

Quantifying selection in immune receptor repertoires

Yuval Elhanati ^{*}, Anand Murugan [†], Curtis G. Callan Jr. [‡], Thierry Mora [§] and Aleksandra M. Walczak ^{*}

^{*}Laboratoire de physique théorique, UMR8549, CNRS and École normale supérieure, 24, rue Lhomond, 75005 Paris, France, [§]Laboratoire de physique statistique, UMR8550, CNRS and École normale supérieure, 24, rue Lhomond, 75005 Paris, France, [†]Department of Applied Physics, Stanford University, California 94305 USA, and [‡]Joseph Henry Laboratories, Princeton University, Princeton, New Jersey 08544 USA

Submitted to Proceedings of the National Academy of Sciences of the United States of America

April 15, 2014

The efficient recognition of pathogens by the adaptive immune system relies on the diversity of receptors displayed at the surface of immune cells. T-cell receptor diversity results from an initial random DNA editing process, called VDJ recombination, followed by functional selection of cells according to the interaction of their surface receptors with self and foreign antigenic peptides. To quantify the effect of selection on the highly variable elements of the receptor, we apply a probabilistic maximum likelihood approach to the analysis of high-throughput sequence data from the β -chain of human T-cell receptors. We quantify selection factors for V and J gene choice, and for the length and amino-acid composition of the variable region. Our approach is necessary to disentangle the effects of selection from biases inherent in the recombination process. Inferred selection factors differ little between donors, or between naive and memory repertoires. The number of sequences shared between donors is well-predicted by the model, indicating a purely stochastic origin of such “public” sequences. We find a significant correlation between biases induced by VDJ recombination and our inferred selection factors, together with a reduction of diversity during selection. Both effects suggest that natural selection acting on the recombination process has anticipated the selection pressures experienced during somatic evolution.

Significance statement

The immune system defends against pathogens via a diverse population of T-cells that display different antigen recognition surface receptor proteins. Receptor diversity is produced by an initial random gene recombination process, followed by selection for a desirable range of peptide binding. Although recombination is well-understood, selection has not been quantitatively characterized. By combining high throughput sequencing data with modeling, we quantify the selection pressure that shapes functional repertoires. Selection is found to vary little between individuals or between the naive and memory repertoires. It reinforces the biases of the recombination process, meaning that sequences more likely to be produced are also more likely to pass selection. The model accounts for “public” sequences shared between individuals as resulting from pure chance.

The T-cell response of the adaptive immune system begins when receptor proteins on the surface of these cells recognize a pathogen peptide presented by an antigen presenting cell. The immune cell repertoire of a given individual is comprised of many clones, each with a distinct surface receptor. This diversity, which is central to the ability of the immune system to defeat pathogens, is initially created by a stochastic process of germline DNA editing (called VDJ recombination) that gives each new immune cell a unique surface receptor gene. This initial repertoire is subsequently modified by selective forces, including thymic selection against excessive (or insufficient) recognition of self proteins, that are also stochastic in nature. Due to this stochasticity and the large T-cell diversity, these repertoires are best described by probability distributions. In

this paper we apply a probabilistic approach to sequence data to obtain quantitative measures of the selection pressures that shape T-cell receptor repertoires.

New receptor genes are formed by choosing at random from a set of genomic templates for several sub-regions (V, D and J) of the complete gene. Insertion and deletion of nucleotides in the junctional regions between the V and D and D and J genes greatly enhances diversity beyond pure VDJ combinatorics [1]. The variable region of the gene lies between the last amino acids of the V segment and the beginning of the J segment; it codes for the Complementarity Determining Region 3 (CDR3) loop of the receptor protein, a region known to be functionally important in recognition [2]. Previous studies have shown that immune cell receptors are not uniform in terms of VDJ gene segment usage [3, 4, 5, 6], or probability of generation [1], and that certain receptors are more likely than others to be shared by different individuals [7, 4]. In other words, the statistical properties of the immune repertoire are rather complex, and their accurate determination requires sophisticated methods.

Over the past few years, advances in sequencing technology have made it possible to sample the T-cell receptor diversity of individual subjects in great depth [8], and this has in turn led to the development of sequence statistics-based approaches to the study of immune cell diversity [9, 10]. In particular, we recently quantitatively characterized the primary, pre-selection diversity of the human T-cell repertoire by learning the probabilistic rules of VDJ recombination from out-of-frame DNA sequences that cannot be subject to functional selection and whose statistics can reflect only the recombination process [1]. After generation, T-cells undergo a somatic selection process in the thymus [11] and later in the periphery [12]. Cells that pass thymic selection enter the peripheral repertoire as ‘naive’ T-cells, and the subset of naive cells that eventually engage in an immune response will survive as a long-lived ‘memory’ pool. Even though we now understand the statistical properties of the initial repertoire of immune receptors [1], and despite some theoretical studies of thymic selection at the molecular level [13, 14], a quantitative understanding of how selection modifies those statistics to produce the naive and memory repertoires is lacking.

Reserved for Publication Footnotes

In this paper, we build on our understanding of the primitive pre-selection distribution of T-cell receptors to derive a statistical method for identifying and quantifying selection pressures in the adaptive immune system. We apply this method to naive and memory DNA sequences of human T-cell β chains obtained from peripheral blood samples of nine healthy individuals (Fig. 1). Our goal is to characterize the likelihood that any given sequence, once generated, will survive selection. Our analysis reveals strong and reproducible signatures of selection on specific amino acids in the CDR3 sequence, and on the usage of V and J genes. Most strikingly, we find significant correlation between the primitive generation probability of a sequence and the probability it will pass selection. This suggests that natural selection, which acts on very long time scales to shape the generation mechanism itself, may have tuned it to anticipate somatic selection, which acts on single cells throughout the lifetime of an individual. The quantitative features of selection inferred from our model vary very little between donors, indicating that these features are universal. In addition, our measures of selection pressure on the memory and naive repertoires are statistically indistinguishable, consistent with the hypothesis that the memory pool is a random subsample of the naive pool.

Methods

We analyzed human CD4+ T-cell β -chain DNA sequence reads (60 or 101 nucleotides long) centered around the CDR3 region. T-cells were obtained from nine individuals and sorted into naive (CD45RO-) and memory (CD45RO+) subsets, yielding datasets of $\sim 200,000$ unique naive and $\sim 120,000$ unique memory sequences per individual, on average. The datasets are the same as those used in [1] and were obtained by previously described methods [15, 16].

In [1] we used the “nonproductive” sequences (those where either the J sequences are out of frame, or the CDR3 sequences have a stop codon) to characterize the receptor generation process. The result of that analysis was an evaluation of the probability $P_{\text{pre}}(\bar{\sigma})$ that a VDJ recombination event will produce a β -chain gene consistent with the specific DNA sequence read $\bar{\sigma}$. In this study we focus instead on the in-frame, productive, sequences (from both the naive and the memory repertoires) with the goal of quantifying how the post-selection probability distribution on sequences is modified from the original distribution $P_{\text{pre}}(\bar{\sigma})$. In what follows we distinguish between the read $\bar{\sigma}$ and the CDR3 region $\bar{\tau}$, the latter defined to run from a conserved cysteine near the end of the V segment to the last amino acid of the read (leaving two amino acids to the conserved Phe). The CDR3 amino acid sequence can be uniquely read off from each in-frame sequence read; by contrast, the specific V- and J-genes responsible for the read may not be uniquely identifiable (because of the relatively short read length). An unambiguous selection effect can be seen by comparing the length distribution of CDR3 regions between the pre-selection ensemble and the naive, or memory, datasets (Fig. 2A): sequences that are longer or shorter than the mean are suppressed resulting in a more peaked distribution.

For each receptor sequence, we define a selection factor $Q(\bar{\sigma})$ that quantifies whether selection (thymic selection or later selection in the periphery) has enriched or impoverished the frequency of $\bar{\sigma}$ compared to the pre-selection ensemble. Since the generation probability of sequences varies over many orders of magnitude, such a comparison is the only way to define selection strength. Denoting by $P_{\text{post}}(\bar{\sigma})$ the distribution of sequences in the selected naive or memory pools, we will set $P_{\text{post}}(\bar{\sigma}) = Q(\bar{\sigma})P_{\text{pre}}(\bar{\sigma})$. Due to the large number of possi-

ble sequences, we cannot sample the post-selection probability P_{post} for each sequence directly from the data; we need a reduced complexity model to estimate it. We propose a simple model, summarized in Fig. 1A, that we will show captures the main features of selection:

$$Q(\bar{\tau}, V, J) = \frac{P_{\text{post}}(\bar{\tau}, V, J)}{P_{\text{pre}}(\bar{\tau}, V, J)} = \frac{1}{Z} q_L q_{VJ} \prod_{i=1}^L q_{i:L}(a_i), \quad [1]$$

where V and J denote the choice of V and J segments in the sequence $\bar{\sigma}$, L is the amino-acid length of the CDR3 specified by the read, $(\tau_1, \dots, \tau_{3L})$ is CDR3 nucleotide sequence, and (a_1, \dots, a_L) its amino-acid sequence. The factors q_L , $q_{i:L}(a)$ and q_{VJ} denote, respectively, selective pressures on the CDR3 length, its composition, and the associated VJ identities. Note that the D segment is entirely included in this junctional region, so selection acting on it is encoded in the $q_{i:L}$ factors. Z enforces the model normalization condition $\sum_{\bar{\tau}, V, J} Q(\bar{\tau}, V, J) P_{\text{pre}}(\bar{\tau}, V, J) = 1$.

It is important to understand why we do not write Q directly as a function of the read $\bar{\sigma}$. While $(\bar{\tau}, V, J)$ determines $\bar{\sigma}$ and $\bar{\sigma}$ determines $\bar{\tau}$, V and J cannot always be inferred deterministically from the read $\bar{\sigma}$. The VJ assignment of any given read will have to be treated as probabilistically defined hidden variables. In addition, because of correlations in P_{pre} , the q factors cannot be identified with marginal enrichment factors (so that, for example, $P_{i:L,\text{data}}(a_i)/P_{i:L,\text{pre}}(a_i)$, cannot be set equal to $q_{i:L}(a_i)$). For all these reasons, we must use a maximum likelihood procedure to learn the q_L , $q_{i:L}$ and q_{VJ} factors

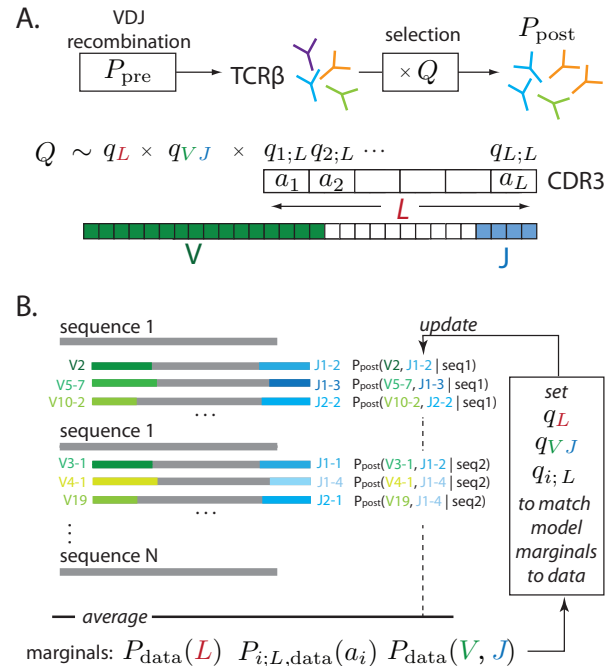


Fig. 1. A. T-cell receptor β chain sequences are formed during VDJ recombination. Sequences from this probability distribution, described by P_{pre} , are then selected with a factor Q defined for each sequence, resulting in the observed P_{post} distribution of receptor sequences. Selection is assumed to act independently on the V and J genes, the length of the CDR3 region and each of the amino acids, a_i , therein. **B.** A schematic of the fitting procedure: the parameters are set so that P_{post} fits the marginal frequencies of amino acids at each position, the distribution of CDR3 lengths and VJ gene choices. Since the latter is not known unambiguously from the observed sequences, it is estimated probabilistically using the model itself in an iterative procedure.

of Eq. 1. We use an expectation maximization algorithm (EM) that iteratively modifies the q 's until the observed marginal frequencies (for CDR3 length, amino acid usage as a function of CDR3 position, and VJ usage) in the data match those implied by the model distribution Eq. 1 (the pre-selection distribution P_{pre} being taken as a fixed, known, input). The procedure is schematically depicted in Fig. 1B (see SI for full details).

One important assumption of the model is that selection factors act independently of each other on the sequence. Consequently, while the model is fit only to single point marginal frequencies, and not to pairwise frequencies. To check the validity of this assumption, we plot the correlation functions of amino acid pairs in the model post-selection repertoire vs

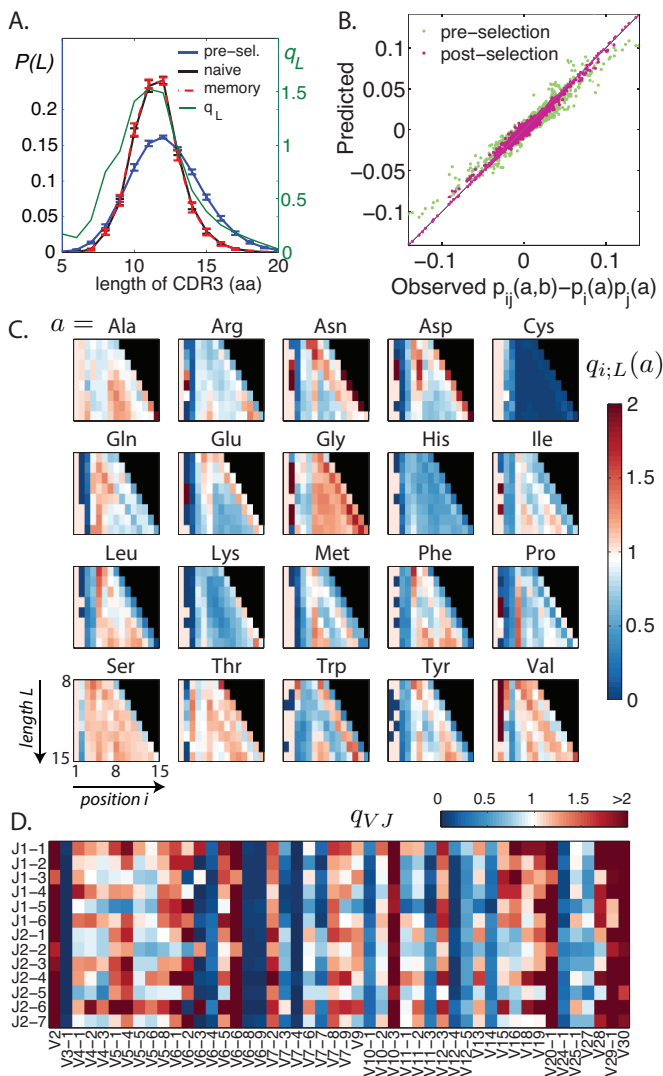


Fig. 2. **A.** CDR3 length distributions, pre- and post-selection and the length selection factor q_L (green). Selection makes the length-distribution of CDR3 regions in the pre-selection repertoire more peaked for the naive and memory repertoires (overlapping). Error bars show standard deviation over 9 individuals. **B.** Comparison between data and model of the connected pairwise correlation functions, which were not fitted by our model. The excellent agreement validates the inference procedure. As a control, the prediction from the pre-selection model (in gray) does not agree with the data as well. **C.** Values of the inferred amino-acid selection factors for each amino acid, ordered by length of the CDR3 region (ordinate) and position in the region (abscissa). **D.** Values of the VJ gene selection factors.

the observed naive ones (Fig 2B). These pairwise correlations are well predicted, even though they are not model inputs. It is also noteworthy that they are nonzero, even though the selection model does not take into account the possibility of interactions in the selection factors $q_{i:L}$. This is because the pre-selection distribution does not factorize over amino acids in the CDR3 region, and has correlations of its own, as shown by the green points of Fig 2B (note that these pre-selection correlations do not agree well with those observed in the post-selection data).

Another assumption of our model is that selection acts at the level of the amino acid sequence, regardless of the underlying codons. To test this, we learned more general models where a represented one of the possible 61 codons, instead of one of the 20 amino acids. We found that codons coding for the same residue had similar selection factors (see Fig. S2), except near the edges of the CDR3 where amino acids may actually come from genomic V and J segments and reflect their codon biases.

To compare the different donors, we learned a distinct model for each donor, as well as a “universal” model for all sequences of a given type from all donors taken together (see SI for details). We also learned models from random subsets of the sequence dataset to assess the effects of low-number statistical noise.

Results

Characteristics of selection and repertoire variability. The length, single-residue, and VJ selection factors, learned from the naive datasets of all donors taken together, are presented

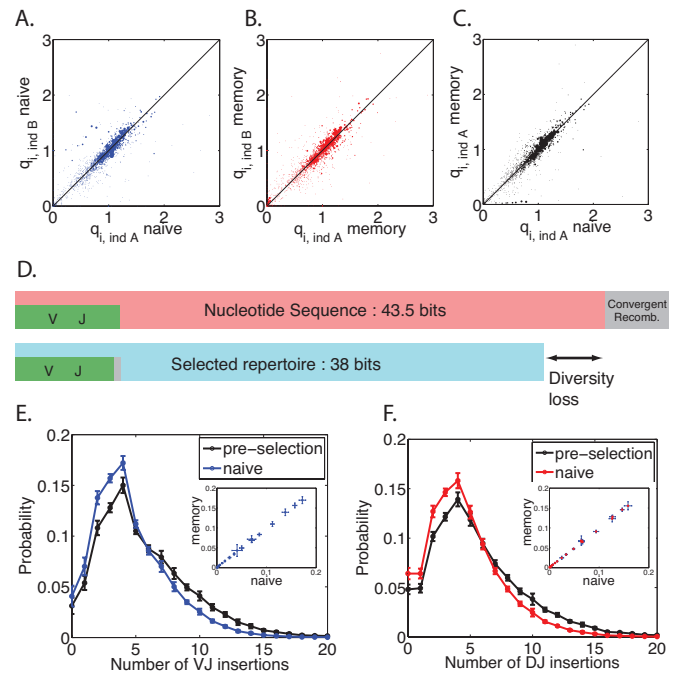


Fig. 3. **A.-C.** Variability between repertoires. The scatter between $q_{i:L}$ selection factors between two sample individuals A and B for naive (**A**) and memory repertoires (**B**) compared to that of memory and naive repertoires for the same individual (**C**) shows great similarity between them. See also Fig. S4. **D.** The entropy of the pre-selection repertoire (top) is reduced in the post-selection repertoire (bottom). **E.-F.** Distribution of VJ (**E**) and DJ (**F**) insertions in the pre-selection and naive repertoires shows elimination of long insertions. Error bars show standard deviations over 9 donors. The insertion distributions for the memory repertoire are the same as for the naive repertoire (see scatter plots in insets).

in Fig. 2A,C,D. The $q_{V,J}$ distribution shows that the different V and J genes are subject to a wide range of selection factors (note that these factors act in addition to the quite varied gene segment usage probabilities in $P_{pre}(\vec{\sigma})$). We looked for correlations between the selection factors $q_{i:L}(a)$ on amino acids and a variety of amino-acid biochemical properties [17]: hydrophobicity, charge, pH, polarity, volume, and propensities to be found in α or β structures, in turns, at the surface of a binding interface, on the rim or in the core [18] (see SI for details and references). We found no significant correlations, save for a negative correlation with amino acid volume and α helix association, as well as a positive correlation with the propensities to be in turns or in the core of an interacting complex (Fig. S7).

To estimate differences between datasets, we calculated the correlation coefficients between the logs of the $q_{V,J}$ and $q_{i:L}(a)$ selection factors (see Fig. S4). Comparing naive *vs.* naive, memory *vs.* memory or naive *vs.* memory between donors (see Fig. 3A-C for an example for $q_{i:L}$, and Fig. S3 for $q_{V,J}$) gave correlation coefficients of ≈ 0.9 in $\log q_{i:L}$, while the naive *vs.* memory repertoires of the same donor gave 0.95. To get a lower bound on small-number statistical noise, we also compared the factors inferred from artificial datasets obtained by randomly shuffling sequences between donors (see SI), yielding an average correlation coefficient of 0.98. Repeating the analysis for $\log q_{V,J}$, we found correlation coefficients of ≈ 0.8 between datasets of different donors, 0.84 for the naive and memory dataset of the same donor, all of which must be compared to 0.94 obtained between shuffled datasets. Thus, the observed variability between donors of $q_{i:L}$ and $q_{V,J}$ are small, and consistent with their expected statistical variability.

We use Shannon entropy, $S = -\sum_{\vec{\sigma}} P_{post}(\vec{\sigma}) \log_2 P_{post}(\vec{\sigma})$, to quantify the diversity of the naive and memory distributions. Entropy is a diversity measure that accounts for non-uniformity of the distribution and is additive in independent components. Since $S = \log_2 \Omega$ when there are Ω equally likely

outcomes, the diversity index 2^S can be viewed as an effective number of states. The entropy of the naive repertoire according to the model is 38 bits (corresponding to a diversity of $\sim 3.0 \cdot 10^{11}$) down from 43.5 bits in the primitive, pre-selection repertoire (Fig. 3D). This is a reduction of ~ 6 bits, or 50-fold in diversity. The majority of the reduction comes from insertions and deletions, which accounted for most of the diversity in the pre-selection repertoire. The entropies of the memory and naive repertoires are the same, indicating that selection in the periphery does not further reduce diversity.

Knowing the post-selection distribution of sequences, we can ask how different features of the recombination scenario fare in the face of selection. This does not imply that selection acts on the scenarios themselves—it acts on the final product—but it is an *a posteriori* assessment of the fitness of particular rearrangements. For example, the distributions of insertions at VD and DJ junctions in the post-selection ensemble have shorter tails (Fig. 3E-F), while the distribution of deletions at the junctions seems little affected by selection (Fig. S5), although large numbers of deletions are selected against.

Selection factor as a measure of fitness. The selection factor Q is a proxy for the probability of a particular sequence to be selected or amplified, and sequences with large Q values should thus be enriched in the observed dataset. To test this, we consider the distributions of Q both in the pre-selection model, $P_{pre}(Q)$, and in the dataset from which Q was learned, $P_{data}(Q)$ (insets of Fig. 4; see SI Text for details on how the distributions are calculated when V and J are hidden variables). This approach is very similar to the one used by Mustonen et al. [19, 20] to characterize the fitness landscape of transcription factor binding sites.

By construction, the distribution of Q in the post-selection model satisfies exactly $P_{post}(Q) = Q P_{pre}(Q)$. In Fig. 4A-B we plot the ratio $P_{data}(Q)/P_{pre}(Q)$ as a function of Q , both for the naive and memory models learned from all donors. We observe that for $Q \leq 5$, *i.e.* for $> 90\%$ of sequences, this ratio

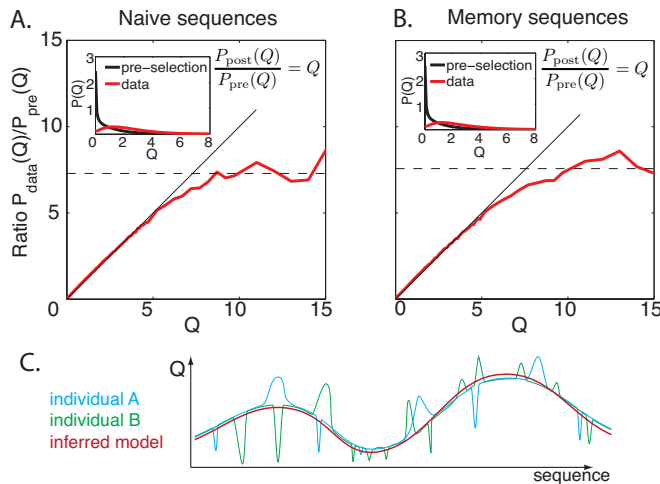


Fig. 4. Probability of passing selection. **A.-B.** Ratio of the distributions of sequence-wide selection factors Q between the observed sequences and the pre-selection ensemble (red line), plotted as a function of Q for naive (**A.**) and memory (**B.**) repertoires. The model prediction $P_{post}(Q)/P_{pre}(Q) = Q$ is shown in black, and the pre-selection and observed distributions of Q are shown in the insets. The selection ratio saturate around ≈ 7 , which may be interpreted as the maximum probability of being selected. Naive and memory repertoires show similar behaviors. **C.** A cartoon of the effective selection landscape captured by our model (red line). Our method does not capture localized selection pressures (such as avoiding self) specific to each individual, but captures general global properties.

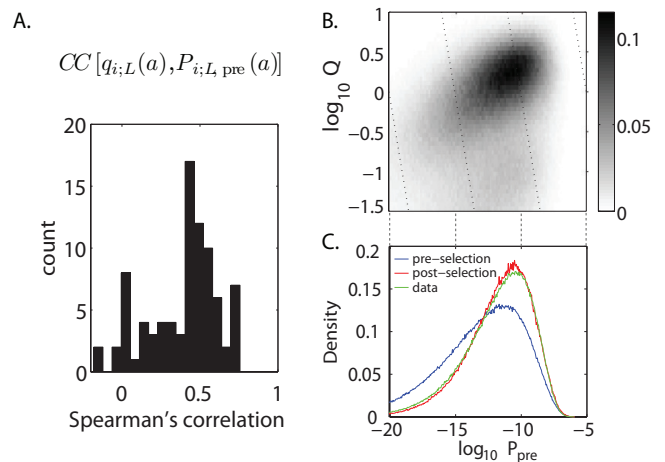


Fig. 5. Correlations between the pre- and post-selection repertoires. **A.** A histogram of Spearman's correlation coefficient values between the $q_{i:L}(a)$ selection factors in the CDR3 region and their generation probabilities $P_{i:L,pre}(a)$ for all i, L shows an abundance of positive correlations. **B.** Heatmap of the joint distribution of the pre-selection probability distribution P_{pre} and selection factors Q for each sequence shows the two quantities are correlated. **C.** Sequences in the observed, selected repertoire (green line) had a higher probability to have been generated by recombination than unselected sequences (blue line). Agreement between the post-selection model (red line) and data distribution (green line) is a validation of the model.

is exactly equal to Q — a validation of our model prediction at the sequence-wide level. For larger values of Q however, this ratio saturates to around $Q_{\max} \approx 7$.

This plateau may be viewed as a limiting value, above which selection is insensitive to Q . A similar plateau was observed in the fitness of transcription factor binding sites below a certain binding energy [20]. In the case considered here, the plateau can be rationalized if we assume that Q is proportional to the probability for a sequence to be selected, $P_{\text{sel}}(\vec{\sigma}) = \alpha Q(\vec{\sigma})$. Since P_{sel} cannot exceed one, Q cannot exceed α^{-1} . The average probability of selection is given by $\sum_{\vec{\sigma}} P_{\text{pre}}(\vec{\sigma}) P_{\text{sel}}(\vec{\sigma}) = \alpha$. The observed plateau gives a lower bound to the true maximum of Q : $\alpha^{-1} \geq Q_{\max}$, and thus the average fraction of cells to pass selection satisfies $\alpha \leq 15\%$. This can be compared to estimates [2] for passing positive and negative thymic selection: 10–30% for positive selection only, and $\approx 5\%$ for both. This analysis only includes the β chain, and including the α chain could further reduce our estimate.

The saturation also seems to indicate that our model may be too simple to describe the very fit (high Q) sequences. Because of its fairly simple factorized structure, our model can only account for the coarse features of selection, and may not capture very individual-specific traits such as avoidance of the self (corresponding to $Q \ll 1$ in localized regions of the sequence space) or response to pathogens ($Q \gg 1$ for particular sequences). This individual-dependent ruggedness of the fitness landscape Q , schematized in Fig. 4C, is probably ignored by our description, and may be hard to model in general.

To check that the saturation does not affect our inference procedure, we relearned our model parameters from simulated data, where sequences were generated from P_{pre} and then selected with probability $\min(Q/Q_{\max}, 1)$ (see SI Text for de-

tails), and we found that the model was correctly recovered (Fig. S6).

Natural selection anticipates somatic selection. Comparing the pre- and post-selection length distributions in Fig. 2A shows that the CDR3 lengths that were the most probable to be produced by recombination are also more likely to be selected. Formally, Spearman’s rank correlation coefficient between $P_{\text{pre}}(L)$ and q_L is 0.76, showing good correlation between the probability of a CDR3 length and the corresponding selection factor. We asked whether this correlation was also present in the other sequence features. The histogram of Spearman’s correlation between the selection factors $q_{i;L}(a)$ and the pre-selection amino-acid usage $P_{i;L,\text{pre}}(a)$ for different lengths and positions (i, L) (Fig. 5A) shows a clear majority of positive correlations. Likewise, the selection factors q_{VJ} are positively correlated with the pre-selection VJ usage $P_{VJ,\text{pre}}$ (Spearman’s rank correlation 0.3, $p < 2 \cdot 10^{-20}$).

The correlations observed for each particular feature of the sequence (CDR3 length, amino acid composition and VJ usage) combine to create a global correlation between the probability $P_{\text{pre}}(\vec{\sigma})$ that a sequence $\vec{\sigma}$ was generated by recombination, and its propensity $Q(\vec{\sigma})$ to be selected (Spearman’s rank correlation 0.4, $p = 0$, see Fig. 5B). Consistent with this observation, the post-selection repertoire is enriched in sequences that have a high probability $P_{\text{pre}}(\vec{\sigma})$ to be produced by recombination (Fig. 5C). This enrichment is well predicted by the model, providing another validation of its predictions at the sequence-wide level.

Taken together, these results suggest that the mechanism of VDJ recombination (including insertions and deletions) has evolved to preferentially produce sequences that are more likely to be selected by thymic or peripheral selection.

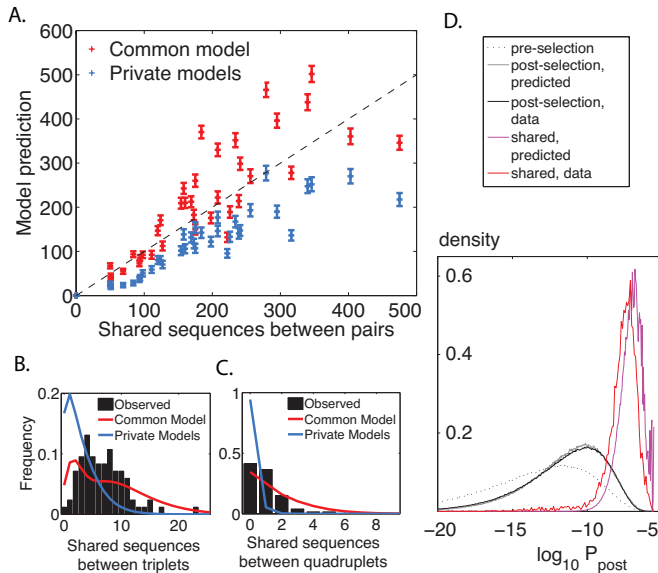


Fig. 6. Shared sequences between individuals. **A.** The mean number of shared sequences between any pair of individuals compared to the number expected by chance (model prediction) for one common model for all individuals (red crosses) and private models learned independently for each individual (blue crosses). Error bars are standard deviations from distributions over pairs. The distribution of shared sequences between triplets (**B.**) and quadruplets (**C.**) of individuals for the data (black histogram), from common (red line) and private models (blue line). **D.** The shared sequences are most likely to be generated and selected: comparison of the P_{post} post-selection distribution for sequences from the pre-selection (dotted line), and post-selection repertoires (according to the model in gray, and to the data in black), as well as the sequences shared by at least two donors (model prediction in magenta, data in red).

Shared sequences between individuals. The observation of unique sequences that are shared between different donors has suggested that these sequences make up a “public” repertoire common to many individuals, formed through convergent evolution or a common source. However, it is also possible that these common sequences are just statistically more frequent, and are likely to be randomly recombined in two individuals independently, as previously discussed by Venturi *et al.* [21, 7, 6]. In other words, public sequences could just be chance events. Here we revisit this question by asking whether the number of observed shared sequences between individuals is consistent with random choice from our inferred sequence distribution P_{post} .

We estimated the expected number of shared sequences between groups of donors in two ways: (i) by assuming that each donor α had its own “private” model learned from his own sequences or (ii) by assuming that sequences are drawn from a “universal” model learned from all sequences together. While the latter ignores small yet perhaps significant differences between the donors, the former may exaggerate them where statistics are poor. For details on how these estimates are obtained from the models, we refer the reader to the SI Text. In Fig. 6A we plot, for each pair of donors, the expected number of shared nucleotide sequences in their naive repertoires under assumptions (i) and (ii), versus the observed number. The number is well predicted under both assumptions, the universal model assumption giving a slight overestimate, and the private model giving a slight underestimate. We repeat the analysis for sequences that are observed to be common to at least three or at least four donors (Fig. 6B-C). The universal model predicts their number better than the private models, although it still slightly overestimates it.

These results suggest that shared sequences are indeed the result of pure chance. If that is so, shared sequences should have a higher occurrence probability than average; specifically, the model predicts that the sequences that are shared between at least two donors are distributed according to $P_{\text{post}}(\bar{\sigma})^2$ (see SI Text). We test this by plotting the distribution of P_{post} for regular sequences, as well as for pairwise-shared sequences, according to the model and in the naive datasets (Fig. 6D), and find excellent agreement. In general, sequences that are shared between at least n individuals by chance should be distributed according to $P_{\text{post}}(\bar{\sigma})^n$. For triplets and quadruplets, this model prediction is not as well verified (see Fig. S8). This discrepancy may be explained by the fact that such sequences are outliers with very high occurrence probabilities, and may not be well captured by the model, which was learned on typical sequences.

We repeated these analyses for sequences shared between the memory repertoires of different individuals, with very similar conclusions, except for donors 2 and 3, and donors 2 and 7, who shared many more sequences than expected by chance (see Fig. S9). We conclude that the vast majority of shared sequences occur by chance, and are well predicted by our model of random recombination and selection.

Discussion

We have introduced and calculated a selection factor $Q(\bar{\sigma})$ that serves as a measure of selection acting on a given receptor sequence $\bar{\sigma}$ in the somatic evolution of the immune repertoire. Our approach accounts for the fact that the pre-selection probabilities of sequences vary over orders of magnitude.

Using this measure, we show that the observed repertoires have undergone significant selection starting from the initial repertoire produced by VDJ recombination. We find little difference between the naive and memory repertoires, in agreement with recent findings showing no correlation between receptor and T-cell fate [22], as well as between the repertoires of different donors. This is perhaps surprising, because the donors have distinct HLA types (which determine the interaction between T-cell receptors and peptide-MHC complexes), and we could expect their positive and negative selective pressures to be markedly different. Besides, memory sequences have undergone an additional layer of selection compared to the naive ones—recognizing a pathogen—and we could also expect to see different signatures of selection there. A possible interpretation is that our model only captures coarse and uni-

versal features of selection related to the general fitness of receptors, and not the fine-grained, individual-specific selective pressures such as avoidance of the self, as illustrated schematically in Fig. 4C. In other words, our selection factors may “smooth out” the complex landscapes of specific repertoires and fail to capture their rough local properties, such as would be expected from specific epitopes that would correspond to very tall peaks or deep valleys in the landscape of selection factors. To really probe these specific deep valleys, we need to develop methods based on accurate sequence counts. Another interesting future direction would be to see whether at this global level the signatures of selection are similar between (relatively) isolated populations. Lastly, comparing data from different species (mice, fish), in particular where inbred individuals with the same HLA type can be compared, would be an interesting avenue for addressing these issues.

Our results suggest that natural selection has refined the generation process over evolutionary time scales to produce a pre-selection repertoire that anticipates the actions of selection. Sequences that are likely to be eliminated and fail selection are not very likely to be produced in the first place. Because of this “rich become richer” effect, the diversity of the repertoire is significantly reduced by selection, by a 50-fold factor in terms of diversity index. This does not mean that only 2% of the sequences pass selection. In fact, our results are consistent with as much as 15% of sequences passing selection. This apparent paradox is resolved by noting that selection, by keeping clones that were likely to be generated, get rids of very rare clones that contributed to the large initial diversity.

Although we did observe sequences that were present in the repertoires of different donors, we showed using our model that their number was broadly compatible with that expected by pure chance. This suggests that the “public” part of the repertoire is made of sequences that are just more likely to be randomly produced and selected.

To summarize, our work clearly shows that thymic selection and later peripheral selection modify the form of the generated repertoire. Our work is a starting point for a description of a mechanism of the two processes.

Acknowledgements. The work of YE, TM and AW was supported in part by grant ERCStG n. 306312. The work of CC was supported in part by NSF grants PHY-0957573 and PHY-1305525 and by W.M. Keck Foundation Award dated 12/15/09.

- Murugan A, Mora T, Walczak AM, Callan CG (2012) Statistical inference of the generation probability of t-cell receptors from sequence repertoires. *Proceedings of the National Academy of Sciences* 109:16161–16166.
- Janeway C (2005) *Immunobiology, the immune system in health and disease* (Garland Pub).
- Weinstein JA, Jiang N, White RA, Fisher DS, Quake SR (2009) High-throughput sequencing of the zebrafish antibody repertoire. *Science* 324:807–10.
- Ndifon W, et al. (2012) Chromatin conformation governs t-cell receptor j gene segment usage. *Proc Natl Acad Sci USA* 109:15865–70.
- Mora T, Walczak AM, Bialek W, Callan CG (2010) Maximum entropy models for antibody diversity. *Proceedings of the National Academy of Sciences of the United States of America* 107:5405–5410.
- Quigley MF, et al. (2010) Convergent recombination shapes the clonotypic landscape of the naive t-cell repertoire. *Proc Natl Acad Sci USA* 107:19414–9.
- Venturi V, Price DA, Douek DC, Davenport MP (2008) The molecular basis for public t-cell responses? *Nat Rev Immunol* 8:231–8.
- Baum PD, Venturi V, Price DA (2012) Wrestling with the repertoire: The promise and perils of next generation sequencing for antigen receptors. *Eur. J. Immunol.* 42:2834–2839.
- Six A, et al. (2013) The past, present, and future of immune repertoire biology – the rise of next-generation repertoire analysis. *Front. Immunol.* 4:1–16.
- Robins H (2013) Immunosequencing: applications of immune repertoire deep sequencing. *Current Opinion in Immunology* 25:646–52.
- Yates AJ (2014) Theories and quantification of thymic selection. *Front. Immunol.* 5:13.
- Jameson SC (2002) Maintaining the norm: T-cell homeostasis. *Nature Reviews Immunology* 2:547–56.
- Detours V, Mehr R, Perelson AS (1999) A quantitative theory of affinity-driven t cell repertoire selection. *J Theor Biol* 200:389–403.
- 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 USA* 105:16671–6.
- Robins HS, et al. (2009) Comprehensive assessment of T-cell receptor beta-chain diversity in alphabeta T cells. *Blood* 114:4099–4107.
- Robins H, et al. (2010) Overlap and effective size of the human CD8+ T cell receptor repertoire. *Science translational medicine* 2:47ra64.
- Stryer L, Berg JM, Tymoczko JL (2002) *Biochemistry*, 5th edition (W.H. Freeman & Co Ltd) Vol. 5th edition.
- Martin J, Lavery R (2012) Arbitrary protein-protein docking targets biologically relevant interfaces. *BMC Biophysics* 1:7.
- Mustonen V, Lässig M (2005) Evolutionary population genetics of promoters: predicting binding sites and functional phylogenies. *Proc Natl Acad Sci USA* 102:15936–41.
- Mustonen V, Kinney J, Callan CG, Lässig M (2008) Energy-dependent fitness: a quantitative model for the evolution of yeast transcription factor binding sites. *Proc Natl Acad Sci USA* 105:12376–81.

21. Venturi V, et al. (2006) Sharing of t cell receptors in antigen-specific responses is driven by convergent recombination. *Proc Natl Acad Sci USA* 103:18691–6.
22. Wang C, et al. (2010) High throughput sequencing reveals a complex pattern of dynamic interrelationships among human t cell subsets. *Proc Natl Acad Sci USA* 107:1518–23.