



Antigen presentation on MHC molecules as a diversity filter that enhances immune efficacy[☆]

Hugo A. van den Berg^{*,1}, David A. Rand¹

Mathematics Institute, University of Warwick, Coventry CV4 7AL, UK

Received 14 February 2002; received in revised form 7 February 2003; accepted 3 April 2003

Abstract

We consider the way in which antigen is presented to T cells on MHC molecules and ask how MHC peptide presentation could be optimized so as to obtain an effective and safe immune response. By analysing this problem with a mathematical model of T-cell activation, we deduce the need for both MHC restriction and high presentation selectivity. We find that the optimal selectivity is such that about one pathogen-derived peptide is presented per MHC isoform, on the average. We also indicate upper and lower bounds to the number of MHC isoforms per individual based on detectability requirements. Thus we deduce that an important role of MHC presentation is to act as a filter that limits the diversity of antigen presentation.

© 2003 Elsevier Ltd. All rights reserved.

Keywords: MHC polymorphism; Antigen presentation; TCR signalling; Negative selection; Positive selection

1. Introduction

T cells recognize and respond to processed protein antigens in the form of peptide fragments bound to specialized peptide receptors, which are membrane glycoproteins encoded by genes belonging to the major histocompatibility (MHC) complex, and hence called MHC molecules. The MHC complex is both highly polygenic and highly polymorphic, as a result of which each individual expresses about 10 MHC isoforms, which is still a modest subset of the hundreds of alleles present in the population (Parham, 2000). The peptide is tightly bound to a binding groove on the outer surface of the MHC molecule, and is thus presented to the TCR molecule, which makes contacts with both the peptide and the surrounding surface of the MHC molecule (Parham, 2000).

The peptides displayed on the MHC molecules derive from proteins, expressed by the cell itself in the case of class I MHC molecules, and proteins taken up by endocytosis in the case of class II MHC molecules (Parham, 2000). However, not all peptides that might be

derived from these proteins are eventually displayed on MHC molecules: antigen presentation is selective (Rammensee et al., 1993; Stevanović and Schild, 1999). This presentation selectivity partly derives from the peptide-binding requirements of the MHC peptide-binding groove. Despite these requirements, MHC peptide-binding remains highly promiscuous (Rothbard and Geftter, 1991) and the binding requirements of a given MHC isoform allow a vastly greater variety of peptides to bind than are actually found as epitopes (Jardetzky et al. 1991; Stevanović and Schild, 1999). For example, Influenza A nucleoprotein was found to give rise to only one cytotoxic T lymphocyte (CTL)-epitope on three murine MHC isoforms despite having several MHC binding motif peptides for each of these isoforms, and to yield not a single epitope on an MHC isoform for which the protein was found to contain 10 motif peptides (Stevanović and Schild, 1999). These and other findings indicate that much of the selectivity of peptide presentation originates elsewhere in the antigen processing pathway (Deng et al., 1997; Rammensee et al., 1993; Stevanović and Schild, 1999): for MHC class I presentation, specificity requirements apply during proteasome cleavage (Kuttler et al., 2000; Yewdell and Bennink, 2001) and transport into the endoplasmic reticulum (Neefjes et al., 1993). In MHC class II presentation, endocytic proteases (Watts, 2001) and

[☆]Supported by the Wellcome Trust.

*Corresponding author. Tel.: +44-02476-523698; fax: +44-02476-524182.

E-mail address: hugo@maths.warwick.ac.uk (H.A. van den Berg).

¹Both authors contributed equally to this work.

the B cell receptor (Siemasko and Clark, 2001) are likely contributors to presentation selectivity.

The combined selectivity introduced by the various stages which a peptide has to go through to become an MHC ligand is such that every protein is represented, on average, by less than one peptide on a given MHC isoform (Falk et al., 1991a; Rammensee et al., 1993). Thus presentation is highly selective. This fact has been interpreted in terms of efficient sampling: for survival of pathogen attacks or malignant transformation, it would be advantageous to present a sample of every protein synthesized in the cell (Rammensee et al., 1993; Stevanović and Schild, 1999). By contrast, we propose that the high selectivity of antigen presentation directly arises from the need to mount an immune response which is reliable as well as self-tolerant.

Antigen selectivity determines the typical diversity of foreign peptides presented on diseased cells and the professional antigen presenting cells which initiate the immune response. It is clear that foreign presentation diversity is minimized if the TCR is able to interact with only one of the MHC isoforms expressed by the host, and when this isoform presents just a single foreign peptide. Thus high presentation selectivity and MHC restriction both contribute to low foreign presentation diversity.

The aim of this paper is to demonstrate that low foreign presentation diversity entails high immune efficacy, that is, an enhanced ability of the immune

system to mount a safe and effective response. Low foreign presentation diversity is a combination of MHC restriction and presentation dominance. Although MHC restriction and presentation dominance are mechanistically distinct, we will show that their enhancement of immune efficacy stems from a common mathematical principle. To this end, we need to define immune efficacy precisely, and we propose to do this in terms of the mathematical model of T cell activation developed in van den Berg et al. (2001).

A possible adverse consequence of high presentation selectivity is that a pathogen may not be presented at all on a given MHC isoform. The expression of multiple MHC isoforms obviously reduces the likelihood of a pathogen going undetected altogether. We show that the likelihood of such immunodetectability decreases exponentially with the number of MHC isoforms expressed by the individual, which bounds this number from below. We also consider various effects that bound the size of the MHC isoform repertoire from above, and propose a bound based on our model of TCR signalling.

Using immune efficacy as a central concept, our model of T-cell activation ties together presentation selectivity, MHC restriction, as well as the sizes of the TCR clonotype and MHC isoform repertoires. In Section 2 we outline the model and our account of immune efficacy. The main argument of the paper is presented in Section 3. In Section 4 we discuss the consequences of minimizing foreign presentation

Table 1
Notation

Symbol	Interpretation
i	index for T cell clonotypes
j	index for pMHC species
ℓ	index for MHC isoforms
W	TCR triggering rate
w_{ij}	TCR triggering rate elicited in clonotype i by a pMHC complex of species j
W_{act}	cellular activation threshold
w_{for}	$\sum_{j \in \text{for}} \zeta_j w_{ij}$, TCR triggering rate due to the foreign peptides
P_{act}	probability that a naïve clone is activated by presentation of foreign peptides
M_T	number of pMHC molecules capable of interacting with the TCR
ζ_j	number of MHC molecules presenting peptide j relative to M_T
ζ_{for}	$\sum_{j \in \text{for}} \zeta_j$ (the sum ranges over all foreign peptides)
n_C	number of clonotypes in the naïve T cell repertoire
n_M	number of MHC isoforms
m	number of MHC isoforms capable of interacting with a given TCR clonotype ($m \leq n_M$)
N_{bg}	number of potential self (background) peptides in the APP
n_{bg}	number of self peptides present in the APP
N_{for}	number of potential foreign (pathogen-derived) peptides in the APP
n_{for}	number of foreign peptides in the APP
p	selectivity of peptide presentation on MHC molecules
μ_W	mean of the distribution F_W of pMHC specific triggering rates w_{ij}
σ_W^2	variance of the distribution F_W of pMHC specific triggering rates w_{ij}

APP: antigen presentation profile; TCR: T cell antigen receptor; pMHC: major histocompatibility complex molecule, presenting a peptide. Not included are a few notations used only locally.

diversity for the size of the MHC isoform repertoire in terms of immunodetectability requirements. Notation used throughout is summarized in Table 1; additional notation is always explained locally.

2. T cell activation

T cell signalling takes place in a region of close contact between the T cell and the *antigen-presenting cell* (APC) (Dustin and Shaw, 1999; Grakoui et al., 1999). Within this region of contact, the TCR molecules on the T cell are able to interact with the peptide/MHC (pMHC) molecules on the APC. The TCR binds to the pMHC to form a *ternary* TCR/pMHC complex. Once docked in such a complex, the TCR may be *triggered* to transmit a signal to the cellular interior (Lanzavecchia et al., 1999). Whereas virtually all TCRs on a given T cell have the same molecular identity (Heath and Miller, 1993), the pMHC ensemble is highly diverse (Chicz et al., 1993; Hunt et al., 1992; Rudensky et al., 1991). In principle, all these pMHC species are able to evoke signalling, though some are far more likely to do so than others. Thus, to be able to distinguish signalling from a relevant pMHC species, the T cell cannot be responsive to individual signalling events; rather, it has to be responsive to the TCR triggering rate, that is, the rate at which TCRs are being triggered.

2.1. TCR triggering and T cell activation

When a T cell of clonotype i conjugates with a given APC, TCR molecules are triggered at a rate W_i which is dependent both on the peptide mixture presented on the APC and the clonotype. According to our model, the TCR triggering rate W_i can be written as a weighed sum of contributions from all pMHC species j present

$$W_i = \sum_{\text{pMHC } j} Z_j w_{ij} \quad (1)$$

where Z_j represents the number of MHC molecules in the contact region which present peptide species j and are capable of interacting with a TCR of clonotype i . A pMHC molecule of species j triggers TCRs at a rate w_{ij} . This MHC-specific triggering rate w_{ij} can be related to molecular parameters:

$$w_{ij} \stackrel{\text{def}}{=} \frac{\exp\{-T_R/T_{ij}\}}{T_{ij} + (k_{ij}R)^{-1}} \quad (2)$$

where T_{ij} is the average lifetime (dissociation time) of the ternary complex composed of a TCR of species i and a pMHC molecule of species j ; k_{ij} is the rate of formation of the complex; and R represents the free TCR count in the contact region.

In deriving Eq. (2), we have assumed that a ternary complex formation results in the triggering of the

TCR/CD3 complex if the complex has lasted longer than the *receptor threshold duration* T_R . This is a simple way of accommodating what is known about the sequence of phosphorylation events that take place in the CD3 complex leading up to the transmission of a signal (Neumeister Kersh et al., 1998; Shores et al., 1997); it has been proposed that this sequence has ‘kinetic proofreading’ properties (McKeithan, 1995; Rabinowitz et al., 1996). We also assume that no more than a single triggering event can be associated with the formation of any ternary complex. Eqs. (1) and (2) are derived in our earlier paper (van den Berg et al., 2001); the derivation is briefly summarized in Appendix A.

The advantage of Eq. (2) is that it allows us to calculate the statistics of the triggering rate from the statistics of association and dissociation times. The latter are more readily accessible to experimental determination than the triggering rate itself (Hudrisier et al., 1998; Lyons et al., 1996).

Interestingly, the specific triggering rate w_{ij} is maximal at an intermediate dissociation time on the order of T_R (Fig. 1); cf. Kalergis et al. (2001); Krummel et al. (2000); Lanzavecchia et al. (1999). A pMHC j for which w_{ij} achieves its maximum value is called an *optimal agonist*. The existence of such an optimum accords well with the serial triggering hypothesis which states that a single pMHC molecule must trigger up to ~ 200 TCRs, and that pMHCs with average dissociation times significantly greater than T_R therefore waste time by staying docked for a longer time than required for signalling (Bongrand and Malissen, 1998; Valitutti et al., 1995; Valitutti and Lanzavecchia, 1997). This effect is most pronounced when the triggering rate is not limited by the free TCR count, $R \gg (k_{ij}T_R)^{-1}$ (top curve in Fig. 1). As R becomes smaller, the effect diminishes, and the difference in triggering rates between optimal and




Fig. 1. Normalized MHC-specific triggering rate $w_{ij}T_R$ as a function of normalized average dissociation time T_{ij}/T_R : $w_{ij}T_R = \exp\{-T_R/T_{ij}\}/(T_{ij}/T_R + (k_{ij}T_R R)^{-1})^{-1}$. Curves are for (from top to bottom): $R \gg (k_{ij}T_R)^{-1}$; $R = 4(k_{ij}T_R)^{-1}$; $R = 2(k_{ij}T_R)^{-1}$; $R = (k_{ij}T_R)^{-1}$; $R = (2k_{ij}T_R)^{-1}$; $R = (4k_{ij}T_R)^{-1}$; $R = (8k_{ij}T_R)^{-1}$; $R = (10k_{ij}T_R)^{-1}$.

suboptimal agonists becomes less important. The absence of the optimum under TCR-limited conditions may explain the findings of Holler et al. (2001) that a much increased dissociation time need not diminish the potency of the agonist.

Let $M_T \stackrel{\text{def}}{=} \sum_j Z_j$ denote the total number of MHC molecules capable of interacting with the TCR of clonotype i (we assume that M_T has the same value for every clonotype i). We define the *relative presentation level* of pMHC species j relative to this total MHC count, as follows:

$$\zeta_j \stackrel{\text{def}}{=} Z_j / M_T. \quad (3)$$

The relative presentation levels $\{\zeta_j\}$ together constitute the *antigen presentation profile* (APP) on the APC. The parameter set $\{\zeta_j\}$ characterizes the given APC relative to the MHC molecules capable of interacting with the TCR. It will be convenient in what follows to refer to some subset of the APP; we call such a subset an *antigen presentation subprofile*.

We have argued that the T cell must be responsive to the triggering rate rather than individual triggering events. The T cell might be able to gauge this rate directly, or respond to some accumulatory process, such as the total number of down-regulated TCR molecules (Viola and Lanzavecchia, 1996). Whatever the precise mechanism, it seems reasonable to assume that the relevant signal transmitted through the T cell's downstream signalling pathways is a monotone increasing function of the TCR triggering rate. We assume that the T cell responds when this signal exceeds a certain threshold. The response can be proliferation, differentiation, cytokine secretion, or delivery of a lethal hit (Parham, 2000). A T cell is typically only capable of a subset of these responses, depending on its state of differentiation (Parham, 2000). Moreover, some responses are more readily evoked than others (Itoh and Germain, 1997; Valitutti et al., 1996). We therefore associate each of the possible responses with a specific threshold. On the monotonicity assumption just mentioned, the threshold translates into a threshold on the triggering rate itself. Thus we arrive at the following *T cell activation hypothesis*:

The T cell becomes activated when the triggering rate exceeds an activation threshold $W_i > W_{\text{act}}$.

This hypothesis leads us to consider the probability that the triggering rate elicited by a given APP will exceed the activation threshold W_{act} in a clonotype chosen at random. This probability is the frequency of responding clonotypes in the TCR repertoire, in other words, the proportion of clones that is capable of responding to the foreign antigen. The probability is given by the triggering rate distribution function, which is the subject of the next section.

2.2. Immune response efficacy and activation curves

The cellular immune response can fail to protect the host from disease in two ways: either the immune response could fail to detect the presence of a pathogen or the immune response could harm healthy cells. An efficacious immune response reacts to the presence of the pathogen and avoids responding to self antigens, both with high probabilities. In particular, given an APC presenting only self-peptides and a clonotype i selected at random from the naïve repertoire, the probability that the triggering rate W_i evoked by the self-only APP exceeds W_{act} should be very much smaller than the reciprocal of the repertoire size. But given an APC which also presents one or more foreign peptides, the probability that W_i exceeds W_{act} should be much larger than the reciprocal of the repertoire size. The former requirement means that the expected number of clones responding to self-only APPs on professional APCs is very much smaller than 1, and the latter means that with very high probability at least one clone will respond to a self-plus-foreign APP.

To calculate these probabilities, we need to know the triggering rate distribution among the clonotypes relative to a given APP. We represent this distribution in terms of the proportion of clonotypes whose triggering rate exceeds ω when conjugating with this APP:

$$P_W(\omega) \stackrel{\text{def}}{=} \mathbb{P}\{W_i > \omega\}. \quad (4)$$

This triggering rate distribution can be calculated by combining the triggering rate model, Eq. (1) with the distribution of average dissociation times, from which each T_{ij} is a realization. The latter distribution can in principle be determined empirically (Hudrisier et al., 1998; Lyons et al., 1996; Savage et al., 1999). The examples shown below are based on the assumptions that the average dissociation times follow an exponential distribution (see Appendix A), and that the TCR triggering rate is MHC-limited, that is, the free TCR count is sufficiently large, $R \gg (k_{ij} T_{ij})^{-1}$.

The triggering rate distribution is depicted schematically in Fig. 2, where it is plotted as a function of W both for an APP containing only self-peptides (dashed line) and for a self-plus-foreign APP (solid line). We call such curves *activation curves*. At $W = W_{\text{act}}$, we can read off the proportion P_{act} of clonotypes which will respond to the self-plus-foreign APP as well as the proportion P_{self} of clonotypes which will respond to the self-only APP. The separation Δ_P between these two probabilities is essential. In particular, if n_C is the number of distinct TCR clonotypes (the TCR repertoire size), its reciprocal $1/n_C$ must be separated by several orders of magnitude from both P_{act} and P_{self} .




Fig. 2. Activation curves: the triggering rate distribution for the self peptide background (dashed line) and the self background plus a foreign component (drawn line). Top panel: prior to negative selection. Bottom panel: following negative selection.

Once the activated naïve cells have differentiated into armed effector cells, they will interact with numerous healthy cells presenting only self peptides. The probability that a CD8 effector will inflict lethal damage on such a healthy cell must stay below a certain bound if the immune response is to be sufficiently self-tolerant. This is expressed by the probability P_{tol} in Fig. 2 (top panel). Thus, W_{act} is not allowed to be smaller than ω defined by $P_W(\omega) = P_{tol}$.

The latter point gains importance when we take into account the combined effects of negative selection in the thymus and peripheral tolerization. Both induce a shift to the left of the self-only activation curve, as shown in Fig. 2 (bottom panel). The separation Δ_W represents the maximum allowed difference between the activation thresholds in the naïve cells and in the effector cells. It is readily seen that the existence of such a separation allows the effector cells to safely detect the foreign antigen at lower presentation levels than the levels at which the professional APCs in the secondary lymphoid tissues present these peptides. The extent to which the background self-only activation curve can be shifted to the left by negative selection depends on the parameters of selection and on the relative importance of variability in the self triggering rate between APCs as compared to the variability between the clonotypes. We intend to give a detailed treatment of these relationships elsewhere.

For the present purposes, it suffices to establish the importance of the separations Δ_P and Δ_W for immune efficacy, since the objective of this paper is to study how the separation between self-only and self-plus-foreign activation curves may be maximized by variations in the foreign antigen presentation subprofile.

Summing the relative presentation levels ζ_j (as defined by Eq. (3)) for all foreign peptides j , we obtain the foreign presentation level ζ_{for} :

$$\zeta_{for} \stackrel{\text{def}}{=} \sum_{j \in \text{for}} \zeta_j. \quad (5)$$

Here ‘for’ is the set of distinct pMHC species in the foreign subprofile, of which there are n_{for} . Let w_{for} denote the MHC-specific TCR triggering rate due to the foreign peptides

$$w_{for} \stackrel{\text{def}}{=} \sum_{j \in \text{for}} \zeta_j w_{ij}. \quad (6)$$

Let F_{bg} denote the distribution function of the contribution made to the triggering rate W_i by the self background. Since the TCR triggering rate is a sum consisting of a term contributed by the self background and a term contributed by the foreign peptides, the distribution function of this triggering rate can be found from the convolution formula

$$P_W(\omega) = \int \mathbb{P}\{w_{for} > (\omega - (1 - \zeta_{for})\omega')/M_T\} dF_{bg}(\omega') \quad (7)$$

whose integrand is a probability of the form

$$\mathbb{P}\{w_{for} > \omega\}.$$

Maximizing probabilities of this form amounts to maximizing P_W , which in turn maximizes the separation between self-only and self-plus-foreign activation curves, that is, immune efficacy. (We will refer to both Δ_P and Δ_W indiscriminately as the ‘separation’ between these activation curves.) This motivates the main objective of this paper, which is to argue that probabilities of the form $\mathbb{P}\{w_{for} > \omega\}$ are maximized if the foreign presentation diversity is minimized.

3. Diversity and presentation dominance

We now argue that the foreign triggering probability $\mathbb{P}\{w_{for} > \zeta_{for} w\}$ is maximized by minimizing the diversity of the foreign presentation subprofile (at least for large enough w). It then follows that the efficacy of immune recognition is greatly improved by *presentation dominance*. A peptide will be said to be presentation dominant for a TCR clonotype when it (almost) entirely occupies the foreign presentation subprofile as registered by that TCR clonotype (i.e. $\zeta_k \approx \zeta_{for}$); the foreign presentation diversity is then minimized.

We analyse presentation dominance as a combination of low foreign presentation diversity and MHC restriction. Presentation dominance in turn affects the immunodetectability of the pathogen, as will be discussed in Section 4.

3.1. Low foreign presentation diversity

Our first step toward presentation dominance is to establish the importance of low diversity in the foreign antigen presentation subprofile, with a view to the analysis of how presentation selectivity p affects immune efficacy. Presentation selectivity is the probability that a randomly chosen peptide species can be presented on a given MHC isoform. We argue in Appendix B that ζ_{for} does not itself vary with presentation selectivity p , essentially because self and foreign peptides are equally affected by variations in p .

Let n_{for} denote the number of foreign peptide species presented. We can characterize the diversity of the foreign subprofile as follows:

$$n_{\text{div, for}} \stackrel{\text{def}}{=} 1 / \sum_{j \in \text{for}} (\zeta_j / \zeta_{\text{for}})^2 \quad (8)$$

so that $1 \leq n_{\text{div, for}} \leq n_{\text{for}}$, where $n_{\text{div, for}} = 1$ corresponds to total presentation dominance of the foreign presentation subprofile by a single pMHC species.

For large n_{for} , we can directly apply a Large Deviations estimate (see Appendix C.1):

$$\mathbb{P}\{w_{\text{for}} > \zeta_{\text{for}} w\} \approx \exp\{-n_{\text{div, for}}(w - \mu_w)^2 / (2\sigma_w^2)\}, \quad (9)$$

which shows immediately that the foreign triggering probability decreases exponentially with increasing diversity (w_{ij} is distributed with mean μ_w and variance σ_w^2).

To extend this result on diversity and foreign triggering probability to small numbers of foreign pMHC species ($n_{\text{for}} \sim 1$), we note that we can always decrease the diversity by eliminating the pMHC species

with the smallest presentation level from the foreign presentation subprofile, keeping ζ_{for} constant (Appendix B) and redistributing the remaining foreign peptides in such a way that the relative proportions remain the same. Thus, the foreign pMHC species k defined by

$$\zeta_k = \min_{j \in \text{for}} \{\zeta_j\} \quad (10)$$

is eliminated in favour of the other foreign pMHC species, where the ratio $\zeta_j / \zeta_{j'}$ is kept constant for each pair of foreign pMHC species (j, j') , $j \neq k, j' \neq k$. Eq. (10) implies $\zeta_k / \zeta_{\text{for}} \leq 1 / n_{\text{div, for}}$. It is readily seen that $\zeta_k / \zeta_{\text{for}} < 1 / n_{\text{div, for}}$ is equivalent to $dn_{\text{div, for}} / d\zeta_k > 0$, which implies that the foreign presentation diversity is reduced by elimination of the foreign pMHC species k with the smallest presentation level. To complete the argument, it remains to show that the foreign triggering probability is increased by eliminating this pMHC species; we then have a decrease in diversity concomitant with an increase in the foreign triggering diversity. To this end, consider the graphs shown in Figs. 3 (for $n_{\text{for}} = 3$) and 4 (for $n_{\text{for}} = 4$), where the foreign triggering probability is shown as a function of $\zeta_k / \zeta_{\text{for}}$ (see Appendix A for simulation details). Here ζ_k is an arbitrarily selected foreign pMHC species, and the remaining presentation levels retain their relative proportions.

We show in Appendix C.2 that, when Eq. (10) is satisfied, $\zeta_k / \zeta_{\text{for}}$ is below the value for which these curves attain their minimum (estimates of which are indicated by $\bar{z}$ in the graphs). Thus the simulations depicted in Figs. 3 and 4 indicate that a decrease in diversity correlates with an increase in foreign triggering probability. In fact, the graphs suggest that this decrease in diversity can lead to a dramatic increase of the foreign triggering probability, provided that w is close enough to the maximum value $(eT_R)^{-1}$. The improvement of the foreign triggering probability by elimination of the




Fig. 3. Foreign triggering rate probability as a function of $\zeta_1 / \zeta_{\text{for}}$, for $n_{\text{for}} = 3$. Estimates for the location of the minimum are indicated (see Appendix C.2). Left panel: $\zeta_2 = \zeta_3$. Middle panel: $\zeta_2 = \frac{1}{4}\zeta_3$. Right panel: $\zeta_2 = 0$; this is a degenerate case coinciding with the $n_{\text{for}} = 2$ case. From top to bottom, curves are for $wT_R = 0.20, 0.21, 0.22, 0.23, 0.24, 0.25, 0.26, 0.27, 0.28, 0.29, 0.30, 0.31, 0.32, 0.33, 0.34, 0.35, 0.36, 0.37$; the mean value of wT_R is 0.017.




Fig. 4. Foreign triggering rate probability as a function of ζ_1/ζ_{for} , for $n_{for} = 4$. Estimates for the location of the minimum are indicated (see Appendix C.2). Left panel: $\zeta_2 = \zeta_3 = \zeta_4$. Middle panel: $\zeta_2 = \frac{1}{5}\zeta_4$; $\zeta_3 = \frac{3}{5}\zeta_4$. Right panel: $\zeta_2 = \frac{1}{7}\zeta_4$; $\zeta_3 = \frac{2}{7}\zeta_4$. From top to bottom, curves are for $wT_R = 0.20, 0.21, 0.22, 0.23, 0.24, 0.25, 0.26, 0.27, 0.28, 0.29, 0.30, 0.31, 0.32, 0.33, 0.34, 0.35, 0.36, 0.37$; the mean value of wT_R is 0.017.

foreign pMHC species with the lowest presentation level is, approximately,

$$\ln \frac{\mathbb{P}\{\text{after}\}}{\mathbb{P}\{\text{before}\}} \approx \frac{(w - \mu_w)^2}{2\sigma_w^2} \quad (11)$$

confirming that for w -values which exceed the mean by several standard deviations, the improvement can be several orders of magnitude (see Appendix C.2 for a derivation of this estimate).

In conclusion, the inverse relationship between foreign presentation diversity and foreign triggering probability obtains both at high and low numbers of foreign peptides in the foreign presentation subprofile. It follows that the separation between self-only and self-plus-foreign activation curves will be greatest when a single foreign peptide dominates, since in that case foreign presentation diversity attains its absolute minimum.

3.2. MHC restriction

The second step towards presentation dominance is to argue that the foreign triggering rate distribution P_w typically increases when a clonotype is able to interact with a smaller number of distinct MHC isoforms. The triggering rate due to the foreign subprofile may be rewritten as follows:

$$\begin{aligned} w_{for} &= \sum_{j \in \text{for}} \zeta_j w_{ij} = \sum_{\ell=1}^m \sum_{j \in \text{for} \cap \ell} \zeta_j w_{ij} \\ &= \sum_{\ell=1}^m \zeta_\ell \sum_{j \in \text{for} \cap \ell} \frac{\zeta_j}{\zeta_\ell} w_{ij} \end{aligned} \quad (12)$$

where ζ_ℓ is the fraction of MHC molecules of isoform ℓ , among the MHC molecules capable of interacting with clonotype i . One would typically expect $\zeta_\ell = 1/m$, barring any biases in isoform usage. It may be seen that ζ_j/ζ_ℓ is an independent random variable (in fact, it is determined by the peptides' statistics with respect to

the various antigen processing steps which determine whether they are presented on an MHC molecule isoform ℓ among peptides; Appendix B makes this idea more precise in terms of presentation propensities). Thus, on the assumption that presentation selectivity p is the same for all MHC isoforms, we may define independent and identically distributed random variables $w_i^{[\ell]}$ as follows:

$$w_i^{[\ell]} \stackrel{\text{def}}{=} \sum_{j \in \text{for} \cap \ell} \frac{\zeta_j}{\zeta_\ell} w_{ij}. \quad (13)$$

The probability of the foreign subprofile's TCR triggering rate exceeding a given value ω becomes, with Eqs. (12) and (13),

$$\mathbb{P}\{w_{for} > \omega\} = \mathbb{P}\left\{\sum_{\ell=1}^m \zeta_\ell w_i^{[\ell]} > \omega\right\}. \quad (14)$$

The triggering probability is now in a form in which the same argument can be applied as in the previous section: the probability increases upon elimination of the isoform ℓ' which satisfies $\min_\ell \{\zeta_\ell\} = \zeta_{\ell'}$. In this instance, the argument does not rely on the constancy of ζ_{for} , by virtue of the way the $w_i^{[\ell]}$ were defined in Eq. (13) (in particular, the ζ_ℓ now play the role of the z_j in Appendix C.2).

Example. An example of the variation of the triggering rate distribution function with the number of MHC isoforms capable of interacting with the given clonotype is depicted in Fig. 5. In the simulation shown, it was assumed that $\zeta_j/\zeta_\ell \equiv z_V$ where z_V is some positive constant (thus, the foreign presentation level restricted to a given MHC isoform was proportional to the number of foreign peptides presented by that isoform). For the sake of simplicity, it was assumed that $\zeta_\ell \equiv 1/m$, with $n_{for,\ell}$ foreign peptides presented on MHC isoform ℓ : the m values in $\{n_{for,\ell}\}_{\ell=1}^m$ were taken to be independently Poisson distributed with expectation 1.




Fig. 5. MHC restriction: $\mathbb{P}\{\frac{1}{m} \sum_{\ell=1}^m \sum_{j \in \text{for} \cap \ell} w_{ij} > \omega\}$ as a function of ωT_R , for $m = 1, 2, 3, 4, 5, 6, 7$ (top to bottom curves). The number of foreign peptides presented by a single MHC isoform is Poisson distributed with mean 1.

The triggering rate distribution function was computed as

$$\begin{aligned} & \mathbb{P}\left\{\frac{1}{m} \sum_{\ell=1}^m \sum_{j \in \text{for} \cap \ell} z_V w_{ij} > z_V \omega\right\} \\ &= \mathbb{P}\left\{\frac{1}{m} \sum_{\ell=1}^m \sum_{j \in \text{for} \cap \ell} w_{ij} > \omega\right\} \end{aligned}$$

and shown as a function of ωT_R in Fig. 5. The graph clearly shows that the gains accrued by MHC restriction can be substantial. We therefore suggest that MHC restriction has a significant role to play in maximizing immune efficacy.

We conclude that the separation between the self-only and self-plus-back-ground activation curves increases with a reduction of m , the number of MHC isoforms capable of interacting with a given TCR clonotype. Carrying this reduction of m through to its logical conclusion, we see that immune efficacy is greatest when each TCR clonotype is *MHC restricted*: able to interact with no more than a single MHC isoform among those expressed. Note that this reasoning does not insist that absolute MHC restriction must ensue, in which each TCR is absolutely incapable of interacting with any expressed pMHC not of the TCR's nominal isoform. Rather, the argument shows that the closer the actual situation approaches the idealization of absolute MHC restriction, the greater the efficacy-enhancing effect is.

3.3. Thymic selection and alloreactivity

Our argument only concerns MHC restriction relative to those isoforms expressed in the host. That is, the argument only leads us to conclude that efficacy is enhanced by restricting each TCR clonotype to a single MHC isoform among those present in the individual

host. In particular, we leave open the possibility that a clonotype might still be able to interact with one or more MHC isoforms not expressed by the individual. This is not surprising: after all, response efficacy is not affected if some TCR is also able to interact with isoforms not expressed in the host. That capability only comes into play when tissue bearing non-host MHC isoforms is introduced, and the occurrence of MHC-driven graft rejection indicates that many clonotypes do in fact have this ability (Janeway and Travers, 1997).

The possibility of alloreactivity immediately raises the question of how the immune system ensures that each clonotype is able to interact with no more than one MHC isoform among those expressed. Such clonotypes may be far more likely to be deleted by negative selection in the thymus. A clonotype which is able to interact with m MHC isoforms will register, during thymic selection, a TCR signal which is typically m times as strong as the signal registered by a clonotype able to interact with only a single isoform. Suppose that a thymocyte is deleted by negative selection if it registers a TCR signal greater than some deletion threshold value W_{thy} . To attain MHC restriction it suffices if the thymic deletion threshold W_{thy} is in the vicinity of the typical TCR signal registered by a clonotype able to interact with a single MHC isoform. Then negative selection will with very high probability delete any clonotype able to interact with two or more isoforms among those expressed by the dendritic cells in the thymus (Elliott, 1993; Ignatowicz et al., 1996).

We thus suggest that negative selection enforces restriction of nearly each of the naïve repertoire's clonotypes to a single MHC isoform among those present. More precisely, it enforces the “no more than one isoform” part of MHC restriction, whereas the “at least one isoform” part is mediated by positive selection. If there are n_M MHC isoforms present of the appropriate class for the TCR at hand, and p_M is the

probability that a thymocyte will interact with a randomly selected MHC isoform of the appropriate class, the number of MHC isoforms recognized by a given pre-selection thymocyte is binomially distributed with parameters p_M and n_M . Zerrahn et al. (1997) estimate p_M to be about 5.7%; this means that—depending on n_M —typically 60–90% of the preselection repertoire would be deleted in positive selection due to failure to recognize any of the n_M MHC isoforms present, and 1–10% would be deleted by negative selection on the basis of MHC promiscuity alone, not accounting for additional deletion due to autoreactivity. The probability that a mature T cell will be able to interact with a non-self MHC isoform (chosen at random from among those not encountered during thymic selection) remains $p_M \approx 5.7\%$. The mechanism we propose here leaves the potential for alloreactivity unaffected: it spares clonotypes able to interact with only one of the MHC isoforms expressed in the individual, in addition to any number of ‘non-self’ MHC isoforms expressed by other individuals in the population.

Tying together the results on low foreign presentation diversity and MHC restriction, we arrive at our main conclusion:

Immune efficacy is maximized when each clonotype is capable of interacting with a *single* MHC isoform which presents a *low diversity* of foreign peptides.

To conclude our argument, we discuss how presentation dominance influences the size of an individual’s MHC isoform repertoire.

4. Immunodetectability: the MHC isoform repertoire

We have argued that immune efficacy is enhanced by presentation dominance, composed of low presentation diversity (presentation selectivity p is such that on average each isoform presents less than a single peptide per foreign protein) and MHC restriction ($m = 1$). However, these properties impose constraints on the sizes of the MHC isoform repertoire. The essential requirement underlying the MHC isoform repertoire size n_M is that the pathogen with high probability be *immunodetectable*, that is, at least one pathogen-derived immunogenic peptide is presented. Here we analyse this requirement in more detail, and hence estimate the isoform repertoire size in terms of lower and upper bounds.

4.1. Lower bounds to the number of MHC isoforms

Suppose for the moment that the pathogen is immunodetectable only if at least one MHC isoform presents exactly one foreign peptide. It follows from the

propensity model of antigen presentation profiles (outlined in Appendix B) that the number of foreign peptides presented in the foreign subprofile of a given MHC molecule is to a very good approximation Poisson distributed with mean pN_{for} , where N_{for} is the number of potential peptides in the foreign (viral) genome. The probability $P(1)$ of presenting exactly one peptide then is

$$P(1) = pN_{\text{for}} \exp\{-pN_{\text{for}}\} \quad (15)$$

which represents the probability that the virus is immunodetectable on a given MHC isoform. This probability attains its maximum $1/e$ when $p = 1/N_{\text{for}}$; the expected number of presented peptides then equals 1. Thus the probability of not being detectable decreases with the number n_M of MHC isoforms present:

$$\begin{aligned} \mathbb{P}\{\text{immunodetectability}\} &= 1 - (1 - P(1))^{n_M} \\ &= 1 - ((e - 1)/e)^{n_M} \\ &\approx 1 - \exp\{-n_M/2\}. \end{aligned} \quad (16)$$

This estimate is based on the conservative assumption that foreign presentation subprofiles with more than one peptide will fail to be immunodetectable. In practice, it may still be possible to mount an effective immune response when more than one peptide is presented in the foreign subprofile, even though this entails a reduced separation between self-only and self-plus-foreign activation curves. For a non-conservative estimate we assume the pathogen to be immunodetectable if at least one MHC isoform presents at least one foreign peptide. In that case we obtain, again with $\mathbb{E}n_{\text{for}} = 1$:

$$\mathbb{P}\{\text{immunodetectability}\} = 1 - \exp\{-n_M\}. \quad (17)$$

Thus, for the probability of not being immunodetectable to be lower than 10^{-x} , n_M has to exceed a lower bound which the two estimates above put between $x \ln\{10\} \approx 2.3x$ and $2x \ln\{10\} \approx 5x$. The estimate $p = 1/N_{\text{for}}$ accords well with the estimate that a typical protein of ~ 400 amino acids will yield, on average, between 0.1 and 1 peptides that will bind to a given MHC isoform (Stevanović and Schild, 1999).

Eqs. (15) and (16) show that the risk of not being immunodetectable decreases exponentially with the number of distinct MHC isoforms n_M (a similar result was established by Nowak et al., 1992). This observation already furnishes a rough lower bound on $\log\{n_M\}$ between 1 and 2.

4.2. Upper bounds to the number of MHC isoforms

Next we consider what effects limit n_M from above. Obvious biological effects are the costs of carrying multiple alleles as well as physico-chemical constraints on the variety of MHC binding sites that can be generated. Although the presence of these effects cannot be doubted, they are probably not the limiting factor.

A disproportionate increase in deletion by negative selection constitutes another detrimental effect of increasing the number of MHC (Janeway and Travers, 1997; Nowak et al., 1992; Parham, 2000). However, the detrimental consequences of expressing more MHC isoforms become less severe when clonotypes show a greater degree of restriction in their ability to interact with various isoforms. As pointed out by Borghans (2000) T cells that fail to be positively selected on a particular MHC molecule run at least a lower risk, and presumably no risk at all, to be negatively selected on that MHC molecule. When each TCR clonotype is able to interact with exactly one MHC isoform, the effect vanishes altogether, and the reduction in repertoire diversity wreaked by negative selection is balanced exactly by the increase due to positive selection on more MHC isoforms; Mason (2001) calculates that the probability of surviving negative selection is not much affected if one allows for a 10% probability that a pMHC interacts with a TCR that is nominally restricted to another MHC isoform.

Yet another effect which limits the number n_M of distinct MHC isoforms is due to stochastic fluctuations in the triggering rate registered by the T cell. Such fluctuations become more important as the number of MHC molecules of a given isoform decreases. We propose that this effect imposes a rather tight bound on n_M , comparable to the lower bound established above, which strongly suggests that this may be the dominating effect in limiting n_M .

Let M_T denote the number of MHC molecules in the contact region with which the TCR is able to interact. By the argument of Section 3.2, all these MHC molecules belong to a single isoform. Our objective is to derive a lower bound on M_T , say $M_{\min}$, based on considerations of signalling noise. From MHC restriction it follows that the total number of MHC molecules is $M_T n_M$, for we may assume that every cell expresses all MHC isoforms (otherwise we only need to re-interpret n_M as the typical number of distinct MHC isoforms). However, the total number of MHC molecules is clearly physically bounded (Stevanović and Schild, 1999), say by some maximum $\hat{M}$. Thus we have $M_T n_M \leq \hat{M}$, or $n_M \leq \hat{M}/M_T \leq \hat{M}/M_{\min}$. In Appendix D we argue that $M_{\min} \sim \zeta_{\text{for}}^{-2}$, where, as above, ζ_{for} is the fraction of MHC molecules presenting foreign peptide(s). This estimate implies a lower bound on M_T between 10^3 and 10^4 . Combining this with a typical value of 10^5 for the total number of MHC molecules in the contact region $\hat{M}$ (Bongrand and Malissen, 1998), we estimate the upper bound to the number of MHC molecules expressed in an individual $\hat{M}/M_{\min}$ to lie between 10^1 and 10^2 . Since $\hat{M}$ may be rather less in some situations, the estimate of 10^1 for n_M is the more realistic one here. We assume here that the decision of a T cell to become activated depends on a substantial sample of the MHC molecules present

on the APC. This assumption is likely to be warranted for a naïve T cell in conjugation with a professional APC in a secondary lymphoid tissue; for a CTL the situation may be very different. (Whether or not the T cell will form a conjugate at all with the APC is of course dependent on a fleeting contact which is subject to a much greater fluctuations; the trade-offs involved in this zipping up decision bear upon another aspect of immune efficacy, which is how efficiently the foreign-presenting APCs are able to scan the repertoire. We will leave this aspect aside.)

The upper bound on n_M coincides with the lower estimate of the lower bound given earlier. We thus arrive at the conclusion that $\log\{n_M\} \approx 1$, arising from immunodetectability requirements at the system level (which suggest that n_M should not be smaller than this) and at the molecular level (which suggest that n_M should not be much larger than this).

In Section 2.1 we briefly alluded to another source of noise in TCR signalling: the rate given by Eq. (1) is merely an expected value about which the actual rate registered by the T cell fluctuates. The instantaneous rate at time t may be defined as the reciprocal of the time difference between the last two triggering events immediately prior to t . Noest (2000) gives an in-depth analysis of the stochastic fluctuations of this instantaneous rate about the mean. Again, the magnitude of such fluctuations is tied up with the MHC count M_T . It can be shown that the bound on M_T obtained by analysing these fluctuations is weaker (the order-of-magnitude estimate becomes $M_{\min} \sim \zeta_{\text{for}}^{-1}$; we omit the details).

5. Discussion

We have argued that the efficacy of the immune system is maximized by reducing the diversity of the foreign component. The underlying intuition is that the triggering rate elicited by the foreign peptide in the extremely rare responding clonotype is allowed to ‘stand out’ most against the background of self peptides when that particular foreign peptide dominates the foreign component in antigen presentation. The low foreign presentation diversity is brought about by (i) MHC restriction and (ii) peptide selectivity that is sufficiently stringent to allow about one peptide through per protein. Thus peptide presentation acts as a diversity filter. An adverse effect of this filter function is that a virus might go undetected altogether. We have argued that this effect is countered by the presence of multiple MHC isoforms (allelic variants) in the individual. These act in parallel, in virtue of MHC restriction. We have examined various effects that limit the MHC isoform repertoire from below, and argued that the avoidance of MHC loading fluctuations provides a tight constraint.

The following assumptions are crucial to our argument: the relevant signal registered by the T cell is a weighed sum of peptide-specific signal strengths, where the presentation densities act as weights; and the T cell responds when this signal exceeds a certain threshold. Since these seem to be rather mild assumptions, our proposal that peptide selectivity and MHC restriction act as a diversity filter is probably quite robust.

5.1. Tunable cellular activation thresholds

The argument of the present paper rests on the principle of separation of activation curves, which we discussed, in Section 2.2, in terms of the MHC-specific activation threshold W_{act} . This normalization to the MHC count implies that the activation threshold of a T cell must be dynamically adapted to MHC expression levels on APCs, which may vary widely (although only the class of recently migrated, tissue-derived APCs with upregulated costimulatory and peptide presentation molecules is relevant to the problem of commitment to clonal expansion, and variation within this class may be smaller). Tunable cellular activation thresholds have been discussed extensively by Grossman and Singer (1996) and by Grossman and Paul (2001). In fact, such a mechanism has profound consequences for the statistical structure of autoreactivity in the TCR repertoire, and in particular the interplay between central and peripheral tolerance; we analyse these consequences in detail in a forthcoming publication. In Section 2.2 we tacitly assumed that all naïve T cells in the repertoire share the same value of the MHC-specific activation threshold. In general, this assumption is not warranted, since it is likely that every clonotype adapts its cellular activation threshold to its own particular across-APP distribution of self stimulation. However, this does not affect the argument of Section 2.2, which can readily be restricted and applied separately to sub-populations of T cells that have their W_{act} -value in common.

5.2. Presentation differences between professional APCs and target cells

Low presentation diversity may, in addition to the effects considered in the present paper, contribute to the coordination of presentation on the professional APCs that instigate the immune response and the diseased cells themselves. If both professional APCs and diseased cells were presenting a diverse foreign profile, the efficacy of the responding cytotoxic T cells could readily be compromised by relatively minor differences between the professional APCs and the target cells in the affected tissue. On the other hand, if the foreign component is dominated by the same single peptide on both cell types, shifts in expression patterns are far less likely to affect recognition by the cytotoxic T cells. The main factor

which determines whether the T cells will be effective in killing the diseased cells is whether presentation patterns are comparable on professional APCs and target cells (Byers and Lindahl, 1999; Zeh III et al., 1999).

5.3. Class I vs Class II

In principle, our argument applies to class I as well as class II recognition. However, in the case of the latter there may well be other effects which influence the level of presentation selectivity (such as DO expression; van Ham et al., 2000). Immune responses based on class II recognition involve interactions with other components of the immune system, such as professional antigen presenting cells. The B cells, in particular, may add a second layer of self vs non-self discrimination.

5.4. Experimental validation

Our argument is essentially a thought experiment in which variation of the parameters p and m leads to the conclusion that immune efficacy is best served when $p \approx 1/N_{\text{for}}$ and $m \approx 1$. These parameters are more or less given in the actual immune system, which poses some problems for direct experimental validation of the argument. Nevertheless, it should be possible to demonstrate the importance of presentation diversity in experiments where controlled antigen presentation profiles are set up on lipid bilayers (Grakoui et al., 1999) or peptide-pulsed APCs (Inaba et al., 1990). To ensure that the observed differences can be ascribed to the diversity effect, care must be taken to ensure that ζ_{for} , the total relative presentation level of the foreign peptides taken together, remains more or less equal among the experiments that are compared. Such precaution is necessary since the individual presentation levels also influence the activation probability P_{act} , which in turn governs response characteristics such as time until a clone has been activated and the quality of the responding clones (as is well attested: Byers and Lindahl, 1999; Reay et al., 2000; Zeh III et al., 1999). An in vivo experiment using pulsed APCs could give the most direct indication that our arguments on presentation diversity are relevant to the actual situation. Alternatively, in vitro experiments can more directly explore the relationship between presentation diversity and triggering rate.

5.5. Immunodominance

It stands to reason that presentation dominance contributes to immunodominance, the phenomenon that most of the responding effector cells recognize only one or a few foreign epitopes. Effects downstream from recognition events, such as competition between responding clonotypes, may also contribute to

immunodominance. The significance of presentation dominance may be questioned on the grounds that peptide-MHC recognition constraints together with competition for loading are simply an intrinsic physico-chemical property of the peptide presentation system. In essence, this argument attacks the treatment of p as an adjustable parameter. However, we do not require that much; our treatment of p as an adjustable parameter serves merely to bring out the functional significance, in terms of immune efficacy, of p being of order $1/N_{\text{for}}$.

5.6. Pathogen escape

We have treated N_{for} as a given, and argued that a selectivity p of order $1/N_{\text{for}}$ maximizes the diversity effect. Of course, pathogens differ in genome size and hence in their peptide count N_{for} , so that p cannot equal $1/N_{\text{for}}$ for all pathogens. Rather perversely, pathogens which are large in the sense that $N_{\text{for}} > p^{-1}$ may confound immune efficacy by increasing foreign diversity, although they will of course also be more likely to be immunodetectable. On the other hand, pathogens for which $N_{\text{for}} < p^{-1}$ have an increased likelihood of escaping immunodetection altogether. Therefore we would expect the actual p^{-1} to reflect the smallest pathogens.

A related issue concerns the set of ligands (self or foreign) that are excluded by the presentation filter. This set should not pose an evolutionary target for pathogens. Each MHC allele is associated with its own particular set of non-immunodetectable ligands, and this set becomes larger as p decreases, which might set another lower bound on MHC selectivity p . As long as the population does not segregate itself into communities of people who tend to share the same MHC alleles, the evolutionary target for a pathogen must consist of the *cross-section* of the non-immunodetectable sets associated with each of the MHC alleles present in the population. It is clear that the size of this cross-section decreases exponentially with this number, which points to the evolutionary advantage of MHC polymorphism.

5.7. Diversity filtering of the self background

The filter function of peptide selectivity combined with MHC restriction affects the self background as well as foreign peptides: indeed a cornerstone of our theory is that self peptides are not a priori biochemically distinguishable from foreign ones. Therefore the filter function may have the adverse effect that the self background activation curve widens, thus diminishing the separation between self only and self plus foreign curves. However, large deviations theory shows that the probability that the triggering rate due to a given subprofile exceeds a given value goes down exponen-

tially with the diversity of that subprofile, and thus this adverse effect may be ignored as long as the diversity of the self background is several orders of magnitude higher than that of the foreign component. Moreover, the effect of negative selection is to move the self background activation curve ‘to the left’ (see Section 2.2).

5.8. MHC restriction and thymic selection

The TCR makes contacts not only with its peptide ligand, but with the MHC molecule as well (Janeway and Travers, 1997). Thus MHC restriction comes about naturally. It is possible that the physico-chemical constraints on TCR/pMHC interaction are relatively lax, such that the ability to interact with more than one isoform is relatively common among thymocyte TCRs (the very existence of alloreactivity would suggest that this is the case; see Detours et al., 1999, for a quantitative analysis). A maximum-likelihood estimate for the probability that a given pre-selection TCR will recognize a randomly selected MHC isoform is 5.7%, based on the data of Zerrahn et al. (1997). These authors also found reactivities to MHC isoforms to be statistically independent, which indicates that up to about 10% of preselection thymocytes recognizes two or more self MHC isoforms. Such thymocytes would be removed by negative selection, by the mechanism discussed in Section 3.2. Thus the ability of the TCR to interact productively with exactly one MHC isoform among those present in the individual depends on both positive and negative selection. Obst et al. (2000) found that T cells tend to be more strongly alloreactive against non-self MHC molecules when the latter are structurally closer to the self MHC molecules.

5.9. Optimal dissociation time

The existence of an optimal average TCR/peptide/MHC dissociation time (Fig. 1) may explain the seeming paradox of low TCR affinity (van den Berg et al., 2001; Lanzavecchia et al., 1999; Mason, 1998). Briefly, the number of triggering events during a period ΔT will only on average equal $W\Delta T$ where W is the triggering rate as given by Eq. (1). The stochasticity inherent in a process that relies on encounters between receptor molecules gives rise to fluctuations. These triggering fluctuations are distinct from the MHC loading fluctuations discussed above. The coefficient of variation of the triggering fluctuations is of order $1/\sqrt{W\Delta T}$, which indicates that triggering fluctuations become less important as the expected rate W increases. Since the latter is of order $1/T_R$ for an optimal agonist, we may expect T_R to be modest, which in fact it is: less than 10 s (Davis et al., 1998; Hudrisier et al., 1998).

5.10. The size of the TCR clonotype repertoire

In Section 2.2 we discussed $1/P_{\text{act}}$ as a lower bound to the TCR clonotype repertoire size n_C (Fig. 2, top panel). The activation probability $P_{\text{act}} = P(W_{\text{act}}|\zeta_{\text{for}} > 0)$ is the proportion of naïve clonotypes that will respond to the foreign peptide(s). We may calculate P_{act} as a mixed distribution, as in the example of Section 3.2, where we need to supply a model of the foreign presentation profile, (e.g. the propensity model, Appendix B) as well as a distribution of isoform specificities among the (restricted) naïve TCR repertoire. Alternatively, following the considerations of immunodetectability (see Section 4), we may conservatively assume that only one MHC isoform (the ℓ th, say) is presenting a foreign peptide, and compute P_{act} conditional on the clonotype being able to interact with that isoform. That is, the conditional probability $P_{\text{act}|\ell}$ expresses the fraction of responding clones among the TCR subrepertoire restricted to the isoform at hand. Then $m/P_{\text{act}|\ell}$ represents the lower bound to the TCR clonotype repertoire size n_C . The total TCR repertoire size has been estimated as $\sim 10^7$ (Arstila et al., 1999), which would mean that the size of a subrepertoire devoted to a single MHC isoform is about 10^5 – 10^6 , suggesting a typical operating value of $P_{\text{act}|\ell}$ in the range 10^{-4} – 10^{-5} .

Acknowledgements

We gratefully acknowledge the contributions made by José Borghans, Nigel Burroughs, Carmen Molina-París, Don Mason, Martin Nowak, and two referees.

Appendix A

A.1. TCR triggering rate model

TCR triggering occurs when the TCR molecule forms a ternary complex with a pMHC molecule. We assume that the rate of formation of ternary complexes composed of a TCR of clonotype i and a pMHC molecule of species j is proportional to R , the number of unbound TCR receptors in the region of contact between T cell and antigen presenting cell, and also to M_j , the number of MHC molecules which (i) present peptide species j ; (ii) are capable of interacting with the TCR; and (iii) are not already bound to a TCR molecule. Thus the on-rate is $k_{ij}RM_j$, where k_{ij} is a rate constant specific for the clonotype i and the pMHC species j at hand. Furthermore, we assume that the ternary complex composed of TCR i and pMHC j can be characterized by an average lifetime T_{ij} . Thus the dissociation rate is C_j/T_{ij} , where C_j denotes the number of MHC molecules in the contact region which present

peptide species j and are bound to a TCR molecule. At kinetic equilibrium, association and dissociation rates are equal:

$$k_{ij}RM_j = C_j/T_{ij}. \quad (\text{A.1})$$

The rate at which TCR molecules are triggered by pMHC species j is the rate of association/dissociation given by Eq. (A.1) times the probability P_{ij} that the ternary complex has resulted in a TCR triggering event (we assume that no more than one triggering event can happen during the lifetime of an individual ternary complex). Thus the triggering rate associated with pMHC species j is

$$P_{ij}C_j/T_{ij}.$$

To calculate P_{ij} , we assume that a ternary complex formation results in the triggering of the TCR/CD3 complex if the complex has lasted longer than a *receptor threshold duration* T_R . A sharp threshold T_R exists if the activation of the TCR/CD3 complex is a steep S-shaped function of the life time of the TCR/pMHC complex: such strongly non-linear behaviour can be justified by postulating ‘kinetic proofreading’ properties for CD3/ ζ -chain phosphorylation, or, alternatively, that the balance between activating and deactivation kinase phosphorylation is tipped towards activation through a positive feedback loop. If lifetimes of the ternary complex are exponentially distributed, the probability that a given ternary complex formation results in triggering is $P_{ij} = \exp\{-T_R/T_{ij}\}$. The total triggering rate is the sum over all pMHC species:

$$W_i = \sum_j \exp\{-T_R/T_{ij}\} C_j/T_{ij}. \quad (\text{A.2})$$

From this sum we may exclude those pMHC species with which the TCR is unable to interact (that is, those species for which $k_{ij} = 0$). It cannot be ruled out in general that a TCR clonotype is able to interact with a number of different MHC isoforms (Janeway and Travers, 1997). In fact, since the main objective of this paper is to explain why immune efficacy requires MHC restriction, our point of departure must be a complete lack of such restriction, that is, we assume MHC promiscuity to begin with.

The sum $C_j + M_j$ is Z_j the total number of MHC molecules presenting peptide j among the MHC molecules capable of interaction with the TCR. Combining this conservation law with Eq. (A.1), we obtain the triggering rate Eqs. (1) and (2). These equations still contain the free receptor count R , implicitly related to the total (free plus bound) R_T by

$$R_T = R \left(1 + \sum_j Z_j / (R + (k_{ij}T_{ij})^{-1})^{-1} \right), \quad (\text{A.3})$$

which has a unique solution for R .

In the simulations shown in Figs. 3–5 of this paper, we have assumed that the T_{ij} for a given clonotype are drawn from an exponential distribution: $\mathbb{P}\{T_{ij} > t\} = \exp\{-t/\bar{T}\}$ where $\bar{T}$ is the expected value, which we took to equal $T_R/10$ in the simulations shown (the results of the present paper do not depend materially on either the choice of this value or the choice of the exponential distribution). The distribution of the triggering rate was derived from the distribution of the T_{ij} by means of Eq. (2), on the assumption that the free TCR count was not limiting (that is, $R \gg (k_{ij}T_R)^{-1}$). The latter assumption is explored in full depth in van den Berg et al. (2002).

Appendix B. Constancy of the foreign presentation level

We explain the propensity model of antigen presentation in more detail and establish that the foreign presentation level ζ_{for} is independent of the selectivity of presentation.

Of all MHC isoforms present in the population, only a limited number (n_M) is expressed in the host. Let $m \leq n_M$ denote the number of MHC isoforms which are (i) expressed by the host, and (ii) typically capable of interacting with a given TCR clonotype. For the sake of exposition, we first assume $m = 1$, that is, all M_T MHC molecules which are capable of interacting with the TCR of the clone at hand belong to a common isoform.

In an earlier paper (van den Berg et al., 2001) we introduced the concept of *presentation propensity* to describe APP variations between APCs with varying protein expression patterns. The relative presentation level of a peptide depends on a number of factors: whether the peptide is expressed in the APC; the level at which it is expressed (Jardetzky et al. 1991); the likelihood that the peptide is properly processed (protein degradation in the proteasome, transport to the intracellular compartment containing the MHC molecules; and its affinity for the MHC binding cleft, which determines its ability to compete with other peptides. We combine these factors into a single, aggregate ‘propensity’ $r \geq 0$. The key idea is that the ratio between presentation levels of every pair of expressed peptides j and j' is given by the ratio of their propensities, $\zeta_j/\zeta_{j'} = r_j/r_{j'}$ where r_j and $r_{j'}$ are the propensities of the two peptides, each of which is only presented on a single MHC isoform. Thus peptides with equal propensities have equal presentation levels, when both are expressed. It follows that the relative presentation level ζ_j of peptide j is found by dividing its propensity r_j by the sum of the propensities of all peptides that are presently expressed in the APC. It may be expected that the presentation propensity r_j will be minute for the vast majority of peptides. (Only ratios of the form $r_j/r_{j'}$ are physically relevant and are interpreted as presentation

ratios $\zeta_j/\zeta_{j'}$: as long as the propensities assigned to the peptides conform to these ratios, the question of a unit is immaterial.)

Peptides unable to bind to the MHC cleft have zero propensity $r_j = 0$ for the MHC isoform at hand. Let p denote the probability that r_j will be non-zero with respect to a given MHC isoform; we call p the *presentation selectivity*. The mean number of peptides present in the self background subprofile then equals $n_{\text{bg}} = pN_{\text{bg}}$, where N_{bg} is the total number of potential peptides expressed in the APC. Similarly, the mean number of peptides present in the foreign subprofile equals $n_{\text{for}} = pN_{\text{for}}$ where N_{for} is the number of potential peptides in the foreign genome. The latter is a characteristic of the pathogen, which we have to regard as given.

Let $\bar{r}_{\text{bg}}$ denote the mean propensity among the self peptides, and $\bar{r}_{\text{for}}$ the mean propensity among the foreign peptides (the two may generally be different due to differing typical expression levels of the proteins involved). Then in infected cells we have

$$\zeta_{\text{for}} \approx \frac{\bar{r}_{\text{for}}N_{\text{for}}}{\bar{r}_{\text{bg}}N_{\text{bg}} + \bar{r}_{\text{for}}N_{\text{for}}} \quad (\text{B.1})$$

The crucial assumption now is that $\bar{r}_{\text{for}}/\bar{r}_{\text{bg}}$ is independent of p . This simply means that the distributions of chemical characteristics relevant to MHC binding do not differ between self peptides and foreign peptides, and that therefore these distributions are altered in the same way whenever p changes. On this assumption, the ratio $\bar{r}_{\text{for}}/\bar{r}_{\text{bg}}$ merely reflects differences between the foreign and self peptides as regards their concentrations in the MHC loading compartment. It then follows that ζ_{for} is independent of presentation selectivity p .

We have assumed $m = 1$ here. It is straightforward to extend the present argument to the general case, where the M_T MHC molecules capable of interacting with the TCR belong to m different isoforms. If ζ_ℓ is the fraction of these M_T MHC molecules belonging to isoform ℓ , then the formula for ζ_{for} in infected cells becomes

$$\zeta_{\text{for}} \approx \sum_{\ell=1}^m \zeta_\ell \frac{\bar{r}_{\text{for},\ell}N_{\text{for}}}{\bar{r}_{\text{bg},\ell}N_{\text{bg}} + \bar{r}_{\text{for},\ell}N_{\text{for}}}$$

and the argument proceeds along the same lines; here $\bar{r}_{\text{for},\ell}$ denotes the mean propensity of the foreign peptides for MHC isoform ℓ . Moreover, it would seem reasonable to assume $\bar{r}_{\text{bg},\ell} \equiv \bar{r}_{\text{bg}}$ and $\bar{r}_{\text{for},\ell} \equiv \bar{r}_{\text{for}}$ and Eq. (B.1) is recovered.

The constancy of ζ_{for} with respect to presentation selectivity p should not be confused with variability of ζ_{for} over different MHC isoforms which occurs when the host is confronted with a given viral strain (that is, a given instance of N_{for}). For the sake of simplicity, we return to the case where $m = 1$ and, furthermore, we introduce the idealization that the propensity r_j of a peptide j is, with probability p , equal to $\bar{r}_{\text{bg}}$, respectively

$\hat{r}_{\text{for}}$, or 0 (with probability $1 - p$). Then we have

$$\zeta_{\text{for}} = \frac{n_{\text{for}}/n_{\text{bg}}}{\hat{r}_{\text{bg}}/\hat{r}_{\text{for}} + n_{\text{for}}/n_{\text{bg}}}. \quad (\text{B.2})$$

When $n_{\text{for}}/n_{\text{bg}} \ll \hat{r}_{\text{bg}}/\hat{r}_{\text{for}}$ the total foreign presentation level ζ_{for} is proportional to the number of foreign peptides presented on the isoform.

A difference between $\bar{r}_{\text{bg}}$ and $\bar{r}_{\text{for}}$ need not reflect intrinsic biochemical differences between self background and foreign peptides, which the MHC molecule would somehow be able to detect, but may instead simply reflect differences in concentrations in the MHC loading compartment. There is some evidence which suggests that such propensity differences may be real. Whereas self peptides are seldom, if ever, expressed at relative presentation levels above 1% (Engelhard, 1994; Hunt et al., 1992), foreign peptides may rise to levels of 5% or even higher (van Bleek and Nathenson, 1990; Falk et al., 1991b; Rudensky et al., 1991; Stevanović and Schild, 1999). The differences may be due to high expression levels of viral proteins per se, or some bias toward foreign products in the presentation pathways (Engelhard, 1994). There is ample evidence for such a bias in dendritic cells, which are professional APCs that initiate the immune response (Hartgers et al., 2000; Kleijmeer et al., 1995; Kurts et al., 1998; Puré et al., 1990; Thomas and Lipsky, 1996; Steinman, 1991).

Appendix C. Large deviations estimates

Large deviations theory can be used to furnish some useful estimates of how the foreign triggering probability varies with the diversity of the foreign presentation subprofile. Let us consider the foreign triggering probability

$$\mathbb{P}\{w_{\text{for}} > \zeta_{\text{for}} w\} = \mathbb{P}\left\{\sum_{j \in \text{for}} (\zeta_j / \zeta_{\text{for}}) w_{ij} > w\right\} \quad (\text{C.1})$$

(the reader may want to note that $\zeta_j / \zeta_{\text{for}} = Z_j / Z_{\text{for}}$ where $Z_{\text{for}} \stackrel{\text{def}}{=} \sum_{j \in \text{for}} Z_j$). Evaluation of this probability involves a convolution of the distribution F_W of the individual triggering rates w_{ij} with the $\zeta_j / \zeta_{\text{for}}$ as weighing coefficients. The distribution F_W can be derived from the triggering rate model, Eq. (2), combined with the distribution of dissociation times

$$F_W(w) \stackrel{\text{def}}{=} \mathbb{P}\{w_{ij} \leq w \mid i, j \text{ chosen at random}\} \\ = 1 - \mathbb{P}\{T_{\text{low}} < T_{ij} < T_{\text{upp}}\}, \quad (\text{C.2})$$

where T_{low} and T_{upp} are, respectively, the lower and upper solutions for T of the equation $w = \exp\{-T_R/T\}/T$. We denote the mean and variance of the distribution F_W by μ_W and σ_W^2 .

Let $F_\Sigma(\cdot)$ and $M_\Sigma(\cdot)$ denote the distribution function and the moment generating function, respectively, of

the sum

$$\sum_{j \in \text{for}} (\zeta_j / \zeta_{\text{for}}) w_{ij}.$$

Of fundamental importance is the *exponential overbound* estimate:

$$\ln\{\mathbb{P}\{w_{\text{for}} > \zeta_{\text{for}} w\}\} \leq S(w) \quad \text{where} \\ S(w) \stackrel{\text{def}}{=} \inf_{t>0} \{-tw + \ln\{M_\Sigma(t)\}\}. \quad (\text{C.3})$$

This is readily derived, as we may write, for $t > 0$:

$$\mathbb{P}\{w_{\text{for}} > \zeta_{\text{for}} w\} \\ = \int_w^{(eT_R)^{-1}} dF_\Sigma(x) = e^{-tw} \int_w^{(eT_R)^{-1}} e^{tx} dF_\Sigma(x) \\ \leq e^{-tw} \int_w^{(eT_R)^{-1}} e^{tx} dF_\Sigma(x) \\ \leq e^{-tw} \int_0^{(eT_R)^{-1}} e^{tx} dF_\Sigma(x) = e^{-tw} M_\Sigma(t). \quad (\text{C.4})$$

By Cramér's Theorem (Dembo and Zeitouni, 1998) the exponential overbound becomes exact as we take n_{for} to infinity. This is why Eq. (9) establishes the relationship between presentation diversity and triggering probability for sufficiently large n_{for} . Cramér's result obviously does not apply for small n_{for} , of order one, but we may still use the exponential overbound to gain insight into the relationship between diversity and triggering probability.

C.1. Second-order approximation to the exponential overbound

By Eq. (C.3) the exponential overbound $S(w)$ equals $\ln M_\Sigma(t) - tw$, where t satisfies

$$\frac{M'_\Sigma(t)}{M_\Sigma(t)} = w. \quad (\text{C.5})$$

Note that $S(\mu_W) = S(w)|_{t=0} = 0$ and $M'_\Sigma(t)/M_\Sigma(t)|_{t=0} = \mu_W$. Now

$$\frac{d}{dt} S(w) = \frac{M'_\Sigma(t)}{M_\Sigma(t)} - w - t \frac{dw}{dt} = -t \frac{dw}{dt}$$

by Eq. (C.5), so that $S'(w) = -t$ and $S''(w) = -dt/dw$. Evaluating these derivatives at $w = \mu_W$ we find $S'(\mu_W) = -t|_{t=0} = 0$ and

$$S''(\mu_W) = -\left(\frac{dw}{dt}\right)^{-1} \Big|_{t=0} = -\left(\frac{M''_\Sigma(t)}{M_\Sigma(t)} \Big|_{t=0} - \mu_W^2\right)^{-1} \\ = -\frac{n_{\text{div, for}}}{\sigma_W^2}$$

with $n_{\text{div, for}}$ as defined by Eq. (8). Thus the second-order Taylor expansion of the large deviations rate function

about $w = \mu_W$ is

$$S(w) \approx -n_{\text{div, for}} \frac{(w - \mu_W)^2}{2\sigma_W^2}. \quad (\text{C.6})$$

This is the estimate used in Eq. (9).

C.2. Diversity and proportional displacement

We can change the diversity of the foreign presentation subprofile by means of an arbitrary alteration of the relative presentation level of one of the peptides in the subprofile, where the remaining peptides in the subprofile are redistributed among the remaining MHC molecules devoted to the subprofile according to their original proportions. Such a *proportional displacement* can be represented as follows:

$$\zeta_j = \begin{cases} z_k \zeta_{\text{for}}, & j = k, \\ z_j^\circ (1 - z_k) \zeta_{\text{for}}, & j \in \text{for} \setminus k, \end{cases} \quad (\text{C.7})$$

where k can be any one of the peptides in the foreign subprofile. The $\{z_j^\circ\}_{j \in \text{for} \setminus k}$ represent the relative proportions among the remaining peptides; these parameters are all non-negative and satisfy $\sum_{j \in \text{for} \setminus k} z_j^\circ = 1$. They can

be derived from a given APP by $z_j^\circ \stackrel{\text{def}}{=} \zeta_j / (\zeta_{\text{for}} - \zeta_k)$. All presentation levels are enslaved to $z_k \in [0, 1]$ by Eq. (C.7); $z_k = 0$ corresponds to proportional displacement of peptide k from the subprofile, while $z_k = 1$ corresponds to the situation where peptide k completely dominates the subprofile: $\zeta_k = \zeta_{\text{for}}$. To see how such proportional displacement affects the foreign triggering probability, we define

$$P(w, z_k) \stackrel{\text{def}}{=} \mathbb{P}\{w_{\text{for}} > \zeta_{\text{for}} w\}. \quad (\text{C.8})$$

With $M_\Sigma(t) = \prod_{j \in \text{for}} M_W(z_j t)$, where $M_W(\cdot)$ is the moment generating function of F_W , the exponential overbound estimate, Eq. (C.3), becomes

$$\ln\{P(w, z_k)\} \leq \inf_{t>0} \left\{ -tw + \ln\{M_W(z_k t)\} + \sum_{j \in \text{for} \setminus k} \ln\{M_W(z_j t)\} \right\}. \quad (\text{C.9})$$

For $w > \mu_W$ the infimum can be found by minimizing the above expression with respect to t , which gives an implicit definition of t as a function of w and of z_k :

$$w = z_k H_W(z_k t) + (1 - z_k) \sum_{j \in \text{for} \setminus k} z_j^\circ H_W(z_j^\circ (1 - z_k) t), \quad (\text{C.10})$$

where $H_W(t) \stackrel{\text{def}}{=} M'_W(t)/M_W(t)$. Let $S(w, z_k)$ denote the infimum of Eq. (C.9), with t as defined by Eq. (C.10), so that $P(w, z_k) \leq \exp\{S(w, z_k)\}$.

We can gain some insight into the behaviour shown in Figs. 3 and 4 by considering the derivative of $S(w, z_k)$ with respect to z_k :

$$\frac{dS}{dz_k} = t \left(H_W(z_k t) - \sum_{j \in \text{for} \setminus k} z_j^\circ H_W(z_j^\circ (1 - z_k) t) \right).$$

The exponential overbound has a minimum ($dS/dz_k = 0$) at $z_k = \bar{z}_k$ and $t = \bar{t} \stackrel{\text{def}}{=} t(\bar{z}_k)$ which implies, together with $t > 0$ and Eq. (C.10),

$$w = H_W(\bar{z}_k \bar{t}) \quad \text{and} \quad w = \sum_{j \in \text{for} \setminus k} z_j^\circ H_W(z_j^\circ \bar{t} - z_j^\circ \bar{z}_k \bar{t}).$$

Since H_W is monotone increasing, the first condition fixes a unique value of the product $\bar{z}_k \bar{t}$, whereupon the second fixes a unique value of $\bar{t}$. The exponential overbound thus has a unique minimum at $\bar{z}_k$. An explicit estimate for $\bar{z}_k$ can be obtained using the first-order approximation

$$H_W(t) \approx \mu_W + \sigma_W^2 t,$$

which gives

$$\bar{z}_k \approx \frac{\sum_{j \in \text{for} \setminus k} z_j^{\circ 2}}{1 + \sum_{j \in \text{for} \setminus k} z_j^{\circ 2}}. \quad (\text{C.11})$$

It is natural to suppose that the z_k -value at which $P(w, z_k)$ attains a minimum with respect to z_k is approximately equal to $\bar{z}_k$. We may rewrite Eq. (C.11) in terms of the diversity of presentation among the peptides other than k in the foreign subprofile. This diversity is defined as follows:

$$n_{\text{div, for}}^\circ(k) \stackrel{\text{def}}{=} \frac{1}{\sum_{j \in \text{for} \setminus k} z_j^{\circ 2}} \quad (\text{C.12})$$

and our estimate becomes

$$\bar{z}_k \approx \frac{1}{1 + n_{\text{div, for}}^\circ(k)}, \quad (\text{C.13})$$

which shows that $\bar{z}_k$ varies between $1/n_{\text{for}}$ (when the remaining foreign peptides are equidistributed) and $\frac{1}{2}$ (when one other foreign peptide dominates the remainder of the foreign subprofile). Figs. 3 and 4 suggest that the estimate of Eq. (C.13) is fairly good, although at w -values near the maximum $(eT_R)^{-1}$ the actual value of $\bar{z}_k$ tends to be somewhat smaller. Similar results are found for $n_{\text{for}} = 4$: a maximum ($n_{\text{div, for}}^\circ(k) = 3$), intermediate ($n_{\text{div, for}}^\circ(k) \doteq 2.31$), and low ($n_{\text{div, for}}^\circ(k) \doteq 1.85$) diversity case are shown in Fig. 4.

If some foreign pMHC species k is such that $\zeta_k \leq \zeta_j$ for all foreign pMHC species j , we have

$$z_k \leq (1 - z_k) \bar{z}_k^\circ,$$

where $\tilde{z}^\circ \stackrel{\text{def}}{=} \min_{j \in \text{for} \setminus k} \{z_j^\circ\}$. Thus

$$\begin{aligned} z_k &\leq (1 - z_k) \tilde{z}^\circ \sum_{j \in \text{for} \setminus k} z_j^\circ \\ &= (1 - z_k) \sum_{j \in \text{for} \setminus k} z_j^\circ \tilde{z}^\circ \leq (1 - z_k) \sum_{j \in \text{for} \setminus k} z_j^{\circ 2} \end{aligned}$$

which implies $z_k \leq \bar{z}_k$, if we take the estimate, Eq. (C.13), for the latter. Putatively identifying $\bar{z}_k$ with the z -value where $P(w, z_k)$ attains its minimum, we conclude that proportional displacement of the foreign pMHC species with the lowest presentation level leads to an increase of the foreign triggering probability. Figs. 3 and 4 suggest that this effect can be quite significant. We may try to estimate this factor by comparison of the exponential overbounds, evaluating Eq. (C.6) at $z_k = 0$ and $z_k = \bar{z}_k$. For $z_k = 0$ we have $n_{\text{div, for}} = n_{\text{div, for}}^\circ(k)$, whereas for $z_k = \bar{z}_k$ we have $n_{\text{div, for}} = n_{\text{div, for}}^\circ(k) + 1$, using approximation (C.13) for $\bar{z}_k$. We thus obtain

$$\begin{aligned} &\ln \frac{\mathbb{P}\{w_{\text{for}} > \zeta_{\text{for}} w \mid z_k = 0\}}{\mathbb{P}\{w_{\text{for}} > \zeta_{\text{for}} w \mid z_k = \bar{z}_k\}} \\ &= \ln \frac{P(w, 0)}{P(w, \bar{z}_k)} \\ &\approx S(w, 0) - S(w, \bar{z}_k) \approx \frac{(w - \mu_w)^2}{2\sigma_w^2} \end{aligned} \quad (\text{C.14})$$

whence Eq. (11).

Since $dn_{\text{div, for}}/dz_k \geq 0$ according as $z_k \leq 1/n_{\text{div, for}}$ in proportional displacement, we would require $P(w, z_k)$ to attain its minimum at $z_k = 1/n_{\text{div, for}}$ for a general inverse relation between presentation diversity and triggering probability. However, this is not to be expected since we have, for the estimate of Eq. (C.13), $(1 + n_{\text{div, for}}^\circ(k))^{-1} \geq 1/n_{\text{div, for}}$ with equality when $z_k = (1 + n_{\text{div, for}}^\circ(k))^{-1}$, which precludes us from establishing such a general result. However, what we have established already suffices: any foreign subprofile dominated by some pMHC species s can be converted into any other APP in which ζ_s is even greater than before, by a series of successive proportional displacements, during each of which $n_{\text{div, for}}$ decreases and $P(w, z_k)$ increases.

Appendix D. MHC loading fluctuations

We consider the fluctuations in the TCR signal registered by a T cell belonging to a fixed clonotype, as that cell interacts with various different APCs. One source of such fluctuations is peptide loading onto MHC molecules, which is stochastic by nature. In particular, whereas ζ_j has thus far been discussed as the fraction of MHC molecules occupied by peptide species j , it should more properly be regarded as the probability that peptide j will be presented by a given MHC molecule of the isoform at hand. If n_{pep} peptides are competing for occupation of the M_T MHC molecules

with which the clonotype can interact, the antigen presentation profile is a multinomial variate parametrized by n_{pep} and the $\{\zeta_j\}$. Thus far we have neglected the fluctuations associated with MHC loading, on the tacit assumption that M_T would be sufficiently large (justified in van den Berg et al., 2001). However, we now require a sharper characterization of ‘sufficiently large,’ as this will lead us to an expression for M_{min} . To this end, we consider the presentation levels as averages over APCs, and, accordingly, the TCR triggering rate $W = \sum_j Z_j w_j$ as an average over APCs. The criterion which furnishes an estimate for W_{min} (given as Eq. (D.4)) compares the expected signal over APCs with the standard deviation over APCs of the TCR signal W . We denote this standard deviation by S_{APC} .

To analyse MHC loading fluctuations we calculate the variance of W , which is a sum over M_T terms, a fraction ζ_j of which are expected (over the APCs) to equal w_j . Therefore, the variance of W equals M_T times the variance of an individual term. The latter is a variate which takes on the value w_j with probability ζ_j , and which thus has expectation

$$\zeta_{\text{for}} w_{\text{for}} + \sum_{j \in \text{bg}} \zeta_j w_j$$

and second moment

$$\zeta_{\text{for}} w_{\text{for}}^2 + \sum_{j \in \text{bg}} \zeta_j w_j^2;$$

here we have assumed that an MHC presenting a foreign peptide contributes an amount w_{for} to the TCR signal. Since the clonotype is fixed, the $\{w_j\}$ are now parameters rather than random variables as before. Thus

$$\begin{aligned} S_{\text{APC}}^2 &= M_T \left[\zeta_{\text{for}} w_{\text{for}}^2 + \sum_{j \in \text{bg}} \zeta_j w_j^2 \right. \\ &\quad \left. - \left(\zeta_{\text{for}} w_{\text{for}} + \sum_{j \in \text{bg}} \zeta_j w_j \right)^2 \right]. \end{aligned} \quad (\text{D.1})$$

The relevant typical value of this variance is found by taking its expectation over all those clonotypes that share a common TCR signal W_{for} against the foreign component of the APP. Let $\bar{w}_{\text{bg}}$ and s_{bg}^2 express mean and variance, respectively, of a self peptide contribution w_j , taken over such clonotypes. It is reasonable to assume that $\bar{w}_{\text{bg}}$ and s_{bg}^2 are both independent of j . Finally, we need the diversity of the self background

$$n_{\text{div, bg}} \stackrel{\text{def}}{=} \frac{(1 - \zeta_{\text{for}})^2}{\sum_{j \in \text{bg}} \zeta_j^2}. \quad (\text{D.2})$$

Straightforward calculations then give the following estimate for loading noise:

$$S_{APC}^2 \approx M_T(1 - \zeta_{\text{for}})(s_{\text{bg}}^2(1 - (1 - \zeta_{\text{for}})/n_{\text{div,bg}}) + \zeta_{\text{for}}(w_{\text{for}} - \bar{w}_{\text{bg}})^2). \quad (\text{D.3})$$

The expected difference between foreign and background signals must exceed S_{APC} if a responsive clonotype is to register a consistent stimulation from various APCs:

$$M_T(\zeta_{\text{for}}w_{\text{for}} - (1 - \zeta_{\text{for}})\bar{w}_{\text{bg}}) \gg S_{APC}. \quad (\text{D.4})$$

Combining this criterion with the estimate for S_{APC} , Eq. (D.3), we obtain

$$\sqrt{M_T} \gg \frac{\sqrt{(1 - \zeta_{\text{for}})(s_{\text{bg}}^2(1 - (1 - \zeta_{\text{for}})/n_{\text{div,bg}}) + \zeta_{\text{for}}(w_{\text{for}} - \bar{w}_{\text{bg}})^2)}}{\zeta_{\text{for}}w_{\text{for}} - (1 - \zeta_{\text{for}})\bar{w}_{\text{bg}}}. \quad (\text{D.5})$$

This sets a lower bound to $\sqrt{M_T}$ of order ζ_{for}^{-1} (take $s_{\text{bg}} \sim w_{\text{for}}$ and $w_{\text{for}} \gg \bar{w}_{\text{bg}}$, with $1 - \zeta_{\text{for}} \approx 1$ and $n_{\text{div,bg}} \gg 1$). The resulting order-of-magnitude estimate is $M_{\text{min}} \sim \zeta_{\text{for}}^{-2}$, which is what was required.

References

- Arstila, T.P., Casrouge, A., Baron, V., Even, J., Kanellopoulos, J., Kourilsky, P., 1999. A direct estimate of the human $\alpha\beta$ T cell receptor diversity. *Science* 286, 958–961.
- van den Berg, H.A., Rand, D.A., Burroughs, N.J., 2001. A reliable and safe T cell repertoire based on low-affinity T cell receptors. *J. Theor. Biol.* 209, 465–486.
- van den Berg, H.A., Burroughs, N.J., Rand, D.A., 2002. Quantifying the strength of ligand antagonism in TCR triggering. *Bull. Math. Biol.* 64, 781–808.
- van Bleek, G.M., Nathanson, S.G., 1990. Isolation of an endogenously processed immunodominant viral peptide from the class I h-2k^b molecule. *Nature* 348, 213–216.
- Bongrand, P., Malissen, B., 1998. Quantitative aspects of T-cell recognition: from within the antigen-presenting cell to within the T cell. *BioEssays* 20, 412–422.
- Borghans, J., 2000. Diversity in the immune system. Ph.D. Thesis, University of Utrecht.
- Byers, D.E., Lindahl, K.F., 1999. Peptide affinity and concentration affect the sensitivity of M3-restricted CTLs in vitro. *J. Immunol.* 163, 3022–3028.
- Chicz, R.M., Urban, R.G., Gorga, J.C., Vignali, D.A.A., Lane, W.S., Strominger, J.L., 1993. Specificity and promiscuity among naturally processed peptides bound to HLA-DR alleles. *J. Exp. Med.* 178, 27–47.
- Davis, M.M., Boniface, J.J., Reich, Z., Lyons, D., Hampl, J., Arden, B., Chien, Y.-H., 1998. Ligand recognition by $\alpha\beta$ T cell receptors. *Annu. Rev. Immunol.* 16, 523–544.
- Dembo, A., Zeitouni, O., 1998. Large Deviations Techniques and Applications. Springer, New York.
- Deng, Y., Yewdell, J.W., Eisenlohr, L.C., Bennink, J.R., 1997. MHC affinity, peptide liberation, T cell repertoire, and immunodominance all contribute to the paucity of MHC class I-restricted peptides recognized by antiviral CTL. *J. Immunol.* 158, 1507–1515.
- Detours, V., Mehr, R., Perelson, A.S., 1999. A quantitative theory of affinity-driven T cell repertoire selection. *J. theor. Biol.* 200, 389–403.
- Dustin, M.L., Shaw, A.A., 1999. Costimulation: building an immunological synapse. *Science* 283, 649–650.
- Elliott, J.I., 1993. Thymic selection reinterpreted. *Immunol. Rev.* 135, 227–242.
- Engelhard, V.H., 1994. Structure of peptides associated with class I and class II mhc molecules. *Annu. Rev. Immunol.* 12, 181–207.
- Falk, K., Rötzschke, O., Deres, K., Metzger, J., Jung, G., Rammensee, H.-G., 1991a. Identification of naturally processed viral nonapeptides allows their quantification in infected cells and suggests an allele-specific T cell epitope forecast. *J. Exp. Med.* 174, 425–434.
- Falk, K., Rötzschke, O., Stevanović, S., Jung, G., Rammensee, H.-G., 1991b. Allele-specific motifs revealed by sequencing of self-peptides eluted from MHC molecules. *Nature* 351, 290–296.
- Grakoui, A., Bromley, S.K., Sumen, C., Davis, M.M., Shaw, A.S., Allen, P.M., Dustin, M.L., 1999. The immunological synapse: a molecular machine controlling T cell activation. *Science* 285, 221–227.
- Grossman, Z., Paul, W.E., 2001. Autoreactivity, dynamic tuning and selectivity. *Curr. Opin. Immunol.* 13, 687–698.
- Grossman, Z., Singer, A., 1996. Tuning of activation thresholds explains flexibility in the selection and development of T cells in the thymus. *Proc. Natl. Acad. Sci. USA* 93, 14747–14752.
- van Ham, M., Lith, M. van, Lillemeier, B., Tjin, E., Grüneberg, U., Pastoors, L., Meijgaarden, K. van, Roucard, C., Trowsdale, J., Ottenhoff, T., Pappin, D., Neefjes, J., 2000. Modulation of the major histocompatibility complex class II-associated peptide repertoire by human histocompatibility leukocyte antigen (HLA)-DO. *J. Exp. Med.* 191, 1127–1135.
- Hartgers, F.C., Figdor, C.G., Adema, G.J., 2000. Towards a molecular understanding of dendritic cell immunobiology. *Immunol. Today* 21, 542–545.
- Heath, W.R., Miller, J.F.A.P., 1993. Expression of two alpha chains on the surface of T cells in TCR transgenic mice. *J. Exp. Med.* 178, 1807–1811.
- Holler, P.D., Lim, A.R., Cho, B.K., Rund, L.A., Kranz, D.M., 2001. CD8⁺ T cell transfectants that express a high affinity T cell receptor exhibit enhanced peptide-dependent activation. *J. Exp. Med.* 194, 1043–1052.
- Hudrisier, D., Kessler, B., Valitutti, S., Horvath, C., Cerottini, J.-C., Luescher, I.F., 1998. The efficiency of antigen recognition by CD8⁺ CTL clones is determined by the frequency of serial TCR engagement. *J. Immunol.* 161, 553–562.
- Hunt, D.F., Henderson, R.A., Shabanowitz, J., Sakaguchi, K., Michel, H., Sevilir, N., Cox, A.L., Appella, E., Engelhard, V.H., 1992. Characterization of peptides bound to the class I MHC molecule HLA-A2.1 by mass spectrometry. *Science* 255, 1261–1263.
- Ignatowicz, L., Kappler, J., Marrack, P., 1996. The repertoire of T cells shaped by a single MHC/peptide ligand. *Cell* 84, 521–529.
- Inaba, K., Metlay, J.P., Crowley, M.T., Steinman, R.M., 1990. Dendritic cells pulsed with protein antigens in vitro can prime antigen-specific, MHC-restricted T cells in situ. *J. Exp. Med.* 172, 631–640.
- Itoh, Y., Germain, R.N., 1997. Single cell analysis reveals regulated hierarchical T cell antigen receptor signaling thresholds and intracellular heterogeneity for individual cytokine responses of CD4⁺ cells. *J. Exp. Med.* 186, 757–766.
- Janeway Jr., C.A., Travers, P., 1997. Immunobiology: The Immune System in Health and Disease, 3rd Edition. Garland Publishing, New York.
- Jardetzky, T.S., Lane, W.S., Robinson, R.A., Wiley, D.C., 1991. Identification of self peptides bound to purified HLA-B27. *Nature* 353, 326–329.

- Kalergis, A.M., Boucheron, N., Doucey, M.-A., Palmieri, E., Goyarts, E.C., Vegh, Z., Luesher, I.F., Nathenson, S.G., 2001. Efficient T cell activation requires an optimal dwell-time of interaction between the TCR and the pMHC complex. *Nat. Immunol.* 2, 229–234.
- Kleijmeer, M.J., Ossevoort, M.A., Veen, C.J.H. van, Hellemond, J.J. van, Neefjes, J.J., Kast, W.M., Melief, C.J.M., Geuze, H.J., 1995. MHC class II compartments and the kinetics of antigen presentation in activated mouse spleen dendritic cells. *J. Immunol.* 154, 5714–5724.
- Krummel, M., Wülfing, C., Sumen, C., Davis, M.M., 2000. Thirty-six views of T-cell recognition. *Philos. Trans. R. Soc. Lond. B* 355, 1071–1076.
- Kurts, C., Miller, J.F.A.P., Subramaniam, R.M., Carbone, F.R., Heath, W.R., 1998. Major histocompatibility complex class I-restricted cross-presentation is biased toward high dose antigens and those released during cellular destruction. *J. Exp. Med.* 188, 409–414.
- Kuttler, C., Nussbaum, A.K., Dick, T.P., Rammensee, H.-G., Schild, H., Haderl, K.-P., 2000. An algorithm for the prediction of proteasomal cleavages. *J. Mol. Biol.* 298, 417–429.
- Lanzavecchia, A., Iezzi, G., Viola, A., 1999. From TCR engagement to T cell activation: a kinetic view of T cell behavior. *Cell* 96, 1–4.
- Lyons, D.S., Lieberman, S.A., Hampl, J., Boniface, J., Chien, Y.-H., Berg, L.J., Davis, M.M., 1996. A TCR binds to antagonist ligands with lower affinities and faster dissociation rates than to agonists. *Immunity* 5, 53–61.
- Mason, D., 1998. A very high level of crossreactivity is an essential feature of the T-cell receptor. *Immunol. Today* 19, 395–404.
- Mason, D., 2001. Some quantitative aspects of T cell repertoire selection: the requirement for regulatory T cells. *Immunol. Rev.* 182, 80–88.
- McKeithan, T.W., 1995. Kinetic proofreading in T-cell receptor signal transduction. *Proc. Natl. Acad. Sci. USA* 92, 5042–5046.
- Neefjes, J.J., Momburg, F., Hämmerling, G.J., 1993. Selective and ATP-dependent translocation of peptides by the MHC-encoded transporter. *Science* 261, 769–771.
- Neumeister Kersh, E., Shaw, A.S., Allen, P.M., 1998. Fidelity of T cell activation through multistep T cell receptor ζ phosphorylation. *Science* 281, 572–575.
- Noest, A.J., 2000. Designing lymphocyte functional structure for optimal signal detection: *Viola*, T cells. *J. theor. Biol.* 207, 195–216.
- Nowak, M.A., Tarczy-Hornch, K., Austyn, J.M., 1992. The optimal number of major histocompatibility molecules in an individual. *Proc. Natl. Acad. Sci. USA* 89, 10896–10899.
- Obst, R., Netuschil, N., Klopfer, K., Stevanović, S., Rammensee, H.-G., 2000. The role of peptides in T cell alloreactivity is determined by self-major histocompatibility complex molecules. *J. Exp. Med.* 191, 805–812.
- Parham, P., 2000. *The Immune System*. Garland Publishing, New York.
- Puré, E., Inaba, K., Crowley, M.T., Tardelli, L., Witmer-Pack, M.D., Ruberti, G., Fathman, G., Steinman, R.M., 1990. Antigen processing by epidermal Langerhans cells correlates with the level of biosynthesis of major histocompatibility complex class II molecules and expression of invariant chain. *J. Exp. Med.* 172, 1459–1469.
- Rabinowitz, J.D., Beeson, C., Lyons, D.S., Davis, M.M., McConnell, H.M., 1996. Kinetic discrimination in T-cell activation. *Proc. Natl. Acad. Sci. USA* 93, 1401–1405.
- Rammensee, H.-G., Falk, K., Rötzschke, O., 1993. Peptides naturally presented by MHC class I molecules. *Annu. Rev. Immunol.* 11, 213–244.
- Reay, P.A., Matsui, K., Haase, K., Wülfing, C., Chien, Y.-H., Davis, M., 2000. Determination of the relationship between T cell responsiveness and the number of MHC-peptide complexes using specific monoclonal antibodies. *J. Immunol.* 164, 5626–5634.
- Rothbard, J.B., Gefter, M.L., 1991. Interactions between immunogenic peptides and MHC proteins. *Annu. Rev. Immunol.* 9, 527–565.
- Rudensky, A.Y., Preston-Hurlburt, P., Hong, S.-C., Barlow, A., Janeway Jr., C.A., 1991. Sequence analysis of peptides bound to MHC class II molecules. *Nature* 353, 622–627.
- Savage, P.A., Boniface, J.J., Davis, M.M., 1999. A kinetic basis for T cell receptor repertoire selection during an immune response. *Immunity* 10, 458–492.
- Shores, E.W., Tran, T., Grinberg, A., Sommers, C.L., Shen, H., Love, P.E., 1997. Role of the multiple T cell receptor (TCR)- ζ chain signaling motifs in selection of the T cell repertoire. *J. Exp. Med.* 185, 893–900.
- Siemasko, K., Clark, M.R., 2001. The control and facilitation of MHC class II processing by the BCR. *Curr. Opin. Immunol.* 13, 32–36.
- Steinman, R.M., 1991. The dendritic cell system and its role in immunogenicity. *Annu. Rev. Immunol.* 9, 271–296.
- Stevanović, S., Schild, H., 1999. Quantitative aspects of T cell activation—peptide generation and editing by MHC class I molecule. *Semin. Immunol.* 11, 375–384.
- Thomas, R., Lipsky, P.E., 1996. Presentation of self peptides by dendritic cells. *Arthritis Rheumatism* 39, 183–190.
- Valitutti, S., Lanzavecchia, A., 1997. Serial triggering of TCRs: a basis for the sensitivity and specificity of antigen recognition. *Immunol. Today* 18, 299–304.
- Valitutti, S., Müller, S., Cella, M., Padovan, E., Lanzavecchia, A., 1995. Serial triggering of many T-cell receptors by a few peptide-MHC complexes. *Nature* 375, 148–151.
- Valitutti, S., Müller, S., Dessing, M., Lanzavecchia, A., 1996. Different responses are elicited in cytotoxic T lymphocytes by different levels of T cell receptor occupancy. *J. Exp. Med.* 183, 1917–1921.
- Viola, A., Lanzavecchia, A., 1996. T cell activation determined by T cell receptor number and tunable thresholds. *Science* 273, 104–106.
- Watts, C., 2001. Antigen processing in the endocytic compartment. *Curr. Opin. Immunol.* 13, 26–31.
- Yewdell, J.W., Bennink, J.R., 2001. Cut and trim: generating MHC class I peptide ligands. *Curr. Opin. Immunol.* 13, 13–18.
- Zeh III, H.J., Perry-Lalley, D., Dudley, M.E., Rosenberg, S.A., Yang, J.C., 1999. High avidity CTLs for two self-antigens demonstrate superior in vitro and in vivo antitumor efficacy. *J. Immunol.* 162, 989–994.
- Zerrahn, J., Held, W., Raulet, D.H., 1997. The MHC reactivity of the T cell repertoire prior to positive and negative selection. *Cell* 88, 627–636.