Effects of Thymic Selection on T-Cell Recognition of Foreign and Tumor Antigenic Peptides

Jason T. George[a,b,d,1], **David A. Kessler**[e], **and Herbert Levine**[a,b,c,1]

[a]Center for Theoretical Biological Physics, [b]Department of Bioengineering, [c]Department of Physics and Astronomy, Rice University, 6100 Main Street, Houston, TX 77005; [d]Medical Scientist Training Program, Baylor College of Medicine, 1 Baylor Plaza, Houston, TX 77030; [e]Department of Physics, Bar-Ilan University, Ramat-Gan IL52900, Israel

This manuscript was compiled on July 25, 2017

The advent of cancer immunotherapy has generated renewed hope for the treatment of many malignancies by introducing a number of novel strategies that exploit various properties of the immune system. These therapies are based on the idea that CD8+ T-lymphocytes (CTLs) directly recognize and respond to tumor-associated neoantigens (TANs) in much the same way as they would to foreign peptides presented on cell surfaces. To date, however, nearly all attempts to optimize immunotherapeutic strategies have been empirical. Here, we develop a model of T-cell selection based on the assumption of random interaction strengths between a self-peptide and the various CTLs. The model enables the analytical study of the effects of selection on the CTL recognition of TANs and of completely foreign peptides and can estimate the number of CTL's that can detect donor-matched transplants. We show that negative selection thresholds chosen to reflect experimentally observed thymic survival rates result in near-optimal production of T-cells that are capable of surviving selection and recognizing foreign antigen. These analytical results are confirmed by simulation.

Cancer Immunotherapy | Immunology | Applied Probability

Immunotherapeutic strategies which can, in principle, evolve with a growing malignancy have gained recent popularity for treating a variety of cancer types (1, 2). Successful therapy depends on cytotoxic CD8+ T-lymphocyte (CTL) recognition of tumor-associated neoantigens (TANs), as well as non-mutated but misexpressed self-peptides to which T-cells are intolerant. These antigens are displayed on the surface of cells via major histocompatibility complexes (MHCs) (3–8).

MHC-displayed TANs arising by somatic point mutations, along with over-expressed or mislocalized self-peptides present in only minimal quantities during thymic negative selection, carry antigenic potential (9–12). TANs resemble self-peptides while over-expressed self-peptides may appear to CTLs as entirely foreign antigens which have never been selected against. Previous quantitative models of T-cell receptor (TCR)-peptide interactions have succeeded in accurately characterizing many aspects of T-cell immunology and thymic selection. Digit string representations have provided insight into foreign peptide recognition and MHC-unmatched alloreactivity rates, as well as T-cell specificity and cross-reactivity(13–17). Modeling TCRs and peptides as amino acid strings has also been successful in studying HIV (18–21). In the context of immunotherapy, characterizing recognition rates of TAN and over-expressed/mislocalized self-peptides is relevant to understanding CTL repertoire targeting efficiency. This recognition process has yet to be mathematically modeled; instead, nearly all attempts to understand CTL-based immunology have been empirical. Clearly, developing such a model is basic to achieving a fuller understanding of the overall functioning of the immune system.

Here we began our analysis with a preexisting model, adapted to isolate negative selection effects. We found anomalous behavior in the ability of this formulation to apportion the contributions of individual self-peptides to the overall selection process. In particular, a small number of 'potent' thymic self-peptides dominated selection. We then formulated an alternative, more general approach for the T-cell receptor (TCR)-peptide interaction that assigns random amino acid binding interactions in a position-independent manner, as opposed to the fixed set of interactions between different amino acid pairs in the previous model. Despite exhibiting improvement, the problems discussed above still remained, rendering this model also unacceptable. This led us to consider a final approach that incorporates position dependence into the random interaction picture. This last model exhibits a realistic selection balance between individual self-peptides and allowed us to then consider issues of detection of altered peptides.

Using this model, we found that antigenic proximity to self-peptide only minimally reduced CTL recognition when compared to recognition of foreign antigen. Moreover, we showed that TCR activation thresholds consistent with empirical selection rates resulted in an near-optimal production of T-cells that both survive negative selection and identify foreign antigen. Lastly, we applied our model to the setting of transplantation, predicting alloreactivity (i.e. detection by host CTLs) rates consistent with empirical observations, given known levels of host and donor single nucleotide dissimilarity.

Significance Statement

We have developed a model of T-cell binding that accurately represents the influence of self-peptides on thymic negative selection. From this, we generated estimates for relevant antigen recognition rates. We found that negative selection only slightly interferes with a T-cell's ability to detect antigens that differ from self-peptide by a single amino acid, and that these peptides may effectively be regarded as foreign. Moreover, negative selection thresholds chosen to reflect experimentally observed thymic survival rates result in optimal production of T-cells that are capable of surviving selection and recognizing foreign antigen. Lastly, our model predicts empirically reasonable donor tissue recognition rates in the context of an HLA-matched transplant.

Author Contributions: J.T.G., D.A.K., and H.L. designed research, performed research, contributed new analytic equations, analyzed the data, and wrote the paper.

The authors declare no conflict of interest.

[1]To whom correspondence may be addressed. E-mail: Herbert.Levine@rice.edu or Jason.George@rice.edu

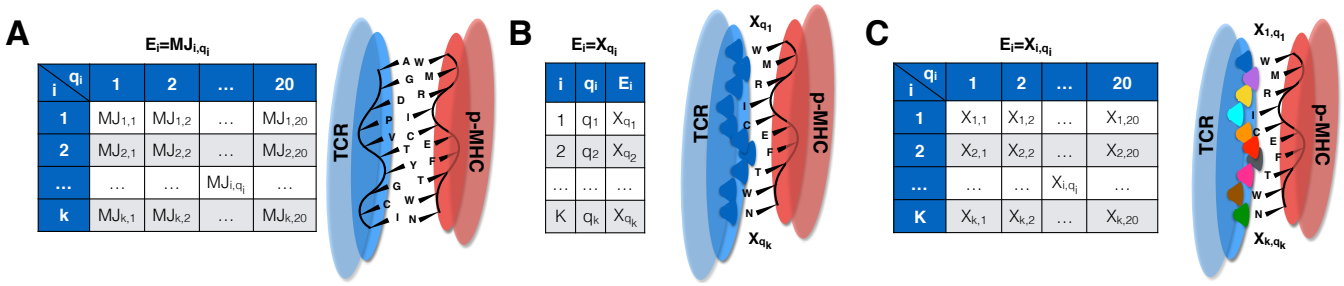

Fig. 1. Alternative TCR-MHC Interactions Formulations: (A) Sequential MJ (S-MJ) - Each TCR is represented by a string of amino acids and the binding energy with a self-peptide is the sum of pairwise binding energies between the TCR amino-acids and those in corresponding positions along the peptides; (B) Position Independent Random Affinity (PIRA) - TCR regions (shaded dark blue in the cartoon) that bind peptide are characterized by TCR-specific binding energies for each given amino acid type, all drawn from a standard Gaussian distribution. These binding energies are assumed to be position independent; (C) Random Interaction between Cell receptor and Epitope (RICE) - The behavior of TCR contact regions is now position-dependent (hence represented by different colors in the cartoon). A TCR is represented by an 20-by-10 array of IID standard Gaussian random variables indicating the binding energy contribution for each amino acid/contact position pair along the peptide.

Model Development

We seek a model of negative thymic selection where the representation of self-peptides is defined explicitly by amino acid sequences and the survivors constitute a representative CTL repertoire. Most importantly, the model should exhibit reasonable thymic behavior. That is, we require that a plurality of the thymic self-peptides, estimated to number around 10^4 (16, 22) for each MHC class, non-trivially participate in negative selection. To this end, we start with a version of an existing model (18–21) built upon the use of an amino acid binding matrix. We found that this model does not yield satisfactory behavior, and so considered in turn two alternative formulations for the TCR-MHC interaction. We studied the nature of thymic selection in each framework, employing both analysis and simulation. In each of these models, peptide-bound MHC (p-MHC) is represented by a sequence, $\{q_i\}_{i=1}^k$, of k amino acids; for definiteness, we will choose k=10. In order to facilitate analytical insight, we consider a single MHC type (in reality there may be up to six) for each individual. Such a framework will allow us to determine the extent to which negative selection diminishes mutated self-peptide recognition. In the following we do not focus on the precise physical mechanism by which TCRs recognize antigen (for example, affinity-driven versus binding lifetime) (13, 23, 24), merely positing an interaction strength governing recognition; for simplicity, we will use throughout the language of binding energy.

Sequential MJ (S-MJ). Our first formulation is an extension of the work of Chakraborty and colleagues, which has been successful in providing insight into HIV-immune dynamics (18–21). There, a TCR t is represented by a complementary string $t = \{t_i\}_{i=1}^k$ of k amino acids that contact p-MHC, and the total binding interaction is the sum of a direct TCR-MHC interaction E_c, and pairwise amino acid binding interactions using the 20×20 Miyazawa-Jernigan (MJ) matrix (25, 26) (Fig. 1A). The interaction between TCR t and p-MHC q was then given by:

$$E(t,q) = E_c + \sum_{i=1}^{k} MJ_{t_i,q_i}, \qquad [1]$$

where MJ_{t_i,q_i} represents the MJ interaction between amino acids t_i and q_i. Since, E_c plays no role when only one MHC class is present, we henceforth set it equal to 0.

Our extension, called the sequential MJ (S-MJ) model, allows us to consider independently and sequentially the effects of positive and negative thymic selection on naïve T-cell generation. The role of positive selection is to filter defective TCRs unable to properly interface with p-MHC (27, 28); this happens separately (both in time and space) from negative selection, which is our sole concern here (see SI for more discussion of positive selection).

As we shall see, our analysis shows that this first model does not yield satisfactory behavior and we are thus obligated to modify the model. In our alternate formulations, self-peptide sequences are represented in the same way as above, but we change the form of the interaction. In reality, the binding site for each amino acid on p-MHC is complex and binding interactions represent the net effects of complicated TCR-p-MHC association in a binding groove. This affords TCRs with a large degree of freedom in their ability to interface with each amino acid in a given p-MHC. We therefore assume that individual amino acid interactions that comprise TCR-p-MHC interface are random variables, which we take to be Gaussian distributed. To form a tractable model, we also assume that the binding energies attributed to each possible amino acid are independent and identically distributed (IID). This type of approach is reminiscent of the random energy model which was used to great effect in studies of protein biophysics (29).

Position-Independent Random Affinity (PIRA). In this first alternative model, we assume that TCR interactions with each amino acid are position-independent (Fig. 1B). That is, a given amino acid interacts with a given TCR identically regardless of its position in the p-MHC. We refer to this as Position-Independent Random Affinity (PIRA). In this case, a TCR t may be represented by its interactions with all $|A| = 20$ amino acids, and so are described by a sequence $\{X_\alpha^t\}_{\alpha=1}^A$ of independent standard Gaussian random variables. The interaction function for negative selection is given by:

$$E(t,q) = \sum_{i=1}^{k} X_{q_i}^t. \qquad [2]$$

Random Interaction Between Cell Receptor and Epitope (RICE). More realistically, each 3-dimensional location in a given TCR handles amino acids very differently. For now, we neglect any correlation in binding interactions that might occur either due to amino acids with similar properties or through

a dependence on adjacent amino acids. We characterize the binding energy of the peptide at the i^{th} position of p-MHC as a function of i itself in addition to the identity of the amino acid at this position (Fig. 1C). We refer to this approach as the Random Interaction between Cell receptor and Epitope (RICE) formulation. Here, we represent a TCR, t, and all of its possible interactions by an $k \times A$ array $\{X^t_{i,\alpha}\}$ of IID standard Gaussian random variables, where $X_{i,\alpha}$ denotes the interaction with which TCR t binds amino acid α located at position i. The interaction function for negative selection is then given by:

$$E(t,q) = \sum_{i=1}^{k} X^t_{i,q_i}. \qquad [3]$$

In the SI we analyze the alternate choice of IID uniform distributions and verify that the important results are independent of this level of detail.

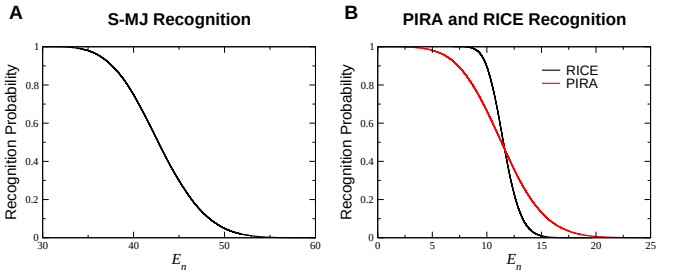

Fig. 2. TCR recognition probability for each model. Recognition occurs whenever the interaction energy is greater than an upper threshold, E_n. Recognition rates as a function of threshold for (A) S-MJ, (B) PIRA and RICE. General recognition behavior is similar among all three formulations. In each case 10^5 thymocytes are simulated to undergo selection.

Results

To study negative selection in these models, a given randomly constructed TCR was tested against a collection of randomly constructed N_n self-peptides. For a given TCR t to survive selection, the interaction energy between t and each and every self-peptide q must not exceed E_n. Conversely, TCR t recognizes (non-self) peptide q whenever the TCR-p-MHC interaction exceeds E_n. Potency and recognition simulations utilize cohorts of 10^5 TCRs.

In the following we present the main findings of our analysis. Full mathematical derivations are provided in the SI.

RICE Yields a Sensible Spectrum of Self-Peptide Contributions to Thymocyte Selection. All three formulations described above exhibit similar empirical survival profiles with respect to the negative selection threshold (Fig. 2). However, in both S-MJ and PIRA a very small number (one, in a typical simulation of S-MJ; 125 in PIRA, both with $N_n = 10^4$) of 'potent' self-peptides dominates nearly all of thymocyte selection, as seen both in simulation and analytically (Figs. 3, S7, S10). This feature does not depend on the assumed size of the training set N_n (see SI). One manifestation of a few peptides dominating the entire selection process is a much higher fluctuation in mean survival rates (Fig. S14). More generally, we do not think that it makes sense for a system

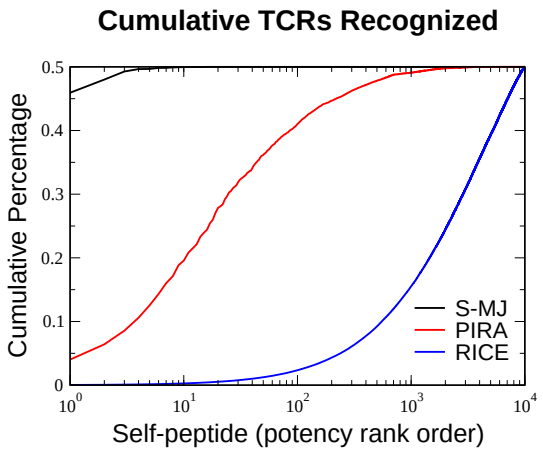

Fig. 3. Peptide potency for each model. The 10^4 self-peptides were ordered by 'potency', or the fraction of (the 10^5) thymocytes recognizing them during selection simulations. 'Potent' self-peptides were those that were recognized most often by the TCRs. The cumulative contributions of each self-peptide to negative selection was plotted in decreasing order of self-peptide potency for the S-MJ, PIRA and RICE models. In all cases, the selection thresholds are chosen in order to give 50% survival. We see that for the S-MJ model, the most potent self-peptide is responsible for roughly 90% of the selection behavior, whereas for the PIRA (resp. RICE) model, 200 (resp. 7100) out of 10^4 self-peptides generate this level of selection.

to employ a large number of peptides for negative selection if only a very small fraction would achieve the same outcome.

In contrast, RICE results in a repertoire sculpted by all the self-peptides (Fig. 3). The extreme potency of some self-peptides in the S-MJ model is a result of the presence in these self-peptides of many amino acids with anomalously large average binding energy, which cause them to be recognized by a large number of TCRs. In PIRA, the potent peptides are those which have a large number of amino acid repeats, so that they are recognized by all TCRs who have a significant binding energy to those particular amino acids (see SI for details). Since there are twenty different amino acids, it takes of order twenty self-peptides to accomplish the bulk of the negative selection (specifically, twenty-five peptides accomplish 60% of the selection). Neither of these issues occurs in RICE.

Given these findings, we have chosen to proceed with the RICE model. Current empirical observations of negative selection vary anywhere from 30% to as high as 90% (30–35). This defines an acceptable range, $11 \leq E_n \leq 13$, for reasonable negative selection thresholds to be used in our analysis (Fig. 2B).

Thymic Selection Minimally Decreases Recognition of Point–Mutated Self-Peptide. By construction, the binding interactions between TCR and self-peptide are sums of k IID random variables. A given TCR t survives negative selection against a collection $\{q^{(j)}\}_{j=1}^{N_n}$ of N_n thymic self-peptides if all of its binding energies are blow threshold. Under RICE, the survival probability p_s may be approximated by neglecting similarities between self-peptides and noting that the amino acids comprising a self-peptide are IID.

$$p_s \approx \mathbb{P}\Big(E(X^t, q) \leq E_n\Big)^{N_n}. \qquad [4]$$

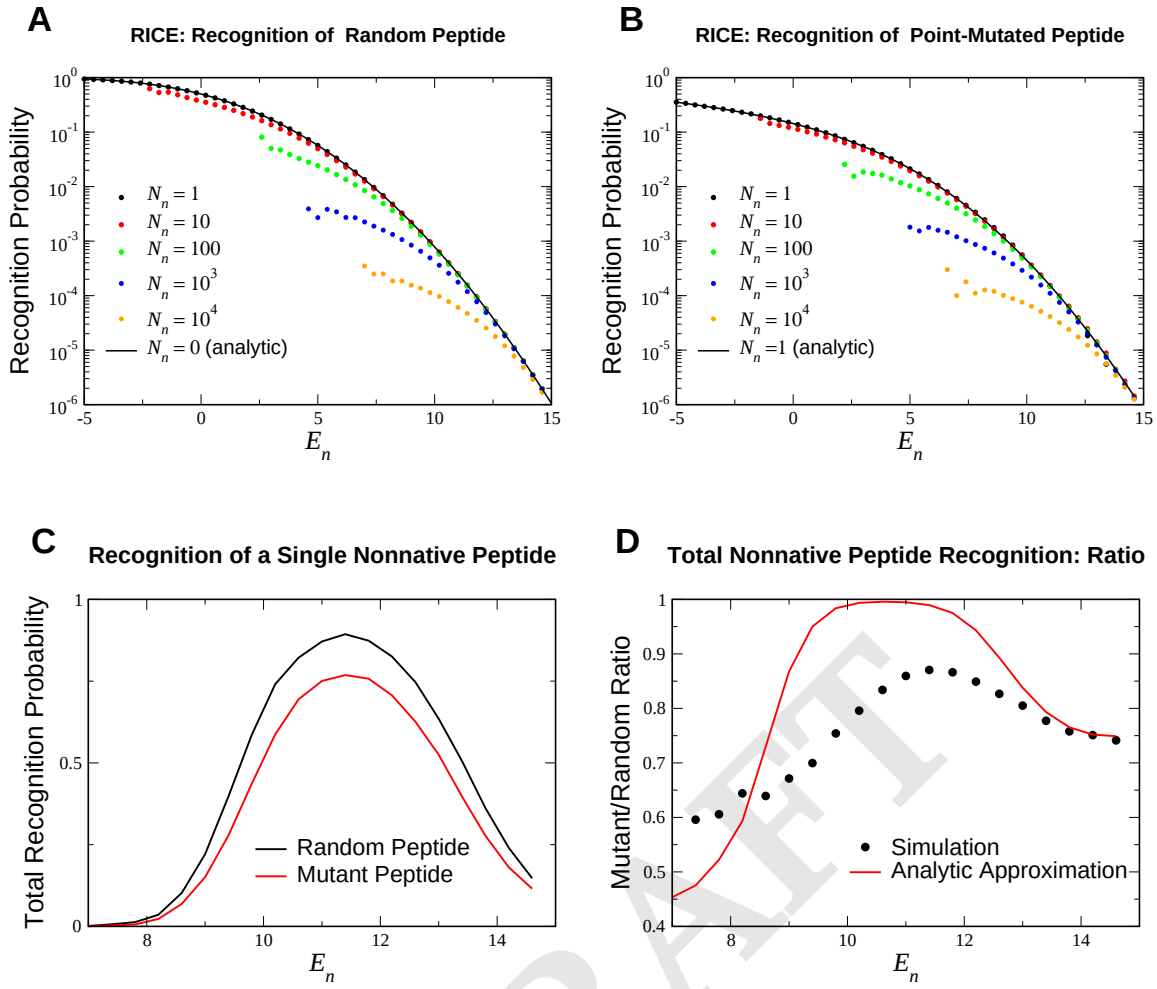

Fig. 4. RICE Recognition Behavior. (A-B) Probabilities for a single selected TCR to recognize (A) random peptides and (B) point-mutants of self-peptides for the analytically tractable limits and for higher values of N_n. In both cases, the effect of increasing the number of negatively selecting self-peptides to N_n=10^4 has a relatively small effect on recognition rates in the range of relevant values (11 to 13) of E_n. The simulation averaged over all the surviving TCRs from the initial cohort of 10^5. 10^4 random and point-mutant variants were tested. (C) The total recognition probability for the surviving TCR cohort to recognize: a random peptide (black curve) and a single-site mutant of a native peptide (red). (D) The ratio of the two recognition probabilities in (C). 10^4 peptides of each class were tested. Included in (D) is the theoretical estimate of the ratio, $(1-(1-\tilde{p}_1)^{N_s})/(1-(1-\hat{p}_0)^{N_s})$, where $N_s(E_n)$ is the number of TCRs that survived selection.

The quantity $E(X^t, q)$ is a sum of k IID Gaussian random variables, and its distribution is given by

$$F_k(x) = \Phi\left(\frac{x}{\sigma}\right). \quad [5]$$

where $\Phi(\cdot)$ is the CDF of the standard zero-mean, unit variance normal distribution and $\sigma^2 = k$. The survival probability is then the probability that the maximum of $E(X^t, q)$ over the set of amino acids q is less than E_n. The distribution of the maximum of a large number of Gaussian random deviates can be approximated by a Gumbel distribution with CDF (see SI for details):

$$p_s(E_n) \approx \exp\left[-e^{-\frac{E_n-\mu}{W}}\right] \quad [6]$$

where

$$\mu(N_n) = \sigma\sqrt{2\ln N_n/N_0(N_n)}$$
$$W(N_n) = \frac{\sigma}{\sqrt{2\ln N_n/N_0(N_n)}} \quad [7]$$

are the N_n-dependent mode and width parameters of the Gumbel distribution and $N_0(N_n)$ is a parameter that is found by solving the implicit equation

$$N_0^2(N_n) = 4\pi \ln N_n/N_0. \quad [8]$$

This approximation for p_s, as well as more detailed analytic approaches which also include the role of the variance of the mean energy for a given TCR due to the finite number of amino acids, are compared to direct simulations of negative selection in the SI.

We now wish to compare two cases of single TCR recognition probabilities, one of a random peptide and one of a point-mutated self-peptide. For the random peptide, we denote this probability by $\hat{p}$. A simple analytical estimate is given by $\hat{p} \approx \hat{p}_0$, the value for the recognition probability obtained by ignoring selection completely as there are not likely to be any self-peptides close to one chosen completely

at random:

$$\hat{p}_0 = \mathbb{P}\Big(E(X,q) \geq E_n\Big) = 1 - F_k(E_n). \quad [9]$$

The '0' subscript here denotes that no selection takes place (or equivalently $N_n = 0$ peptides) for this event. The analytic expressions for $\hat{p}_0$ is compared to simulations for N_n from 1 to 10^4 (Fig. 4A). Note that in the selection threshold range of interest ($E_n \geq 11$), the agreement is semi-quantitative even for $N_n = 10^4$.

One can similarly construct an analytic estimate for the single TCR recognition probability for a point mutated self-peptide, which we will refer to as a tumor neoantigen (TAN) assuming that a tumor may be detected by its creation of singly mutated peptides. We will denote this TAN as $\tilde{q}$ which is a mutated version of self-peptide q, differing only at position i^*. We let $\tilde{p}$ be the probability that TCR t negatively trained on N_n self-peptides would nonetheless recognize TAN $\tilde{q}$. We estimate $\tilde{p}$ by evaluating the probability that TCR t trained exclusively on the single (non-mutated) self-peptide q can detect $\tilde{q}$. We denote this probability by $\tilde{p}_1$ to indicate that it survived selection against a single self-peptide. This is motivated by the fact that q is most closely related to $\tilde{q}$, and therefore should account for a significant amount of the dependency of TCR t recognition ability on t's survival under full thymic selection. The probability, $\tilde{p}_1$, that TCR t recognizes $\tilde{q}$, conditioned on surviving thymic selection by q, is given by:

$$\tilde{p}_1 = \mathbb{P}\Bigg(\sum_{i=1}^{n} X_{i,\tilde{q}_i} > E_n \;\Bigg|\; \sum_{i=1}^{n} X_{i,q_i} \leq E_n\Bigg)$$

This probability is computed in the SI; the result is approximately

$$\tilde{p}_1 = \frac{e^{-E_n^2/2(k-1)}}{\pi\sqrt{2k}\,\Phi\Big(\frac{E_n}{\sqrt{k-1}}\Big)} \quad [10]$$

This expression for $\tilde{p}_1$ estimating the probability of TAN recognition is compared to simulations with N_n ranging from 1 to 10^4 (Fig. 4B). We find that recognition estimates are again reasonably accurate even for large N_n in the regime of realistic negative selection thresholds ($E_n \geq 11$), despite the potential influence of many negative selectors. This suggests that cross-dependencies between self-peptides due to sharing amino acids in the same position have a weak effect on the overall selection of a repertoire. Moreover, reductions in the ability of a repertoire to detect closely-related (TAN) peptides are quite modest when compared to foreign or non-mutated self-peptides (Fig. 4C). The ratio ranges from 0.6 to 0.9 in the range $7 < E_n < 15$. Also shown in the figure is the theoretical estimate deriving from $\tilde{p}_1$ and $\hat{p}_0$, which shows the same trends, with slightly larger variation with E_n. These findings support the hypothesis that TCR selection against self-peptides has a minimal influence on the recognition of peptides which are 'close' to self and that these peptides are detected with rates similar to those of completely random (foreign) antigens. In the RICE model, the immune system appears to simply memorize the list of self-antigens and by doing so generates a surprising level of immune protection against peptides not included on that precise list.

Observed Thymic Selection is Close to Optimality for TCR Recognition Ability. The above analysis provides a convenient context for framing negative selection as an optimization problem. Aside from maximally producing thymocytes, the immune system could be attempting to choose the TCR interaction threshold (E_n) in such a way as to encourage recognition of foreign antigen. In other words, the host benefits from producing TCRs that have the ability to survive negative selection *and* subsequently recognize random (currently unknown) foreign antigens. We again approximate the detection probability by the probability of recognition by TCRs undergoing no selection ($\hat{p}_0$), and obtain

$$\begin{aligned}\hat{p}p_s &\approx \hat{p}_0 p_s \\ &\approx \big[1 - F_k(E_n)\big] \cdot \big[F_k(E_n)\big]^{N_n}. \end{aligned} \quad [11]$$

Then, existence of an extremum, E_n^*, requires

$$\frac{d(\hat{p}_0 p_s)}{dE_n} = F_k' F_k^{N_n - 1}\big[N_n - (N_n+1)F_k\big] = 0. \quad [12]$$

Thus, at the optimal threshold,

$$F_k(E_n^*) = \frac{N_n}{N_n + 1}. \quad [13]$$

The rate of optimal negative selection for large numbers of peptides ($N_n \gg 1$) is characterized by

$$p_s(E_n^*) = \Big(\frac{N_n}{N_n+1}\Big)^{N_n} \approx 1/e. \quad [14]$$

This value, consistent with the (low end) of measured survival probabilities, corresponds to a system optimized for recognition and agrees with an independent analysis that considered the optimal diversity of the T and B cell repertoires (16). We should note of course that we expect there to be slight differences in the optimum threshold for $\tilde{p}p_s$ case as opposed to our estimate based on $\hat{p}_0 p_s$. We do expect these to be close in general and in fact can prove that the true maximum is always no less than E_n^* (Proposition S1) in selection regimes of interest.

Effects of Host-Donor Sequence Differences on Alloreactivity Percentages. Let us now turn to the setting of transplants with MHC-matched host and donor. Even with the matching, there will be some single nucleotide polymorphisms (SNPs) between host and donor. These SNPS may give rise to amino acid differences that are detectable by the host T-cells. Let us denote by Y the number of such differences. Since $\tilde{p}$ is the probability of a given TCR detecting a peptide with a single amino-acid difference, the probability P_A of a host TCR recognizing a difference in the donor tissue (termed alloreactivity) is given by

$$P_A = 1 - \big(1 - \tilde{p}\big)^Y. \quad [15]$$

We now need to calculate the distribution of Y. We consider MHC matched donor and host and use the frequency of SNPs in the genome (≈300/bp) (36) to estimate the number of 10-mer peptides which contain a SNP. Since each peptide comes from 30bp of sequence, the chance that this sequence will contain a mutation is $30/300 = 0.1$. Assuming that all donor peptides being probed by the immune system are contained in the size N_n training set, the number of peptides that exhibit

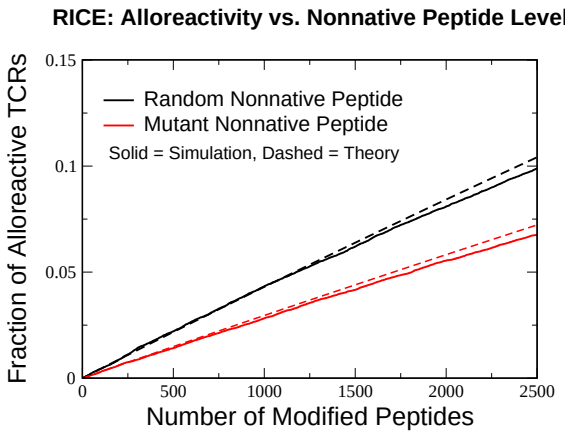

Fig. 5. The Effects of Increasing Differences in Host and Donor Thymic Self-Peptides on Alloreactivity percentages in the RICE model. The simulation was performed with an original cohort of 10^5 TCRs, of which 50% survived selection under 10^4 self-peptides ($E_n = 11.52$). Increasing numbers of nonnative peptides (either random (black curves) or single-difference mutants (red)) were introduced and the numbers of TCRs reacting to these were recorded. The theoretical estimate is from Eq. (16) with single TCR recognition probabilities for random and single-mutant peptide taken from Fig. 4A,B

differences from the training set is distributed according to $Z{\sim}\text{Poisson}(\lambda{=}N_n/10)$. As alluded to above, a point mutation may or may not actually manifest as an amino acid difference between host and donor. We assume approximately equal frequencies of DNA base pairs, and calculate the probability of an amino acid difference given a SNP as $p_d \approx 0.6$ by considering the likelihood of missense mutations arising from DNA codons. Thus, Y, the number of self-peptides that actually manifest a different amino acid is distributed as $[Y|Z{=}z]{\sim}\text{Binomial}(z, p_d)$. By use of the probability generating functions of Z, it can be shown that $Y{\sim}\text{Poisson}(\lambda p_d)$ (see SI). From this, we may obtain the first and second moments of P_A, i.e., the mean and variance of the fraction of TCRs exhibiting alloreactivity:

$$\mathbb{E}\big[P_A\big] = 1 - e^{-\lambda p_d \hat{p}}. \quad [16]$$

$$\mathbb{E}\big[P_A^2\big] = 1 - 2e^{-\lambda p_d \tilde{p}} + e^{-\lambda p_d \tilde{p}(2-\tilde{p})}. \quad [17]$$

$$\text{Var}(P_A) = \mathbb{E}\big[P_A^2\big] - \mathbb{E}\big[P_A\big]^2 = e^{-2\lambda p_d \tilde{p}}\big(e^{\lambda p_d \tilde{p}^2} - 1\big). \quad [18]$$

The percentage of alloreactive TCRs given by above equations, in other words the allogeneic CTL response, equals $2.02{\pm}0.08\%$ in MHC-matched host and donor pairs due to SNPs (Fig. 5). This response is obtained from contributions of roughly 600 potential allogeneic p-MHC. This number is on the low end of experimentally observed estimates of MHC-unmatched alloreactivity falling between 1% to 24% (13). We note in passing that the case of maximal single amino acid sequence differences in our model with $N_n = 10^4$ would correspond to an alloreactive rate of 26%, while maximal numbers of random peptides would correspond to rates as high as 38%; see Fig. S12.

Discussion

The development of a generative model relating T-cell repertoires to thymic selection against individual self-peptides represents an important theory-driven milestone to better understand CTL cancer immunotherapy. Here, we were primarily interested in studying the influence of thymic negative selection on CTL repertoire recognition of relevant non-self-peptides, with applications to TAN recognition and SNP recognition by MHC-matched CTLs. It was, therefore, important that the analysis be sensitive to small differences in individual thymic self-peptides that sculpt T-cell repertoires. This, in turn, required that the model appropriately capture the behavior of thymic negative selection on an individual self-peptide level.

We started by comparing an adaptation of the previously proposed MJ discrete model of thymic selection, focusing on negative selection effects. We discovered that this model does not behave in a statistically reasonable manner. Specifically, single peptides can have inordinate consequences and there is concomitant fluctuations in the selection behavior; these are the result of correlations in the MJ matrix. Instead of trying to modify the form of a peptide-peptide matrix, perhaps following the shape space ideas of (13), we introduced a more general perspective on how the t-cell sequence creates binding pockets for the p-MHC. This then allowed us to formulate two alternative models, PIRA and RICE, and found that wide-spread participation by thymic self-peptides action in T-cell selection was observed only in the latter alternative, which supposed a position-dependent character of TCR-p-MHC interaction.

Using RICE, we analytically characterized events of relevance to the problem of immune action including T-cell survival during negative selection, SNP detection, and non-self peptide recognition probabilities. We observed that TCR negative selection by host peptides has only a weak suppressive effect on detecting peptides which closely resemble self. This finding suggests that self-education during central tolerance in the thymus is a strategy that seeks to memorize as many of the self-peptides commonly found in the periphery (Fig. 1B) as possible, as opposed to selection by a few self-peptides capable of mitigating autoimmunity, and is a testament to the level of specificity exhibited by TCRs. Using the RICE model, we showed that parameter selection which generates realistic survival percentages also results in an near-optimal generation of thymocytes best suited to survive selection and most effectively identify foreign peptides. Finally, the model produced realistic characterizations of alloreactivity when applied to the setting of MHC-matched individuals. A potential advantage of an immune system designed to follow the RICE model over MJ is that the latter presents only a static challenge that foreign peptides must undergo in order to evade detection; evasion of this detection which might then be evolutionarily selected by pathogens. In contrast, there is no a priori strategy assumed within the RICE model with its energy landscape that varies randomly from TCR to TCR.

We cannot expect a simple hypothesis such as RICE to fully capture every detail of actual TCR-pMHC binding. In the absence of a quantitatively reliable molecular biophysics approach, we have chosen to work backwards and illustrate the type of statistical model that makes functional sense and that allows for new questions (such as the penalty imposed by negative selection on tumor neoantigen detection) to be addressed. One criticism of RICE might be that it does not allow for similar peptide recognition by very similar TCR's or conversely for TCR activation by very similar peptides.

There are two reasons why this does not immediately concern us. First, the coverage of the entire possible space of TCR sequences by actual TCR clones is so sparse that the chance of getting two TCRs with (nearly) identical chemical sequences should be very small. (Note; if there are a few "public" clones which are specifically programmed into the TCR formation rules, these would presumably also be programmed to automatically survive negative selection; our considerations apply to all the others). Second, the results of the RICE model are not significantly changed if we first group together chemically similar amino acids (37, 38) (i.e. use a reduced amino acid alphabet) and then proceed with the repertoire construction (see SI).

The overall objective of optimizing CTL therapy is complex, and may require future analysis that incorporates additional relevant aspects of acquired immunity and T-cell tolerance. Understanding this complex process holds the promise of one day optimizing and extending CTL immunotherapy to additional therapeutic contexts.

ACKNOWLEDGMENTS. We thank Philip A. Ernst for critical reading of the manuscript, and Haven R. Garber and Jeffrey J. Molldrem for helpful discussions. J.T.G. is supported by the National Cancer Institute of the National Institutes of Health (F30CA213878). D.A.K. is supported by the United States-Israel Binational Science Foundation (Grant 2015619). H.L. is supported by Cancer Prevention and Research Institute of Texas Scholars program (R1111).

1. Couzin-Frankel J (2013) Breakthrough of the year 2013. Cancer immunotherapy. *Science* 342(6165):1432–1433.
2. McGranahan N, et al. (2016) Immune Checkpoint Blockade. *Science* 351(November 2015):1463–1469.
3. Robinson J, Soormally AR, Hayhurst JD, Marsh SGE (2016) The IPD-IMGT/HLA Database - New developments in reporting HLA variation. *Human Immunology* 77(3):233–237.
4. Ding L, et al. (2012) Clonal evolution in relapsed acute myeloid leukaemia revealed by whole-genome sequencing. *Nature* 481(7382):506–510.
5. Gubin MM, Artyomov MN, Mardis ER, Schreiber RD (2015) Tumor neoantigens: building a framework for personalized cancer immunotherapy. 125(9):3413–3421.
6. Verdegaal EME, et al. (2016) Neoantigen landscape dynamics during human melanoma-T cell interactions. *Nature* 536(7614):91–5.
7. Abbas AK, Lichtman AH, Shiv P (2015) *Cellular and Molecular Immunology*. (Elsevier Saunders, Philadelphia, PA), 8 edition.
8. Schumacher TN, Schreiber RD (2015) Neoantigens in cancer immunotherapy. *Science* 348(6230):69–74.
9. Molldrem JJ, Komanduri K, Wieder E (2002) Overexpressed differentiation antigens as targets of graft-versus-leukemia reactions. *Current opinion in hematology* 9(6):503–508.
10. Cai A, et al. (2012) Mutated BCR-ABL generates immunogenic T-cell epitopes in CML patients. *Clinical Cancer Research* 18(20):5761–5772.
11. Restifo N, Dudley M, Rosenberg SA (2012) Adoptive immunotherapy for cancer: harnessing the T cell response. *Nature Reviews Immunology* 12(April):269–281.
12. Strønen E, et al. (2016) Targeting of cancer neoantigens with donor-derived T cell receptor repertoires. *Science* 352(June):1337–41.
13. Detours V, Perelson AS (1999) Explaining high alloreactivity as a quantitative consequence of affinity-driven thymocyte selection. *Proceedings of the National Academy of Sciences of the United States of America* 96(9):5153–8.
14. Detours V, Mehr R, Perelson AS (2000) Deriving Quantitative Constraints on T Cell Selection from Data on the Mature T Cell Repertoire. *Journal of Immunology* 164:121–128.
15. Chao DL, Davenport MP, Forrest S, Perelson AS (2005) The effects of thymic selection on the range of T cell cross-reactivity. *European Journal of Immunology* 35:3452–3459.
16. De Boer RJ, Perelson AS (1993) How diverse should the immune system be? *Proceedings: Biological Sciences* 252(1335):171–175.
17. Frankild S, De Boer RJ, Lund O, Nielsen M, Kesmir C (2008) Amino Acid Similarity Accounts for T Cell Cross- Reactivity and for "Holes" in the T Cell Repertoire. *PloS one* 3(3):1–8.
18. Kosmrlj A, Jha AK, Huseby ES, Kardar M, Chakraborty AK (2008) How the thymus designs antigen-specific and self-tolerant T cell receptor sequences. *Proceedings of the National Academy of Sciences of the United States of America* 105(43):16671–16676.
19. Chakraborty AK, Kosmrlj A (2010) Statistical mechanical concepts in immunology. *Annual review of physical chemistry* 61:283–303.
20. Kosmrlj A, et al. (2010) Effects of thymic selection of the T-cell repertoire on HLA class I-associated control of HIV infection. *Nature* 465(7296):350–4.
21. Košmrlj A, Chakraborty AK, Kardar M, Shakhnovich EI (2009) Thymic selection of T-cell receptors as an extreme value problem. *Physical Review Letters* 103(6):3–6.
22. Yates AJ (2014) Theories and quantification of thymic selection. *Frontiers in immunology* 5.
23. Feinerman O, Germain RN, Altan-Bonnet G (2008) Quantitative challenges in understanding ligand discrimination by $\alpha\beta$ T cells. *Molecular Immunology* 45(3):619–631.
24. François P, Voisinne G, Siggia ED, Altan-bonnet G, Vergassola M (2013) Phenotypic model for early T-cell activation displaying sensitivity , speci fi city , and antagonism. *Proceedings of the National Academy of Sciences of the United States of America* 110(10):E888–E897.
25. Miyazawa S, Jernigan RL (1985) Estimation of effective interresidue contact energies from protein crystal structures: quasi-chemical approximation. *Macromolecules* 18(3):534–552.
26. Miyazawa S, Jernigan RL (1996) Residue-residue potentials with a favorable contact pair term and an unfavorable high packing density term, for simulation and threading. *Journal of molecular biology* 256(3):623–644.
27. Hernandez JB, Newton RH, Walsh CM (2011) Life and death in the thymus - cell death signaling during T cell development. *Curr Opin Biol* 22(6):865–871.
28. Klein L, Kyewski B, Allen PM, Hogquist K (2014) Positive and negative selection of the T cell repertoire: what thymocytes see (and don't see). *Nature reviews. Immunology* 14(6):377–91.
29. Bryngelson JD, Wolynes PG (1989) Intermediates and barrier crossing in a random energy model (with applications to protein folding). *The Journal of Physical Chemistry* 93(19):6902–6915.
30. Sinclair C, Bains I, Yates AJ, Seddon B (2013) Asymmetric thymocyte death underlies the CD4:CD8 T-cell ratio in the adaptive immune system. *Proceedings of the National Academy of Sciences of the United States of America* 110(31):E2905–14.
31. Itano a, Robey E (2000) Highly efficient selection of CD4 and CD8 lineage thymocytes supports an instructive model of lineage commitment. *Immunity* 12(4):383–389.
32. Merkenschlager M, et al. (1997) How many thymocytes audition for selection? *The Journal of experimental medicine* 186(7):1149–58.
33. Ignatowicz L, et al. (1997) T cells can be activated by peptides that are unrelated in sequence to their selecting peptide. *Immunity* 7(2):179–186.
34. Tourne S, et al. (1997) Selection of a Broad Repertoire of CD4+ T Cells in H-2Ma 0 / 0 Mice. *Immunity* 7:187–195.
35. Zerrahn J, Held W, Raulet DH (1997) The MHC Reactivity of the T Cell Repertoire Prior to Positive and Negative Selection. *Cell* 88(5):627–636.
36. Kruglyak L, Nickerson Da (2001) Variation is the spice of life. *Nature genetics* 27(3):234–236.
37. Truong HH, Kim BL, Schafer NP, Wolynes PG (2013) Funneling and frustration in the energy landscapes of some designed and simplified proteins. *The Journal of Chemical Physics* 139:1–15.
38. Murphy LR, Wallqvist A, Levy RM (2000) Simplified amino acid alphabets for protein fold recognition and implications for folding. *Protein Engineering* 13(3):149–152.