

How nonuniform contact profiles of T cell receptors modulate thymic selection outcomes

Hanrong Chen,¹ Arup K. Chakraborty,^{2,3} and Mehran Kardar³

¹*Harvard John A. Paulson School of Engineering and Applied Sciences,
Harvard University, Cambridge, Massachusetts 02138, USA*

²*Departments of Chemical Engineering, Chemistry,
and Biological Engineering, Institute for Medical Engineering and Science,
Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

³*Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

(Dated: January 10, 2018)

T cell receptors (TCRs) bind foreign or self-peptides attached to major histocompatibility complex (MHC) molecules, and the strength of this interaction determines T cell activation. Optimizing the ability of TCRs to discriminate between foreign and self-peptides and recognize foreign peptides with high specificity is crucial for the adaptive immune system to properly function. This is accomplished through selection of TCRs in the thymus, where immature T cells with stochastically generated TCRs interact with a large number of self-peptide-MHC; if a TCR does not bind strongly enough to any self-peptide-MHC, or too strongly with at least one self-peptide-MHC, it is eliminated. Past theoretical work cast thymic selection as an extreme value problem, and characterized the statistical enrichment or depletion of amino acids in the post-selection TCR repertoire, showing how this process selects for T cells that are able to specifically recognize peptides derived from diverse pathogens, yet have limited self-reactivity. Here, we investigate how the degree of enrichment of amino acids in the post-selection repertoire is modified by nonuniform contacts of a TCR. Specifically, we were motivated by recent experiments showing that amino acids at certain positions of a TCR sequence are largely responsible for thymic selection outcomes, and crystal structure data that reveal a nonuniform contact profile between TCRs and their peptide-MHC ligands. Using a representative contact profile as an illustration, we show that the degree of enrichment now varies by position according to the contact profile, and, importantly, it depends on the interpretation of how nonuniform contacts are mediated during selection. We explain these nontrivial results analytically. Our study has implications for understanding the selection forces that shape the functionality of the post-selection TCR repertoire.

I. INTRODUCTION

T cell receptors (TCRs) bind peptides loaded onto major histocompatibility complex (MHC) molecules (abbreviated as peptide-MHC) on the surface of antigen presenting cells (APCs), and the strength of this interaction determines T cell activation [1, 2]. Such peptides are derived from either the host itself (self-antigens) or, potentially, pathogens infecting the host (foreign antigens), and optimizing the ability of T cells to discriminate between them and recognize foreign antigens with high specificity is essential for the proper functioning of the adaptive immune system. Immature T cells (or thymocytes) express a distinct TCR on their surface assembled through a stochastic process of gene rearrangement, generating a highly diverse repertoire [3] that is acted upon by selection in the thymus. There, thymocytes are screened against a large number of self-peptide-MHC; those that do not productively bind to self-peptide-MHC are deleted (this is called positive selection), and those that bind too strongly are also eliminated (this is called negative selection). T cells that survive thymic selection are exported to the body's periphery where they participate in the adaptive immune response. In this way, thymic selection shapes the post-selection TCR repertoire to potentially recognize diverse foreign antigens,

yet limit self-reactivity [4]. While detailed statistics of the post-selection repertoire are now available [5–7], how thymic selection achieves this outcome is less well-understood.

Mathematical models have shed light on various aspects of thymic selection (for a review, see Ref. [8]); in particular, models that represented TCR-peptide-MHC interactions as pairwise interactions between digit strings [9, 10] were useful for studying how TCR cross-reactivity can result. In Refs. [11, 12], some of the authors considered a more explicit representation of inter-amino acid interaction strengths [13], and characterized the statistical enrichment (or depletion) of amino acids in the post-selection TCR repertoire, computationally [11] and analytically [12], as a function of parameters such as the number of self-peptide-MHC encountered during selection. These studies also provided insight into how post-selection T cells are specific for unknown foreign peptides.

In the past few years, more detailed information about thymic selection outcomes have emerged. Advances in high-throughput sequencing have allowed researchers to quantify the statistics of post-selection TCR sequences in detail, revealing positional differences in the enrichment of amino acids [5, 6]. Furthermore, recent experiments found that pre-selection thymocytes that were activated by self-peptide-MHC (and hence would fail nega-

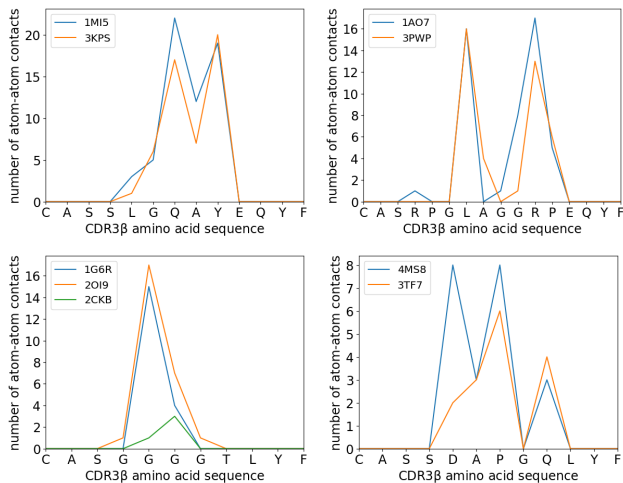


Figure 1: Number of atom-atom contacts (within 4\AA) between TCR amino acids and peptide-major histocompatibility complex (MHC) molecules, measured from crystal structures of TCR-peptide-MHC complexes (contact data from [7]; Protein Data Bank IDs are shown on top left of figures). Each x-axis shows a single TCR β -chain complementarity-determining region (CDR3 β) amino acid sequence, and each colored line plots the number of contacts made by that CDR3 β amino acid when bound to separate peptide-MHC molecules. These plots are a representation of the nonuniform contact interface between a TCR and its peptide-MHC ligand.

tive selection) were enriched in hydrophobic amino acids at positions 6 and 7 of the TCR β -chain complementarity-determining region (CDR3 β), while thymocytes passing both positive and negative selection were enriched in amino acids with moderate hydrophobicity at these positions [7]. These results are consistent with theoretical predictions made in previous studies [11, 12], but additionally find varying levels of enrichment at different CDR3 β positions. Furthermore, a large number of crystal structures of TCR-peptide-MHC complexes have been analyzed, that reveal a nonuniform contact interface between TCR and peptide-MHC (quantified, for example, by atom-atom contact profiles — see Fig. 1), and show that positions 6 and 7 of the CDR3 β sequence make the strongest contacts on average. These findings show that certain TCR positions are more important than others for influencing thymic selection outcomes; and furthermore, this information can be captured by nonuniform contact profiles such as those in Fig. 1.

While Refs. [11, 12] modeled thymic selection outcomes depending on properties of inter-amino acid interactions, they did not account for nonuniform contacts that a TCR makes with peptide-MHC. In this paper, we develop a formalism to do so, and investigate how this affects thymic selection outcomes. In particular, we consider two possible mechanisms by which nonuniform con-

tacts are mediated during thymic selection, one of which we term *deterministic*, and the other *stochastic*. We perform numerical simulations and analytical computations to characterize the degree of enrichment of amino acids in the post-selection TCR repertoire, and show that positions making stronger contacts end up with greater degrees of enrichment. While this may appear an expected outcome, the degree of enrichment depends nontrivially on the entire contact profile, as well as the positive and negative selection thresholds. In addition, we find that the interpretation of nonuniform contacts remarkably affects the degree, and even sign, of enrichment. Our study suggests a mechanistic origin for positional differences in post-selection TCR amino acid enrichment that has been observed in statistical analyses [5, 6] and experiments [7], and has implications for understanding how the functionality of the post-selection repertoire emerges.

The paper is organized as follows: in Section II, we develop a mathematical model of thymic selection that incorporates nonuniform contact profiles. In Section III, we consider two possible interpretations of nonuniform contacts, and study their effects on thymic selection outcomes. The stark differences between levels of enrichment from the two interpretations are explained analytically. In Section IV, we discuss further work and conclude.

II. MODEL DESCRIPTION

Immature T cells (or thymocytes) undergo positive and negative selection in the thymus before maturation. Each thymocyte expresses a distinct TCR on its surface, generated stochastically through V(D)J gene recombination and insertions and deletions of nucleotides (whose probabilities have been inferred from high-throughput sequencing data [3]), creating a diverse pre-selection repertoire. In the thymic cortex, thymocytes are presented with self-peptide-MHC by thymic antigen presenting cells (APCs). Thymocytes that do not bind strongly enough to any self-peptide-MHC die of insufficient survival signals; this is called positive selection. Thymocytes that survive positive selection migrate to the thymic medulla, where they are further screened against self-peptide-MHC, and those that bind too strongly with at least one self-peptide-MHC receive apoptotic signals and are eliminated; this is called negative selection.

We cast this process in a mathematical model as follows: following Ref. [11], TCR sequences, $\mathbf{t} = (t_1, \dots, t_N)$, of length N are generated by sampling amino acids t_i , $i = 1, \dots, N$, independently from a distribution $p_{\text{pre}}(t_i)$, that is taken to be the amino acid distribution of the human proteome [14]. While this is highly simplified compared to how the actual pre-selection repertoire is generated [3], it does not affect our later results, which concern the action of thymic selection on this distribution. Unlike Ref. [11], here a TCR has the capacity to make nonuniform contacts, which we capture by specifying $\mathbf{c} = (c_1, \dots, c_N)$. We describe values and

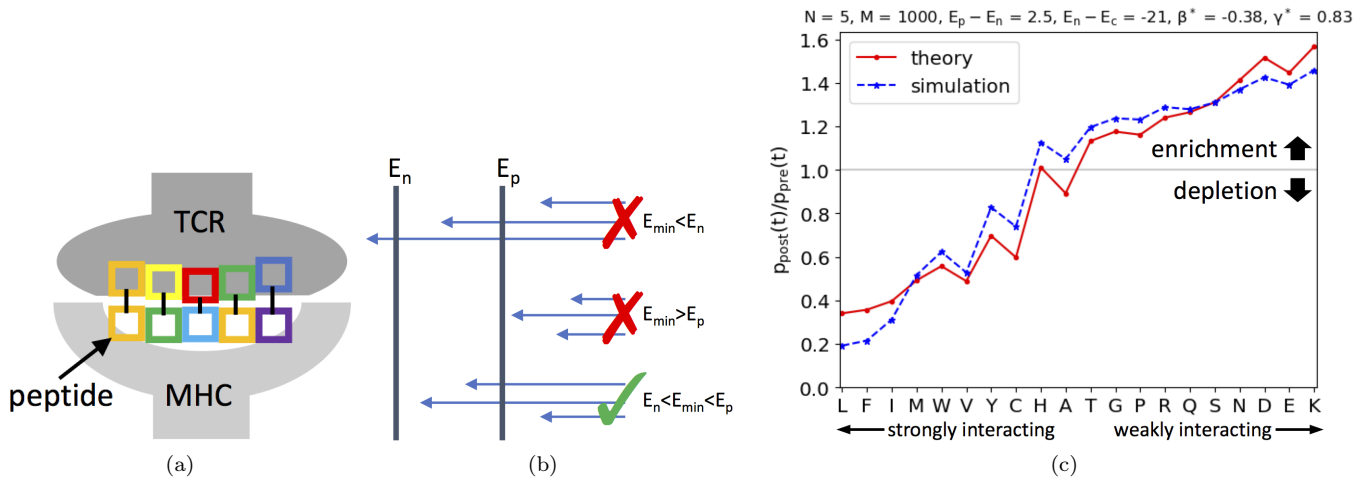


Figure 2: Mathematical model of positive and negative selection in the thymus. (a) String representation of TCR–peptide–MHC binding. A TCR sequence of length N (here $N = 5$) interacts with a peptide sequence, also of length N , bound to a MHC molecule. Each colored square represents a different amino acid. TCR and peptide amino acids interact in a pairwise fashion (see Eq. (1)). (b) Schematic of possible thymic selection outcomes. Top: thymocyte failing negative selection. Middle: thymocyte failing positive selection. Bottom: thymocyte surviving positive and negative selection. (c) Enrichment curve showing the amino acid distribution of TCR sequences surviving thymic selection, divided by the pre-selection distribution $p_{\text{pre}}(t)$. Parameter values (same as Ref. [12]): all $f(c_i) = 1$, $N = 5$, $M = 10^3$, $E_p - E_n = 2.5$, $E_n - E_c = -21$ (in units of $k_B T$). Blue dotted line: results of numerical simulations of the model for 5×10^6 independently generated TCR sequences, each encountering $M = 10^3$ independent self-peptide sequences. Red solid line: prediction of Eq. (8) with β^* obtained from Eqs. (9)–(11). Amino acids are arranged in order of increasing $[J(t, a)]_a$ (see Eq. (6)). Enrichment curves for all $i = 1, \dots, N$ have been averaged together. A value of 1 implies that the amino acid is observed equally often pre- and post-selection; values greater (less) than 1 imply enrichment (depletion) in the post-selection repertoire. Theory gave $\beta^* = -0.38$, implying selection for weakly interacting amino acids (right side of figure).

interpretations of \mathbf{c} in the next section.

During thymic selection, a TCR interacts with M independent self-peptide sequences, $\mathbf{s} = (s_1, \dots, s_N)$, of length N bound to MHC, that are also randomly generated according to p_{pre} . Because these are mostly linear peptides, they do not have an associated \mathbf{c} . Following Ref. [11], we model the binding strength between a TCR and self-peptide as pairwise interactions between TCR amino acid t_i and self-peptide amino acid s_i , for $i = 1, \dots, N$ (see Fig. 2a). Thus, the overall binding energy, $E(\mathbf{t}, \mathbf{s}, \mathbf{c})$, is

$$E(\mathbf{t}, \mathbf{s}, \mathbf{c}) = E_c + \sum_{i=1}^N f(c_i) J(t_i, s_i), \quad (1)$$

where E_c captures interactions between TCR and MHC, $J(t, s)$ is an interaction potential that in principle captures biochemical and other properties of inter-amino acid interactions, and $f(c_i)$ accounts for nonuniform contact at position i . Following Ref. [11], we use the Miyazawa–Jernigan (MJ) matrix for $J(t, s)$ [13], whose structure largely arises from hydrophobic forces [7, 15]; hydrophobic amino acids are *strongly* interacting (more negative $J(t, s)$ values), while hydrophilic amino acids are *weakly* interacting (less negative $J(t, s)$ values). (Note

that the results that follow do not qualitatively depend on the potential, but Ref. [7] noted the importance of hydrophobicity for the strength of TCR–peptide–MHC interactions, implying that using the MJ matrix is reasonable.) Reference [11] implicitly assumed that all $f(c_i) = 1$. Henceforth, we will specialize to one MHC type (i.e. E_c a constant) because its diversity is much lower than that of self-peptides (a human for example expresses 6 different MHC class I molecules), and we are not going to focus on how TCR cross-reactivity to other MHC molecules can arise [9, 10].

Positive and negative selection are carried out as follows: if the strongest (minimum) interaction energy between a TCR and M independent self-peptide–MHC is weaker (greater) than a positive selection threshold E_p , or if it is stronger (less) than a negative selection threshold E_n , the thymocyte dies. These hard constraints are consistent with experiments that found relatively small differences in TCR–ligand affinity at the negative selection threshold [16] (although they found that the positive selection threshold is softer). Thymocytes that interact with M self-peptide–MHC with strongest (minimum) binding energy within $[E_n, E_p]$ survive thymic selection and mature into naïve T cells (see Fig. 2b).

Parameter values: The CDR3 loops of a TCR typically make the greatest contacts with peptide (as opposed to MHC) [17]. Here, the TCR amino acid string \mathbf{t} represents the CDR3 β sequence, which is typically of length 10–18 [6]. During thymic selection, a thymocyte typically interacts with 10^3 – 10^4 self-peptide-MHC. Realistic values of the difference between positive and negative selection thresholds, and the T cell activation free energy with self-peptide alone (without MHC), are given by $E_p - E_n = 2.5 k_B T$ and $E_n - E_c = -21 k_B T$, respectively [11, 12].

Model without nonuniform contacts: References [11, 12] considered a model with all $f(c_i) = 1$, $N = 5$, $M = 10^3$, and $E_p - E_n$ and $E_n - E_c$ given above. Numerical simulations of this model resulted in a distribution of TCR sequences surviving selection that was statistically different from the pre-selection one, in that the former was enriched in weakly interacting amino acids and depleted in strongly interacting ones (see Fig. 2c). This result was consistent with later experiments [7].

A. Theory of the post-selection TCR repertoire distribution

Reference [12] developed an analytical theory for the post-selection TCR repertoire distribution, valid in the limit of $N, M \rightarrow \infty$. Here, we extend this theory to include nonuniform contact profiles. A self-contained derivation is in the Appendix; we state below the essential results.

A TCR with sequence $\mathbf{t} = (t_1, \dots, t_N)$ and contact profile $\mathbf{c} = (c_1, \dots, c_N)$ experiences a distribution of binding energies during thymic selection, with mean

$$\mu(\mathbf{t}, \mathbf{c}) = E_c + \sum_{i=1}^N \mu(t_i, c_i), \quad (2)$$

and variance

$$\nu(\mathbf{t}, \mathbf{c}) = \sum_{i=1}^N \nu(t_i, c_i), \quad (3)$$

where there are no cross-correlation terms in Eq. (3) because sites are independent. In Ref. [12] (where all $f(c_i) = 1$), $\mu(t_i, c_i)$ and $\nu(t_i, c_i)$ are given by

$$\mu(t_i) = [J(t_i, a)]_a, \quad (4)$$

and

$$\nu(t_i) = [J(t_i, a)^2]_a - [J(t_i, a)]_a^2, \quad (5)$$

where $[J(t, a)]_a$ and $[J(t, a)^2]_a$ are the first and second moments of interaction strength experienced by amino acid t when interacting with self-peptide sequences, i.e.

$$[J(t, a)]_a \equiv \sum_{a=1}^{20} J(t, a) p_{\text{pre}}(a). \quad (6)$$

In this paper, the forms of $\mu(t_i, c_i)$ and $\nu(t_i, c_i)$ depend on the interpretation of \mathbf{c} and will be specified in the next section.

For large N and M , the minimum binding energy experienced by a TCR during selection tends to the Gumbel distribution (see Appendix), whose peak is at

$$\rho_G(\mathbf{t}, \mathbf{c}) = \mu(\mathbf{t}, \mathbf{c}) - \sqrt{2\nu(\mathbf{t}, \mathbf{c})}\alpha, \quad (7)$$

where $\alpha = \sqrt{\log M}$, and whose variance is $\nu_G(\mathbf{t}, \mathbf{c}) = \pi^2 \nu(\mathbf{t}, \mathbf{c}) / (12 \log M)$. In the limit of $N, M \rightarrow \infty$ (keeping $N \propto \log M$), this distribution concentrates around its peak, which lies somewhere between E_n and E_p . In this limit, the post-selection TCR repertoire distribution becomes (see Appendix)

$$p_{\text{post}}(t_i | c_i) = \frac{1}{Z_{\beta, \gamma, c_i}} e^{-\beta[\mu(t_i, c_i) - \gamma \nu(t_i, c_i)]} p_{\text{pre}}(t_i), \quad (8)$$

where Z_{β, γ, c_i} ensures normalization,

$$\gamma = \frac{\alpha}{\sqrt{2 \sum_{i=1}^N \langle \nu(t_i, c_i) \rangle_{\beta, \gamma, c_i}}}, \quad (9)$$

and

$$\langle \nu(t_i, c_i) \rangle_{\beta, \gamma, c_i} \equiv \sum_{t_i=1}^{20} \nu(t_i, c_i) p_{\text{post}}(t_i | c_i). \quad (10)$$

The value of β in Eq. (8) is chosen such that $\langle \rho_G(\mathbf{t}, \mathbf{c}) \rangle$ (from Eq. (7)) lies within $[E_n, E_p]$. $p_{\text{post}}(t_i | c_i)$ depends on c_i explicitly, and on the full contact profile implicitly through γ . Also in this limit, $\langle \rho_G(\mathbf{t}, \mathbf{c}) \rangle$ becomes

$$\langle \rho_G(\mathbf{t}, \mathbf{c}) \rangle_{\beta, \gamma, \mathbf{c}} = \sum_{i=1}^N \langle \mu(t_i, c_i) - \gamma \nu(t_i, c_i) \rangle_{\beta, \gamma, c_i} - \frac{\alpha^2}{2\gamma}. \quad (11)$$

In practice, for each value of β , one iterates between Eqs. (9) and (10) until a self-consistent value of γ is obtained. Values of β corresponding to $\langle \rho_G(\mathbf{t}, \mathbf{c}) \rangle_{\beta, \gamma, \mathbf{c}} = E_p$ and $\langle \rho_G(\mathbf{t}, \mathbf{c}) \rangle_{\beta, \gamma, \mathbf{c}} = E_n$ are found, and the one closer to zero is taken to be β^* (or $\beta^* = 0$ if they straddle 0). Intuitively, β^* parametrizes the degree to which weakly or strongly interacting amino acids are enriched in the post-selection TCR repertoire; positive β^* implies selecting for strongly interacting amino acids, and negative β^* implies selecting for weakly interacting ones.

III. INTERPRETATIONS OF NONUNIFORM TCR CONTACT PROFILES

How is the picture of Fig. 2c modified when nonuniform contact profiles are taken into account? We were initially inspired by crystal structures of TCR-peptide-MHC complexes, that show a variation in the number of contacts made between TCR and peptide-MHC along the TCR sequence for every structure examined (see Fig. 1);




Figure 3: Schematic of deterministic (a) and stochastic (b) interpretations of nonuniform TCR contact profiles.

(a) Interactions between TCR amino acid t_i and peptide amino acid s_i at position i are screened by a factor c_i , here represented as contacts of different sizes. (b) Interactions between TCR and peptide amino acids at position i are made with probability c_i , $0 \leq c_i \leq 1$; this is represented as contacts of the same size that are either present or absent.

this contact profile need not even have a single maximum. However, when many contact profiles were added together, it appeared that positions 6 and 7 of the CDR3 β sequence made the greatest contacts with peptide-MHC on average (see Fig. 4a), which is consistent with experimental findings [7] and suggests that nonuniform contact profiles are useful in capturing positional differences that influence thymic selection outcomes. Figure 4a was obtained by rescaling the total number of contacts measured from 53 TCR-peptide-MHC structures such that the maximum is 1. We will make use of this representative, “average” contact profile as \mathbf{c} in the following.

A. Deterministic screening interpretation

Motivated by crystal structure analyses, we first considered a model where every interaction between TCR amino acid t_i and self-peptide amino acid s_i was weakened by a factor c_i , $0 \leq c_i \leq 1$. This is represented by setting $f(c_i) = c_i$ in Eq. (1). In principle, there is no reason to exclude values of $c_i > 1$ (depending on how interactions are modified by structure), but here we only considered reduced interactions, for example due to screening by intervening water molecules.

We performed numerical simulations of our model using the deterministic interpretation of the contact profile of Fig. 4a, with TCRs of length $N = 15$ and keeping the other parameter values the same as in Ref. [12] and Fig. 2c (i.e. $M = 10^3$, $E_p - E_n = 2.5 k_B T$, and $E_n - E_c = -21 k_B T$). Statistics of TCR amino acids in the post-selection repertoire are shown in Figs. 4b–c (dotted lines). Enrichment curves for each site are plotted in a different color. Now, the degree of enrichment depends on position; sites corresponding to the largest c_i (sites 6 and 7) experience the greatest degree of enrichment, while sites making no contact (sites 1–3, 14 and

15) have enrichment values close to 1. Furthermore, sites making contacts are now enriched in *strongly interacting* amino acids, which is opposite of Fig. 2c!

To understand this result, we repeated the previous theoretical analysis. The mean and variance of interaction energies for amino acid t_i at position i are now modified to

$$\mu_A(t_i, c_i) = c_i [J(t_i, a)]_a, \quad (12)$$

and

$$\nu_A(t_i, c_i) = c_i^2 [J(t_i, a)]_a^2 - c_i^2 [J(t_i, a)]_a^2, \quad (13)$$

by replacing $J(t_i, a) \rightarrow c_i J(t_i, a)$ in Eqs. (4) and (5). Using these expressions in Eqs. (9)–(11), we found $\beta^* = 0.79 > 0$, which indeed implies enrichment of strongly interacting amino acids. Plots of Eq. (8) for each position (solid lines in Figs. 4b–c) also resemble enrichment curves obtained from simulations.

Upon reflection, this change is not surprising. The contact profile of Fig. 4a has an “effective length” of $\sum_{i=1}^{15} c_i \approx 3.2 < 5$, implying that the mean binding energy experienced by a TCR during selection is roughly 64% of that experienced in the model with $N = 5$ and without nonuniform contacts which produced Fig. 2c (and indeed, we find $\beta^* > 0$ for $N = 3$ and without nonuniform contacts). Thus, it is plausible that here, TCR sequences need to be enriched in more strongly interacting amino acids in order to have their strongest binding energy during selection fall between E_n and E_p . Indeed, the analytical theory makes this intuition concrete.

Note that the analytical curves systematically underpredict the degree of enrichment, i.e. the value of β^* obtained from theory was slightly too small. For finite M , the extreme value distribution in fact has a finite width, and so one might underestimate β^* by matching $\langle \rho_G(\mathbf{t}, \mathbf{c}) \rangle_{\beta, \gamma, \mathbf{c}}$ to E_p . Indeed, when we include the next-order correction to $\rho_G(\mathbf{t}, \mathbf{c})$, which is positive (see after Eq. (A.2) in the Appendix), the analytical curves match those from simulations much better (not shown).

However, these results contradict the experiments of Ref. [7], as TCRs with strongly interacting amino acids at sites 6 and 7 should fail negative selection. This implies that some of our assumptions, for example that the interaction potential diminishes proportionally with the contact profile (i.e. $f(c_i) \propto c_i$), are incorrect.

B. Stochastic binding interpretation

In fact, crystal structures are merely static pictures of TCR-peptide-MHC binding, whereas the events leading to T cell activation following encounters with peptide-MHC are much more dynamic [17]. Indeed, it is known that the CDR3 β loop of TCRs is relatively flexible, with a range of conformations that can bind different ligands [18]. There also exist crystal structures of the same

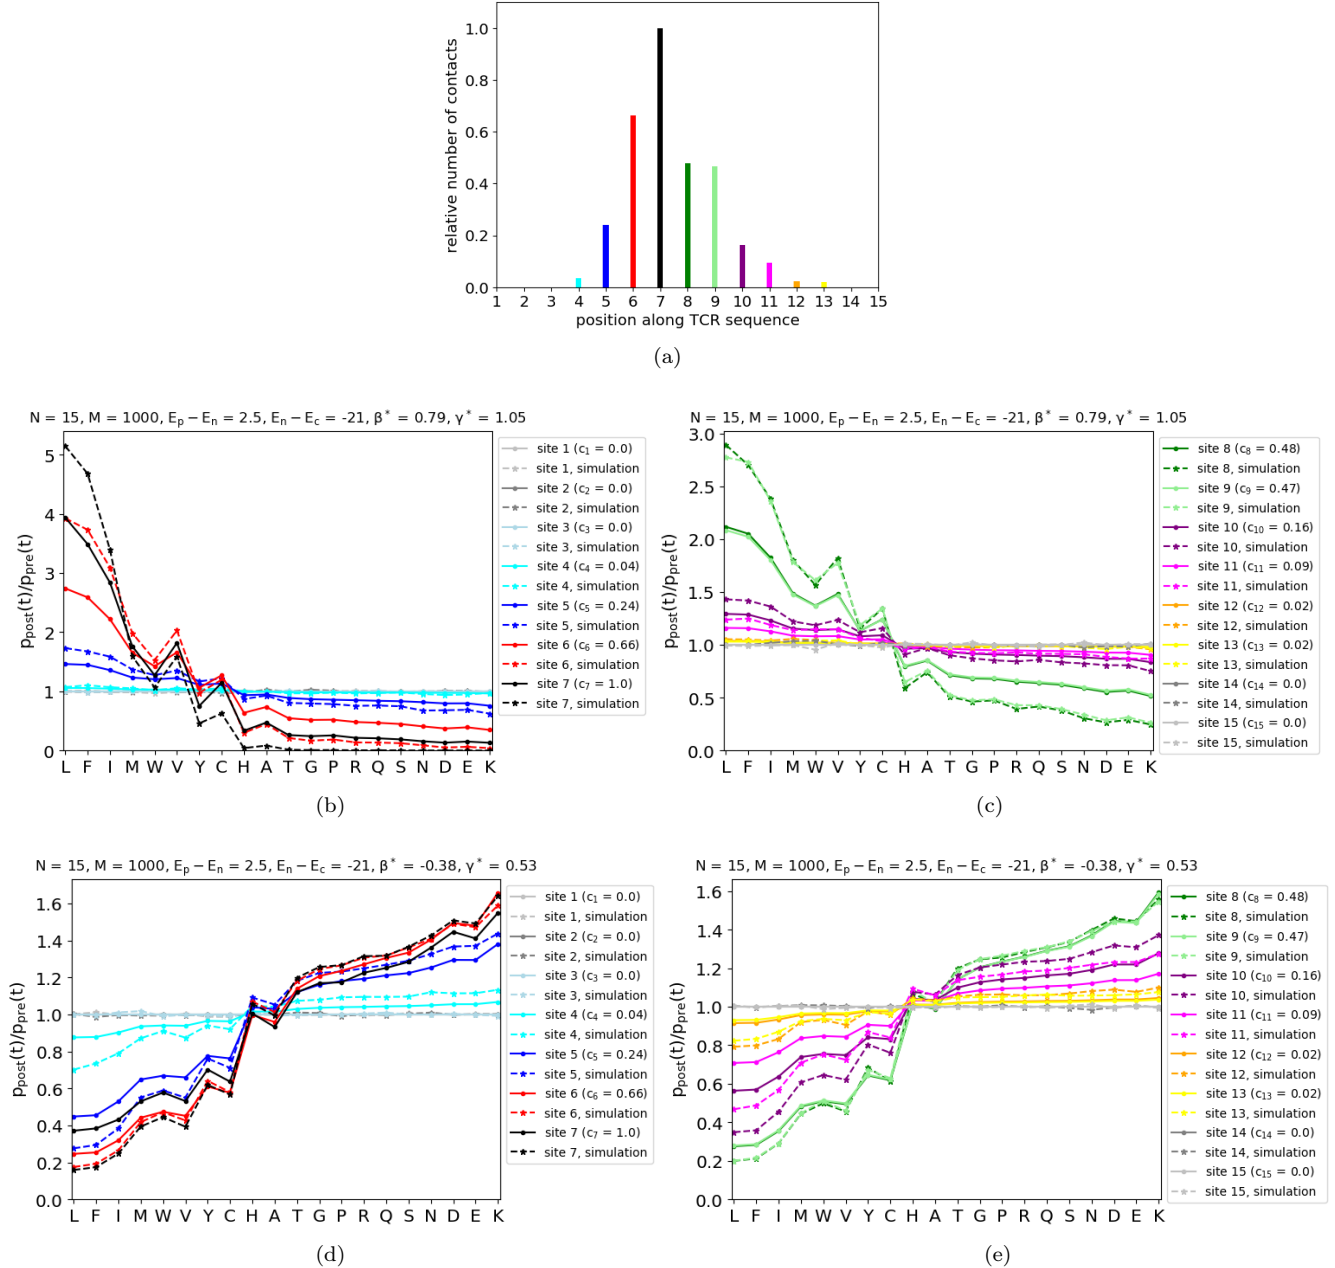


Figure 4: Thymic selection with nonuniform contact profiles. (a) Total number of atom-atom contacts for 53 TCR-peptide-MHC crystal structures, rescaled such that the maximum is 1 (contact data from Ref. [7]; see Protein Data Bank IDs therein). (b-e) Numerical simulations (dotted lines) and analytical computations (solid lines) for the deterministic (b-c) and stochastic (d-e) interpretations of nonuniform contacts. Each colored line represents the degree of enrichment of amino acids at a position along TCR sequences surviving thymic selection. For ease of viewing, sites 1–7 are plotted in (b,d), and sites 8–15 in (c,e). To create the dotted lines, 5×10^6 simulations were performed for each model with independently drawn TCR and self-peptide sequences. Solid lines are plots of Eq. (8) with β^* obtained from Eqs. (9)–(11). Strongly interacting amino acids were enriched ($\beta^* > 0$) for the deterministic model, but depleted ($\beta^* < 0$) for the stochastic model. Parameter values: $N = 15$, $M = 10^3$, $E_p - E_n = 2.5$, $E_n - E_c = -21$ (in units of $k_B T$). Note: in (c) and (e), the curves for sites 8 and 9 almost coincide as c_8 and c_9 are almost equal.

TCR bound to different peptide-MHC that show different parts of the TCR binding the different ligands [19].

These facts motivate another interpretation of nonuniform contact profiles: for every encounter during thymic selection, TCR amino acid t_i binds to self-peptide amino acid s_i with probability c_i , $0 \leq c_i \leq 1$. This is represented by setting $f(c_i) = X_i$, where X_i is a Bernoulli random variable with parameter c_i . This *stochastic binding* interpretation gives the same average interaction energy as the deterministic one:

$$\mu_B(t_i, c_i) = c_i [J(t_i, a)]_a, \quad (14)$$

while its variance is instead

$$\nu_B(t_i, c_i) = c_i [J(t_i, a)^2]_a - c_i^2 [J(t_i, a)]_a^2, \quad (15)$$

which has one factor of c_i multiplying $[J(t_i, a)^2]_a$ (as opposed to two in Eq. (13)).

We performed numerical simulations of this model, and obtained a post-selection repertoire shown in Figs. 4d–e (dotted lines). Again, the degree of enrichment is greater for sites making greater contacts. However, unlike the deterministic interpretation, here sites making contacts are enriched in *weakly interacting* amino acids. Using Eqs. (14) and (15), the analytical computations of Eqs. (9)–(11) also give $\beta^* = -0.38 < 0$. Thus, the argument that we made before was incomplete; even though the “effective length” here is approximately 3.2 as well because the mean interaction energies are equal (Eqs. (12) vs. (14)), their variances are different, and it is the *variance* that nontrivially modifies thymic selection outcomes! Because of the larger variance in this case, it is plausible that weakly interacting amino acids are sufficient for having the strongest of M binding energies reach below E_p . Again, the analytical theory makes this intuition concrete.

Another way to think of this is in terms of extreme values in *lengths*. During each interaction, a subset of positions of a TCR interact with self-peptide amino acids. In the negative selection-dominated regime, TCRs are more likely to be eliminated because of one interaction that is too strong, rather than all M interactions being too weak. Using the analytical theory for $N \geq 7$ and without nonuniform contacts gives $\beta^* = -\infty$, implying that no TCR that makes 7 or more contacts during selection should survive. Also, the analytical theory for $N = 6$ and without nonuniform contacts gives $\beta^* = -1.7$, implying that TCRs making 6 contacts during selection have a high probability of failing negative selection too. Thus, the majority of TCRs that survive selection will have made their strongest interaction with ~ 5 contacts, and indeed, the enrichment curves for sites 6–9 resemble that of Fig. 2c.

The predicted curves from Eq. (8) (solid lines in Figs. 4d–e) agree well with numerical results. Note that the theory predicts that site 7 is less enriched than site 6, even through $c_7 > c_6$ (compare black and red solid curves in Fig. 4d)! That the degree of enrichment is not

monotonically related to c_i can be seen by differentiating $\mu_B(t_i, c_i) - \gamma \nu_B(t_i, c_i)$ w.r.t. c_i , which reveals that according to this theory, there is an optimum $0 < c_i < 1$ that gives the greatest enrichment, which is different for each amino acid (because this depends on $[J(t, a)]_a$ and $[J(t, a)^2]_a$). (Repeating this calculation for the deterministic model however gives an optimal c_i that is negative, implying that there is monotonicity for that case.) Numerical simulations however do not show this non-monotonicity, and so we believe this is a manifestation of finite M and N .

To summarize, we have shown how two possible interpretations of nonuniform TCR contacts modify thymic selection outcomes in different ways. We have also explained our results using an analytical theory valid in the limit of large N and $M \propto \ln N$.

IV. DISCUSSION

In this paper, we presented a formalism to incorporate information about TCR structure, through its nonuniform contact profile, into a model of thymic selection. We showed how this leaves statistical signatures at different positions of post-selection TCR sequences. Importantly, we showed how these signatures depend on implementation of nonuniform contacts, as a deterministic screening of interactions, or as probabilities of stochastic binding events. In the actual thymus, these and other scenarios probably play a role, and it would be interesting to quantify their relative roles.

While we have added a further degree of realism to modeling TCR–peptide-MHC interactions, many features have been left out, such as:

- We did not account for nonuniform contacts with MHC. This is important because a TCR binding more strongly with MHC might require weaker interactions with peptide in order to be activated; this has been studied previously to explain TCR cross-reactivity with foreign MHC molecules [9, 10]. Mathematically, this may be included into the model by modulating E_c by $f(c_c)$.
- Stochastic binding events may be correlated between neighboring TCR positions. Mathematically, this introduces cross-correlation terms into the variance of interaction energies (Eq. (15)), and makes the Legendre transform and self-averaging within the theory more complicated (see Appendix).
- The pre-selection TCR distribution resulting from V(D)J recombination is not factorizable into $\prod_{i=1}^N p_{\text{pre}}(t_i)$, which also introduces correlations into $P_{\text{post}}(\mathbf{t})$.
- The same TCR amino acid may contact more than one peptide-MHC amino acid, and the same peptide-MHC amino acid may contact more than one TCR amino acid.

Nonuniform contact profiles are a step towards modeling the complicated, structure-dependent nature of TCR-peptide-MHC interactions. However, a more informed method should be developed to infer such a profile from crystal structures. For example, we made use of measurements of the number of peptide-MHC atoms a distance of 4\AA away from TCR amino acids, but characteristic distances should depend on the kind of interaction (hydrogen bonding, van der Waals, etc.).

We limited computations in this paper to one contact profile to illustrate its effect on thymic selection outcomes. Separately, it would be interesting to characterize the statistics of contact profiles from crystal structures. This task however is limited by the relatively small number of crystal structures known, as opposed to the large quantities of high-throughput sequencing data available. We note that an additional step needs to be taken to connect the model in this paper to statistics of aligned TCR sequences that appear in, e.g., Refs. [5, 6] (which do not focus on peptide contact residues but rather the entire aligned CDR3 region, and hence do not find enrichment at TCR positions corresponding to recent experiments [7]): the model should be run separately for different contact profiles, and the enrichment curves should be averaged together according to the *passing rates* for the different contact profiles. Note that this is a possible mechanism for obtaining enrichment of an amino acid at one position in a sequence alignment, and depletion of the *same* amino acid at another, because these positions might feature in *different* contact profiles that have different β^* . This also implies that the results from an *average* contact profile are, in general, different from running the model for separate contact profiles and averaging the results together, because the former does not account for different passing rates. Thus, it would be interesting to attempt the inverse problem of inferring differential contacts and binding tendencies at different positions of a TCR sequence from positional differences in the post-selection TCR repertoire, but this is complicated by the under-determinacy of the problem.

Inferring overall patterns determining a TCR's specificity to peptide-MHC from knowledge of TCR-peptide-MHC crystal structures is challenging because there is no one canonical way by which a TCR interacts with peptide-MHC [17, 20, 21]; the same TCR may bind different peptide-MHC in very different ways [19]. Predicting TCR sequences that recognize a given set of peptide-MHC by inferring "sequence motifs" of TCRs has been achieved very recently [22, 23]. We believe that analyzing thymic selection outcomes has implications for antigenic

specificity, because surviving thymic selection involves interactions with a large number of self-peptide-MHC, and thus features relevant for thymic selection outcomes are also relevant for antigenic specificity. The features we have studied here have measurable effects on the post-selection repertoire, and thus they probably contribute to antigenic specificity as well, which perhaps gives a mechanistic basis for the sequence motifs discovered in the recent studies [22, 23].

Recently, a paper that also modified the thymic selection model of Refs. [11, 12] to include positional differences in TCR-peptide-MHC interactions appeared [24]. In essence, for a given TCR, they drew the values of $f(c_i)J(t_i, s_i)$ in Eq. (1), $i = 1, \dots, N$, independently from a Gaussian distribution. Thus, they reduced thymic selection to an extreme value problem with a random energy model, in which "TCRs" and "amino acids" lose their meaning; this contradicts studies that find predictive features determining specificity that are based on TCR amino acid sequences [7, 22, 23]. However, their model is simpler to analyze and may be a useful null model, and it would be interesting to compare it with the results from averaging many different contact profiles together. The authors also commented that the model of Refs. [11, 12] fell short in that very few self-peptides (i.e. the most strongly interacting ones) perform the job of negative selection equally effectively as the full panel of $M = 10^3\text{--}10^4$ self-peptides. While this is true (which follows directly from specifying an inter-amino acid interaction potential), it is not known how large this fraction is for the real thymic selection process. Experimentally, this could be tested by engineering the thymus to contain peptides consisting of only strongly interacting amino acids [4]. Also, it is possible that these peptides are somehow found rarely or not at all in the thymus, because self-peptides are chopped-up versions of actual proteins. Large values of M are likely still required to randomly generate such special peptides. We note that the model we study here moves away from the limitations they raised, as different TCRs with different contact profiles need not bind equally strongly with the same, strongly interacting peptide.

ACKNOWLEDGMENTS

This work was supported by the Ragon Institute of MGH, MIT and Harvard and an A*STAR Scholarship (to H.C.). M.K. acknowledges support from NSF through grant number DMR-1708280.

[1] S. M. Alam, P. J. Travers, J. L. Wung, W. Nasholds, S. Redpath, S. C. Jameson, and N. R. J. Gascoigne, *Nature* **381**, 616 (1996).
 [2] M. Krogsgaard and M. M. Davis, *Nat. Immunol.* **6**, 239 (2005).

[3] A. Murugan, T. Mora, A. M. Walczak, and C. G. Callan, Jr., *Proc. Natl. Acad. Sci. U.S.A.* **109**, 16161 (2012).
 [4] E. S. Huseby, J. White, F. Crawford, T. Vass, D. Becker, C. Pinilla, P. Marrack, and J. W. Kappler, *Cell* **122**, 247 (2005).

- [5] Y. Elhanati, A. Murugan, C. G. Callan, Jr., T. Mora, and A. M. Walczak, *Proc. Natl. Acad. Sci. U.S.A.* **111**, 9875 (2014).
- [6] Z. Sethna, Y. Elhanati, C. R. Dudgeon, C. G. Callan, Jr., A. J. Levine, T. Mora, and A. M. Walczak, *Proc. Natl. Acad. Sci. U.S.A.* **114**, 2253 (2017).
- [7] B. D. Stadinski, K. Shekhar, I. Gómez-Touriño, J. Jung, K. Sasaki, A. K. Sewell, M. Peakman, A. K. Chakraborty, and E. S. Huseby, *Nat. Immunol.* **17**, 946 (2016).
- [8] A. J. Yates, *Front. Immunol.* **5**, 13 (2014).
- [9] V. Detours, R. Mehr, and A. S. Perelson, *J. Theor. Biol.* **200**, 389 (1999).
- [10] V. Detours and A. S. Perelson, *Proc. Natl. Acad. Sci. U.S.A.* **96**, 5153 (1999).
- [11] A. Košmrlj, A. K. Jha, E. S. Huseby, M. Kardar, and A. K. Chakraborty, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 16671 (2008).
- [12] A. Košmrlj, A. K. Chakraborty, M. Kardar, and E. I. Shakhnovich, *Phys. Rev. Lett.* **103**, 068103 (2009).
- [13] S. Miyazawa and R. L. Jernigan, *J. Mol. Biol.* **256**, 623 (1996).
- [14] B. L. Aken et al., *Nucleic Acids Res.* **45**, D635 (2017).
- [15] H. Li, C. Tang, and N. S. Wingreen, *Phys. Rev. Lett.* **79**, 765 (1997).
- [16] M. A. Daniels, E. Teixeira, J. Gill, B. Hausmann, D. Roubaty, K. Holmberg, G. Werlen, G. A. Holländer, N. R. J. Gascoigne, and E. Palmer, *Nature* **444**, 724 (2006).
- [17] K. C. Garcia and E. J. Adams, *Cell* **122**, 333 (2005).
- [18] B. M. Baker, D. R. Scott, S. J. Blevins, and W. F. Hawse, *Immunol. Rev.* **250**, 10 (2012).
- [19] L. A. Colf, A. J. Bankovich, N. A. Hanick, N. A. Bowerman, L. L. Jones, D. M. Kranz, and K. C. Garcia, *Cell* **129**, 135 (2007).
- [20] X. Ysern, H. Li, and R. A. Mariuzza, *Nat. Struct. Biol.* **5**, 412 (1998).
- [21] D. Housset and B. Malissen, *Trends Immunol.* **24**, 429 (2003).
- [22] P. Dash, A. J. Fiore-Gartland, T. Hertz, G. C. Wang, S. Sharma, A. Souquette, J. C. Crawford, E. B. Clemens, T. H. O. Nguyen, K. Kedzierska, N. L. La Gruta, P. Bradley, and P. G. Thomas, *Nature* **547**, 89 (2017).
- [23] J. Glanville, H. Huang, A. Nau, O. Hatton, L. E. Wagar, F. Rubelt, X. Ji, A. Han, S. M. Krams, C. Petrus, N. Haas, C. S. Lindestam Arlehamn, A. Sette, S. D. Boyd, T. J. Scriba, O. M. Martinez, and M. M. Davis, *Nature* **547**, 94 (2017).
- [24] J. T. George, D. A. Kessler, and H. Levine, *Proc. Natl.*

- Acad. Sci. U.S.A.* **114**, E7875 (2017).
- [25] M. Kardar, *Phys. Rev. Lett.* **51**, 523 (1983).

Appendix: Derivation of the post-selection TCR repertoire distribution for large N , M

Here, we provide a self-contained derivation of a theory for the post-selection TCR repertoire distribution, extending Ref. [12] to capture a nonuniform contact profile. The probability that a TCR with sequence \mathbf{t} and contact profile \mathbf{c} survives selection, $P(\text{post}|\mathbf{t}, \mathbf{c})$, is equal to the probability that the minimum of M binding energies it encountered lies within $[E_n, E_p]$. Now, the binding energy $E(\mathbf{t}, \mathbf{s}, \mathbf{c})$ between a TCR and self-peptide sequence \mathbf{s} bound to MHC will be Gaussian distributed for large N , by the Central Limit Theorem (because Eq. (1) contains a sum of N independent, but not identically distributed, $f(c_i)J(t_i, s_i)$), with mean $\mu(\mathbf{t}, \mathbf{c})$ and variance $\nu(\mathbf{t}, \mathbf{c})$ given by Eqs. (2) and (3), respectively. And if $E(\mathbf{t}, \mathbf{s}, \mathbf{c})$ is Gaussian distributed, then the limiting distribution of $\min_{k=1}^M E(\mathbf{t}, \mathbf{s}^{(k)}, \mathbf{c})$ as $M \rightarrow \infty$ will be the Gumbel distribution, which has cumulative distribution function

$$P_G(\min_k E^{(k)} < E) = 1 - \exp \left[- \exp \left(\frac{E - a_M(\mathbf{t}, \mathbf{c})}{b_M(\mathbf{t}, \mathbf{c})} \right) \right], \quad (\text{A.1})$$

where

$$a_M(\mathbf{t}, \mathbf{c}) = \mu(\mathbf{t}, \mathbf{c}) - \sqrt{2\nu(\mathbf{t}, \mathbf{c})}\alpha, \quad (\text{A.2})$$

with $\alpha = \sqrt{\log M} - \frac{1}{4\sqrt{\log M}}(\log \log M + \log 4\pi) + \mathcal{O}((\log M)^{-3/2})$ (depending only on M), and

$$b_M(\mathbf{t}, \mathbf{c}) = \sqrt{\frac{\nu(\mathbf{t}, \mathbf{c})}{2 \log M}} + \mathcal{O}((\log M)^{-3/2}). \quad (\text{A.3})$$

The peak of this distribution is $\rho_G(\mathbf{t}, \mathbf{c}) = a_M(\mathbf{t}, \mathbf{c})$, and its variance is $\nu_G(\mathbf{t}, \mathbf{c}) = \frac{\pi^2}{6} b_M^2(\mathbf{t}, \mathbf{c})$. In the main text of this paper, we used the leading-order value of α , $\alpha = \sqrt{\log M}$ (see after Eq. (7)).

Using Bayes' rule, the posterior distribution of TCR sequences surviving selection, $P(\mathbf{t}|\text{post}, \mathbf{c})$, is given by

$$\begin{aligned} P(\mathbf{t}|\text{post}, \mathbf{c}) &= \frac{P(\text{post}|\mathbf{t}, \mathbf{c})P_{\text{pre}}(\mathbf{t})}{\sum_{\mathbf{t}} P(\text{post}|\mathbf{t}, \mathbf{c})P_{\text{pre}}(\mathbf{t})} \\ &\propto [P_G(E < E_p) - P_G(E < E_n)] \times P_{\text{pre}}(\mathbf{t}) \\ &= \left\{ \exp \left[- \exp \left(\frac{E_n - a_M(\mathbf{t}, \mathbf{c})}{b_M(\mathbf{t}, \mathbf{c})} \right) \right] - \exp \left[- \exp \left(\frac{E_p - a_M(\mathbf{t}, \mathbf{c})}{b_M(\mathbf{t}, \mathbf{c})} \right) \right] \right\} \times P_{\text{pre}}(\mathbf{t}). \end{aligned} \quad (\text{A.4})$$

While this is exact within the Gaussian approximation,

it is not immediately obvious how to make progress

quantifying the enrichment of post-selection TCR amino acids by direct marginalization, i.e. $P(t_i|\text{post}, \mathbf{c}) = \sum_{\{t_j\}_{j=1 \dots N \setminus i}} P(\mathbf{t} = (t_1 \dots t_N) | \text{post}, \mathbf{c})$.

Reference [12] made progress in the limit of $N, M \rightarrow \infty$ (keeping $N \propto \log M$), when the extreme value distribution concentrates around its peak, $\rho_G(\mathbf{t}, \mathbf{c})$, which lies somewhere between E_n and E_p . Now, we ask: what is the probability distribution that minimizes the relative entropy (or Kullback–Liebler divergence) to the pre-selection distribution $P_{\text{pre}}(\mathbf{t})$, given the constraint that $\rho_G(\mathbf{t}, \mathbf{c})$ lies between E_n and E_p ? The answer is

$$P(\mathbf{t} | \text{post}, \mathbf{c}) = \frac{1}{Z_{\beta, \mathbf{c}}} e^{-\beta(\mu(\mathbf{t}, \mathbf{c}) - \sqrt{2\nu(\mathbf{t}, \mathbf{c})\alpha})} P_{\text{pre}}(\mathbf{t}), \quad (\text{A.5})$$

where $Z_{\beta, \mathbf{c}}$ ensures normalization, and we have used Eq. (A.2). Here, β is a Lagrange multiplier constraining the value of $\langle \mu(\mathbf{t}, \mathbf{c}) - \sqrt{2\nu(\mathbf{t}, \mathbf{c})\alpha} \rangle \equiv \sum_{\mathbf{t}} (\mu(\mathbf{t}, \mathbf{c}) - \sqrt{2\nu(\mathbf{t}, \mathbf{c})\alpha}) P(\mathbf{t} | \text{post}, \mathbf{c})$ to lie between E_n and E_p . The

optimal value of β , β^* , is the one as close to 0 as possible that satisfies this constraint. The mapping from hard constraints on the extreme value to a constraint on its mean is analogous to that from the microcanonical to the canonical ensemble in the thermodynamic limit [12].

The marginal distribution of amino acid t_i at position i of post-selection TCR sequences, $P(t_i | \text{post}, \mathbf{c})$, may be obtained from Eq. (A.5) by taking a sum over 20^{N-1} terms. However, in the $N \rightarrow \infty$ limit, $\nu(\mathbf{t}, \mathbf{c})$ self-averages, i.e. $\sum_{i=1}^N \nu(t_i, c_i) \rightarrow \sum_{i=1}^N \langle \nu(t, c_i) \rangle$; thus, performing the double Legendre transform on $\rho_G(\mathbf{t}, \mathbf{c})$ w.r.t. $\nu(\mathbf{t}, \mathbf{c})$ (equivalently Hamiltonian minimization [25]), and replacing $\nu(\mathbf{t}, \mathbf{c})$ by its self-averaged value $\sum_{i=1}^N \langle \nu(t, c_i) \rangle$, Eq. (A.5) factorizes into Eq. (8) of the main text, where γ (given by Eq. (9)) is the conjugate variable to $\nu(\mathbf{t}, \mathbf{c})$. After the Legendre transforms w.r.t. $\nu(\mathbf{t}, \mathbf{c})$, $\langle \rho_G(\mathbf{t}, \mathbf{c}) \rangle$ becomes Eq. (11). How β^* is found in practice is described after Eq. (11) of the main text.