Science Progress (2011), 94(3), 245–264
doi: 10.3184/003685011X13139280383942

Specific T-cell activation in an unspecific T-cell repertoire

HUGO A. VAN DEN BERG, CARMEN MOLINA-PARÍS AND ANDREW K. SEWELL

ABSTRACT

T-cells are a vital type of white blood cell that circulate around our bodies, scanning for cellular abnormalities and infections. They recognise disease-associated antigens via a surface receptor called the T-cell antigen receptor (TCR). If there were a specific TCR for every single antigen, no mammal could possibly contain all the T-cells it needs. This is clearly absurd and suggests that T-cell recognition must, to the contrary, be highly degenerate. Yet highly promiscuous TCRs would appear to be equally impossible: they are bound to recognise self as well as non-self antigens. We review how contributions from mathematical analysis have helped to resolve the paradox of the promiscuous TCR. Combined experimental and theoretical work shows that TCR degeneracy is essentially dynamical in nature, and that the T-cell can differentially adjust its functional sensitivity to the salient epitope, "tuning up" sensitivity to the antigen associated with disease and "tuning down" sensitivity to antigens associated with healthy conditions. This paradigm of continual modulation affords the TCR repertoire, despite its limited numerical diversity, the flexibility to respond to almost any antigenic challenge while avoiding autoimmunity.

Keywords: *T-cell activation, T-cell antigen receptor, mathematical modelling, TCR repertoire, co-receptor, costimulation*

Hugo van den Berg developed mathematical models of energy budgets at the Free University of Amsterdam, where he obtained a PhD in 1998. His current research interests in the Mathematics Institute, University of Warwick, include endocrinology and energetics, mathematical models of the immune system, and the development of the excitability of the uterine muscle wall during gestation.
E-mail: hugo@maths.warwick.ac.uk

Professor Molina-París is at the School of Mathematics of the University of Leeds. Her research interests include stochastic dynamics at the molecular scale, with a focus on molecules such as DNA and microtubules, as well as the stochastic population dynamics of various classes of T-cells in the body.
E-mail: carmen@maths.leeds.ac.uk

Professor Sewell is at the Department of Medical Biochemistry and Immunology of the Cardiff University School of Medicine. His research is aimed at understanding how the T-cell immune response is regulated, and to apply these insights to diagnostics and treatments for cancer, HIV, EBV, tuberculosis, and many other human diseases. E-mail: SewellAK@cf.ac.uk

1. Introduction

T-cells are a type of white blood cell that circulate around our bodies, scanning for cellular abnormalities and infections (Figure 1). Without T-cells, human immunity does not function optimally[1]. For instance, in AIDS, one particular type of T-cell is present in lower than normal numbers and the devastating effects are all too well-known. Almost every aspect of the adaptive immune response is controlled, in some way, by T-cells. It is important to understand how T-cells are regulated, not only to enhance the beneficial responses, but also to suppress unwanted actions of T-cells. The latter include rejection of a transplanted organ, virtually all auto-immune disease (diabetes, multiple sclerosis, rheumatoid arthritis), as well as certain allergic reactions, such as gluten intolerance[1].

There are several different kinds of T-cell. They can be divided into two different types, called killer T-cells and helper T-cells. The

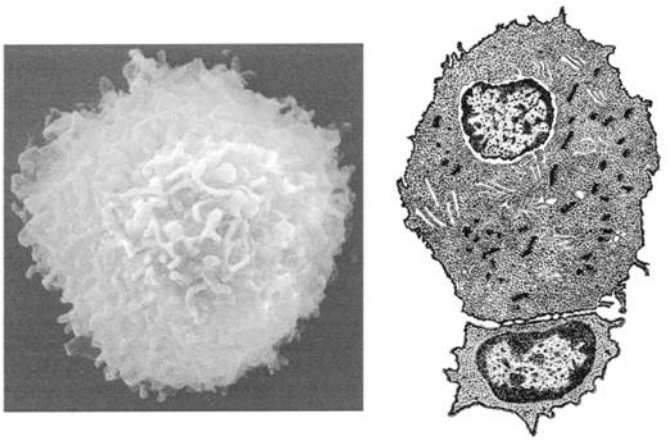

Figure 1 *T-cells. Left: a T-cell in a petri dish, seen from above. Right: a T-cell (small cell at the bottom) killing a tumour cell (large cell).*

latter are the ones that fail in AIDS. T-cells have a way of telling what is happening inside our bodies' own cells simply by scanning their surface. This mechanism allows killer T-cells to hunt down and destroy cells that are infected with germs or that have become cancerous[1]. Closely related to the killer T-cells that mediate this immune response are memory cells, which remain in the system, so that it can remember a pathogen that was encountered years or even decades ago[2,3].

1.1 How T-cells work

T-cells recognise molecular markers of disease, called antigens, via a surface receptor known as the T-cell antigen receptor (TCR). TCR recognition of antigens takes place in the contact area between a T cell and an antigen-presenting cell (APC)[4]. TCR molecules bind to ligands on the surface of the APC, as shown in Figure 2. 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[5]. As the size of these peptide fragments increases, the probability that a nonself (pathogen-derived) peptide also occurs in the human self decreases[6].

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

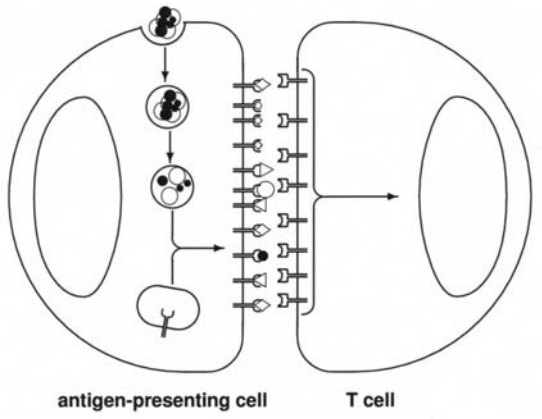

Figure 2 *Antigen presentation and the T-cell. The antigen-presenting cell acquires proteins derived from the pathogen, cuts them into small fragments called peptides, and presents these on its outer surface on MHC molecules; these interact with the T-cell antigen receptor located on the outer surface of the T-cell.*

intracellular signalling domains associated with the TCR/CD3 complex[7]. This process of TCR triggering leads to various cellular responses, such as changes in gene expression or the killing of a target cell[1].

The pMHC ligands found on a single APC form a mixed population of thousands of different species[5,8]. Most of these are 'harmless', that is, derived from host proteins not associated with disease, or from non-pathogenic organisms.

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[9].

Before they become activated, T-cells reside in the lymphoid tissues as quiescent cells. This is the TCR repertoire. It is created by random mutation and contains about a million different TCR clonotypes[10].

1.2 The controversial idea of a promiscuous TCR repertoire

We have discovered that a lack of specificity is a hallmark of the T-cell antigen receptor repertoire. This might seem a startling, if not

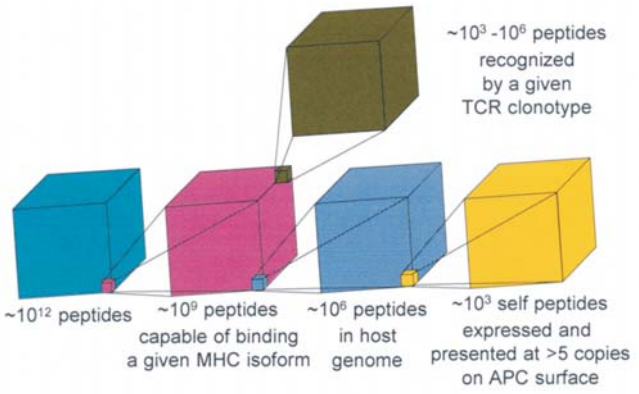

Figure 3 *The TCR repertoire must cover a vast epitope space. The numbers indicate the approximate orders of magnitude.*

foolhardy, claim in view of the undeniable fact any given TCR clonotype is capable of engaging productively with only a minute subset of all possible pMHC species. The favoured ligands for a given TCR are said to be the agonists of that TCR. Since T-cells must avoid responding to antigens not associated with harm (such as autoantigens), an intuitively obvious starting point is the assumption that TCR recognitions is one-to-one, that is, there is a specific TCR for each antigen. However, if this were the case, the repertoire would struggle for adequate coverage of the space of possible epitopes. Mason[11] pointed out that there cannot possibly be a single TCR for each antigen, since this would require an immune system vastly larger than the host organism (Figure 3). This led him to the conclusion that T-cell recognition must be highly degenerate, whatever perplexities this might entail. In other words, the minute subset is, nonetheless, very large, possibly containing many thousands of strong agonists[12,13].

The concept of a highly promiscuous TCR is controversial and has met with skepsis that continues until the present day–assuaged at least partly by the accumulating evidence for tuning mechanisms, which allow the T cell to modulate its antigen sensitivity by adjusting the expression of surface molecules and intracellular signalling machinery[14–20]. Mathematical modelling of the TCR/peptide-MHC interaction has shown just how powerful these modulatory mechanisms are, and has led to a new concept of dynamic differential regulation of functional sensitivity. Here, we give an overview of the main ideas in non-technical language; references to mathematical publications are given below and in Van

den Berg and Rand[21]. We shall concentrate mainly on our own efforts since the focus is on our controversial claim of TCR promiscuity; a more complete survey of the work of several theoretical groups has been published elsewhere[22], and the reader is urged to consult this reference for further key citations to the field.

2. TCR degeneracy

The time-honoured ''lock-and-key'' model might be taken to suggest that the TCR key either fits the pMHC lock, or fails to fit. In reality, the strength of recognition lies on a continuous scale. This is well known from experiments in which antigen-presenting cells (APCs) are incubated with various doses of the agonist, and an EC_{50} is determined. Taking the logarithm of the reciprocal of this EC_{50}, one obtains the pEC_{50} which can be viewed as a quality parameter: strong recognition is characterised by a high pEC_{50}, whereas a low pEC_{50} indicates weak recognition. The name *functional avidity* has become entrenched in the literature (as in ''high-avidity'' *versus* ''low-avidity'' clones see Alexander-Miller[23]), but given the other meanings of this term, we prefer *functional sensitivity*.

2.1 Functional sensitivity and degeneracy

The riddle of the promiscuous TCR starts to dissolve when one comes to terms with the idea that degeneracy is in essence a statistical concept. In particular, TCR degeneracy is defined completely and precisely by the statistical distribution of a given TCR's functional sensitivity values for all possible pMHC epitopes. TCR degeneracy is not a number, but a curve, as illustrated in Figure 4. This curve is a quantification of the idea that for a fixed TCR, there are few very strong agonists, perhaps scores of moderate agonists, and hundreds to thousands of weak agonists.

The degeneracy distribution can be experimentally determined and can also be calculated from the underlying statistics on kinetic parameters[24], exploiting the relation between kinetics and functional sensitivity which is explained in more detail in Box 1. The T-cell can modulate the shape of this distribution curve (Figure 4). Thus, under some circumstances, a T-cell may behave very much as predicted by the classic lock-and-key concept, with a distribution that concentrates its probability mass at the very end points of the functional sensitivity range, whereas under other circumstances, the

Box 1: Functional sensitivity and TCR/pMHC kinetics

The functional sensitivity of a pMHC ligand can be represented by the following formula:

$$x = \log \frac{e^{1-t}}{t_c + t}$$

where t is the average duration of the TCR/pMHC interaction relative to the TCR triggering threshold and t_c is a parameter which depends on the density of free TCRs on the T-cell surface (see Van den Berg *et al*.[58] for a full derivation). The graph shows functional sensitivity x as a function of mean scaled residence time t for various values of t_c. At high surface densities of free TCR molecules ($t_c \ll 1$) the curve has an optimum at $t = 1$ (this is the point where the half-life of the TCR/pMHC interaction equals the receptor triggering threshold time, multiplied by log 2). This optimum behaviour reflects the serial triggering effect, which becomes less limiting at low surface densities of free TCR molecules ($t_c \gg 1$); accordingly, the optimum becomes less pronounced. This means that when the MHC density is not limiting, the serial triggering effect disappears and ligands with long interaction times become potent agonists.

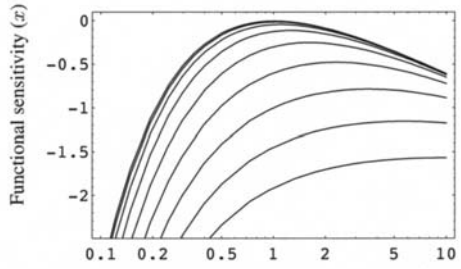

Mean TCR/pMHC interaction time relative to receptor threshold (t)

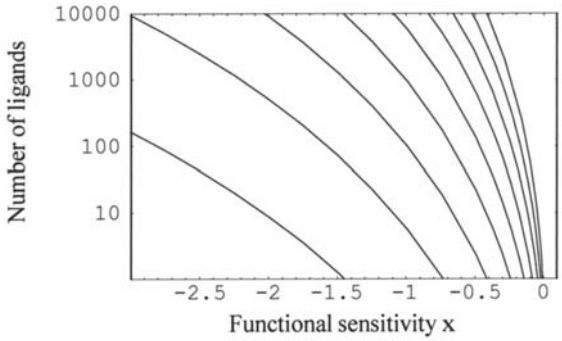

Figure 4 *The T cell can modulate the degeneracy of its TCR. Each curve represents the number of ligands with functional sensitivity at least as great as the value on the* x*-axis, for a fixed T-cell. Functional sensitivity is expressed here as the difference between the* pEC_{50} *of the ligand and that of the optimal ligand; thus less negative corresponds to a more avid ligand, 0 being the optimum. Thus, as* x *increases, fewer and fewer ligands have functional specificities greater than or equal to* x*, which is why the curves are monotone decreasing. Differences between the curves reflect different expression levels of the co-receptor and kinases/phosphorylases in the T-cell : APC contact area. It is apparent that the T-cell can modulate its degeneracy to a very large degree. Reading off the curves at the value* -0.5*, we see that the expected number of pMHC ligands whose functional sensitivity is within half a decade of the optimum can vary from less than 1 to several thousands.*

T-cell can have a range of moderate functional sensitivities (implying that the T-cell response can be modulated to be narrow or broad spectrum, or anything in between[25]).

2.2 Dynamics of TCR degeneracy modulation

The mechanisms through which the T-cell modulates its degeneracy are complex. There are two threshold levels: the triggering threshold which is a property of the individual TCR/CD3 complex, and the cellular activation threshold which is the property of the whole T-cell, and which expresses the rate and/or number of TCRs which need to be triggered to elicit a T-cell response (Figure 5). A T-cell may be capable of various responses, and each of these will have its own threshold value, giving rise to a response hierarchy[26,27]. We have recently succeeded in estimating these cellular thresholds in experiments where various types of cellular responses were measured as a function of antigen presentation levels.

Both types of thresholds are modulated by changing expression levels of CD45 isoforms, adhesion molecules, CD2, CD28/CD152 and CD4 or CD8 co-receptors, kinases and phosphorylases[28–30].

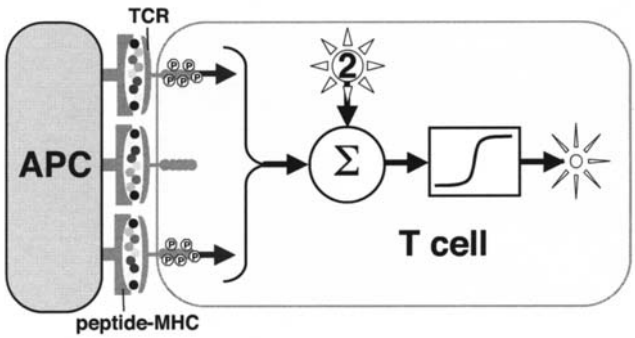

Figure 5 *The dual threshold model of T cell activation. Individual TCRs are triggered by pMHC ligands; collectively they stimulate downstream activation. In this way, the signals from the TCRs are summed and the sum serves as input into a decision function at the cellular level; modulating the sum is "signal 2" arising from costimulation.*

Details of the mathematical analysis are given by Van den Berg and Sewell[31]; here we highlight two novel findings.

2.2.1 Dual-receptor signalling in costimulation

The first finding concerns the modulation of the cellular activation threshold. This threshold is modulated by signals carried by the receptors CD28 and CD125, but with opposite effects (the former positive, the latter negative; Figure 6). It has been observed that both CD28 and CD125 can interact with both CD80 and CD86 on the APC[32–34]. This cross-talk allows the APC to influence the activation threshold of the T-cell, but only if the T-cell expresses a suitable mixture of CD28 and CD125. This is an example of a cell-to-cell signalling system where the recipient determines how it decodes the incoming signal. We have proposed that this effect is exploited by the immune system to control the oligo/polyclonality of a T-cell response, as well as the range of functional sensitivity values in the response[31].

2.2.2 TCRs tuning in on particular antigens

The second result concerns modulation by the co-receptor CD8. A combination of experimental and modelling studies[18,35,36] indicates that CD8 is not only involved in controlling degeneracy, but also in differential regulation of functional sensitivity. This means that the T-cell can increase its functional sensitivity for one ligand while

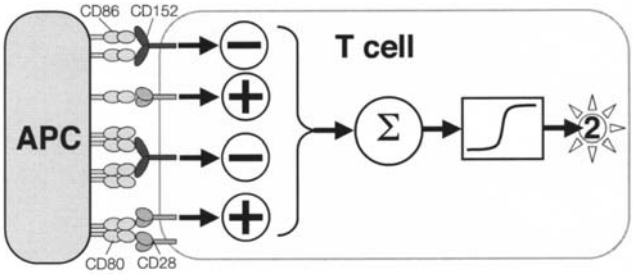

Figure 6 *Costimulation. The APC can encode information (for instance, about copy numbers of salient pMHC on its surface) in its expression levels of two receptors (CD80 and CD86). Likewise, the T-cell can determine how it decodes this information by varying the expression levels of two receptors (CD152 and CD28). The signals from CD152 and CD28 are combined to determine the strength of "signal 2", the costimulation required by naïve T-cells in addition to TCR triggering.*

simultaneously reducing the sensitivity for all others; this concept goes beyond the notion that the coreceptor generally up or down-regulates the sensitivity of the TCR[17,37]. By varying the CD8 expression level, as well as CD8 isoforms, the T-cell can ''tune in'': *i.e.*, select a sensitive epitope from among a wider group of potential agonists. This gives the T-cell repertoire a degree of flexibility which alleviates the paradox of TCR promiscuity[38].

3. Genesis of the TCR repertoire

The dynamic picture of TCR degeneracy prompts a radical rethink of the role of the selection processes in the thymus[39]. If avoidance of inappropriate selectivity is achieved through continual modulation of individual T-cells' sensitivity in the mature repertoire, where does this leave the thymic selection processes, which are generally assumed to fulfil this function?

3.1 Enforcement of MHC restriction

It is generally held that negative selection serves to eliminate self-reactive TCR clonotypes[40]. One way of being too self-reactive is violating MHC restriction. If a thymocyte recognizes two or more of the host's MHC isoforms, it will receive elevated stimulation by autoantigens. This enhanced self background signal may give rise to spurious activation (autoimmunity). The importance of allorecognition as a clinical phenomenon shows that TCRs are commonly

capable of engaging multiple MHC isoforms, implying that safeguarding this restriction with respect to the autoisoforms must be a major function of negative selection in the thymus.

An important operating parameter of the immune system is the signal/noise ratio between the autoantigen background and the pathogen-derived signal[41]. Notwithstanding its degeneracy, a TCR is highly unlikely to have high functional sensitivity to more than one of the pathogen-derived antigens displayed by the APC, and thus most of the relevant epitopes will tend to spoil the signal/noise ratio. This adverse effect is counteracted by the presence of highly selective steps in the MHC presentation pathway. This is the essence of the ''diversity filter'' theory of MHC presentation; a more detailed argument hinges on Large Deviations Theory[42]. A drawback of the filter is that any given isoform has a non-negligible probability of presenting no signal at all. The immune system compensates for this by having several TCR repertoires in parallel, kept separated by their restrictedness to MHC isoforms which is enforced by negative selection.

3.2 The target of negative selection: not merely autorecognition

One major function of negative selection is to enforce MHC restriction. Apart from this, would there be an additional need to eliminate autorecognition, given the battery of modulation mechanisms which will operate on the mature T-cell? After all, a naïve T-cell, experiencing normal self stimulation in a non-alarm context, could simply be instructed to increase these thresholds, or differentiate into a regulatory-type cell[39,43]. This suggests that the function of negative selection is not to eliminate autorecognition *per se*, but to eliminate the wrong kind of autorecognition – the kind that could easily evade control by the peripheral tuning mechanisms.

Is there such a wrong kind of autorecognition, specifically targetted by negative selection? Whereas some immunologists believe so[44], this has yet to become the consensus theory[1]. However, regardless of the merits of either point of view, the theorists can make an important contribution to the debate: *viz*., that it is possible to characterise, in an objective way, the ''kind of autorecognition'' that is targetted by negative selection. In particular, it can be shown that the target of negative selection is encoded unambiguously in the presentation patterns of autoantigens in the thymus[45].

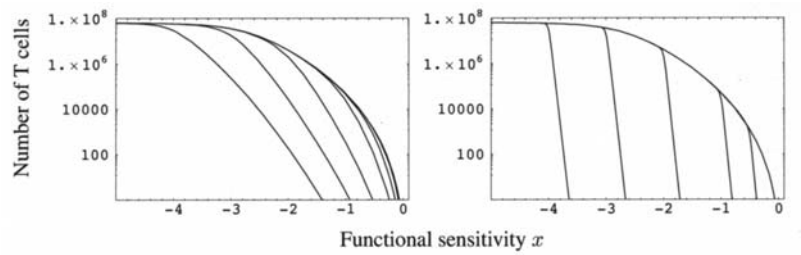

***Figure** 7 The functional sensitivity of self-recognition is determined by thymic presentation statistics. Each curve represents the number of TCR clonotypes with functional sensitivity at least as great as the value on the* x*-axis for a* fixed self peptide *(note how this differs from the previous figure). Left panel: numbers of reactive T-cells are much reduced at high functional specificities, and more strongly so as the presentation levels on negatively-selecting cells in the thymus increase (the top curve corresponds to the absence of selection). Right panel: the truncation is sharper when more negatively-selecting cells present the autoantigen.*

Important statistics that define the target of negative selection are (i) how often an autoantigen is presented on a negatively-selecting cell, and (ii) its presentation level, *i.e*. the MHC copy number at which it occurs when presented. These presentation statistics determine together the functional sensitivity distribution among the naïve T-cells in the mature TCR repertoire, as illustrated in Figure 7: the thymic presentation level determines the truncation point of the degeneracy curves, whereas the frequency of presentation determines the sharpness of the truncation[46].

The statistical structure of functional sensitivity in the TCR repertoire as a whole is determined by the thymic presentation statistics of all self antigens. The general theory is quite involved, but we can illustrate it by considering two examples of particular interest. The first is the case where the thymic presentation statistics are a truthful reflection of those in the peripheral lymphoid tissues. This thymic microcosms corresponds most closely to the classical scenario, in that most likely to be deleted are those thymocytes that would go on to register the most autorecognition as naïve T-cells. On the second scenario, thymic presentation is biased towards those autoantigens which occur only in one or a few locations in the body, but at very high levels where they do. We would then have ‘‘thymic selection for tunability,’’ where most likely to be deleted are those thymocytes which, as naïve T-cells, could not be reliably down-modulated to avert autoimmunity. The survivors are those that can adjust their cellular activation threshold dynamically[47]. On the latter scenario,

central (thymic) and peripheral (tuning) anti-autoimmune mechanisms serve distinct, complementary functions.

4. Maintaining T-cell repertoire diversity

The functional efficiency of the naïve T-cell repertoire depends on a number of factors. One is how rapidly a response, once expanded, is able to clear a pathogenic challenge. Another is how reliably an appropriate clone (one whose TCR has high functional sensitivity to one of the pathogen's epitopes) can be identified and expanded within a critical time frame following an pathogenic challenge. Intuitively we expect this reliability to be correlated with the clonal diversity in the TCR repertoire. Ecologists use a diversity measure called Simpson's diversity index[48]. This is a dimensionless number which equals 1 when all clones are present with equal numbers of cells, and which equals 0 when one clonotype numerically dominates all others in terms of T-cells belonging to that type. The typical size of an antigen-inexperienced clone is almost certainly less than several thousand cells (in humans), in sharp contrast to expanded antigen-experienced clones which may have millions of cells each.

Whereas the link between this diversity index and reliability is intuitively plain, a more careful mathematical treatment invokes the Moderate Deviations Theorem[49], which says that, as the Simpson's diversity index becomes large, the time required to activate any TCR clone becomes vanishingly unlikely to be substantially longer than the average waiting time across the TCR repertoire[50].

4.1 Stochastic dynamics of clone sizes

The next step is to understand how sensitively the diversity (which can be determined experimentally, in principle) depends on the various random influences that drive the size of any given clone up or down. A convenient quantity is the coefficient of variation (CV, *i.e.* the mean scaled by the standard deviation) of the distribution of clonal sizes in the naïve repertoire. Keeping track of millions of clones poses a formidable computational task. Fortunately, an idea known as the Ergodic Hypothesis[51] brings succour. In the present context, this idea essentially means that a series of snapshots over time of the cell numbers of any given clone gives the same general impression as a snapshot at a given moment of time of all the clones taken together.

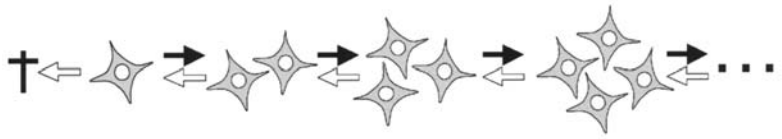

Figure 8 *Stochastic dynamics of a T-cell clone. Solid arrows correspond to cell division events, open arrows to cell deaths.*

Using this ergodicity trick, we can estimate the CV-value from the corresponding statistic of a single clone. To obtain this value, we treat the clone as a stochastic birth-and-death process. This is quite realistic, more so in fact than the deterministic treatment which has been more commonly followed in mathematical immunology. The number of cells of the clone is then represented as a continuous-time Markov chain (Figure 8). The stationary distribution of this Markov chain expresses, roughly speaking, how much time a typical clone spends at any given size. Ergodicity then allows us to apply the CV of this stationary Markov chain distribution to the entire naïve repertoire, and thus calculate its diversity from the parameters of the Markov chain[50,52,53].

A technical difficulty is that the individual clone-related Markov chains do not actually have a stationary distribution, because every clone must eventually go extinct. Fortunately, there is a theory which shows that each doomed clone will nevertheless spend most of its limited lifetime following the so-called Quasi-Stationary Distribution (QSD; *e.g.* van Doorn[54]), and the ergodicity trick still goes through when we use the QSD to furnish the coefficient of variation, as explained in Box 2.

4.2 Self stimulation and survival

Naïve T-cells are thought to receive survival stimuli that are dependent on the degree to which their TCRs are able to engage autoantigens on a specialized class of APCs[55,56]. This hypothesis is attractive since it indicates, at least on an intuitive level, a means for the repertoire to maximize functional diversity: clones would be favoured to the extent that their TCR has a recognition profile unlike that of most other clones. However, the hypothesis raises two questions. The first is how a huge number of clones can subsist on a much more modest number of autoantigens. The second is how diversity can be protected since the ecological metaphor of autoantigens as niches suggests that a system is prone to take-over by the most promiscuous, versatile TCRs through a general mechanism

Box 2: Birth and death model of the number of T-cells of a particular clone

The number of T-cells of a particular clonotype is modelled as a type of Markov process called a birth and death process[59]. The birth rate from a state with n cells is given by λ_n and the death rate from a state with n cells is given by μ_n. This is illustrated in the following diagram:

$$0 \underset{\mu_1}{\longleftarrow} 1 \underset{\mu_2}{\overset{\lambda_1}{\rightleftharpoons}} 2 \underset{\mu_3}{\overset{\lambda_2}{\rightleftharpoons}} 3 \cdots n-1 \underset{\mu_n}{\overset{\lambda_{n-1}}{\rightleftharpoons}} n \underset{\mu_{n+1}}{\overset{\lambda_n}{\rightleftharpoons}} n+1 \cdots$$

There can be no births or deaths from a state with zero cells so $\lambda_0 = \mu_0 = 0$. We say that zero is an *absorbing state*; once this state is reached, the T-cell clonotype becomes extinct. If the birth and death rates satisfy the condition

$$\sum_{n=1}^{+\infty} \frac{\mu_1 \mu_2 \cdots \mu_n}{\lambda_1 \lambda_2 \cdots \lambda_n} = +\infty$$

the process is guaranteed to reach the zero state eventually. We have shown that in our case this condition is satisfied for all values of the parameters. The mean time to extinction from an initial state with n cells can be computed from the birth and death rates and is always finite.

We denote by $p_n(t)$ the probability that there are n cells at time t. Since extinction is certain, eventually we have $\lim_{t \to +\infty} p_0(t) = 1$. Therefore, it is useful to introduce the conditional probability, $q_n(t)$, which is the probability that there are n cells at time t, given that extinction has not yet occurred. Formally, this is defined as

$$q_n(t) = \frac{p_n(t)}{1 - p_0(t)} n \geq 1.$$

The limit of this probability distribution as time goes to infinity is called the quasi-stationary probability distribution (QSD) of the process, and is denoted by q_n. The mean and variance of this distribution are given by

$$\overline{n} = \sum_{n=1}^{+\infty} n \overline{q}_n \qquad \text{and} \qquad \sigma^2 = \sum_{n=1}^{+\infty} n^2 \overline{q}_n - \overline{n}^2,$$

respectively. Simpson's diversity index for the T-cell repertoire can be computed in terms of these two quantities as follows

$$D_S = \frac{1}{1 + (\sigma/\overline{n})^2}.$$

called competitive exclusion[48]. These two problems are best addressed together since the answers are closely related.

We start from the assumptions that a TCR recognises multiple autoantigens (at relatively low functional sensitivity) and that a minimum number of these must be present in the APC's autoantigen presentation profile (APP) if the T-cell is going to receive a survival stimulus from the APC. Applying combinatorics we find an enormous richness of distinct clonotypes (compared to the – possibly quite small – set of participating autoantigens) as well as a very dilute mode of competition, as is explained in more detail in Box 3. A T-cell will compete for access to survival stimuli with almost all other clones, but for any particular APP, it will be competing with a very small subset of competitors. Moreover, the clone will find itself competing with a different group of clones for each of the APPs that can potentially furnish a survival stimulus. As a result, the impact of any one clone on the others is negligible, which means that competitive exclusion is, perhaps surprisingly, not a major factor. We see rather that a clone's chances of survival depend on a quantity called the mean niche overlap which expresses the expected number of competing clones for any given APP. Thus, while a clone typically competes with millions of other clones, for any given APP it only competes with a limited number of clones – in fact, just one or none at all, in the case of a typical clone: detailed calculations have shown that a clone with an expected niche overlap greater than 1 has a dramatically shortened life time in the repertoire. Again, this one (if any) competitor will generally be a different clone for each different APP, which is why any given clone competes only vanishingly weakly with any other clone, unless their TCRs are highly similar. The selection for low average niche overlap favours TCRs with different recognition signatures. At the same time, the sharp cut-off near unity overlap promotes diversity.

5. Perspectives

Mathematical immunology as a discipline is only justified if mathematicians can provide immunologists with a theoretical support that affords deeper insights, rigorous quantification, as well as experimentally testable hypotheses. We have tried to argue the case by exhibiting various instances where such modelling leads to new immunological ideas or experiments. Working together as mathematical modellers and as experimentalists, we have shown that TCR recognition is a graded phenomenon,

Box 3: Combinatorics creates an abundance of "niches" with a limited set of autoantigens

The number of different antigen presentation profiles can be very large, even if the number of self-peptides that provide survival signals to naïve T-cells is relatively small. We denote the number of self-peptides which provide survival signals to the naïve repertoire by N_A and the number of self-peptides presented on a single APP by N_{APP}, with $N_A \geq N_{\mathrm{APP}}$. Then to form an APP, N_{APP} of the N_A self-peptides are chosen and therefore, the number of distinct APPs is given by

$$\binom{N_A}{N_{\mathrm{APP}}} = \frac{N_A!}{N_{\mathrm{APP}}!(N_A - N_{\mathrm{APP}})!}.$$

For example, if $N_A = 1000$ and $N_{\mathrm{APP}} = 100$, the number of distinct APPs is approximately 6×10^{139} .

Combinatorics can also be used to calculate the probability that a T-cell of a given clonotype will receive a survival signal from an APP chosen at random. This probability is denoted by p. A given TCR will only interact productively with a small subset of the N_A self-peptides. This subset is referred to as the stimulatory subset of the given TCR (or clonotype) and the number of self-peptides it contains is denoted by N_a, with $1 \leq N_a \leq N_A$. For a T-cell to receive a survival signal, it may be necessary for the APP to present more than one selfpeptide from the stimulatory subset (of the given clonotype or TCR). Hence we introduce the stimulation threshold, s, which is the number of self-peptides from the stimulatory subset that need to be presented on the APP for the T-cell to receive a survival signal (note that $1 \leq s \leq N_a$). The number of distinct APPs that present s self-peptides from the stimulatory subset is $\binom{N_a}{s}\binom{N_A - N_a}{N_{\mathrm{APP}-s}}$. For an APP to provide a survival signal to T-cells of this clonotype it must present at least s of the N_a self-peptides from the stimulatory subset, so the total number of APPs providing a survival signal is $\sum_{z=s}^{N_a}\binom{N_a}{z}\binom{N_A - N_a}{N_{\mathrm{APP}-z}}$. The probability that a T-cell of this clonotype receives a survival signal from a random APP is given by the number of APPs that can provide a survival signal divided by the total number of APPs. Thus

$$p = \sum_{z=s}^{N_a} \frac{\binom{N_a}{z}\binom{N_A - N_a}{N_{\mathrm{APP}-z}}}{\binom{N_A}{N_{\mathrm{APP}}}}.$$

Similar combinatoric arguments can be used to calculate the mean niche overlap of a given clonotype.

expressible as functional sensitivity (pEC_{50}). A clone's degeneracy is exactly characterised by its functional sensitivity distribution. The fact that T-cells can modulate this degeneracy using costimulatory receptors and TCR-coreceptors suggests numerous interesting new experiments as well as immune system-based therapies. The analysis indicates that T-cells are capable of differential regulation of functional sensitivity – the key to the paradox of TCR promiscuity.

In the life sciences, the interface of traditional modelling, statistical analysis, and bioinformatics is now the purview of ''systems biologists,'' a vogue driven by an awareness that mathematics is no longer avoidable in analysis (as well as planning) of large-scale, high-throughput experiments that generate huge amounts of data[57]. No doubt this interface is crucial to immunology as well; hence systems immunology. We envisage that it will become feasible to determine the functional specificity of many TCR clonotypes for huge numbers of different peptide ligands; that it will become possible to follow the clonal sizes of naive clones over the course of weeks, months, and years; that we will be able to chart the changes of the functional specificity of a given TCR clone for a given ligand during the course of its clonal expansion and response, and simultaneously follow the expression levels of co-receptors and costimulatory receptors. To extract meaning from such data, and to be able to formulate sensible clinical interventions, systems immunologists will require the tools provided by our integrative theory.

References

1. Parham, P. (2000) *The immune system*. Garland Publishing, New York.
2. Utzny, C. and Burroughs, N.J. (2001) *Bull. Math. Biol*., **63**, 685 – 713.
3. van denBerg, H.A. and Kiselev, Y.N. (2004) *Bull. Math. Biol*., **66**, 1345 – 1369.
4. Davis, S.J., Ikemizu, S., Evans, E.J., Fugger, L., Bakker, T.R. and van der Merwe, P.A. (2003) *Nature Immunol*., **4**, 1 – 8.
5. Stevanovi, S. and Schild, H. (1999) *Semin. Immunol*., **11**, 375 – 384.
6. Burroughs, N.J., Kesmir, C. and de Boer, R. (2004) *Immunogenetics*, **56**, 311 – 320.
7. Werlen, G. and Palmer, E. (2002) *Curr. Opin. Immunol*., **14**, 299 – 305.
8. Rammensee, H.-G., Falk, K. and Rötzschke, O. (1993) *Ann. Rev. Immunol*., **11**, 213 – 244.
9. Matzinger, P. (2002) *Science*, **296**, 301 – 305.
10. Arstila, T.P., Casrouge, A., Baron, V., Even, J., Kanellopoulos, J. and Kourilsky, P. (1999) *Science*, **286**, 958 – 961.
11. Mason, D. (1998) *Immunol. Today*, **19**, 395 – 404.

12. Ignatowicz, L., Rees, W., Pacholczyk, R., Ignatowicz, H., Kushnir, E., Kappler, J. and Marrack, P. (1997) *Immunity*, **7**, 179–186.
13. Jameson, S.R. and Bevan, M.J. (1998) *Curr. Opin. Immunol*., **10**, 214–219.
14. Acuto, O. and Michel, F. (2003) *Nature Rev. Immunol*., **12**, 939–951.
15. Hutchinson, S.L., Wooldridge, L., Tafuro, S., Laugel, B., Glick, M., Boulter, J.M., Jakobsen, B.K., Price, D.A. and Sewell, A.K. (2003) *J. Biol. Chem*., **278**, 24285–24293.
16. Pecht, I. and Gakamsky, D.M. (2005) *FEBS Lett*., **579**, 3336–3341.
17. Maile, R., Siler, C.A., Kerry, S.E., Midkiff, K.E., Collins, E.J. and Frelinger, J.A. (2005) *Peripheral. J. Immunol*., **174**, 619–627.
18. Laugel, B., van den Berg, H.A., Gostick, E., Cole, D.K., Wooldridge, L., Boulter, J., Milicic, A., Price, D.A. and Sewell, A.K. (2007) *J. Biol. Chem*., **282**, 23799–23810.
19. Wooldridge, L., Lissina, A., Vernazza, J., Gostick, E., Laugel, B., Hutchinson, S.L., Mirza, F., Dunbar, P.R., Boulter, J.M., Glick, M., Cerundolo, V., van den Berg, H.A., Price, D.A. and Sewell, A.K. (2007) *Eur. J. Immunol*., **37**, 1323–1333.
20. Wooldridge, L., Laugel, B., Ekeruche, J., Clement, M., van den Berg, H.A., Price, D.A. and Sewell, A.K. (2010) CD8 controls T cell cross-reactivity. *J. Immunol*., **185**, 4625–4632.
21. van den Berg, H.A. and Rand, D.A. (2007) *Immunol. Rev.*, **216**, 81–92.
22. Molina-París, C. and Lythe, G. (eds). (2011) *Mathematical models and immune cell biology.* Springer Verlag, New York.
23. Alexander-Miller, M.A. (2000) *Cell. Immunol*., **201**, 58–62.
24. Zint, N., Baake, E. and den Hollander, F. (2008) *J. Math. Biol*., **57**, 841–861.
25. Lipsmeier, F. and Baake, E. (2009) *J. Stat. Phys*., **134**, 537–566.
26. Itoh, Y. and Germain, R.N. (1997) *J. Exp. Med*., **186**, 757–766.
27. Bachmann, M.F., Barner, M. and Kopf, M. (1999) CD2 sets quantitative thresholds in T cell activation. *J. Exp. Med*., **190**, 1383–1391.
28. Ostergaard, H.L. and Trowbridge, I.S. (1990) Coclustering CD45 with CD4 or CD8 alters the phosphorylation and kinase activity of $p56^{lck}$. *J. Exp. Med*., **172**, 347–350.
29. Azzam, H.S., DeJarnette, J.B., Huang, K., Emmons, R., Park, C.S., Sommers, C.L., El-Khoury, D., Shores, E.W. and Love, P.E. (2001) *J. Immunol*., **166**, 5464–5472.
30. Hermiston, M.L., Xu, Z. and Weiss, A. (2003) *Annu. Rev. Immunol*., **21**, 107–137.
31. van den Berg, H.A. and Sewell, A.K. (2011) In: Molina-París, C. and Lythe, G. (eds), *Mathematical models and immune cell biology.* Springer-Verlag, Berlin.
32. Collins, A.V., Brodie, D.W., Gilbert, R.J.C., Iaboni, A., Manso-Sancho, R., Stuart, D.I., van der Merwe, P.A. and Davis, S.J. (2002) *Immunity*, **17**, 201–210.
33. Sansom, D.M., Manzotti, C.N. and Zheng, Y. (2003) *TRENDS Immunol*., **24**, 313–318.
34. Zheng, Y., Manzotti, C.N., Liu, M., Burke, F., Mead, K.I. and Sansom, D.M. (2004) *J. Immunol*., **172**, 2778–2784.
35. Wooldridge, L., van den Berg, H.A., Glick, M., Gostick, E., Brenchley, J.M., Douek, D.C., Price, D.A. and Sewell, A.K. (2005) *J. Biol. Chem*., **280**, 27491–27501, 200.

36. van den Berg, H.A., Wooldridge, L., Laugel, B. and Sewell, A.K. (2007) *J. Theor. Biol*., **249**, 395–408.
37. Purbhoo, M.A., Boulter, J.M., Price, D.A., Vuidepot, A.-L., Hourigan, C.S., Dunbar, P.R., Olson, K., Dawson, S.J., Phillips, R.E., Jakobsen, B.K., Bell, J.I. and Sewell, A.K. (2001) *J. Biol. Chem*., **276**, 32786–32792.
38. Park, J.-H., Adoro, S., Lucas, P.J., Sarafova, S.D., Alag, A.S., Doan, L.L., Erman, B., Liu, X., Ellmeier, W., Bosselut, R., Feigenbaum, L. and Singer, A. (2007) *Nature Immunol*., **8**, 1049–1058.
39. Anderton, S.M. and Wraith, D.C. (2002) *Nature Rev. Immunol*., **2**, 487–497.
40. Kappler, J.W., Roehm, N. and Marrack, P. (1987) *Cell*, **49**, 273–280.
41. Noest, A.J. (2000) *J. Theor. Biol*., **207**, 195–216.
42. van den Berg, H.A. and Rand, D.A. (2003) *J. Theor. Biol*., **224**, 249–267.
43. Mason, D. (2001) *Immunol. Rev*., **182**, 80–88.
44. Seddon, B. and Mason, D. (2000) *Immunol.Today*, **21**, 95–99.
45. van den Berg, H.A. and Molina-París, C. (2003) *J. Theor. Med*., **5**, 1–22.
46. van den Berg, H.A. and Rand, D.A. (2004) *J. Theor. Biol*., **231**, 535–548.
47. van den Berg, H.A. and Rand, D.A. (2004) *J. Theor. Biol*., **228**, 397–416.
48. McGlade, J. (ed.), (1999) *Advanced ecological theory.* Blackwell, New York.
49. Dembo, A. and Zeitouni, O. (1998) *Large deviations techniques and applications*. Springer Verlag, New York.
50. Stirk, E.R., Molina-París, C. and van den Berg, H.A. (2008) *J. Theor. Biol*., **255**, 237–249.
51. Toda, M., Kubo, R. and Saitô, N. (1998) *Statistical physics*. Springer Verlag, New York.
52. Stirk, E.R., Lythe, G., van den Berg, H.A., Hurst, G.A.D. and Molina-París, C. (2010) *Math. Biosci*., **224**, 74–86.
53. Stirk, E.R., Lythe, G., van den Berg, H.A. and Molina-París, C. (2010) *J. Theor. Biol*., **265**, 396–410.
54. van Doorn, E.A. (1991) *Adv. Appl. Prob*., **23**, 683–700.
55. Tanchot, C., Lemonnier, F.A., Pérarnau, B., Freitas, A.A. and Rocha, B. (1997) *Science*, **276**, 2057–2062.
56. Veiga-Fernandes, H., Walter, U., Bourgeois, C., McLean, A. and Rocha, B. (2000) *Nature Immunol*., **1**, 47–53.
57. van den Berg, H.A. (2011) *Mathematical models of biological systems*. Oxford University Press.
58. van den Berg, H.A., Burroughs, N.J. and Rand, D.A. (2002) *Bull. Mater. Biol*., **64**, 781–808.
59. Taylor, H.M. and Karlin, S. (1998) *An introduction to stochastic modelling*. Academic Press.