Quantitative theories of T cell responsiveness

Hugo A. van den Berg and David A. Rand

WARWICK SYSTEMS BIOLOGY CENTRE
Coventry House, University of Warwick, Coventry CV4 7AL, United Kingdom
hugo@maths.warwick.ac.uk

HAB is the corresponding author

Acknowledgments The authors acknowledge the useful discussions with many of the people whose contributions to the field are discussed in this review. HAB was supported by the Wellcome Trust and the EPSRC Life Sciences Interface. DAR was supported by the EPSRC, the EU (BioSim Network Contract No. 005137) and the BBSRC

To appear in *Immunological Reviews*

Summary: We review recent advances toward a comprehensive mathematical theory of T cell immunity. A key insight is that the efficacy of the T cell response is best analysed in terms of TCR avidity and the distribution of this avidity across the TCR repertoire (the "avidity spectrum"). Modification of this avidity spectrum by a wide range of tuning and tolerance mechanisms allows the system to adapt cross-reactivity and specificity to the challenge at hand whilst avoiding inappropriate responses against non-pathogenic cells and tissues. Theoretical models relate molecular kinetics and cellular properties to systemic-level statistics such as avidity spectra. Building such bridges is crucial for rational clinical manipulation of T cells at the molecular level.

Running title: T cell responsiveness: Quantitative theories

Keywords: T cell activation; Functional diversity; T cell specificity

Introduction

Due to the diversity and rapid evolution of pathogenic organisms, the T cell immune system cannot anticipate precisely which epitopes will be encountered during future pathogenic challenges. Instead the T cell system relies on a vast *repertoire* of different T cell antigen receptor (TCR) clonotypes which are generated at random in advance, prior to infection. The TCR repertoire comprises millions of clonotypes, a diversity which virtually ensures that at least one T cell clone will attack a given pathogen (1–4). At the same time, the T cells must guard against inappropriate immune responses directed against non-pathogenic proteins (*horror autotoxicus*).

Here we review theoretical work that has helped us understand how the T cell repertoire succeeds in responding to pathogens while remaining non-responsive to the host's own antigens. The recurring theme of this review is that T cell recognition, rather than being a discrete phenomenon, has to be defined on a *continuous* scale. The most important consequence is that immune tolerance also becomes a matter of degree. Instead of absolute responsiveness to foreign or 'nonself' and an absolute non-responsiveness to 'self', it will become apparent that immune tolerance finds a natural place on this continuous scale.

T cell responsiveness: degenerate and graded, or specific and discrete?

One might expect that the requirements of immune fidelity (i.e. reliability, efficacy, and appropriate targetting) are fulfilled by ensuring *stringent specificity* of TCR antigen recognition, such that each TCR clonotype derives strong stimulation from only one epitope. Recognition would then be a dichotomous, all-or-none phenomenon: appropriate responses could then be ensured simply by removal from the repertoire of TCR clonotypes that recognize non-pathogenic epitopes. However, while clonal deletion of autoreactive cells is a fundamental process in shaping the TCR repertoire (5–7), controlled autoreactivity remains an essential part of the physiology of the immune system (8), and a plethora of mechanisms is involved in the regulation of T cell activation (9–13). The reason that controlled autoreactivity remains an essential property of T cell immunity is that stringent specificity is unattainable: it would require an unfeasibly large number of distinct clonotypes (14). Instead, TCR antigen recognition is *degenerate*, with each TCR clonotype recognizing, typically with varying strengths, a wide variety of epitopes (15–22).

The efficacy of TCR epitope recognition is often described in terms of *avidity*: a clonotype is said to have high avidity for a given epitope if it is able to respond even when the epitope is present at very low levels, whereas a low-avidity clone requires higher presentation levels (23–25). We here review mathematical models of the molecular basis of TCR avidity as well as its role in immunity at the systemic level. Thus we gain an insight into the central problem of TCR recognition, which is how immune fidelity arises out of TCR degeneracy and controlled autoreactivity.

The T cell receptor recognizes peptide/MHC ligands

TCR recognition of antigen ligands takes place in the contact area between a T cell and an antigen-presenting cell (APC) (26). TCR molecules bind to ligands on the surface of the APC, as shown in *Fig. 1*. Each ligand molecule consists of a peptide (a protein

fragment) attached to a specialized antigen-presenting receptor encoded by the major histocompatibility complex (MHC) of genes. The peptides are generated by cutting up proteins into small fragments (27). As the size of these peptide fragments increases, the probability that a nonself (pathogen-derived) peptide also occurs in the human self decreases (28).

Upon binding to a peptide/MHC (pMHC) molecule, transmembrane molecules associated with the TCR acquire the ability to transmit signals to the cellular interior by phosphorylation of intracellular signalling domains associated with the TCR/CD3 complex (29). This *TCR triggering* leads to various cellular responses, such as changes in gene expression or the killing of a target cell (30).

The pMHC ligands found on a single APC form a mixed population of thousands of different species (27, 31). Most of these are 'harmless', that is, derived from host proteins not associated with disease, or from non-pathogenic organisms. Probabilistic arguments suggest that TCR repertoire diversity is not constrained by the number of ligands against which the system is protecting, but by the within-host diversity of harmless ligands (32); the presentation statistics of such ligands is discussed in Box I.

T cell efficacy rests on the ability to respond to 'harmful', that is, disease-associated, antigens while remaining non-responsive to the harmless background against which these 'harmful' ligands are presented. A key problem in immunology is to understand precisely how the immune system achieves such discrimination. This problem is known as the *self-nonself* problem, a misleading misnomer since self-derived ligands may be associated with disease, while the body harbours many nonself proteomes which are perfectly harmless and are normally ignored by the immune system (33).

The TCR triggering rate as a measure of avidity

The signalling strength of a given pMHC species is related to the rate at which it triggers TCR molecules; that is, the rate at which TCRs are induced to acquire the ability to transmit intracellular signals during an encounter between a T cell and an APC. A given pMHC ligand triggers TCR molecules at different rates in different clonotypes. Similarly, a TCR of a given clonotype is triggered at different rates by different ligands. The TCR triggering rate per pMHC molecule thus depends on both the clonotype (labelled i) and the ligand species (labelled j). The dependence of the TCR triggering rate on the mean lifetime of a TCR/pMHC interaction is shown in *Fig. 2*. The top curve shows that the TCR triggering rate is maximal at an optimal mean TCR/pMHC residence time; this prediction has been experimentally confirmed (34).

We propose that w_{ij}, the triggering rate elicited in clonotype i by a molecule of pMHC species j, quantitates the avidity of the TCR/pMHC interaction (Box II and ref. 35). This MHC-specific TCR triggering rate w_{ij} can be measured in a standard bioassay which determines the percentage of T cells of a defined clonotype responding to a range of APCs that have been incubated at various antigen concentrations. The mean residence time T_{ij} can also be experimentally determined; combining the two data, one can deduce the receptor triggering threshold (T_R; see Box II).

Regulated T cell activation: the responder avidity spectrum

An immune response is initiated in the lymphoid organs (lymph nodes, spleen) with the presentation of disease-associated pMHC ligands to *naive T cells* which constitute the working repertoire of the cellular immune system. Naive T cells interact with professional APCs presenting the relevant ligand from the diseased tissue to the lymphoid organ. Lanzavecchia and colleagues (36) proposed that a naive T cell is activated when a sufficient number of its TCR molecules is triggered, that is, when the rate at which the T cells' TCR molecules are being triggered is above a certain *activation threshold* which is a cellular property. The activated T cell then commits itself to proliferation and differentiation, leading to the *expansion* of sufficient numbers of T cells to combat the disease. A combination of models and *in vivo* data has shown that competition for 'antigenic sites' on APCs is a limiting factor in T cell expansion (37, 38).

Only a minute fraction of the naive repertoire is activated (14, 39, 40): efficient responders exhibit a much higher avidity for the relevant ligand than the general repertoire. This is illustrated in *Fig. 3*, which compares the *TCR avidity spectrum* of the naive repertoire to the TCR avidity spectrum among responding clones: responder spectra are shown for various values of the naive T cell activation threshold. The TCR avidity spectrum is a statistical distribution, defined in detail in Box III and *Fig. 3*.

Taken together, the responder spectra in *Fig. 3* demonstrate that the naive T cell activation threshold has to be sufficiently high to ensure that there is a significant enrichment of avidity among the responding T cells. In general, there will be 'false positives', low-avidity responders which occur due to random fluctuations in the population of harmless ligands on the APC, which will occasionally lead to activation even if the T cell has low avidity for the relevant pathogenic ligand. The condition for an effective immune response is that the frequency at which these false positives occur must be much lower than the frequency at which 'true positives' occur. The latter frequency is just the within-repertoire proportion of T cells with high avidity for the relevant ligand. This condition is satisfied if activation thresholds are sufficiently high, as the responder avidity spectra in *Fig. 3* demonstrate.

One way to evoke an immune response against a selected antigen is to incubate professional APCs with the antigenic peptide, which is thereby loaded onto the APCs' MHC molecules. *Fig. 4* shows predicted avidity spectra among responders for various incubation concentrations. These spectra indicate progressively less probability mass is concentrated at high avidities as the antigen incubation concentration increases. Thus we have the perhaps counter-intuitive prediction that *less* efficient responses are to be expected when the antigen is present at higher levels on the APC. This prediction agrees with experimental findings (24, 41). The need for high-avidity responses therefore furnishes an upper bound to the antigen incubation concentration.

Selected T cells among the responders differentiate into long-lived *memory T cells*, which ensures a vigorous response to subsequent infections with the same (or a closely related) pathogen (42). The dynamics of the memory formation affects the long-term evolution of the diversity of the memory T cell pool, which is crucial in maintaining significant immunity to any given challenge over the long term, (43–46). An important factor in memory diversity is cross-reactivity, which is a function of specificity of the naive repertoire T cells (47).

From molecular properties to repertoire statistics

A central theoretical objective is to relate the avidity spectrum to underlying cellular and molecular parameters, since this would help us link therapeutic intervention at the molecular level to the intended curative effects at the systemic level.

Combining the triggering rate model (Box II) with the statistical distribution of the mean TCR/pMHC residence time (Box III), one can derive the statistical distribution of the MHC-specific TCR triggering rate for a given foreign pMHC ligand. A foreign ligand is a ligand to which the repertoire T cells are yet to be exposed (a self peptide derived from the host proteome can be 'foreign' by this definition). Graphs of this TCR triggering rate distribution are shown in *Fig. 5* for various parameter values. These curves plot the probability that the given foreign ligand elicits, in a randomly selected naive T cell, a triggering rate per pMHC molecule that is at least as great as the abscissa.

Among the possibilities shown in *Fig. 5*, there is a curve with a broad central plateau, indicating that probability mass is concentrated at the extremes of TCR triggering rate zero and the maximum possible triggering rate. This particular case corresponds to the dichotomous (all-or-none) character alluded to above. Interestingly, a virtually discrete triggering rate distribution can be realized despite the intrinsically continuous nature of the underlying distribution of the off-rate (Box III). The vertical location of the plateau (10^{-5} in the figure) indicates the probability that a particular clonotype recognizes a randomly chosen ligand—this probability is a measure of *TCR degeneracy* (14), also termed the *crossreactivity parameter* (48). For other parameter combinations the plateau exhibits a pronounced slope, which endows the immune system with the opportunity to fine-tune the degeneracy of recognition. For instance, the number of clones responding to a given pathogenic challenge can be adjusted by a global (system-level) signal which induces all naive T cells in the repertoire to alter their activation threshold. Such global signals can be transmitted through cytokines, the hormones of the immune system.

We conclude that variation of molecular parameters allows the immune system to adjust the extent to which recognition is an all-or-none phenomenon, and, moreover, to tune the level of degeneracy. Thus, the system can dynamically adapt to circumstances which might, at one extreme, require highly targetted, high-avidity effectors only, or, at the other extreme, a broad-spectrum polyclonal response consisting of fairly cross-reactive effectors.

Immunological tolerance as a multi-component system

Tolerance, defined as the lack of responsiveness to non-disease-associated antigens, is an essential feature of the immune system. Such antigens include peptides derived from the host-proteome, often referred to as *self*. A detailed understanding of tolerance at the molecular and cellular levels is a crucial prerequisite for clinical manipulation aimed either at reinforcing tolerance to combat autoimmune disease which develops when tolerance is lost, or to suppress tolerance for selected ligands such as tumour-associated antigens.

Tolerance is traditionally divided into *central tolerance*, acquired by immature lymphocytes in the central lymphoid organs (notably the thymus), and *peripheral tolerance*, acquired by mature lymphocytes in the peripheral tissues. From a functional point of view one may also classify tolerance as *deletional* or *adaptive*, depending on whether

self-reactive T cells are removed from the repertoire or whether they tune their responsiveness so that self reactivity is lost. Extreme non-responsiveness (anergy) effectively amounts to removal from the repertoire. Moreover, such anergic cells may differentiate into *regulatory T cells* whose high reactivity to non-disease-associated antigens plays a positive role in safeguarding tolerance (9, 49). Although regulatory T cells are without doubt an essential component of immunological tolerance (50), theorists were quick to rule out the necessity of such a component when the concept of regulatory T cells fell out of favour with immunologists (46). More recently, a model of T cell-mediated tolerance correctly predicted the inverse correlation between infection and the incidence of autoimmune disease (51) .

Deletional tolerance can also take the form of programmed cell death (apoptosis) predicated on recognition of non-disease-associated antigens. Death of strongly self-recognizing T cells takes place in the thymus and is known as *negative selection* (30). Yet virtually every selected clonotype will retain some residual avidity for several self antigens; this residual effect is succinctly encoded in the *tolerized avidity spectrum.*

Tolerized avidity spectra: tolerance and ignorance are shades of grey

Tolerance to a given antigen can be characterized by investigating the avidity spectrum over the T cell repertoire for that antigen. A range of such spectra are shown in *Fig. 6*. The top curve is the spectrum for an antigen against which absolutely no tolerance has been acquired; such an antigen is foreign by definition. Antigens derived from the proteome of a pathogen are almost invariably foreign in this sense. The remaining curves demonstrate various degrees of tolerance: a truncation at high TCR triggering rates is evident. For any given ligand, its truncation point depends on its presentation level in the thymus. The frequency with which the thymocyte encounters the ligand during central tolerization determines the degree to which the repertoire is depleted as regards avidities (this frequency is zero for foreign ligand, and no avidity-depletion is achieved for such ligands). These theoretical results (52) predict that negative selection will tend to allow T cells with high avidity for subdominant or cryptic epitopes to persist while effectively deleting high-avidity T cells specific for dominant epitopes. Experimental findings are in keeping with this prediction (53). The repertoire is sometimes said to be *ignorant* about such cryptic epitopes. The relationships between the thymic presentation level of an antigen, its presentation level during the induction of a response, and experimental findings on avidity spectrum of the response (25) agree fully with the theoretical predictions.

An obvious clinical importance attaches to host proteome-derived antigens that remain immunogenic even after tolerization, as a result of low ‘immunovisibility’. When such a self antigen is associated with proteins involved in normal physiological functions, it poses a risk for autoimmune disease (54). As soon as the antigen becomes immunologically more ‘visible’, clonotypes reactive to the antigen may be expanded, leading to damage of the healthy tissue in which it is expressed. On the other hand, when an essentially foreign antigen derives from proteins that are expressed at high levels only in diseased cells (such as neoplastic cells) it constitutes a likely target for immune surveillance of tumours, and a good candidate for so-called *T cell vaccination* (55).

An interesting corollary is that information about thymic selection is contained in the avidity spectra of the working repertoire. Comparison of an antigen’s avidity spectrum

to the foreign spectrum allows one to determine (i) the frequency with which T cells are exposed to that antigen during central (deletional) tolerization; as well as (ii) the typical levels at which the antigen is presented during that process. Thus, key parameters of deletional tolerance can be experimentally determined from tolerized avidity spectra across the naive repertoire.

Tuning T cell responsiveness

To obtain an insight into adaptive tolerance, we plot the probability that a given naive T cell is activated as a function of the cellular activation threshold, obtaining an *activation curve*. *Fig. 7* shows two such activation curves for a given T cell: one where the professional APC is presenting only harmless self antigens, and one where pathogen-derived antigens are presented as well. Marking off a maximum allowed probability of autoactivation and a minimum allowed probability of activation by the relevant ligand, we find, respectively, minimum and maximum allowed cellular activation threshold values, indicated by $W_{\rm act}^{\rm min}$ and $W_{\rm act}^{\rm max}$. The business of adaptive tolerance mechanisms is to set the activation threshold to a value between these two extremes (see Box IV for more details). Recalling the discussion surrounding the responder avidity spectra shown in *Fig. 3*, we infer that the optimal cellular activation threshold lies at the upper extreme.

The difference $\Delta W = W_{\rm act}^{\rm max} - W_{\rm act}^{\rm min}$ is the separation of the activation curves. With a greater separation ΔW it is less likely that adaptive tolerance fails to set the cellular activation threshold to a value in between $W_{\rm act}^{\rm min}$ and $W_{\rm act}^{\rm max}$.

Three factors determine the separation of the activation curves: (i) the presentation levels of the foreign ligands; (ii) the presentation diversity of the pathogen-derived ligands; and (iii) the variability of the background formed by harmless self antigens. We discuss these three factors in turn.

First, the activation curve separation ΔW increases with the foreign presentation levels (56). Consequently, a minimum foreign presentation level is defined by the condition $\Delta W = 0$. As we saw, the responder avidity spectra (*Fig. 4*) bound the foreign presentation level from above; the activation curve separation principle bounds it from below. Combining these two bounds, we predict that there is a window of effective antigen presentation levels on APCs.

Second, activation curve separation decreases with the presentation diversity of the foreign antigens (57). This diversity is defined as the Simpson's diversity index on the presentation levels. The separation ΔW is maximized by minimizing the foreign diversity. This provides a deep theoretical explanation for the stringent selectivity of peptide presentation on MHC molecules (Box V and ref. 57). We can also explain why there are several distinct MHC isoforms, with each T cell restricted to interaction with only one of these isoforms (Box V).

Finally, activation curve separation decreases with the statistical variability of the harmless background (these fluctuations are discussed in Box I). This is intuitive: the noisier the background, the larger the signal has to be unequivocally detected (cf. ref. 58).

An intriguing result is that deletional tolerance can specifically target T cells with high self background APC-to-APC variability (59). Such targetting requires a special relationship between self antigen presentation statistics in the thymus and the periphery. More generally and abstractly, the relationship between thymic and peripheral presen-

tation statistics can be said to *encode the target* of negative selection (this idea is made more precise in Box VI).

Discussion

The theoretical developments reviewed here elucidate and tie together numerous phenomena in the adaptive cellular immune system: epitope presentation levels, relative immunogenicity, the avidity of expanded T cell clones, MHC isoform diversity within the individual and TCR restriction to MHC isoforms, deletional and adaptive tolerance, tuning of the T cell activation threshold, and the functional significance of altered patterns of self antigen presentation in the thymus. As we have pointed out, the connections expressed by theory are amenable to empirical verification; the key parameters of the theory can be determined experimentally.

Further layers of complexity

We attempted to cover many aspects of T cell responsiveness while keeping the theory as simple as possible, and many important aspects have inevitably been left out. For instance, we have ignored the crucial role of intercellular signalling in the orchestrated evolution of the immune response, which has been extensively modelled (60–63). Most notably, we have neglected the evolution of dedicated receptor interaction zones within the T cell:APC contact area, in which the kinase/phosphatase balance is drastically altered (64–67). We think of the TCR/CD3 complex as a bistable system, which is normally stably locked in the untriggered 'ground' state by negative regulatory interactions (68), and, provided the pMHC/TCR interaction lasts longer than the threshold duration T_R, becomes locked into the signalling state which is similarly stabilized by positive interactions (29). We have assumed that the TCR quickly reverts to the untriggered state if the triggered state is not attained; this assumption may be false in phosphatase-depleted zones of the contact area, where the TCR would revert slowly, allowing subsequent interactions with distinct ligand molecules to contribute to the triggering of a given TCR/CD3 complex (69). To model this situation, intermediate states of TCR/CD3 phosphorylation must be taken into account (refs. 70, 71 and Burroughs and Van der Merwe, this volume). Detailed computational models of immune-receptor signalling machinery are reviewed in ref. 72. A key enzyme in the transition of the TCR from the untriggered to the triggered state is Lck, a kinase that phosphorylates specific intracellular components of the TCR/CD3 complex, and whose activation state exhibits bistability (68). TCR bistability (untriggered/triggered) interacts with Lck bistability (inactive/active), leading to nonlinear dynamics (73, 74). Another aspect is the role of the co-receptor, which stabilizes the TCR/pMHC interaction and thus modulates T_{ij} (75) and may also profoundly affect the receptor threshold T_R, for instance, by modulating its recruitment to kinase-rich membrane microdomains (76–79).

Conclusion

Immunologists are traditionally perhaps most comfortable with the concept of affinity, but TCR recognition raises many perplexities when addressed in terms of affinity: it may seem paradoxical that TCR recognition achieves such exquisite ligand discrimination

when even potent interactions have very low biochemical affinity, and one may wonder how affinity-based selection of a TCR repertoire is possible at all. As the theoretical developments reviewed here show, such questions really are red herrings. This becomes clear once one realizes the key parameter to focus on is *avidity*, represented here in terms of the TCR triggering rate. This quantity depends primarily on the mean TCR/pMHC dwell time, while it is modulated by the dissociation constant (i.e. affinity) and combines with antigen presentation levels to produce additive multiple-site effects (Box II). Moreover, mathematical immunology now affords an integrated understanding of immune tolerance, response efficiency, and MHC restriction by considering statistical distributions of avidity, both across the naive repertoire and the set of responders to a given pathogenic challenge.

References

1. Goldrath AW, Bevan MJ. Selecting and maintaining a diverse T-cell repertoire. *Nature* : 1999; 402:255 – 262.

2. Perelson AS, Weisbuch G. Immunology for physicists. *Reviews of Modern Physics* : 1997; 69:1219 – 1267.

3. Arstila TP, Casrouge A, Baron V, Even J, Kanellopoulos J, Kourilsky P. A direct estimate of the human α–β T cell receptor diversity. *Science* : 1999; 286:958 – 961.

4. Keşmir C, Borghans JAM, Boer RJde Diversity of human $\alpha\beta$ T cell receptors. *Science* : 2000; 288:1135a – 1135a.

5. Kappler JW, Roehm N, Marrack P. T cell tolerance by clonal elimination in the thymus. *Cell* : 1987; 49:273 – 280.

6. Kisielow P, Blüthmann H, Staerz UD, Steinmetz M, von Boehmer H. Tolerance in T-cell-receptor transgenic mice involves deletion of nonmature $CD4^+8^+$ thymocytes. *Nature* : 1988; 333:742 – 746.

7. Laufer TM, Fan L, Glimcher LH. Self-reactive T cells selected on thymic cortical epithelium are polyclonal and are pathogenic in vivo. *J. Immunol.* : 1999; 162:5078 – 5084.

8. Annacker O, Pimenta-Araujo R, Burlen-Defranoux O, Bandeira A. On the ontology and physiology of regulatory T cells. *Immunological Reviews* : 2001; 182:5 – 17.

9. Roncarolo MG, Levings MK. The role of different subsets of T regulatory cells in controlling autoimmunity. *Current Opinion in Immunology* : 2000; 12:676 – 683.

10. Zheng Y, Manzotti CN, Liu M, Burke F, Mead KI, Sansom DM. CD86 and CD80 differentially modulate the suppressive function of human regulatory T cells. *J. Immunol.* : 2004; 172:2778 – 2784.

11. Salomon B, Bluestone JA. Complexities of CD28/B7: CTLA-4 costimulatory pathways in autoimmunity and transplantation. *Annu. Rev. Immunol.* : 2001; 19:225 – 252.

12. Thompson CB, Allison JP. The emerging role of CTLA-4 as an immune attenuator. *Immunity* : 1997; 7:445 – 450.

13. Davis MM, Krogsgaard M, Huppa JB, Sumen C, Purbhoo MA, Irvine DJ, Wu LC, Ehrlich L. Dynamics of cell surface molecules during T cell recognition. *Annu. Rev. Biochem.* : 2003; 72:717 – 742.

14. Mason D. A very high level of crossreactivity is an essential feature of the T-cell receptor. *Immunol. Today* : 1998; 19:395 – 404.

15. Anderson AC, Waldner H, Turchin V, Jabs C, Prabhu Das M, Kuchroo VK, Nicholson LB. Autoantigen-responsive T cell clones demonstrate unfocused TCR cross-reactivity toward multiple related ligands: Implications for autoimmunity. *Cell. Immunol.* : 2000; 202:88 – 96.

16. Rudolph MG, Wilson IA. The specifi city of TCR/pMHC interaction. *Current Opinion in Immunology* : 2002; 14:52 – 65.

17. Garcia KC, Degano M, Pease LR, Huang M, Peterson PA, Teyton L, Wilson IA. Structural basis of plasticity in T cell receptor recognition of a self peptide-MHC antigen. *Science* : 1998; 279:1166 – 1172.

18. Hemmer B, Fleckenstein BT, Vergelli M, Jung G, McFarland H, Martin R, Wiesmüller KH. Identifi cation of high potency microbial and self ligands for a human autoreactive class II-restricted T cell clone. *J. Exp. Med.* : 1997; 185:1651 – 1659.

19. Hemmer B, Vergelli M, Pinilla C, Houghten R, Martin R. Probing degeneracy in T-cell recognition using peptide combinatorial libraries. *Immunol. Today* : 1998; 19:163 – 168.

20. Ignatowicz L, Kappler J, Marrack P. The repertoire of T cells shaped by a single MHC/peptide ligand. *Cell* : 1996; 84:521 – 529.

21. Kersh GJ, Allen PM. Structural basis for T cell recognition of altered peptide ligands: A single T cell receptor can productively recognize a large continuum of related ligands. *J. Exp. Med.* : 1996; 184:1259 – 1268.

22. Mazza G, Housset D, Piras C, Gregoire C, Lin SY, Fontecilla-Camps JC, Malissen B. Glimpses at the recognition of peptide/MHC complexes by T-cell antigen receptors. *Immunological Reviews* : 1998; 163:187 – 196.

23. Ashton-Rickardt PG, Tonegawa S. A differential-avidity model for T-cell selection. *Immunol. Today* : 1994; 15:362 – 366.

24. Alexander-Miller MA. Differential expansion and survival of high and low avidity cytotoxic T cell populations during the immune response to a viral infection. *Cellular Immunology* : 2000; 201:58 – 62.

25. Gross DA, Graff-Dubois S, Opolon P, Cornet S, Alves P, Bennaceur-Griscelli A, Faure O, Guillaume P, Firat H, Chouaib S, Lemonnier FA, Davoust J, Miconnet I, Vonderheide RH, Kosmatopoulos K. High vaccination effi ciency of low-affi nity epitopes in antitumor immunotherapy. *J. Clin. Invest.* : 2004; 113:425 – 433.

26. Davis SJ, Ikemizu S, Evans EJ, Fugger L, Bakker TR, Merwe PAvan der. The nature of molecular recognition by T cells. *Nature Immunol.* : 2003; 4:1 – 8.

27. Stevanović S, Schild H. Quantitative aspects of T cell activation—peptide generation and editing by MHC class I molecule. *Seminars in Immunology* : 1999; 11:375 – 384.

28. Burroughs N, Kesmir C, de Boer R. Discriminating self from nonself with short peptides from large proteomes. *Immunogenetics* : 2004; 56:311 – 320.

29. Werlen G, Palmer E. The TCR signalosome: A dynamic structure with expanding complexity. *Current Opinion in Immunology* : 2002; 14:299 – 305.

30. Parham P. *The Immune System*. Garland Publishing, New York: 2000.

31. Rammensee HG, Falk K, Rötzschke O. Peptides naturally presented by MHC class I molecules. *Annu. Rev. Immunol.* : 1993; 11:213 – 244.

32. Boer RJde, Perelson AS. How diverse should the immune system be? *Proc. R. Soc. Lond. B* : 1993; 252:171 – 175.

33. Matzinger P. The Danger model: A renewed sense of self. *Science* : 2002; 296:301 – 305.

34. Kalergis AM, Boucheron N, Doucey MA, Palmieri E, Goyarts EC, Vegh Z, Luesher IF, Nathenson SG. Efficient T cell activation requires an optimal dwell-time of interaction between the TCR and the pMHC complex. *Nature Immunol.* : 2001; 2:229 – 234.

35. Berg HAvan den, Burroughs NJ, Rand DA. Quantifying the strength of ligand antagonism in TCR triggering. *Bull. Mat. Biol.* : 2002; 64:781 – 808.

36. Lanzavecchia A, Sallusto F. Dynamics of T lymphocyte responses: Intermediates, effectors, and memory cells. *Science* : 2000; 290:92 – 97.

37. Pilyugin S, Mittler J, Antia R. Modeling T-cell proliferation: An investigation of the consequences of the Hayflick limit. *J. Theor. Biol.* : 1997; 186:117–129.

38. Borghans JAM, Taams LS, Wauben MHM, Boer Rde. Competition for antigenic sites during T cell proliferation: A mathematical interpretation of *In Vitro* data. *Proc. Natl. Acad. Sci. USA* : 1999; 96:10782 – 10787.

39. Butz EA, Bevan MJ. Massive expansion of antigen-specific $CD8^+$ T cells during an acute virus infection. *Immunity* : 1998; 8:167 – 175.

40. Gavin MA, Bevan MJ. Increased peptide promiscuity provides a rationale for the lack of N regions in the neonatal T cell repertoire. *Immunity* : 1995; 3:793 – 800.

41. Alexander-Miller MA, Leggatt GR, Berzofsky JA. Selective expansion of high- or low-avidity cytotoxic T lymphocytes and efficacy for adoptive immunotherapy. *Proc. Natl. Acad. Sci. USA* : 1996; 93:4102 – 4107.

42. Dutton RW, Bradley LM, Swain SL. T cell memory. *Annu. Rev. Immunol.* : 1998; 16:201 – 223.

43. Antia R, Pilyugin SS, Ahmed R. Models of immune memory: On the role of cross-reactive stimulation, competition, and homeostasis in maintaining immune memory. *Proc. Natl. Acad. Sci. USA* : 1998; 95:14926 – 14931.

44. Utzny C, Burroughs NJ. Stability of a diverse immunological memory is determined by T cell population dynamics. *Bull. Math. Biol.* : 2001; 63:685 – 713.

45. Yates AJ, Callard RE. Cell death and the maintenance of immunological memory. *Discr. Cont. Dynam. Sys. B* : 2001; 1:43 – 60.

46. Boer RJde, Hogeweg P. Immunological discrimination between self and non-self by precursor depletion and memory accumulation. *J. Theor. Biol.* : 1987; 124:343 – 369.

47. Borghans JAM, Noest A, Boer RJde. How specific should immunological memory be? *J. Immunol.* : 1999; 163:569 – 575.

48. Borghans J, Boer RJde. Crossreactivity of the T-cell receptor. *Immunol. Today* : 1998; 19:428 – 429.

49. Jordan MS, Boesteanu A, Reed AJ, Petrone AL, Holenbeck AE, Lerman MA, Naji A, Caton AJ. Thymic selection of $CD4^+CD25^+$ regulatory T cells induced by an agonist self-peptide. *Nature Immunol.* : 2001; 2:301 – 306.

50. Mason D. Some quantitative aspects of T cell repertoire selection: The requirement for regulatory T cells. *Immunological Reviews* : 2001; 182:80 – 88.

51. León K, Faro J, Lage A, Carneiro J. Inverse correlation between the incidences of autoimmune disease and infection predicted by a model of T cell mediated tolerance. *J. Autoimmunol.* : 2004; 22:31 – 42.

52. Berg HAvan den, Rand DA. Foreigness as a matter of degree: The relative immunogenicity of peptide/MHC ligands. *J. Theor. Biol.* : 2004; 231:535 – 548.

53. Slifka MK, Blattman JN, Sourdive DJD, Liu F, Huffman DL, Wolfe T, Hughes A, Oldstone MBA, Ahmed R, von Herrath MG. Preferential escape of subdominant $CD8^+$ T cells during negative selection results in an altered antiviral T cell hierarchy. *J. Immunol.* : 2003; 170:1231 – 1239.

54. Thomas R, Lipsky PE. Could endogenous self-peptides presented by dendritic cells initiate rheumatoid arthritis? *Immunol. Today* : 1996; 17:559 – 564.

55. Bona CA, Casares S, Brumeanu TD. Towards development of T-cell vaccines. *Immunol. Today* : 1998; 19:126 – 132.

56. Berg HAvan den, Rand DA, Burroughs NJ. A reliable and safe T cell repertoire based on low-affinity T cell receptors. *J. Theor. Biol.* : 2001; 209:465 – 486.

57. Berg HAvan den, Rand DA. Antigen presentation on MHC molecules as a diversity filter that enhances immune efficacy. *J. Theor. Biol.* : 2003; 224:249 – 267.

58. Noest AJ. Designing lymphocyte functional structure for optimal signal detection: *Voilà*, T cells. *J. Theor. Biol.* : 2000; 207:195 – 216.

59. Berg HAvan den, Molina-París C. Thymic presentation of autoantigens and the efficiency of negative selection. *J. Theor. Med.* : 2003; 5:1 – 22.

60. Segel LA, Lev Bar-Or R. On the role of feedback in promoting conflicting goals of the adaptive immune system. *J. Immunol.* : 1999; 163:1342 – 1349.

61. Callard RE, George AJT, Stark J. Cytokines, chaos, and complexity. *Immunity* : 1999; 11:507 – 513.

62. Yates AJ, Bergmann Claudia van Hemmen L, Stark J, Callard RE. Cytokine-modulated regulation of helper T cell populations. *J. Theor. Biol.* : 2000; 205:539 – 560.

63. Utzny C, Burroughs NJ. Perturbation theory analysis of competition in a heterogenous population. *Physica D* : 2003; 175:109 – 126.

64. Monks CRF, Freiberg BA, Kupfer H, Sciaky N, Kupfer A. Three-dimensional segregation of supramolecular activation clusters in T cells. *Nature* : 1998; 395:82 – 86.

65. Dustin ML, Shaw AA. Costimulation: Building an immunological synapse. *Science* : 1999; 283:649 – 650.

66. Grakoui A, Bromley SK, Sumen C, Davis MM, Shaw AS, Allen PM, Dustin ML. The immunological synapse: A molecular machine controlling T cell activation. *Science* : 1999; 285:221 – 227.

67. Freiberg BA, Kupfer H, Maslanik W, Delli J, Kappler J, Zaller DM, Kupfer A. Staging and resetting T cell activation in SMACs. *Nature Immunol.* : 2002; 3:911 – 917.

68. Veillette A, Latour S, Davidson D. Negative regulation of immunoreceptor signalling. *Annu. Rev. Immunol.* : 2002; 20:669 – 707.

69. Wülfing C, Sumen C, Sjaastad MD, Wu LC, Dustin ML, Davis MM. Costimulation and endogenous MHC ligands contribute to T cell recognition. *Nature Immunol.* : 2002; 3:42 – 47.

70. Chan C, George AJT, Stark J. Cooperative enhancement of specificity in a lattice of T cell receptors. *Proc. Natl. Acad. Sci. USA* : 2001; 98:5758 – 5763.

71. Hlavacek WS, Redondo A, Wofsy C, Goldstein B. Kinetic proofreading in receptor-mediated transduction of cellular signals: Receptor aggregation, partially activated receptors, and cytosolic messengers. *Bull. Math. Biol.* : 2002; 64:887 – 911.

72. Goldstein B, Faeder JR, Hlavacek W. Mathematical and computational models of immune-receptor signalling. *Nature Rev. Immunol.* : 2004; 4:445 – 456.

73. Kaufman M, Andris F, Leo O. A model for antigen-induced T cell unresponsiveness based on autophosphorylative protein tyrosine kinase activity. *International Immunology* : 1996; 8:613 – 624.

74. Chan C, Stark J, George AJT. Feedback control of T-cell receptor activation. *Proc. R. Soc. Lond. B* : 2004; 271:931 – 939.

75. Wooldridge L, van den Berg HA, Glick M, Gostick E, Brenchley JM, Douek DC, Price DA, Sewell AK. Interaction between the CD8 coreceptor and MHC class I stabilizes TCR-antigen complexes at the cell surface. *J. Biol. Chem.* : 2005; 280:27491 – 27501.

76. Arcaro A, Grégoire C, Bakker TR, Baldi L, Jordan M, Goffin L, Boucheron N, Wurm F, Merwe PAvan der, Malissen B, Luescher IF. $CD8\beta$ endows CD8 with efficient coreceptor function by coupling T cell receptor/CD3 to raft-associated $CD8/p56^{lck}$ complexes. *J. Exp. Med.* : 2001; 194:1485 – 1495.

77. König R. Interactions between MHC molecules and co-receptors of the TCR. *Current Opinion in Immunology* : 2002; 14:75 – 83.

78. Garcia KC, Scott CA, Brunmark A, Carbone FR, Peterson PA, Teyton L. CD8 enhances formation of stable T-cell receptor/MHC class I molecule complexes. *Nature* : 1996; 384:577 – 581.

79. Luescher IF, Vivier E, Layer A, Mahiou J, Godeau F, Malissen B, Romero P. CD8 modulation of T-cell antigen receptor-ligand interactions on living cytotoxic T lymphocytes. *Nature* : 1995; 373:353 – 356.

80. McKeithan TW. Kinetic proofreading in T-cell receptor signal transduction. *Proc. Natl. Acad. Sci. USA* : 1995; 92:5042 – 5046.

81. Rabinowitz JD, Beeson C, Lyons DS, Davis MM, McConnell HM. Kinetic discrimination in T-cell activation. *Proc. Natl. Acad. Sci. USA* : 1996; 93:1401 – 1405.

82. Schodin BA, Tsomides TJ, Kranz DM. Correlation between the number of T cell receptors required for T cell activation and TCR-ligand affinity. *Immunity* : 1996; 5:137 – 146.

83. Girgis L, Davis MM, Fazekas de St. Groth B. The avidity spectrum of T cell receptor interactions accounts for T cell anergy in a double transgenic model. *J. Exp. Med.* : 1999; 189:265 – 277.

84. Chan C, George AJT, Stark J. T cell sensitivity and specificity — kinetic proofreading revisited. *Discrete and Continuous Dynamical Systems—Series B* : 2003; 3:343 – 360.

85. Valitutti S, Lanzavecchia A. Serial triggering of TCRs: A basis for the sensitivity and specificity of antigen recognition. *Immunol. Today* : 1997; 18:299 – 304.

86. Gonzales PA, Carreño LJ, Coombs D, Mora JE, Palmeiri E, Goldstein B, Nathenson SG, Kalergis AM. T cell receptor binding kinetics required for T cell activation depend on the density of the cognate ligand of the antigen-presenting cell. *Proc. Natl. Acad. Sci. USA* : 2005; 102:4824 – 4829.

87. Detours V, Mehr R, Perelson AS. A quantitative theory of affinity-driven T cell repertoire selection. *J. Theor. Biol.* : 1999; 200:389 – 403.

88. Berg HAvan den, Rand DA. Dynamics of T cell activation threshold tuning. *J. Theor. Biol.* : 2004; 228:397 – 416.

89. Nowak MA, Tarczy-Hornoch K, Austyn JM. The optimal number of major histocompatibility molecules in an individual. *Proc. Natl. Acad. Sci. USA* : 1992; 89:10896 – 10899.

90. Borghans J, Noest A, Boer RJde Thymic selection does not limit the individual MHC diversity. *Eur. J. Immunol.* : 2003; 33:3353 – 3358.

91. Müller V, Bonhoeffer S. Quantitative constraints on the scope of negative selection. *TRENDS Immunol.* : 2003; 24:132 – 135.

Figure 1

TCR interacts with pMHC in T cell:APC contact area

The antigen presenting cell (APC) takes up proteins, breaks them down into fragments called peptides, and presents selected peptides on its MHC molecules. These interact with the T cells antigen receptors (TCRs). The resulting signal arising at these TCRs is integrated by the T cell's intracellular signalling pathways, leading to a cellular response if a sufficient number of TCRs is triggered.

Figure 2

Mean TCR/pMHC residence time determines the MHC-specific TCR triggering rate

Normalized MHC-specific triggering rate $w_{ij}T_R$ as a function of normalized average residence time T_{ij}/T_R (equation (II.1)). The numbers express the free TCR density as a multiple of ligand's association constant.

Figure 3

The cellular activation threshold affects avidity spectra among responding clones

The avidity spectrum is the statistical distribution of the TCR triggering rate for a given immunogenic pMHC species j, over the population of responding TCR clonotypes i. These responders have been activated by a professional antigen presenting cell bearing the ligand j on its surface. The different curves obtain for various values of the cellular activation threshold. As the threshold increases, the mode (peak) of the distribution shifts to the right. A higher cellular activation threshold thus selects more stringently for high-avidity clones. When this threshold is too low, most of the probability mass is concentrated near zero (which is also the case, naturally, for the preselection, naive TCR repertoire). Mathematically, the avidity spectrum corresponds to the probability density function $d\mathbb{P}(w_{ij}T_R \leq \omega \mid i \text{ activated by ligand } j)/d\omega$.

Figure 4

Responder avidity spectra for various ligand presentation levels

Distribution of the normalized MHC-specific TCR triggering rate $w_{ij}T_R$ for a given immunogenic ligand j, among the responding TCR clonotypes i, as in the previous figure. Curves obtain for various amounts of immunogenic ligand present on the professional APCs that activate the naive T cells (numbers indicate the dose, normalized with respect to the lowest one shown). The probability mass shifts to the right as the dose decreases, indicating a more stringent selection for high-avidity clones at lower presentation levels on the professional antigen presenting cell.

Figure 5

Statistical distribution of the MHC-specific TCR triggering rate

The statistical survivor function (defined as complement of the cumulative distribution function) of $w_{ij}T_R$, the normalized MHC-specific TCR triggering rate, for various values of the parameters T_0 and ν (see Box III), chosen so the typical precursor probability equals 10^{-5} at $w_{ij} = 0.2/T_R$ for all curves. This statistical survivor curve plots the probability that the given foreign ligand elicits, in a randomly selected naive T cell, a triggering rate per pMHC molecule that is at least as great as the abscissa. The top curve in *Fig. 2* shows how w_{ij} depends on T_{ij}. A range of behaviours is possible; parameter values: $\nu = 1$, $T_0/T_R = 0.00000438$; $\nu = 2$, $T_0/T_R = 0.00125$; $\nu = 5$, $T_0/T_R = 0.03933$; $\nu = 10$, $T_0/T_R = 0.1244$; $\nu = 15$, $T_0/T_R = 0.1826$; the first pair giving a Bernoulli-like, nearly horizontal plateau, the last pair giving nearly exponential behaviour.

Figure 6

Tolerance alters the triggering rate distribution

Curves represent the statistical survivor function (defined as complement of the cumulative distribution function) of $w_{ij}T_R$, the normalized MHC-specific TCR triggering rate. The top curve is one of the curves in the previous figure; it applies for a ligand species j that has never been presented in deletional tolerization, as is typically the case for pathogen-derived ligands, but may also be true for certain 'cryptic' ligands derived from the host proteome. The remaining curves apply for ligands that have been presented at various presentation levels during tolerization. A truncation point is evident; this point shifts to the left as the tolerizing presentation level increases.

Figure 7

Principle of separation of activation curves

The dashed curve represents the probability of autoactivation as a function of the cellular activation threshold W_{act}. The maximum allowable probability of autoactivation, $P_{\text{auto}}^{\max}$, defines a minimum allowable activation threshold $W_{\text{act}}^{\min}$, as indicated. The bold curve represents the probability of activation by an APC carrying the immunologically relevant ligand(s). The minimum allowable probability of activation, $P_{\text{act}}^{\min}$, defines a maximum allowable activation threshold $W_{\text{act}}^{\max}$. The difference ΔW is the separation of activation curves. This separation increases as the presentation level of antigen on the APC increases. An effective immune response is possible if, at physiological presentation levels, ΔW is at least an order of magnitude wider than the natural variation of cellular activation levels in the T cell repertoire.

BOX I ***Describing fluctuations in antigen presentation***

Numerous factors impinge on the surface density of a given pMHC ligand on an APC: the quantity of protein from which the ligand is derived, or the amount of this protein taken up by endocytoses; the susceptibility of the protein to digestion and cleavage at the appropriate positions; the likelihood of the peptide being translocated to the MHC loading compartment; and the peptide's affinity with the MHC binding cleft. All these factors may vary between different tissues, between different life stages, and according to the cytokine background.

In our calculations we often use a very simple model in which all these factors are condensed into an aggregate quantity called the presentation propensity (56). If two ligands have presentation propensities r_j and $r_{j'}$, respectively, their surface densities are related by $[Z_j]/[Z_{j'}] = r_j/r_{j'}$. Moreover, we define a ligand's ubiquity u_j as the probability that it is presented on the surface of the APC. Different ligands generally have different presentation propensities and ubiquities. Propensity and ubiquity are not only specific for a ligand, but for a class of APCs as well: thus, the same ligand can have different propensities (and/or ubiquities) in different classes of APC; see equation (VI.2) for an example.

BOX II *TCR/pMHC avidity determined by mean residence time and affinity*

Peptide/MHC molecules continually associate and dissociate with TCR molecules in the T cell:APC interface. At equilibrium, the TCR triggering rate per pMHC molecule of species j is given by

$$w_{ij} = \frac{[R]}{[R] + K_{D,ij}} \exp\{-T_R/T_{ij}\}/T_{ij} \qquad \text{(II.1)}$$

where we have assumed that a TCR is triggered when the duration of a docking to a pMHC ligand exceeds the receptor threshold T_R (80, 81). In the above formula, $[R]$ denotes the surface density of free TCR molecules, T_{ij} the mean docking time of TCR clonotype i with pMHC ligand j ($1/T_{ij}$ is the off rate), and $K_{D,ij}$ is the two-dimensional dissociation constant of this interaction. The concept of TCR avidity is sometimes understood to include TCR density $[R]$ (82, 83); while our theory covers this aspect, we prefer to regard w_{ij} as a measure of TCR avidity.

A high-affinity ligand has a low dissociation constant. Such ligands typically satisfy $K_{D,ij} \ll [R]$, and then the MHC-specific TCR triggering rate is a function of mean residence time alone:

$$w_{ij} = \exp\{-T_R/T_{ij}\}/T_{ij} \; ; \qquad \text{(II.2)}$$

(see ref. 84 for a discussion of this formula) moreover, the dependence is non-monotone as was predicted on the basis of mechanistic considerations (85). The maximum rate $(eT_R)^{-1}$ is attained at an optimal mean residence time which equals the receptor triggering threshold T_R (ref. 56 and *Fig. 2*), in line with experimental data (34).

The free TCR density in the T cell:APC contact area is determined by the following implicit equation:

$$[R_T]/[R] = 1 + \sum_j \frac{[Z_j]}{[R] + K_{D,ij}} \qquad \text{(II.3)}$$

where $[Z_j]$ is the surface density of pMHC species j in the T cell:APC contact area, and $[R_T]$ denotes the total TCR density (free plus bound). As the presentation level of a high-affinity ligand increases, the free TCR density decreases and may become lower than the dissociation constant of the ligand. The MHC-specific TCR triggering rate of this ligand then becomes a monotone increasing function of the mean residence time T_{ij} and becomes proportional to the association rate, as shown in *Fig. 2*. This predicted shift from a unimodal to a monotone increasing dependence of w_{ij} on T_{ij} was confirmed experimentally (86). The T_{ij}–w_{ij} plot can again become unimodal, with an optimum at T_R, if $[M_T]$ is sufficiently high and the affinity of the ligand dominates the antigen presentation profile ($[M_T]$ denotes the total (free plus bound) surface density of MHC molecules). In this case, the MHC-specific TCR triggering rate becomes inversely proportional to $[M_T]$, which means that the total contribution made by the ligand ($[Z_j]w_{ij}$) is in fact independent of the presentation density $[Z_j]$. The possibility of presentation level-independent signalling increases the scope of repertoire selection in the thymus (35).

BOX III *Statistical distribution of mean TCR/pMHC residence times*

Let U_{ij} denote the dissociation energy barrier (in Boltzmann units) for the binding of a TCR molecule of clonotype i to a pMHC molecule of species j. The mean residence time of an interaction between these two species is given by the Arrhenius equation:

$$T_{ij} = \exp\{U_{ij}\}/f_0 \tag{III.1}$$

where f_0 is the frequency factor, which we assume independent of i or j. The dissociation energy U_{ij} is determined by a large number of interactions between the atoms at the TCR/pMHC interface. If these interactions contribute additively to the energy barrier, and if their number is large enough, the central limit theorem states that U_{ij} is normally distributed (alternatively, U_{ij} may be modelled in more detail as a complementarity score, e.g. ref. 87). Thus, transforming according to equation (III.1), we find that T_{ij} follows the log-normal distribution. It is convenient to approximate the normal distribution with a logistic distribution, which yields a two-parameter quasi-Pareto distribution for T_{ij}:

$$\mathbb{P}(T_{ij} \leq t) = \frac{1}{1 + (T_0/t)^{\nu}} \tag{III.2}$$

where $T_0 \stackrel{\text{def}}{=} \exp\{\mathbb{E}[U_{ij}]\}/f_0$ and $\nu \stackrel{\text{def}}{=} \pi/\sqrt{3\mathbb{V}[U_{ij}]}$ where $\mathbb{E}[U_{ij}]$ and $\mathbb{V}[U_{ij}]$ are mean and variance of U_{ij}.

BOX IV *Dynamics of the T cell activation threshold*

We assume that the total amount W of TCR stimulation registered by a T cell is additive, that is

$$W = \sum_j w_{ij}[Z_j] \tag{IV.1}$$

and that a given response is elicited by this stimulation if W exceeds a threshold value associated with the response (w_{ij} and $[Z_j]$ are defined in Box II). The problem is to give a mechanistic account of such a threshold, and relate it to its molecular underpinnings (co-stimulation, intracellular signalling components such as kinases and phosphatases). We describe this signalling machinery as a dynamic system with state variable $\mathbf{x}(t)$. Moreover, let us stipulate that there is a stereotypical intrinsic cellular response threshold; for instance, a gene involved in the response is activated when the cellular concentration of the relevant transcription factor exceeds a critical value, and this value is the same for all T cells. Generally: the cellular response results when one of the state variables (which we arbitrarily take to be x_1) exceeds a stereotypical value κ. On a time scale of minutes, the following dynamics is plausible for this 'excitation' variable:

$$\frac{d}{dt}x_1 = \Phi(W, \mathbf{x}) - \delta(\mathbf{x})x_1 \ . \tag{IV.2}$$

Both Φ and δ depend on the state of the signalling machinery $\mathbf{x}$: Φ is the transduction function linking the TCR signal W to increase of x_1, and δ represents the intensity of processes leading to decrease of x_1 (either degradation of the 'x_1 molecules' or conversion of these molecules into an inactive form).

The simplest model in the general class described here has just one additional state variable x_2 with 'slow' dynamics, which allows the gain and loss terms in equation (IV.2) to be adjusted to the average stimulation the T cell receives in a secondary lymphoid tissue. One example is the following system (88):

$$\frac{d}{dt}x_1 = \beta x_2 W - \delta x_1 \tag{IV.3}$$

$$\frac{d}{dt}x_2 = \nu\,(\xi - x_1) \ . \tag{IV.4}$$

This dynamics lets x_1 fluctuate about a stereotypical baseline value (ξ) *regardless of the average TCR signal* $\langle W \rangle$; with $\xi < \kappa$, activation can only ensue if a signal occurs that is significantly larger than the background average $\langle W \rangle$. Thus, as κ represents the threshold of signalling machinery downstream from the TCR, this machinery can be identical for all T cells regardless of clonotype and input history. The slow variable x_2 represents the signal transduction gain that links the TCR signal W to the excitation variable x_1. It can be shown that the cellular activation threshold equals $\kappa\langle W \rangle/\xi$ for this model, showing that the T cell activation threshold is adapted to the recent TCR input history of the individual T cell.

BOX V *The MHC presentation filter*

An important determinant of the separation of activation curves is the Simpson's diversity of the mixture of relevant, foreign, pMHC species presented on the professional APC. As this presentation diversity increases, the two curves in *Fig. 7* lie closer together, and it becomes more difficult to set the cellular activation threshold to a value that avoids autoimmunity (Box IV discusses cellular mechanisms that set this threshold). Therefore, the risk of autoimmunity is minimized when the activation threshold is adjusted to the maximum activation curve separation achieved when foreign presentation diversity is minimal, that is, only one foreign species is presented. The probability $P(1)$ of presenting exactly one peptide is

$$P(1) = pN_{\text{foreign}} \exp\{-pN_{\text{foreign}}\} \tag{V.1}$$

where N_{foreign} is the number of potential peptides in the foreign (viral) genome and p is the *presentation selectivity* of MHC presentation (this Poisson approximation follows from the ubiquity model, see Box I). Probability $P(1)$ is maximized when $p = 1/N_{\text{foreign}}$. This 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 (27).

With the activation threshold set to respond when foreign presentation diversity is minimal, a pathogen can evade detection by increasing this diversity: the pathogen can make itself invisible by making *more* of its peptides presentable. The immune system circumvents this difficulty by dividing the TCR repertoire into n_M parallel repertoires, each restricted to interacting with only one of the n_M MHC isoforms present. The probability that the pathogen escapes detection then decreases exponentially with n_M (32, 89, 90):

$$\mathbb{P}\{\text{immunodetectability}\} \approx 1 - \exp\{-n_M/2\}\,. \tag{V.2}$$

Restriction of TCR recognition to a single MHC isoform is required to preserve the maximum separation of activation curves. In sum, the phenomena of MHC presentation selectivity, MHC isoform diversity within the individual, and MHC restriction all stem from the principle of separation of activation curves.

BOX VI *Thymic antigen presentation determines how negative selection shapes the TCR repertoire*

Stringent intrinsic bounds on thymic presentation statistics preclude negative selection from fully eradicating TCR autoreactivity (91). Some autoreactivity unavoidably remains after negative selection. How can we quantify this residual autoreactivity?

We start by observing that a TCR clonotype i is fully specified by $\mathbf{w}_i \stackrel{\text{def}}{=} \{w_{ij}\}_j$, the MHC-specific TCR triggering rates for all pMHC species j (equation (II.2)). Let the parameter vector ϑ_{lymph} characterize the statistics of antigen presentation in the secondary lymphoid tissue, and consider a function X of $\mathbf{w}_i$ and ϑ_{lymph}, such that $x_i = X(\mathbf{w}_i, \vartheta_{\text{lymph}})$ represents a statistical property of autorecognition by clonotype i. The effect of thymic selection (deletional tolerance) on the statistical structure of the TCR repertoire as regards property X can then be expressed as follows:

$$\frac{\mathbb{P}(x_i > \xi \mid \text{before selection})}{\mathbb{P}(x_i > \xi \mid \text{after selection})} = S_X(\xi; \vartheta_{\text{lymph}}, \vartheta_{\text{thymus}}) \qquad \text{(VI.1)}$$

where $\vartheta_{\text{thymus}}$ characterizes the statistics of antigen presentation on negatively-selecting APCs in the thymus, and S_X is a selection function which expresses how strongly property X is affected by thymic selection.

If we know ϑ_{lymph} and $\vartheta_{\text{thymus}}$, we can determine which properties X are most strongly affected by selection. Conversely, given a property X and statistics ϑ_{lymph}, we can determine $\vartheta_{\text{thymus}}$ such that S_X is large (for a certain range of ξ-values). For instance, it can be shown (59) that the choice $\vartheta_{\text{thymus}} = \vartheta_{\text{lymph}}$ strongly selects against thymocytes that would register, if allowed to mature into naive T cells, a large *average* self stimulation in the secondary lymphoid tissues. By contrast, strong selection against the *variability* of the self signal (registered over different APCs) occurs when propensity and ubiquity (see Box I) satisfy the following condition (59):

$$\frac{u_j^{[\text{thy}]} r_j^{[\text{thy}]}}{u_{j'}^{[\text{thy}]} r_{j'}^{[\text{thy}]}} = \frac{u_j^{[\text{lym}]}(1 - u_j^{[\text{lym}]})}{u_{j'}^{[\text{lym}]}(1 - u_{j'}^{[\text{lym}]})} \left(\frac{r_j^{[\text{lym}]}}{r_{j'}^{[\text{lym}]}} \right)^2 \qquad \text{(VI.2)}$$

for every pair of self antigens j and j' (thy=thymus, lym=secondary lymphoid tissues).

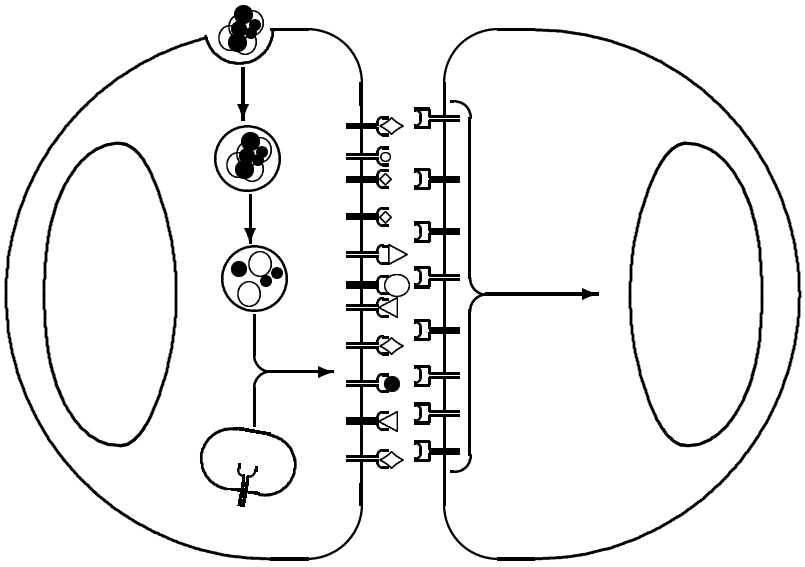

antigen presenting cell T cell

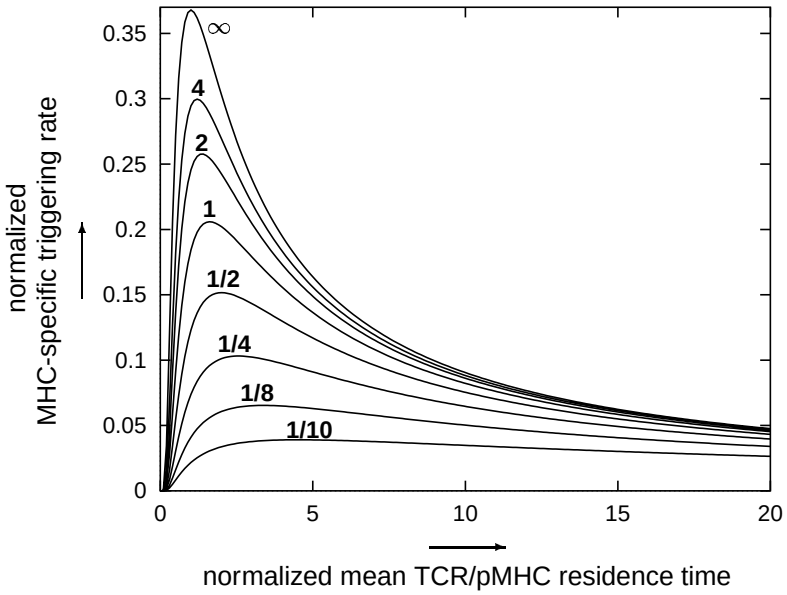
normalized
MHC-specific triggering rate
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
0
5
10
15
20
∞
4
2
1
1/2
1/4
1/8
1/10
normalized mean TCR/pMHC residence time

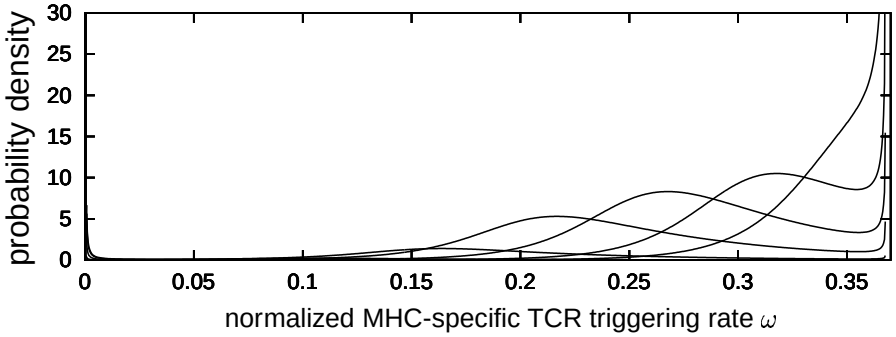
probability density
30
25
20
15
10
5
0
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35
normalized MHC-specific TCR triggering rate ω

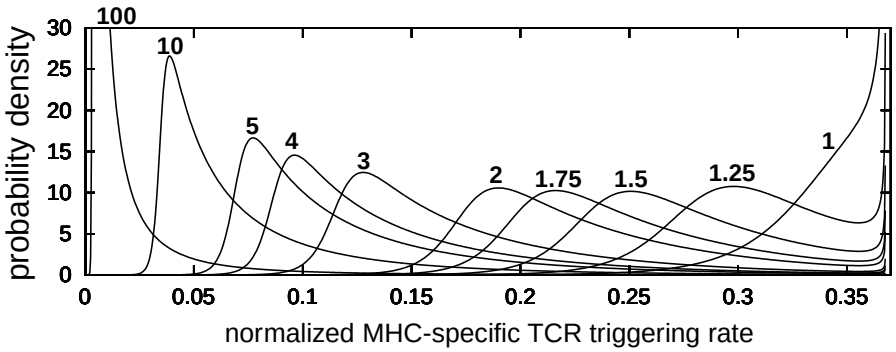
probability density
30
25
20
15
10
5
0
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35
normalized MHC-specific TCR triggering rate
100
10
5
4
3
2
1.75
1.5
1.25
1

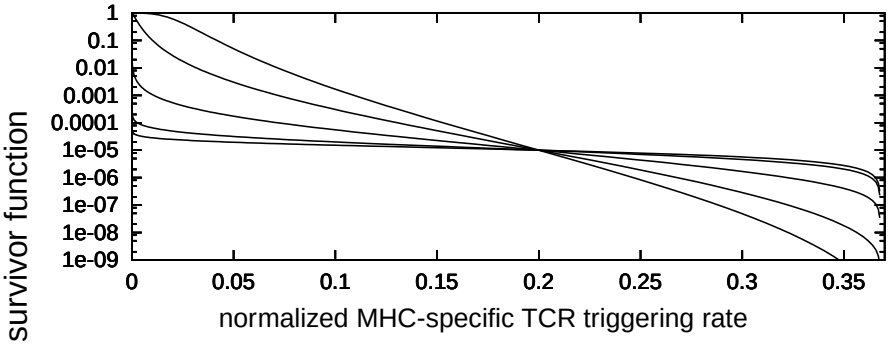
survivor function
1
0.1
0.01
0.001
0.0001
1e-05
1e-06
1e-07
1e-08
1e-09
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35
normalized MHC-specific TCR triggering rate

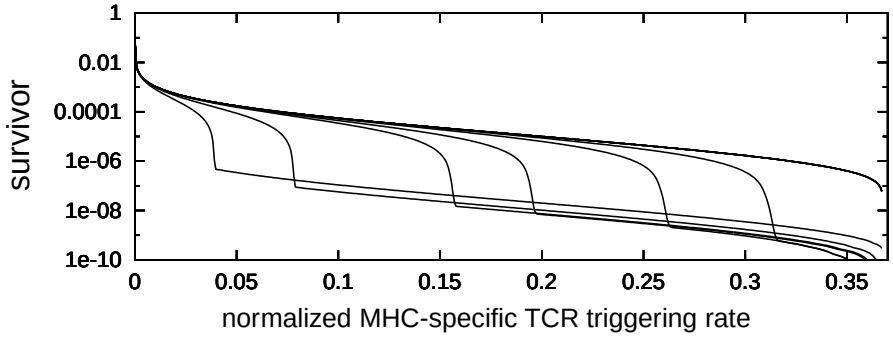
1
0.01
0.0001
1e-06
1e-08
1e-10
survivor
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35
normalized MHC-specific TCR triggering rate

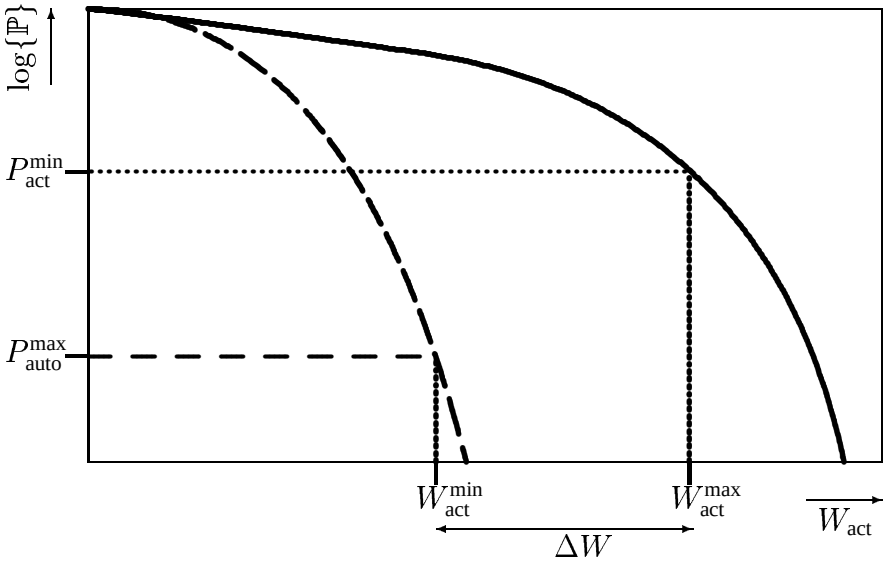
log{ℙ}
P_{act}^{min}
P_{auto}^{max}
W_{act}^{min}
W_{act}^{max}
W_{act}
ΔW