{\mathcal{T}}^{\vartheta^*})$.

Searching the table would be the speed- (or precision-) limiting step, requiring $\mathcal{O}(\log K)$ operations if K is the number of discretisation steps. This can be remedied by applying the so-called *alias method* to quickly generate random variables according to the discretised probability distribution. For a description of the method, we refer the reader to [125, pp. 25–27], [113], or [165, p. 248]. Let us just summarise here that, after a preprocessing step, which is done once for a given distribution, the method only requires one $\text{Uni}_{[0,1]}$ random variable together with one multiplication, one cutoff and one subtraction (or two $\text{Uni}_{[0,1]}$ random variables together with one multiplication, one cutoff and one comparison, depending on the implementation) to generate one realisation of $\tilde{\mathcal{T}}^{\vartheta^*}$, regardless of K (in particular, it does without searching altogether).

We now have everything at hand to formulate the algorithm to simulate (realisations of) $G(z^{(f)})$ of (3.3). (For notational convenience, we will not distinguish between random variables and their realisations here).

Algorithm 1: Estimation of the activation probabilities in the BRB model.

Input: activation threshold g_{act} , foreign antigen copy number $z^{(f)}$, number of samples $N \in \mathbb{N}$

Result: estimate of $\mathbb{P}(G(z^{(f)}) \geq g_{\text{act}})$

1 compute ϑ^* by solving Eq. (5.10) numerically

2 calculate the tilted densities $\tilde{f}_{\alpha}^{\vartheta^*}$, $\alpha \in \{qz^{(c)}, qz^{(v)}, z^{(f)}\}$, via (5.13)

3 **for** $i = 1, \dots, N$ **do**

4 for every summand j of (3.3) generate a sample $(\tilde{\mathcal{T}}_j^{\vartheta^*})^{(i)}$ according to its density $\tilde{f}_{\alpha(j)}^{\vartheta^*}$ with the help of the alias method (here, the upper index (i) is added to reflect sample i , and $\alpha(j)$ is the weighting factor of the sum to which j belongs)

5 calculate

$$(G(z^{(f)}))^{(i)} = \left(\sum_{j=1}^{m^{(c)}} qz^{(c)} w((\tilde{\mathcal{T}}_j^{\vartheta^*})^{(i)}) \right) + \left(\sum_{j=m^{(c)+1}^{m^{(c)}+m^{(v)}} qz^{(v)} w((\tilde{\mathcal{T}}_j^{\vartheta^*})^{(i)}) \right) + z^{(f)} w((\tilde{\mathcal{T}}_{m^{(c)}+m^{(v)}+1}^{\vartheta^*})^{(i)})$$

6 calculate the indicator function times the reweighting factor (i.e., the i -th summand in Eq. (4.9))

7 **if** $(G(z^{(f)}))^{(i)} \geq g_{\text{act}}$ **then**

8 $R^{(i)} = \prod_{j=1}^m \frac{f_{\alpha(j)}((\tilde{\mathcal{T}}_j^{\vartheta^*})^{(i)})}{\tilde{f}_{\alpha(j)}^{\vartheta^*}((\tilde{\mathcal{T}}_j^{\vartheta^*})^{(i)})}$

9 **else**

10 $R^{(i)} = 0$

11 **end**

12 **end**

13 calculate $(\widehat{P_{P_m}^{\vartheta^*}}(A))_N = \frac{\sum_{i=1}^N R^{(i)}}{N}$, as estimate of $\mathbb{P}(G(z^{(f)}) > g_{\text{act}})$

Note that simulation of the extended model (3.4) is a straightforward generalisation (see also [232] for the explicit LD theory).

5.2 RESULTS

Let us now present the results of our simulations in two steps. We first investigate the performance of the method, and then use it to gain more insight into the underlying phenomenon of statistical recognition.

5.2.1 Performance of the simulation method

We will examine the performance of the importance-sampling method in three respects: we will compare it to simple sampling (the previously-used simulation method) and to the results of exact asymptotics (the previously-used analytic method); finally, we will quantify the efficiency in terms of the relative error (and thus get back to the theory of Sect. 4.2). In any case, we will consider $\mathbb{P}(G(z^{(f)}) \geq g_{\text{act}})$ as a function of g_{act} (and for various values of the parameter $z^{(f)}$). Of course, this probability is just one minus the distribution function of $G(z^{(f)})$; in immunobiology, the corresponding graph is known as the activation curve.

Evaluating this graph by LD simulation requires, for each value of g_{act} to be considered, a fresh sample, simulated with its individual tilting factor ϑ^* (recall that this depends on g_{act} via (5.10)). At first sight, this looks like an enormous disadvantage relative to simple sampling, where no threshold needs to be specified in advance; rather, the outcomes of the simulation directly yield an estimate over the entire range of the activation curve. However, it will turn out that this disadvantage is offset manifold by the specific efficiency of “hitting” the rare events in LD sampling. (There is still room for improvement: The samples that do not “hit” a given rare event could be used to improve the estimates of the more likely events.)

Comparison with simple sampling

Clearly, both the simple-sampling and the importance-sampling estimates are unbiased and converge to the true values as $N \rightarrow \infty$. It is therefore no surprise that they yield practically identical results wherever they can be compared – and this yields a first quick consistency check for our method.

This is demonstrated in Fig. 5.2, which shows simple sampling (SS) and importance sampling (IS) activation curves, each for $z^{(f)} = 1000$ and $z^{(f)} = 2000$. For SS, $N = 1.3 * 10^8$ samples, $G^{(i)}(z^{(f)}), 1 \leq i \leq N$, were generated altogether for every graph, whereas for IS, $N = 10000$ samples were generated for every threshold value considered (from $g_{\text{act}} = 100$ to $g_{\text{act}} = 1000$ in steps of 50), i.e. $1.9 * 10^5$ samples altogether. Beyond $g_{\text{act}} = 450$ and $g_{\text{act}} = 800$ (for $z^{(f)} = 1000$ and $z^{(f)} = 2000$, respectively), no estimates could be obtained via SS due to the low probabilities involved, whereas with IS, it is easy to get beyond $g_{\text{act}} = 900$ in either case, although the probabilities can get down to 10^{-20} (note, however, that this far end of the distribution is no longer biologically relevant). In terms of runtime, determining an activation curve (over its entire range) by SS took 48 hours of CPU time (Intel Pentium M 1.4 GHz 512MB RAM), whereas IS required only about 2 minutes (in the threshold regime where the methods are comparable), that is, a speedup by a factor of nearly 1500 is achieved.

We also applied our method to the extended model (3.4) with binomially distributed copy numbers. Figure 5.3 shows the simulation results for two values of $z^{(f)}$, each for SS and IS. Again, the curves agree, as they must. As to runtime, it took about 130 hours to generate the $2 * 10^7$ samples for SS, whereas for IS it took 10 min. to generate the $9.5 * 10^4$ samples.

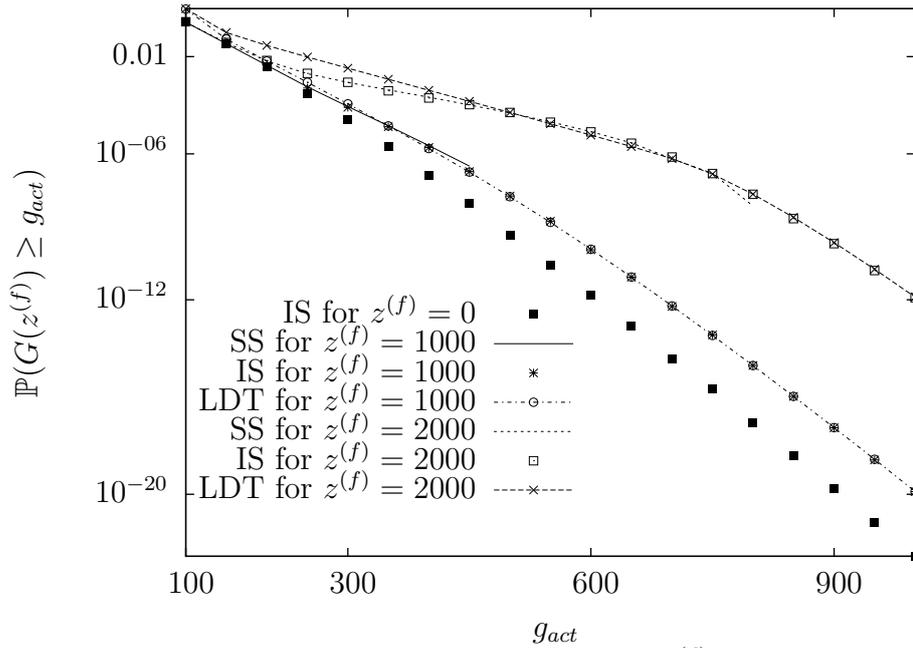


Figure 5.2: Estimates of the activation curve, $\mathbb{P}(G(z^{(f)}) \geq g_{\text{act}})$, in the basic model (3.3) for $z^{(f)} = 1000$ and $z^{(f)} = 2000$, as well as for the self background ($z^{(f)} = 0$), on logarithmic scale. The probabilities were estimated independently with simple sampling (SS), importance sampling (IS), and exact asymptotics based on large deviation theory (LDT) as used in [232]. For IS, 19 values of g_{act} were considered (from 100 to 1000 in steps of 50), and $N = 10000$ samples were generated for each value (i.e., $1.9 * 10^5$ samples altogether), whereas for the SS simulation, $N = 1.3 * 10^8$ samples were used over the entire range. The SS curves end at $g_{\text{act}} = 400$ and $g_{\text{act}} = 800$, respectively, because larger values were not hit in the given sample. The IS and SS graphs agree perfectly until the SS simulation lacks precision. For larger threshold values, we see a perfect agreement of the IS and LDT graphs. Note the general feature that, for threshold values that are not too small, the activation probability in the presence of foreign antigens is several orders of magnitude larger than the self background, i.e. Eq. (3.6) is satisfied.

Comparison with exact asymptotics

A pillar of the previous analysis of Zint et al. [232] (and its precursor BRB [205]) has been so-called exact asymptotics. This is a refinement of large deviation theory which yields estimates for the probabilities $P_n(A)$ themselves, rather than just their exponential decay rates obtained via the LDP in Def. 4.1. With standard large deviation theory (and our simulation method), it shares the tilting parameter which is calculated according to Eq. (5.10); for more details, we refer to [232]. A comparison of IS simulation with exact asymptotics is also included in Fig. 5.2. For small values of g_{act} , exact asymptotics is slightly imprecise. This is due to the asymptotic nature ($n \rightarrow \infty$) of the method, which yields more precise results in the very tail of the distribution, where the deviations are truly “large”. Note that, although our tilting factors agree with those in exact asymptotics, rare event simulation does not suffer from this accuracy problem since, due to the reweighting, it is always a valid importance sampling scheme that yields unbiased estimates for every finite n ; the finite-size effects will only manifest themselves as a

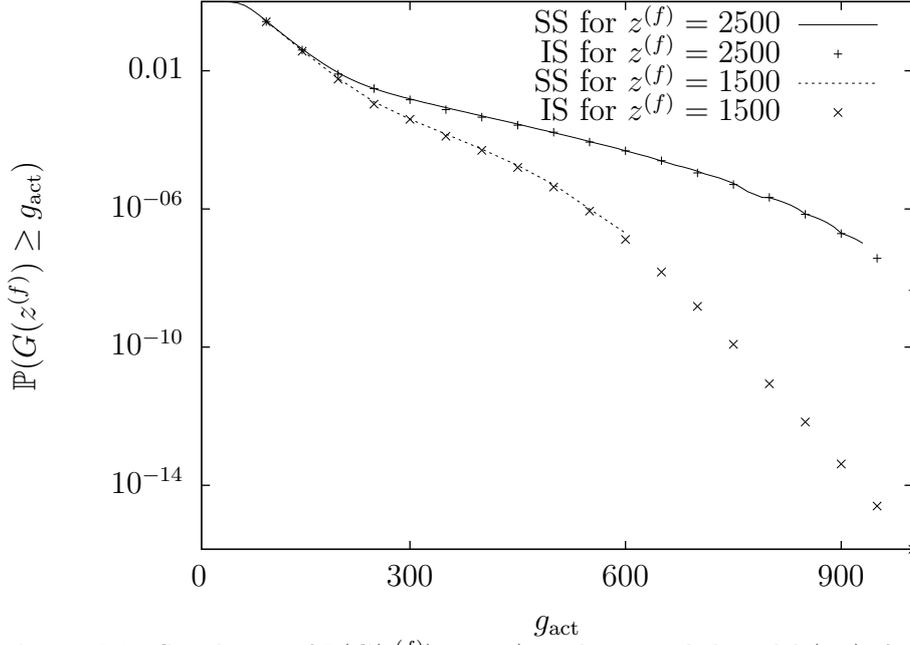


Figure 5.3: Simulation of $\mathbb{P}(G(z^{(f)}) \geq g_{\text{act}})$ in the extended model (3.4), for $z^{(f)} = 1500$ and $z^{(f)} = 2500$. The probabilities were estimated independently with simple sampling, and with importance sampling at 19 different threshold values (from 100 to 1000 in steps of 50). For IS, 9.5×10^4 samples were generated (5000 per threshold); for SS, 2×10^7 samples were used. No estimates are obtained with SS at thresholds beyond 600 or 920, respectively, in analogy with the situation in Fig. 5.2.

certain loss of efficiency, as will be seen below.

Asymptotic efficiency and relative error

In order to investigate the relative error of $(\widehat{P}_{P_n^{\vartheta^*}}(A))_N$, we first note that the variance of the estimator is given by

$$\mathbb{V}\left(\widehat{(P_{P_n^{\vartheta^*}}(A))}_N\right) = \frac{1}{N} \mathbb{V}\left(\widehat{(P_{P_n^{\vartheta^*}}(A))}_1\right) = \frac{1}{N} \mathbb{E}\left[\left(\mathbb{1}\{(T_n^{\vartheta^*})^{(1)} \in A\} \frac{dP}{dP_n^{\vartheta^*}}((T_n^{\vartheta^*})^{(1)}) - P_n(A)\right)^2\right], \quad (5.14)$$

where we have used (4.9) for $N = 1$. $\mathbb{V}\left(\widehat{(P_{P_n^{\vartheta^*}}(A))}_1\right)$ can be estimated via the given number N of samples in a single simulation run, i.e., as the sample variance

$$\widehat{\mathbb{V}}\left(\widehat{(P_{P_n^{\vartheta^*}}(A))}_1\right) = \frac{1}{N-1} \sum_{i=1}^N \left(\mathbb{1}\{(t_n^{\vartheta^*})^{(i)} \in A\} \frac{dP}{dP_n^{\vartheta^*}}((t_n^{\vartheta^*})^{(i)}) - \widehat{(P_{P_n^{\vartheta^*}}(A))}_N\right)^2, \quad (5.15)$$

where the $(t_n^{\vartheta^*})^{(i)}$ are now considered as realisations of $(T_n^{\vartheta^*})^{(1)}$. We can thus estimate the squared relative error as

$$\widehat{\eta}_N^2(P_n^{\vartheta^*}, A) = \frac{1}{N} \frac{\widehat{\mathbb{V}}\left(\widehat{(P_{P_n^{\vartheta^*}}(A))}_1\right)}{\left(\widehat{(P_{P_n^{\vartheta^*}}(A))}_N\right)^2}. \quad (5.16)$$

For simple sampling, one proceeds in the obvious analogous way (without tilting and reweighting).

In line with the limit discussed in Sec. 5.1.3, we now considered $G_n(z^{(f)})$ for system sizes $n = n_i$, where $n_i = n_i^{(c)} + n_i^{(v)} + n_i^{(f)}$, $0 \leq i \leq 10$, and we choose $n_i^{(\alpha)} = im^{(\alpha)}$, $\alpha \in \{c, v, f\}$, for $1 \leq i \leq 10$, as well as $n_0^{(c)} = m^{(c)}/2$, $n_0^{(v)} = m^{(v)}/2$, and $n_0^{(f)} = m^{(f)}$ (i.e., we simply ‘multiply’ the system, except for $i = 0$, which corresponds to ‘half’ a system except for the foreign peptide, which cannot be split into two). We then simulate $\mathbb{P}(G_{n_i}(z^{(f)}) \geq g_{\text{act}}n_i/m)$ for two values of $z^{(f)}$ and a fixed value of g_{act} with our importance sampling method, as shown in Fig. 5.4.

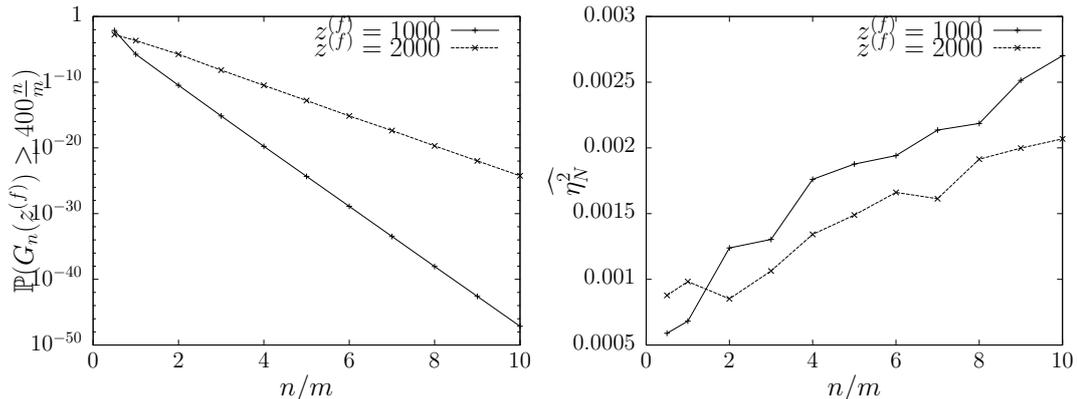


Figure 5.4: Importance sampling simulations for $\mathbb{P}(G_n(z^{(f)}) \geq g_{\text{act}}n/m)$ for $n = n_i$, $0 \leq i \leq 10$, for $g_{\text{act}} = 400$ and two values of $z^{(f)}$. Left: Estimate of the probability (note that the vertical axis is on logarithmic scale). Right: estimated squared RE.

Obviously, the (estimated) probabilities decay to zero at an exponential rate with increasing n , as they must by their LDP. In contrast, the (estimated) squared RE only increases linearly – this even beats the prediction of the theory (asymptotic efficiency only guarantees a subexponential increase).

So far, we have considered the n -dependence of the method for a fixed value of g_{act} , in the light of the available asymptotic theory. For the practical simulation of the given T-cell problem, we now take the given system size $n = m$ and numerically investigate the relative error as a function of g_{act} . Here, the exponential decay of $\mathbb{P}(G(z^{(f)}) \geq g_{\text{act}})$ as a function of g_{act} is decisive, which we have already observed in Fig. 5.2, and which goes together with the at-least-linear *increase* of I with g_{act} (recall that I is convex, and see Fig. 5.1). Fig. 5.5 shows the relative error of both SS and IS. It does not come as a surprise that, again, IS does extremely well and ‘beats’ the exponential decay of the probabilities: Whereas, on the log scale of the vertical axis, the squared RE of SS grows roughly linearly, it remains more or less constant for IS. (The very low squared RE of the simple sampling graphs for low thresholds in the right panel is due to the fact that the probability to reach this threshold is quite high and the huge sample of $N = 1.3 \cdot 10^8$ contributes to estimating it, that is, the sample sizes are not comparable. A simple sampling simulation run with the total sample size of a corresponding IS simulation (i.e., $N = 10000$ times the number of steps contained in the interval considered) results

in higher relative errors than for importance sampling even for the low threshold values (left panel). We would like to note, however, that the runtime of simple sampling for these small sample sizes is shorter than the runtime for IS, even if one does not count the overhead required to get the tilting parameters for importance sampling.)

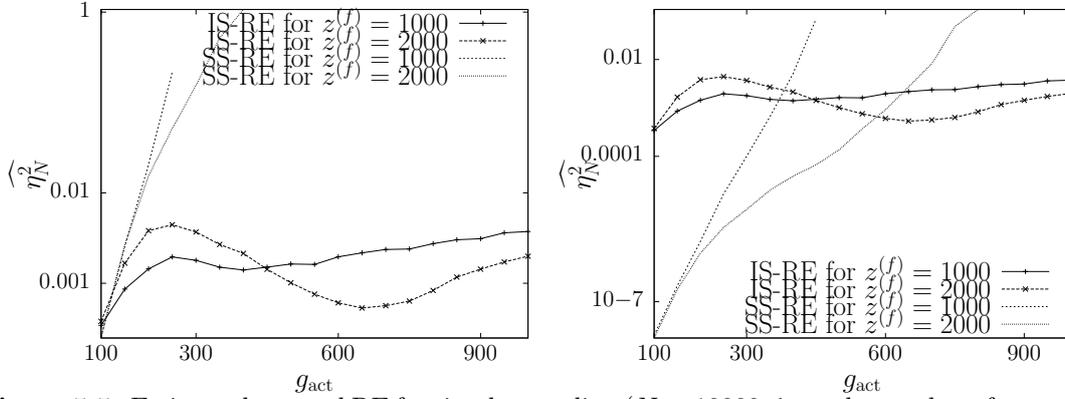


Figure 5.5: Estimated squared RE for simple sampling ($N = 10000$ times the number of steps contained in the considered interval (left), $N = 1.3 * 10^8$ (right)), and importance sampling ($N = 10000$ per threshold value in either panel) simulations of $\mathbb{P}(G(z^{(f)}) > g_{\text{act}})$ of the basic model, Eq. (3.3). Note that the vertical axis is on logarithmic scale.

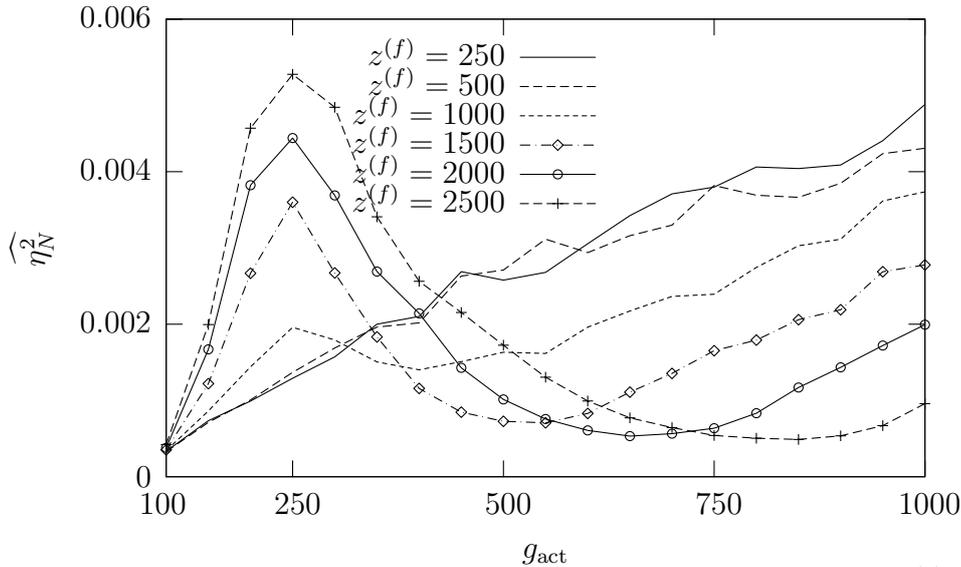


Figure 5.6: Estimated squared RE of our IS estimate, for various frequencies $z^{(f)}$ of the foreign antigen. Details are as in Fig. 5.5, but now the vertical axis is on linear scale.

Figure 5.6 sheds more light on the behaviour of the relative error of the IS simulation. It shows the squared RE for 6 distinct $z^{(f)}$ -values and reveals the finite-size effects. The wave-like behaviour for larger $z^{(f)}$ is due to the fact that, for very low threshold values, there is no real need for tilting, because the original distribution P_n is already close to optimal and the tilting factor is very small. For increasing thresholds, substantial

tilting is required, but there are still visible deviations from the $n \rightarrow \infty$ limit (as already discussed in the context of Fig. 5.2), so the tilted distributions are not optimal. This produces the “hump” in the squared RE curves, which is more pronounced for larger $z^{(f)}$ values because, for the case $n = m$ considered here, the foreign antigens come as a single term that may stand out. For large g_{act} , finally, one gets close enough to the limit, and the expected sub-exponential increase sets in (in our case, it is, in fact, roughly linear). Nevertheless, it should be clear that, in spite of the slight non-optimality at small threshold values, our tilted distributions still yield a far lower squared RE than does simple sampling.

5.2.2 Analysis of the T-cell model

In this Section, we use our simulation method to obtain more detailed insight into the phenomenon of statistical recognition in the T-cell model. As discussed before, the task is to discriminate one foreign antigen type against a “noisy” background of a large number of self antigens. We already know from Fig. 5.2 that, for threshold values that are not too small, the activation probability in the presence of foreign antigens is several orders of magnitude larger than the activation probability of the self-background, i.e. Eq. (3.6) is satisfied. As discussed in [232], this distinction relies on $z^{(f)} > z^{(c)}, z^{(v)}$ – basically, what happens is that larger copy numbers of the foreign antigen “thicken” the tail of the distribution of $G(z^{(f)})$ (without changing its mean), so that the threshold is more easily surpassed. The self-nonself distinction may, according to this model, be roughly described as follows. For a given antigen (foreign or self), finding a “highly-stimulating” T-cell receptor is a rare event; but if it occurs to a foreign antigen, it occurs manifold since there are numerous copies, which all contribute the same large signal, since all receptors of the T-cell involved are identical; the resulting stimulation rate is thus high. In contrast, if it is a self antigen that finds a highly-stimulating receptor, the effect is less pronounced due to the smaller copy numbers. Put this way, the toy model “explains” the distinction solely on the basis of copy numbers; but see the Discussion for more sophisticated effects that alleviate this requirement.

Following these intuitive arguments, we now aim at a more detailed picture of how the self background looks, and how the foreign type stands out against it. To investigate this, it is useful to consider the histograms of the total constitutive, variable, and foreign stimulation rates, i.e., the contributions of the “constitutive sum”, the “variable sum”, and the individual “foreign term” in the sum (3.3), either for all samples or for the subset of samples for which $G(z_f) \geq g_{\text{act}}$, for various g_{act} (normalising by the number of “successful” samples would result in an estimate of the conditional distribution). Since this requires a higher resolution (and thus larger sample size) than the calculation of the activation probabilities alone, such analysis would be practically impossible with simple sampling. With IS, we again generated 10000 samples per g_{act} value, from which between 30 and 70 percent turned out to reach the threshold.

Figure 5.7 shows the resulting histograms when all samples are included, and Figs. 5.8 and 5.9 show the histograms for the subset of samples that have surpassed four represen-

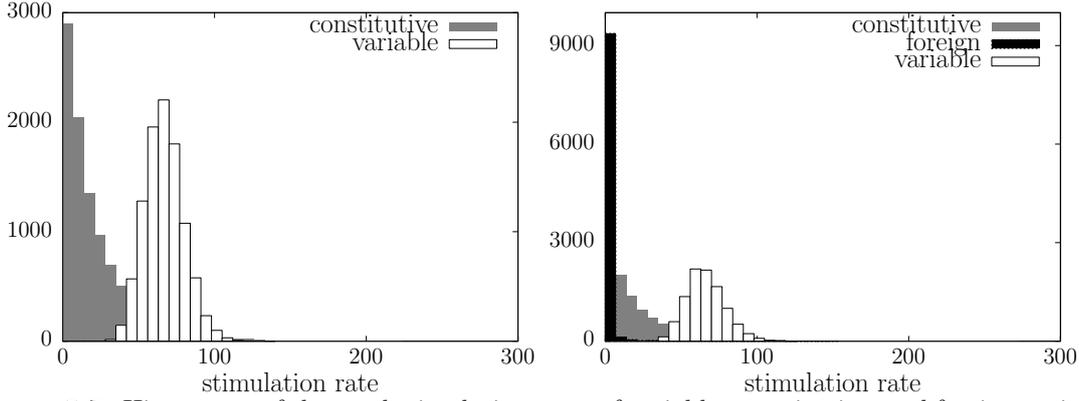


Figure 5.7: Histograms of the total stimulation rates of variable, constitutive, and foreign antigens, for $z^{(f)} = 0$ (left) and $z^{(f)} = 1000$ (right), in the basic model (3.3), when all samples are included. Sample size is 10000, and the vertical axis holds the number of samples whose total constitutive (variable, foreign) stimulation rates fall into given intervals. Note that the scaling of the vertical axis varies across diagrams.

rate \ g_{act}	0	100	250	500	1000
variable	66.6	74.9	77.1	78.8	80.0
constitutive	22.2	59.2	277.7	590.6	1160

rate \ g_{act}	0	100	250	500	1000
variable	12.7	13.9	14.5	14.9	15.1
constitutive	23.1	35.6	88.8	134.9	191.3

Table 5.1: Sample means (up) and sample standard deviations (below) of the histograms in Fig. 5.7 (left) and Fig. 5.8 (i.e., the self-only case).

tative threshold values, without and with foreign antigen. Tables 5.1 and 5.2 summarise these results in terms of means and standard deviations. Finally, Fig. 5.10 shows the joint empirical distribution for all pairs of variable, constitutive, and foreign stimulation rates, again for various threshold values.

Let us start with the situation without foreign antigens, as displayed in Figs. 5.7 (left) and 5.8 as well as Table 5.1. This already illustrates the fundamental difference between variable and constitutive antigens. Judging from the large number ($m^{(v)} = 1500$) of individual terms in the sum at low copy number ($z^{(v)} = 50$), the variable stimulation rate is expected to be approximately normally distributed and fairly closely peaked around its mean – at least as long as no restriction on $G(z^{(f)})$ is involved – and, as the Figure shows, this feature persists when $G(z^{(f)}) > g_{act}$, practically independently of the threshold involved. So, the variable antigens form a kind of background that poses no difficulty to foreign-self distinction: It is not very noisy, and it does not change with the threshold.

In contrast, the distribution of the constitutive activation rates is wider; this is due to the large copy numbers ($z^{(c)} = 500$), the effect of which is not compensated by the smaller

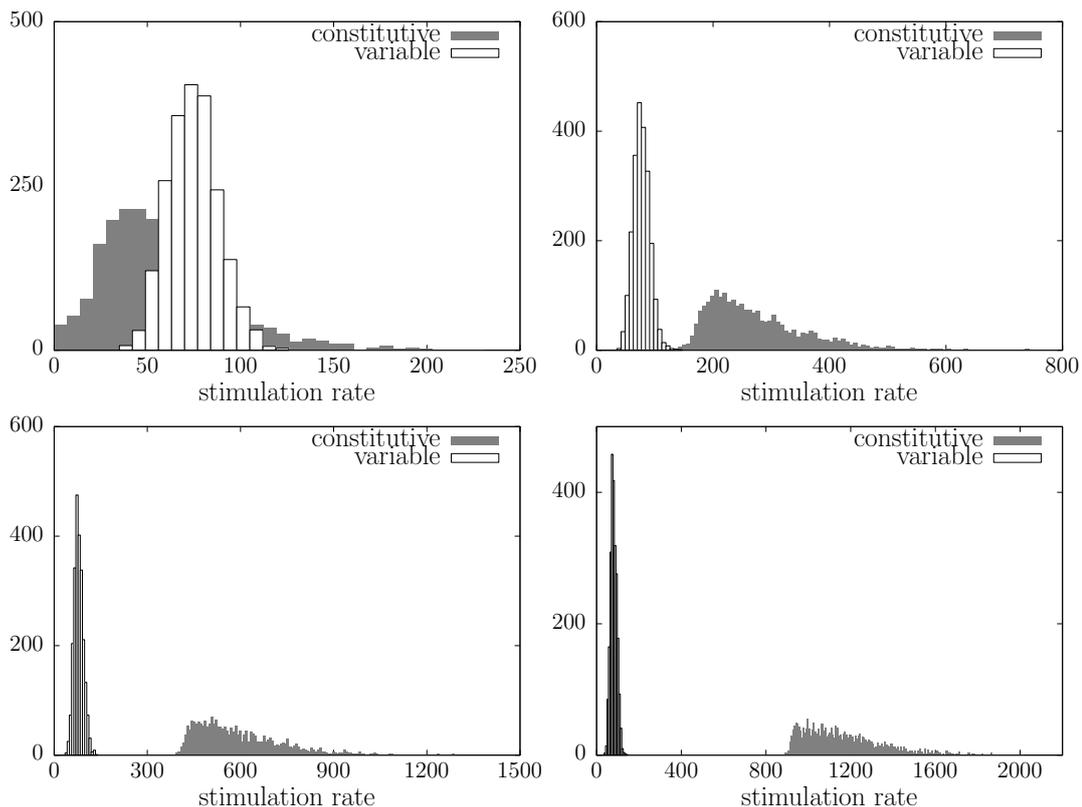


Figure 5.8: Histograms of the total stimulation rates of variable and constitutive antigens, for $z^{(f)} = 0$, in the basic model (3.3), for samples that reach a given threshold value ($g_{\text{act}} = 100$ (upper left), $g_{\text{act}} = 250$ (upper right), $g_{\text{act}} = 500$ (lower left), $g_{\text{act}} = 1000$ (lower right)). Sample size is 10000, and the vertical axis holds the number of samples that reach g_{act} and whose total constitutive (variable, foreign) stimulation rates falls into given intervals. Note that the scaling of both axes varies across diagrams.

number of terms, $m^{(c)} = 50$. Furthermore, the normal approximation is not expected to be particularly good for the constitutive antigens – given the extreme asymmetry of the W -distribution (see Fig. 3.2), the central limit theorem will not average out the deviations at only $m^{(c)} = 50$. In particular, the distribution remains asymmetric. With increasing threshold, this distribution moves to the right. The reason for this is that, in order to reach an increasing g_{act} , the “tail events” of the constitutive or the variable sum or both must be used, but it is “easier” (that is, more probable) to use the constitutive one because it contains more atypical events. In the language of large deviation theory, this is an example of the general principle that “large deviations are always done in the the least unlikely of all the unlikely ways” [88, Ch. I]. In the language of biology, the constitutive antigens are the “problem” of foreign-self distinction : Due to their high copy numbers and incomplete averaging, fluctuations persist that occasionally induce an immune response even in the absence of foreign antigens. This occurs if a T-cell receptor happens to fit particularly well to one, or a number of, constitutive antigen types on an APC; due to their large copy numbers, these few highly-stimulating types are then sufficient to surpass the threshold (in contrast, several highly-stimulating types would

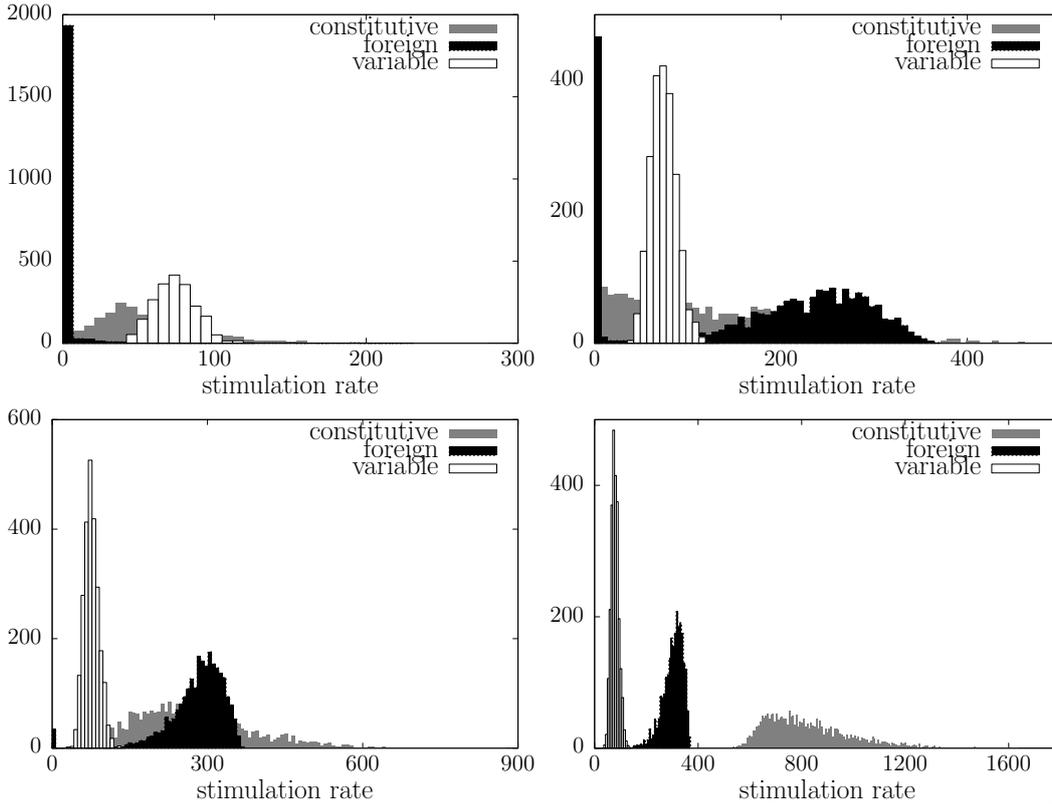


Figure 5.9: Histograms of the total constitutive, variable and foreign stimulation rates for $z^{(f)} = 1000$ in the basic model (3.3). Sample size is 10000, and the vertical axis holds the number of samples that reach the threshold g_{act} and whose total constitutive (variable, foreign) stimulation rate falls into a given interval, for $g_{\text{act}} = 100$ (upper left), $g_{\text{act}} = 250$ (upper right), $g_{\text{act}} = 500$ (lower left), $g_{\text{act}} = 1000$ (lower right). The maximal stimulation rate for the foreign antigens is $z^{(f)}w(1) = 367.9$. Note that the scaling of both axes varies across diagrams.

be required for the variable antigens to elicit a reaction, which is too improbable).

Let us now turn to the picture with foreign antigen present (Figs. 5.7 (right), 5.9, 5.10, and Table 5.2). One salient feature here is that the variable stimulation rate behaves exactly as in the “self-only” case: closely peaked around a small mean, unchanged when $G(z^{(f)}) > g_{\text{act}}$ is imposed. The picture is thus dominated by the interplay of constitutive and foreign types. In line with Fig. 5.2, the situation is similar in the case without restriction on $G(z^{(f)})$ (Fig. 5.7, right) and the case when $G(z^{(f)}) \geq 100$ (Fig. 5.9, upper left). In particular, the foreign stimulation rate is closely peaked at 0; only the constitutive background has moved slightly to the right, exactly as in the “self-only” case. For $g_{\text{act}} = 250$ (Fig. 5.9, upper right), where, according to Fig. 5.2, foreign-self distinction sets in, the foreign stimulation rate becomes prominent: The right branch of the W -distribution now becomes populated, and the associated stimulation rates are large due to the large copy numbers $z^{(f)}$ involved.

Still, for $g_{\text{act}} = 250$, the foreign stimulation rate is close to 0 in a sizeable fraction of the cases in which an immune reaction occurs – here, the reaction is brought about by the constitutive background, which moves to the right just as in the “self-only” case

rate \ g_{act}	0	100	250	500	1000
variable	65.9	74.1	74.2	76.2	78.4
constitutive	21.8	55.9	129.5	270.4	821.1
foreign	0.9	4.0	184.8	279.6	302.2

rate \ g_{act}	0	100	250	500	1000
variable	12.7	14.1	13.9	14.2	14.7
constitutive	22.4	42.0	90.4	109.1	163.7
foreign	6.7	18.5	112.2	54.5	39.2

Table 5.2: Sample means (top) and sample standard deviations (bottom) of the histograms in Fig. 5.7 (right) and Fig. 5.9 (i.e., the case with foreign antigens).

(but less pronounced). Fig. 5.10 shows that the constitutive and foreign stimulation rates are, indeed, negatively correlated: as is to be expected, low foreign rates are compensated by high constitutive rates and vice versa (in contrast, the variable background hardly correlates with either the constitutive or the foreign stimulation rate). As in the “self-only” case, therefore, the level of “unwanted” activation (“self-only” or “mainly self, without appreciable foreign activation”) is set by the tail behaviour of the constitutive background. However, if g_{act} is increased further (Fig. 5.9, lower left), every T-cell beyond the threshold displays high stimuli for the foreign antigen, their distribution shifting even further to the right and concentrating near the maximal stimulation rate given by the maximum of the function w of Eq. (3.1), more precisely, by $z^{(f)}w(1)$. This maximum can, of course, not change by imposing restrictions on $G(z^{(f)})$; thus, any further increase of g_{act} (Fig. 5.9, lower right) must then be matched by the by now familiar shift of the constitutive background. (This last panel is, however, less biologically realistic since the probabilities involved are too small to be relevant – after all, with about 10^7 different T-cell types, threshold values that yield activation probabilities far below 10^{-7} even in the presence of foreign antigens cannot be very healthy.)

A further illustration of the onset of self-nonsel distinction is presented in Fig. 5.11. Here we consider

$$\mathbb{P}(G(z^{(f)}) - z^{(f)}W_{n^{(c)}+n^{(v)}+1} > g_{\text{act}} \mid G(z^{(f)}) > g_{\text{act}}) = \frac{\mathbb{P}(G(z^{(f)}) - z^{(f)}W_{n^{(c)}+n^{(v)}+1} > g_{\text{act}})}{\mathbb{P}(G(z^{(f)}) > g_{\text{act}})}, \quad (5.17)$$

i.e., the probability that, in a T-cell that is activated in the presence of foreign antigen, the self component alone would have been sufficient for the activation. From $z^{(f)} = 1000$ onwards, this probability decreases to 0 quickly with increasing g_{act} . Put differently, in large parameter regions, the foreign antigens do indeed make the difference, which is the decisive feature of self-nonsel distinction.

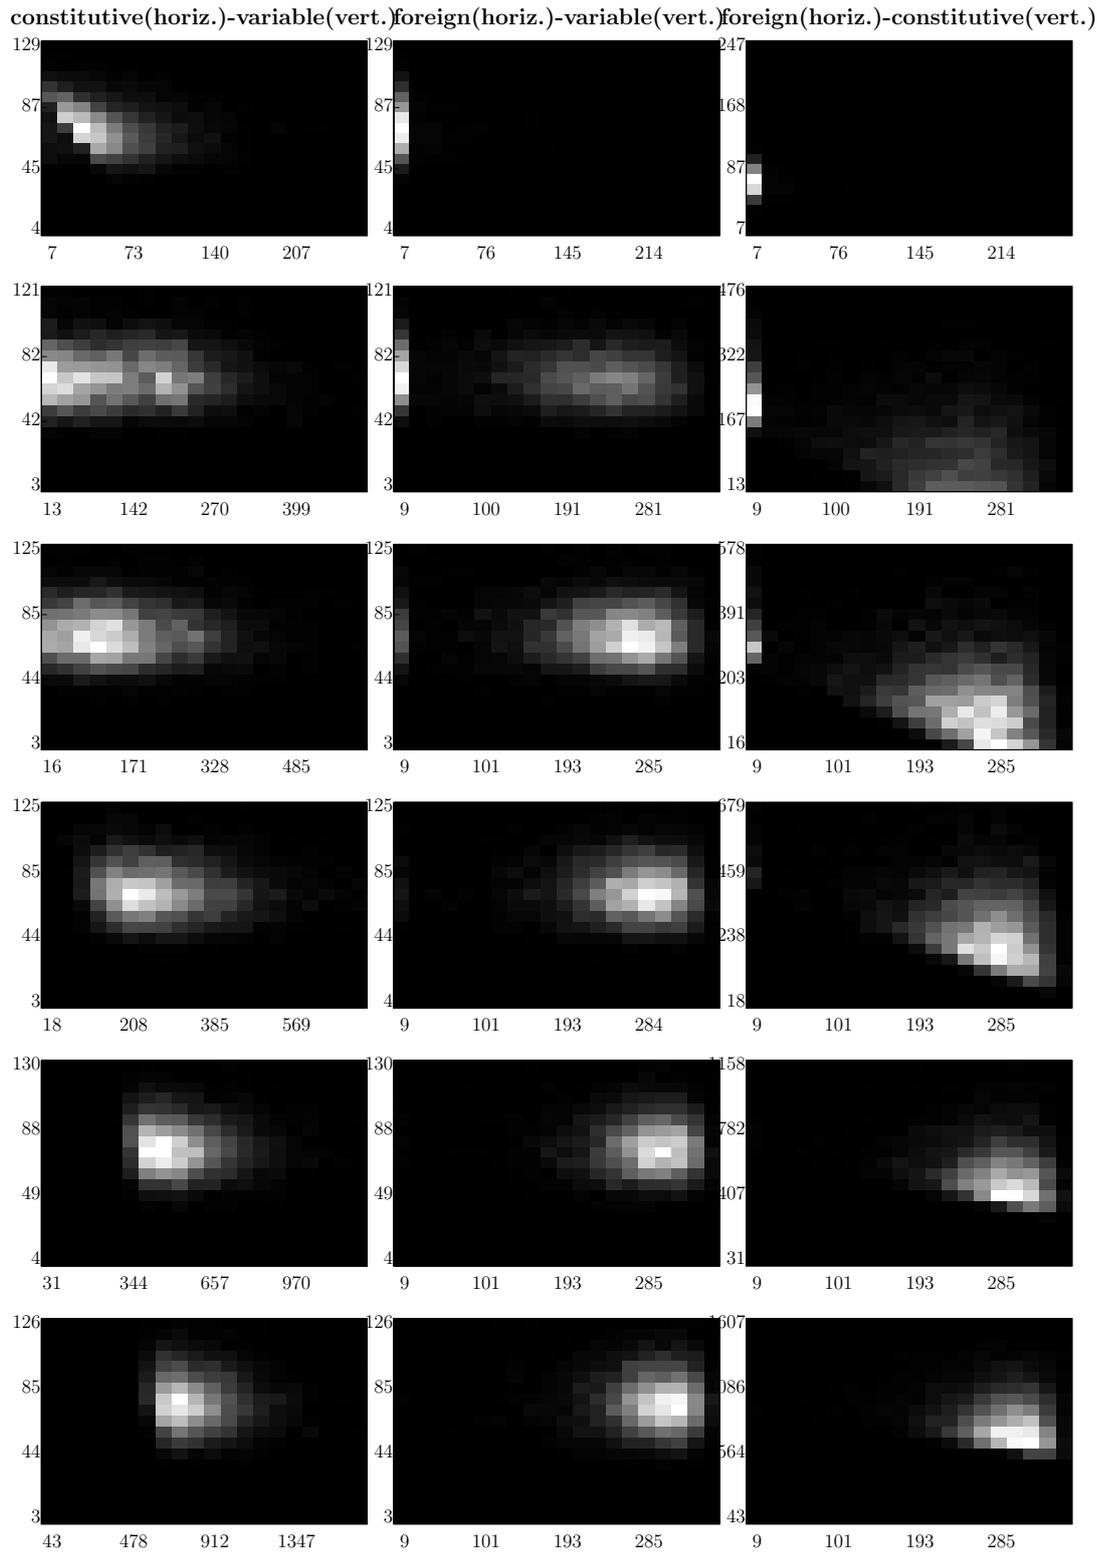


Figure 5.10: Pairwise joint empirical distributions of the total constitutive, variable, and foreign stimulation rates, for those samples with $G(z^{(f)}) > g_{act}$ in the basic model (3.3) (with $z^{(f)} = 1000$). Greyscales correspond to number of samples falling into 2D-intervals defined by total stimulation rates of pairs of antigen types. Rows (from top to bottom): $g_{act} = 100, 250, 350, 500, 750, 1000$; columns (from left to right): constitutive (horizontal) – variable (vertical); foreign (horizontal) – variable (vertical); foreign (horizontal) – constitutive (vertical). Lighter shading corresponds to higher frequencies.

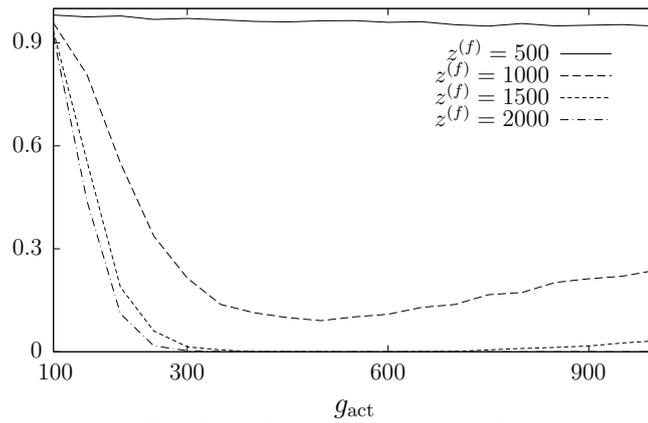


Figure 5.11: Fraction of samples whose self-component alone is above threshold, among those that reach the threshold in the presence of $z^{(f)}$ foreign molecules, for various $z^{(f)}$ (i.e., IS simulation of the probability in Eq. (5.17)). Sample size is 10000 for each g_{act} value considered.

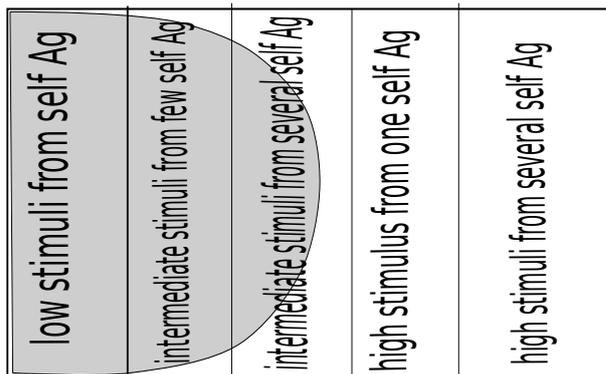


Figure 5.12: Different categories of T-cells in dependence to their reaction to self antigens. The grey areas mark the part of the T-cell repertoire that would be the desirable outcome of negative selection.

5.3 NEGATIVE SELECTION

In this second section of the T-cell activation analysis chapter we motivate and develop extensions of the basic BRB model that incorporate negative selection. Again we use our importance sampling approach to analyse these extensions.

From the analysis of the basic BRB model as presented in the previous section and in [205, 232, 119] it follows that with this kind of model foreign-self discrimination is possible if enough foreign antigens are presented and an adequate activation threshold is chosen. As described in the introductory chapter and for example in [150, 132], there are indications that dendritic cells have an antigen processing mechanism that is stimulated by the parallel triggering of an innate receptor. Consequently, antigenic material from pathogens is more efficiently processed and more pathogenic antigens are presented on the DC surface. This would clearly speak for the relevance of the basic BRB model. However, it is still unclear if similar mechanisms are in place for other APC such as B cells or macrophages. Furthermore, it is thinkable that the impact of this positive stimulation does not lead to an elevated presence of foreign antigens in comparison to self antigens on the DC surface but is a necessity to have at least some of them presented. This is the case if an (abundant) number of self antigens compete with the foreign antigens for MHC molecules. Moreover, experiments show that only very few cognate antigens are necessary to activate a T-cell [156]. The basic BRB model has therefore to be modified in order to better reflect the reality.

One crucial assumption in the basic BRB model is that from the point of view of every antigen the dissociation rate over the pool of all possible T-cell receptors follows an exponential distribution. All the more so, this exponential distribution is the same for all antigens. However, if we take a closer look on the T-cell repertoire we can intuitively split it into different sets, see therefore Figure 5.12.

There is the set of T-cells that react, if at all, only very weakly to self antigens. Other T-cells react very strongly to at least one self antigen and there are T-cells that react with various intermediate strengths. Intuitively speaking, a T-cell repertoire which consists only of T-cells from the first set and perhaps of parts of the last set, but definitely not

of the second set would be preferable if not necessary in comparison to the original, unconstrained T-cell repertoire. T-cells with a high probability of self activation are the cause of autoimmune diseases and a T-cell repertoire without such cells is beneficial. The question is how the peripheral T-cell repertoire is shaped such as to consist only of beneficial T-cells.

In the introductory part of the thesis we described two selection processes in the thymus. A T-cell undergoes positive selection by which it is tested if it is capable of binding to a pMHC complex at all. This guarantees that only suitable T-cells enter the periphery. In the basic BRB model it is assumed that this process is already completed. The second process, negative selection, deletes all T-cells that react too strongly to self antigens [184]. This process should, if working properly, provide for a peripheral T-cell repertoire as described above. Negative selection is not included in the basic BRB model as presented in the previous sections but in an extended version which is also introduced in [205, 232]. Additionally, there exists a mathematical analysis of a related T-cell activation model that uses a Bernoulli distribution instead of the W distribution for the description of the stimulation rates [204]. We revisit the extended BRB model and propose a second alternative extension. Both new models are then analysed.

5.3.1 BRB model with negative selection

In simple words negative selection in our modeling context works as follows. In the thymic medulla a T-cell meets an APC loaded with a random ensemble of self antigens. If the T-cell is stimulated too much by these antigens, the T-cell is removed from the pool of T-cells. T-cells reside for at least 4 to 5 days in the medulla. Therefore they encounter several APCs with random antigen loadings. Hence, all surviving T-cells should have a very low probability of self activation. This process is included in the BRB model and thereby changes the stimulation rate distribution of the self antigens.

In the BRB model, it is assumed that negative selection acts only upon constitutive antigens. Originally, this was motivated by the idea that variable antigens, which are tissue specific or only expressed under certain stress conditions, are not presented in the thymus. This was proven to be wrong (see e.g. [115]). Nevertheless, it is justifiable to concentrate on the constitutive antigens, as the stimulation rates of the variable antigens play no major role in the basic BRB model (see our discussion on p. 58). The probability of a T-cell to survive one round of negative selection in this model can be calculated as

$$\mathbb{P}(\text{survival of a T-cell}) = \mathbb{P}\left(\sum_{j=1}^{m^{(c)}} z^{(c)} W_j + \sum_{j=m^{(c)}+1}^{m^{(c)}+m^{(v)}} z^{(v)} W_j \leq g_{\text{thy}}\right), \quad (5.18)$$

where g_{thy} is the thymic activation threshold [143]. This is the maximal stimulation rate allowed to be induced by a random collection of self antigens and thus poses a constraint to the stimulation distribution by self antigens. In accordance to Zint et al. g_{thy} is set to 140 [232]. This is equivalent to the assumption that the probability of a T-cell to survive one meeting with an APC is about 95%. In the course of negative selection a

T-cell encounters several APCs and is thereby challenged with different sets of antigens. Henceforth we term these different meetings as rounds of negative selection and we consider k rounds of negative selection for different values of k . The other variables are the same as in the basic BRB model, that is $m^{(c)} = 50$, $m^{(v)} = 1500$, $z^{(c)} = 500$, $z^{(v)} = 50$.

The question at this point is how negative selection can be modelled in the context of the basic BRB model. In the following we present two different approaches.

1st variant of negative selection for the BRB model

In the proposed extension of van den Berg et al. and Zint et al. the number of different antigen types or more exactly the number of classes of antigen types that induce the same stimulus to a given T-cell is confined to $m^{(c)}$ [205, 232]. That is, a class of antigen is defined as such every T-cell that encounters a member of the antigen class is stimulated by this member in the same way as all the other members would do. Thereby we effectively reduce the size of the set of possible self antigens from nearly infinite to only $m^{(c)}$.

Although the idea of introducing such classes is biologically plausible if we remember that TCRs 'see' antigens as small amino acid strings and there are amino acids that are relatively similar in their composition (atomic composition, charging ...), in our case the number of classes $m^{(c)}$ is very small. It is a rather extreme case. Here, we put a very strong constraint on the constitutive antigen repertoire but we consequently do not have to deal with the setback of the need to meet to many different antigens. There is only one possible mixture. Several rounds k of negative selection are then only necessary to search through the space of variable antigen mixtures.

If we recall the definition of the BRB model (see Section 3), the stimulation rate distributions are the basic ingredients in describing the meeting of a T-cell with a random APC in the periphery. In order to incorporate negative selection into the framework of the BRB model, T-cell deletion has to be reflected in the stimulation rate distributions of the constitutive antigens. In line with the basic T-cell model it is impossible to define a stimulation rate distribution for every self antigen, due to the vast amount of possible self antigens. Therefore, van den Berg et al. and Zint et al. chose not to look at single stimulation rate distributions at all but define negative selection in a way to work on the distribution of the total sum of the constitutive stimulation rates [205, 232]. Mathematically speaking, they defined the random variables $C := \sum_{j=1}^{m^{(c)}} z^{(c)} W_j$ (random variable representing the total stimulation rate induced by all constitutive antigens on the APC) and $V := \sum_{j=1}^{m^{(v)}} z^{(v)} W_j$ (random variable representing the total stimulation rate induced by all variable antigens on the APC) with distribution functions $F(c)$ and $H(v)$. Negative selection only acts on the constitutive stimulation rate, such that C is transformed to a new random variable C^{neg} . This random variable describes the total constitutive stimulation rate after negative selection. Its distribution function can be calculated following Bayes theorem, with the events $A = \{C \leq c\}$ and $B = \{C + V \leq$

$g_{\text{thy}}\}$ as

$$F^{\text{neg}}(c) = \mathbb{P}(A|B) = \frac{\mathbb{P}(A \cup B)}{\mathbb{P}(B)} = \frac{\int \mathbb{1}_{\{c' \leq c\}} H(g_{\text{thy}} - c') dF(c')}{\mathbb{P}(B)} \quad (5.19)$$

Building upon these considerations we can reformulate the total stimulation rate equation from the basic BRB model (eq. (3.3)).

$$G^{\text{neg}}(z^{(f)}) := C^{\text{neg}} + \left(\sum_{j=1}^{m^{(v)}} qz^{(v)} W_j \right) + z^{(f)} W_{m^{(v)}+1}, \quad (5.20)$$

This means that a T-cell that survives negative selection and is released into the periphery there meets APCs which are equipped with antigens, again. We assume that every APC is equipped with antigens from all constitutive antigen classes and in total these can only induce a stimulus that follows the newly calculated stimulation rate distribution $F^{\text{neg}}(c)$. Additionally the APC is also equipped with variable and foreign antigen types whose stimulation rates all follow the original W distribution.

From the biological point of view this perception of negative selection can be also interpreted in another way. Without the introduction of $m^{(c)}$ classes of constitutive antigens we can assume to have a very large finite set of constitutive antigens. Their stimulation rates are all identically distributed and the subset that is presented by an APC is generated by choosing them without replacement. The resulting total stimulation rate always follows the same distribution as our random variable C . So, if we let negative selection act upon C , this can also be interpreted as if a given T-cell meets all mixtures of antigens, that induce the same stimulus as the current realisation of C .

To illuminate this important point it is helpful to think of an example. With g_{thy} we introduce a selection threshold and with equation (5.19) it is clear that the total stimulation rate after negative selection never surpasses this threshold as we only consider events where the stimulation rate stays below g_{thy} . From the perspective of a surviving T-cell this implies that it has seen all possible antigen mixtures almost surely. It is safe to say that for this T-cell there is no possibility to stimulate it with constitutive antigens more than g_{thy} . This assertion is true for every possible T-cell of the surviving T-cell repertoire. Such a negative selection is of course impossible because of the time and space constrains in the thymus if we consider the numbers in section 2.2.4. To introduce several rounds k of negative selection is more or less useless as already implicitly a nearly infinite k is assumed. Nevertheless, from a mathematical point of view it is possible and biologically this would imply an even more thorough search not only through the space of constitutive antigen mixtures but also through the additional variable antigen space.

2nd variant of negative selection for the BRB model

In this first approach to negative selection we comprised all individual constitutive stimulation rate random variables in one total stimulation rate random variable. There is of course also the second approach to negative selection, which is the opposite of this first approach. Instead of working on the total constitutive sum, negative selection is

here seen from the perspective of an individual stimulation rate W with distribution function $F(w)$ (by abuse of notation we again use F as the symbol for the distribution function, but mark the difference to the previous F by the use of another argument for the function).

We assume that for every possible antigen the stimulation rate it induces to all possible T-cells follows the W distribution. Here, we sort of mark one antigen and observe its meetings with different T-cells together with different other antigen types on different APCs. If we repeat that for infinitely many meetings and keep track of the stimulation rate this antigen induces to a T-cell during such a meeting and if the T-cell survives the meeting with the APC we can calculate a new W^{neg} distribution. The introduction of more rounds of negative selection just means that we assume that a given T-cell meets k APCs which are equipped with our particular antigen and induces always the same stimulus to the T-cell. Only the other antigen types and thereby the total stimulus changes. Of course the probability that this T-cell survives several meetings is much lower. Hence, the W^{neg} distribution should change even more. We assume that this happens independently for every self antigen in the thymus. Therefore, we can transfer the result for one particular self antigen to all other self antigens.

If we define the random variable $R := \sum_{j=2}^{m^{(c)}} z^{(c)} W_j + \sum_{j=m^{(c)+1}^{m^{(c)}+m^{(v)}} z^{(v)} W_j$ with distribution function $H(r)$ we can again use Bayes theorem to calculate the distribution of the stimulation rate induced by one constitutive antigen to a random T-cell conditioned on the survival of the T-cell during negative selection.

$$F^{\text{neg}}(w) = \frac{\int \mathbb{1}_{\{w' \leq w\}} H(g_{\text{thy}} - z^{(c)} w') dF(w')}{\mathbb{P}(B)} \quad (5.21)$$

with $B := \{z^{(c)} W + R \leq g_{\text{thy}}\}$

Implicitly this leads to a modification of the mean binding time \mathcal{T} between a random TCR and a random constitutive antigen as $W = w(\mathcal{T})$. This is the third way to think of negative selection. For constitutive antigens the mean binding time to a randomly chosen TCR of the mature T-cell repertoire is not any longer exponentially distributed. If we define $F(t)$ as the distribution function of \mathcal{T} , the distribution of \mathcal{T}^{neg} can be calculated by a minor modification of eq. (5.21)

$$F^{\text{neg}}(t) = \frac{\int \mathbb{1}_{\{t' \leq t\}} H(g_{\text{thy}} - z^{(c)} w(t')) dF(t')}{\mathbb{P}(B)} \quad (5.22)$$

Either $F^{\text{neg}}(w)$ or $F^{\text{neg}}(t)$ are used as the general distribution for either the individual constitutive stimulation rates or the individual constitutive binding times. That is, we assume that the independence condition of the stimulation rates also persists after negative selection. Every constitutive antigen has the same modified stimulation rate distribution.

These calculations enable us again to reformulate the basic total stimulation rate equation:

$$G^{\text{neg}}(z^{(f)}) := \left(\sum_{j=1}^{m^{(c)}} q z^{(c)} W_j^{\text{neg}} \right) + \left(\sum_{j=m^{(c)+1}^{m^{(c)}+m^{(v)}} q z^{(v)} W_j \right) + z^{(f)} W_{m^{(c)}+m^{(v)}+1}, \quad (5.23)$$

with $W_j^{\text{neg}} = w(\mathcal{T}^{\text{neg}})$

For a better discrimination we denote the first approach to negative selection as comprised in equation (5.20) by case 1 and the second approach (eq. (5.23)) by case 2. The calculations for the underlying distribution functions can be readily modified if we want a T-cell to undergo k rounds of negative selection instead of just one.

In both cases, for a good foreign-self discrimination

$$1 \gg \mathbb{P}(G^{\text{neg}}(z^{(f)}) \geq g_{\text{act}}) \gg \mathbb{P}(G^{\text{neg}}(0)) \quad (5.24)$$

has to hold for biologically relevant $z^{(f)}$. As discussed previously, the activation probabilities have to be very low, because otherwise a T-cell would be in a constant attack mode, not only against foreign intruders but also against the own body.

5.3.2 Simulation method

This section describes the development of a suitable simulation method, which makes use of the previously developed importance sampling method.

In the extended BRB model two main problems arise with regard to an efficient simulation of this model: The estimation of the stimulation rate distributions depending on the survival of negative selection and the efficient simulation of eqs. (5.20) or (5.23). In case 1 a straightforward naive approach would be a simulation with a huge number of trials, where in every trial the constitutive and variable stimulation rates are generated and summed up. If this total sum does not exceed g_{thy} , the constitutive stimulation rates are kept and new variable and foreign stimulation rates are generated and it is evaluated if the sum of these stimulation rates exceeds the activation threshold g_{act} . As mentioned previously in the analysis of the basic BRB model, the second step of this simulation is very inefficient because of the exponentially decreasing probabilities with increasing activation thresholds. The first simulation step where it is evaluated if the hypothetical T-cell survives leads to even more inefficiency as a percentage of trials is lost at that point because the T-cell might be deleted. This is even more true if more rounds of negative selection are included.

For case 2 the situation is similar. In fact, the straightforward naive simulation approach for only one round of negative selection is the same as for case 2. However, if a newly generated T-cell survives this step the situation in the periphery is different. Here, only one constitutive stimulation rate is kept and all other constitutive stimulation rates and the variable and foreign stimulation rates are newly generated. Introducing more rounds of negative selection also changes the negative selection process in comparison to case 1. Only one constitutive stimulation rate is kept through all hypothetical T-cell-APC meetings. All other stimulation rates (constitutive and variables) are newly generated for every round of negative selection. Obviously this whole simulation procedure would be very inefficient and thus a good estimation of activation probabilities would be prohibited. Therefore, we have to develop other simulation schemes.

Generally, in the context of importance sampling with exponential tilting two ways of efficient simulation are thinkable. We could put eqs. (5.19) and (5.20) together and

try to develop an importance sampling scheme with probably two tilting factors which have to be numerically determined, one for the generation of constitutive stimulation rates that do not exceed the thymic activation threshold, the other for the generation of variable and foreign stimulation rates in the second step. This would pose a rather involved problem especially as the second tilting factor would be dependent on the first. For case 2 the formulation of the problem in such a way would be even more complicated. We therefore chose an alternative approach in which we considered both problems individually.

In the case of eq. (5.19) we needed the distributions for the variable and constitutive sum before negative selection. Unfortunately, there exists no closed form expression for either of them. Therefore, we chose to estimate $F^{\text{neg}}(c)$ via simple sampling. This was possible since reaching the thymic activation threshold g_{thy} is not a rare event. The estimated distribution $\hat{F}^{\text{neg}}(c)$ is no longer a continuous but a discrete distribution (in our calculations represented by the according discrete density in form of a histogram). We evaluated it at 700 points from 0 to 140.

For the estimation of $F^{\text{neg}}(t)$ in eq. (5.22) we chose a different way. Here, we were in a slightly better position. We needed the original distribution $F(t)$ of the random variable \mathcal{T} , which is the exponential distribution by definition, over which we had to integrate. We also needed the distribution $H(r)$ of the random variable $R := \sum_{j=2}^{m^{(c)}} z^{(c)} W_j + \sum_{j=m^{(c)}+1}^{m^{(c)}+m^{(v)}} z^{(v)} W_j$. This was easy to estimate via simple sampling for the range from 0 to 140. The estimated distribution $\hat{H}(r)$ was then used for a numerical integration of (5.22).

The estimation of the stimulation rate distributions after negative selection enabled us to reuse our importance sampling method (see sec. 5.1.1 or [119]) for equations (5.20) or (5.23). In both formulas we have sums of independent but not identically distributed random variables. So they are part of the category of models for which our method is applicable.

To apply the framework of our importance sampling approach to the extended BRB model, we consider $n = m := 1 + m^{(v)} + m^{(f)}$ (case 1) or $n = m := m^{(c)} + m^{(v)} + m^{(f)}$ (case 2), where $m^{(f)} = 0$ or $m^{(f)} = 1$ depending on whether a foreign antigen is absent or present, and identify S_m with $G^{\text{neg}}(z^{(f)})/m$ and a with g_{act}/m . Tilting S_m with $m\vartheta$ then corresponds to tilting $G^{\text{neg}}(z^{(f)})$ with ϑ . This, in turn, is equivalent to tilting every summand in (5.20) or (5.23) with ϑ (since these summands are independent). Tilting and sampling from the distribution $F^{\text{neg}}(c)$ of the constitutive sum C^{neg} (case 1) is easy as it is a discrete distribution. In case 2 in accordance to the simulation method for the basic BRB model we do not directly generate random variables of the tilted W^{neg} but 'pull back' the tilting parameter to the underlying distribution of \mathcal{T}^{neg} . Thereby we avoid the numerical difficulties of W^{neg} and get a precise simulation method. To circumvent the speed limiting step of searching through the table of the discrete distributions we again apply the so-called alias method for discrete random number generation (see [119, 165, 113]).

Now with everything at hand we formulate the algorithm to simulate realisations of $G^{\text{neg}}(z^{(f)})$ and estimate $\mathbb{P}(G^{\text{neg}} \geq g_{\text{act}})$. (For notational convenience, we will not

distinguish between random variables and their realisations here). We will restrict us to case 2, the other case follows readily

Algorithm 2: Estimation of the activation probabilities for the the second variant of the extended BRB model

Input: activation threshold g_{act} , foreign antigen copy number $z^{(f)}$, number of samples $N \in \mathbb{N}$

Result: estimate of $\mathbb{P}(G^{\text{neg}}(z^{(f)}) \geq g_{\text{act}})$

- 1 calculate \mathcal{T}^{neg} (case 2)
 - 2 compute ϑ numerically such that (5.10) is satisfied; see [119] for the explicit procedure
 - 3 calculate the tilted densities $\tilde{f}_{qz^{(c)}}^{\vartheta, \text{neg}}$ and $\tilde{f}_{\alpha}^{\vartheta}$, $\alpha \in \{qz^{(v)}, z^{(f)}\}$, via (5.13)
 - 4 **for** $i = 1, \dots, N$ **do**
 - 5 **for** every summand j of (5.23) generate a sample $(\tilde{\mathcal{T}}_j^{\vartheta, \text{neg}})^{(i)}$ or $(\tilde{\mathcal{T}}_j^{\vartheta})^{(i)}$ according to its density $\tilde{f}_{qz^{(c)}}^{\vartheta, \text{neg}}$ or $\tilde{f}_{\alpha(j)}^{\vartheta}$ with the help of the alias method (here, the upper index (i) is added to reflect sample i , and $\alpha(j)$ is the weighting factor of the sum to which j belongs)
 - 6 calculate

$$(G^{\text{neg}}(z^{(f)}))^{(i)} = \left(\sum_{j=1}^{m^{(c)}} qz^{(c)} w((\tilde{\mathcal{T}}_j^{\vartheta, \text{neg}})^{(i)}) \right) + \left(\sum_{j=m^{(c)+1}^{m^{(c)}+m^{(v)}} qz^{(v)} w((\tilde{\mathcal{T}}_j^{\vartheta})^{(i)}) \right) + z^{(f)} w((\tilde{\mathcal{T}}_{m^{(c)}+m^{(v)}+1}^{\vartheta})^{(i)})$$
 - 7 calculate the indicator function times the reweighting factor (i.e., the i -th summand in Eq. (4.9))
 - 8 **if** $(G^{\text{neg}}(z^{(f)}))^{(i)} \geq g_{\text{act}}$ **then**
 - 9
$$R^{(i)} = \prod_{j=1}^m \frac{f_{\alpha(j)}((\tilde{\mathcal{T}}_j^{\vartheta})^{(i)})}{\tilde{f}_{\alpha(j)}^{\vartheta}((\tilde{\mathcal{T}}_j^{\vartheta})^{(i)})}$$
 - 10 **else**
 - 11 $R^{(i)} = 0$
 - 12 **end**
 - 13 **end**
 - 14 calculate $(\widehat{P}_{P_m}^{\vartheta}(A))_N = \frac{\sum_{i=1}^N R^{(i)}}{N}$, as estimate of $\mathbb{P}(G^{\text{neg}}(z^{(f)}) > g_{\text{act}})$
-

5.3.3 Results

We used our simulation method to analyse the extended BRB model and the results are shown and explained in this section. At first, we briefly show the impact of negative selection on the total constitutive stimulation distribution and the constitutive

binding-time distribution, respectively. Afterwards, we show and analyse the activation probabilities of both model variants keeping in mind eq. (5.24). Finally we explain how these activation probabilities and the differences between the basic model and our two variants of the extended model come about.

To get a first impression on the effect of negative selection, we present some results from the estimations of (5.19), (5.22) and their counterparts for $k = 100$ rounds of negative selection. During negative selection a certain amount of T-cells is killed. From section 2.2.4 we know these are about 35 – 70% of all T-cells surviving positive selection. If we analyse our negative selection calculation in order to assert that they reflect this percentage we come into troubles. First of all, case 1 has two biological interpretations. If we use the restricted repertoire idea then the estimation of the percentage of deleted T-cells is straightforward. With one round of negative selection 5% of all T-cells are deleted and with 100 rounds of negative selection about 20% vanish. However, if we adopt the other interpretation of case 1 there is no way to calculate such numbers. In the stimulation rate distribution all possible stimulations of mixtures of constitutive antigens are hidden. Trimming of this distribution is equivalent to deleting all T-cells that are stimulated so much by these different mixtures of constitutive antigens. A look at case 2 reveals similar problems. From the point of view of one antigen and 1 round of negative selection, again, 5% of all T-cells get killed. With 100 rounds of negative selection we have about 99% killed T-cells. But this is only half of the truth, as this calculation has to be done for all possible constitutive antigens. Even if there would be exact numbers for the amount of constitutive antigens this would not help, because the individual antigens are always shown in mixtures. Therefore, we would have to take into account that in the 5% killed T-cells of one antigen there are T-cells that would have also been killed by another antigen. All in all it suffices to say that we cannot draw a connection from our models to the actual number of deleted T-cells .

	$k = 1$	$k = 100$
experimental estimates	0.00043%	4.28%
case 1 restricted repertoire	5%	20%
case 1 whole repertoire	? (many more)	? (many more)
case 2 one antigen	5%	99%
case 2 all antigens	?	?

Table 5.3: Estimation of T-cell deletion by negative selection in the two model variants. The experimental estimates follow from the estimated number of DCs a T-cell meets in the thymus and the overall survival percentage (see sec. 2.2.4).

For both models the first step for the simulation is to calculate the distributions of the constitutive sum C^{neg} or the binding time \mathcal{T}^{neg} after negative selection. Figures 5.13 and 5.14 show the estimated density functions before and after negative selection. For case 1 the effect of one round of negative selection is striking but also consequential. The complete tail down to a stimulation rate of 140 is cut off, because of the conditioning on not surpassing g_{thy} . Values in the neighbourhood of g_{thy} are quite improbable because of

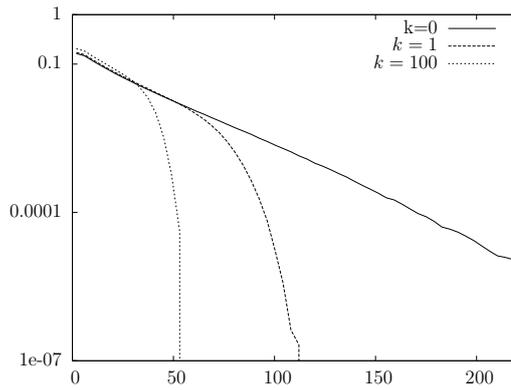


Figure 5.13: case 1: estimated densities for the stimulation rate of the total constitutive sum before negative selection ($k = 0$) and after $k = 1$ and $k = 100$ rounds of negative selection.

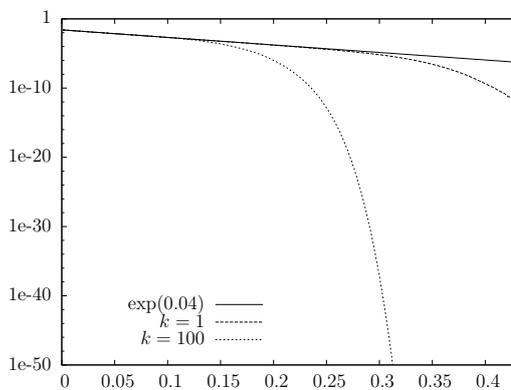


Figure 5.14: case 2: densities of \mathcal{T} before and after negative selection, for $k = 1$ and $k = 100$ rounds of negative selection.

the effect of the variable sum. C has not to surpass $g_{\text{thy}} - V$ and very low values of V are improbable. While the differences between no negative selection at all and introducing one round of negative selection is enormous, the introduction of more rounds of negative selection is not so eye-catching, but nevertheless recognisable. This is because already with one round of negative selection the harshest condition, not to reach g_{thy} has to be met. More rounds of negative selection only explore more combinations of a realisation of C with k different realisations of V . Consequently, the effect of negative selection in terms of cutting the tail of the distribution diminishes with a growing number of negative selection rounds.

The situation is different for case 2. One round of negative selection cuts only a small part of the tail of \mathcal{T} . This is a consequence of only looking at the role of one antigen in a random mixture. Although the overall probability of such a mixture to induce a stimulation rate g_{thy} is relatively high as Fig. 5.13 clearly shows, this is not true for a single antigen. High stimulation rates of one antigen are easily compensated by low stimulation rates of other antigens in the mixture. With the introduction of more negative selection rounds the situation changes. More and more combinations of

this individual antigen together with different mixtures of other antigens are tested. It becomes more probable that a high stimulation rate of the individual antigen is not compensated by the other antigens at least one time in the many negative selection rounds. This leads to a cut off of a considerable part of the tail of the exponentially distributed binding time.

To summarise the first impressions of the 2 negative selection models, we have to admit that they do not fit to the biologically realistic parameter range of T-cell deletion. The effect of different rounds of negative selection is quite different for the 2 models and it is interesting to see how this is reflected in the activation probabilities.

Activation probabilities

This section deals with the estimation and interpretation of the activation probabilities of the two extended BRB models. In line with section 5.2 and [119] we estimated the activation probabilities for the extended BRB model for $g_{\text{act}} \in \{100, 150, \dots, 1000\}$ and various values of $z^{(f)}$. We present our estimations and compare them to the results of the basic BRB model.

We start with the analysis of case 1. Figure 5.15 shows the activation probabilities for three different copy numbers of foreign antigen and only one round of negative selection estimated via simple sampling (SS) and our simulation method (IS). The simple

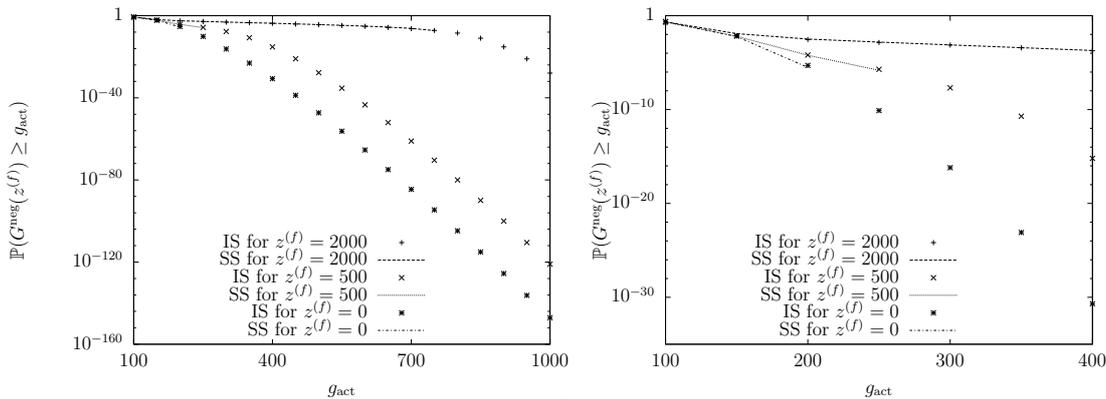


Figure 5.15: case 1: Activation curves for $z^{(f)} = 0, 500, 2000$. Estimated via simple sampling (SS) and our simulation method based on importance sampling (IS). The right figure is a clipping of the left one, only showing the probabilities for thresholds between 0 and 400.

sampling estimations serve as confirmation of our results. Obviously, we estimated only activation probabilities for low thresholds with SS, because of the exponentially decreasing probabilities. For this relevant regime they support our results from the importance sampling.

As expected, all three probability curves for three different values $z^{(f)}$ differ much with increasing thresholds. Therefore, the condition (5.24) is fulfilled, as was already shown for the basic BRB model. But if we compare the results from both models we can see a great change in the order of magnitude of activation probability (see Fig. 5.16). In

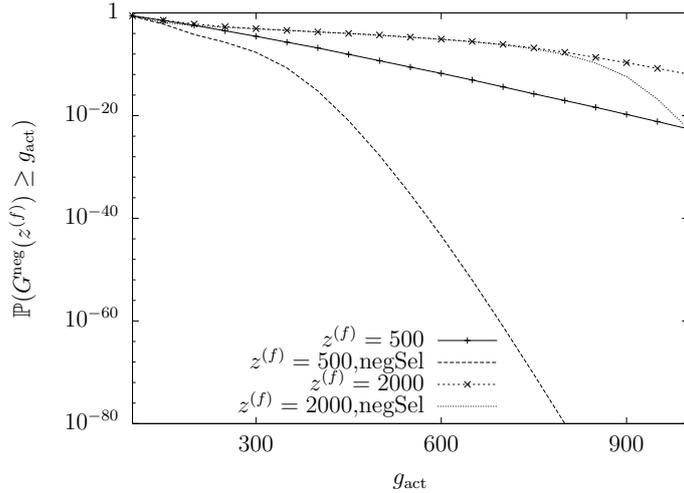


Figure 5.16: basic vs. case 1: A comparison of the activation probabilities for two different numbers of foreign antigen with and without negative selection.

fact, what we can learn from this Figure is that the effect of negative selection on the activation probability increases with decreasing copy number $z^{(f)}$. Consequently, with negative selection foreign-self discrimination is possible for a whole different range of parameters. In the basic model discrimination is hardly possible for $z^{(f)} = 500$, whereas in the extended model, we have a big difference in the order of magnitude of activation probability for $z^{(f)} = 500$ and $g_{\text{act}} = 250$. This shows that foreign-self discrimination is now possible for cases where foreign antigen copies are only as abundant as copies of constitutive antigens. This is biologically much more plausible. These results can be optimised even further if we assume more than one round of negative selection. In Fig. 5.17 we show the activation probabilities for $k = 1$ and $k = 100$ rounds of negative selection. The effect of allowing more than one negative selection round is again bigger

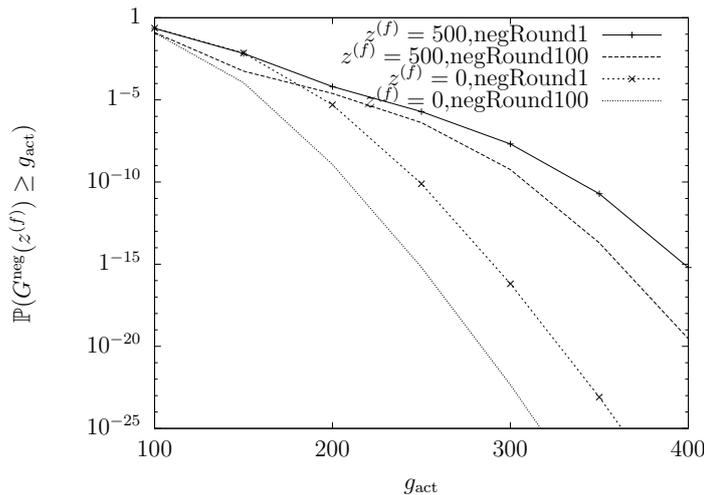


Figure 5.17: case 1: A comparison of the activation probabilities for two different numbers of foreign antigen with two different numbers of negative selection rounds ($k = 1$ and 100).

for low $z^{(f)}$. All in all compared to the non-selection case the differences between 1 or 100 are not so big, but still relevant. This had to be expected if we think of the effect of more rounds of negative selection on the stimulation rate distribution. Nevertheless, we can see that there is a good foreign-self discrimination for even lower threshold values. The same holds also if we decrease $z^{(f)}$, which we do not show here. It follows that, assuming case 1 is a reasonable explanation for negative selection, the BRB model is able to explain foreign-self activation for more biologically relevant parameters.

Next, we turn to negative selection as in case 2. In Figure 5.18 one can see the activation probabilities for three different values of $z^{(f)}$ and $k = 1$, whereas in Fig. 5.19 you can see the same for $k = 100$. Evidently, the effect of negative selection is very low for only one round compared to case 1.

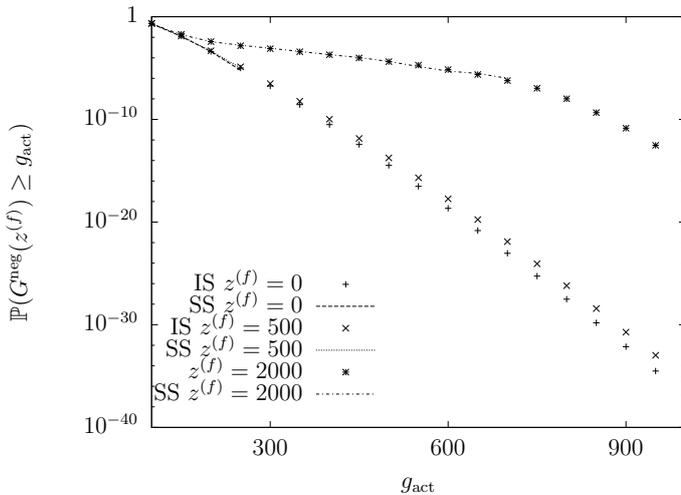


Figure 5.18: case 2: Activation curves for $z^{(f)} = 0, 500, 2000$ and one round of negative selection estimated via simple sampling (SS) and our simulation method (IS).

There is nearly no discrimination for $z^{(f)} = 500$ and negative selection has, therefore, no effect. This goes in line with the observations made at the beginning of this section, where we showed that there is hardly a change for the binding time densities of the constitutive antigens. The situation is extremely different for a high number of negative selection rounds. This can also be seen in the comparison with the basic model (Fig. 5.20). Foreign-self discrimination is easily possible, even for lower threshold values and a $z^{(f)}$ lower than 500. The results are comparable to case 1, as can be seen in Figure 5.21. The effect of negative selection is even slightly greater for case 2. Hence, also negative selection as in case 2 contributes significantly to a better foreign-self discrimination.

Stimulation rate histograms

As in section 5.2 and [119], we explain how this results come about by using stimulation rate histograms. These are histograms of the empirical total stimulation rates of

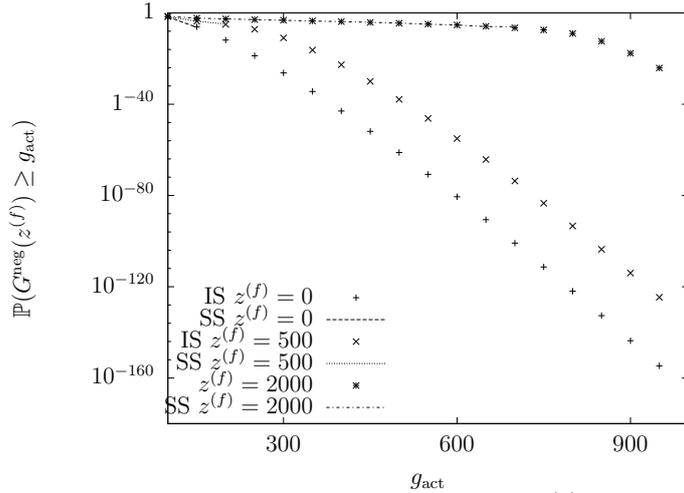


Figure 5.19: case 2: Activation curves for $z^{(f)} = 0, 500, 2000$ and 100 rounds of negative selection estimated via simple sampling (SS) and our simulation method (IS).

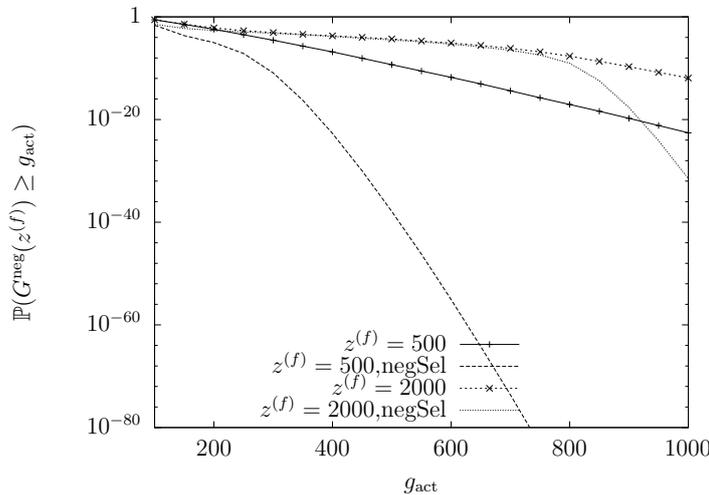


Figure 5.20: case 2 vs. basic: Comparison of activation curves with no negative selection and 100 rounds of negative selection.

the constitutive, variable or foreign antigens. Again, we start with case 1. For better comparability we choose in both cases $k = 100$. Figure 5.22 shows the histograms for 4 different activation thresholds. For the basic model we have shown that with increasing thresholds, more and more probability mass is moved to higher stimulation rates for the constitutive and foreign antigens. This movement is much faster for the foreign antigens. This is just a consequence of the fact, that it is far more likely to get one high activation by foreign, than to get several higher stimulation rates of the constitutive antigens, which would have the same effect as the foreign. We deal with rare events and they happen in the least unlikely of all unlikely ways.

The histograms here are quite different in comparison to the basic model. Again, much probability mass is moved to higher foreign stimulation rates, but there is hardly a change

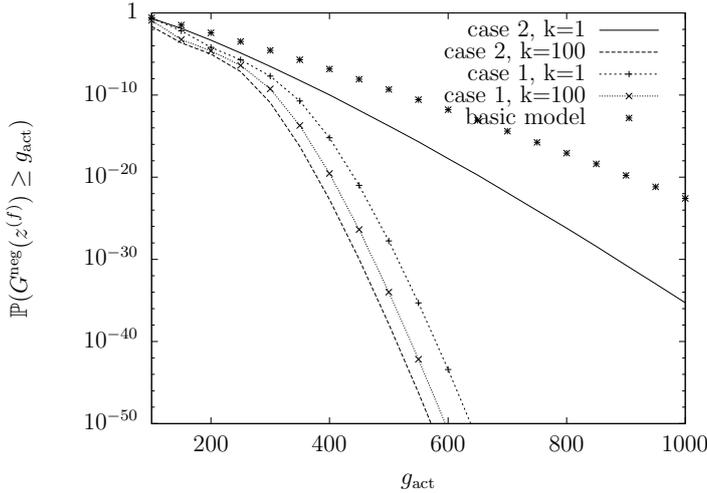


Figure 5.21: basic vs. case 1 vs. case 2: Comparison of the activation probabilities for $z^{(f)} = 500$ and model variant, that is, no negative selection, case 1 with $k = 1$ or $k = 100$, and case 2 with $k = 1$ and $k = 100$.

of the constitutive stimulation rate distribution. Instead, the variable stimulation rate distribution starts to move to higher stimulation rates. We know that this movement is very improbable, otherwise it would have also occurred in the basic model. Consequently, this is the explanation why the extended negative selection model allows for a far better foreign-self discrimination. The reason for the different behaviour of the constitutive stimulation rate follows from the process of negative selection. The probability to reach a higher stimulation rate than 140 is zero after negative selection in case 1 and with more and more rounds of negative selection the probability to reach values near 140 is also reduced, which could also be seen in the estimations of the stimulation rate distribution.

As case 2 follows a quite different interpretation of negative selection, it is interesting to see what happens here. Thus, in Figure 5.23 we show the stimulation rate histograms for the same activation thresholds for the second case. The general behaviour of the stimulation rates is similar to the ones for case 1, but there is also a big difference. The constitutive stimulation rates are lower. Nearly all probability mass is situated at the left border and this does not change much for increasing threshold values. The constitutive antigens are taken out of the game nearly completely. This is why case 2 works better than case 1 for 100 rounds of negative selection. But this also illustrates that the model is unrealistic: In biological reality it is not possible to take the constitutive antigens out completely.

The analysis of both extended BRB models reveals that negative selection can have a huge impact on the foreign-self discrimination ability of a mature T-cell repertoire. The activation probability and the number of foreign antigen copies $z^{(f)}$ can be significantly reduced without losing the ability of foreign-self discrimination. The activation rate histograms show that by negative selection the peripheral mature T-cell repertoire is changed in such a way that these T-cells have a higher probability of reacting to variable antigens than to constitutive ones. It is important to underline that this higher

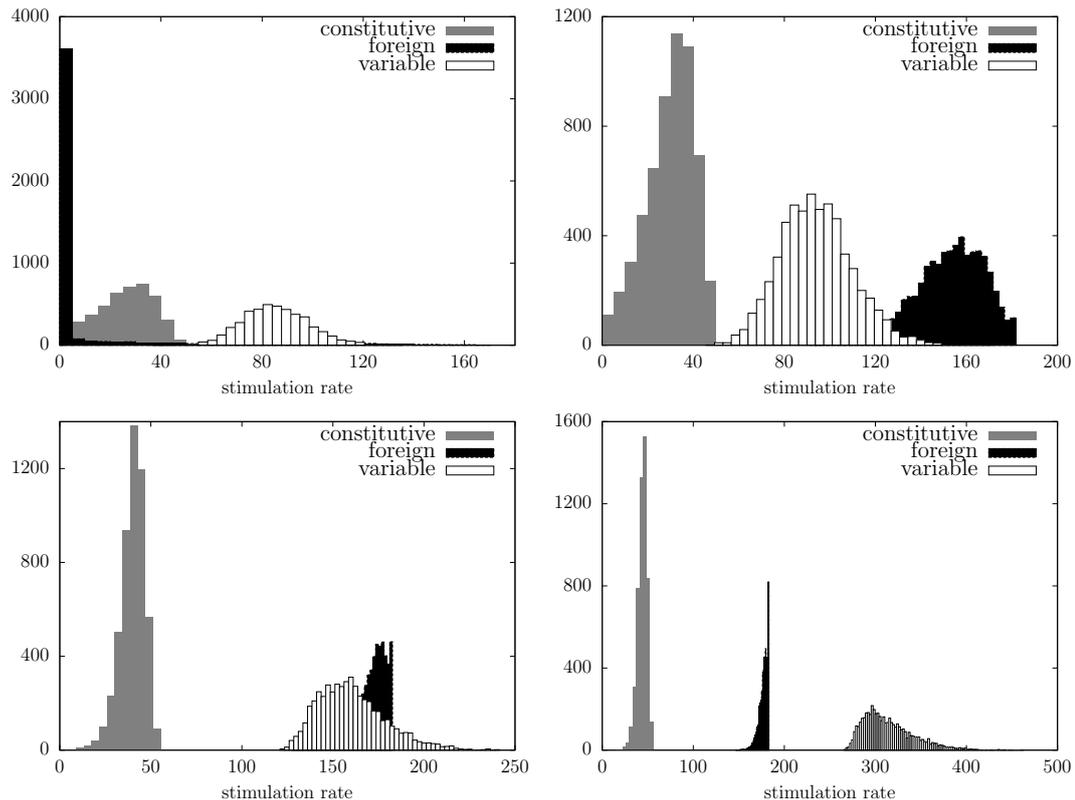


Figure 5.22: Case 1: Histograms of the total constitutive, variable and foreign activation rates for for $z^{(f)} = 500$ in the negative selection model (5.20). Sample size is 10000, and the vertical axis holds the number of samples that reach the threshold g_{act} and whose total constitutive (variable, foreign) activation rate falls into a given interval, for $g_{act} = 100$ (upper left), $g_{act} = 250$ (upper right), $g_{act} = 350$ (lower left), $g_{act} = 500$ (lower right). The maximal activation rate for the foreign antigens is $z^{(f)}w(1) = 183.95$. Note that the scaling of both axes varies across diagrams.

probability is only relative to the probabilities for constitutive stimulation rates but not to the foreign stimulation rate. Variable and foreign stimulation rates are the same as in the basic BRB model.

5.3.4 Discussion

The results of the analysis of the two extended BRB models are very promising. However, it is important to put them into the right context. The previous work in [205, 232, 119] and sec. 5.2 has already shown that with a probabilistic description of a stimulation mixture it is possible to describe foreign-self discrimination. Unfortunately it does not work for biologically more relevant parameters. Nevertheless, it is a good starting point for further investigations. Here, we concentrated on extensions of the basic model using negative selection. It was previously shown in [205] and [232] that this extension promises more biologically relevant results. As, until now, the mechanism of negative selection is a field of active research, we chose to adopt two different views on negative selection

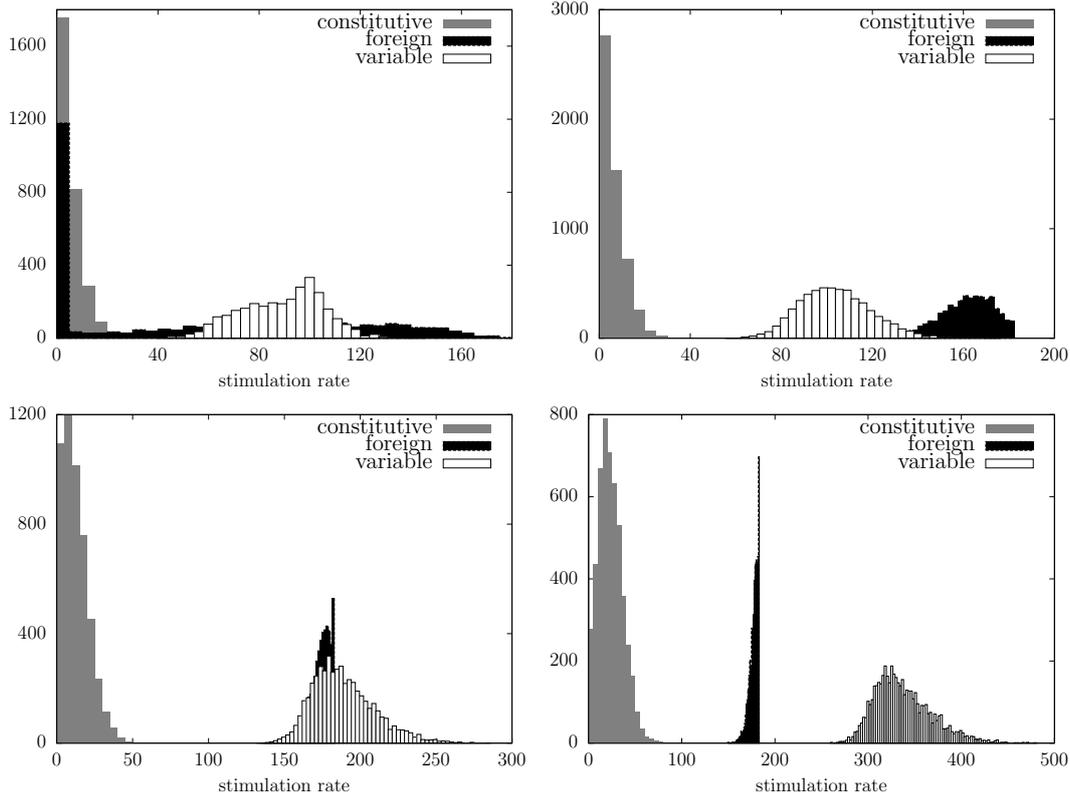


Figure 5.23: Case 2: Histograms of the total constitutive, variable and foreign activation rates for for $z^{(f)} = 500$ in the negative selection model (5.20). Sample size is 10000, and the vertical axis holds the number of samples that reach the threshold g_{act} and whose total constitutive (variable, foreign) activation rate falls into a given interval, for $g_{\text{act}} = 100$ (upper left), $g_{\text{act}} = 250$ (upper right), $g_{\text{act}} = 350$ (lower left), $g_{\text{act}} = 500$ (lower right). The maximal activation rate for the foreign antigens is $z^{(f)}w(1) = 183.95$. Note that the scaling of both axes varies across diagrams.

which are both simplifications, again.

We explained before that one interpretation of case 1 implies a pool of constitutive antigens which is very restricted to 50 classes with similar binding behaviours for all T-cells. This is of course very unlikely, especially if we think of the many different TCRs and antigen structures. In fact, under this assumption the very reason for the adaptive immune system to exist would be undermined. The T-cell repertoire loses all its flexibility and specificity. The second interpretation of case 1 is more promising. Negative selection works on the total unrestricted T-cell repertoire. However, there is also a big drawback. It is impossible to estimate how much of the pre-selection T-cell repertoire will be deleted and how many rounds of negative selection a T-cell needs to make sure it is not too autoreactive. Both parameters are important as they naturally restrict the scope of negative selection. Experimental and theoretical research shows that about 35% to 70% of the pre-selection T-cells survive negative selection. From recent experiments the mean binding time of a T-cell to a DC could be estimated to about 3min and the time of a T-cell in the medulla is restricted to about 5 days assume a mean scanning time of 3min, which is the mean binding time between a T-cell and a

DC during stage 1 in a lymph node (see Section 2.2.4). This leaves a T-cell with about 2400 DC meetings during negative selection. It is more than questionable if the extended BRB model as in case 1 would meet this conditions.

Case 2 on the other hand only looks at one individual constitutive antigen and its binding time and assumes that even after negative selection these binding times are independent and identical. This is, although more realistic than case 1, still unlikely. We estimated that for one antigen alone the survival probability of a T-cell is much lower than the biologically realistic one. This estimation is true for every constitutive antigen which reduces the survival probability again. This is true even if we consider that there are many antigen mixtures that can delete one individual T-cell and therefore there are dependencies between the killing rates of the individual antigens, that will reduce the overall killing rate. Also the second condition of 2400 DC meetings is hard to meet. We assumed 100 rounds of negative selection for one constitutive antigen in different mixtures. This means that a T-cell has to meet every individual constitutive antigen in 100 mixture. There will be of course much overlap between these mixture, but, given the huge number of possible self antigens 2400 DCs will not be sufficient to fulfill this assumption.

Thus, it is obvious that the two extended BRB models are oversimplifications of the biological reality that make the analyses of the model tractable. Nevertheless, there is much to learn on the effects negative selection can have on a T-cell repertoire. We have shown that both versions of the extended BRB model produce foreign-self discrimination for $z^{(f)} = 500$ and even lower. In that range this parameter becomes biologically plausible. This holds also for the activation probabilities themselves. We can find suitable activation thresholds, such that activation by foreign is not too improbable and on the other hand activation only by self is improbable enough (see eq. (5.24)). From the activation rate histograms we learn that through negative selection the impact of constitutive antigens on the stimulation of peripheral T-cells is minimised drastically. The formerly uninteresting variable antigens have to take their place which leads to the very good foreign-self discrimination capabilities. In section 2.2.3 we mentioned that the old assumption that this variable antigens are not present in the thymus is wrong. A new extension of our models to include negative selection to variable antigens would therefore be appropriate and could enhance foreign-self activation. However, a careful modeling would be needed, such that the discussed problems with these models are not facilitated.

At this point we do not see how with this class of models the problem of biologically unrealistic behaviour can be resolved and the estimated parameters for negative selection be met. Both, case 1 and case 2 are extreme cases as the restriction to several classes of antigens is not true as well as the assumption of total independence between all the antigens concerning their stimulation ability. As the idea of restricting the T-cell repertoire does not suffer from the drawback of killing too much T-cells and needing too much rounds of negative selection, it is appealing to use this idea. T-cells are crossreactive, especially to antigens with similar amino acid sequences. This could pose a possibility to introduce antigen classes. These will be considerably more than the few we assumed here and negative selection will therefore not work as great, but foreign-self discrimination should still be enhanced. It is also interesting to have a closer look on

the effect of different presentation probabilities for different antigens inside and outside the thymus. In [204] this problem was already investigated for a similar model but with stimulation rates that are Bernoulli distributed. The authors could show that the intuitive answer, that the ratio of presentation probabilities between the different antigens should be similar inside and outside the thymus is right. It would be interesting to see if this is true for our models and how robust foreign-self discrimination is, if there is a sudden increase of presentation probability of some antigens in the periphery for example due to cellular stress reactions.

A DISCRETE T-CELL ACTIVATION MODEL

As a consequence of our previous results with the extended BRB models we formulate a new model for the investigation of T-cell activation and negative selection. The objectives are to keep the basic ideas of the BRB model but try a different approach that helps us to assess or include other important parameters such as the T-cell survival rate and the number of possible APC meetings during negative selection.

The standard BRB model and its extensions work with two families of self antigen types. These are the constitutive antigens resulting from all regularly produced proteins/molecules (for example housekeeping proteins from the cells) and variable antigens from proteins that are produced only under special circumstances or in certain tissues. The difference between these two families in the model is reflected by a high copy number of constitutive antigens in contrast to a low copy number of variable antigens (see page 33). Furthermore it is assumed that there are many more variable than constitutive antigen types. By our previous analysis we could show that these variable antigens pose no threat to foreign-self discrimination. But, thereby we tempered with the biological reality. In Section 2.2.3 we introduced the tissue-restricted antigens. These are about 3000 antigens that are only expressed in certain tissues or during certain developmental stages (see Section 2.2.4). Hence, they are typical members of the variable antigen family. However, there is a flaw in these considerations. Taking the body as a whole this might be true, but in the special tissue they belong to, they are expressed in large numbers and therefore qualify to be termed constitutive. Ignoring these antigens in negative selection leads to severe autoimmune reactions and we know that via pGE these antigens are actually expressed randomly in the thymus [114]. Consequently, in our new model we should not make a distinction between antigens that are expressed in low or high copy numbers, as there are always situations in which antigens that are usually expressed rarely are then expressed excessively.

There have to be tolerance mechanisms in place to prevent T-cell reaction to any self antigen. It is still unclear how big the contributions of central tolerance and peripheral tolerance are. For example, Muller and Bonhoffer argue that from the possible amount of more than 10^6 self antigen types in a mouse only about $2 \cdot 10^5$ are presented in the thymic medulla whereas for all the other types mechanisms of peripheral tolerance are responsible [138].

Even if we take this lower number for granted we have a problem with the model of negative selection in the extended BRB model. One idea there was to introduce classes of antigens types, where every class comprises all antigen types that induce a similar stimulus to the TCRs (see page 65). Members of each of these 50 classes are presented

by each APC during and after negative selection. With the above considerations and numbers in mind this number of classes is clearly too small.

However, the idea of introducing antigen classes and thereby reducing the space of distinguishable antigen types is very tempting. It is estimated that every T-cell can bind to about 10^6 antigen types, where we have to keep in mind that this number refers to the theoretical total amount of possible antigens [67]. Hence, the number of actual antigens in a host to which a T-cell can bind should be much lower. Furthermore not all these antigens bind long enough to induce activation stimuli. If we assume that all the antigens activating one individual T-cell belong to one class of antigen types and these types induce similar stimuli to the other T-cells we are back at a model with a discrete set of antigen classes.

It is biologically reasonable to assume that there are 10 – 100 antigen types that can activate an individual T-cell (because they are very similar at the molecular level). Consequently, the amount of antigen classes can be restricted to 1000 – 10000 if every class has similar numbers of members and we assume that 10^5 self antigens are presented in the thymus. We have to ignore the other self antigens due to a lack of knowledge on peripheral tolerance. This poses no problem if we assume that either peripheral tolerance is really effective for these self antigens or if these self antigens are only expressed in low numbers at all times and hence really belong to the family of variable antigens.

In the previous models we assumed the T-cell to 'see' all antigens on an APC, which are at least for the case of dendritic cells about 300000 [205, 52, 128]. However, if we assume to investigate negative selection in combination with tissue restricted antigens, these antigens are mostly presented by mTECs and only by some DCs that get the antigenic material via mechanisms of crosspresentation from mTECs. These epithelial cells express about $1 - 5 \cdot 10^4$ MHC molecules and present about 90 different TRA types besides other antigens [138, 52]. Hence, the number of antigens a T-cell sees is much smaller than before. Even if we only investigate antigens that are regularly expressed in all tissues and are mostly presented by DCs, we have to note that over the course of time DCs lose pMHC molecules from their surface [85]. Therefore, at least thymus-homing DCs that enter the thymus from the periphery might have lost antigens on their way into the thymus and lose even more in the thymus. It is also true that T-cells scan only parts of a dendritic cell and therefore do not see all antigens [85]. Consequently, a reduction of the number of presented antigens is not only adequate for the thymus but also for the periphery.

6.1 THE MODEL

With these facts at hand we are now able to establish a new, discrete model of T-cell activation, where we speak of discrete because we assume to have a discrete set of potentially immunogenic antigen classes in an organism.

The essence of BRB-like models is that a T-cell receives a total stimulation that is composed of a mixture of signals induced by various antigens and has then to decide if one of the stimulating antigens is foreign. One major difficulty for the analysis of

these models is the stimulation rate W , where no closed form of its distribution exists. Unlike Molina-Paris and van den Berg [204], we do not simplify this distribution in order to make it mathematically tractable, but rely again on our previously described importance sampling approach that should be suitable for the analysis of all BRB-like models. In order to introduce a model that enables us to keep track of the negative selection parameters, we follow our ideas from the previous section, that is we do not assume an infinite space of different antigen types and we introduce classes of antigen types. Each class consists of antigen types that induce similar stimuli to a given T-cell. We furthermore assume that all these antigens appear in high frequencies and we ignore all antigen types that are only expressed in low frequencies. An APC presents only members of some of these antigen classes and not from all at once as we assumed for the extended BRB model in Sec. 5.3.

This assumption leads to the definition of two variants of our model. The antigens presented by an APC can either be sampled without replacement from the antigen classes, such that each presented antigen type belongs to another class, or they can be sampled with replacement, whereby we do allow for multiple antigen types from the same antigen class being presented by an APC. Accordingly we now define our T-cell activation model for the sampling without replacement variant (abbreviated by 'swor') and afterwards for the sampling with replacement variant (abbreviated by 'swr').

'swor' variant of the discrete T-cell activation model

1. Let $S := \{1, 2, \dots, K\}$ be the set of (high-frequency) self antigen classes. (We do not consider low-frequency self antigens at all because we know they do not matter).
2. Let T-cell i be defined by the individual stimulation rates induced by *all* self antigen classes, i.e., $T_i := (W_{i1}, \dots, W_{iK})$. The W_{ij} are i.i.d. $\sim g$ (see page 33 and Figure 3.2), drawn once and fixed for the entire life of the T-cell. (This is a novel aspect of the model.)
3. An APC r presents (and is defined by) a subset of size n_s of all self antigen classes, i.e., $A_r \subseteq S$, $|A_r| = n_s$, where the A_r are independent of each other, and the elements of every A_r are drawn from S independently and without replacement. Every antigen is displayed at the same copy number z_s . There are $R + 1$ APCs.
4. When T_i meets A_r , it adds together the stimulation rates it assigns to this APC's antigens, i.e., $G_i^{(A_r)}(z_f) = (qz_s \sum_{a \in A_r} W_{ia}) + qz_f W_{i, n_s + 1}$. The factor $q = M / (z_f + z_s n_s)$ is the displacement factor to ensure a constant total antigen number M on an APC if we add foreign antigens.
5. T_i survives negative selection if $G_i^{(A_r)}(0) < g_{\text{thy}}$, $1 \leq r \leq R$.
6. A surviving T cell is then sent to the periphery and is activated if $G_i^{(A_{R+1})}(z_f) > g_{\text{act}}$.

'sur' variant of the discrete T-cell activation model

In a similar way we can introduce our second model variant.

1. Let $S := \{1, 2, \dots, K\}$ be the set of (high-frequency) self antigen classes.
2. Let T-cell i be defined by the individual stimulation rates induced by *all* self antigen classes, i.e., $T_i := (W_{i1}, \dots, W_{iK})$. The W_{ij} are i.i.d. $\sim g$ (see page 33 and Figure 3.2), drawn once and fixed for the entire life of the T-cell.
3. An APC r presents (and is defined by) a multiset of size n_s of all self antigen classes, i.e., $A_r = (S, m)$, where m is a function from S to \mathbb{N} such that for each $s \in S$, $m(s)$ is the multiplicity of s and $\sum_{s \in S} m(s) = n_s$. In this case m follows a multinomial distribution on the set S , where the probability to choose any $s \in S$ is $1/K$ and is realised by the sampling of n_s antigens from the set S independently and with replacement. Every antigen is displayed at the same copy number z_s . There are $R + 1$ APCs.
4. When T_i meets A_r , it adds together the stimulation rates it assigns to this APC's antigens, i.e., $G_i^{(A_r)}(z_f) = (qz_s \sum_{a \in A_r} m(a)W_{ia}) + qz_f W_{i, n_s + 1}$. The factor $q = M/(z_f + z_s n_s)$ is the displacement factor to ensure a constant total antigen number M on an APC if we add foreign antigens.
5. T_i survives negative selection if $G_i^{(A_r)}(0) < g_{\text{thy}}, 1 \leq r \leq R$.
6. A surviving T cell is then sent to the periphery and is activated if $G_i^{(A_{R+1})}(z_f) > g_{\text{act}}$.

Choice of parameters

We analyse the model under the assumption that the number of copies of foreign antigens z_f is either 0 if no foreign antigen is present or z_s if a foreign antigen is present. Note that this is a significant restriction in contrast to the basic BRB model where we achieved foreign-self discrimination by an increase in the copy number of the foreign antigen type.

We know from the previous analysis of the basic BRB model that under the assumption of similar antigen copy numbers for self and foreign, foreign-self discrimination is not possible. Although it does not follow exactly the basic BRB model this holds also for our model before negative selection, as we will show in Section 6.3. Hence, we are confined to negative selection in order to achieve foreign-self discrimination.

For the introduction of negative selection in this model we need two different parameters, the T-cell survival probability and the number of APC meetings. As already mentioned previously the former is estimated to be in the range of 35–60% and the latter is at most about 2400 (see Section 2.2.4). For our model we therefore choose $R = 2000$ APC meetings during negative selection for every T-cell and adjust the thymic activation threshold g_{thy} such that 50%(or 40%) of all T-cells survive the process (the adjustment

of g_{thy} is done via simulation. Note that this is not a difficult task, since neither survival nor death is a rare event.). Especially the number of APC meetings R (also known as rounds of negative selection) is still under discussion and there is experimental and theoretical evidence that this number might either be much too high or that the variance is quite high, that is some T-cells see much less APCs whereas some much more (see therefore chapter 7). However, as this estimate results from the most recent experiments we chose to adhere to them.

According to the previous section, we chose the number of antigen classes K to be either 1000 or 10000. The number of sampled antigen classes n_s is either 50 or 100. These numbers will suffice to give us a good idea on how they affect negative selection and therewith foreign-self discrimination in our model as we will see in Section 6.3.

6.2 SIMULATION APPROACH

For the proper analysis of this model we have to solve the problem to estimate the probability of T-cell activation in the periphery with and without foreign antigen. The assumption that T-cell activation in general has to be a rare event also holds for this model and an analytical or numerical calculation of these probabilities is impossible because of the distribution of the stimulation rate W . Hence, a simulation method is necessary for the estimation. Again, the easiest way is the straightforward simple sampling approach, which consists just of sampling of a T-cell, letting this T-cell undergo negative selection and afterwards meet an APC in the periphery. If this T-cell survives and is activated we increase a counter and start the procedure again for a new T-cell. At the end we divide our counter by the total number of procedure rounds and get the estimate for the probabilities. This would be time-consuming even without negative selection as we deal with rare events and to estimate the probability of such events a huge amount of simulation trials is needed. By the inclusion of negative selection the whole simulation procedure slows down even more. We therefore chose to reuse central ideas of our importance sampling and create new heuristic simulation methods that fit the model. Obviously, because we have two different variants of our model depending on the way how the self antigens are sampled, we had to develop two different simulation methods. We start with the method for the swor variant.

IS simulation for the swor variant

First of all we have to note that our model consists of 3 steps, creation of a T-cell, checking whether it survives negative selection and testing its activation capability. If we introduce the further assumption that every T-cell is only subjected to one APC meeting in the periphery, we can interchange the order of these steps to introduce a more efficient sampling scheme than simple sampling. The creation of a T-cell, the random choice of antigens an APC presents in each selection step and for the activation test in the periphery are independent events. Therefore, we can also choose the first n_s antigens for the activation test, then sample the stimulation rates of these antigens, check if the

stimulation rate exceeds a given threshold g_{act} and afterwards add the remaining antigen stimulation rates and start negative selection. Note, that we arbitrarily chose to use the first n_s stimulation rates. We could have chosen any other combination of n_s indices, because one crucial assumption in our model is that we draw all stimulation rates W independently and we do not superimpose any other restriction on the antigen repertoire (such as, that some antigen classes are presented more often). The consequence is a permutation invariance of the indices. This change of order enables us to reuse our initial simulation method for the original BRB model, because the total stimulation rate formula describing the stimulus that is induced to a T-cell T_i can be expressed as

$$G_i(z_f) = qz_s \sum_{j=1}^{n_s} W_{ij} + qz_f W_{i,n_s+1} \quad (6.1)$$

Equation (6.1) is the same as equation (3.3) without the variable antigens. The first step for the estimation of the activation probability of a random T-cell that has survived negative selection, $\mathbb{P}(G^{\text{neg}}(z_f) \geq g_{\text{act}})$, is therefore to use the importance sampling scheme we developed for this model (see algorithm 1). To include negative selection we just have to add a minor modification. We want to estimate

$$\begin{aligned} \mathbb{P}(G^{(R+1)}(z_f) \geq g_{\text{act}} \mid \Omega) &= \frac{\mathbb{P}(G^{(R+1)}(z_f) > g_{\text{act}}, \Omega)}{\mathbb{P}(\Omega)} \\ &= \mathbb{P}(\Omega \mid G^{(R+1)}(z_f) > g_{\text{act}}) \frac{\mathbb{P}(G^{(R+1)}(z_f) > g_{\text{act}})}{\mathbb{P}(\Omega)}, \end{aligned} \quad (6.2)$$

where Ω is the event of surviving negative selection. The probability to survive negative selection $\mathbb{P}(\Omega)$ is predefined and we estimate the probability $\mathbb{P}(G_i^{(R+1)}(z_f) > g_{\text{act}})$ with our usual importance sampling algorithm. The only new component is the factor $\mathbb{P}(\Omega \mid G^{(R+1)}(z_f) > g_{\text{act}})$. We can estimate this probability if we extend our simulation. In our algorithm we at first only generate the first n_s stimulation rates according to the tilted W^θ s, check whether $G(z_f)$ exceeds g_{act} and calculate the reweighting factor. In addition if the stimulus exceeds the threshold we now generate $K - n_s$ stimulation rates from the untilted stimulation rate distribution and thereby get a complete T-cell. Next we let this T-cell undergo negative selection. By counting the number of T-cells that survive negative selection and dividing this number by the number of samples that reach the activation threshold we get an estimate of $\mathbb{P}(\Omega \mid G^{(R+1)}(z_f) > g_{\text{act}})$.

An approximative IS simulation for the swor variant

Unfortunately, it turns out that this importance sampling scheme can be quite inefficient for certain parameter combinations in our model. The higher the activation threshold, the higher the n_s stimuli have to be in order to induce a total stimulus which is higher than the threshold. For the case of 1000 antigen classes it is quite probable that during negative selection some of these antigens are seen together and hence the thymic activation threshold is reached. In fact, with increasing threshold values the number of T-cells

that survive decreases exponentially. This effect is even worse for the T-cells that induce stimuli which exceed g_{act} . Therefore we have to increase the number of sampled T-cells drastically for higher threshold values to get a good estimator.

As a consequence we developed a second, more efficient but approximative importance sampling scheme. The idea behind it is rather simple. At first we estimate a new stimulation rate distribution after negative selection. We therefore generate T-cells T_i , let them undergo negative selection and take all W_{ij} of the surviving T-cells to create a new empirical distribution of the post-selection stimulation rate. As T-cell survival is not a rare event, this simple sampling approach is not too time consuming.

We now assume that generating T-cells from the original stimulation rate distribution and letting them undergo negative selection leads to the same repertoire of surviving T-cells as if we directly generate them from the post-selection stimulation rate distribution. Hence, we can use this distribution instead of the original one in order to estimate T-cell activation probabilities in our IS approach. Of course, this can only be approximatively true, because we condition on the survival of negative selection and therefore the W_{ij} are not independent anymore. However, we will see that for our model the results of this simulation approach are comparable to the ones of the IS scheme describe before.

In the result section 6.3 we will show that the estimations from both importance sampling schemes are similar at least for the range of activation thresholds where we get good estimates for our first method. As the second method does not have to deal with the drawback of the first, we can use it to estimate probabilities for far higher activation thresholds.

IS simulation for the swr variant

Antigen sampling with replacement changes the situation. All the considerations regarding the independence of the different steps in the model and the permutation invariance of the antigen indices still hold, but we do not have a constant copy number z_s for every antigen type presented. Instead we have random copy numbers. Equation (6.1) therefore changes to

$$G_i(z_f) = qz_s \sum_{s \in S: m(s) > 0} K_s W_{is} + qz_f W_{i, n_s + 1}, \quad (6.3)$$

where $K_s = m(s)$ are (dependent) random variables with $\sum_{s=1}^{n_s} K_s = n_s$ and follow a multinomial distribution. While it might seem a notational overload to introduce these random variables K_s as they are covered by the random multiplicity function $m(\cdot)$, we feel that this will help to clarify our simulation approach which we develop in the following.

In principle, equation (6.3) is also covered by the general theory of rare event simulation via exponential tilting. However, in our original importance sampling scheme we assumed independence of the different summands (see section 5.1.1). Thereby, the generating function of the total stimulus random variable in eq. (6.3) would factorise. Here we have a dependence between the different summands and no factorisation of the moment-generating function. Numerical calculations would require the knowledge

of the joint probabilities, especially the (small) probabilities we are looking for in our simulation.

We could of course just use the importance sampling scheme which we derived for equation (6.1). This would be rather inefficient, depending on how the K_s actually look like (only for an increasing set S of self antigen classes, the probability of sampling antigen indices severalfold decreases, such that ultimately $\mathbb{P}(K_s = 1) \rightarrow 1$, $s = 1, \dots, n_s$, which would be the case without replacement). In order to develop an efficient importance sampling scheme we can make use of the permutation invariance of the antigen indices, again, as we did for the 'swor' variant. In the context of this model variant this means also that it does not matter to which antigen classes the multiplicity factors actually belong, once they are generated. Consequently, it also does not matter if we create these K_s first and then sort them from highest to lowest, such that we have $K_{\pi(s_1)} > K_{\pi(s_2)} > \dots > K_{\pi(s_{n_s})}$ with $s_j \in S$ and π a permutation of the indices. Only afterwards we generate the stimulation rates W .

A first idea for a better IS scheme would be to define a related model for which it is easy to compute tilting factors. We therefore estimate mean values $\bar{K}_{\pi(s_j)}$ for the sorted multiplicity factors $K_{\pi(s_j)}$ over a huge amount of T-cell APC meetings and define a modified equation for the total stimulus induced to the T-cell T_i :

$$G_i(z_f) = qz_s \sum_{j=1}^{n_s} \bar{K}_{\pi(s_j)} W_{ij} + qz_f W_{i, n_s + 1}. \quad (6.4)$$

Computation of the tilting parameters is straightforward as the $\bar{K}_{\pi(s_j)}$ are constants. We can then use these for the simulation of our original model.

However, we thereby ignore the stochastic nature of the $K_{\pi(s_j)}$, which can, again, lead to problems in our sampling scheme. To clarify this point we have to think about the nature of the events we are about to estimate. We want to estimate rare events, in this case reaching high activation stimuli. For this to happen we need either high stimulation rates or (and this is the crucial point) not so high stimuli multiplied by high $K_{\pi(s)}$.

So, we do not only have the possibility to increase the probability to reach high stimuli by tilting the stimulation rate distributions, but also by changes in the discrete Uniform distribution \mathcal{U} on S which guides the sampling with replacement of the antigen indices and therewith the generation of the $K_{\pi(s_j)}$.

In order to develop an IS scheme based on these considerations we, at first, go one step further. According to the importance sampling procedure, tilting of the stimulation rate distributions and changes of the uniform distribution both lead to the need of reweighting back to the original distributions. We now consider these reweighting factors as the 'costs' to get high stimuli and we want to minimise these costs. Equation (6.4) would be useful if the costs for tilting the stimulation rate distribution are much lower than the costs of changing \mathcal{U} .

Let us now imagine we are in the opposite situation that the 'changing costs' to enhance the probability to choose a high stimulus antigen severalfold are much lower than the 'tilting costs' to elevate the probability to create a T-cell with many high stimuli. In this situation it is appropriate to only change the discrete Uniform distribution \mathcal{U} .

One possible importance sampling scheme that could be used in this situation works as follows.

We just generate a large number L of T-cells T_i , sort the stimuli for every T-cell from lowest to highest, such that we have T-cells $\tilde{T}_i = (W_{i,\pi(1)}, \dots, W_{i,\pi(K)})$ with $W_{i,\pi(1)} < W_{i,\pi(2)} < \dots < W_{i,\pi(K)}$ and π a permutation of the indices. We can then define a vector $V = (V_1, \dots, V_K)$ with $V_j = \frac{1}{L} \sum_{i=1}^L W_{i,\pi(j)}$, which we call mean stimulation rate vector. This vector can be used to calculate an appropriate change of \mathcal{U} . Again, we propose a modified version of the exponential tilting. That is we are looking for an optimal ζ to get a new distribution \mathcal{U}^ζ via tilting of the $p_j = \frac{1}{K}$ from the discrete uniform distribution with $e^{\zeta V_j}$:

$$p_j^\zeta = \exp(\zeta V_j) \frac{1}{K} / C, \quad (6.5)$$

as usual, C is the normalisation factor. The optimal tilting parameter can be calculated by solving

$$\mathbb{E} \left[qz_s \sum_{i=j}^{n_s} V_{Q_i} + qz_f W_{n_s+1}^\zeta \right] = g_{\text{act}}, \quad Q_i \sim \mathcal{U}^\zeta \quad (6.6)$$

for ζ . The probability of T-cell activation is then calculated by generating and sorting of a T-cell, choosing stimuli by sampling from \mathcal{U}^ζ and then checking if g_{act} is exceeded and reweight back to the original uniform distribution. So, the trick here is to enhance the probability that a T-cell encounters an APC which primarily presents antigens that induce high stimuli to the T-cell. See algorithm 3 for a summary of the method.

We now have importance sampling schemes for the two extreme cases, where tilting of the stimulation rate distributions is either much more expensive or much cheaper than tilting of \mathcal{U} . However, in reality we do not know if any of these situations applies. For many parameter settings it will be rather something in between. Hence, we have to go one step further to get a really efficient importance sampling scheme and combine tilting of the stimulation rate distributions and tilting of \mathcal{U} . The main problem for this is that we have to calculate optimal parameter pairs (ϑ, ζ) . But the question is how we can describe efficiency or optimality in the given context. In Section 5.1.3 we used ideas from large deviation theory to derive an optimal IS scheme, where the optimality criterion is based on the relative error η . For optimal tilting parameter pairs we therefore also want to use this criterion, that is we are looking for (ϑ, ζ) given an activation threshold g_{act} that lead to the best possible reduction of η for a predefined number of samples.

In order to find such pairs we need two functions $L^1(\zeta)$ and $L^2(\vartheta)$. The first function $L^1(\zeta)$ is the logarithm of the moment-generating function of $\mathcal{U}_{\otimes n_s}^\zeta$:

$$L^1(\zeta) := n_s \ln M^1(z_s q \zeta) \quad (6.7)$$

where $M^1(\cdot)$ is the moment-generating function of \mathcal{U} .

The second function $L^2(\vartheta)$ is defined as

$$L^2(\vartheta) := \sum_{j=1}^{n_s} \ln M^2(z_s q \bar{K}_{\pi(s_j)}^\zeta \vartheta) + \ln M^2(z_f q \vartheta), \quad (6.8)$$

Algorithm 3: Estimation of the activation probabilities when only ζ is changed

Input: activation threshold g_{act} , foreign copy number z_f , number of antigen classes per APC n_s , total number of antigen classes K , density g of W , number of samples $N \in \mathbb{N}$, size L of the (large enough) T-cell repertoire

Result: estimate of $\mathbb{P}(G(z_f) \geq g_{\text{act}})$

1 **Preprocessing:**

2 generate a sorted T-cell repertoire $\{\tilde{T}_i\}_{1 \leq i \leq L}$ and calculate the mean stimulation rate vector V

3 compute ζ for g_{act} by solving eq (6.6)

4 calculate \mathcal{U}^ζ and g^ζ

5 initialise estimators $E_{g_{\text{act}}}$

6 **Simulation start:**

7 **for** $i = 1, \dots, N$ **do**

8 choose a T-cell \tilde{T}_i randomly and uniformly from the repertoire

9 draw random numbers Q_1, \dots, Q_m from the tilted distribution \mathcal{U}^ζ and W_{n_s+1} from g^ζ

10 **if** $\left(\sum_{j=1}^{n_s} q \cdot z_s \cdot W_{iQ_j} + q \cdot z_f \cdot W_{n_s+1} \geq g_{\text{act}}\right)$ **then**

11 calculate the reweighting factor

12 increase estimator $E_{g_{\text{act}}}$ by the reweighting factor

13 **end**

14 **end**

15 estimate $\mathbb{P}(G(z_f) \geq g_{\text{act}})$ by $\mathbb{P}(G(\widehat{z_f}) \geq g_{\text{act}}) = E_{g_{\text{act}}}/N$

resulting from the distribution of the a total stimulus random variable which is a modified version of eq. (6.4):

$$G^\vartheta(z_f) = qz_s \sum_{i=1}^{n_s} \bar{K}_{\pi(s_j)}^\zeta W_i^\vartheta + qz_f W_{n_s+1}^\vartheta, \quad (6.9)$$

where $M^2(\cdot)$ is the moment-generating function of the stimulation rate distribution and the $\bar{K}_{\pi(s_j)}^\zeta$ are calculated as described previously but using \mathcal{U}^ζ instead of \mathcal{U} . Note that ϑ is implicitly dependent on ζ via this connection and we calculate ϑ by solving

$$\mathbb{E}[G^\vartheta(z_f)] = g_{\text{act}}. \quad (6.10)$$

This follows from our original importance sampling scheme (see Section 5.1.3).

With these two functions we can show via simulation that we minimise η if we estimate parameters ζ and ϑ such that

$$L^1(\zeta) = L^2(\vartheta) \quad (6.11)$$

holds. Unfortunately a mathematical proof of this result is still lacking.

However, this result enables us to develop a straightforward numerical method to estimate the tilting parameters. We have to stepwise search through the space of possible (ϑ, ζ) pairs. At first we start with a low value of ζ . We then generate n_s antigen indices

from U^ζ and calculate how many times the individual indices are chosen and sort these numbers K_s from highest to lowest to get the $K_{\pi(s_j)}$. This is repeated several times and we calculate the mean $(\bar{K}_{\pi(s_1)}, \dots, \bar{K}_{\pi(s_{n_s})})$. Afterwards, we have to solve eq. (6.10) for ϑ in the same way as we do in the original importance sampling scheme. Now we can calculate $L^1(\zeta)$ and $L^2(\vartheta)$. The whole process is then started again with a higher ζ . We stop if in step $t - 1$ $L^1 < L^2$ and in step t $L^1 > L^2$ or vice versa. At this point it would be necessary to decrease ζ slightly, check again and so on. However, in order to speed up the whole procedure we make an extra assumption. For not too big step sizes of ζ and consequently ϑ we assume that L^1 and L^2 behave almost linearly. Therefore we can calculate the intersection of the line $\overline{L_{t-1}^1}, \overline{L_t^1}$ with $\overline{L_{t-1}^2}, \overline{L_t^2}$ and from this result we can extrapolate an almost optimal ζ and ϑ . This informal description is formalised in algorithm 4

Algorithm 4: Calculation of the optimal tilting parameters

Input: activation threshold g_{act} , foreign copy number z_f , number of antigen classes per APC n_s , total number of antigen classes K , density g of W

Result: optimal ζ , optimal $\vartheta, (\bar{K}_{\pi(s_1)}, \dots, \bar{K}_{\pi(s_{n_s})})$

```

1 initialise  $t = 0$ 
2 set  $L_0^1 = 0, L_0^2 = 0$ 
3 for  $\zeta$  stepwise increasing do
4    $t = t + 1$ 
5   calculate  $U^\zeta$ 
6   initialise  $(\bar{K}_{\pi(s_1)}, \dots, \bar{K}_{\pi(s_{n_s})})$ 
7   for large number of samples  $M$  do
8     generate  $n_s$  indices via  $U^\zeta$ 
9     count the abundances with which the indices appear, sort them and denote
      them by  $(K_{\pi(s_1)}, \dots, K_{\pi(s_{n_s})})$ 
10     $(\bar{K}_{\pi(s_1)}, \dots, \bar{K}_{\pi(s_{n_s})}) = (\bar{K}_{\pi(s_1)}, \dots, \bar{K}_{\pi(s_{n_s})}) + (K_{\pi(s_1)}, \dots, K_{\pi(s_{n_s})})$ 
11  end
12   $(\bar{K}_{\pi(s_1)}, \dots, \bar{K}_{\pi(s_{n_s})}) = (\bar{K}_{\pi(s_1)}, \dots, \bar{K}_{\pi(s_{n_s})})/M$ 
13  estimate  $\vartheta$  from eq. (6.10)
14  calculate  $L_t^1$  and  $L_t^2$ 
15  if  $L_{t-1}^1 < L_{t-1}^2$  and  $L_t^1 > L_t^2$  or vice versa then
16    | break
17  end
18 end
19 calculate the optimal  $\zeta$  by calculation of the crossing point of the lines  $\overline{L_{t-1}^1}, \overline{L_t^1}$  and
       $\overline{L_{t-1}^2}, \overline{L_t^2}$ 
20 use  $\zeta$  to estimate the optimal  $\vartheta$ 

```

If we calculate the optimal tilting parameters for several g_{act} , we can use the previous optimal ζ as the starting point for our stepwise search for the next optimal ζ . Thereby only very few search steps are needed. Alternatively, also a binary search for ζ would be possible.

Now that we have established a method to compute good tilting parameters for a given activation threshold, we are ready to formulate the actual importance sampling scheme for the T-cell model and antigen sampling with replacement. It is essentially similar to our usual importance sampling scheme but with an additional step where we simulate the $K_{\pi(s_j)}$ out of the tilted uniform distribution as already described before. In the end we therefore have to multiply the reweighting factors resulting from tilting the W -distributions by the reweighting factors resulting from this first step. For the whole algorithm see Alg.. 5

Algorithm 5: Estimation of the activation probabilities for the second model variant

Input: activation threshold g_{act} , foreign copy number z_f , number of antigen classes per APC n_s , total number of antigen classes K , density g of W , number of samples $N \in \mathbb{N}$

Result: estimate for $\mathbb{P}(G(z_f) \geq g_{\text{act}})$

- 1 preprocessing: use Alg. 4 to calculate ζ, ϑ and $(\bar{K}_{\pi(s_1)}, \dots, \bar{K}_{\pi(s_{n_s})})$
- 2 initialise estimator $E = 0$
- 3 **for** $i = 1, \dots, N$ **do**
- 4 generate $(K_{\pi(s_1)}, \dots, K_{\pi(s_{n_s})})$ from \mathcal{U}^ζ
- 5 generate W_1, \dots, W_{n_s+1} from $g^{\bar{K}_{\pi(s_1)}\vartheta}, \dots, g^{\bar{K}_{\pi(s_{n_s})}\vartheta}, g^\vartheta$
- 6 **if** $qz_s \sum_{i=1}^{n_s} K_{\pi(s_j)} W_i + qz_f W_{n_s+1} > g_{\text{act}}$ **then**
- 7 calculate reweighting factors:
- 8 $r_1 = \prod_{i=1}^{n_s} (p_{s_j} / p_{s_j}^\zeta)^{K_{\pi(s_j)}}$
- 9 $r_2 = [\prod_{i=1}^{n_s} g(W_i) / g^{\bar{K}_{\pi(s_j)}\vartheta}(W_i)] \cdot [g(W_{n_s+1}) / g^\vartheta(W_{n_s+1})]$
- 10 $E = E + r_1 r_2$
- 11 **end**
- 12 **end**
- 13 estimate activation probability by $\mathbb{P}(G(\widehat{z_f}) \geq g_{\text{act}}) = E/N$

So far we have ignored the new part of our discrete T-cell model, the negative selection process. To estimate the activation probabilities after negative selection we propose to use the same two approaches which we use for the swor variant.

First of all we can readily modify our simulation method to include negative selection directly. Therefore for every sample, that is for every T-cell T_i , we generate n_s stimulation rates according to the algorithm and then generate the missing $K - n_s$ stimulation rates from the untilted stimulation rate distribution in order to get a full T-cell. This T-cell is then subjected to negative selection and only if it survives the reweighting factor is added to the estimate.

This is exactly the same procedure as for the swor variant and it also has the same problems. As a second simulation method, we therefore propose to again calculate a new stimulation rate distribution after negative selection and use this distribution with algorithm 5 for the stimuli induced by self antigens. Again, this is only an approximate method which is not unbiased.

In Section 6.3 we will show that also for this model variant both simulation methods produce similar estimates, especially in light of the fact that we are not too concerned with absolute accuracy but that we are seeking for differences in foreign-self discrimination that are measurable in powers of ten.

In this section we have established simulation methods for the estimation of activation probabilities for our discrete T-cell activation model either assume antigen sampling with or without replacement and with or without negative selection. We could reuse much of the ideas we already used in the previous chapter and thereby ensure to have an efficient sampling scheme. The next section is now devoted to a rigorous analysis of the model by means of our simulation methods.

6.3 RESULTS

As pointed out before, there are different parameters in the T-cell model that can be varied and additionally we have two different ways to model the antigen selection process by the APCs. Hence, in this section we will try to get an understanding of how these parameters affect the outcome of the model and contrast the results of the model in the swor variant with the results in the swr variant. For the probability estimations we used 10000 samples per threshold in the swor variant and 100000 samples per threshold in the swr variant. We start with a basic parameter set. For this set we not only examine foreign-self discrimination but also the negative selection process itself. As this process is also stochastic it is interesting to see how this stochasticity effects the outcome of negative selection. Furthermore we have examine the simulation methods with regard to their efficiency. Afterwards we investigate the effect on foreign-self discrimination if we introduce more classes of self antigens, increase the number of antigens classes presented on an APC and finally combine both.

The basic parameter set

At first we start with a set S of $K = 1000$ antigen classes out of which $n_s = 50$ are presented on an APC with a total of $M = 25000$ antigens. Consequently, the number of self antigen copies is $z_s = 500$. In Figure 6.1 we show the results of our estimation of the activation curves for the swor model (left) and the swr model (right) before negative selection.

From here on we will proceed in the same manner throughout this section. That is, for a given set of parameters we show the estimates for both model variants next to each other, where the left one is always for sampling without replacement and the right one for sampling with replacement.

Coming back to the activation curves, in both cases we can see that foreign-self discrimination is not possible. In fact, activation only by self is slightly more probable than activation with foreign included. This is a consequence of the displacement of self antigens by foreign antigens. Obviously, the probability to get medium to high stimuli from some presented self antigens is higher than to get a similar stimulus from the one

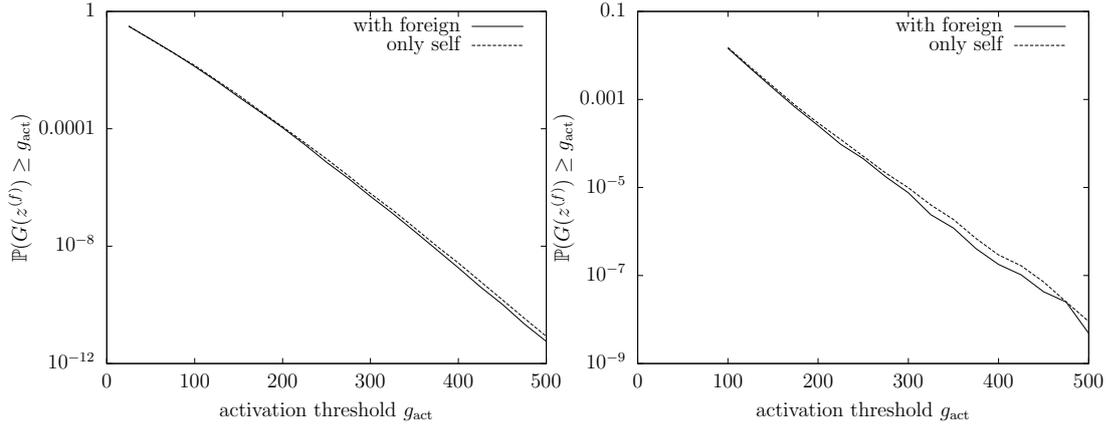


Figure 6.1: Activation curves before negative selection for $K = 1000$ and $n_s = 50$. In the swor case (left) as well as in the swr case (right) no foreign-self discrimination is possible. A comparison of both Figures also reveals that for this parameter combination the effect of sampling antigens severalfold in the swr case is quite high. This is reflected in the much higher activation probability in the swr case for a given threshold compared to the swor case. Note that the y-axis have different scales.

foreign stimulus. This is just due to the fact that self and foreign antigen stimulation rates not only follow the same distribution but are also multiplied by the same copy number and there are much more self stimuli generated than the single foreign one. By the displacement the total stimulation rate loses some of the high self stimuli in favour of (often) lower foreign stimuli. Mathematically speaking we distribute the expectation on 51 instead of 50 i.i.d. random variables and this reduces the variability. For the swor case we expected as much, since there is no difference from the actual BRB model besides the lack of variable antigens. Moreover, if you compare these estimations with the ones from our previous analysis of the BRB model, we can not only see qualitative but also quantitative similarities (see Figure 5.2. This proves once again the minor importance of the variable antigens on the activation curves (at least before negative selection).

For the swr case the differences between only self and with foreign activation probabilities in favour of self activation are even a bit more accentuated. The reason for this is of course also the displacement of self antigens. By construction of our simulation scheme and as a consequence of dealing with rare events, there exists a correlation between the two events of having a high stimulus antigen and presenting this antigen severalfold. Introducing a foreign antigen to the APC means that we have to displace self antigens. By the way we constructed the model this means that we remove more of these high stimuli antigens that are presented severalfold than low stimuli antigens. This cannot be compensated by the foreign antigen. Hence, the probability of getting a high enough total stimulus is lowered a bit.

It is also necessary to compare the activation rates of the swr model with the swor model. At a first glance on Figure 6.1, they seem similar. However, if we note the actual values of the y -axis we can see that for a given activation threshold, the probability of activation is much higher in the swr case in comparison to the swor case. To a certain extent this result could have been deduced from the precalculated tilting parameters

which we have not presented here. At least for this set of parameters it is obvious that the randomness coming from sampling with replacement has a major influence on the outcome of the model. In fact, here we reach similar estimates if we use algorithm 3 instead of algorithm 5 (data not shown).

It will now be interesting to see how negative selection affects the activation curves in both cases. We chose to introduce two different thymic activation thresholds for each model variant. One ensures that a T-cell has a probability of 50% to survive, the other lowers this probability to 40%. Evidently, these thresholds have to be different for each model variant due to the differences in the antigen sampling. To get a first idea on how effective negative selection might be, we have a look on empirical estimates of the W density before and after negative selection. We estimated them as probability mass functions, that is as normalised histograms with a step-size of $3.6 \cdot 10^{-4}$. For reasons of simplicity we keep on calling them densities in the following. They are shown in Figure 6.2

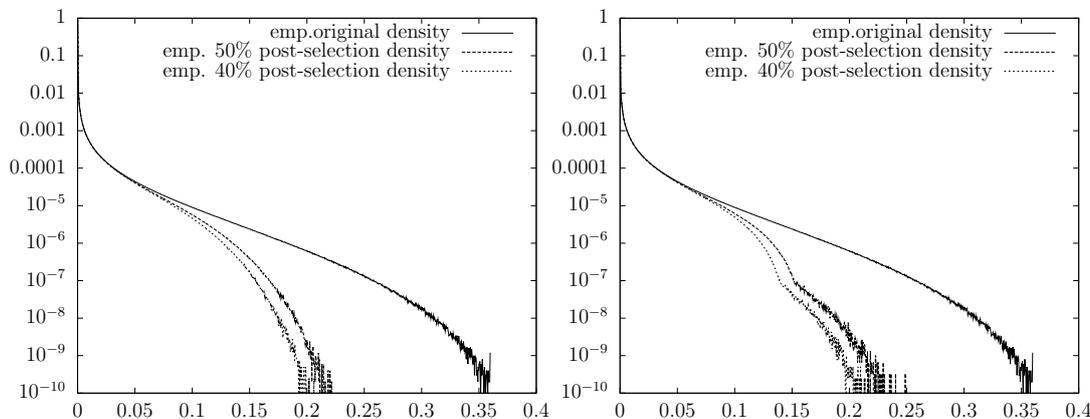


Figure 6.2: Empirical estimates of the pre- (original) and post-selection densities of W (swor (left) - swr (right)). We estimated them as probability mass functions (pmfs), that is as normalised histograms with a step-size of $3.6 \cdot 10^{-4}$. These estimates result from letting 10^7 randomly generated T-cells undergo the negative selection process. The stimulation rate vectors of the surviving T-cells are then used to compute the pmfs. In both cases the densities are compressed considerably in the horizontal direction by negative selection. However, there are both qualitative and quantitative differences between the post-selection densities of both model variants. In general negative selection in the swor case seems to work a little better, that is we estimate a compression factor of about 0.61 (0.56 for 40%). The overall compression factor in the swr case is about 0.67 (0.61 for 40%). Besides this quantitative difference, there is also the qualitative differences that the post-selection densities in the swr case show two phases. Up to a certain point they follow one super-exponentially decreasing slope and from this point on they follow a different not so fast super-exponentially decreasing slope. We seem to have two different compression factors depending on the stimulation rate. Thereby for a certain range of stimulation rates their probability is lower in the swr case than it is in the swor case.

There are not too big differences in the effect of negative selection on the W density if we compare both model variants. This is true if we either negatively select 50% or 60% of all T-cells. The densities are horizontally compressed considerably. This points to a strong effect of negative selection which should elevate the foreign-self discrimination

capability in the model. However, it is also of note that there are some qualitative differences between both variants. In the swr case there seem to be two different compression 'processes' at work. From a certain point on, the graph of the post-selection density is not declining so fast as before. Thereby the effect of negative selection is reduced. Instead of being compressed by more than a factor of 0.5 which is suggested by the initial slope of the post-selection densities, for higher stimulation rates the density is compressed by a factor of about 0.67, only (we calculate the compression factors just by comparing the x -values of the densities at a y -value of 10^{-10}). In the swor case the post-selection density is continuous with an estimated compression factor of 0.61. Our analysis of the BRB model indicates, however, that we need a compression factor of 0.5 or greater to enable safe foreign-self discrimination. This follows from the fact that if we set $z_f = 2z_s$, we have a good foreign-self discrimination in that model. But if it suffices to have $2W_f$, this should also hold for $\frac{1}{2}W_s$ instead of W_s and W_f instead of $2W_f$.

We assume that the two-phasic behaviour of the swr post-selection density results from the fact that we have the possibility to present an intermediate strength stimulus severalfold and thereby also reach g_{thy} . For the swor variant we do not have this advantage and therefore less selection pressure on intermediate stimuli. However, we do not know why we have a kink at this special point in the swr post-selection density.

Let us now turn to the results which are the very reason why we constructed this model, the activation curves after negative selection. In Figure 6.3 we present our estimations if we negatively select 50% of all T-cells.

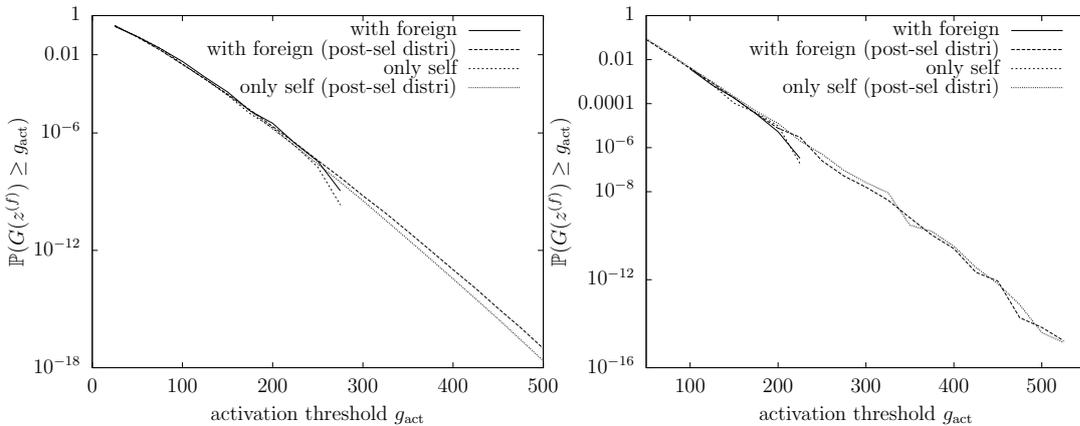


Figure 6.3: Activation curves after negative selection (swor (left) - swr (right)). We estimate the activation probability of a random T-cell under the assumption that it survives the negative selection process. The probabilities are estimated either by direct inclusion of negative selection in the original sampling scheme or by making use of the post-selection distribution. With regard to foreign-self discrimination, activation by foreign becomes more probable if we compare the results with the activation probabilities before negative selection. But, there is still no foreign-self discrimination possible.

We see that negative selection has a recognisable effect. Nevertheless, it is far from sufficient. For the swor case activation with foreign antigen is slightly more probable than activation only by self. For the swr case activation with or without foreign antigen are equiprobable. But we seek differences in probability in several orders of magnitude.

Otherwise the model is not working properly. These results prove our assumption we got from the post-selection densities. The situation does not change much if we negatively select 60% of all T-cells as you can see in Figure 6.4

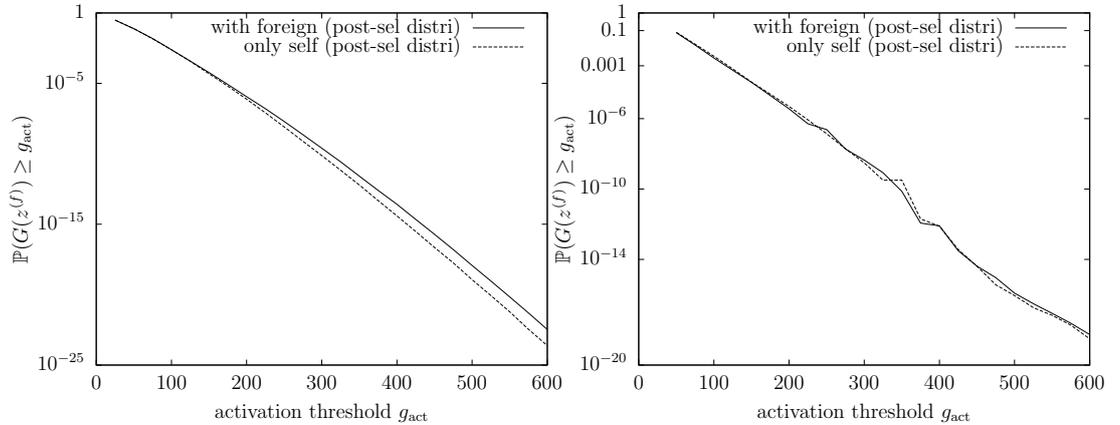


Figure 6.4: Activation curves after negative selection (swor (left) - swr (right)). We assume that the survival probability for a random T-cell is 40%. There is only a minor enhancement of the foreign-self discrimination capability in comparison to the assumption that 50% of all T-cells survive.

The impact of deleting an additional 10% of the T-cells is not vital, as was to be expected if we consider Figure 6.2. All in all we have to conclude that at least for our given basic set of parameters and our simple model of negative selection, foreign-self discrimination cannot be explained.

Brief analysis of the simulation methods

In this paragraph we analyse the efficiency of the simulation methods for the swor and the swr variant. In case of the swor variant we do not expect surprising results as we use the same algorithm as for the basic BRB model. For the swr variant we had to develop a new algorithm so it will be interesting to see how efficiently it works in comparison. For all the results presented in the following keep in mind that we used 10000 samples per threshold for the swor variant and 100000 samples per threshold for the swr variant.

In line with our theoretical considerations in Section 4.2 and our analysis in Section 5.2 we use the relative error as the criterion to evaluate the efficiency of the simulations. Recall that the relative error is just the standard deviation of our estimator divided by the estimator (see equation (4.10)). Note that we used a ten times higher sample size for the swr variant than we used for the swor variant. In Figure 6.5 we show the relative errors for the estimation of the activation probabilities before negative selection for both model variants. The relative error in the swor variant behaves very similar to the relative error in the simulations of the basic BRB model (see Figure 5.5) and shows the typical signs of asymptotic efficiency (subexponentially increasing relative error). In the swr variant the relative error increases much faster and the method seems not to be asymptotically efficient. However, the simulation method is still much more efficient than simple sampling.

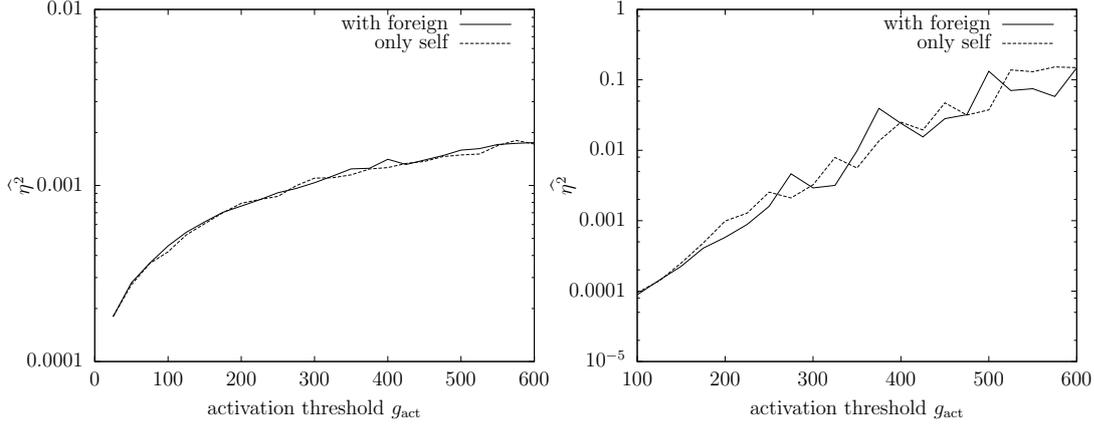


Figure 6.5: Relative errors of the estimation of the activation probabilities before negative selection (swor (left) - swr (right)). We used 10000 samples per threshold for the swor variant and 100000 samples for the swr variant. The randomness in the multiplicities in the swr variant introduces a much higher variance, which is reflected in the relative errors.

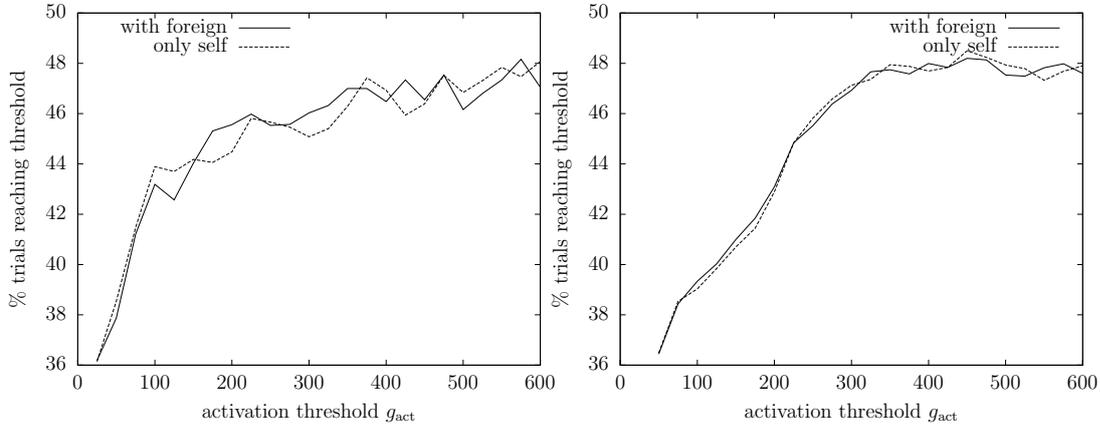


Figure 6.6: Percentage of samples that reach the activation threshold before negative selection (swor (left) - swr (right)). With increasing activation threshold the percentage of trials that reach the threshold increases up to a value of 48%. In our simulation methods we tilt such that reaching the activation threshold is the typical event, so these values are in the range of what was to be expected.

The core of both simulation methods is the tilting of the distributions such that reaching the activation threshold becomes the typical event. This should be reflected in the percentage of samples that reach a given activation threshold. In Figure 6.6 we show these percentages for both model variants. With increasing threshold values the percentages of samples that reach the threshold increases up to a value of 48% (swor). This reflects that with increasing activation thresholds both sample mean and sample median converge.

We already mentioned that including negative selection directly into our algorithms is possible but has the defect that many samples are lost. Hence, if we want to estimate activation probabilities for higher threshold values we need more samples. This was the motivation to develop the approximate importance sampling algorithms that use

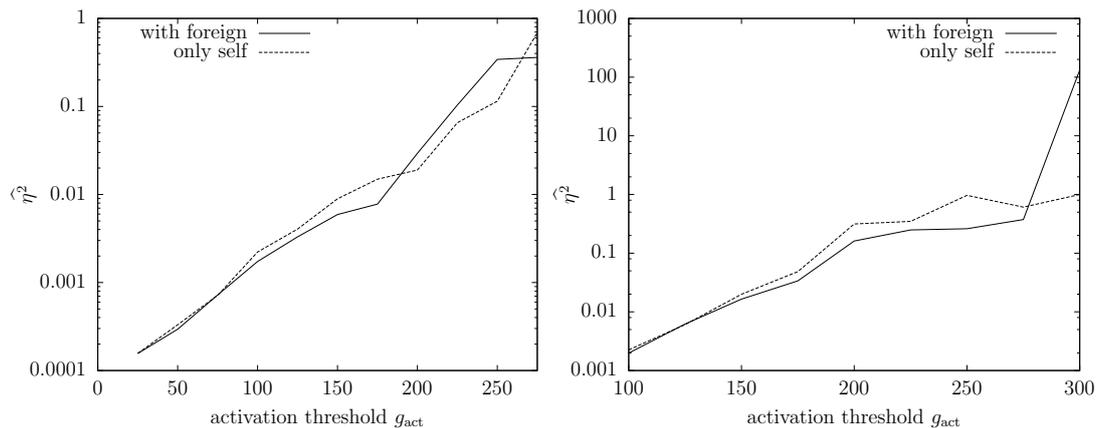


Figure 6.7: Relative errors of the estimation of the activation probabilities after negative selection (swr (left) - swr (right)). Negative selection is directly included in the simulation. We used 10000 samples per threshold value for the swr variant and 100000 samples for the swr variant. For both variants the relative error increases very fast. The reason for this is the exponentially decreasing number of samples that reach the activation threshold (see Figure 6.8).

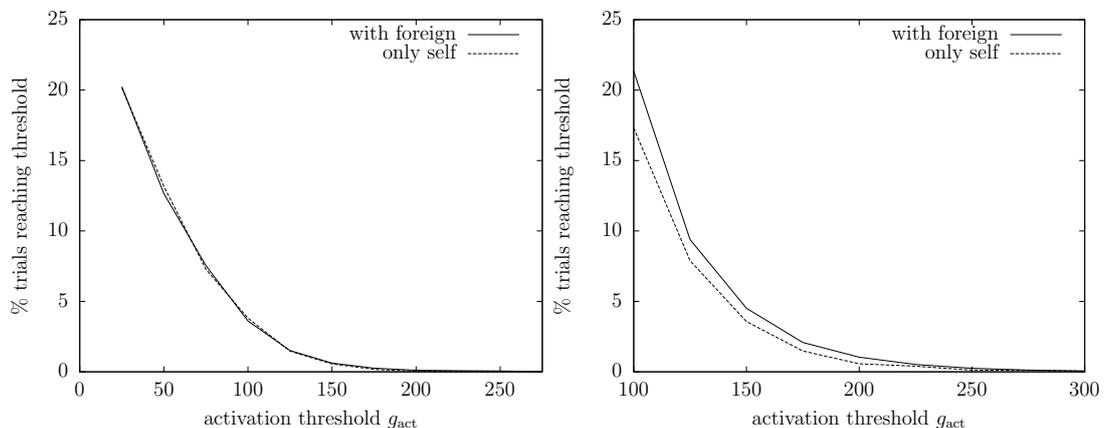


Figure 6.8: Percentage of samples that reach the activation threshold after negative selection (swr (left) - swr (right)). For both variants the percentage decreases to zero exponentially fast. This is the reason why the approximate importance sampling method is needed for the estimation of activation probabilities for higher threshold values.

the stimulation rate distributions after negative selection. To underline these claims we show in Figures 6.7 and 6.8 the relative error and the percentage of samples that reach a given activation threshold if we include negative selection directly.

Note that the percentage of samples that survive negative selection and reach the activation threshold decreases exponentially. This leads to a faster increase of the relative error in both model variants. Consequently, the sample size also needs to be exponentially increasing.

Finally we examine the reweighting factors of the samples that reach a given activation threshold. This is interesting because the idea behind our importance sampling algorithms is to generate the least unlikely events of all unlikely events. The reweight-

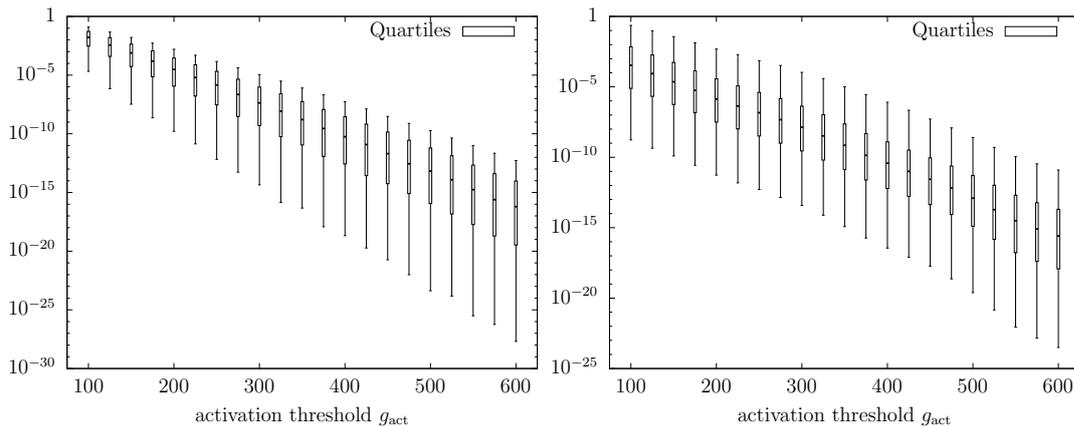


Figure 6.9: Boxplots of the reweighting factors of all samples that reach a given activation threshold (swor (left) - swr (right)). The boxes represent the quartiles and the whiskers the 0.025– and 0.975–quantiles.

ing factors can show us how well we achieve this goal. If our methods work as we intend them to, most of the reweighting factors should have values not too much smaller than the actual activation probabilities. In Figure 6.9 we present our analysis of the reweighting factors per threshold in form of boxplots. These plots show the differences between our simulation method for the swor and the swr variant and help to explain why the former works much better than the latter. In case of the swor variant, the length of the box (representing the quartiles) and the whiskers (representing the 0.025– and 0.975–quantiles) is first very small and then increases slowly. Furthermore, the length of the upper whisker is much smaller than the length of the lower whisker. If we translate this, it means that with increasing activation thresholds the events we generate are more and more spread over the space of all events reflecting stimulation rates higher than the activation threshold. Therefore we lose precision in our estimations. However, the bulk of the events is still in the part of the space with the least unlikely events, therefore we only lose precision very slowly.

In case of the swr variant, the situation looks different. Even for low activation thresholds we have long boxes and whiskers and they get longer with increasing threshold values. But, in contrast to the swor variant, this increase is slower. For very high threshold values, the box and whisker length in both variants is comparable. The events are not only more scattered, but also the bulk of our generated events does not fall into the region with the least unlikely events. This explains why the precision of our activation probability estimations in the the swr variant is much worse than for the swor variant.

Brief analysis of the negative selection process

Before going on to change the parameters in our model we briefly have a closer look at the random effects of negative selection on the peripheral T-cell repertoire. This is interesting in itself as it should give us more insights into the scope of negative selection

and perhaps lead to new ideas to improve the process (in a biologically meaningful way). For a proper analysis we randomly generated a set of 100000 T-cells (via simple sampling, that is without tilting of the W distributions). These T-cells had then to pass negative selection. We did this for 100 repetitions and compared the sets of surviving T-cells.

These sets differed every time due to the stochasticity of negative selection, but these differences had almost no impact on the estimated post-selection W^{neg} density, that is Figure 6.2 stays almost the same (data not shown). This is also reflected in the real percentage of surviving T-cells. We chose a thymic activation threshold, such that 50% of all T-cells survive. Nevertheless, this is only true for infinitely many T-cells. If we repeat negative selection for a finite set of T-cells the survival percentage varies due to stochastic effects. Luckily, the variance in the percentage of survivors is quite low. Figure 6.10 shows that for all 100 repetitions the survival rate differed only up to 0.2 (swor) and 0.4 (swr) from our aim of a 50% survival rate.

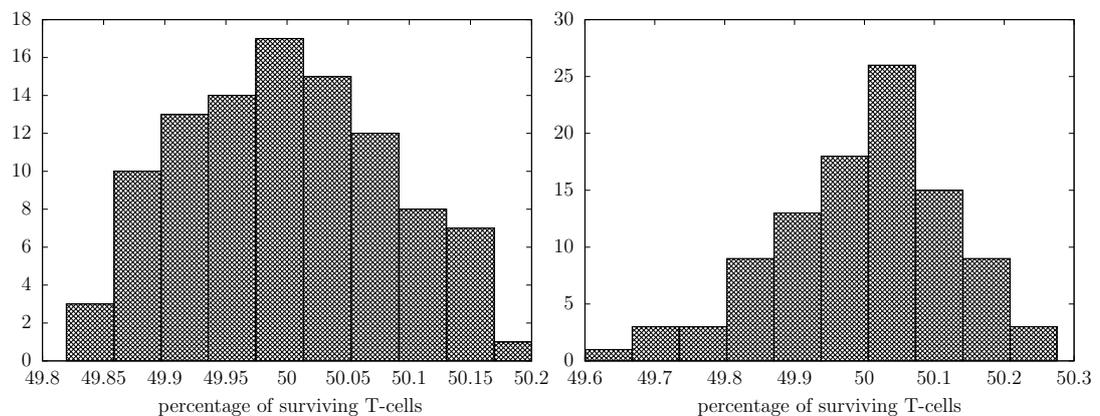


Figure 6.10: The histograms show the fraction of the original T-cell repertoire that survives negative selection (swor (left) - swr (right)): 100 repetitions of negative selection on the same repertoire of 100000 T-cells. For both model variants the differences from the desired survival probability of 50% are very low. Note that for both variants we estimate the threshold g_{thy} only once, that is for a different repertoire of T-cells we use the same threshold. The results which we present here are reproducible for any new repertoire of T-cells that is generated via the same method.

However, with regard to the actual sets of surviving T-cells the situation looks different. By a comparison of the 100 different T-cell sets we calculated that in the mean the sets differed by 14% (swor) and 17% from each other (see Figure 6.11). Taking into account the negligible differences in W^{neg} , it follows that there is a certain interchangeable subset of T-cells whose survival has no impact on the post-selection density. With a view to the foreign-self discrimination capability of the peripheral T-cell repertoire these T-cells pose a potential problem. They can neither be assigned unambiguously to the class of T-cells that are too self reactive nor the class of T-cells that have a low self reactivity. Given the fact that there is such a possible variance in the set of surviving T-cells the question also arises how probable it is that too self reactive T-cells survive negative selection.

With these first observations we have established some basic ideas on the scope and effect of negative selection as well as on the differences and their implications of our

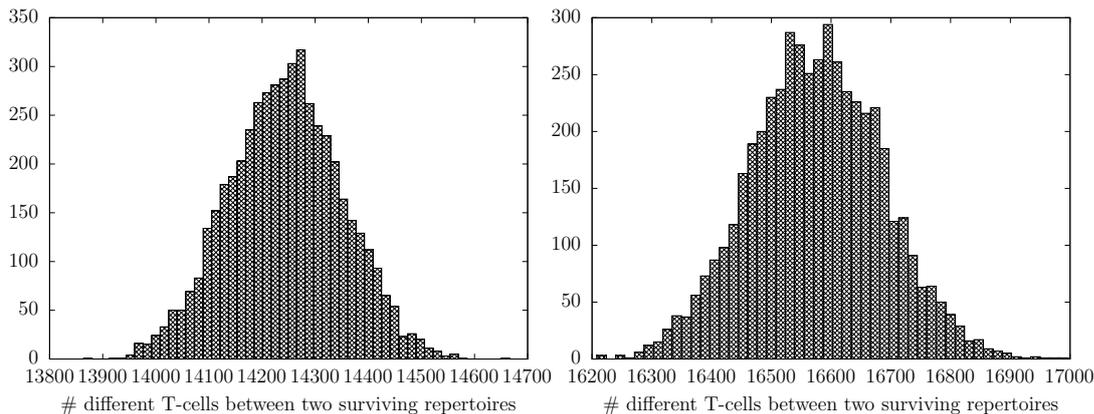


Figure 6.11: The histograms show the number of T-cells that are different between the survival repertoires if we compare them pairwise (swor (left) - swr (right)): 100 repetitions of negative selection on the same repertoire of 100000 T-cells. In the mean two surviving T-cell repertoires differ by about 14% (swor) or 17% (swr). This number indicates that there is a big subset of T-cells which cannot be unambiguously characterised as too self reactive.

two model variants. Next we have a closer look on the inherent interactions between the different parameters that guide our model. We restrict ourselves to a 50% negative selection process, but change either the total number of antigen classes K , the number of different classes n_s per APC or both together. Again, we investigate the activation curves before and after negative selection as well as our estimates of the empirical densities of the stimulation rate W .

1st variant: increasing the total number of antigen classes K

We start by a 10-fold increase of the total antigen class number, that is we now assume that there are $K = 10000$ antigen classes available. It is clear that this change of parameter should not have any effect for our model in the swor variant before negative selection. However, the negative selection process should be affected and hence also the activation curves after negative selection. For the swr variant of our model we expect notable changes also before negative selection as the probabilities to sample some antigen types severalfold decrease notably. Figure 6.12 shows our estimates of the activation curves for both model variants before negative selection.

Evidently, while the activation curves for the swor variant remain unchanged, those of the swr variant change quite a bit. As was to be expected, the activation probabilities approach those of the swor variant. Although the stochastic effects of sampling with replacement still play a role, this role becomes a minor one in comparison to the outcome for our model with the original parameter set.

The next question is again, what happens if we introduce negative selection. A first idea of the effect of this process gives us Figure 6.13 with the post-selection densities.

We can see that the changes in the density of W are much less pronounced than they were for the original parameter set. We have compression factors of 0.92 (swor) and

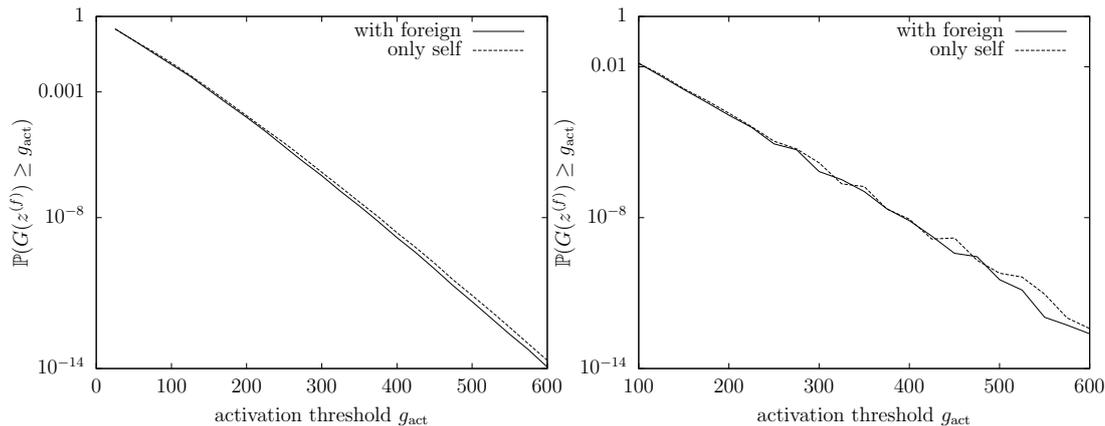


Figure 6.12: Activation curves before negative selection if we assume an antigen set size of $K = 10000$ and keep the number of chose antigen types the same ($n_s = 50$) (swor (left) - swr (right)). There is no foreign-self discrimination possible. The activation curves of both model variants are not only qualitatively but also quantitatively much more similar than for our original parameter set. This is due to the fact that the probabilities to choose antigen types severalfold is much more reduced.

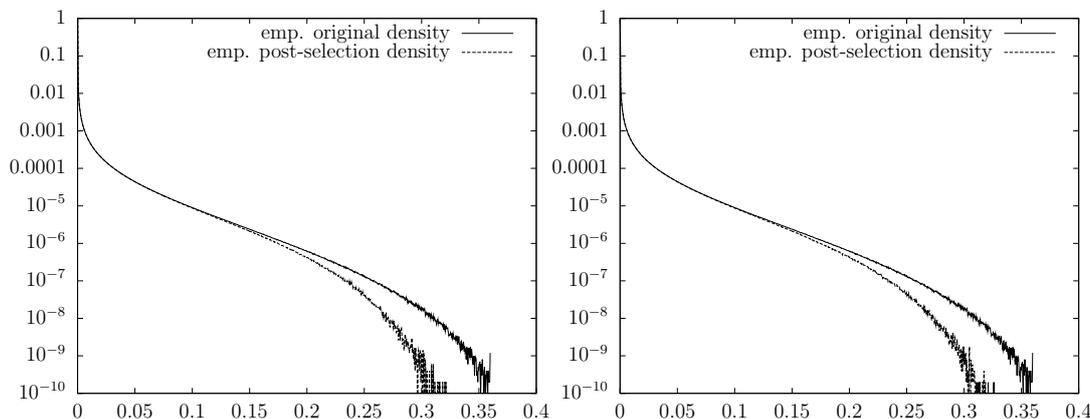


Figure 6.13: Comparison of the empirical post-selection densities with the original density for both model variants and $K = 10000$, $n_s = 50$ (swor (left) - swr (right)). The compression effect of negative selection on the post-selection densities is considerably reduced in comparison to the model with the original parameter set. The compression factors are estimated as 0.92 (swor) and 0.89 (swr).

0.89 (swr). The explanation for this has two reasons. For one it lies in the fact that with 10 times more antigens, the empirical distribution of the stimulation rates of every single T-cell resembles the W distribution much more than before. So even if the T-cells with extremely high stimuli do not survive, the rest is not so different at all. Or to put it another way, if we use some measure on how good a T-cell approximates the W distribution, the variance in the outcome of the measurements would be decreased quite a bit in comparison to the previous case. This is just a consequence of the law of large numbers, by which one can show that if we construct an empirical W distribution from the stimuli vector representing a T-cell, this empirical distribution converges to the original W distribution almost surely for increasing size of antigen classes.

Additionally, we only sample 50 antigens types (in the swr case even less) out of these 10000 and have only 2000 different APC meetings. This implies that some potentially harmful stimuli combinations might be missed during negative selection. Taken together this means, that if we perform the same analysis of the negative selection process as we did before, the differences between the sets of surviving T-cells would be much greater, because the role of pure chance in contrast to the stimuli composure of the single T-cells is much enhanced.

All these considerations might also explain the fact that in the swr case we do not have this special form of post-selection density anymore, where we could recognise a two-phasic behaviour. We supposed that more intermediate strength antigens were sorted out because they appear more often in higher copy numbers. This would be much harder for this parameter set.

The less pronounced effect of negative selection is of course also reflected in the activation curves after negative selection.

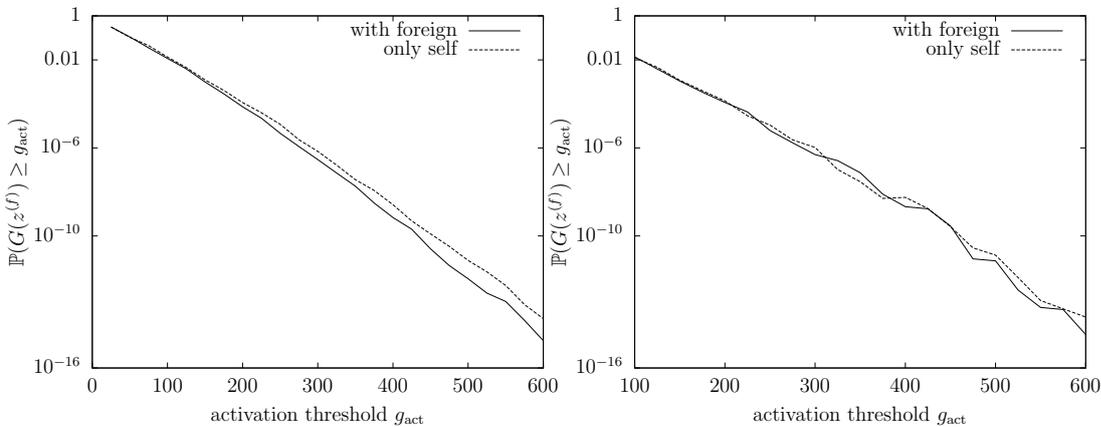


Figure 6.14: Activation curves after negative selection for the parameter set of $K = 10000$ antigens out of which $n_s = 50$ are chosen (swor (left) - swr (right)). These estimates were obtained by the direct inclusion of negative selection in our simulation method, that is we do not make use of the post-selection distribution. For both model variants negative selection does not enhance foreign-self discrimination.

In Figure 6.14 you can see that only the slope of the activation curves change, while the differences between the activation probabilities with or without foreign remain similar. At this point we have to say some words on our estimation of these probabilities, as these also shed some light on the negative selection process again. Here, we chose to only use our estimates from the simulation where we included negative selection into the algorithm, that is we did not use the empirical densities after negative selection. The reason for this is just that we did not have to deal with the drawback of losing too many samples, especially samples that reach g_{act} , during negative selection. In the original setup with increasing threshold values the number of surviving T-cells decreased and even more the number of T-cells that were stimulated to reach the threshold. Now, the situation is different because of the decreased probability to have an APC that presents a notable amount of the antigen types that induce the high stimuli to the T-cell. These high (tilted) stimuli get lost in the background of the rest of the (untilted) stimuli. They

might even be counteracted by very low stimuli that are presented much more often as a consequence of the original W distribution. This illustrates the drawback of negative selection for this parameter setup quite nicely, as obviously antigen combinations that are potentially harmful are often not seen during negative selection, while they have a good chance to appear in the periphery during the lifetime of a T-cell. Consequently, our estimates seem not only to be good enough by our chosen simulation method, they also might reflect the reality a little better. This information would be lost if we just used the empirical density estimate for W after negative selection. Anyway, the results are not expected to be very different, therefore we skip these additional simulations.

2nd variant: increasing the number of presented antigen classes n_s

Instead we change the parameters in the model again. Out of a set of $K = 1000$ antigen types, $n_s = 100$ are chosen to be present on an individual APC. As we keep the total number antigens per APC constant, this implies that the copy number z_s has to be reduced from 500 to 250. From our previous simulations it is clear that we have to expect that the activation curves are decreasing much faster than for the case with $n_s = 50$ as the individual antigens do not contribute as much to the total stimulation rate anymore. Furthermore the possibility for the foreign stimulation rate to stand out against the self background should be reduced even more, even after negative selection. The first assumption is correct as can be seen in Figure 6.15 where we present the activation curves for both model variants before negative selection. Interestingly, this effect also

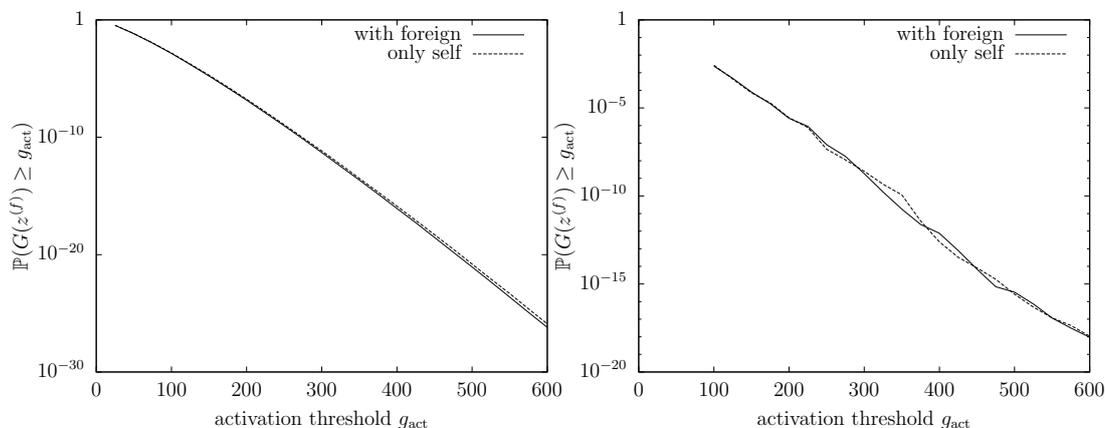


Figure 6.15: Activation curves for our model with $K = 1000$ and $n_s = 100$ (swor (left) - swr (right)). For both model variants the activation curves of 'only self' and 'with foreign' practically coincide. Hence, foreign self discrimination is not possible. Although there is this qualitative similarity, both models are quantitatively very different. The activation probabilities for the swor variant are much smaller than for the swr variant. This is of course due to the fact that choosing antigen types severalfold is quite probable in the swr variant and hence plays a major role.

appears for the swr variant. This implies than an increased probability to choose self antigens severalfold cannot compensate for the effect of reducing the copy number z_s . In contrast to Figure 6.1 where the activation by 'only self' was slightly higher than

activation by 'with foreign', here both activation curves coincide. This is an effect of the central limit theorem. By introducing more antigens and reducing the copy number the total stimulation rate distribution becomes more like a normal distribution, which also implies that individual antigens that have high stimulation rates are compensated by antigens with very low stimulation rates. In fact, because of the W distribution we have an exponentially larger pool of very low stimulation rates to choose from in comparison with few high stimulation rates.

Our second assumption is that negative selection has a minor effect on the foreign-self discrimination capability in this model. Figure 6.16 shows the empirical post-selection densities. We see that again these densities are compressed considerably in comparison to the original density of W . Here, it actually seems to make more of a difference if we use the swor or the swr variant. For the latter we can see that at least for our estimations down to probabilities of 10^{-10} the two different slopes that can be seen in Figure 6.2 do not appear. But it is also evident that the overall shrinking factor is larger (0.5). This

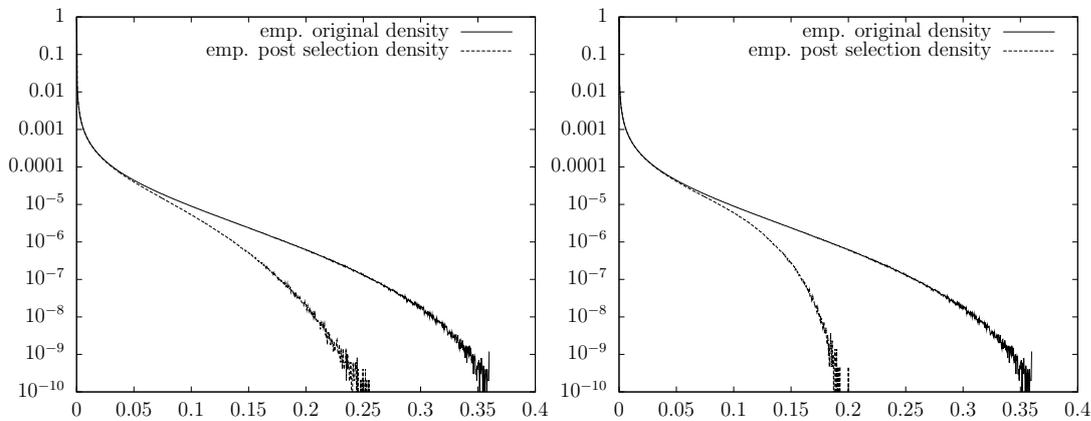


Figure 6.16: Comparison of the empirical original W density with the empirical post-selection density (swor (left) - swr (right)). In both model variants negative selection has a major effect. But there are interesting quantitative differences. Compared to the original density the post-selection density of the swr variant is compressed by a factor of 0.5, which is the best factor we achieved so far. The results for the swor variant are much less promising. The estimated compression factor is only about 0.69.

is different for the swor model. Here, the density is only compressed by a factor of 0.69. One explanation for these findings is that for the swr model choosing 100 antigen types enhances the probability of choosing individual types severalfold is greatly increased. Thereby also stimuli of lower strength can become important as they are multiplied by the multiplicity factor. We assume that this effect is even more enhanced than it is for the original parameter set, such that we do not see the two phases in the density but only have the first slope with the bigger compression factor. Perhaps the kink could be seen for even lower y -values.

For the swor model on the other hand, choosing more antigen types has quite the opposite effect. Negative selection gets worse. The reason for this is that in the model variant we are forced to choose 100 different antigens, which enhances the probability to choose many antigens with low stimuli which counteract some possible high stimulus

antigens.

It is now interesting to see whether these differences in the outcome of negative selection are also reflected in the activation curves after negative selection. Unfortunately, there are no differences (see Figure 6.17). For both cases there activation only by self or with foreign antigen is equiprobable. The explanation for this goes into the same direction in both cases. To get a good foreign-self discrimination the probability of having a high stimulus by foreign has to be higher than the probability to get a similar high stimulus by any of the self antigens or a combination thereof. However, with the possibility of choosing 100 antigens, we might have a low probability to have a T-cell with such a stimulus in the peripheral repertoire, but we can compensate this by choosing several antigens with medium strength stimuli. This is even more true in the swr variant, where we have a quite high probability of choosing antigen types severalfold. Hence, the only effect of negative selection under those circumstances is that the activation probabilities decrease faster than in the case without negative selection.

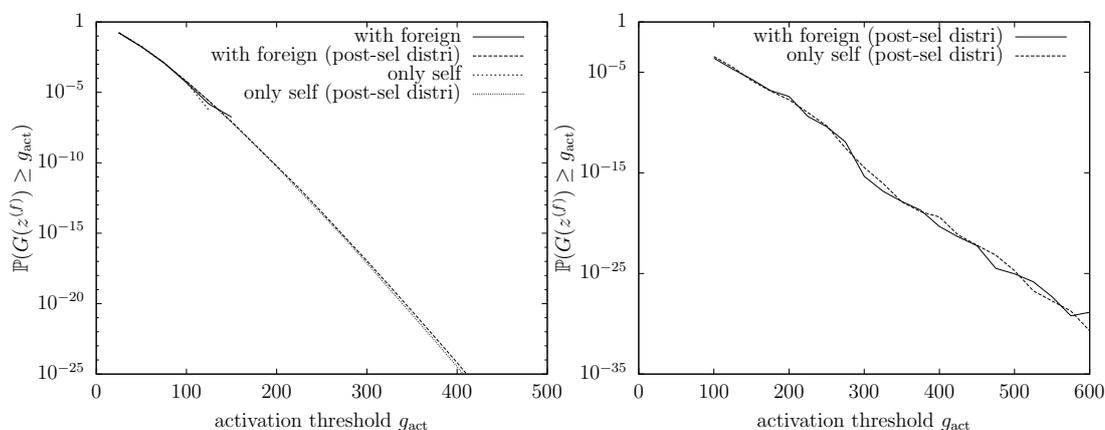


Figure 6.17: Activation curves after negative selection for the parameter set $K = 1000$, $n_s = 100$, $z_s = 250$ (swr (left) - swr (right)). For both model variants foreign-self discrimination is impossible.

3rd variant: increasing the total number of antigen classes K and the number of presented antigen classes n_s

For the sake of completeness we finally combine the changes of both parameters and now assume to have $K = 10000$ antigen types, out of which $n_s = 100$ are chosen. As a consequence of our previous results we have to expect that this parameter combination presents the worst case scenario. Which means similar results for both model variants as the swr variant approaches the swr variant, nearly no effect on the post-selection densities and thereby also no enhancement of foreign-self discrimination capability. Figures 6.18, 6.19 and 6.20 support these expectations.

In summary, we can see that for either model variant a change in the parameters to different (more biologically plausible?) values decreases the, even before, almost negligible foreign-self discrimination capability of the model. This leaves us with different

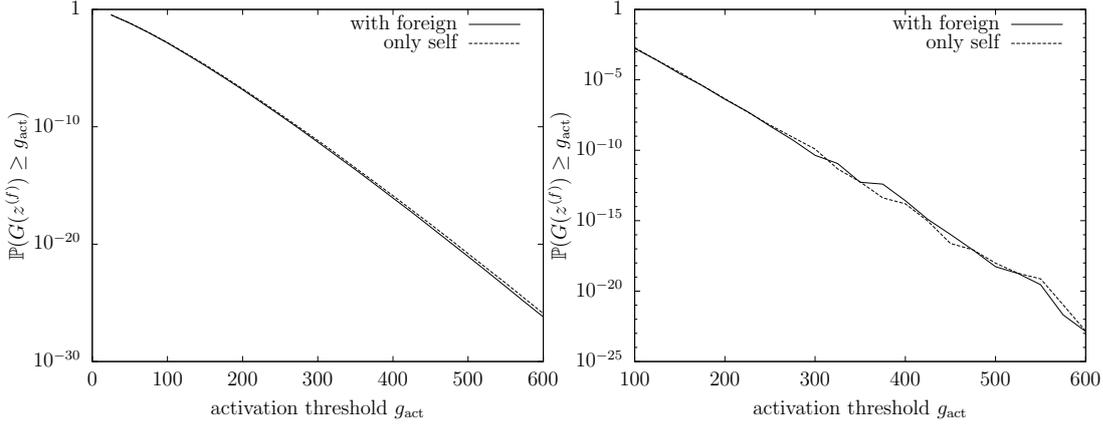


Figure 6.18: Activation curves for $K = 10000$ antigens out of which $n_s = 100$ are chosen (swor (left) - swr (right)). For both model variants there is no difference in activation probability. Hence, foreign-self discrimination is not possible.

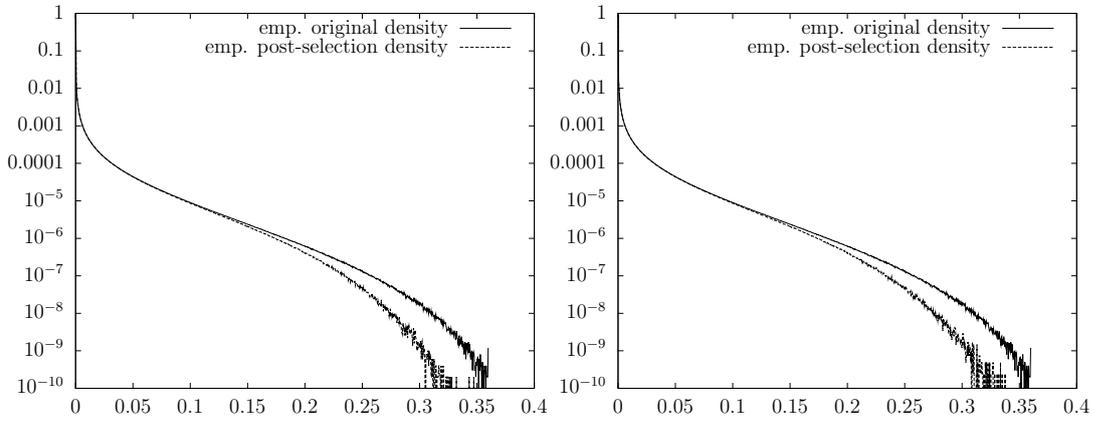


Figure 6.19: Empirical densities before and after negative selection for the model with $K = 10000$ and $n_s = 100$ (swor (left) - swr (right)). There is no noticeable difference between both model variants.

possibilities to move on, some of which we will mention in the discussion in the next section. However, we chose to pick out one before, as it turns out that this has quite an impact on foreign-self discrimination without the need to change the model any further and therewith the simulation method.

4th variant: changing the parameter $\bar{\tau}$ of the exponentially distributed binding time

Throughout this thesis we stuck to the original parameter for the heart of the BRB model, the exponential distribution describing the binding event between antigen and T-cell receptor (see page 32). But there are of course other values thinkable for this distribution than the one we adopted for our model. So, we now change this parameter from $\bar{\tau} = 0.04$ to $\bar{\tau} = 0.03$. Figure 6.21 shows that before negative selection we have the same qualitative behaviour as we have for the original parameters. For both model

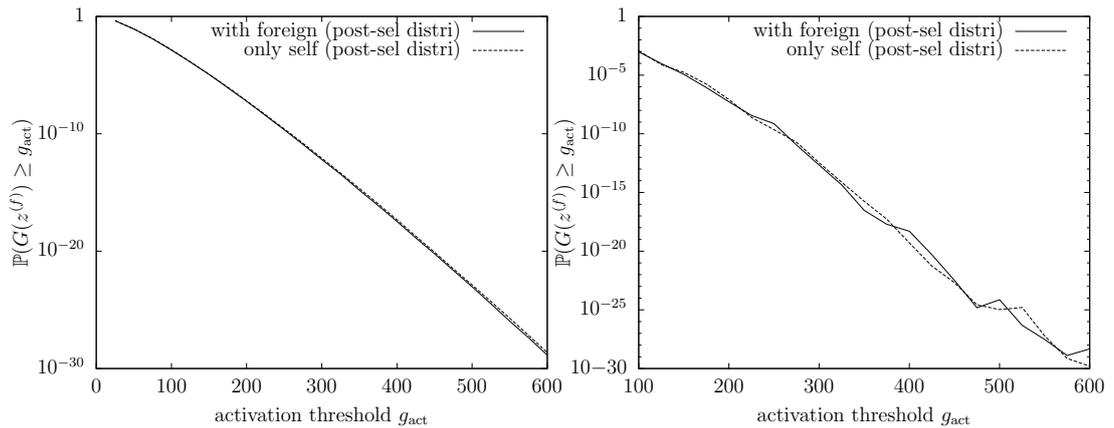


Figure 6.20: Activation curves after negative selection for $K = 10000$ and $n_s = 100$ (swr (left) - swr (right)). Activation with and without foreign antigen present has the same probability. This is true for both model variants. Furthermore even the actual probabilities are very similar in both variants given the same activation threshold.

variants there is no foreign-self discrimination possible.

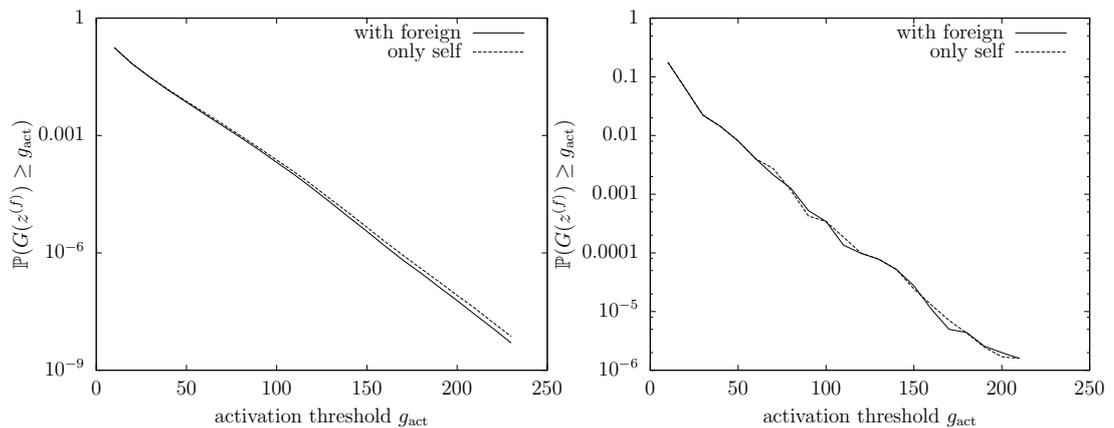


Figure 6.21: Activation curves before negative selection for $K = 1000$, $n_s = 50$ and $\bar{\tau} = 0.03$ (swr (left) - swr (right)). Foreign-self discrimination is not possible.

However, if we come to negative selection the situation changes drastically. This can be already deduced from the post-selection densities (see Figure 6.22). In comparison to the original density the post-selection densities are compressed considerably. For the swr variant we estimate a compression factor of 0.25 and for the swr variant we estimate a compression factor of 0.31. Already from these numbers we can assume that this has a noticeable effect on the foreign-self discrimination capability in the model. This is confirmed by our estimates of the activation curves after negative selection (see Fig. 6.23).

For both variants activation by foreign is much more probable than activation without foreign. These differences suffice to enable foreign-self discrimination.

A comparison of the densities in Figure 6.22 and Figure 6.2 reveals why negative se-

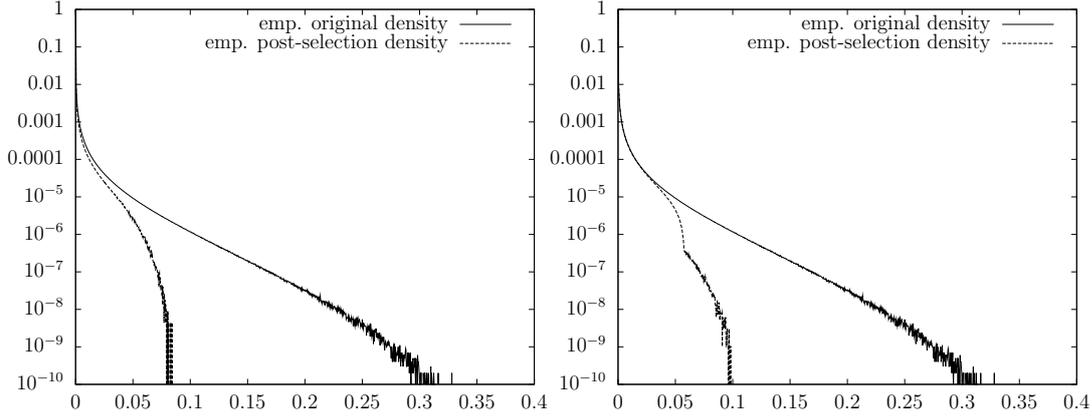


Figure 6.22: Comparison of the densities before and after negative selection in the case of $\bar{\tau} = 0.03$ (swor (left) - swr (right)). We have a much bigger effect of negative selection. For both model variants the shrinking factors are considerably higher than for the original parameters. (compression factors: 0.25 (swor) and 0.31 (swr)).

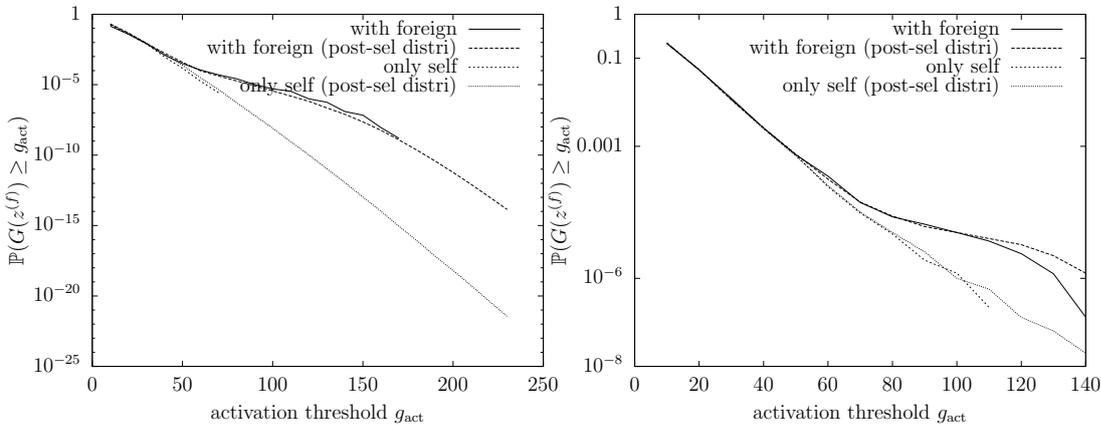


Figure 6.23: Activation curves after negative selection for $\bar{\tau} = 0.03$ (swor (left) - swr (right)). For both model variants we see a big difference in activation probabilities with and without foreign antigen present. These differences are great enough to finally enable foreign-self discrimination.

lection is doing a much better job here than for the model with the basic parameter set. The long tail of the pre-selection density with $\bar{\tau} = 0.03$ is thinned out considerably in comparison to the pre-selection density with $\tau = 0.04$. This enables negative selection not only to cut off the part of the tail with the high stimulus events (as was already possible before) but also the part of the tail with the intermediate stimulus events. Before, it was quite probable that several self antigens with intermediate stimuli are presented together such that, if added together, they exceed a high foreign antigen stimulus. Thereby foreign-self discrimination is obscured. This probability is now significantly reduced.

6.4 DISCUSSION

In this section we discuss the results of the previous section and formulate consequences.

The unquestionable main result of our simulation is the fact that our new model as such is not capable of foreign-self discrimination in our tested parameter range if we do not change the parameters of the original BRB model. With the basic and extended BRB models in mind one could argue that an increase in the copy number of foreign antigens could improve the situation. This is certainly right. However, our restriction of the same copy number for foreign and self antigens is carefully considered. With this new model our intention was to become biologically more plausible. As already mentioned before, T-cell activation is possible for very low numbers of foreign antigen copies. A good T-cell activation model should therefore be capable of explaining foreign-self discrimination in such settings. Hence, if anything our restriction should be much harder, allowing less foreign antigen copies than the other way around. Another important argument in this context is discussed at the end of this section.

However, there are different other starting points to change our model in order to improve foreign-self discrimination. Before changing something in the model itself, it is necessary to revisit all details of the existing model. Here, this is, above all, the stimulation rate distribution which is the key element of the T-cell activation model. It starts with the question if the exponential distribution is really the best description for the dissociation process. If so, another question is whether the parameter for the underlying exponential distribution describing the dissociation probability of a TCR and a pMHC can be changed and if this has any effect. Already in the original paper of the BRB model a lognormal distribution is proposed as it can be motivated by the Arrhenius law [205]. Zint et al. could show that even without negative selection this distribution leads to better discrimination capabilities for smaller z_f [232]. Our change of the parameter $\bar{\tau}$ therefore had two different motivations. First of all there is no experimental evidence that only points to the original factor of $\bar{\tau} = 0.04$. The resolution is simply not sufficient. Moreover, many models, that assume a Bernoulli distribution instead of our W distribution, estimate the probability of a T-cell recognising an antigen to be in the range of $10^{-5} - 10^{-4}$ [127]. If we compare this with our original W distribution, it suggests that the tail of the W distribution is too thick. Secondly, our change to $\bar{\tau} = 0.03$ leads to a stimulation rate distribution that is quite similar to the stimulation rate distribution resulting from an underlying lognormal distribution. Therefore, we could explore this idea without having to recalculate too much for our simulations. Although this change of parameter seems to be a minor one, the effect on foreign-self discrimination is a major one. This is the first time we really have a model that is capable of foreign-self discrimination where all the antigen types appear with the same copy number (swor variant) and even better, if there is the possibility to have more copies of some self antigens (swr variant).

There are of course also changes in the equation $w(\tau)$ thinkable. First of all we have to note, that here changes are only relevant that concern the leading, increasing part of $w(\tau)$. The other part barely plays any role because of the very fast decreasing exponential distribution of \mathcal{T} . All these changes have to be supported by experimental evidence if possible. One recent promising result in this context is a remodelling of the underlying kinetic-proofreading concept. In the original model a TCR is dephosphorylated instantly after the dissociation of a pMHC. However, if a short time-delay is introduced and the

briefly dissociated pMHC molecule has a high association rate it probably rebinds to the TCR and could complete the phosphorylation chain in order to trigger the TCR [62].

Recent experiments have also revealed another important mechanism. It seems that a T-cell is able to integrate activation signals over several APC meetings [85]. This opens up a new interesting possibility to change our model. One interpretation of signal integration is the introduction of a new sum of signals over different APC meetings of a single T-cell. Thereby, we would cover all three levels of interactions that could influence T-cell activation, as demanded for example in [35]. We have the first level of a single TCR that meets a pMHC, the second level of all TCRs on a T-cell and all pMHCs on an APC and finally the third level of one T-cell and several APCs. We tested this idea in a model where we just assumed that a T-cell meets k APCs and is totally activated if it is activated by l of these APCs. Some first results obtained by simple sampling and a T-cell repertoire of 20000 T-cells without negative selection are shown in Figure 6.24. We estimated the activation probabilities if a T-cell meets 10 or 20 APCs in a

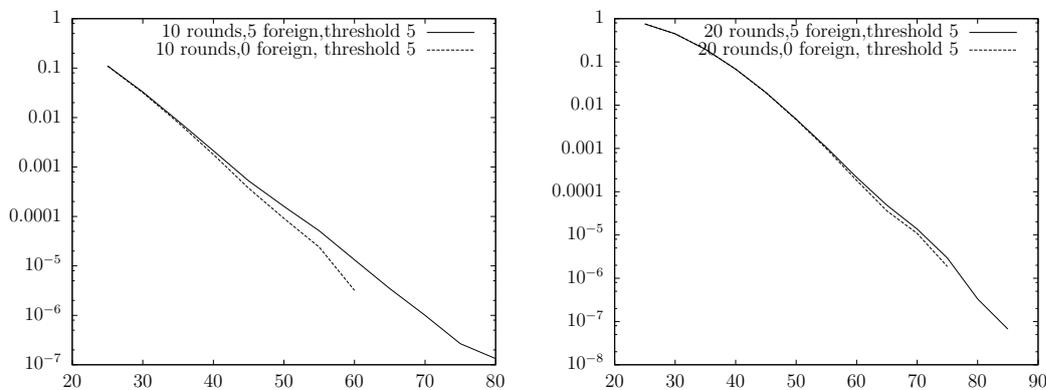


Figure 6.24: A simple sampling estimation of activation probabilities for a T-cell repertoire of 20000 T-cells under the signal integration model. For the left graph we allowed for 10 APC meetings out of which 5 carried the same foreign antigen type and a T-cell was fully activated if it was activated in 5 APC meetings. For the right graph we changed this parameters to 20 APC meetings.

row and integrates the single activation signals which might occur as a consequence of a T-cell-APC meeting. Five of these APCs were equipped with the same foreign antigen type. The other APCs presented only self antigens. We assumed that the T-cell is totally activated if it has received 5 activation signals. In both situations this T-cell repertoire that has not been negative selected shows signs of foreign-self discrimination capability. This is of course a very crude model with flaws. A more realistic model could look more like a queueing model, where for every time step (or every APC meeting) one signal leaves the queue. The queue length here would be 10 or 20 and we have two different signals, activation or not activation. If there are 5 or more activation signals in the queue the T-cell is finally activated. However, even our model gives us a first idea of the potential of signal integration.

Another open question is how we can change negative selection such that its outcome does not vary so much and is, if possible, the optimal peripheral T-cell repertoire. Some preliminary results indicate that the introduction of signal integration at least reduces

the variance in the outcome if we run the same original T-cell repertoire through negative selection several times and compare the resulting sets. In order to make more improvements it is necessary to understand this process better both experimentally and theoretically. This is the topic of the next chapter.

Furthermore there is one interesting point to learn from the empirical post-selection densities in the last chapter for both model variants. Evidently, it seems to lead to a more effective T-cell deletion during negative selection if we have the possibility to present antigen types severalfold. This seems to be the case because it puts more selection pressure on stimuli of intermediate strength. A thymic environment which favours such an antigen expression should therefore be favourable with respect to our model.

It is of course also possible that the idea of TCR triggering as used in the thesis is wrong. Another approach that has recently become popular is the triggering of TCRs by receptor deformation [123, 124, 197]. One idea in this context is that the TCR is triggered by the pulling force that is induced by the movements of the APC and the T-cell on a TCR-pMHC binding. It is thinkable to adapt serial-triggering and our type of mixture models in order to capture this assumption.

Finally, we have to discuss one open question that has not been investigated in this thesis and the modeling literature it is based on. In all these mixture models we define foreign-self discrimination by the difference between activation curves for activation with and without foreign antigen. Thereby, we ignore one very crucial point. If activated a T-cell attacks or helps to attack the pathogen that is the source of its cognate antigen. It is therefore necessary that the T-cell receives the strongest activation stimulus from the foreign antigen otherwise it would cause an autoimmune reaction. This seems obvious. However, the way we defined T-cell activation does not capture this. We have to analyse in our models if, when a T-cell is activated, the foreign stimulus really was the strongest. It would be even better if this stimulus would be much stronger than any other self stimulus. This is especially true for the basic BRB model, but also for the extended BRB models because here the stimulus intensity is intensified by raising the foreign antigen copy number z_f .

In summary the analysis of our new model brought us many new insights on the effect of negative selection in mixture models. We could describe its flaws and benefits and finally present a parameter set for which we could show foreign-self discrimination for biologically more plausible parameters than in all models before. This is a new result which underlines the strength of this type of model. It is clear that this is only the starting point for further investigations. These include the need for better experimental estimates on all the different parameters and many different possibilities to expand or edit the model to incorporate new ideas such as presented in this section.

A MODEL FOR T-CELL MIGRATION IN THE THYMIC MEDULLA

The important site for T-cell development is the thymus. In this organ every T-cell develops its unique receptor and has to survive different selection processes. We already highlighted the whole developmental process in the background section 2.2. There and also in the other sections of this thesis we pointed out the importance of the negative selection process as the supposed key element in the creation of a peripheral T-cell repertoire that is anergic if challenged with self antigens but reacts on an encounter with foreign antigens. Despite its importance much of the experimental and theoretical research is very recent and there are still many things unknown. One prominent example for this is the discovery of the promiscuous gene expression mechanism for tissue restricted antigens [107]. This discovery directly affects models that try to explain foreign-self discrimination and tolerance induction in the periphery. Generally, it is insufficient to include negative selection as a kind of black box that somehow shapes the T-cell repertoire. Models that try to explain T-cell activation in the periphery should also be applicable to the negative selection process and with this two key ingredients put together the foreign-self discrimination capability of peripheral T-cells should be elucidated. A general assumption in many T-cell activation models is that activation depends on the binding time between a TCR and a pMHC molecule and the TCRs of all T-cells that survived negative selection have short binding times to all self antigens. A crucial test for such a model is therefore to reproduce this assumption when the T-cell activation model is applied to the negative selection process. In the last chapter we already did this with our new model and could show that the outcome is very much dependent on certain parameters. In this model we included a very simplistic form of negative selection. Although we assume that changes in the negative selection process cannot improve the situation too much it would be preferable to have a more detailed and realistic model of negative selection to totally exclude this argument.

In this section we therefore start to develop a model of T-cell migration in the thymic medulla. These dynamics primarily include T-cell movement and interaction with dendritic and medullary epithelial cells. The reason for this is that they are the key-players in the negative selection process and also constitute the majority of cells in the thymic medulla. In its basic features our model shares similarities with a model of T-cell movement and decision making in a lymph node from Zheng et al. [231] and is different from a much more detailed model of T-cell migration in the lymph node based on the cellular Potts model [14, 15, 13].

For reasons of computational tractability we model T-cell movement in the medulla

not in a continuous 3-dimensional space but on a 3-dimensional lattice. Every node in the lattice can either be uninhabited or occupied by a T-cell, a DC or an mTEC alone or occupied by a DC or mTEC together with a T-cell. In a first version we even allow the occupation of a node by more than one T-cell. We furthermore restrict our simulation to a small section of the thymic medulla, again, due to computational tractability but also because information might be lost in a too big and complicated model. Primary, in our very basic model we are only interested in the number of mTECs and/or DCs an individual T-cell meets while migrating through the medulla. We can also equip the mTECs with tissue-restricted antigens. Thereby we can get some important information such as the probability of a T-cell to explore all tissue-restricted antigens if these are only presented by mTECs. In the course of modeling we omit any T-cell activation and negative selection mechanisms. These can be introduced at a later stage.

A model as we have in mind needs some basic parameters and assumptions. We have to create a realistic 3-dimensional setting, that is we have to define the number of nodes and their distance in our lattice and need to put the appropriate numbers of DCs and mTECs in an appropriate spatial distribution on our lattice. Furthermore, we have to describe the T-cell movement and the event of a T-cell-DC/mTEC meeting. As most of the experimental data is based on research on mice we built up an *in silico* section of a murine thymic medulla.

Recent multi-photon microscopy experiments on living mice or explanted organs helped to characterise T-cell movement at least in lymph nodes and the thymic cortex [134, 135, 136, 11, 24]. As there is no reason to assume that T-cell movement in the thymic medulla differs too much from the movement in these other tissues we use their results. It was determined that T-cell movement follows a random walk, where a T-cell moves into the same direction on the order of $10\mu\text{m}$ and then changes its direction randomly. The average speed was determined to be $9 - 12\mu\text{m}/\text{min}$ [133, 134]. This leads to a straightforward definition of our lattice. The distance between two nodes is just this $10\mu\text{m}$ and one time step in our model is 1 minute. We chose to use two different lattice sizes ($300 \times 300 \times 300\mu\text{m}$ and $500 \times 500 \times 500\mu\text{m}$ which corresponds to $30 \times 30 \times 30$ or $50 \times 50 \times 50$ nodes), for two reasons. On the one hand simulations with these two lattices should lead to similar results, otherwise there are inconsistencies in the simulation. The smaller lattice which can be used for much faster simulations can therefore be used as a comparison with the bigger lattice. On the other hand there are more modeling possibilities in the bigger model. If we start to expand the model, we could, for example, introduce a more directed fashion of T-cell movement as a new assumption. This would lead to difficulties for the smaller model variant as we assume periodic boundary conditions in our model. If a T-cell leaves our lattice it enters again on the other site. Zheng et al. argue that this kind of lattice presentation does not lead to artifacts on long time and length scales [231]. This condition is certainly met with our lattice size and a simulation period that takes about 5 days, the time a T-cell spends in the medulla during negative selection.

In order to create a realistic biological setting we cannot just put mTECs and DCs randomly into our lattice, representing a section of the medulla. Instead, there are distinct mTEC and DC areas [162, 199, 18]. For the number of mTECs and the number

of mTEC areas we take the numbers from section 2.2.4. We randomly create these mTEC areas and fill the rest of the grid with the appropriate number of DCs. As we only simulate a section of the murine thymic medulla we had to obtain estimates of all these numbers. In order to get suitable numbers we had to use different sources of experimental data and merge them. As we could not obtain exact data on the size of the murine thymic medulla we used data obtained from dissections of the human thymus and assumed similar proportions for the murine thymic medulla. It is important to note that we assume to simulate a medulla of a young individual as the thymus size changes much with ageing. From experimental data in [186] we estimate the size of the human thymic medulla for a child of 1 – 10y as $3 - 5\text{cm}^3$. The weight of the thymus at that age is between 22 – 30g [186, 31]. The weight of the murine thymus is about 100mg in [98] and 50mg in [148, 99]. We thus calculated the size of the murine thymic medulla as $0.01 - 0.02\text{cm}^3$.

In a next step we had to calculate the number of DCs and mTECs in our section of the medulla. The smaller section fills an area of 0.000027cm^3 , the bigger one fills 0.000125cm^3 . The number of DCs per 10mg thymus is estimated to be $1 - 5 \cdot 10^5$ [161]. Most of these DCs are in the medulla. We estimate this number as $4 \cdot 10^5$. It follows that in our smaller section there are 5400 and in the bigger 25200 DCs. These will be randomly placed on the free nodes after the mTECs are placed.

For the mTECs we have a different situation. These form clusters, called mTEC islets and these islets again form bigger clusters called mTEC areas. The number of mTECs in the medulla is estimated to be in the range of 300000 and there are about 1800 mTEC islets [162, 199, 18]. Thus, there are about 166 mTECs per islet. The numbers from the literature seem to be much smaller (5 – 45 mTECs), but we have to be careful, since these were estimated for 2 dimensions (Rodewald, personal communication).

The (2 dimensional) size of an mTEC islet is measured to be between $60 \times 40\mu\text{m}$ and $170 \times 170\mu\text{m}$. We here assume a general diameter of $100\mu\text{m}$. An mTEC islet is modeled as a cuboid. This fits better with the lattice model and a simple calculation shows that assuming the islets as a sphere would not allow for enough space for 166 mTECs given an mTEC radius of $10 - 15\mu\text{m}$. Hence, the volume filled with islets is about 0.0018cm^3 , which is 9 – 18% of the total volume of the medulla. For our smaller medullary section this means that we have a total mTEC area of $2.43 - 4.86 \cdot 10^{-6}\text{cm}^3$. It follows that we have 2 – 5 mTECs islets filled with 330 – 830 mTECs. The total mTEC area thus shapes a $14 \times 14 \times 14 - 17 \times 17 \times 17$ sub-lattice of our small lattice. For our bigger medullary section the mTEC area is about $1.12 - 2.25 \cdot 10^{-5}\text{cm}^3$. Hence, there are 10 – 23 mTEC islets filled with 1660 – 3818. This translates into a sub-lattice of $23 \times 23 \times 23 - 28 \times 28 \times 28$. If we compare our calculated numbers of mTECs per section and medullary volume to a simple percentile calculation given the total amount of mTECs and the ratios of total medullary volume to the volume of the section the results are very much comparable. Therefore our assumption of mTEC islets as cuboids seems to be reasonable. All important numbers are summarised in tables 7.1 and 7.2.

	small lattice	big lattice
mTECs	330	1660
mTEC islets	2	10
sub-lattice size	$14 \times 14 \times 14$	$23 \times 23 \times 23$
DCs	5400	25000

Table 7.1: Estimated parameter values for a medulla size of 0.02cm^3 .

	small lattice	big lattice
mTECs	830	3818
mTEC islets	5	23
sub-lattice size	$17 \times 17 \times 17$	$28 \times 28 \times 28$
DCs	10800	50000

Table 7.2: Estimated parameter values for a medulla size of 0.01cm^3 .

7.1 SIMULATION METHOD

For the simulation of T-cell migration in the thymic medulla we have to consider two different steps. In a preprocessing step the artificial medulla is generated and afterwards the main procedure, the T-cell migration simulation, starts.

At first we randomly distribute the mTECs over the sub-lattice representing the total mTEC area in our artificial medullary section. We then assume that 3 mTEC islets form an mTEC area and divide the sub-lattice into equally sized parts. These parts are the actual mTEC areas and are randomly placed into our model. It is not allowed to have more than one mTEC per node. Hence, we check for this condition and place all mTECs that do not fulfill this condition randomly somewhere else on the lattice. Afterwards all dendritic cells are also placed randomly such that no node is filled with two DCs or an mTEC and a DC.

There are of course different ways of generation thinkable. We could try to place the mTEC areas not just randomly but by trying to reproduce microscopic images of medullary sections from experiments [162]. However, as we do not know if there is a specific structure involved in the development of the medulla and we furthermore do not assume any specific position of our section in the medulla, the random arrangement is justified. The same is true for the replacement of mTECs that share a node with another mTEC. These are placed randomly everywhere and not only in an mTEC area. This mirrors the fact that there may be very small mTEC spots or single mTECs that are not in an mTEC area. A restriction to an mTEC area is of course also possible as well as a replacement only to the borders of the mTEC areas, whereby such an area would lose its artificial cuboid form.

The main step of the simulation, the T-cell migration, is relatively simple to establish. A T-cell can be in one of two different modes. In the 'scanning mode' it can either share a node with a DC/mTEC or be at a node next to a DC. A T-cell can scan a DC at neighboring nodes because of the long dendrites of a dendritic cell. In our first primitive implementation we do not superimpose any kind of scanning and activation mechanism

but only assume a mean scanning time of 3min, which is the mean binding time between a T-cell and a DC during stage 1 in a lymph node [85, 231]. Thus, for every time-step we generate a Bernoulli random variable with mean 0.25 and if it is 1 the T-cell moves on (if T is the waiting time to generate a 1 then $T - 1$ is geometrically distributed and one can show that $\mathbb{E}(T) = \frac{1}{0.25}$). We have to point out that in the literature one can find different values for this mean binding/scanning time and this question is currently under investigation [14, 15, 13]. An exponential or Poisson distribution might be more realistic than the Bernoulli distribution. The other mode of a T-cell is the 'movement mode'. The T-cell is not bound to any cell and moves freely from node to node. This is implemented by randomly choosing one of the three coordinates of the T-cell and randomly in- or decrementing this coordinate by 1. Thereby, a kind of random walk is established. As mentioned before we assume periodic boundary conditions. If a T-cell leaves the lattice on one side it enters the lattice on the opposite.

We now have a basic T-cell migration model for the thymic medulla at hand. It is easily possible to change this model by changing the mean scanning time of T-cells, change T-cell movement from totally random to perhaps a directed random walk, change the settings for the environment or allow the simultaneous scanning of several DCs/mTECs by one T-cell. It is also easy to extend the model by for example introducing a special scanning and/or activation model and/or equip the mTECs and DCs with self antigens.

7.2 RESULTS

Although the model is very basic, there are already questions that can be asked and answered by it. We can compare our artificial medullary section to real microscopic images of medullary sections from experiments and we can estimate the number of (different) DCs and mTECs a single T-cell meets during a 5 day cycle of migration. We also introduce a first extension of the model by equipping the mTECs with tissue restricted antigens. This allows us to test if the mTECs alone are sufficient to guarantee that a T-cell meets all these TRAs during negative selection. This is crucially important because otherwise autoimmune reactions are quite probable as shown in several experiments, see for example [117, 194, 199].

In our model we did not specify an environment in the medulla but generate it randomly under some constrains. In Figure 7.1 one such realisation is shown from two different perspectives. We do not show the dendritic cells as these are distributed uniformly over the free nodes, but concentrate on the mTECs. It is evident that the desired structure of different mTEC areas and rare single mTECs outside these areas is generated. This kind of visualisation gives a good overview over the three-dimensional composition of our artificial medullary section. However, it is not suitable for comparisons with results from experiments. In experiments the results are visualised by microscopic images of thin slices of the thymic medulla. Therefore, we imitate this procedure by keeping one coordinate constant and visualising the other coordinates in a 2D image as can be seen in Figure 7.2 for different constant x_1 coordinates. A comparison of our four example images with images from [162] is difficult. If we take Figures 1g and 1j from [162] we get

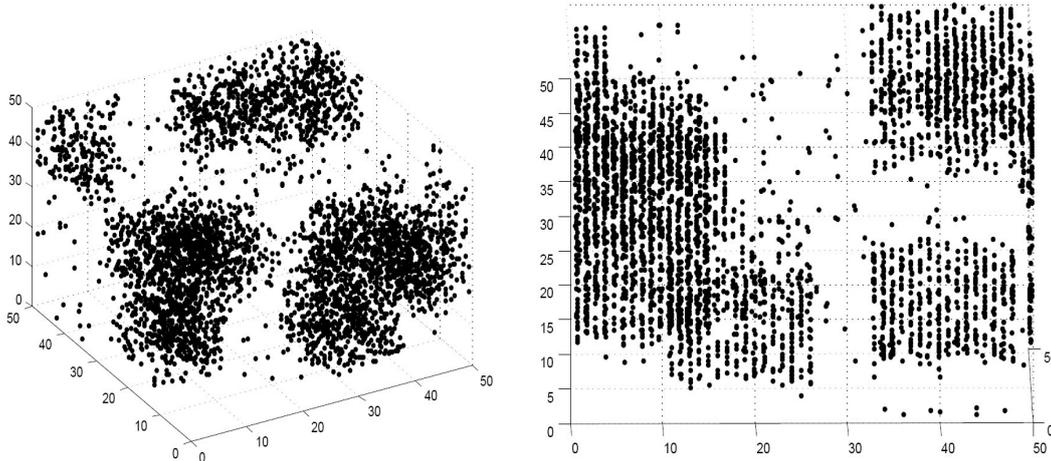


Figure 7.1: Example of a medullary section generated randomly by our simulation from two different perspectives. Only the mTECs are shown.

a picture of the mTEC/DC composition of a medullary area in two dimensions. This looks roughly qualitatively similar to medullary areas in our images. For future work a more specific comparison with experimentally generated images is of course desirable but for a first model our generated environments seem to be sufficient.

In the second step we introduce the T-cells. These are randomly placed on the lattice and start to move in a random fashion when the simulation is started. Figure 7.3 shows the movement of a T-cell during 7200 time steps. This corresponds to a migration time of 5 days. For reasons of visibility we again did not plot the dendritic cells. The right figure is a zoomed-in version of the left figure. There, we visualised the points where a T-cell was connected to an mTEC/DC for 3 time steps or longer by circles. The figures show the random walk like movement of the T-cells for the timesteps where it is not bound to mTECs or DCs as well as the jumps at the boundaries because of the periodic boundary conditions. These assumptions are therefore met and the simulation behaves in the intended way. Although the dendritic cells are not shown it is evident from the circles that a T-cells does not move freely very much but is bound to either DCs or mTECs. The T-cells seem to use the time they have efficiently. Long times of movement without DC/mTEC encounters would lead to a very inefficient negative selection process. T-cell migration and the development of the medullary microenvironment should be guided in a way to enable as many encounters as possible such that a T-cell sees as many self antigens as possible. However, we can also see that a T-cell revisits some positions and thereby some mTECs or DCs. At first glance this seems inefficient, but if we assume that a T-cell only scans parts of a DC this looks different. Most probably the T-cell just sees another part of the self antigen repertoire on the DC surface.

We repeated our simulations for 10 randomly generated microenvironments and 100 randomly placed T-cells and estimated the number of mTECs/DCs a T-cell encounters. The results can be seen in Figure 7.4. We compare the number of mTEC/DC hits for the four different scenarios emerging from taking either the big or small lattice and

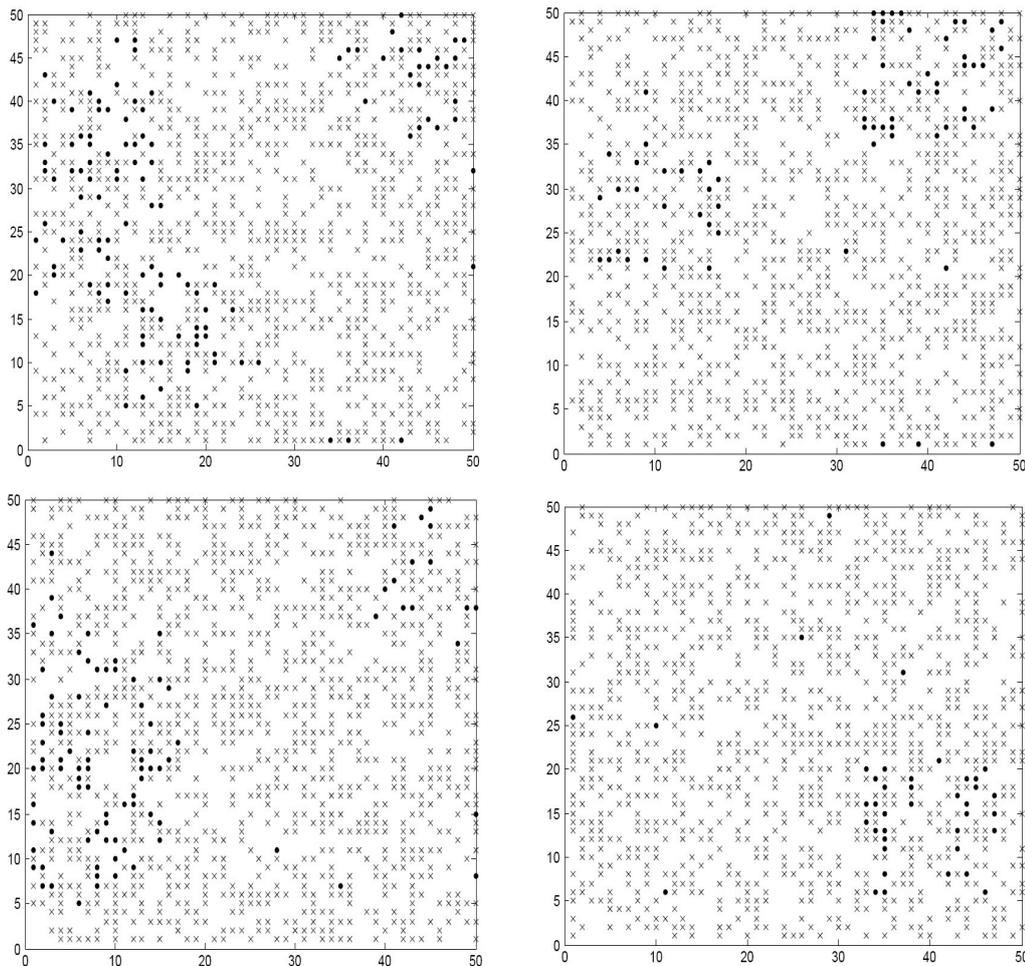


Figure 7.2: Four cuts through the example model in Figure 7.1. The cuts were done at $x_1 = 10, 20, 30, 40$. The black dots are again the mTECs, whereas the black crosses are dendritic cells. A comparison to the Figures in [162] is difficult. If we take Figures 1g and 1j together we get a picture of the distribution of mTECs and DCs in an islet and this seems qualitatively similar to mTEC islet regions in our figures.

assuming a total medulla size of 0.1cm^3 or 0.2cm^3 . The histograms show that in all cases the number of hits seem to approach a Gaussian curve. The only difference is the mean of this curve, which is for the smaller medulla about 1800 and for the bigger medulla about 1700. This is an interesting result, because by doubling the size of the medulla, the number of mTECs and DCs are halved for our section. It follows, that although our lattice is less crowded with mTECs and DCs, it is crowded enough such that a single T-cell meets only 100 fewer mTECs/DCs. Furthermore, we see that the maximal average number of hits, 2400, is never reached in our small test simulations. This might be important for estimations in models of negative selection. For our model in the previous sections we assumed for example 2000 APC meetings during negative selection. Another lesson to learn from the histograms is that the results for the small and the big lattice are not too different. Therefore, the smaller lattice, which leads to a

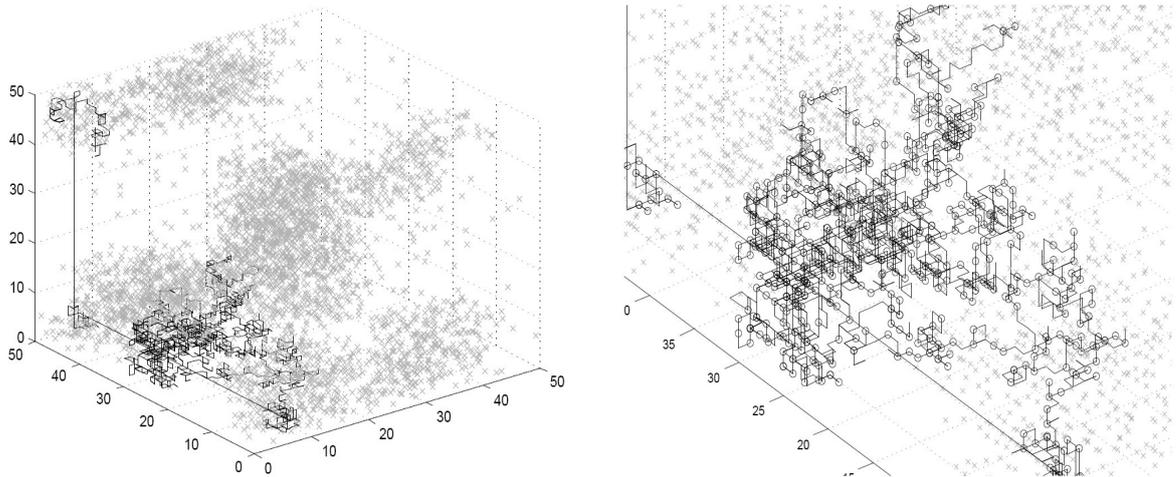


Figure 7.3: Visualisation of one T-cell movement trajectory through our artificial medullary environment. The right picture is a zoomed version of the left picture.

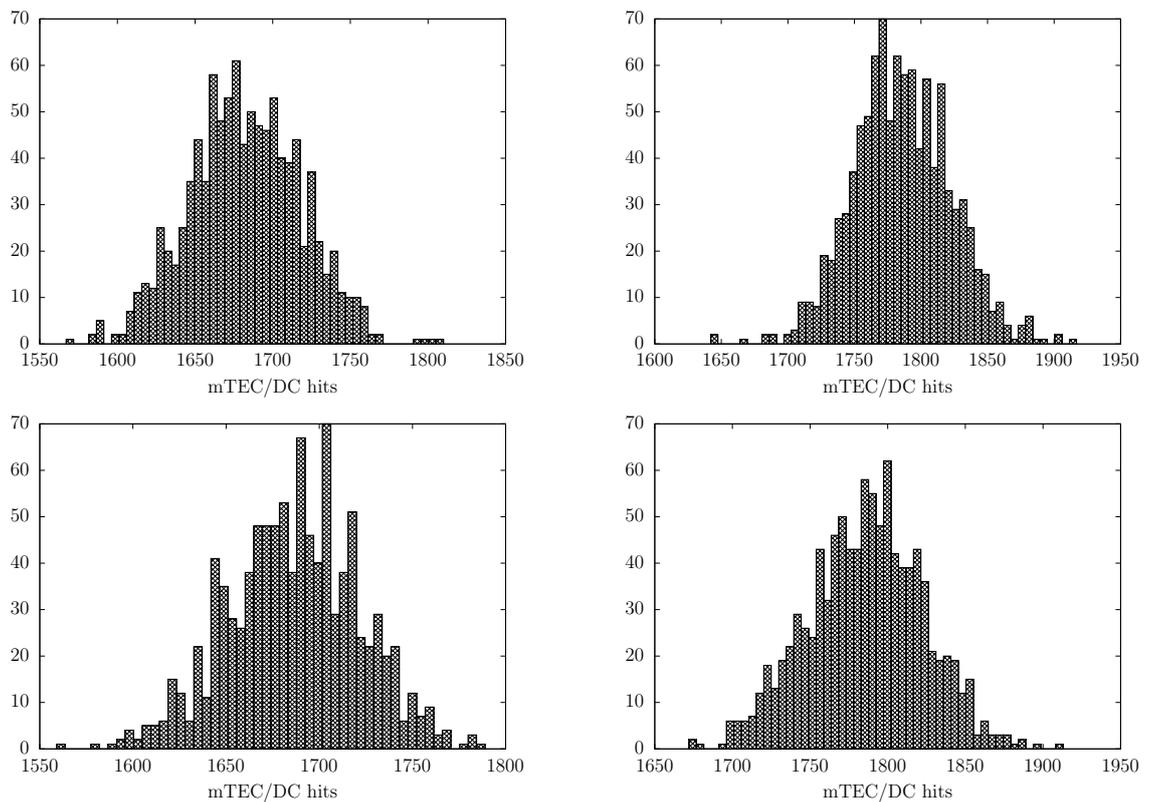


Figure 7.4: Histograms of the number of mTEC/DC hits of a T-cell calculated from 10 different microenvironments with 100 T-cells that migrated for 5 days. Upper left: 30×30 lattice, 0.02 total medulla size; Upper right: 30×30 lattice, 0.01 total medulla size; Lower left: 50×50 lattice, 0.02 total medulla size; Lower right: 50×50 lattice, 0.01 total medulla size.

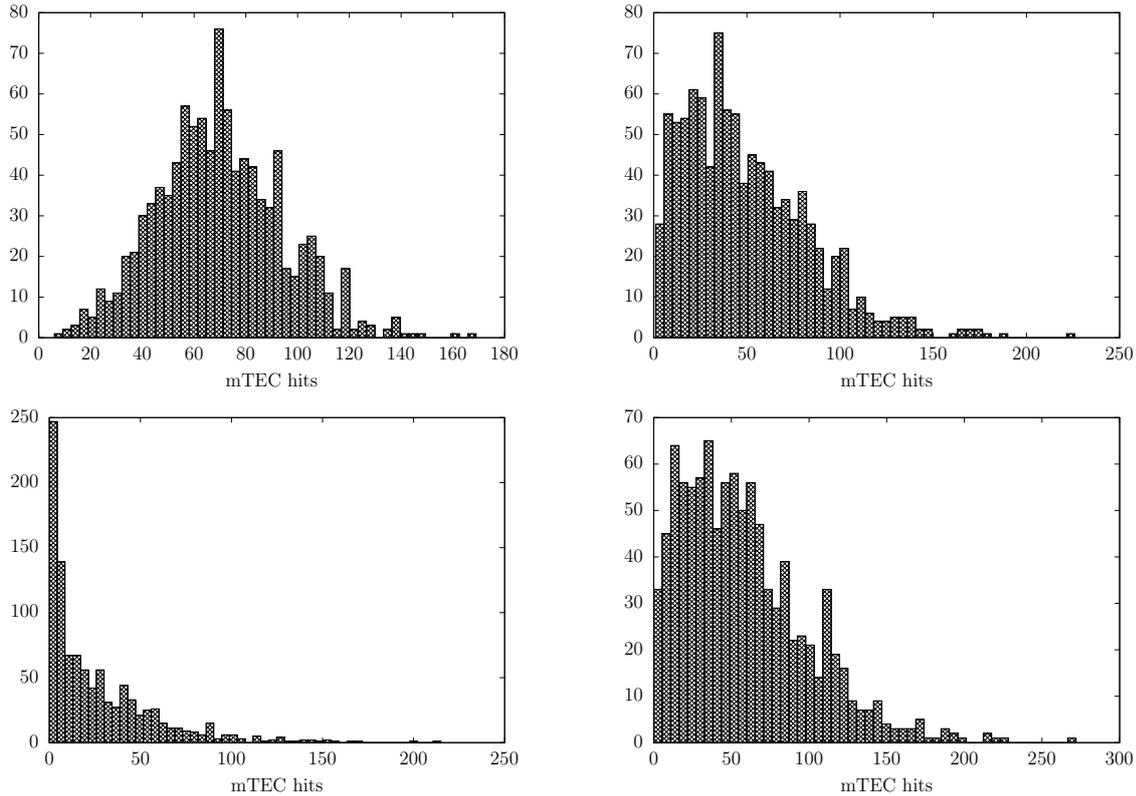


Figure 7.5: Histograms of the number of mTEC hits of a T-cell calculated from 10 different microenvironment with 100 T-cells that migrated for 5 days. Upper left: 30×3 lattice, 0.02 total medulla size; Upper right: 30×3 lattice, 0.01 total medulla size; Lower left: 50×3 lattice, 0.02 total medulla size; Lower right: 50×3 lattice, 0.01 total medulla size.

speed up of the simulation, should be suitable for many simulatory test scenarios.

Having clarified these first general facts of our simulation, we now turn to the investigation of promiscuous gene expression in our model. In a first attempt, we assume that only the mTECs are involved in presenting tissue restricted antigens. In Figure 7.5 we show the number of only the mTEC hits for the same experimental settings as in the last paragraph. In contrast to Figure 7.4 the number of mTEC hits follows quite different distributions for the different parameter settings. For a medulla size of 0.01 (upper and lower right figure) the hit distribution looks quite similar with a mean hit number of about 50. Here, a Gaussian-like shape is not met, in contrast to the case where we have the 30×3 lattice and a 0.02cm^3 medulla size. The mean number of mTEC hits is also higher with about 70. The biggest discrepancy occurs for the 50×3 lattice and 0.02cm^3 medulla size. The mTEC hit distribution is shaped similar to an exponential distribution with a mean hit number of about 10.

Given the fact, that under our assumptions a T-cell has to meet as many mTECs as possible, the variance in all hit distributions and the differences of the mean number of mTEC hits are significant. Obviously, the number of mTEC hits is very much dependent on the model parameters and the positioning of the individual T-cells on the lattice.

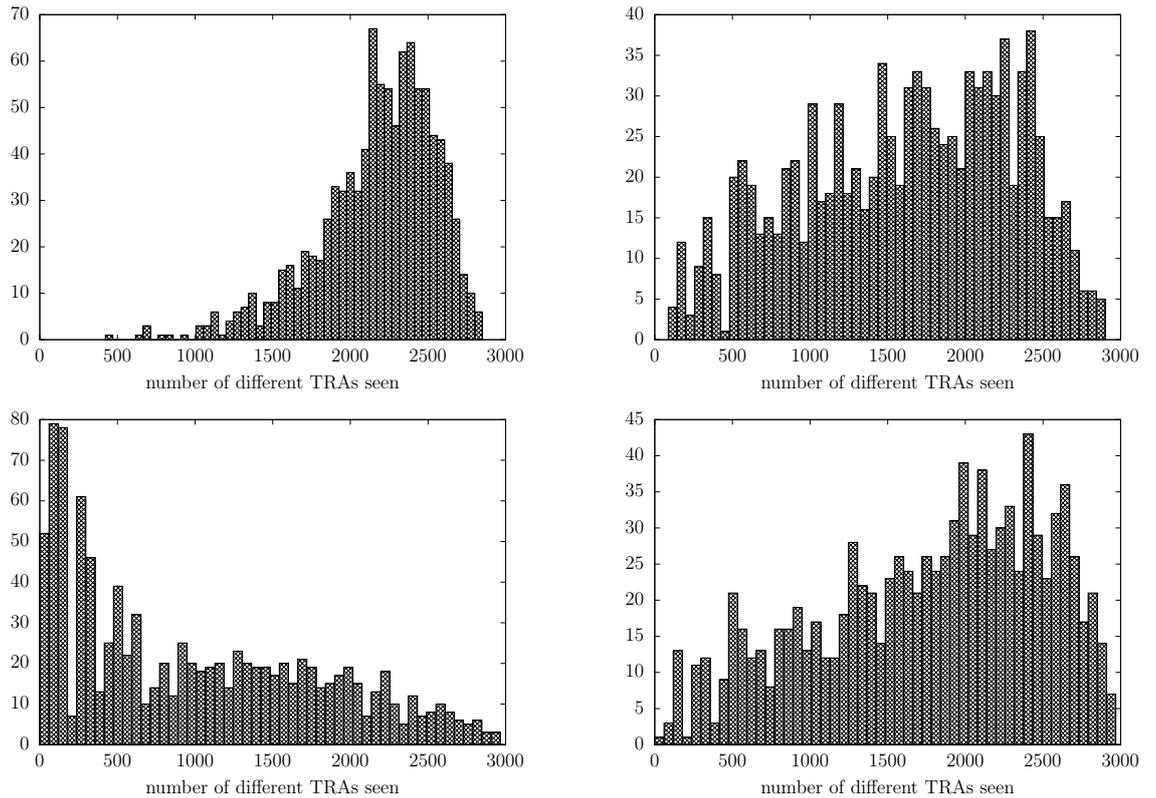


Figure 7.6: Histograms of the number different tissue restricted antigens a T-cell sees. Calculated from 10 different microenvironment with 100 T-cells that migrated for 5 days. Upper left: 30×3 lattice, 0.02 total medulla size; Upper right: 30×3 lattice, 0.01 total medulla size; Lower left: 50×3 lattice, 0.02 total medulla size; Lower right: 50×3 lattice, 0.01 total medulla size.

This can be even better illustrated if we go one step further. Until now we only observed mTEC and DC encounters. However, we can equip the mTECs with tissue restricted antigens and estimate how many of these a T-cell sees during negative selection. From the literature we know that every TRA is expressed by about 3% of all mTECs. It follows that one mTEC expresses about 90 randomly chosen TRAs. In our simulation we can follow how many of the 3000 different TRAs one T-cell observes. Furthermore every TRA on an mTEC is replaced by another every 20 hours. This is especially helpful for the purpose of our simulations, because the T-cells cannot leave the lattice but reenter when leaving and will most probably meet several mTECs/DCs more than one time on their journey. The results can be seen in Figure 7.6. The scenario settings are again the same as described before. The most eye-catching message from this Figure is that (nearly) no T-cell sees all tissue restricted antigens, given that only mTECs present them. Even more, there is a wide variance in the number of seen TRAs over all T-cells. This contradicts the overall goal to see all TRAs almost surely. In any circumstances it is obvious that promiscuous gene expression just by mTECs does not work. This speaks strongly for the mechanism of crosspresentation of TRAs from mTECs to dendritic cells. By this mechanism TRAs are transferred to DCs and presented by them.

Despite its simple structure without any specific T-cell activation mechanism, this simulation gives already some first answers. We could estimate the number of meetings between T-cells and mTECs/DCs together and for mTECs alone. This can help to specify where the boundaries of self antigen detection in the thymus are. Furthermore, we could give evidence why a mechanism of TRA delivery to DCs as discussed in the literature is really necessary. To ensure a meeting of a T-cell with all TRAs the number of mTECs is simply too low.

7.3 DISCUSSION

The model of T-cell migration which we developed in this chapter is still quite simple. However, it has the potential to be developed further in different directions in order to reflect the reality as well as possible. It is the kind of model that is well suited for information transfer from and to biological experiments. You can test hypotheses in the model, integrate new findings and later on try to show the outcomes of the simulations in experiments.

However, before expanding the model we have to get better estimates for the parameters. We extrapolated the numbers for the medullary microenvironment from different numbers of different sources. It should not be too hard to measure all of them in a single experimental setting. Nevertheless, from our perspective the estimated numbers seem to be quite reasonable. We cover a range of values, such that between about 20% and 50% of the thymic medulla is filled with DCs and mTECs. Much larger or lower numbers are unrealistic given the experimental evidence, like images of sections of the medulla, and also the fact that there are also many T-cells, other cell types and structural fibers in the medulla.

There are different possibilities to go on from here. In general, it would be interesting to see if the clustering of mTECs has a measurable effect on the rate of T-cell-mTEC meetings. It could also have an effect on the crosspresentation of TRAs. The model is also suited to test different hypotheses of how crosspresentation works. In this context it should be also possible to compare the effect of random TRA presentation in contrast to tissue emulation by mTECs or mTEC groups. The former was shown by experiments. However, it would be interesting to quantify the benefits of this kind of TRA presentation. We assume that the reason lies in the random walk migration of a single T-cell. This kind of migration does not guarantee that the T-cell meets with enough mTECs to see all different tissues. This directly leads to the question, if a T-cell really follows such a random walk or if this movement is somehow directed. Without superimposing a specialised structure on the microenvironment, such as tissue emulation, we do not see any benefit in a directed movement, but this question is definitely worth it to explore. One further important point is to investigate changes in the binding time behaviour between T-cell and mTECs/DCs and measure how these affect all the other values. Especially for shorter binding times, it seems important to us to ask how many self antigens that are presented by a given mTEC/DC are actually seen by the T-cell. This should be of great help for pursuing our overall goal, the quantification of the effect of negative selection on

the peripheral T-cell repertoire. We should be able to evaluate by the simulation how many antigens a T-cell sees and how this number is distributed. These are important information for the consideration of any T-cell activation mechanism and mechanisms of peripheral tolerance.

All the ideas formulated so far are still very general. If we have answered these, it is time to really include models of T-cell activation into the T-cell migration model. In line with our modeling approach in the previous chapter, we propose to use a (big enough) artificial self antigen repertoire and compare the outcome of negative selection for various T-cell activation models from the literature. This should give insights into the negative selection process itself and it should help to develop, prove or disprove ideas on T-cell activation.

CHAPTER 8

SUMMARY

In this chapter we recapitulate the results of this thesis, highlight their significance and finally give an outlook on how further research should be based on these results. As we already have independent discussion sections in most of the chapters, we do not go into the details again, but elucidate the results on a more general level.

Chapter 2 and 3

Let us start with the introductory chapters 2 and 3. These chapters are a review on the recent experimental and theoretical findings with regard to T-cell development and T-cell activation in order to facilitate modeling approaches. Especially, we concentrate on one very important point if it comes to modeling, namely experimental estimates for possible parameters which can be used in a model and the introduction of already existing models from which one can go on.

This review clarifies some points in particular. For one, T-cell development and T-cell activation actually are highly complex processes if all molecular details are included. This complexity magnifies if one furthermore tries to include the influence of possible interactions between a T-cell and other cells of the immune system. Hence, it is necessary to find the adequate level of abstraction. It is neither helpful to become too obsessed with the details (Occam's razor) nor is it helpful if by accident important facets of T-cell activation are ignored. In this thesis we try to explain foreign-self discrimination and we therefore took the liberty to ignore many of these molecular interactions. Only further experimental research will be able to show if this can be justified.

Furthermore, we highlighted why we think that T-cell activation should not be explained via deterministic models. Instead we proposed to define it as a problem of statistical recognition. This also motivated why we introduced the BRB model of T-cell activation and used this as the starting point for our investigations.

Finally we showed that the BRB model can really only be seen as a starting point, because there are recent experimental findings that do not contradict the model itself but need to be included in extensions of this model.

Chapter 5

In chapter 5 we developed our importance sampling method for the simulation of a special type of probabilistic models. We proved that this method fulfills a certain efficiency criterion, such that we only need subexponentially more samples to get a good

estimate for exponentially decreasing probabilities. In fact, for the BRB model we only have to increase the sample size linearly and gain a speed-up by a factor of 1000. Again, we want to point out that this method is not restricted in its use only to the BRB model of T-cell activation but to a much more general set of models. It is therefore a result in its own that can be seen separate from our results with regard to foreign-self discrimination of T-cells.

We used this importance sampling method in order to estimate T-cell activation probabilities in the BRB model. These results were already obtained before with the help of other methods. However, our results are much more exact and, what is more important, we could extract more information out of our simulations in order to really explain how foreign-self discrimination comes about in the BRB model. This was possible because with the help of our simulation results we could 'zoom' into the tail events which cause the T-cell activation. This was important, because it led to new ideas for the further modeling process.

Foreign-self discrimination in the basic BRB model is only achieved if there are much more copies of the foreign antigen than copies of the individual self antigen types on an APC. The reality looks different and therefore we need additional mechanisms which influence the foreign-self discrimination capability of the peripheral T-cell repertoire. One such mechanism can be negative selection, the process that tries to sort out T-cells that are too self-reactive before they are released into the periphery. A first extension of the BRB model therefore also includes this process and we could back up the already established results that thereby foreign-self discrimination is possible for much lower copy numbers of the foreign antigen. Furthermore we introduced a second way to include negative selection into the BRB model and showed that thereby also a better foreign-self discrimination can be established. For both extensions we used our simulation method to explain the effect of negative selection and thereby why foreign-self discrimination works better. However, we also came to the conclusion that both extensions do not reflect the biological reality well enough.

Chapter 6

Hence, in chapter 6 we tried to overcome the defects of the BRB model and its extensions by proposing a new model of T-cell activation. Several of our conclusions from the introductory chapters combined with the central aspects of the BRB model resulted in this new model. The essential new aspect in the model is, that we reduce the space of self antigen types to a discrete set. Every T-cell can then be represented by a vector of the different stimulation rates induced by the set of self antigens to this T-cell. Our model includes negative selection and we assume the same copy number for both, foreign and self antigen types, but have several other parameters which can be varied. We developed two different model variants, because we assume two different ways how an antigen presenting cell collects and presents antigens. For the estimation of the activation probabilities in our model we were able to adapt our original importance sampling method when it comes to one model variant. For the other model variant we had to

develop another importance sampling method, which is related to our previous method but the parameter estimation here is heuristic and much more involved.

We investigated the foreign-self discrimination capability and the effect of negative selection for different values of our model parameters. Thereby, we got important insights into the effect of negative selection in our model and the defects of certain parameter value combinations on the the foreign-self discrimination capability of the model.

Finally we could show that we have to change the value of the exponential binding time distribution in order to enhance foreign-self discrimination in our model. This is also biologically reasonable as there is no conclusive evidence on what the exact value of this parameter really is. In fact, the probability that a T-cell recognizes a random antigen is often estimated in the range of $10^{-5} - 10^{-4}$ [127]. If we adopt this number, than this argues for the new parameter value. Otherwise the tail of the stimulation rate distribution is not thin enough. It is interesting to see that by our change we actually reduce the probability for an antigen to induce a high stimulus to a random T-cell. This is true for both, foreign and self antigens. However, we could show that thereby the effect of negative selection is drastically increased and thus a potential peripheral T-cell repertoire is much more depleted of too self-reactive T-cells than it is for the original parameter value.

Chapter 7

As a consequence of our review in chapter 2 and our results in chapter 6 we concluded that negative selection in itself is a process that has to be investigated more thoroughly. Until now the scope of negative selection, that is its capability to sort our self-reactive T-cells, is only estimated by measurements of the number of T-cells before and after negative selection. Until very recently, the thymic medulla and thereby all processes in there had to be seen as a black box, because it was impossible to experimentally investigate it. This is gradually changing and therefore new possibilities for modeling negative selection open up. We used these recent results to model T-cell migration in the thymic medulla. This model is only a very simple first proposition. Exemplarily we estimated the number of APCs and the number of tissue-restricted antigens a T-cell encounters during negative selection. Our estimates show a great variance. As this would greatly decrease the efficacy of negative selection, we assume that our first model is too simple and we are missing important facts. However, this is exactly the kind of model whose further development benefits on the one hand from further experimental research and can on the other hand also influence further experimental research. A deeper understanding of negative selection is necessary to develop new therapies for many different diseases.

8.1 OUTLOOK

Our work delivers answers on how foreign-self discrimination of T-cells can be explained by means of probabilistic models. With our new model we have the possibility to explain

the mechanism for biological reasonable parameter values. However, our research also points out where we have to investigate in more detail in order to either develop this kind of model further or reject it ultimately.

First of all further research with regard to the mean binding time distribution is needed. Besides the exponential distribution also other distributions are possible. Van den Berg et al. already argue for a log-normal distribution because it can be derived from certain physical properties [205]. Zint et al. could show that foreign-self discrimination is enhanced if we change from exponential to log-normal in the BRB model [232]. On the theoretical level, we have to investigate how changes in the parameter $\bar{\tau}$ of the exponential distribution or a change to another distribution and all the other model parameters together can explain foreign-self discrimination. On the experimental level, we have to try to get better estimates for all of these parameters, because this will ultimately show if our model can really work. We are aware of the fact that for this the development of new experimental methods is needed. However, we think that this effort is justified as the results would not only be helpful for our model but also in general and bring us a step nearer to the true explanation of foreign-self discrimination by T-cells.

Furthermore, we have to revisit our model in the light of new experimental findings. In the sections 3.1 and 6.4 we already mentioned several of these. Often it is possible to even include them in our model as we exemplarily showed in section 6.4 for the mechanism of signal integration. Other findings might lead to new models which can nevertheless include central ideas gained in this thesis, as we argued for the mechanism of T-cell activation via receptor deformation.

Finally, we opened up a second line of research with our model of T-cell migration in the thymic medulla. This is worthwhile in its own because it can help to get ideas on how we can manipulate negative selection in order to have a peripheral T-cell repertoire that reacts to pathogens to which they, at the moment, hardly react. It is also helpful for the understanding of autoimmune reactions and how we perhaps can prevent them. As already pointed out, our model should be the opener to a fruitful discussion between experimentalist and modelers in order to pursue this target together. However, we must not forget that negative selection is also an integral part of our T-cell activation model and therefore all new knowledge on negative selection directly influences our attempts to explain foreign-self discrimination of T-cells.

BIBLIOGRAPHY

- [1] O. ACUTO AND F. MICHEL, *CD28-mediated co-stimulation: a quantitative support for TCR signalling.*, Nat Rev Immunol **3**, no. 12 (2003), pp. 939–951.
- [2] D. AIT-AZZOUZENE, P. SKOG, M. RETTER, V. KOUSKOFF, M. HERTZ, J. LANG, J. KENCH, M. CHUMLEY, D. MELAMED, J. SUDARIA, A. GAVIN, A. MARTENSSON, L. VERKOCZY, B. DUONG, J. VELA, D. NEMAZEE, AND C. ALFONSO, *Tolerance-induced receptor selection: scope, sensitivity, locus specificity, and relationship to lymphocyte-positive selection.*, Immunol Rev **197** (2004), pp. 219–230.
- [3] S. M. ALAM, G. M. DAVIES, C. M. LIN, T. ZAL, W. NASHOLDS, S. C. JAMESON, K. A. HOGQUIST, N. R. GASCOIGNE, AND P. J. TRAVERS, *Qualitative and quantitative differences in T-cell receptor binding of agonist and antagonist ligands.*, Immunity **10**, no. 2 (1999), pp. 227–237.
- [4] F. W. ALT, E. M. OLTZ, F. YOUNG, J. GORMAN, G. TACCIOLI, AND J. CHEN, *VDJ recombination.*, Immunol Today **13**, no. 8 (1992), pp. 306–314.
- [5] G. ANDERSON, P. J. L. LANE, AND E. J. JENKINSON, *Generating intrathymic microenvironments to establish T-cell tolerance.*, Nat Rev Immunol **7**, no. 12 (2007 Dec), pp. 954–963.
- [6] M. S. ANDERSON, E. S. VENANZI, L. KLEIN, Z. CHEN, S. P. BERZINS, S. J. TURLEY, H. VON BOEHMER, R. BRONSON, A. DIERICH, C. BENOIST, AND D. MATHIS, *Projection of an immunological self shadow within the thymus by the aire protein.*, Science **298**, no. 5597 (2002), pp. 1395–1401.
- [7] K. M. ARMSTRONG, K. H. PIEPENBRINK, AND B. M. BAKER, *Conformational changes and flexibility in T-cell receptor recognition of peptide-MHC complexes.*, Biochem J **415**, no. 2 (2008), pp. 183–196.
- [8] T. P. ARSTILA, A. CASROUGE, V. BARON, J. EVEN, J. KANELLOPOULOS, AND P. KOURILSKY, *A direct estimate of the human alphabeta T-cell receptor diversity.*, Science **286**, no. 5441 (1999), pp. 958–961.
- [9] S. ASMUSSEN, *Applied Probability and Queues*, Springer, New York, 2nd ed., 2003.
- [10] M. F. BACHMANN, M. SALZMANN, A. OXENIUS, AND P. S. OHASHI, *Formation of TCR dimers/trimers as a crucial step for T-cell activation.*, Eur J Immunol **28**, no. 8 (1998), pp. 2571–2579.
- [11] C. BEAUCHEMIN, N. M. DIXIT, AND A. S. PERELSON, *Characterizing T-cell movement within lymph nodes in the absence of antigen.*, J Immunol **178**, no. 9 (2007), pp. 5505–5512.

- [12] C. BEERS, A. BURICH, M. J. KLEIJMEER, J. M. GRIFFITH, P. WONG, AND A. Y. RUDENSKY, *Cathepsin S controls MHC class II-mediated antigen presentation by epithelial cells in vivo.*, *J Immunol* **174**, no. 3 (2005 Feb 1), pp. 1205–1212.
- [13] J. B. BELTMAN, S. E. HENRICKSON, U. H. VON ANDRIAN, R. J. DE BOER, AND A. F. M. MARÉE, *Towards estimating the true duration of dendritic cell interactions with T-cells.*, *J Immunol Methods* **347**, no. 1-2 (2009), pp. 54–69.
- [14] J. B. BELTMAN, A. F. M. MARÉE, AND R. J. DE BOER, *Spatial modelling of brief and long interactions between T-cells and dendritic cells.*, *Immunol Cell Biol* **85**, no. 4 (2007), pp. 306–314.
- [15] J. B. BELTMAN, A. F. M. MAREE, J. N. LYNCH, M. J. MILLER, AND R. J. DE BOER, *Lymph node topology dictates T-cell migration behavior.*, *J Exp Med* **204**, no. 4 (2007), pp. 771–780.
- [16] N. R. BHAKTA AND R. S. LEWIS, *Real-time measurement of signaling and motility during T-cell development in the thymus.*, *Semin Immunol* **17**, no. 6 (2005), pp. 411–420.
- [17] P. BILLINGSLEY, *Probability and Measure*, Wiley, New York, 3rd ed., 1995.
- [18] C. C. BLEUL, T. CORBEAUX, A. REUTER, P. FISCH, J. S. MONTING, AND T. BOEHM, *Formation of a functional thymus initiated by a postnatal epithelial progenitor cell.*, *Nature* **441**, no. 7096 (2006 Jun 22), pp. 992–996.
- [19] B. BLOM AND H. SPITS, *Development of human lymphoid cells.*, *Annu Rev Immunol* **24** (2006), pp. 287–320.
- [20] B. BLOM, M. C. VERSCHUREN, M. H. HEEMSKERK, A. Q. BAKKER, E. J. VAN GASTEL-MOL, I. L. WOLVERS-TETTERO, J. J. VAN DONGEN, AND H. SPITS, *TCR gene rearrangements and expression of the pre-T-cell receptor complex during human T-cell differentiation.*, *Blood* **93**, no. 9 (1999 May 1), pp. 3033–3043.
- [21] T. BOEHM AND C. C. BLEUL, *Thymus-homing precursors and the thymic microenvironment.*, *Trends Immunol* **27**, no. 10 (2006), pp. 477–484.
- [22] Z. BOROVSKY, G. MISHAN-EISENBERG, E. YANIV, AND J. RACHMILEWITZ, *Serial triggering of T-cell receptors results in incremental accumulation of signaling intermediates.*, *J Biol Chem* **277**, no. 24 (2002), pp. 21529–21536.
- [23] N. BOSCO, J. KIRBERG, R. CEREDIG, AND F. AGENES, *Peripheral T-cells in the thymus: have they just lost their way or do they do something?*, *Immunol Cell Biol* **87**, no. 1 (2009 Jan), pp. 50–57.
- [24] P. BOUSSO AND E. ROBEY, *Dynamics of CD8+ T-cell priming by dendritic cells in intact lymph nodes.*, *Nat Immunol* **4**, no. 6 (2003), pp. 579–585.
- [25] J. A. BUCKLEW, *Introduction to Rare Event Simulation*, Springer, 2004.
- [26] N. J. BURROUGHS AND P. A. VAN DER MERWE, *Stochasticity and spatial heterogeneity in T-cell activation.*, *Immunol Rev* **216** (2007 Apr), pp. 69–80.

- [27] S. R. CARDING AND P. J. EGAN, *Gammadelta T-cells: functional plasticity and heterogeneity.*, Nat Rev Immunol **2**, no. 5 (2002 May), pp. 336–345.
- [28] L. J. CARRENO, P. A. GONZALEZ, AND A. M. KALERGIS, *Modulation of T-cell function by TCR/pMHC binding kinetics*, Immunobiology **211** (2006), pp. 47–64.
- [29] R. CASETTI AND A. MARTINO, *The plasticity of gamma delta T-cells: innate immunity, antigen presentation and new immunotherapy.*, Cell Mol Immunol **5**, no. 3 (2008 Jun), pp. 161–170.
- [30] A. CASROUGE, E. BEAUDOING, S. DALLE, C. PANNETIER, J. KANELLOPOULOS, AND P. KOURILSKY, *Size estimate of the alpha beta TCR repertoire of naive mouse splenocytes.*, J Immunol **164**, no. 11 (2000), pp. 5782–5787.
- [31] D. CAVALLOTTI, M. ARTICO, G. IANNETTI, AND C. CAVALLOTTI, *Quantification of acetylcholinesterase-positive structures in human thymus during development and aging.*, Neurochem Int **36**, no. 1 (2000), pp. 75–82.
- [32] S. CELLI, Z. GARCIA, H. BEUNEU, AND P. BOUSSO, *Decoding the dynamics of T-cell-dendritic cell interactions in vivo.*, Immunol Rev **221** (2008), pp. 182–187.
- [33] S. CELLI, F. LEMAÎTRE, AND P. BOUSSO, *Real-time manipulation of T-cell-dendritic cell interactions in vivo reveals the importance of prolonged contacts for CD4+ T-cell activation.*, Immunity **27**, no. 4 (2007), pp. 625–634.
- [34] C. CHAN, A. J. GEORGE, AND J. STARK, *T-cell sensitivity and specificity - kinetic proofreading revisited*, Discrete and Continuous Dynamical Systems - Series B **3** (2003), pp. 343–360.
- [35] C. CHAN, J. STARK, AND A. J. GEORGE, *The impact of multiple T-cell-APC encounters and the role of anergy*, Journal of Computational and Applied Mathematics **184** (2005), pp. 101–120.
- [36] C. CHAN, J. STARK, AND A. J. T. GEORGE, *Feedback control of T-cell receptor activation.*, Proc Biol Sci **271**, no. 1542 (2004), pp. 931–939.
- [37] D. D. CHAPLIN, *1. Overview of the human immune response.*, J Allergy Clin Immunol **117**, no. 2 Suppl Mini-Primer (2006), pp. S430–S435.
- [38] T. A. CHATILA, *Molecular mechanisms of regulatory T-cell development.*, Chem Immunol Allergy **94** (2008), pp. 16–28.
- [39] C. CHIRATHAWORN, J. E. KOHLMEIER, S. A. TIBBETTS, L. M. RUMSEY, M. A. CHAN, AND S. H. BENEDICT, *Stimulation through intercellular adhesion molecule-1 provides a second signal for T-cell activation.*, J Immunol **168**, no. 11 (2002), pp. 5530–5537.
- [40] K. CHOUDHURI AND P. A. VAN DER MERWE, *Molecular mechanisms involved in T-cell receptor triggering.*, Semin Immunol **19**, no. 4 (2007), pp. 255–261.
- [41] M. CIOFANI AND J. C. ZUNIGA-PFLUCKER, *The thymus as an inductive site for T lymphopoiesis.*, Annu Rev Cell Dev Biol **23** (2007), pp. 463–493.

- [42] D. K. COLE, N. J. PUMPHREY, J. M. BOULTER, M. SAMI, J. I. BELL, E. GOSTICK, D. A. PRICE, G. F. GAO, A. K. SEWELL, AND B. K. JAKOBSEN, *Human TCR-binding affinity is governed by MHC class restriction.*, J Immunol **178**, no. 9 (2007), pp. 5727–5734.
- [43] D. COOMBS, A. M. KALERGIS, S. G. NATHENSON, C. WOFYSY, AND B. GOLDSTEIN, *Activated TCRs remain marked for internalization after dissociation from pMHC.*, Nat Immunol **3**, no. 10 (2002), pp. 926–931.
- [44] P. J. COOPER, *Interactions between helminth parasites and allergy.*, Curr Opin Allergy Clin Immunol **9**, no. 1 (2009 Feb), pp. 29–37.
- [45] G. M. DAVEY, S. L. SCHOBER, B. T. ENDRIZZI, A. K. DUTCHER, S. C. JAMESON, AND K. A. HOGQUIST, *Preselection thymocytes are more sensitive to T-cell receptor stimulation than mature T-cells.*, J Exp Med **188**, no. 10 (1998 Nov 16), pp. 1867–1874.
- [46] M. M. DAVIS AND P. J. BJORKMAN, *T-cell antigen receptor genes and T-cell recognition.*, Nature **334**, no. 6181 (1988 Aug 4), pp. 395–402.
- [47] M. M. DAVIS AND ET AL., *Ligand recognition by alpha beta T-cell receptors*, Annual Review of Immunology **16** (1998), pp. 523–544.
- [48] M. M. DAVIS, M. KROGSGAARD, J. B. HUPPA, C. SUMEN, M. A. PURBHOO, D. J. IRVINE, L. C. WU, AND L. EHRLICH, *Dynamics of cell surface molecules during T-cell recognition.*, Annu Rev Biochem **72** (2003), pp. 717–742.
- [49] S. J. DAVIS, S. IKEMIZU, E. J. EVANS, L. FUGGER, T. R. BAKKER, AND P. A. VAN DER MERWE, *The nature of molecular recognition by T-cells*, Nat Immunol **4**, no. 3 (2003), pp. 217–224.
- [50] S. J. DAVIS AND P. A. VAN DER MERWE, *The kinetic-segregation model: TCR triggering and beyond.*, Nat Immunol **7**, no. 8 (2006), pp. 803–809.
- [51] E. C. DE JONG, H. H. SMITS, AND M. L. KAPSENBERG, *Dendritic cell-mediated T-cell polarization.*, Springer Semin Immunopathol **26**, no. 3 (2005 Jan), pp. 289–307.
- [52] J. R. DELANEY, Y. SYKULEV, H. N. EISEN, AND S. TONEGAWA, *Differences in the level of expression of class I major histocompatibility complex proteins on thymic epithelial and dendritic cells influence the decision of immature thymocytes between positive and negative selection.*, Proc Natl Acad Sci U S A **95**, no. 9 (1998), pp. 5235–5240.
- [53] A. DEMBO AND O. ZEITOUNI, *Large Deviations Techniques and Applications*, Springer, 1998.
- [54] J. DERBINSKI, J. GABLER, B. BRORS, S. TIERLING, S. JONNAKUTY, M. HERGENHAHN, L. PELTONEN, J. WALTER, AND B. KYEWSKI, *Promiscuous gene expression in thymic epithelial cells is regulated at multiple levels.*, J Exp Med **202**, no. 1 (2005 Jul 4), pp. 33–45.
- [55] J. DERBINSKI, S. PINTO, S. ROSCH, K. HEXEL, AND B. KYEWSKI, *Promiscuous gene expression patterns in single medullary thymic epithelial cells argue for a stochastic mechanism.*, Proc Natl Acad Sci U S A **105**, no. 2 (2008 Jan 15), pp. 657–662.

- [56] V. DETOURS, R. MEHR, AND A. S. PERELSON, *Deriving quantitative constraints on T-cell selection from data on the mature T-cell repertoire.*, J Immunol **164**, no. 1 (2000), pp. 121–128.
- [57] J. DEVOSS, Y. HOU, K. JOHANNES, W. LU, G. I. LIOU, J. RINN, H. CHANG, R. R. CASPI, L. FONG, AND M. S. ANDERSON, *Spontaneous autoimmunity prevented by thymic expression of a single self-antigen.*, J Exp Med **203**, no. 12 (2006 Nov 27), pp. 2727–2735.
- [58] A. DIEKER AND M. MANDJES, *On Asymptotically Efficient simulation of large deviation probabilities*, Advances in applied probability **37** (2005), pp. 539–552.
- [59] W. A. DIK, K. PIKE-OVERZET, F. WEERKAMP, D. DE RIDDER, E. F. E. DE HAAS, M. R. M. BAERT, P. VAN DER SPEK, E. E. L. KOSTER, M. J. T. REINDERS, J. J. M. VAN DONGEN, A. W. LANGERAK, AND F. J. T. STAAL, *New insights on human T-cell development by quantitative T-cell receptor gene rearrangement studies and gene expression profiling.*, J Exp Med **201**, no. 11 (2005 Jun 6), pp. 1715–1723.
- [60] E. DONSKOY, D. FOSS, AND I. GOLDSCHNEIDER, *Gated importation of prothymocytes by adult mouse thymus is coordinated with their periodic mobilization from bone marrow.*, J Immunol **171**, no. 7 (2003 Oct 1), pp. 3568–3575.
- [61] O. DUSHEK AND D. COOMBS, *Analysis of serial engagement and peptide-MHC transport in T-cell receptor microclusters*, Biophysical Journal (2008), pp. biophysj.107.116897–.
- [62] O. DUSHEK, R. DAS, AND D. COOMBS, *A role for rebinding in rapid and reliable T-cell responses to antigen.*, PLoS Comput Biol **5**, no. 11 (2009), p. e1000578.
- [63] M. EGERTON, R. SCOLLAY, AND K. SHORTMAN, *Kinetics of mature T-cell development in the thymus.*, Proc Natl Acad Sci U S A **87**, no. 7 (1990 Apr), pp. 2579–2582.
- [64] L. K. ELY, K. J. GREEN, T. BEDDOE, C. S. CLEMENTS, J. J. MILES, S. P. BOTTOMLEY, D. ZERNICH, L. KJER-NIELSEN, A. W. PURCELL, J. MCCLUSKEY, J. ROSSJOHN, AND S. R. BURROWS, *Antagonism of antiviral and allogeneic activity of a human public ctl clonotype by a single altered peptide ligand: implications for allograft rejection.*, J Immunol **174**, no. 9 (2005), pp. 5593–5601.
- [65] K. ESHIMA, H. SUZUKI, AND N. SHINOHARA, *Cross-positive selection of thymocytes expressing a single TCR by multiple major histocompatibility complex molecules of both classes: implications for CD4+ versus CD8+ lineage commitment.*, J Immunol **176**, no. 3 (2006 Feb 1), pp. 1628–1636.
- [66] D. L. FOSS, E. DONSKOY, AND I. GOLDSCHNEIDER, *The importation of hematogenous precursors by the thymus is a gated phenomenon in normal adult mice.*, J Exp Med **193**, no. 3 (2001 Feb 5), pp. 365–374.
- [67] S. FRANKILD, R. J. DE BOER, O. LUND, M. NIELSEN, AND C. KESMIR, *Amino acid similarity accounts for T-cell cross-reactivity and for "holes" in the T-cell repertoire.*, PLoS One **3**, no. 3 (2008), p. e1831.

- [68] A. M. GALLEGOS AND M. J. BEVAN, *Central tolerance to tissue-specific antigens mediated by direct and indirect antigen presentation.*, J Exp Med **200**, no. 8 (2004 Oct 18), pp. 1039–1049.
- [69] K. C. GARCIA, J. J. ADAMS, D. FENG, AND L. K. ELY, *The molecular basis of TCR germline bias for MHC is surprisingly simple.*, Nat Immunol **10**, no. 2 (2009 Feb), pp. 143–147.
- [70] A. J. T. GEORGE, J. STARK, AND C. CHAN, *Understanding specificity and sensitivity of T-cell recognition.*, Trends Immunol **26**, no. 12 (2005), pp. 653–659.
- [71] R. N. GERMAIN AND I. STEFANOVA, *The dynamics of T-cell receptor signaling: complex orchestration and the key roles of tempo and cooperation.*, Annu Rev Immunol **17** (1999), pp. 467–522.
- [72] M. GIRAUD, R. TAUBERT, C. VANDIEDONCK, X. KE, M. LEVI-STRAUSS, F. PAGANI, F. E. BARALLE, B. EYMARD, C. TRANCHANT, P. GAJDOS, A. VINCENT, N. WILLCOX, D. BEESON, B. KYEWSKI, AND H.-J. GARCHON, *An IRF8-binding promoter variant and AIRE control CHRNA1 promiscuous expression in thymus.*, Nature **448**, no. 7156 (2007 Aug 23), pp. 934–937.
- [73] A. W. GOLDRATH AND M. J. BEVAN, *Selecting and maintaining a diverse T-cell repertoire.*, Nature **402**, no. 6759 (1999 Nov 18), pp. 255–262.
- [74] P. A. GONZALEZ, L. J. CARRENO, D. COOMBS, J. E. MORA, E. PALMIERI, B. GOLDSTEIN, S. G. NATHENSON, AND A. M. KALERGIS, *T-cell receptor binding kinetics required for T-cell activation depend on the density of cognate ligand on the antigen-presenting cell.*, Proc Natl Acad Sci U S A **102**, no. 13 (2005), pp. 4824–4829.
- [75] J. GOTTER, B. BRORS, M. HERGENHAHN, AND B. KYEWSKI, *Medullary epithelial cells of the human thymus express a highly diverse selection of tissue-specific genes colocalized in chromosomal clusters.*, J Exp Med **199**, no. 2 (2004 Jan 19), pp. 155–166.
- [76] D. GRAY, J. ABRAMSON, C. BENOIST, AND D. MATHIS, *Proliferative arrest and rapid turnover of thymic epithelial cells expressing aire.*, J Exp Med **204**, no. 11 (2007), pp. 2521–2528.
- [77] D. H. D. GRAY, N. SEACH, T. UENO, M. K. MILTON, A. LISTON, A. M. LEW, C. C. GOODNOW, AND R. L. BOYD, *Developmental kinetics, turnover, and stimulatory capacity of thymic epithelial cells.*, Blood **108**, no. 12 (2006), pp. 3777–3785.
- [78] J. E. GRETZ, A. O. ANDERSON, AND S. SHAW, *Cords, channels, corridors and conduits: critical architectural elements facilitating cell interactions in the lymph node cortex.*, Immunol Rev **156** (1997), pp. 11–24.
- [79] J. GUO, A. HAWWARI, H. LI, Z. SUN, S. K. MAHANTA, D. R. LITTMAN, M. S. KRANGEL, AND Y.-W. HE, *Regulation of the TCRalpha repertoire by the survival window of CD4(+)CD8(+) thymocytes.*, Nat Immunol **3**, no. 5 (2002 May), pp. 469–476.
- [80] J. S. HALE AND P. J. FINK, *Back to the thymus: peripheral T-cells come home.*, Immunol Cell Biol **87**, no. 1 (2009 Jan), pp. 58–64.

- [81] Y. HAMAZAKI, H. FUJITA, T. KOBAYASHI, Y. CHOI, H. S. SCOTT, M. MATSUMOTO, AND N. MINATO, *Medullary thymic epithelial cells expressing Aire represent a unique lineage derived from cells expressing claudin.*, *Nat Immunol* **8**, no. 3 (2007 Mar), pp. 304–311.
- [82] K. J. HARE, J. PONGRACZ, E. J. JENKINSON, AND G. ANDERSON, *Modeling TCR signaling complex formation in positive selection.*, *J Immunol* **171**, no. 6 (2003 Sep 15), pp. 2825–2831.
- [83] S. HASSLER, C. RAMSEY, M. C. KARLSSON, D. LARSSON, B. HERRMANN, B. ROZELL, M. BACKHEDEN, L. PELTONEN, O. KAMPE, AND O. WINQVIST, *Aire-deficient mice develop hematopoietic irregularities and marginal zone b-cell lymphoma.*, *Blood* **108**, no. 6 (2006 Sep 15), pp. 1941–1948.
- [84] A. C. HAYDAY AND D. J. PENNINGTON, *Key factors in the organized chaos of early T-cell development.*, *Nat Immunol* **8**, no. 2 (2007 Feb), pp. 137–144.
- [85] S. E. HENRICKSON, T. R. MEMPEL, I. B. MAZO, B. LIU, M. N. ARTYOMOV, H. ZHENG, A. PEIXOTO, M. P. FLYNN, B. SENMAN, T. JUNT, H. C. WONG, A. K. CHAKRABORTY, AND U. H. VON ANDRIAN, *T-cell sensing of antigen dose governs interactive behavior with dendritic cells and sets a threshold for T-cell activation.*, *Nat Immunol* **9**, no. 3 (2008), pp. 282–291.
- [86] W. S. HLAVACEK, A. REDONDO, C. WOFESY, AND B. GOLDSTEIN, *Kinetic proofreading in receptor-mediated transduction of cellular signals: receptor aggregation, partially activated receptors, and cytosolic messengers.*, *Bull Math Biol* **64**, no. 5 (2002), pp. 887–911.
- [87] K. A. HOGQUIST, T. A. BALDWIN, AND S. C. JAMESON, *Central tolerance: learning self-control in the thymus.*, *Nat Rev Immunol* **5**, no. 10 (2005 Oct), pp. 772–782.
- [88] F. D. HOLLANDER, *Large Deviations*, AMS, RI, 2000.
- [89] P. D. HOLLER, L. K. CHLEWICKI, AND D. M. KRANZ, *TCRs with high affinity for foreign pMHC show self-reactivity.*, *Nat Immunol* **4**, no. 1 (2003), pp. 55–62.
- [90] K. HONEY AND A. Y. RUDENSKY, *Lysosomal cysteine proteases regulate antigen presentation.*, *Nat Rev Immunol* **3**, no. 6 (2003 Jun), pp. 472–482.
- [91] V. HOREJSÍ, *Lipid rafts and their roles in T-cell activation.*, *Microbes Infect* **7**, no. 2 (2005), pp. 310–316.
- [92] D. HUDRISIER, B. KESSLER, S. VALITUTTI, C. HORVATH, J. C. CEROTTINI, AND I. F. LUESCHER, *The efficiency of antigen recognition by CD8+ CTL clones is determined by the frequency of serial TCR engagement.*, *J Immunol* **161**, no. 2 (1998), pp. 553–562.
- [93] D. F. HUNT, R. A. HENDERSON, J. SHABANOWITZ, K. SAKAGUCHI, H. MICHEL, N. SEVILIR, A. L. COX, E. APPELLA, AND V. H. ENGELHARD, *Characterization of peptides bound to the class I MHC molecule HLA-A2.1 by mass spectrometry.*, *Science* **255**, no. 5049 (1992), pp. 1261–1263.

- [94] E. S. HUSEBY, B. SATHER, P. G. HUSEBY, AND J. GOVERMAN, *Age-dependent T-cell tolerance and autoimmunity to myelin basic protein.*, *Immunity* **14**, no. 4 (2001 Apr), pp. 471–481.
- [95] E. S. HUSEBY, J. WHITE, F. CRAWFORD, T. VASS, D. BECKER, C. PINILLA, P. MARRACK, AND J. W. KAPPLER, *How the T-cell repertoire becomes peptide and MHC specific.*, *Cell* **122**, no. 2 (2005 Jul 29), pp. 247–260.
- [96] L. IGNATOWICZ, J. KAPPLER, AND P. MARRACK, *The repertoire of T-cells shaped by a single MHC/peptide ligand.*, *Cell* **84**, no. 4 (1996 Feb 23), pp. 521–529.
- [97] D. IRVINE AND ET AL., *Direct observation of ligand recognition by T-cells*, *Nature* **419** (2002), pp. 854–859.
- [98] M. ITO, K. NISHIYAMA, S. HYODO, S. SHIGETA, AND T. ITO, *Weight reduction of thymus and depletion of lymphocytes of T-dependent areas in peripheral lymphoid tissues of mice infected with francisella tularensis.*, *Infect Immun* **49**, no. 3 (1985), pp. 812–818.
- [99] T. ITO AND T. HOSHINO, *Influences on the gonad on the thymus in the mouse.*, *Z Anat Entwicklungsgesch* **123** (1963), pp. 490–497.
- [100] Y. ITOH, B. HEMMER, R. MARTIN, AND R. N. GERMAIN, *Serial TCR engagement and down-modulation by peptide:MHC molecule ligands: relationship to the quality of individual TCR signaling events.*, *J Immunol* **162**, no. 4 (1999), pp. 2073–2080.
- [101] C. A. JANEWAY AND R. MEDZHITOV, *Innate immune recognition.*, *Annu Rev Immunol* **20** (2002), pp. 197–216.
- [102] J. B. JOHNNIDIS, E. S. VENANZI, D. J. TAXMAN, J. P.-Y. TING, C. O. BENOIST, AND D. J. MATHIS, *Chromosomal clustering of genes controlled by the aire transcription factor.*, *Proc Natl Acad Sci U S A* **102**, no. 20 (2005 May 17), pp. 7233–7238.
- [103] A. M. KALERGIS, N. BOUCHERON, M. A. DOUCEY, E. PALMIERI, E. C. GOYARTS, Z. VEGH, I. F. LUESCHER, AND S. G. NATHENSON, *Efficient T-cell activation requires an optimal dwell-time of interaction between the TCR and the pMHC complex.*, *Nat Immunol* **2**, no. 3 (2001), pp. 229–234.
- [104] S. KANDULA AND C. ABRAHAM, *LFA-1 on CD4+ T-cells is required for optimal antigen-dependent activation in vivo.*, *J Immunol* **173**, no. 7 (2004), pp. 4443–4451.
- [105] M. KARIN, *Rsk Tolls the bell for endocytosis in DCs.*, *Nat Immunol* **8**, no. 11 (2007), pp. 1197–1199.
- [106] B. M. KESSLER, P. BASSANINI, J. C. CEROTTINI, AND I. F. LUESCHER, *Effects of epitope modification on T-cell receptor-ligand binding and antigen recognition by seven h-2kd-restricted cytotoxic t lymphocyte clones specific for a photoreactive peptide derivative.*, *J Exp Med* **185**, no. 4 (1997), pp. 629–640.
- [107] L. KLEIN AND B. KYEWSKI, *"Promiscuous" expression of tissue antigens in the thymus: a key to T-cell tolerance and autoimmunity?*, *J Mol Med* **78**, no. 9 (2000), pp. 483–494.

- [108] L. KLEIN, B. ROETTINGER, AND B. KYEWSKI, *Sampling of complementing self-antigen pools by thymic stromal cells maximizes the scope of central T-cell tolerance.*, Eur J Immunol **31**, no. 8 (2001 Aug), pp. 2476–2486.
- [109] T. KORN, E. BETTELLI, M. OUKKA, AND V. KUCHROO, *IL-17 and Th17 Cells.*, Annu Rev Immunol (2009 Jan 8).
- [110] A. KOSMRLJ, A. K. JHA, E. S. HUSEBY, M. KARDAR, AND A. K. CHAKRABORTY, *How the thymus designs antigen-specific and self-tolerant T-cell receptor sequences.*, Proc Natl Acad Sci U S A **105**, no. 43 (2008 Oct 28), pp. 16671–16676.
- [111] M. KROGSGAARD, Q.-J. LI, C. SUMEN, J. B. HUPPA, M. HUSE, AND M. M. DAVIS, *Agonist/endogenous peptide-MHC heterodimers drive T-cell activation and sensitivity.*, Nature **434**, no. 7030 (2005), pp. 238–243.
- [112] M. KROGSGAARD, N. PRADO, E. J. ADAMS, X.-L. HE, D.-C. CHOW, D. B. WILSON, K. C. GARCIA, AND M. M. DAVIS, *Evidence that structural rearrangements and/or flexibility during TCR binding can contribute to T-cell activation.*, Mol Cell **12**, no. 6 (2003), pp. 1367–1378.
- [113] R. A. KRONMAL AND A. J. PETERSON, *On the Alias Method for Generating Random Variables from a Discrete Distribution*, The American Statistician **33**, no. 4 (1979), pp. 214–218.
- [114] B. KYEWSKI AND J. DERBINSKI, *Self-representation in the thymus: an extended view.*, Nat Rev Immunol **4**, no. 9 (2004 Sep), pp. 688–698.
- [115] B. KYEWSKI, J. DERBINSKI, J. GOTTER, AND L. KLEIN, *Promiscuous gene expression and central T-cell tolerance: more than meets the eye.*, Trends Immunol **23**, no. 7 (2002), pp. 364–371.
- [116] B. KYEWSKI AND L. KLEIN, *A central role for central tolerance.*, Annu Rev Immunol **24** (2006), pp. 571–606.
- [117] B. KYEWSKI AND R. TAUBERT, *How promiscuity promotes tolerance: the case of myasthenia gravis.*, Ann N Y Acad Sci **1132** (2008), pp. 157–162.
- [118] Y. LI, Y. HUANG, J. LUE, J. A. QUANDT, R. MARTIN, AND R. A. MARIUZZA, *Structure of a human autoimmune TCR bound to a myelin basic protein self-peptide and a multiple sclerosis-associated MHC class ii molecule.*, EMBO J **24**, no. 17 (2005), pp. 2968–2979.
- [119] F. LIPSMEIER AND E. BAAKE, *Rare Event Simulation for T-cell Activation*, Journal of Statistical Physics **134** (2009), pp. 537–566.
- [120] S.-J. LIU, J.-P. TSAI, C.-R. SHEN, Y.-P. SHER, C.-L. HSIEH, Y.-C. YEH, A.-H. CHOU, S.-R. CHANG, K.-N. HSIAO, F.-W. YU, AND H.-W. CHEN, *Induction of a distinct CD8 Tnc17 subset by transforming growth factor-beta and interleukin-6.*, J Leukoc Biol **82**, no. 2 (2007 Aug), pp. 354–360.
- [121] G. M. LORD, R. I. LECHLER, AND A. J. GEORGE, *A kinetic differentiation model for the action of altered TCR ligands.*, Immunol Today **20**, no. 1 (1999), pp. 33–39.

- [122] J. MA, T. CHEN, J. MANDELIN, A. CEPONIS, N. E. MILLER, M. HUKKANEN, G. F. MA, AND Y. T. KONTTINEN, *Regulation of macrophage activation.*, Cell Mol Life Sci **60**, no. 11 (2003 Nov), pp. 2334–2346.
- [123] Z. MA, P. A. JANMEY, AND T. H. FINKEL, *The receptor deformation model of TCR triggering.*, FASEB J **22**, no. 4 (2008), pp. 1002–1008.
- [124] Z. MA, K. A. SHARP, P. A. JANMEY, AND T. H. FINKEL, *Surface-anchored monomeric agonist pMHCs alone trigger TCR with high sensitivity.*, PLoS Biol **6**, no. 2 (2008), p. e43.
- [125] N. MADRAS, *Lectures on Monte-Carlo Methods*, Am. Math. Soc., Providence, 2002.
- [126] E. MARKET AND F. N. PAPAVALIOU, *V(D)J recombination and the evolution of the adaptive immune system.*, PLoS Biol **1**, no. 1 (2003), p. E16.
- [127] D. MASON, *A very high level of crossreactivity is an essential feature of the T-cell receptor.*, Immunol Today **19**, no. 9 (1998), pp. 395–404.
- [128] D. MASON, *Some quantitative aspects of T-cell repertoire selection: the requirement for regulatory T-cells.*, Immunol Rev **182** (2001), pp. 80–88.
- [129] T. M. MCCAUGHTRY AND K. A. HOGQUIST, *Central tolerance: what have we learned from mice?*, Semin Immunopathol **30**, no. 4 (2008), pp. 399–409.
- [130] T. M. MCCAUGHTRY, M. S. WILKEN, AND K. A. HOGQUIST, *Thymic emigration revisited.*, J Exp Med **204**, no. 11 (2007 Oct 29), pp. 2513–2520.
- [131] T. W. MCKEITHAN, *Kinetic proofreading in T-cell receptor signal transduction.*, Proc Natl Acad Sci U S A **92**, no. 11 (1995), pp. 5042–5046.
- [132] R. MEDZHITOV, *Approaching the asymptote: 20 years later.*, Immunity **30**, no. 6 (2009), pp. 766–775.
- [133] T. R. MEMPEL, S. E. HENRICKSON, AND U. H. VON ANDRIAN, *T-cell priming by dendritic cells in lymph nodes occurs in three distinct phases.*, Nature **427**, no. 6970 (2004), pp. 154–159.
- [134] M. J. MILLER, A. S. HEJAZI, S. H. WEI, M. D. CAHALAN, AND I. PARKER, *T-cell repertoire scanning is promoted by dynamic dendritic cell behavior and random T-cell motility in the lymph node.*, Proc Natl Acad Sci U S A **101**, no. 4 (2004), pp. 998–1003.
- [135] M. J. MILLER, O. SAFRINA, I. PARKER, AND M. D. CAHALAN, *Imaging the single cell dynamics of CD4+ T-cell activation by dendritic cells in lymph nodes.*, J Exp Med **200**, no. 7 (2004), pp. 847–856.
- [136] M. J. MILLER, S. H. WEI, I. PARKER, AND M. D. CAHALAN, *Two-photon imaging of lymphocyte motility and antigen response in intact lymph node.*, Science **296**, no. 5574 (2002), pp. 1869–1873.
- [137] K. H. G. MILLS, *Induction, function and regulation of IL-17-producing T-cells.*, Eur J Immunol **38**, no. 10 (2008 Oct), pp. 2636–2649.

- [138] V. MULLER AND S. BONHOEFFER, *Quantitative constraints on the scope of negative selection.*, Trends Immunol **24**, no. 3 (2003), pp. 132–135.
- [139] S. MURATA, K. SASAKI, T. KISHIMOTO, S.-I. NIWA, H. HAYASHI, Y. TAKAHAMA, AND K. TANAKA, *Regulation of CD8+ T-cell development by thymus-specific proteasomes.*, Science **316**, no. 5829 (2007 Jun 1), pp. 1349–1353.
- [140] S. MURATA, Y. TAKAHAMA, AND K. TANAKA, *Thymoproteasome: probable role in generating positively selecting peptides.*, Curr Opin Immunol **20**, no. 2 (2008 Apr), pp. 192–196.
- [141] K. M. MURPHY AND S. L. REINER, *The lineage decisions of helper T-cells.*, Nat Rev Immunol **2**, no. 12 (2002 Dec), pp. 933–944.
- [142] K. M. MURPHY, P. TRAVERS, AND M. WALPORT, *Immunobiology*, Garland Scienc, 2007.
- [143] D. NAEHER, M. A. DANIELS, B. HAUSMANN, P. GUILLAUME, I. LUESCHER, AND E. PALMER, *A constant affinity threshold for T-cell tolerance.*, J Exp Med **204**, no. 11 (2007), pp. 2553–2559.
- [144] D. NEMAZEE, V. KOUSKOFF, M. HERTZ, J. LANG, D. MELAMED, K. PAPE, AND M. RETTER, *B-cell-receptor-dependent positive and negative selection in immature B cells.*, Curr Top Microbiol Immunol **245**, no. 2 (2000), pp. 57–71.
- [145] J. NIKOLICH-ZUGICH, M. K. SLIFKA, AND I. MESSAOUDI, *The many important facets of T-cell repertoire diversity.*, Nat Rev Immunol **4**, no. 2 (2004 Feb), pp. 123–132.
- [146] L. J. C. NO, S. M. BUENO, P. BULL, S. G. NATHENSON, AND A. M. KALERGIS, *The half-life of the T-cell receptor/peptide-major histocompatibility complex interaction can modulate T-cell activation in response to bacterial challenge.*, Immunology **121**, no. 2 (2007), pp. 227–237.
- [147] M. OUKKA, *Th17 cells in immunity and autoimmunity.*, Ann Rheum Dis **67 Suppl 3** (2008 Dec), pp. iii26–9.
- [148] Y. OZEKI, K. KANEDA, N. FUJIWARA, M. MORIMOTO, S. OKA, AND I. YANO, *In vivo induction of apoptosis in the thymus by administration of mycobacterial cord factor (trehalose 6,6'-dimycolate).*, Infect Immun **65**, no. 5 (1997), pp. 1793–1799.
- [149] R. PACHOLCZYK AND J. KERN, *The T-cell receptor repertoire of regulatory T-cells.*, Immunology **125**, no. 4 (2008 Dec), pp. 450–458.
- [150] N. W. PALM AND R. MEDZHITOV, *Pattern recognition receptors and control of adaptive immunity.*, Immunol Rev **227**, no. 1 (2009), pp. 221–233.
- [151] Z. PANCER AND M. D. COOPER, *The evolution of adaptive immunity.*, Annu Rev Immunol **24** (2006), pp. 497–518.
- [152] *Fundamental Immunology*, Lippincott Williams and Wilkins, 6th ed., 2008.

- [153] H. T. PETRIE, F. LIVAK, D. G. SCHATZ, A. STRASSER, I. N. CRISPE, AND K. SHORTMAN, *Multiple rearrangements in T-cell receptor alpha chain genes maximize the production of useful thymocytes.*, J Exp Med **178**, no. 2 (1993 Aug 1), pp. 615–622.
- [154] H. E. PORRITT, L. L. RUMFELT, S. TABRIZIFARD, T. M. SCHMITT, J. C. ZÁŽÁŠIGA-PFLÄJCKER, AND H. T. PETRIE, *Heterogeneity among DN1 prothymocytes reveals multiple progenitors with different capacities to generate T-cell and non-T-cell lineages.*, Immunity **20**, no. 6 (2004), pp. 735–745.
- [155] A. I. PROIETTO, S. VAN DOMMELEN, AND L. WU, *The impact of circulating dendritic cells on the development and differentiation of thymocytes.*, Immunol Cell Biol **87**, no. 1 (2009 Jan), pp. 39–45.
- [156] M. A. PURBHOO, D. J. IRVINE, J. B. HUPPA, AND M. M. DAVIS, *T-cell killing does not require the formation of a stable mature immunological synapse*, Nat Immunol **5**, no. 5 (2004), pp. 524–530.
- [157] J. D. RABINOWITZ, C. BEESON, D. S. LYONS, M. M. DAVIS, AND H. M. MCCONNELL, *Kinetic discrimination in T-cell activation.*, Proc Natl Acad Sci U S A **93**, no. 4 (1996), pp. 1401–1405.
- [158] J. D. RABINOWITZ, C. BEESON, C. WULFING, K. TATE, P. M. ALLEN, M. M. DAVIS, AND H. M. MCCONNELL, *Altered T-cell receptor ligands trigger a subset of early T-cell signals.*, Immunity **5**, no. 2 (1996), pp. 125–135.
- [159] C. RAMSEY, O. WINQVIST, L. PUHAKKA, M. HALONEN, A. MORO, O. KAMPE, P. ESKELIN, M. PELTO-HUIKKO, AND L. PELTONEN, *Aire deficient mice develop multiple features of APECED phenotype and show altered immune response.*, Hum Mol Genet **11**, no. 4 (2002 Feb 15), pp. 397–409.
- [160] J. ROBINSON, M. J. WALLER, P. PARHAM, N. DE GROOT, R. BONTROP, L. J. KENNEDY, P. STOEHR, AND S. G. E. MARSH, *IMGT/HLA and IMGT/MHC: sequence databases for the study of the major histocompatibility complex.*, Nucleic Acids Res **31**, no. 1 (2003 Jan 1), pp. 311–314.
- [161] H. R. RODEWALD, T. BROCKER, AND C. HALLER, *Developmental dissociation of thymic dendritic cell and thymocyte lineages revealed in growth factor receptor mutant mice.*, Proc Natl Acad Sci U S A **96**, no. 26 (1999), pp. 15068–15073.
- [162] H. R. RODEWALD, S. PAUL, C. HALLER, H. BLUETHMANN, AND C. BLUM, *Thymus medulla consisting of epithelial islets each derived from a single progenitor.*, Nature **414**, no. 6865 (2001), pp. 763–768.
- [163] P. ROSENSTIEL, P. E. R. EVA, S. SCHREIBER, AND T. C. G. BOSCH, *Evolution and Function of Innate Immune Receptors - Insights from marine invertebrates*, J Innate Immun **1** (2009), pp. 291–300.
- [164] C. ROSETTE, G. WERLEN, M. A. DANIELS, P. O. HOLMAN, S. M. ALAM, P. J. TRAVERS, N. R. GASCOIGNE, E. PALMER, AND S. C. JAMESON, *The impact of duration versus extent of TCR occupancy on T-cell activation: a revision of the kinetic proofreading model.*, Immunity **15**, no. 1 (2001), pp. 59–70.

- [165] S. M. ROSS, *Simulation*, Academic Press, 2002.
- [166] E. V. ROTHENBERG, *How T-cells Count*, Science **273**, no. 5271 (1996), pp. 78–0.
- [167] V. RUBIO-GODOY, V. DUTOIT, D. RIMOLDI, D. LIENARD, F. LEJEUNE, D. SPEISER, P. GUILLAUME, J. C. CEROTTINI, P. ROMERO, AND D. VALMORI, *Discrepancy between ELISPOT IFN-gamma secretion and binding of A2/peptide multimers to TCR reveals interclonal dissociation of ctl effector function from TCR-peptide/MHC complexes half-life.*, Proc Natl Acad Sci U S A **98**, no. 18 (2001), pp. 10302–10307.
- [168] M. G. RUDOLPH, R. L. STANFIELD, AND I. A. WILSON, *How TCRs bind MHCs, peptides, and coreceptors.*, Annu Rev Immunol **24** (2006), pp. 419–466.
- [169] J. S. SADOWSKY AND J. A. BUCKLEW, *On Large Deviations Theory and Asymptotically Efficient Monte Carlo estimation*, IEEE Transactions on Information Theory **36** (1990), pp. 579–588.
- [170] S. SAKAGUCHI, *Naturally arising CD4+ regulatory T-cells for immunologic self-tolerance and negative control of immune responses.*, Annu Rev Immunol **22** (2004), pp. 531–562.
- [171] S. SAKAGUCHI AND F. POWRIE, *Emerging challenges in regulatory T-cell function and biology.*, Science **317**, no. 5838 (2007 Aug 3), pp. 627–629.
- [172] S. SAKAGUCHI AND N. SAKAGUCHI, *Regulatory T-cells in immunologic self-tolerance and autoimmune disease.*, Int Rev Immunol **24**, no. 3-4 (2005 May-Aug), pp. 211–226.
- [173] W. W. A. SCHAMEL, R. M. R. NO, S. MINGUET, A. R. ORTÍZ, AND B. ALARCÓN, *A conformation- and avidity-based proofreading mechanism for the TCR-cd3 complex.*, Trends Immunol **27**, no. 4 (2006), pp. 176–182.
- [174] N. SEACH, T. UENO, A. L. FLETCHER, T. LOWEN, M. MATTESICH, C. R. ENGERWERDA, H. S. SCOTT, C. F. WARE, A. P. CHIDGEY, D. H. D. GRAY, AND R. L. BOYD, *The lymphotoxin pathway regulates Aire-independent expression of ectopic genes and chemokines in thymic stromal cells.*, J Immunol **180**, no. 8 (2008 Apr 15), pp. 5384–5392.
- [175] K. SHORTMAN, M. EGERTON, G. J. SPANGRUDE, AND R. SCOLLAY, *The generation and fate of thymocytes.*, Semin Immunol **2**, no. 1 (1990), pp. 3–12.
- [176] T. N. SIMS, T. J. SOOS, H. S. XENIAS, B. DUBIN-THALER, J. M. HOFMAN, J. C. WAITE, T. O. CAMERON, V. K. THOMAS, R. VARMA, C. H. WIGGINS, M. P. SHEETZ, D. R. LITTMAN, AND M. L. DUSTIN, *Opposing effects of PKCtheta and WASp on symmetry breaking and relocation of the immunological synapse.*, Cell **129**, no. 4 (2007), pp. 773–785.
- [177] A. SINGER AND R. BOSSELUT, *CD4/CD8 coreceptors in thymocyte development, selection, and lineage commitment: analysis of the CD4/cd8 lineage decision.*, Adv Immunol **83** (2004), pp. 91–131.

- [178] D. SKOKOS, G. SHAKHAR, R. VARMA, J. C. WAITE, T. O. CAMERON, R. L. LINDQUIST, T. SCHWICKERT, M. C. NUSSENZWEIG, AND M. L. DUSTIN, *Peptide-MHC potency governs dynamic interactions between T-cells and dendritic cells in lymph nodes.*, Nat Immunol **8**, no. 8 (2007), pp. 835–844.
- [179] J. SLOAN-LANCASTER, B. D. EVAVOLD, AND P. M. ALLEN, *Induction of T-cell anergy by altered T-cell-receptor ligand on live antigen-presenting cells.*, Nature **363**, no. 6425 (1993), pp. 156–159.
- [180] E. Y. SO, H. H. PARK, AND C. E. LEE, *IFN-gamma and IFN-alpha posttranscriptionally down-regulate the IL-4-induced il-4 receptor gene expression.*, J Immunol **165**, no. 10 (2000 Nov 15), pp. 5472–5479.
- [181] M. SOSPEDRA, X. FERRER-FRANCESCH, O. DOMINGUEZ, M. JUAN, M. FOZ-SALA, AND R. PUJOL-BORRELL, *Transcription of a broad range of self-antigens in human thymus suggests a role for central mechanisms in tolerance toward peripheral antigens.*, J Immunol **161**, no. 11 (1998 Dec 1), pp. 5918–5929.
- [182] J. SOUSA AND J. CARNEIRO, *A mathematical analysis of TCR serial triggering and down-regulation.*, Eur J Immunol **30**, no. 11 (2000), pp. 3219–3227.
- [183] H. SPITS, B. BLOM, A. C. JALECO, K. WEIJER, M. C. VERSCHUREN, J. J. VAN DONGEN, M. H. HEEMSKERK, AND P. C. RES, *Early stages in the development of human T, natural killer and thymic dendritic cells.*, Immunol Rev **165** (1998 Oct), pp. 75–86.
- [184] T. K. STARR, S. C. JAMESON, AND K. A. HOGQUIST, *Positive and negative selection of T-cells.*, Annu Rev Immunol **21** (2003), pp. 139–176.
- [185] I. STEFANOVÁ, B. HEMMER, M. VERGELLI, R. MARTIN, W. E. BIDDISON, AND R. N. GERMAIN, *TCR ligand discrimination is enforced by competing ERK positive and shp-1 negative feedback pathways.*, Nat Immunol **4**, no. 3 (2003), pp. 248–254.
- [186] G. G. STEINMANN, *Changes in the human thymus during aging.*, Curr Top Pathol **75** (1986), pp. 43–88.
- [187] S. STEVANOVIC AND H. SCHILD, *Quantitative aspects of T-cell activation–peptide generation and editing by MHC class I molecules.*, Semin Immunol **11**, no. 6 (1999), pp. 375–384.
- [188] B. STOCKINGER AND M. VELDHOEN, *Differentiation and function of Th17 T-cells.*, Curr Opin Immunol **19**, no. 3 (2007 Jun), pp. 281–286.
- [189] J. D. STONE, A. S. CHERVIN, AND D. M. KRANZ, *T-cell receptor binding affinities and kinetics: impact on T-cell activity and specificity.*, Immunology **126**, no. 2 (2009), pp. 165–176.
- [190] Y. SYKULEV, Y. VUGMEYSTER, A. BRUNMARK, H. L. PLOEGH, AND H. N. EISEN, *Peptide antagonism and T-cell receptor interactions with peptide-MHC complexes.*, Immunity **9**, no. 4 (1998), pp. 475–483.

- [191] S. J. SZABO, S. T. KIM, G. L. COSTA, X. ZHANG, C. G. FATHMAN, AND L. H. GLIMCHER, *A novel transcription factor, T-bet, directs Th1 lineage commitment.*, *Cell* **100**, no. 6 (2000 Mar 17), pp. 655–669.
- [192] Y. TAKAHAMA, *Journey through the thymus: stromal guides for T-cell development and selection.*, *Nat Rev Immunol* **6**, no. 2 (2006 Feb), pp. 127–135.
- [193] Y. TAKAHAMA, K. TANAKA, AND S. MURATA, *Modest cortex and promiscuous medulla for thymic repertoire formation.*, *Trends Immunol* **29**, no. 6 (2008 Jun), pp. 251–255.
- [194] R. TAUBERT, J. SCHWENDEMANN, AND B. KYEWSKI, *Highly variable expression of tissue-restricted self-antigens in human thymus: implications for self-tolerance and autoimmunity.*, *Eur J Immunol* **37**, no. 3 (2007 Mar), pp. 838–848.
- [195] A. THEDREZ, C. SABOURIN, J. GERTNER, M.-C. DEVILDER, S. ALLAIN-MAILLET, J.-J. FOURNIE, E. SCOTET, AND M. BONNEVILLE, *Self/non-self discrimination by human gammadelta T-cells: simple solutions for a complex issue?*, *Immunol Rev* **215** (2007 Feb), pp. 123–135.
- [196] S. TIAN, R. MAILE, E. J. COLLINS, AND J. A. FRELINGER, *CD8+ T-cell activation is governed by TCR-peptide/MHC affinity, not dissociation rate.*, *J Immunol* **179**, no. 5 (2007), pp. 2952–2960.
- [197] P. TOLAR AND S. K. PIERCE, *Change we can believe in—of the conformational type. Workshop on the initiation of antigen receptor signaling.*, *EMBO Rep* **10**, no. 4 (2009), pp. 331–336.
- [198] A. TRAUTMANN AND C. RANDRIAMAMPITA, *Initiation of TCR signalling revisited.*, *Trends Immunol* **24**, no. 8 (2003), pp. 425–428.
- [199] L.-O. TYKOCINSKI, A. SINEMUS, AND B. KYEWSKI, *The thymus medulla slowly yields its secrets.*, *Ann N Y Acad Sci* **1143** (2008 Nov), pp. 105–122.
- [200] D. M. UNDERHILL AND A. OZINSKY, *Phagocytosis of microbes: complexity in action.*, *Annu Rev Immunol* **20** (2002), pp. 825–852.
- [201] C. UTZNY, D. COOMBS, S. MULLER, AND S. VALITUTTI, *Analysis of peptide/MHC-induced TCR downregulation: deciphering the triggering kinetics.*, *Cell Biochem Biophys* **46**, no. 2 (2006), pp. 101–111.
- [202] S. VALITUTTI AND A. LANZAVECCHIA, *Serial triggering of TCRs: a basis for the sensitivity and specificity of antigen recognition.*, *Immunol Today* **18**, no. 6 (1997), pp. 299–304.
- [203] S. VALITUTTI, S. MULLER, M. CELLA, E. PADOVAN, AND A. LANZAVECCHIA, *Serial triggering of many T-cell receptors by a few peptide-MHC complexes.*, *Nature* **375**, no. 6527 (1995), pp. 148–151.
- [204] H. A. VAN DEN BERG AND C. MOLINA-PARIS, *Thymic Presentation of Autoantigens and the Efficiency of negative selection*, *Journal of Theoretical Medicine* (2003).

- [205] H. A. VAN DEN BERG, D. A. RAND, AND N. J. BURROUGHS, *A reliable and safe T-cell repertoire based on low-affinity T-cell receptors.*, J Theor Biol **209**, no. 4 (2001), pp. 465–486.
- [206] H. A. VAN DEN BERG, L. WOOLDRIDGE, B. LAUGEL, AND A. K. SEWELL, *Coreceptor CD8-driven modulation of T-cell antigen receptor specificity.*, J Theor Biol **249**, no. 2 (2007), pp. 395–408.
- [207] P. A. VAN DER MERWE, *The TCR triggering puzzle.*, Immunity **14**, no. 6 (2001), pp. 665–668.
- [208] P. A. VAN DER MERWE AND S. J. DAVIS, *Molecular interactions mediating T-cell antigen recognition.*, Annu Rev Immunol **21** (2003), pp. 659–684.
- [209] W. VAN EWYJK, G. HOLLANDER, C. TERHORST, AND B. WANG, *Stepwise development of thymic microenvironments in vivo is regulated by thymocyte subsets.*, Development **127**, no. 8 (2000 Apr), pp. 1583–1591.
- [210] W. VAN EWYJK, E. W. SHORES, AND A. SINGER, *Crosstalk in the mouse thymus.*, Immunol Today **15**, no. 5 (1994 May), pp. 214–217.
- [211] N. S. VAN OERS, N. KILLEEN, AND A. WEISS, *Lck regulates the tyrosine phosphorylation of the T-cell receptor subunits and zap-70 in murine thymocytes.*, J Exp Med **183**, no. 3 (1996), pp. 1053–1062.
- [212] R. VARMA, *TCR triggering by the pMHC complex: valency, affinity, and dynamics.*, Sci Signal **1**, no. 19 (2008), p. pe21.
- [213] V. VENTURI, K. KEDZIERSKA, D. A. PRICE, P. C. DOHERTY, D. C. DOUEK, S. J. TURNER, AND M. P. DAVENPORT, *Sharing of T-cell receptors in antigen-specific responses is driven by convergent recombination.*, Proc Natl Acad Sci U S A **103**, no. 49 (2006 Dec 5), pp. 18691–18696.
- [214] V. VENTURI, D. A. PRICE, D. C. DOUEK, AND M. P. DAVENPORT, *The molecular basis for public T-cell responses?*, Nat Rev Immunol **8**, no. 3 (2008), pp. 231–238.
- [215] J. VILLASENOR, W. BESSE, C. BENOIST, AND D. MATHIS, *Ectopic expression of peripheral-tissue antigens in the thymic epithelium: probabilistic, monoallelic, misinitiated.*, Proc Natl Acad Sci U S A **105**, no. 41 (2008 Oct 14), pp. 15854–15859.
- [216] A. VIOLA AND A. LANZAVECCHIA, *T-cell Activation Determined by T-cell Receptor Number and Tunable thresholds*, Science **273**, no. 5271 (1996), pp. 104–106.
- [217] N. WATANABE, Y.-H. WANG, H. K. LEE, T. ITO, Y.-H. WANG, W. CAO, AND Y.-J. LIU, *Hassall's corpuscles instruct dendritic cells to induce CD4+CD25+ regulatory T-cells in human thymus.*, Nature **436**, no. 7054 (2005 Aug 25), pp. 1181–1185.
- [218] K. S. WEBER, D. L. DONERMEYER, P. M. ALLEN, AND D. M. KRANZ, *Class II-restricted T-cell receptor engineered in vitro for higher affinity retains peptide specificity and function.*, Proc Natl Acad Sci U S A **102**, no. 52 (2005), pp. 19033–19038.

- [219] A. WEISS AND D. R. LITTMAN, *Signal transduction by lymphocyte antigen receptors.*, Cell **76**, no. 2 (1994), pp. 263–274.
- [220] C. M. WITT AND E. A. ROBEY, *Thymopoiesis in 4 dimensions.*, Semin Immunol **17**, no. 1 (2005 Feb), pp. 95–102.
- [221] B. WOELFING, A. TRAUlsen, M. MILINSKI, AND T. BOEHM, *Review. Does intra-individual major histocompatibility complex diversity keep a golden mean?*, Philos Trans R Soc Lond B Biol Sci (2008).
- [222] C. WOFsy, D. COOMBS, AND B. GOLDSTEIN, *Calculations show substantial serial engagement of T-cell receptors.*, Biophys J **80**, no. 2 (2001), pp. 606–612.
- [223] D. L. WOODLAND AND R. W. DUTTON, *Heterogeneity of CD4(+) and CD8(+) T-cells.*, Curr Opin Immunol **15**, no. 3 (2003 Jun), pp. 336–342.
- [224] L. WOOLDRIDGE, H. A. VAN DEN BERG, M. GLICK, E. GOSTICK, B. LAUGEL, S. L. HUTCHINSON, A. MILICIC, J. M. BRENCHLEY, D. C. DOUEK, D. A. PRICE, AND A. K. SEWELL, *Interaction between the CD8 coreceptor and major histocompatibility complex class I stabilizes T-cell receptor-antigen complexes at the cell surface.*, J Biol Chem **280**, no. 30 (2005), pp. 27491–27501.
- [225] L. C. WU, D. S. TUOT, D. S. LYONS, K. C. GARCIA, AND M. M. DAVIS, *Two-step binding mechanism for T-cell receptor recognition of peptide MHC.*, Nature **418**, no. 6897 (2002), pp. 552–556.
- [226] X. O. YANG, R. NURIEVA, G. J. MARTINEZ, H. S. KANG, Y. CHUNG, B. P. PAPPU, B. SHAH, S. H. CHANG, K. S. SCHLUNS, S. S. WATOWICH, X.-H. FENG, A. M. JETTEN, AND C. DONG, *Molecular antagonism and plasticity of regulatory and inflammatory T-cell programs.*, Immunity **29**, no. 1 (2008 Jul), pp. 44–56.
- [227] M. YAZDANBAKHSH, A. VAN DEN BIGGELAAR, AND R. M. MAIZELS, *Th2 responses without atopy: immunoregulation in chronic helminth infections and reduced allergic disease.*, Trends Immunol **22**, no. 7 (2001 Jul), pp. 372–377.
- [228] D. ZEHN AND M. J. BEVAN, *T-cells with low avidity for a tissue-restricted antigen routinely evade central and peripheral tolerance and cause autoimmunity.*, Immunity **25**, no. 2 (2006), pp. 261–270.
- [229] D. ZEHN, S. Y. LEE, AND M. J. BEVAN, *Complete but curtailed T-cell response to very low-affinity antigen.*, Nature **458**, no. 7235 (2009), pp. 211–214.
- [230] J. ZERRAHN, W. HELD, AND D. H. RAULET, *The MHC reactivity of the T-cell repertoire prior to positive and negative selection.*, Cell **88**, no. 5 (1997 Mar 7), pp. 627–636.
- [231] H. ZHENG, B. JIN, S. E. HENRICKSON, A. S. PERELSON, U. H. VON ANDRIAN, AND A. K. CHAKRABORTY, *How Antigen Quantity and Quality Determine T-Cell Decisions in Lymphoid tissue*, Mol Cell Biol. **28** (2008), pp. 4040–4051.
- [232] N. ZINT, E. BAAKE, AND F. DEN HOLLANDER, *How T-cells use large deviations to recognize foreign antigens.*, J Math Biol (2008).