Models of Self-Peptide Sampling by Developing T Cells Identify Candidate Mechanisms of Thymic Selection

Iren Bains[1,2], Hisse M. van Santen[3], Benedict Seddon[1], Andrew J. Yates[2,4]*

1 Immune Cell Biology, MRC National Institute for Medical Research, Mill Hill, London, United Kingdom, 2 Department of Systems and Computational Biology, Albert Einstein College of Medicine, New York, New York, United States of America, 3 Centro Biologia Molecular Severo Ochoa, CSIC/Universidad Autonoma de Madrid, Madrid, Spain, 4 Department of Microbiology and Immunology, Albert Einstein College of Medicine, New York, New York, United States of America

Abstract

Conventional and regulatory T cells develop in the thymus where they are exposed to samples of self-peptide MHC (pMHC) ligands. This probabilistic process selects for cells within a range of responsiveness that allows the detection of foreign antigen without excessive responses to self. Regulatory T cells are thought to lie at the higher end of the spectrum of acceptable self-reactivity and play a crucial role in the control of autoimmunity and tolerance to innocuous antigens. While many studies have elucidated key elements influencing lineage commitment, we still lack a full understanding of how thymocytes integrate signals obtained by sampling self-peptides to make fate decisions. To address this problem, we apply stochastic models of signal integration by T cells to data from a study quantifying the development of the two lineages using controllable levels of agonist peptide in the thymus. We find two models are able to explain the observations; one in which T cells continually re-assess fate decisions on the basis of multiple summed proximal signals from TCR-pMHC interactions; and another in which TCR sensitivity is modulated over time, such that contact with the same pMHC ligand may lead to divergent outcomes at different stages of development. Neither model requires that T_{conv} and T_{reg} are differentially susceptible to deletion or that the two lineages need qualitatively different signals for development, as have been proposed. We find additional support for the variable-sensitivity model, which is able to explain apparently paradoxical observations regarding the effect of partial and strong agonists on T_{conv} and T_{reg} development.

Introduction

Conventional $CD4^+$ T cells (T_{conv}) and $CD4^+$ T regulatory cells (T_{reg}) are essential components of the adaptive immune system. Conventional T cells develop effector function in response to foreign antigens, while natural T regulatory cells produced in the thymus play a key role in the maintenance of tolerance to self-antigens and prevent autoimmune diseases (reviewed in, for example, [1]). Both populations are derived from precursors in the thymus that develop, undergo selection and differentiate into different T cell lineages. The differentiation of a thymocyte into the mature $\alpha\beta$ T cell repertoire is dependent on the engagement of its T cell receptor (TCR) with endogenous peptides presented by major histocompatibility complex (MHC) molecules on thymic antigen presenting cells. Continued very weak or null interactions between the TCR and peptide-MHC ligands (pMHC) lead to failure to positively select ('death by neglect') while excessively strong TCR-pMHC interactions lead to negative selection, removing highly autoreactive cells from the T cell repertoire. However, the precise rules underlying T cell precursor fate are not well understood; based on its exposure to a sample of pMHC, how and when does a thymocyte decide to become a T_{conv}, a T_{reg}, or be deleted?

Studies using fetal thymic organ cultures have shown that there exists a sharp avidity threshold between positive and negatively selecting ligands [2,3]. There is substantial evidence indicating that T_{reg} are induced by TCR signals that lie below this negative selection threshold, but above that required for selection into the conventional T cell pool [4–8]. However, many uncertainties remain. It has been shown that expression of cognate antigen (which we loosely refer to 'agonist peptide') in the thymic epithelium is required for the generation of T_{reg} [9–13], but a recent study showed that T_{reg} commitment occurs over a wide range of TCR affinities for a ubiquitously expressed self antigen [14]. Further, the partitioning of fates with increasing strength of recognition for self (deletion→T_{conv}→T_{reg}→deletion) appears to be questioned by a study in which both expression of an agonist and a weaker partial-agonist could enhance deletion, but only the agonist was able to induce the formation of $Foxp3^+$ regulatory T cells [15], suggesting that either the mapping of avidity to fate is more complex or that qualitatively different signals are required for T_{reg} and T_{conv} selection.

Many experimental models using TCR transgenic cells (clonal populations of T cells with identical TCR) have shown that these cells can develop into both the regulatory and conventional lineages together in the same environment. This observation implies that there is stochasticity in fate determination. This stochasticity can be partitioned conceptually into two sources that are not mutually exclusive. First, there may be heterogeneity at the

Author Summary

T cells develop in the thymus, where they are vetted – they must respond weakly to self-antigens, but not so strongly as to risk causing autoimmunity. This selection process involves developing T cells being exposed to a large sample of self-peptides presented on specialised cells in the thymus, and deciding to die or to differentiate into mature T cells of either conventional or regulatory lineages. The rules by which T cells assimilate information from these interactions to make these decisions are not known. In this study we use previously published data to assess and discriminate between different models of thymic selection and find the most support for a model in which T cells vary their sensitivity to self-peptides during their development. This allows fate decisions to be made on the basis of as few as one peptide at a time, which allows for fine specificity in the selection process.

early double positive ($CD4^+CD8^+$) stage of development, even within a clonal population, that pre-disposes cells to different fates. This heterogeneity might derive, for example, from differences in expression of factors determining the baseline levels or dynamic range of TCR signalling, or other signalling proteins related to lineage commitment. Second, stochasticity may be present later in the selection process, arising at least in part because each thymocyte encounters an independent sample of self-peptide ligands. Evidence for the latter comes from observations that probabilities of deletion and T_{reg} generation have been shown to vary with levels of agonist-peptide expression; *in-vivo* studies in TCR transgenic mice [16–18] and *in-vitro* fetal thymic organ culture [19–21] have shown that the efficiency of T_{reg} selection increases with modest increases in agonist-peptide expression, but drops when expression is high. (We use the term *efficiency* here interchangeably with the probability of experiencing a given fate.) The efficiency of selection into the T_{reg} lineage also decreases in the presence of increasing numbers of cells of the same specificity [14,15,22–24]. Thus the availability of relevant ligands, either in absolute terms or through competition, can influence fate decisions.

The timing of an interaction with a ligand may also influence fate. There is evidence that the sensitivity of thymocytes to TCR stimulation is increased during maturation through the subcellular localisation of signalling molecules such as tyrosine kinase Lck [25]; the inhibition of extracellular signal-regulated kinase (ERK) activation and increased expression of inhibitory tyrosine phosphatase SHP-1 [26]; the upregulation of the negative regulator CD5 [27]; and the increased expression of ZAP-70, a downstream target of TCR signalling [28]. However, the expression of miR-181a, a microRNA that enhances sensitivity to TCR stimulation, is reduced during thymic development [29,30], and TCR signalling in response to low-affinity pMHC ligands is strongest in immature thymocytes [31]. The net effect of changes in TCR signal activating and inhibiting factors is not clear, but it is possible that stimulation with the same ligand will lead to different levels of activation in the same thymocyte at different stages of development.

The challenge of synthesising these observations and describing quantitatively how the affinity, number and timing of pMHC contacts shape the developing T cell repertoire invites a mathematical modelling approach. Models of thymic selection have been successful in providing insight into the relationship between diversity of self peptides sampled in the thymus and the cross-reactivity [32,33], alloreactivity [34], size [35,36] and

CD4SP/CD8SP ratio [37] of the selected repertoire. Models have also helped us understand the relation of HLA phenotype to viral epitope recognition [38] and the trade-off between MHC and T cell receptor diversity [39]. In this study we use stochastic (probabilistic) models to describe previously published *in vivo* data describing T_{conv} and T_{reg} commitment of a transgenic cell population in the presence of varying densities of agonist peptide in the thymus [16]. These data allow us to test and discriminate between models of how developing thymocytes might integrate signals received from pMHC ligands to make lineage decisions.

Models of thymic selection must relate the physical interaction between a TCR and a pMHC ligand to the signal interpreted or integrated by the thymocyte. There is evidence to support competing models of TCR-pMHC interactions in which the level of T activation is determined by TCR-pMHC dwell times, through the kinetic proofreading model [40,41], TCR occupancy [42–44] and overall quantitative pMHC ligand affinity [45]. Previous approaches to quantitative modelling of TCR-pMHC interactions can be divided into three broad categories: (i) detailed modelling of signal transduction immediately downstream of TCR-pMHC engagement [46,47]; (ii) kinetic models of binding events using measured rates of TCR-pMHC association and disassociation [48,49]; and (iii) the use of a 'string model' framework in which the strength of an interaction is determined by pairwise interaction energies between peptides and the aligned residues of amino acids on the variable CDR3 loop of randomly generated TCRs [32,33]. However, binding kinetic parameters are not available for the full range of endogenous peptides that are encountered during thymic development, and uncertainty remains in the relation between avidity and the signalling thresholds determining fate decisions. Here, we abstract from the mechanistic model of signal strength derived from molecular interactions. Instead we assume a distribution of signal strengths that a given TCR derives from pMHC ligands, in which low-strength signalling events occur with the greatest likelihood and, in line with our knowledge of the specificity of T cell recognition, stronger signalling events occur with decreasing probability. We show that our conclusions are insensitive to the precise form of this distribution.

We explore candidate mechanisms of T_{reg} and T_{conv} selection using the canonical hypothesis that signals associated with T_{reg} commitment are stronger than those required for T_{conv} commitment but are below the threshold for negative selection. We use the data of van Santen *et al.* [16] to reject a simple model of the selection process in which thymocyte fate is based on testing sequential single TCR-pMHC interactions. Instead, we find the data can be explained with two generalisations of this model in which perceived TCR signal strength correlates to a strict hierarchy of cell fates (neglect→T_{conv}→T_{reg}→ negative selection). In both models, thymocytes are continuously initiating fate decisions based on measuring the strength of binding to self peptide-MHC ligands in a series of encounters with antigen-presenting cells. In one class of model, the $p-$sum model, cells measure the avidity of each encounter, each of which comprises binding to a sample of multiple self-peptide-MHC ligands simultaneously. The $p-$sum model is motivated by studies implicating the integration of signals from multiple pMHC interactions in the priming of mature T cells by antigen [50–53]. In the second class of model, the two-phase model, we examine the consequences of TCR sensitivity of thymocytes varying during development. In this model a cell's interpretation of the signal derived from a given ligand depends on whether it occurs early or late in selection. We show that both models are able to describe the data, and also make predictions that are consistent with studies quantifying the efficiency of T_{reg} selection with avidity [14].

However, we argue that variable TCR sensitivity is required to explain the effect of partial and full agonist peptide expression in the thymus on T_{reg} generation and negative selection reported by Cozzo Picca and colleagues [15].

Methods

Experimental data

We use data from van Santen et al. (2004) [16] in which the frequency of high affinity intra-thymic ligands was manipulated in vivo. Briefly, a mouse line was used that employs the tetracycline inducible system to conditionally express an invariant chain mutant, bearing the T cell epitope from moth cytochrome c (MCC) in place of the class-II associated invariant chain peptide (CLIP)-encoding region (TIM). TIM was expressed in both cortical and medullary thymic epithelial cells (cTEC and mTEC), and at controllable and graded levels. The mice also contained a transgene encoding a TCR specific for this peptide, such that in the absence of induced TIM, these cells differentiated efficiently into mature AND^+ CD4 single positive thymocytes. Expression of TIM, measured by TIM RNA transcripts via real-time PCR, influenced both T_{conv} and T_{reg} formation in a non-linear fashion (Figure 1).

In Figure 1 we see that (i) low frequencies of a strong agonist (TIM) do not affect the selection of TCR-specific (AND) thymocytes into the conventional T cell pool; (ii) moderate increases in agonist expression lead to increasing efficiency of selection of AND cells into T_{reg} ($\log_{10}(T_{reg})$ against $\log_{10}$(Relative TIM RNA) between 10^{-5} and 10^{-2}; Pearson correlation

$r = 0.59$, $p < 0.002$) and a concurrent drop in the efficiency of T_{conv} selection; and (iii) high frequencies of a strong agonist lead to the deletion of AND T cells. A very similar trend was observed by Cozzo Picca et al. [17] using TCR transgenic cells specific for an epitope of influenza virus in the presence of different levels of expression of this agonist. Atibalentja et al. [18] also observed this trend following intravenous injection of varying concentrations of hen egg-white lysozyme (HEL), which was rapidly processed and presented in the thymus, resulting in the negative selection of specific TCR transgenic T_{conv} and an increase in TCR transgenic T_{reg} at low, but loss at higher, HEL concentrations.

Mathematical models

Developing thymocytes survey pMHC ligands presented on the surface of thymic epithelial cells. In all models we assume that fate decisions are continually reassessed based on 'encounters', each of which is the sum of p interactions with pMHC (Figure 2), where $p \geq 1$. Each thymocyte participates in n encounters at most, where a thymocyte might undergo negative selection, or initiate development into the T_{conv} or T_{reg} lineages, before reaching its n-th encounter. We assume that each encounter with one or more pMHC can be divided into four categories determined by its affinity or avidity and the resulting signal through the T cell receptor(s). These are (i) a weak or null signal below that required for positive selection; (ii) a signal sufficient for selection into the T_{conv} lineage; (iii) a signal that initiates selection into the T_{reg} lineage; and (iv) a strong signal that leads to deletion.

We considered two classes of models. In one, the distribution of signal strengths resulting from encounters is constant throughout
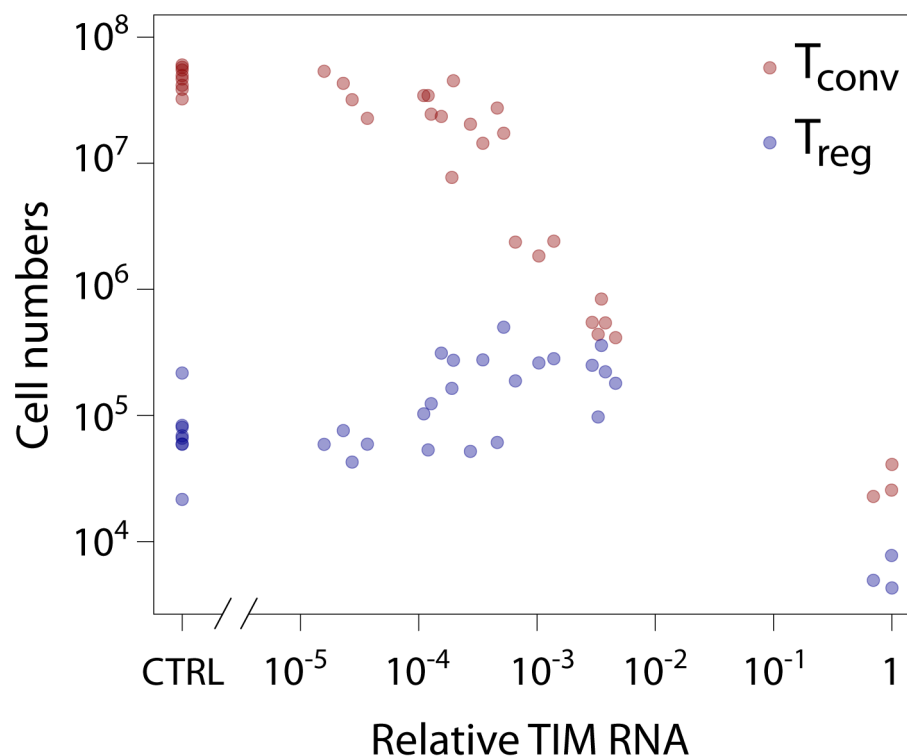
Figure 1. Data taken from van Santen et al. [16]. Absolute number of clonotype positive $CD4^+$ $CD25^-$ (T_{conv}, red circles) and $CD4^+$ $CD25^+$ (T_{reg}, blue circles) thymocytes in tetracycline-treated TAND mice, as a function of the relative expression level of TIM RNA in the thymus of these animals. Control animals lacked either the transactivator or reporter transgene.
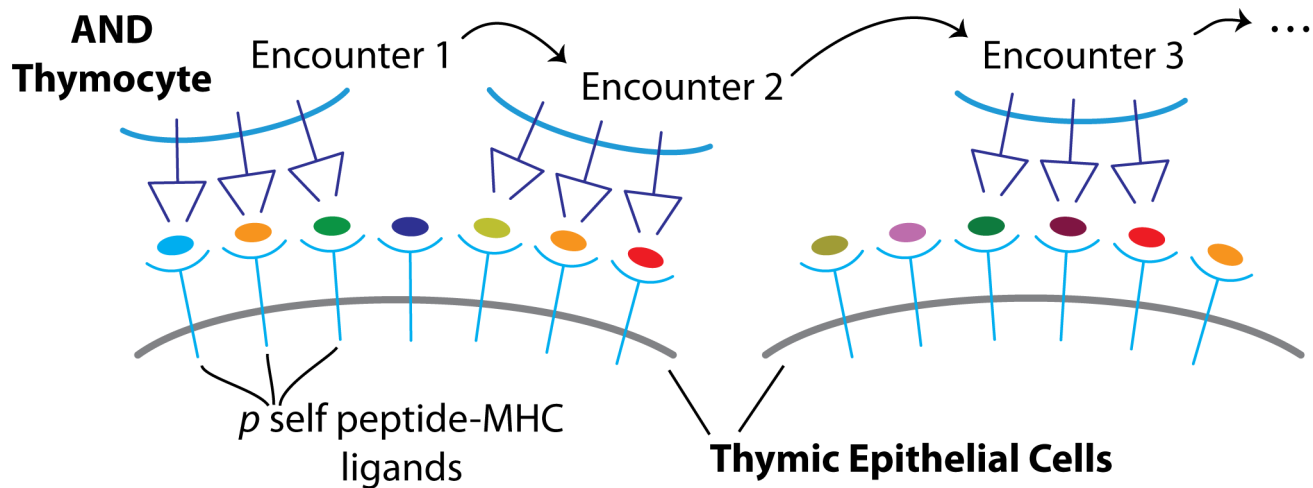doi:10.1371/journal.pcbi.1003102.g001

Figure 2. Thymocyte encounters with self-peptides. An *encounter* is defined as the simultaneous or temporally proximal binding of p TCR to p pMHC ligands on a thymic epithelial cell. Here, $p = 3$.
doi:10.1371/journal.pcbi.1003102.g002

the selection period - the $p-$sum model. In the other, the two-phase model, we allow for the possibility that this distribution shifts during selection as a result of temporal changes in TCR sensitivity.

Fate decisions made by integrating TCR signals; the $p-$sum model. The p interactions constituting an encounter may occur simultaneously, or sequentially within a time interval that is short compared to the decay time for TCR signals transduced by binding to pMHC. We consider an encounter to be the unit of information that can influence fate decisions. When one TCR binds to one randomly chosen pMHC, the contact results in a signal of strength drawn from an unknown probability distribution. One correlate of 'strength' might be the affinity of binding. Indeed affinity of binding to selecting pMHC ligands has been demonstrated to be linearly proportional to selection efficiency [14]. Similarly, when signals from multiple, proximal TCR-pMHC binding events are integrated in each encounter ($p > 1$), the resulting signal strength might be related to the avidity of the interaction. However we allow freedom in the interpretation of the term *strength* to allow for non-linear relationships between the off-rate of a TCR-pMHC complex and the signal transduced by the TCR. It is simply the quantity resulting from each encounter that the T cell uses in its fate-determination machinery. We assume a log-normal distribution of signal strengths, for reasons we discuss below.

To illustrate the calculation of selection efficiencies in this model, consider the case $p = 1$ (Figure 3A). Selection into the conventional T cell lineage requires:

1. at least one encounter with strength greater than a positive selection threshold, k_1;
2. all n encounters below a higher threshold k_2.

Experimental evidence suggests that T_{reg} development requires agonist peptide to be presented in the thymus [9,10,12,24,54]. The canonical explanation is that T_{reg} are induced by TCR signals that lie below the threshold of negative selection, but above that required for selection into the conventional T cell pool. So we define a negative selection threshold k_3, above k_2, such that an encounter between k_2 and k_3 triggers divergence into the T_{reg} lineage. For T_{reg} selection, then,

1. at least one of the n encounters is above a positive selection threshold, k_1 (to pass positive selection),
2. all n encounters are below the threshold k_3 (to avoid negative selection),
3. at least one encounter occurs between thresholds k_2 and k_3.

While this model does not contain time explicitly, the n encounters are considered to occur sequentially and so negative selection (deletion) can be initiated at any time by an encounter with strength $> k_3$. This also means that it is possible, for example, for a cell to receive a signal within the T_{reg} region and initiate development into that lineage, but later to have an encounter above k_3 and be deleted. It also means than n is an upper limit on the number of thymocyte encounters; the mean number of encounters will be fewer than n due to early termination through negative selection.

We then calculate the probability of each fate (fail positive selection, T_{conv}, T_{reg}, deletion) after n encounters. These depend simply on the probabilities α_i of an encounter falling in each region (Figure 3A, blue curve);

$$P[T_{conv} \text{ (CTRL)}] = P[\text{all } n \text{ encounters fall below } k_2]$$
$$- P[\text{all } n \text{ encounters fall below } k_1] \quad (1)$$
$$= (\alpha_1 + \alpha_2)^n - \alpha_1^n$$

$$P[T_{reg} \text{ (CTRL)}] = P[\text{all } n \text{ encounters fall below } k_3]$$
$$- P[\text{all } n \text{ encounters fall below } k_2] \quad (2)$$
$$= (\alpha_1 + \alpha_2 + \alpha_3)^n - (\alpha_1 + \alpha_2)^n$$

Now assume a proportion t of endogenous peptides are replaced by the agonist peptide TIM. Each TCR-pMHC contact will involve an endogenous peptide with probability $(1-t)$ or TIM with probability t. The signal strength derived from this contact will respectively be lognormally-distributed or with fixed strength k_{TIM}. Agonist ligands appear to induce deletion as well as T_{reg}
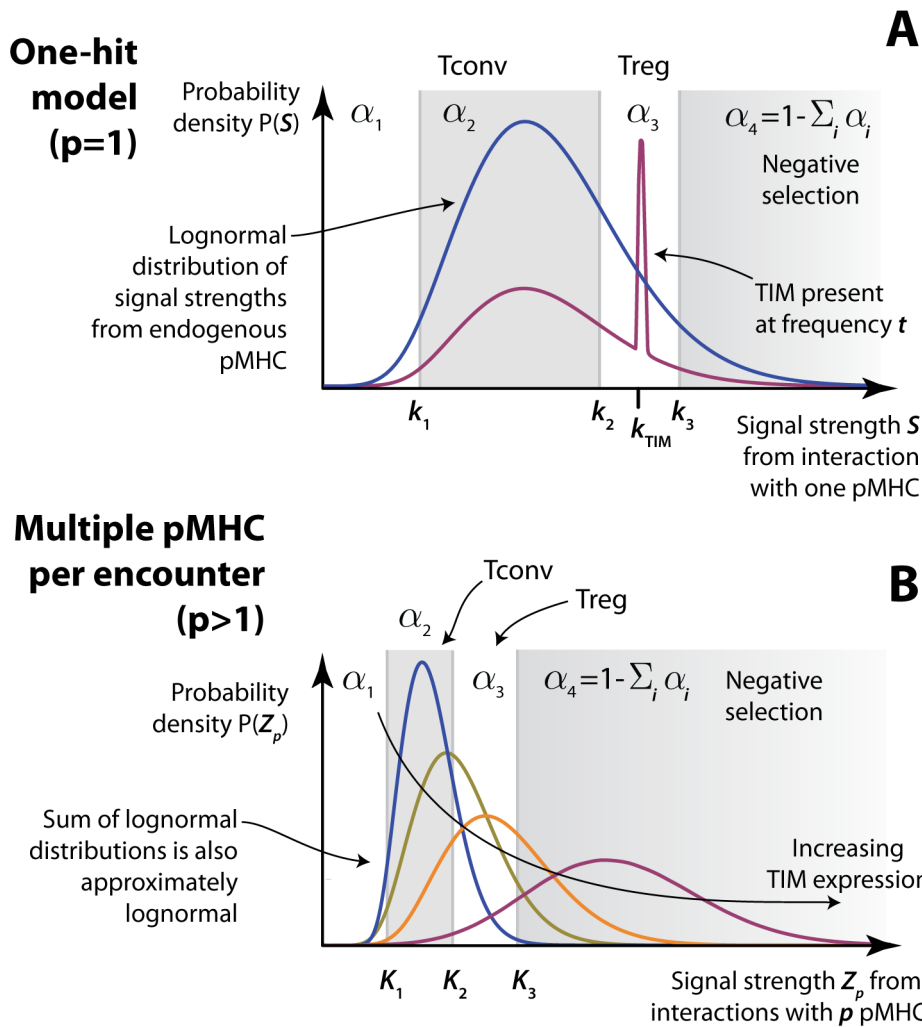
Figure 3. The influence of TIM expression on the distribution of signal strengths resulting from thymocyte encounters. A: The case $p=1$. The probabilities $\alpha_i(i=1\ldots4)$ are those of a signal lying within the different selecting regions. Blue curve; log-normally distributed signal strengths S from self pMHC in the absence of TIM expression. Area under the curve = 1. Red curve; TIM expressed at frequency t superimposed on this wild-type (endogenous pMHC) distribution. The spike at k_{TIM} (shown for convenience here with finite width and height) is a point mass in the probability distribution, of area t; the remainder is of area $(1-t)$. **B**: For $p>1$ the different selecting regions lie beween different signal strength thresholds K_1, K_2 and K_3. Increasing TIM expression t (the proportion of pMHC within an encounter, on average) shifts the distribution of signal strengths rightwards.
doi:10.1371/journal.pcbi.1003102.g003

commitment [9–13,16] and so it is likely that k_{TIM} lies above k_2. To illustrate, assume it lies within the window that triggers T_{reg} commitment, the region bounded by (k_2,k_3). At expression level t, the probabilities α_i change as follows (Figure 3A, red curve):

$$\alpha_1 \rightarrow (1-t)\alpha_1, \qquad \alpha_2 \rightarrow (1-t)\alpha_2, \qquad \alpha_3 \rightarrow (1-t)\alpha_3 + t. \quad (3)$$

For any $k_{TIM} > k_2$, the probability of selection into T_{conv} is

$$P[T_{conv}(t)] = (1-t)^n((\alpha_1 + \alpha_2)^n - \alpha_1^n) = (1-t)^n P[T_{conv}(CTRL)] \quad (4)$$

and the probability of selection into T_{reg} is

$$P[T_{reg}(t)] = \begin{pmatrix} (1-t)^n P[T_{reg}(CTRL)] \\ + \left[\left(1-t+\dfrac{t}{(P[T_{conv}(CTRL)] + P[T_{reg}(CTRL)])^{1/n}} \right)^n - (1-t)^n \right] \\ \times (P[T_{conv}(CTRL)] + P[T_{reg}(CTRL)]), \qquad \text{if } k_2 < k_{TIM} < k_3 \\ (1-t)^n P[T_{reg}(CTRL)], \qquad \text{if } k_{TIM} > k_3, \end{pmatrix} \quad (5)$$

which is a function of the maximum number of encounters n but is independent of the selection thresholds or the probabilities α_i.

When $p > 1$, we assume each encounter is of strength $Z_p = \sum_{j=1}^{p} X_j$. The X_j are identically distributed random variables representing the strength of a single TCR-pMHC binding. Each binding generates either (i) a signal arising from a randomly selected endogenous pMHC, with probability $(1-t)$, or (ii) a signal of strength k_{TIM} resulting from an interaction with TIM pMHC with probability t. We denote the selection thresholds as K_i. When p is small, the distribution of signal strengths contains point masses at k_{TIM}, $2\,k_{\text{TIM}}$, and so on. As p increases, the distribution becomes smoother and shifts rightwards with increasing t (Figure 3B). Text S1 contains the calculation of the selection probabilities for $p > 1$.

Given the complexities of TCR signalling, individual TCR binding events may not contribute linearly to an encounter's strength, however 'strength' is defined. However for our arguments all we require is (i) that endogeneous pMHC provide a smooth distribution of signal strengths arising from encounters, (ii) when agonist is present at frequency t, this distribution shifts rightwards, and (iii) the more pMHC involved in an encounter, the smoother this perturbed distribution is. The additive model is a minimal model that gives this biologically reasonable behaviour.

Variable TCR sensitivity: The two-phase model. This model is an extension of the one-hit ($p = 1$) model that also allows the sensitivity of immature thymocytes to TCR stimulation to vary during maturation. For simplicity, the selection process is divided into two distinct phases, A and B, each with maximum number of interactions n_A and n_B. In each phase the TCR-pMHC encounters are divided into distinct selecting categories, as before, with probabilities $(\alpha_1, \alpha_2, \alpha_3, \alpha_4)$ in the A-phase, where $\alpha_4 = 1 - \alpha_1 - \alpha_2 - \alpha_3$, and probabilities $(\beta_1, \beta_2, \beta_3, \beta_4)$ in the B-phase (Figure 4).

As above, we assume that a thymocyte participates in n encounters at most. At any time during development a signal above the positive selection threshold k_1 triggers T_{conv} development; this decision can be superceded by a signal above the threshold k_2 which triggers differentiation into the T_{reg} lineage; and a thymocyte can be negatively selected at any time by a signal greater than k_3.

The predictions of the model are independent of the order of the phases A and B. Here we discuss the situation in which TCR sensitivity increases during thymic development. In this scenario, k_{TIM} initially lies within the T_{reg}-selecting region, meaning that an encounter with agonist in phase A initiates T_{reg} commitment and causes deletion in phase B (Figure 3). The changes in the quantities α_i and β_i with TIM expression t follow the form of equation 3, and we then express the probabilities of selection into each lineage as a function of t and the parameters $n_A, n_B, k_{\text{TIM}}, \alpha_i$ and $\beta_i, (i = 1, 2, 3)$. Details are in Text S2.

The choice of the distribution of signal strengths. In both models, a thymocyte's fate is determined by the maximum
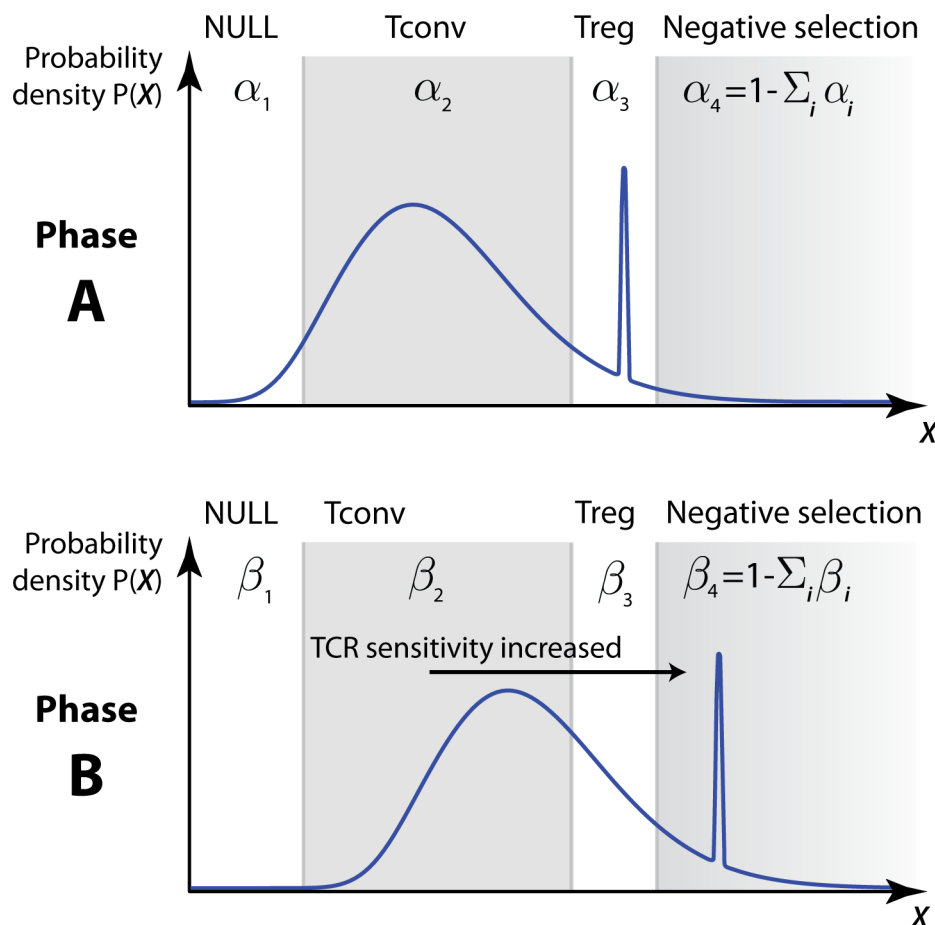

Figure 4. The two-phase model. In this instance, TCR sensitivity is assumed to increase during development, such that an encounter with the agonist ligand TIM delivers a signal that initiates T_{reg} development early in selection (phase A) but causes deletion if encountered later (phase B).
doi:10.1371/journal.pcbi.1003102.g004

signal strength experienced over a large number of encounters. We assume the strength of a single TCR-pMHC binding is log-normally distributed. The strength of an encounter (the sum of p proximal TCR-pMHC interactions) is then also approximately log-normally distributed [55]. We choose the log-normal distribution because it is ubiquitous in cell biology, and arises naturally when a random variable is derived from multiplying random variables from arbitrary distributions – such as concentrations of different signalling molecules in signal transduction pathways. However the maximum value of a large sample drawn from any heavy-tailed distribution converges to the same (Fréchet) distribution [56,57]. Each thymocyte is indeed expected to participate in a large number of encounters, and so our results hold for any heavy-tailed distribution of TCR-pMHC interaction strengths. Further, relative, not absolute, values of these signal strengths are key to the modelling of fate decisions and so we can set the mean of this distribution to be 1. The variance of the distribution is a free parameter which also does not influence our conclusions, but we discuss its influence on some parameter estimates in the Results.

Relating absolute peptide abundance to relative RNA expression. We model agonist abundance t as the fraction of endogenous peptides replaced by the agonist peptide, while the measure of agonist abundance used in ref. [16] is the relative expression of TIM RNA compared to that in control thymi. The relationship between t and TIM RNA is unknown, although we would expect it to increase monotonically. Further, a saturating level of TIM RNA is unlikely to achieve exclusive TIM expression ($t = 1$), either due to competition for loading onto MHC from endogenous peptides and/or the presence of dendritic cells in the thymus that express endogenous but not TIM peptide MHC complexes. In the absence of more information we assume a sigmoid linear-log relation that is approximately linear at low TIM expression levels and saturates at t_{max},

$$t = \frac{t_{max}}{1 + \left(\frac{\log([\text{RNA}])}{\log(B)}\right)^C} \quad (6)$$

where [RNA] is the TIM RNA expression level relative to controls and B is the expression level at which t is half-maximum. C is a measure of the steepness of the function around B and is the slope of $\log(t)$ versus $\log[\text{RNA}]$ at low TIM expression levels (Figure 5). Despite our uncertainty in the relation between relative TIM RNA and t, we will show that we can make robust statements regarding the ability of different models to describe the data.

Parameter estimation

The key parameters of interest were the encounter size p for the $p-\text{sum}$ model, and the number of encounters in each phase in the two-phase model. Other parameters were estimated simultaneously, but several quantities were taken as inputs to the models because the data from [16] did not allow us to parameterise them directly. These were (i) the parameters specifying the relation between relative RNA expression and absolute peptide abundance; (ii) the distribution of signal strengths obtained by the AND TCR from randomly sampled self pMHC ligands; and (iii) the relation between selection probabilities and absolute cell numbers. First, we explored ranges of parameters defining the mapping function (equation 6); $t_{max} \in (10^{-6}, 1)$, $B \in (10^{-6}, 1)$, and $C \in (2, 10)$. We chose to use this generic sigmoid dose-response curve given our ignorance of the mechanistic relation between RNA expression and peptide-MHC abundance on thymic epithelial cells. However, we were able to partially validate this choice of function, and the region of parameter space that we explored, using data from the study by Obst *et al.* [58]. They characterised the relation between the degree
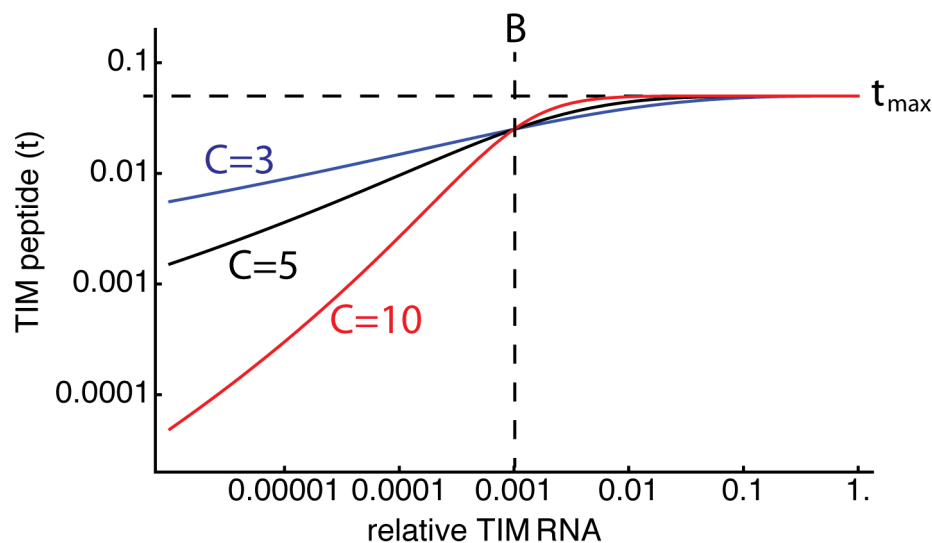

Figure 5. Relating TIM RNA to peptide abundance. We use equation 6 to connect the TIM peptide abundance t to the expression of TIM RNA relative to controls. Shown are three representative functions using different values of the breadth-parameter C, with location parameter $B = 10^{-3}$ and saturating TIM abundance $t_{max} = 0.05$.
doi:10.1371/journal.pcbi.1003102.g005

of activation of adoptively transferred AND CD4$^+$ T cells and the relative TIM RNA expression on MHC class II-expressing cells, using a similar tetracycline-inducible expression system to that used in [16]. Their readout of immune activation was the fraction of AND cells that had divided 60 h following induction of TIM expression. Assuming this fraction is linearly related to peptide availability we used the data from Obst *et al.* to estimate the parameters of the mapping function (equation 6). We found that both the recruited fraction and an alternative measure of immune activation, the estimated *per capita* rate of recruitment into division, yielded mappings within the envelope of functions generated with our parameter ranges. These mappings also lay well within the 95% uncertainty envelope generated by the best-fitting parameters from our analysis of the data from [16]. For details, see Text S3, Figure S1 and Table S1. Second, we assumed the logarithm of the signal strength derived from a single AND-TCR endogenous-pMHC interaction is normally distributed with zero mean and unit variance. The scale of the distribution of signal strengths is arbitrary and its coefficient of variation does not influence our conclusions (see Results). Third, the models provide the probabilities of selection into the T_{conv} and T_{reg} lineages and the data are absolute numbers of these populations in the thymus. We relate the numbers to probabilities through a scaling constant derived from the proportion of AND TCR cells that fail negative selection in control mice (Text S1).

Parameter estimation in the $p-$sum model. The $p-$sum model is characterised by a further six parameters $(n,p,k_{TIM},K_1,K_2,K_3)$ but three could be eliminated or constrained. First, because the AND TCR is strongly selecting we assumed that the probability of any one encounter falling below the positive selection threshold, $\varepsilon = P[Z_p < K_1]$, is small. Second, the parameter n is the upper limit on the number of encounters made by a thymocyte during selection, and is expected to be large. Thymocytes move through the medulla and cortex at similar speeds (15 μm/min and 10 μm/min, respectively) [59]. In the medulla, these speeds were shown to be associated with DC contacts at a rate of between 4 and 7 per hour, respectively. If we assume that additional contacts with TECs will contribute up to 50 contacts per hour, and that the time-spent in the thymus is between 5–10 days, then 10000 is a plausible upper bound on n. We used a conservative lower bound of $n = 500$. Thus the probability of all n encounters falling below the positive selection threshold, ε^n, is vanishingly small and we set $K_1 = 0$. Further, for a given choice of n and p, the thresholds for T_{reg} commitment (K_2) and negative selection (K_3) are determined by the observed probabilities of selection of conventional and regulatory T cells in control mice (Text S1). Selection in TIM transgenic mice using the $p-$sum model is then described by three free parameters (n, p, k_{TIM}). Only two of these can be identified uniquely, so we explored a discrete set of values of $n \in (500,1000,10000)$ and for each used a maximum likelihood approach to identify values of p and k_{TIM}, the signal strength derived from a single AND TCR contact with TIM agonist. The process was repeated across randomly sampled parameters characterising the mapping function. The residual sum of squares (RSS) and the Akaike information criterion (AIC), where with N observations $AIC = N \log(RSS)$ up to an additive constant, were used to identify the best fitting parameter values. Approximate 95% confidence intervals were generated from the parameter sets that yielded AIC values within 2 units of the lowest value.

Parameter estimation in the two-phase model. The predictions of the two-phase model are determined by the TIM mapping function and the three parameters n_A, n_B, and α_4 (Text S2). We varied α_4, the probability that a randomly sampled

pMHC in phase A will lead to negative selection, between 10^{-4} and 0.1, and used a maximum likelihood approach to identify n_A and n_B. As above, the process was repeated for a wide range of mapping functions and AIC used to identify best-fitting parameter combinations and approximate 95% confidence intervals.

Results

Without TCR sensitivity varying during development, a model in which fate decisions derive from single TCR-pMHC contacts is unable to explain the data

The key features of the data are (i) T_{conv} numbers decline monotonically with agonist expression and (ii) modest increases in agonist expression lead to an increase in the absolute number of AND T_{reg}, with numbers then decreasing at higher levels of TIM expression (Figure 1). Assuming there is a positive relationship between TIM peptide presentation (t) and relative TIM RNA expression, equation 4 shows that a model in which fate decisions are re-evaluated after single TCR-pMHC contacts ($p = 1$) can describe the T_{conv} data, which falls progressively with t.

However, we can see using a graphical argument (Figure 6, upper panel) that the $p-$sum model with constant TCR sensitivity will only be able to capture the trend in T_{reg} numbers if encounters comprise TCR signals integrated over multiple pMHC bindings ($p > 1$). If the strength of an interaction between a single AND-TCR and agonist TIM (k_{TIM}) lies within the T_{reg}-selecting range (k_2,k_3), we would expect to see a monotonic increase in T_{reg} numbers with increasing agonist peptide expression; as agonist becomes more abundant, progressively more probability mass is contained within this area, boosting the probability of T_{reg} selection (Figure 5, upper panel; Figure 7A, dotted-blue curve). Here, the one-hit model predicts that the absolute increase in T_{reg} numbers is greater than or equal to the absolute decline in T_{conv} numbers. Conversely, if k_{TIM} is above the threshold for negative selection, k_3, then we would predict a continuous decrease in T_{reg} as agonist peptide becomes more abundant and increases the probability of deletion (Figure 5, upper panel; Figure 7A, dashed red curve). Neither of these trends are what is observed and so we rule out these scenarios. Finally, we can exclude the possibility that k_{TIM} lies within the T_{conv}-selecting range; if $k_{TIM} < k_2$, increasing TIM expression would then increase the probability of selection into T_{conv}, which we do not observe. Thus we can reject the simple one-hit model for selection of AND thymocytes.

The data are consistent with a model in which fate decisions derive from integrating multiple TCR-pMHC encounters

Extending the argument above, to explain the rise and fall of T_{reg} numbers with agonist peptide expression (t) requires the probability mass within the T_{reg}-selecting region to increase then decrease with t. This becomes possible when thymocytes read multiple TCR-pMHC bindings simultaneously ($p > 1$). Qualitatively, this is because when $p > 1$, replacing an increasing fraction of endogenous peptides with TIM (t) right-shifts the distribution of encounter strengths and, in contrast to the $p = 1$ case, increases the probability of an encounter within both the T_{reg} and negative-selection regions (Figure 3B). The probability contained below the T_{reg} selection threshold falls with t, consistent with T_{conv} numbers falling; the probability α_3 of an encounter occurring within the T_{reg} zone first increases then decreases with t, as required to explain the data; and the probability of negative selection continually increases (Figure 6, middle panel).
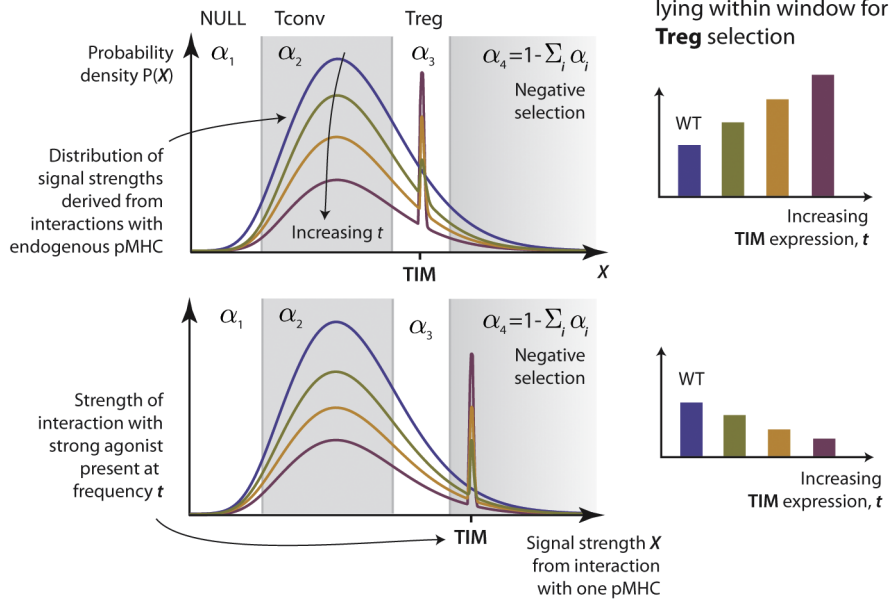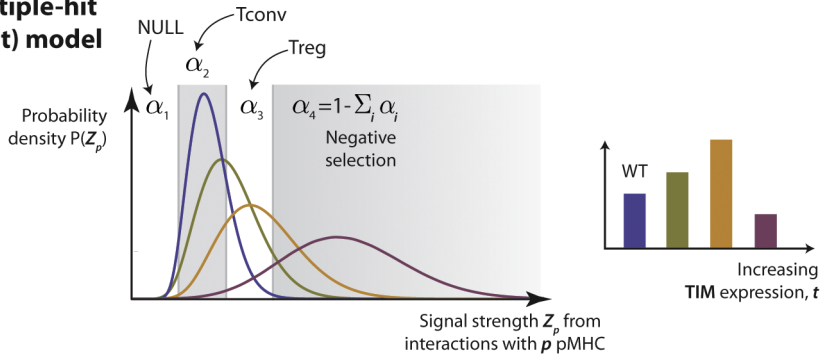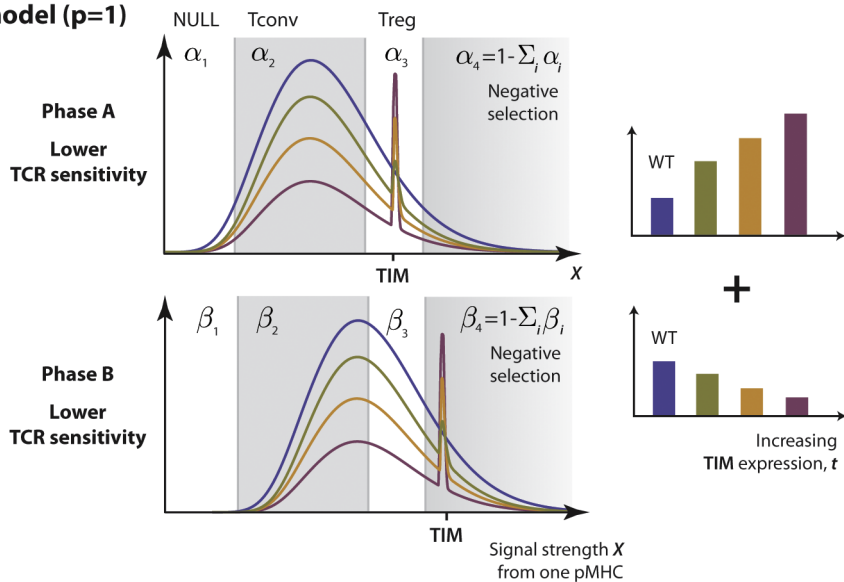
Figure 6. Modelling T_reg selection as a function of agonist abundance. Upper panel If TCR sensitivity remains static throughout thymic development, the simplest one-hit model fails to explain the rise and fall of T_reg numbers with agonist expression, because the predicted probability of receiving a T_reg selecting signal either increases or decreases monotonically. **Middle panel**. Again with static thresholds, if encounters comprise p contacts with pMHC, and $p > 1$, the distribution of signal strengths from each encounter Z_p is smoother and shifts rightwards with TIM expression, first increasing then decreasing the probability α_3 of triggering T_reg development, as required. **Lower panel**. The two-phase mode also explains the data and allows for encounters comprising single ($p = 1$, illustrated here) or multiple ($p > 1$) TCR-pMHC engagements to dictate fate. The trend in T_reg numbers arises from the balance between an increasing probability of receiving a T_reg -selecting signal with agonist expression in the low-sensitivity phase, and a decreasing probability in the higher-sensitivity phase.
doi:10.1371/journal.pcbi.1003102.g006

We explored this quantitatively and sought to identify the parameters of the $p-$sum model from the data. They cannot all be identified uniquely. As described in Methods we took the approach of exploring a range of plausible parameters governing the function mapping RNA expression to endogenous peptide replacement by TIM, and a range of values of the maximum number of encounters, n.

Remarkably, all values of n yielded equivalent descriptions of the data, and the encounter size p was highly insensitive to other parameters; it lay between 2 and 5 for all models, with best fitting value $p = 3$, independent of n. We also found that a range of mapping functions were able to describe the data equally well (Table S1). In particular, we predict that at maximum RNA expression, TIM replaces beween 0.1% and 12% of endogenous peptides. Representative fits to the data are shown in Figure 7. Panel A illustrates the failure of the one-hit model, with the best fit obtained by forcing $p = 1$. Panel B shows the fit achieved with the $p-$sum model with p a free parameter.

The estimate of p is also independent of the variance of the TCR-pMHC signal strength distribution σ^2. This also derives from the fact that the key quantities are just the probabilities α_i of interactions lying between the different thresholds K_i. However these thresholds become increasingly spaced with σ^2 (that is, as the log-normal distribution becomes increasingly fat-tailed). The less heavy-tailed the distribution of signal strengths, the smaller is the window of affinity/avidity for triggering T_reg development with respect to the mean signal strength. Small increases in affinity can shift TCR signals from positively to negatively selecting [2,3], and so if signal strength relates linearly to affinity or avidity [14], our model predicts that the distribution of encounter strengths with self may not be strongly heavy-tailed.

The estimated encounter size p increases in the presence of null peptides

Anything between ten and a few hundred pMHC have been shown to be required for T cell activation (see for example, [60]) and as few as 3–5 for pMHC recognition by cytotoxic T cell effector function [61], although with the extent of TCR binding influencing the degree of activation [42]. However, data interpreted using the kinetic proofreading model suggest that multiple interactions with very weak ligands may not lead to activation at the whole cell level (see, for example, [40,41,62]). Therefore we wanted to test whether the low estimates of p are an artefact of the assumption that every TCR-pMHC interaction generates a signal and so an encounter comprising p weak TCR-pMHC bindings might still lead to strong signalling.

To do this, we extended the $p-$sum model such that only a fraction (q) of self-peptides are capable of inducing a signal through the AND TCR, and the remaining fraction $1-q$ are classifed as null. This introduces a stochastic element to the number of TCR contributing to the signal from each encounter. We found that increasing the abundance of null ligands increases the estimated TCR engagements per encounter (Table S1). For example, we estimate the number of proximal TCR-pMHC engagements per encounter (p) to be between 20–190 if 99% of peptides fail to trigger the TCR, and between 350–1000 when 99.9% of peptides are null. Intuitively, the increase in p derives from the dilution of the information content of each encounter by the presence of null peptides. For each encounter to be a unit of sufficient information with which fate decisions can be triggered, the sample size p must increase in the presence of null interactions. As for the simpler model ($q = 1$) the estimate of p is also independent of the number of encounters, n.

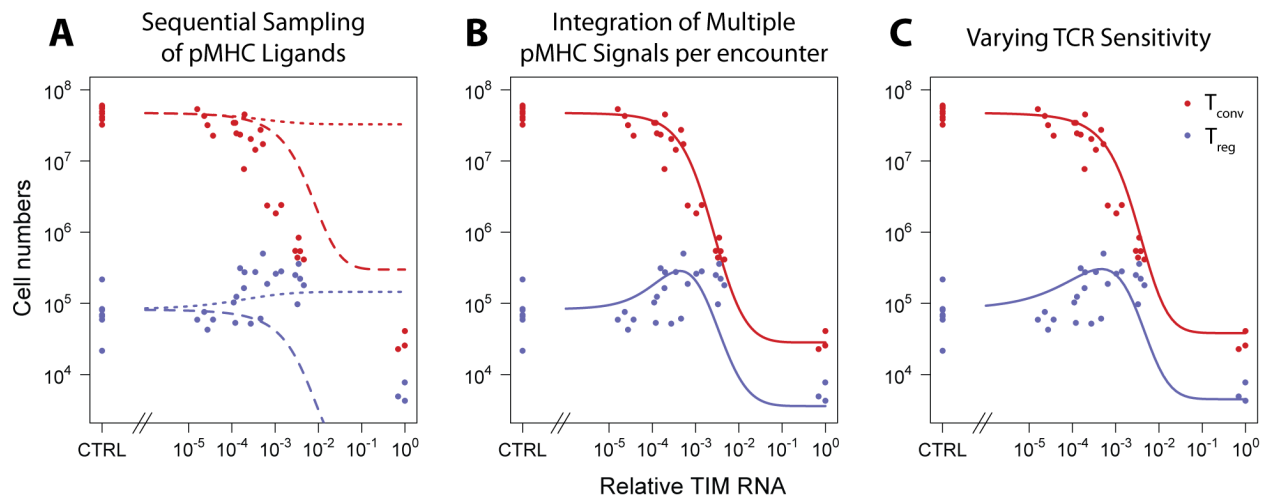**Figure 7. Model descriptions of the data.** Representative descriptions of the data by the models. T_reg in blue, T_conv in red. **Panel A**; the one-hit model in which fate decisions are re-evaluated after single TCR-pMHC contacts. The dotted curves, $k_2 < k_{TIM} < k_3$; dashed curves, $k_{TIM} > k_3$. **Panel B**; the $p-$sum model; **Panel C**; the two-phase model with $p = 1$.
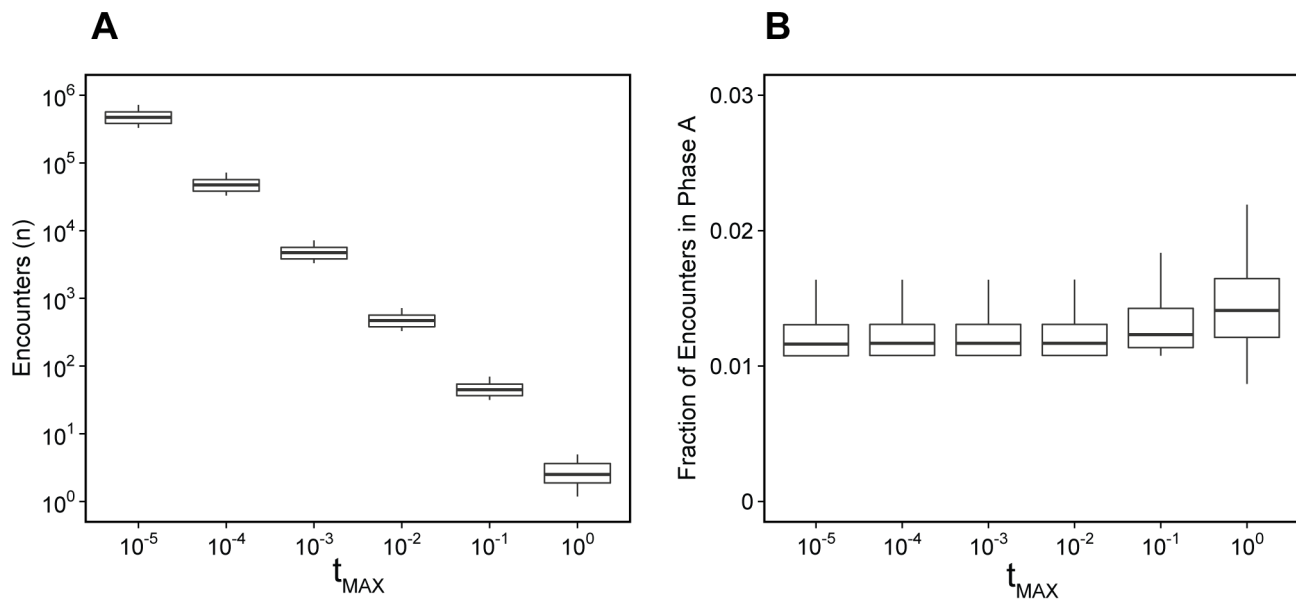doi:10.1371/journal.pcbi.1003102.g007

Figure 8. Two phase model. (A) Total number of encounters (n_A+n_B) and (B) the proportion of all encounters that take place in phase A, for plausible ranges of t_{max}, the proportion of endogenous peptides replaced by TIM at maximum RNA expression.
doi:10.1371/journal.pcbi.1003102.g008

Therefore, this extended model predicts that in the AND TCR system the expected number of productive TCR-peptide MHC interaction per encounter remains remarkably small (of the order 1). This is perhaps unsurprising, as low values of p will allow thymocytes to discriminate between ligands with small differences in affinity.

The two-phase model also explains the dependence of T_{reg} and T_{conv} on agonist expression

Next we explored the implications of a time-varying sensitivity of thymocytes to TCR stimulation during maturation. The two phase model, as described in Methods, extends the one-hit model to include time-varying TCR sensitivity. Its predictions are independent of the direction of variation, but to illustrate we assume an interaction with agonist leads to T_{reg} commitment during phase A early in development, but causes deletion in phase B when the same peptide is capable of inducing a stronger downstream TCR signal (Figure 4). Selection into the T_{reg} lineage is still possible in both phases; what changes between phase A and phase B is a right-shift in the distribution of signal strengths with respect to the selection thresholds. This shift in probabilities within the different fate-determining affinity ranges yields the required trends in T_{conv} and T_{reg} production with TIM expression (Figure 6, lower panel).

The details of parameter estimation for this model are in Methods and in Text S2. The unknowns are n_A, the number of encounters in phase A (Since $p=1$, this is the maximum number of pMHC sampled in phase A), n_B, the number of encounters in phase B, and α_4, the probability that a randomly sampled pMHC in the low-sensitivity phase A will lead to negative selection.

As for the $p-$sum model, a range mapping functions described the data equally well. A representative fit using the two-phase model is shown in Figure 7C. We found a clear inverse relationship between the value of t_{max} and the total number of encounters, n_A+n_B (Figure 8A). The model predicts that between

1–2% of encounters occur in the lower sensitivity phase (Figure 8B).

Model discrimination: The two-phase model is required to explain the effect of partial and full agonists on T_{reg} selection

We have used data in which there is a profound loss of conventional T cells in the presence of relatively low frequency of agonist peptide, while T_{reg} numbers are maintained and even initially increase with moderate increases in agonist frequency (Figure 1). These observations suggested the hypothesis that regulatory T cells are intrinsically more resilient to deletion by agonist peptides than conventional T cells [16]. Further, the study by Cozzo Picca *et al.* [15] showed that a partial agonist can induce deletion of conventional T cells but only an agonist could boost regulatory T cell generation. This led to a hypothesis that agonist peptide may deliver a qualitatively different signal that induces regulatory T cells.

We argue that neither of these hypotheses need be invoked. We have shown that both models can explain the first set of observations within a single affinity/avidity framework with different thresholds, without the need to assume differential susceptibilities of T_{reg} and T_{conv} to deletion. Further, we can see immediately that the $p-$sum model will not explain the partial/full agonist observations in ref. [15]. Their observation that partial agonist increases the probability of deletion with no increase in T_{reg} suggests that the presence of the partial agonist shifts the distribution of the sum of p interactions far to the right of the wild-type distribution, such that the bulk of the distribution is contained above the negative selection threshold. It follows that strong agonist must push this distribution even further rightwards, and so the probability of signals lying within the T_{reg}-inducing zone must fall. This is inconsistent with the observed increase in T_{reg} with agonist strength.

In contrast, the simple two-phase model can explain the effect (Figure 9). Assume that the partial agonist is not strong enough to
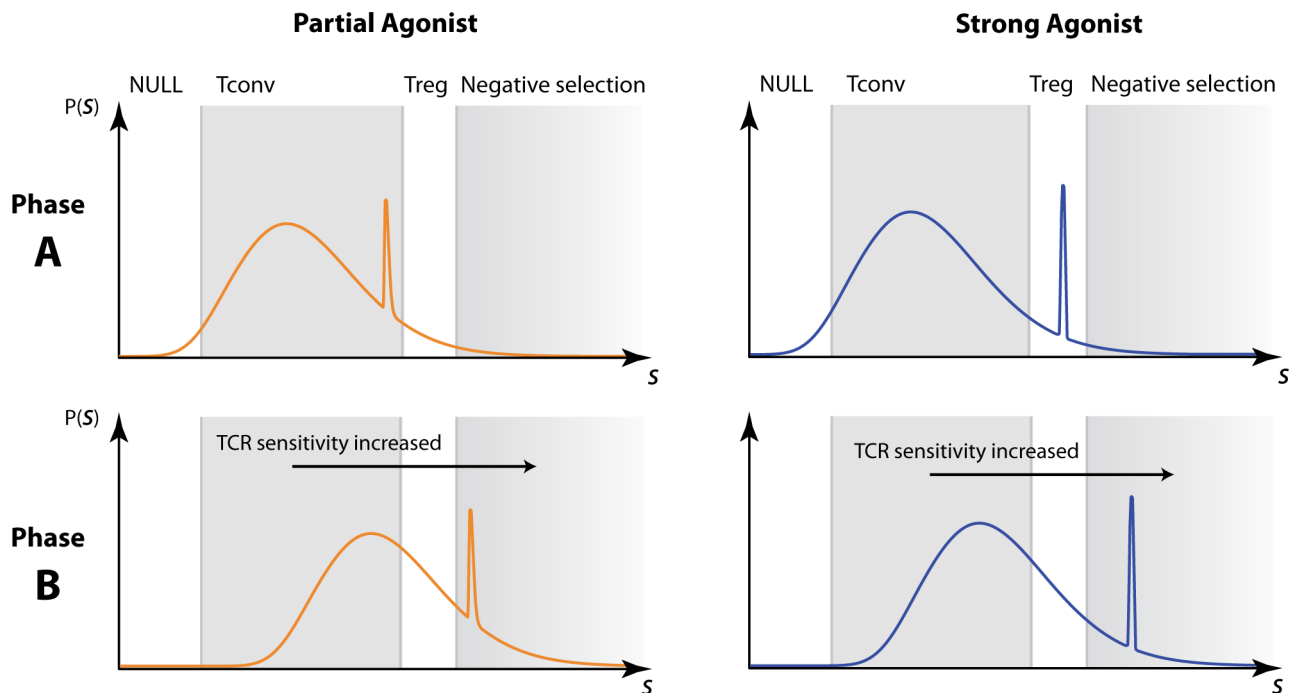
Figure 9. Using the two-phase model to explain the dependence of T_{reg} development on agonist strength. Partial agonist (left panels) may drive T_{conv} commitment in the early phase but induce deletion when TCR sensitivity is increased. In contrast, strong agonist (right panels) may drive T_{reg} commitment early, and despite triggering deletion later, the net effect is still a net increase in T_{reg} numbers. For clarity we illustrate here with the signal-strength distribution derived from $p = 1$, as in the basic two-phase model, but the argument holds also for a generalised two-phase model with encounters of size $p > 1$ (Figure S2).
doi:10.1371/journal.pcbi.1003102.g009

induce T_{reg} commitment in phase A when the TCR is relatively insensitive, but in the more sensitive phase B delivers a signal that lies above the negative selection threshold. Then the net effect of introducing a weak agonist is to increase deletion and have little effect on T_{reg} numbers, as is observed. In contrast, suppose the strong agonist triggers T_{reg} commitment when the TCR in phase A, but is negatively selecting in phase B (Figure 9, right hand columns). Then (i) expression of the strong agonist will always lead to a fall in conventional T cell numbers, and (ii) moderate levels of strong agonist, while increasing the overall probability of negative selection, can drive a net increase in T_{reg} production by boosting the probability of receiving a T_{reg}-inducing signal in phase A.

Figure 9 illustrates this effect for the case $p = 1$, but the same argument holds for a generalisation of the two-phase model with $p > 1$ (Figure S2).

Discussion

How does a thymocyte decide to become a conventional or regulatory T cell, or die, and when are these decisions most likely to take place? To address these questions, we have used experimental data to test models of how thymocytes interpret TCR interactions with self-peptide MHC to make fate decisions. We showed that the data cannot be explained with a model in which (from a given TCR's perspective) individual self-peptide MHC ligands are classed as positively selecting, negatively selecting or promoting T_{reg} development, and in which a single engagement with a strong agonist pMHC is sufficient to influence cell fate. Instead we found stronger and roughly equivalent support for two alternative models, in which (i) a thymocyte bases decisions on TCR signals summed from multiple engagements with pMHC ligands, and/or (ii) a two-phase model in which TCR sensitivity

alters during development, and so an engagement with the same pMHC ligand may lead to divergent outcomes at different stages of development. A robust prediction of the $p-$sum model is that the number of non-null TCR-self-peptide interactions per encounter that contribute to fate decisions is remarkably small. The two-phase model predicts that the initiation of differentiation into the T_{reg} lineage is most efficient during a relatively short temporal window during which the TCR is less sensitive to stimulation. Importantly, the models express the probabilistic aspect of thymic selection that likely underlies the heterogeneity in lineage decisions within a clonotypic population.

Explaining selection efficiency as a function of TCR affinity

We focused here on a system in which agonist availability could be manipulated. We can also use the models to make predictions regarding selection efficiency as a function of TCR affinity. Lee et al. [14] quantified the efficiency of T_{reg} selection in the rat insulin promoter (RIP)-mOVA system for a range of $CD4^+$ TCR clones with varying affinity for OVA. They found T_{reg} were generated across a broad range of reactivities and found an increase in T_{reg} selection efficiency with increasing affinity for a self peptide. Both the $p-$sum and variable-sensitivity models are able to explain these observations (Text S4 and Figure S3), but they make distinct predictions. The $p-$sum model predicts that the probability of the summed-proximal signals at each encounter exceeding the T_{reg} threshold is lower for weaker clones; but this probability remains constant throughout development. A variable-sensitivity model predicts that T_{reg} selection efficiency is determined by the duration of the window in which agonist contact triggers T_{reg} commitment. In an increasing-sensitivity scenario, clones that recognise OVA

more weakly will take longer to reach a level of TCR sensitivity that can induce T_{reg}, and so a prediction of the model is that lower-affinity TCR clones will commit to the T_{reg} lineage later in development.

These models might be distinguished by manipulating thymocytes' TCR sensitivity, for example through altered expression of signalling proteins, at different stages of thymocyte development. The $p-\text{sum}$ model predicts that the efficiency of T_{reg} selection would be altered equally for all clones, whereas a model of increasing TCR sensitivity predicts that damping of TCR signalling later in development would most strongly reduce T_{reg} selection efficiency for clones with low affinity for self peptide.

Robustness of results despite parameter uncertainties

We deliberately did not use model selection criteria to discriminate between models, in part because it is not possible to identify all parameters uniquely. Instead we identified regions of parameter space for each model that provided reasonable descriptions of the data. Importantly, the predicted values of the encounter size p for different abundances of null peptides (q) are insensitive to the remaining parameters (Table S1). We consider the $p-\text{sum}$ and two-phase models capable of describing the data equally well because they are able to capture the decline in T_{conv} and increase then decrease in T_{reg} with TIM abundance. Both models capture this behaviour provided TIM abundance increases monotonically with RNA expression, which we expect to be the case. Further, model selection criteria are not required to reject the simplest one-hit model, nor to compare the abilities of the $p-\text{sum}$ and two-phase models to describe the partial agonist observations.

The effect of integrating many intermediate affinity binding events

One prediction of an avidity-based model is that thymocytes may be deleted if they interact simultaneously with several pMHC at moderate affinities. We believe this prediction is not necessarily unreasonable; such events may be an inevitable byproduct of a selection process that is inherently probabilistic and which can result in cells with identical TCR experiencing divergent fates.

Generalisations of the two models

The two mechanisms we explore here are not mutually exclusive. We illustrated the two-phase model assuming that decisions are made based on single, rather than summed interactions with pMHC. However it can be generalised to an extended version of the $p-\text{sum}$ model in which encounters are interpreted differently as TCR sensitivity increases. This model will be able to explain the observations a least as well as the $p-\text{sum}$ or two phase models, at the cost of extra parameters. Also, we illustrated the impact of increasing TCR sensitivity with a simple model that divided development into two discrete phases, while increases in TCR sensitivity are likely to be continuous. Modelling smooth changes in TCR sensitivity will introduce additional parameters but we expect such a model to yield qualitatively similar results. Importantly, our analysis does not exclude the possibility that T_{reg} are more resistant to deletion and/or that qualitatively different signals are involved in T_{reg} and T_{conv} commitment; we simply show that these mechanisms need not be invoked to explain the observations.

The nature of positively selecting signals

There is evidence that positive selection requires multiple low-affinity contacts with self-peptides in the thymus [63–65]. In contrast, in our model, a single encounter above the threshold K_1

is sufficient for positive selection. We assumed this threshold is very low for the AND thymocytes, which are strongly selecting under normal conditions. These cells will presumably receive essentially continuous positively selecting signals. In the general case, and particularly for weakly signalling TCRs that are near the threshold for death by neglect, the threshold K_1 would need to be included in the parameter search and the model would be extended to include a memory of recent interactions; one possibility is a model in which levels of survival or fate-determining proteins are increased by TCR signalling but decay in its absence.

The role of precursor frequency in limiting T_{reg} development

There is substantial evidence that increasing the frequency of a given clonal (single TCR specificity) population reduces its efficiency of selection and in particular the probability of being directed into the T_{reg} lineage [14,22,24]. Our models treat cells as independent entities and do not explicitly incorporate the possibility that competition between thymocytes of similar TCR specificities might influence the availability of selecting ligands. However one mechanism of competition can be represented quite straightforwardly in the models. If the strength of an encounter correlates with its duration, or perhaps increases the probability of internalisation of the pMHC ligand by the thymocyte, the TCR-specific cells will compete for and possibly sequester agonist and other high-avidity pMHC ligands. This will shift the apparent distribution of signal strengths leftwards towards lower avidity (and more available) interactions, reducing the probability of acquiring T_{reg}-selecting signals. This model of competition for higher-avidity pMHC ligands may also explain the observation that the efficiency of T_{conv} selection can increase with precursor frequency [22]. However it remains an open question whether competition for pMHC plays an important role in selection at physiological precursor frequencies.

Cytokine requirements for T_{reg} development

Signalling through the IL-2 receptor is a requirement for T_{reg} development [66–68]. It is thought that strong TCR signalling below the negative selection threshold may sensitise cells to IL-2, licensing progression towards the T_{reg} lineage. Whatever the precise role for IL-2, it must operate downstream of fate-determining signals if selection is governed by a hierarchy of TCR avidity thresholds. Nevertheless, if IL-2 is limiting it may provide an upper bound on the total rate of production of T_{reg}, either by redirection of cells to the T_{conv} lineage or through loss. We argue however that competition for non-specific factors is unlikely to play a significant role in the system we are working with. First, the source of the IL-2 is unclear but we can reasonably suppose that IL-2 production in this system is independent of TIM expression. T_{reg} numbers increase with TIM at low expression levels, indicating that IL-2 cannot be limiting in this region (Figure 1). Similarly it cannot be limiting at higher TIM levels as T_{reg} decrease. It remains possible that a capping of T_{reg} production through competition for IL-2 may be occuring in a small flat region near the peak in T_{reg} numbers, but competition for non-specific factors alone cannot explain the key aspects of T_{reg} development we are attempting to describe.

Predictions of the models regarding the timing of regulatory T cell commitment

Early neonatal thymectomy experiments suggested that the development of T_{reg} is delayed compared to conventional T cells [69–71]. A key T_{reg} marker, the transcription factor Foxp3, is

predominantly observed in the mature CD4 single positive stage of thymocyte development [72]. However, there may be a lag between initiation of T_{reg} development and the stable expression of Foxp3; and it is possible that factors required for T_{reg} development such as IL-2 [66–68] or medullary thymic epithelial cells [73] may only be required later in the maturation process. Thus the timing of T_{reg} commitment remains unclear. The two models explored here make different predictions regarding this timing. The $p-$sum model suggests that diversion into the regulatory T cell lineage occurs with constant probability per encounter throughout selection. In contrast, the key prediction of the two-phase model is that the development of AND TCR T_{reg} is triggered most efficiently within a relative short window during which thymocytes are relatively insensitive to TCR signalling. This window is estimated to span roughly 2% of the total pMHC ligands sampled, with the caveat that the two-phase model is an abstraction of what is more likely to be a temporally graded shift in sensitivity.

The two-phase model's predictions are identical whether the shorter, less sensitive phase occurs early or late in development. However, expression of the downstream TCR-signalling protein Zap70 increases progressively between the double positive (DP) and single positive stages of thymocyte development, and is associated with increasing sensitivity to TCR stimulation [28]; immature DP thymocytes display lower surface expression of TCRs, as compared to mature single positive (CD4$^+$ or CD8$^+$) thymocytes, and TCR signalling may be actively inhibited in immature DPs [74]. Thymocytes' signalling environment may also change as they develop. Selection begins in the thymic cortex, where pMHC are encountered on cortical thymic epithelial cells, before cells migrate to the medulla where they encounter pMHC on medullary thymic epithelial cells and dendritic cells. It is thought that levels of co-stimulation and antigen presentation are generally lower in the cortex than in the medulla [75–77], suggesting that there may be an effective increase in TCR signalling during development. Using these observations, the two-phase model predicts that T_{reg} development begins predominantly, but not exclusively, during a short window at the earliest double positive stage of selection. Clearly, definitive experimental identification of when T_{reg} development is initiated will substantially increase our understanding of how thymocytes process information.

Ligand discrimination — a role for time-varying TCR sensitivity in the thymus?

Reliable recognition and discrimination of self and nonself ligands requires both TCR sensitivity and specificity. Specificity will decrease as the number of pMHC integrated per encounter (p) increases — when p is large, many pMHC engagements are integrated at each encounter, and so the thymocyte is then just sampling the mean of the distribution of pMHC affinities, and information is lost. This may explain why the optimum values of p in the $p-$sum model are at the lower end of the reported numbers of pMHC engagements required for T cell activation; T cell activation may invole a relatively small number of informative TCR recognition events, together with many brief engagements with null or very low affinity peptide ligands. Our analysis of the $p-$sum model places small lower bounds on the number of non-null TCR engagements per encounter; but the two-phase model explains the data by allowing even a single non-null pMHC recognition event to influence fate. We speculate that varying TCR sensitivity with time in the thymus may allow for increased specificity in self-nonself discrimination.

Supporting Information

Figure S1 Fitting the TIM mapping function to various readouts of immune activation derived from the data in Obst _et al._ [58]. Upper row: Our proposed sigmoid mapping function yielded reasonable descriptions of three readouts of immune activation, all assumed to be proportional to peptide abundance on APC. **Lower panel**: we show the best fitting functions derived from the three measures superimposed on the 95% uncertainty envelope in mapping functions derived from the best-fitting parameters of the $p--$sum model. Assuming TIM abundance was proportional to the recruited fraction at 60 h (red curve) or to the per capita rate of recruitment (orange curve) yielded functions that lay well within this region.
(EPS)

Figure S2 Explaining the effect of partial and strong agonists on T_{reg} selection efficiency. Figure 9 in the main text illustrates how the observations of Cozzo Picca _et al._ [15] regarding the effect of partial and strong agonists on T_{reg} selection efficiency can be explained with the two phase model. This model can be extended to include multiple pMHC per encounter and with it a similar schematic can be used to explain the observations. As before, the temporal order of the low and high TCR sensitivity phases has no effect on our results; here we illustrate the argument for an increasing-sensitivity model. The presence of partial agonist may make T_{conv} commitment far more likely than T_{reg} commitment in phase of lower TCR sensitivity (upper left panel). Increasing TCR sensitivity then may predominantly induce deletion (lower left panel). In contrast, strong agonist (right panels) may yield a higher probability of T_{reg} commitment when TCR sensitivity is lower but, as with a partial agonist, predominantly trigger deletion when TCR sensitivity increases. The net effect is an increase in T_{reg} numbers as we shift from partial to strong agonist.
(EPS)

Figure S3 Alternative models of T_{reg} selection for TCR clonotypes with varying sensitivity to endogenously expressed ovalbumin (OVA) peptide, as used in Lee _et al._ [14]. In the $p-$sum model (**upper panel**), each curve represents the distribution of summed signal strengths (Z_p) for TCRs with varying affinity for endogenously expressed OVA peptide. The probability that the summed signal from multiple pMHC contacts leads to T_{reg} commitment is higher for TCR with higher affinity for OVA, and this probability remains constant for each encounter throughout development. In the varying-TCR sensitivity model (**lower panel**), each curve represents the strength of signal derived from a single contact with OVA for each TCR as a function of time. There is a window of susceptibility in which contact with an OVA pMHC will lead to T_{reg} commitment; the probability of a contact with OVA is equal for all TCR, but the duration and timing of this window will vary for each TCR as a function of its initial ability to respond to OVA. (The blue line represents TCR with highest affinity for OVA; and orange represents TCR with the weakest affinity for OVA).
(EPS)

Table S1 Plausible combinations of parameters of the $p-$sum model. We used discrete combinations of n, the maximum number of APC encounters made by a thymocyte, and q, the fraction of endogenous peptides that are capable of inducing a TCR signal. Mean (minimum, maximum) values correspond to parameter combinations that described the data within ΔAIC<2 of the lowest AIC achieved for each (n,q) combination. p is the number of peptides contacted per APC

encounter; k_{TIM} reflects the signal derived from a single TCR contact with agonist TIM (as a percentile of signal strengths derived from contacts with non-null endogenous peptides); k_2 represents the minimal signal required for T_{reg} selection; k_3 represents the minimal signal required for negative selection (percentile of signal strengths received per encounter (with functional and null endogenous peptides)); and t_{max}, B and C parameterise the mapping function from relative TIM RNA to peptide abundance.
(PDF)

Text S1 The calculation of the selection probabilities in the $p-$sum model.
(PDF)

Text S2 The calculation of the selection probabilities in the two-phase model.
(PDF)

Text S3 Exploring the relation between relative TIM RNA expression and immune activation.
(PDF)

Text S4 Explaining variation in T_{reg} selection efficiency with TCR affinity.
(PDF)

Acknowledgments

Author Contributions

Conceived and designed the study: IB AJY. Contributed data: HMvS. Performed analysis and modelling of data: IB AJY. Wrote the paper: IB HMvS BS AJY.

References

1. Wing K, Sakaguchi S (2010) Regulatory T cells exert checks and balances on self tolerance and autoimmunity. Nat Immunol 11: 7–13.
2. Williams CB, Engle DL, Kersh GJ, Michael White J, Allen PM (1999) A kinetic threshold between negative and positive selection based on the longevity of the T cell receptor-ligand complex. J Exp Med 189: 1531–44.
3. Daniels MA, Teixeiro E, Gill J, Hausmann B, Roubaty D, et al. (2006) Thymic selection threshold defined by compartmentalization of Ras/MAPK signalling. Nature 444: 724–9.
4. Maloy KJ, Powrie F (2001) Regulatory T cells in the control of immune pathology. Nat Immunol 2: 816–22.
5. Hsieh CS, Liang Y, Tyznik AJ, Self SG, Liggitt D, et al. (2004) Recognition of the peripheral self by naturally arising CD25+ CD4+ T cell receptors. Immunity 21: 267–77.
6. Bettini ML, Vignali DAA (2010) Development of thymically derived natural regulatory T cells. Ann N Y Acad Sci 1183: 1–12.
7. Romagnoli P, van Meerwijk JPM (2010) Thymic selection and lineage commitment of CD4(+)Foxp3(+) regulatory T lymphocytes. Prog Mol Biol Transl Sci 92: 251–77.
8. Wong J, Obst R, Correia-NevesM, Losyev G,Mathis D, et al. (2007) Adaptation of TCR repertoires to self-peptides in regulatory and nonregulatory CD4+ T cells. J Immunol 178: 7032–41.
9. Jordan MS, Boesteanu A, Reed AJ, Petrone AL, Holenbeck AE, et al. (2001) Thymic selection of CD4+CD25+ regulatory T cells induced by an agonist self-peptide. Nat Immunol 2: 301–6.
10. Apostolou I, Sarukhan A, Klein L, von Boehmer H (2002) Origin of regulatory T cells with known specificity for antigen. Nat Immunol 3: 756–63.
11. Kawahata K, Misaki Y, Yamauchi M, Tsunekawa S, Setoguchi K, et al. (2002) Generation of CD4(+)CD25(+) regulatory T cells from autoreactive T cells simultaneously with their negative selection in the thymus and from nonautoreactive T cells by endogenous TCR expression. J Immunol 168: 4399–405.
12. Aschenbrenner K, D'Cruz LM, Vollmann EH, Hinterberger M, Emmerich J, et al. (2007) Selection of Foxp3+ regulatory T cells specific for self antigen expressed and presented by Aire+ medullary thymic epithelial cells. Nat Immunol 8: 351–8.
13. Larkin J 3rd, Rankin AL, Picca CC, Riley MP, Jenks SA, et al. (2008) CD4+CD25+ regulatory T cell repertoire formation shaped by differential presentation of peptides from a self-antigen. J Immunol 180: 2149–57.
14. Lee HM, Bautista JL, Scott-Browne J, Mohan JF, Hsieh CS (2012) A broad range of self-reactivity drives thymic regulatory T cell selection to limit responses to self. Immunity 37: 475–86.
15. Cozzo Picca C, Simons DM, Oh S, Aitken M, Perng OA, et al. (2011) CD4+CD25+Foxp3+ regulatory T cell formation requires more specific recognition of a self-peptide than thymocyte deletion. Proc Natl Acad Sci U S A 108: 14890–5.
16. van Santen HM, Benoist C, Mathis D (2004) Number of T reg cells that differentiate does not increase upon encounter of agonist ligand on thymic epithelial cells. J Exp Med 200: 1221–30.
17. Cozzo Picca C, Oh S, Panarey L, Aitken M, Basehoar A, et al. (2009) Thymocyte deletion can bias Treg formation toward low-abundance self-peptide. Eur J Immunol 39: 3301–6.
18. Atibalentja DF, Byersdorfer CA, Unanue ER (2009) Thymus-blood protein interactions are highly effective in negative selection and regulatory T cell induction. J Immunol 183: 7909–18.
19. Sebzda E,Wallace VA, Mayer J, Yeung RS, Mak TW, et al. (1994) Positive and negative thymocyte selection induced by different concentrations of a single peptide. Science 263: 1615–8.
20. Lerman MA, Larkin J 3rd, Cozzo C, Jordan MS, Caton AJ (2004) CD4+ CD25+ regulatory T cell repertoire formation in response to varying expression of a neo-self-antigen. J Immunol 173: 236–44.
21. Feuerer M, Jiang W, Holler PD, Satpathy A, Campbell C, et al. (2007) Enhanced thymic selection of FoxP3+ regulatory T cells in the NOD mouse model of autoimmune diabetes. Proc Natl Acad Sci U S A 104: 18181–6.
22. Bautista JL, Lio CWJ, Lathrop SK, Forbush K, Liang Y, et al. (2009) Intraclonal competition limits the fate determination of regulatory T cells in the thymus. Nat Immunol 10: 610–7.
23. Leung MWL, Shen S, Lafaille JJ (2009) TCR-dependent differentiation of thymic Foxp3+ cells is limited to small clonal sizes. J Exp Med 206: 2121–30.
24. Moran AE, Holzapfel KL, Xing Y, Cunningham NR, Maltzman JS, et al. (2011) T cell receptor signal strength in Treg and iNKT cell development demonstrated by a novel fluorescent reporter mouse. J Exp Med 208: 1279–89.
25. Eck SC, Zhu P, Pepper M, Bensinger SJ, Freedman BD, et al. (2006) Developmental alterations in thymocyte sensitivity are actively regulated by MHC class II expression in the thymic medulla. J Immunol 176: 2229–37.
26. Stephen TL, Tikhonova A, Riberdy JM, Laufer TM (2009) The activation threshold of CD4+ T cells is defined by TCR/peptide-MHC class II interactions in the thymic medulla. J Immunol 183: 5554–62.
27. Azzam HS, DeJarnette JB, Huang K, Emmons R, Park CS, et al. (2001) Fine tuning of TCR signaling by CD5. J Immunol 166: 5464–72.
28. Saini M, Sinclair C, Marshall D, Tolaini M, Sakaguchi S, et al. (2010) Regulation of Zap70 expressionduring thymocyte development enables temporal separation of CD4 and CD8 repertoire selection at different signaling thresholds. Sci Signal 3: ra23.
29. Neilson JR, Zheng GXY, Burge CB, Sharp PA (2007) Dynamic regulation of miRNA expression in ordered stages of cellular development. Genes Dev 21: 578–89.
30. Li QJ, Chau J, Ebert PJR, Sylvester G, Min H, et al. (2007) miR-181a is an intrinsic modulator of T cell sensitivity and selection. Cell 129: 147–61.
31. Davey GM, Schober SL, Endrizzi BT, Dutcher AK, Jameson SC, et al. (1998) Preselection thymocytes are more sensitive to T cell receptor stimulation than mature T cells. J Exp Med 188: 1867–74.
32. Chao DL, Davenport MP, Forrest S, Perelson AS (2005) The effects of thymic selection on the range of T cell cross-reactivity. Eur J Immunol 35: 3452–9.
33. Kosmrlj A, Jha AK, Huseby ES, Kardar M, Chakraborty AK (2008) How the thymus designs antigen-specific and self-tolerant T cell receptor sequences. Proc Natl Acad Sci U S A 105: 16671–6.
34. Detours V, Perelson AS (1999) Explaining high alloreactivity as a quantitative consequence of affinity-driven thymocyte selection. Proc Natl Acad Sci U S A 96: 5153–8.
35. Detours V, Mehr R, Perelson AS (2000) Deriving quantitative constraints on T cell selection from data on the mature T cell repertoire. J Immunol 164: 121–8.
36. Faro J, Velasco S, González-Fernández A, Bandeira A (2004) The impact of thymic antigen diversity on the size of the selected T cell repertoire. J Immunol 172: 2247–55.
37. Souza-e Silva H, Savino W, Feijóo RA, Vasconcelos ATR (2009) A cellular automata-based mathematical model for thymocyte development. PLoS One 4: e8233.
38. Kosmrlj A, Read EL, Qi Y, Allen TM, Altfeld M, et al. (2010) Effects of thymic selection of the T-cell repertoire on HLA class I-associated control of HIV infection. Nature 465: 350–4.
39. Borghans JAM, Noest AJ, De Boer RJ (2003) Thymic selection does not limit the individual MHC diversity. Eur J Immunol 33: 3353–8.
40. McKeithan TW (1995) Kinetic proofreading in T-cell receptor signal transduction. Proc Natl Acad Sci U S A 92: 5042–6.

41. Germain RN, Stefanová I (1999) The dynamics of T cell receptor signaling: complex orchestration and the key roles of tempo and cooperation. Annu Rev Immunol 17: 467–522.
42. Valitutti S, Müller S, Dessing M, Lanzavecchia A (1996) Different responses are elicited in cytotoxic T lymphocytes by different levels of T cell receptor occupancy. J Exp Med 183: 1917–21.
43. Rosette C, Werlen G, Daniels MA, Holman PO, Alam SM, et al. (2001) The impact of duration versus extent of TCR occupancy on T cell activation: a revision of the kinetic proofreading model. Immunity 15: 59–70.
44. Labrecque N, Whitfield LS, Obst R, Waltzinger C, Benoist C, et al. (2001) How much TCR does a T cell need? Immunity 15: 71–82.
45. Stone JD, Chervin AS, Kranz DM (2009) T-cell receptor binding affinities and kinetics: impact on T-cell activity and specificity. Immunology 126: 165–76.
46. Prasad A, Zikherman J, Das J, Roose JP, Weiss A, et al. (2009) Origin of the sharp boundary that discriminates positive and negative selection of thymocytes. Proc Natl Acad Sci U S A 106: 528–33.
47. Altan-Bonnet G, Germain RN (2005) Modeling T cell antigen discrimination based on feedback control of digital ERK responses. PLoS Biol 3: e356.
48. Currie J, Castro M, Lythe G, Palmer E, Molina-París C (2012) A stochastic T cell response criterion. J R Soc Interface 9: 2856–70.
49. Govern CC, Paczosa MK, Chakraborty AK, Huseby ES (2010) Fast on-rates allow short dwell time ligands to activate T cells. Proc Natl Acad Sci U S A 107: 8724–9.
50. Underhill DM, Bassetti M, Rudensky A, Aderem A (1999) Dynamic interactions of macrophages with T cells during antigen presentation. J Exp Med 190: 1909–14.
51. Friedl P, Gunzer M (2001) Interaction of T cells with APCs: the serial encounter model. Trends Immunol 22: 187–91.
52. Gunzer M, Schäfer A, Borgmann S, Grabbe S, Zänker KS, et al. (2000) Antigen presentation in extracellular matrix: interactions of T cells with dendritic cells are dynamic, short lived, and sequential. Immunity 13: 323–32.
53. Mempel TR, Henrickson SE, Von Andrian UH (2004) T-cell priming by dendritic cells in lymph nodes occurs in three distinct phases. Nature 427: 154–9.
54. Kawahata K, Misaki Y, Yamauchi M, Tsunekawa S, Setoguchi K, et al. (2002) Generation of CD4(+)CD25(+) regulatory T cells from autoreactive T cells simultaneously with their negative selection in the thymus and from nonautoreactive T cells by endogenous TCR expression. J Immunol 168: 4399–405.
55. Nadarajah S (2008) A review of results on sums of random variables. Acta Appl Math 103: 131–140.
56. Fisher RA, Tippet LHC (1928) Limiting forms of the frequency distribution of the largest or smallest member of a sample. Proc Cambridge Phil Soc 24: 180–190.
57. Gnedenko PB (1943) Sur la distribution limite du terme maximum d'une serie aleatoire. Annals of Mathematics 44: 423–453.
58. Obst R, van Santen HM, Mathis D, Benoist C (2005) Antigen persistence is required throughout the expansion phase of a CD4(+) T cell response. J Exp Med 201: 1555–65.
59. Le Borgne M, Ladi E, Dzhagalov I, Herzmark P, Liao YF, et al. (2009) The impact of negative selection on thymocyte migration in the medulla. Nat Immunol 10: 823–30.
60. Irvine DJ, Purbhoo MA, Krogsgaard M, Davis MM (2002) Direct observation of ligand recognition by T cells. Nature 419: 845–9.
61. Brower RC, England R, Takeshita T, Kozlowski S, Margulies DH, et al. (1994) Minimal requirements for peptide mediated activation of CD8+ CTL. Mol Immunol 31: 1285–93.
62. Naeher D, Daniels MA, Hausmann B, Guillaume P, Luescher I, et al. (2007) A constant affinity threshold for T cell tolerance. J Exp Med 204: 2553–9.
63. Hogquist KA, Jameson SC, Heath WR, Howard JL, Bevan MJ, et al. (1994) T cell receptor antagonist peptides induce positive selection. Cell 76: 17–27.
64. Liu CP, Crawford F, Marrack P, Kappler J (1998) T cell positive selection by a high density, low affinity ligand. Proc Natl Acad Sci U S A 95: 4522–6.
65. Nikolić-Zugić J, Bevan MJ (1990) Role of self-peptides in positively selecting the T-cell repertoire. Nature 344: 65–7.
66. Lio CWJ, Hsieh CS (2008) A two-step process for thymic regulatory T cell development. Immunity 28: 100–11.
67. Burchill MA, Yang J, Vang KB, Moon JJ, Chu HH, et al. (2008) Linked T cell receptor and cytokine signaling govern the development of the regulatory T cell repertoire. Immunity 28: 112–21.
68. Cheng G, Yu A, Dee MJ, Malek TR (2013) IL-2R signaling is essential for functional maturation of regulatory T cells during thymic development. J Immunol 190: 1567–75.
69. Nishizuka Y, Sakakura T (1969) Thymus and reproduction: sex-linked dysgenesia of the gonad after neonatal thymectomy in mice. Science 166: 753–5.
70. Asano M, Toda M, Sakaguchi N, Sakaguchi S (1996) Autoimmune disease as a consequence of developmental abnormality of a T cell subpopulation. J Exp Med 184: 387–96.
71. Fontenot JD, Dooley JL, Farr AG, Rudensky AY (2005) Developmental regulation of Foxp3 expression during ontogeny. J Exp Med 202: 901–6.
72. Lee HM, Hsieh CS (2009) Rare development of Foxp3+ thymocytes in the CD4+CD8+ subset. J Immunol 183: 2261–6.
73. Cowan JE, Parnell SM, Nakamura K, Caamano JH, Lane PJL, et al. (2013) The thymic medulla is required for Foxp3+ regulatory but not conventional CD4+ thymocyte development. J Exp Med 210: 675–81.
74. Nakayama T, June CH, Munitz TI, Sheard M, McCarthy SA, et al. (1990) Inhibition of T cell receptor expression and function in immature CD4+CD8+ cells by CD4. Science 249: 1558–61.
75. Lorenz RG, Allen PM (1989) Thymic cortical epithelial cells lack full capacity for antigen presentation. Nature 340: 557–9.
76. Mizuochi T, Kasai M, Kokuho T, Kakiuchi T, Hirokawa K (1992) Medullary but not cortical thymic epithelial cells present soluble antigens to helper T cells. J Exp Med 175: 1601–5.
77. Gray DHD, Seach N, Ueno T, Milton MK, Liston A, et al. (2006) Developmental kinetics, turnover, and stimulatory capacity of thymic epithelial cells. Blood 108: 3777–85.