



Inferring the T cell repertoire dynamics of healthy individuals

Meriem Bensouda Koraichi^a, Silvia Ferri^{a,b}, Aleksandra M. Walczak^{a,1} , and Thierry Mora^{a,1}

Edited by Arup Chakraborty, Massachusetts Institute of Technology, Cambridge, MA; received May 1, 2022; accepted December 13, 2022

The adaptive immune system is a diverse ecosystem that responds to pathogens by selecting cells with specific receptors. While clonal expansion in response to particular immune challenges has been extensively studied, we do not know the neutral dynamics that drive the immune system in the absence of strong stimuli. Here, we learn the parameters that underlie the clonal dynamics of the T cell repertoire in healthy individuals of different ages, by applying Bayesian inference to longitudinal immune repertoire sequencing (RepSeq) data. Quantifying the experimental noise accurately for a given RepSeq technique allows us to disentangle real changes in clonal frequencies from noise. We find that the data are consistent with clone sizes following a geometric Brownian motion and show that its predicted steady state is in quantitative agreement with the observed power-law behavior of the clone-size distribution. The inferred turnover time scale of the repertoire increases with patient age and depends on the clone size in some individuals.

immune repertoire dynamics | T-cell receptor | stochastic evolution | longitudinal tracking | statistical inference

1. Introduction

The adaptive immune system protects us from many infections including those caused by pathogenic challenges that did not exist when we were born. This amazing plasticity is encoded, in part, in a diverse repertoire of T cells carrying surface receptors capable of recognizing different antigens, which trigger an immune response. About 10^8 new T cells are estimated to be generated and enter the periphery in human adults every day (1, 2), where they undergo specific proliferation due to antigen stimulation but also nonspecific divisions (3, 4) and death. These processes together result in clone sizes of different T cells that differ over a few orders of magnitude, forming long-tailed distributions (5, 6). The total number of different T cell clones is estimated between 10^8 and 10^{10} (7–9). Qualitatively describing the T cell clonal dynamics in the periphery is important for predicting long- as well as short-term immune response and to understand the maintenance of immune memory.

A lot of effort has been put into describing antigen-specific response and memory formation (10–12). At any given time point, the majority of the T cell repertoire is not always directly involved in fighting the current antigenic challenge. Yet, processes such as homeostasis (3) and unspecific signals in both naive and memory subrepertoires result in frequency changes of background clones. Many first-principles models of naive T cell dynamics have been proposed to study the balance between thymic output and peripheral proliferation and death (2, 4, 13–15). The role of competition for antigens between T cells has been pointed out (16) as well as the effect of cross-reactivity (17) (the ability for one T cell to recognize different antigens), the relative size of a primary vs. secondary response to similar antigens (18), or the effect of heritable changes affecting the homeostatic rate of thymic exports (19). These studies highlight the importance of the naive repertoire of clonal expansions that are not necessarily linked to specific challenges. While these models were instrumental in advancing our understanding of bulk repertoire dynamics, and allowed for the interpretation of deuterated water and bromide staining experiments that describe cellular lifetimes (20), the class of models that are consistent with the data is still large and unexplored.

Thanks to advances in immune repertoire sequencing (RepSeq) (21–23), dynamical models can now be assessed directly against repertoire data at the clonal level. RepSeq experiments isolate and sequence the T cell receptors (TCRs) in a blood sample of individuals. By counting reads with the same TCR sequence, one can estimate the frequency of the corresponding clone (defined as the set of cells with the same receptor) in blood. Even single repertoire snapshots can be informative about the dynamics: the distribution of clone sizes follows a power law (6, 24–27), in accord with proposed models of stochastic growth and death (5). Taking samples from the same individual at

Significance

Immune repertoires carry a diverse set of antigen-specific receptors that protect us from disease and keep the memory of past infections. How stable are those repertoires over time? We combine longitudinal repertoire sequencing data, stochastic evolutionary modeling, and tailored inference techniques to determine the characteristic time scales of T cell repertoire evolution. These timescales are in quantitative agreement with long tails observed in the distribution of clone sizes. We show that repertoire turnover slows down with age and is maximal in the largest clones of young individuals.

Author affiliations: ^aLaboratoire de physique de l'École normale supérieure, CNRS, Paris Sciences & Lettres (PSL) University, Sorbonne Université, and Université Paris Cité, Paris 75005, France; and ^bDipartimento di Fisica e Astronomia dell'Università di Bologna, Bologna 40126, Italy

Author contributions: A.M.W. and T.M. designed research; M.B.K. and S.F. performed research; M.B.K. contributed new reagents/analytic tools; and analyzed data; and M.B.K., A.M.W., and T.M. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2023 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹To whom correspondence may be addressed Email: aleksandra.walczak@phys.ens.fr or thierry.mora@phys.ens.fr.

This article contains supporting information online at <http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2207516120/-/DCSupplemental>.

Published January 20, 2023.

different time points allows for tracking the evolution of TCR clone sizes in time. The longitudinal experiments that have been performed in healthy donors (28, 29) suggest that the repertoire is relatively stable over years.

Our main goal in this article is to characterize the dynamics of the unstimulated background repertoire. We use an inverse approach to learn models of stochastic TCR clonal dynamics directly from data, collecting human TCR RepSeq datasets where we could identify at least two time points between which there was no reported specific acute antigenic stimulation (28–32). A key aspect of our method is the treatment of experimental noise, which confounds naive analyses of stochastic time traces. The method first quantifies both the sampling and natural biological noise thanks to replicate RepSeq experiments (33, 34) and then infers the parameters of a stochastic dynamical model to describe the trajectories of each TCR clone population in a healthy individual, i.e., who did not have medical conditions or known infections during the sampled interval. We explicitly show how correcting for noise allows us to robustly learn the underlying dynamics. A recent study (35) has investigated the formation of the T cell repertoire during development and its maintenance into adulthood. Here, we focus on healthy adult repertoires that are already shaped during the first years of an individual's life and ask how they evolve and are renewed. We extract clonal (and not cellular) turnover time scales and describe how these time scales depend on the person's age. Characterizing these baseline dynamics is an important step toward interpreting TCR dynamics in the presence of antigenic stimuli.

2. Results

A. Longitudinal Sampling of TCR Repertoires of Healthy Individuals. T cell repertoires are large ecosystems in which each species is a clone of T cells carrying the same TCR i formed by a unique pair of α and β chains. The dynamics of this system are characterized by the time course of the number of cells carrying each receptor, $n_i(t)$. This number can be accessed indirectly through TCR repertoire sequencing (RepSeq), obtained by sequencing the TCR of small samples of peripheral blood mononuclear cells (PBMCs), gives us a read count $\hat{n}_i(t)$ for a given chain at different time points (Fig. 1A). Note that our focus will be the dynamics of clones, rather than cells, because repertoire profiling cannot track individual cells. Because the two chains are not paired in the data, from here on, we define clones as collections of cells having the same α or β chain, which we will refer to as clonotypes. This approximation is justified by the low occurrence of TCRs that share one chain but not the other (36), even in specific disease-associated subrepertoires, which are more restricted (32).

We collected repertoire data from nine individuals P1 to P9, aged 18 to 57, sampled at various time points from 1 mo up to 3 y apart, with and without biological replicates. No specific acute antigenic stimulus was reported between time points defining these RepSeq longitudinal datasets. We do not rule out that these individuals may have undergone asymptomatic or mild infections during the duration of the studies, but we call them healthy since they have no particular medical condition and have not experienced medical interventions involving a major perturbation of their immune repertoires.

The repertoires of TCR β chains were sequenced for all samples and α chains for individual P6 only. The properties of the datasets, including their number of clones N_c , total read counts $N_r = \sum_{i=1}^{N_c} \hat{n}_i$, age, library preparation (from genomic

DNA or from mRNA), and chain, are summarized in Fig. 1B and in *SI Appendix, Table S1*. Because P3, P4, P6, and P9 were included in a vaccination study, they had received a shot of the YFV 17D yellow fever vaccine (P3, P4, and P6) or of the influenza vaccine (P9) 45 d prior to the first time point, after the decay of their T cell response, so we assume that the dynamics of vaccine-specific T do not affect much our analysis of the global repertoire.

A major challenge when analyzing RepSeq data is that the measured abundances $\hat{n}_i(t)$ only provide a noisy reflection of the true ones $n_i(t)$. Observed differences between datasets thus result from a combination of the repertoire dynamics and biological and experimental noise. The magnitude of that noise can be assessed by comparing the normalized clonotype frequencies $\hat{f} = \frac{\hat{n}_i}{N_r}$ between two biological replicates obtained at the same time point in the same individual (Fig. 1C, blue dots). By contrast, comparing those frequencies between two time points separated by 1 y (Fig. 1C, red dots) shows a larger dispersion and a slight overall decrease of clonotype frequencies. Our goal is to measure this difference quantitatively.

Another difficulty arises from the observation that clonotype frequencies are highly heterogeneous, with their distribution following a power law $P(\hat{n}) \propto \hat{n}^{-1-\alpha}$ spanning four orders of magnitude despite small expected deviations at low clone sizes due to PCR amplification noise (37) and extreme-value stochasticity in the large-clone tail. The exponent $\alpha \approx 1$ is largely invariant across individuals and time points (Fig. 1D), as previously reported (5, 6, 35). This implies that most clonotypes have a very low abundance and are thus particularly subject to sampling and experimental noise.

B. Mathematical Model of Stochastic Clonal Dynamics. The dynamics of T cell clones are driven by the proliferation and death of cells belonging to them. In addition, new clones with their distinct TCR are continually produced and released by the thymus, although the rate of thymic exports decays rapidly with age (1). Cell division, death, and introduction of new clones constitute the basis of our model (Fig. 2A). Cell division may be caused by antigen stimulation (both self and foreign) or by cytokine and growth factors, and cells die by lack of stimulation or by apoptotic signals. Even in the absence of strong and chronic antigenic stimuli, T cell clonotype abundances display stochastic trajectories due to either weak stimulation, repertoire homeostasis, and demographic fluctuations, produced by the stochastic effects of divisions and deaths in small populations. In addition, individuals may get mild infections over the course of months and years. Since these events are numerous and unknown, we model them by an effectively random net growth rate (divisions minus deaths). It can be shown (5, 38) that on time scales much longer than the typical resolution time of infections, each clonotype size $n_i(t)$ may then be modeled by a geometric Brownian motion (GBM). Its evolution is governed by an effective mean net growth, to which random fluctuations are added to account for bursts of proliferation and decay:

$$\frac{d \ln n_i(t)}{dt} = -\tau^{-1} + \theta^{-1/2} \eta_i(t), \quad [1]$$

where $\eta_i(t)$ is a clonotype-specific white noise of zero mean and unit amplitude. Note that the mean growth rate of clones, $-\tau^{-1}$, is typically negative. On average, each clone should decay to make room for new thymic exports because of homeostatic pressures that control the total number of cells. In this interpretation, τ is the typical decay time of each clone, which would evolve with

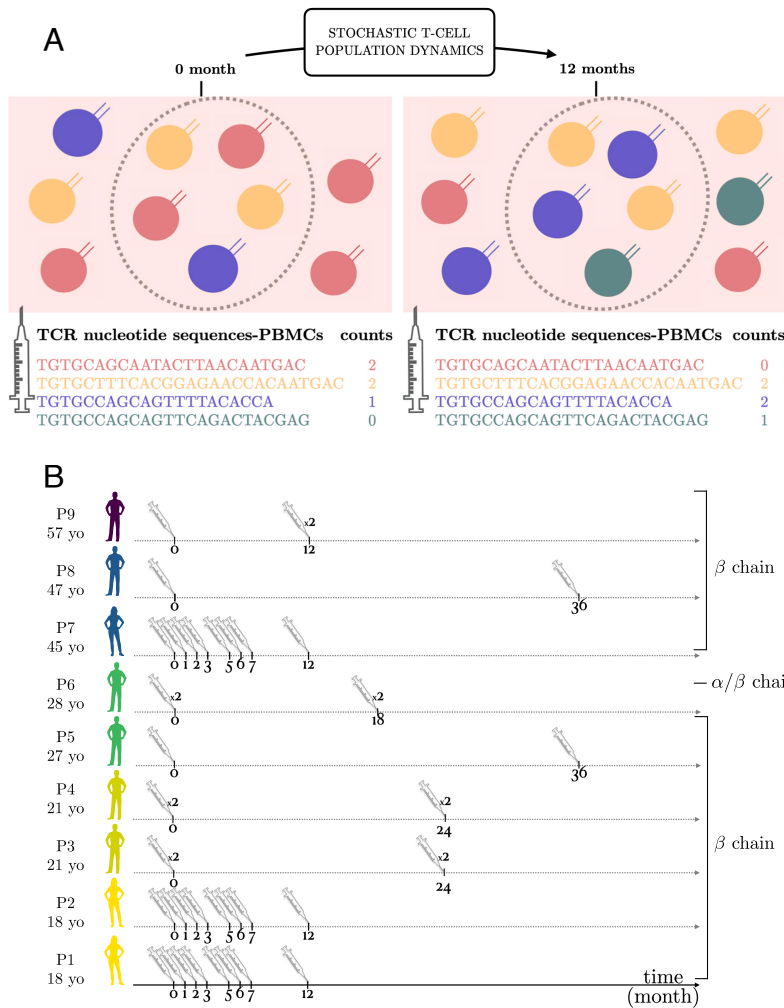


Fig. 1. Longitudinal tracking of T cell repertoires. (A) Experimental workflow. PBMCs from a healthy individual are extracted at two time points, and their TCR repertoire sequenced, yielding lists of clonotypes with count numbers corresponding to the number of individual measurements (or reads). The way in which the two sampled repertoires has changed between the two time points is predicted by a stochastic model of the dynamics of T cell clones. (B) Summary of the TCR α and β repertoire data used in this study. Included were nine individuals from five studies, aged 18 to 57, male and female. When available, replicate experiments are annotated with $\times 2$. Datasets were produced using two different sequencing technologies based on cDNA and gDNA. (C) Typical scatter plot of frequencies of TCR clones in two samples from the same individual P9. Blue: Two biological replicates obtained on the same day show the effect of experimental noise. Red: Two samples taken 1 y apart show a larger spread, resulting from a combination of real changes and noise. The goal of the analysis is to disentangle real changes from the noise. (D) Cumulative distributions of TCR frequencies, which follow a universal power law in all samples and donors, with exponent ≈ 1 .

time as $n_i(t) = n_i(0)e^{-t/\tau}$ in the absence of fluctuations, and is linked to its typical lifetime or turnover time $\sim \tau \ln(n_i(0))$. It should be distinguished from the decay rate of individual cells, which may be faster, as the clone may be sustained by a balance between cell deaths and divisions. Also recall that this is just an average; many clones do not decay, but instead undergo episodes of large growth and decay, as illustrated by simulations of Eq. 1 in Fig. 2B. The typical amplitude of these fluctuations grows with time as $\sqrt{t/\theta}$ (dashed lines). Thus, θ may be interpreted as the typical time it takes for a clone to rise or decay above or below the typical behavior by one log-unit.

In addition to being biologically motivated, the proposed dynamics have the desirable property that, in the presence of a constant rate of thymic exports, the distribution of clone sizes is predicted to evolve in time toward a perfect power law, $P(n) \propto n^{-1-\alpha}$, with exponent $\alpha = 2\theta/\tau$ given by twice the ratio of the two time scales of the model (5). This is illustrated in Fig. 2C on simulated repertoires at steady state and agrees well with the empirical distributions of Fig. 1D.

Our goal is to capture the parameters of these dynamics that are informative about the repertoire turnover timescales, while constraining the experimentally observed clone size frequency distribution. Our approach assumes that on the timescales of the analysis, we do not observe signals of strong and specific antigenic stimulation. It also ignores potential dependences on the size of the clone, which could be mediated by phenotypic differences between clones. This last assumption will be revisited later.

C. Model Inference. We estimate the parameters (θ , τ) of the dynamics in Eq. Eq. 1 from the observed clonotype abundance trajectories using a Bayesian approach for the posterior distribution of parameters given the data:

$$(\tau^*, \theta^*) = \arg \max_{\tau, \theta} \prod_{i=1}^{N_c} P(\hat{n}_i(t_1), \hat{n}_i(t_2) | \tau, \theta), \quad [2]$$

where t_1 and t_2 are the times of the two samples.

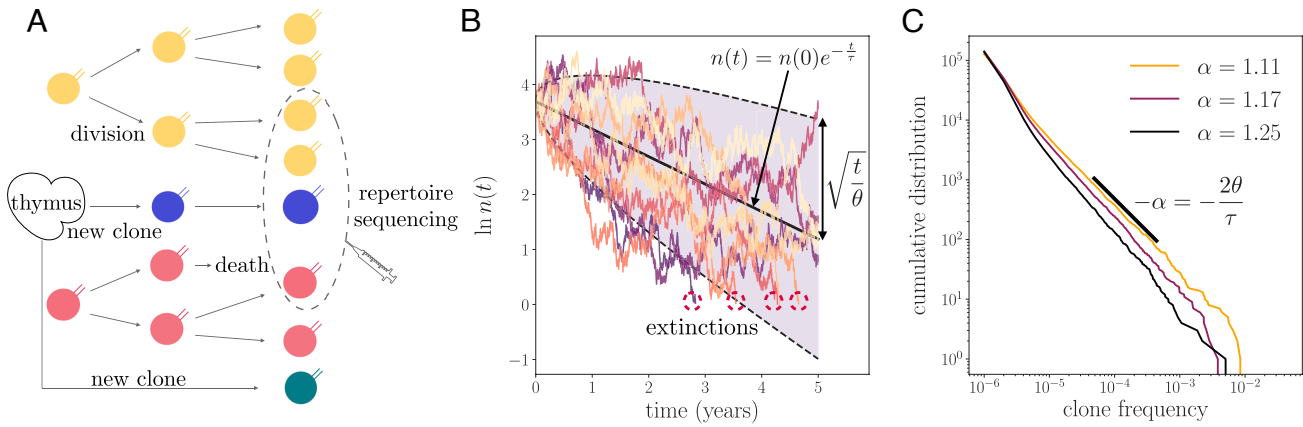


Fig. 2. Stochastic model of repertoire dynamics. (A) T cells are introduced in the peripheral immune system by thymic export, providing a source of new TCR clones. T cells belonging to a specific TCR clone (labeled by their color) divide and die depending on their interactions with the antigenic environment, increasing or reducing the abundance of its TCRs in the repertoire. This process is modeled by a geometric Brownian motion. (B) Example traces of TCR abundances simulated from the model Eq. 1 with $n(0) = 40$, with $\tau = 2$ y and $\theta = 1.11$ y. Clones that reach abundance < 1 go extinct (red circles). The typical trend is for clones to decay exponentially with time scale τ (black solid line). Stochastic events of growth and decay account for a broad variability of individual traces, whose magnitude grows as $\sqrt{t/\theta}$ with time (shaded area) in logarithmic scale. (C) Cumulative frequencies distributions of synthetic TCR clone abundances. The model predicts a power law of exponent $\alpha = 2\theta/\tau$. Different values of τ and θ were used to lead to different values of the exponent α . Parameters: $\tau = 2$ y, $N_{\text{cell}} = 10^{10}$, $n_0 = 40$.

We use two methods to learn the model parameters: naive inference and full inference. The naive inference assumes that the empirical abundances faithfully represent the real clonal abundances n_i through a simple proportionality rule, $\hat{n}_i \approx (N_r/N_{\text{cell}})n_i$, where $N_{\text{cell}} = \sum_i n_i$ is the total number of T cells in the body. In practice, we work with clonotype frequencies $\hat{f}_i = \hat{n}_i/N_r$ and $f_i = n_i/N_{\text{cell}}$, so that this assumption becomes $\hat{f}_i = f_i$. Further assuming that the total number of cells N_{cell} is approximately constant in time at steady state, \hat{f}_i is then governed by the same equation Eq. 1 as n_i . We take advantage of the closed solution available for the propagator associated with the GBM, Eq. 4, to maximize the log-likelihood (*Methods*). This maximization is equivalent to plotting a histogram of the change in log-frequencies between the two time points and simply read off τ^{-1} and θ^{-1} as the negative mean and the variance of the distribution divided by $t = t_2 - t_1$ (Fig. 3A), consistent with their biological interpretation.

The full inference incorporates the fact that the observed clonotype abundances are contaminated by biological (mRNA expression) and experimental noise sources (sequencing errors, stochastic PCR amplification, and sampling), which means that they do not correspond exactly to the clonotype abundances. To give a sense of just the sampling issue, a PBMC sample of ~ 1 mL contains about 1 million T cells, yielding about 1 million reads. By comparison, the organism contains the order of 10^{11} T cells. TCR clonotype frequencies are thus extrapolated from observing a fraction $10^6/10^{11} \approx 10^{-5}$, or 0.001%, of the whole repertoire (9). In addition, not all cells are captured, and each cell may be represented by multiple reads, either through sequencing of multiple mRNA from the same cell or from PCR amplification, depending on the context. To address these sources of uncertainty, in the full inference approach, we introduce an error model (33) relating observed frequencies \hat{f} to their true value f probabilistically through the transfer function $P(\hat{f}|f)$ (Fig. 3B). We use the previously introduced software tool, NoisET (34), which learns such a noise model from replicate RepSeq experiments (*Methods*).

We applied NoisET to individuals P3, P4, P6, and P9 for whom replicates were available. The noise model assumes that

the read count \hat{n} of each clone is drawn from a negative binomial distribution, whose variance grows with the frequency as $\text{Var}(\hat{n}) = fN_r + a(fN_r)^b$, with two learnable parameters a, b . In addition, since true frequencies are unknown, we assume as a prior that frequencies are distributed according to a power law $\rho(f) \propto f^{-1-\alpha}$ with a cut-off $f > f_{\text{min}}$, with α and f_{min} two other parameters. These parameters are reported for all individuals and time points in *SI Appendix, Fig. S1*.

Once the noise model has been learned using NoisET, the likelihood of the data is computed by summing over the latent variables $f_1 = f_i(t_1)$ and $f_2 = f_i(t_2)$:

$$P(\hat{n}_i(t_1), \hat{n}_i(t_2)|\tau, \theta) = \iint_{f_{\text{min}}}^1 df_1 df_2 \rho(f_1)P(f_2|f_1, \tau, \theta) \times P(\hat{f}_1|f_2)P(\hat{f}_2|f_2), \quad [3]$$

where $P(f_2|f_1; \tau, \theta)$ is the propagator of the geometric Brownian motion Eq. 1, and $\hat{f}_j = \hat{n}_i(t_j)/N_r(t_j)$, $j = 1, 2$.

To explore the dependence of the τ and θ parameters on the frequency of clonotypes and to eliminate clones that are not seen at both time points, we can generalize the formulas above to include only clonotypes with frequencies larger than a specific threshold \hat{f}_{th} , which modifies the normalization of the maximum likelihood estimator (*Methods*).

D. Validation of the Inference Methods on Synthetic Data. We first test the naive and full model inference on simulated RepSeq samples. We simulate 10^{10} cells corresponding to $\sim 10^8$ synthetic longitudinal trajectories designed to mimic as closely as possible the features of the real repertoire data at time points 2 y apart. The initial size $n_i(t_1)$ of each clone is drawn from the steady-state distribution of the GBM with a constant source (*Methods*). Then, Eq. 1 is simulated between times t_1 and t_2 with an extinction condition when $n_i < 1$ and with a source of new clones whose rate of introduction is matched to the mean extinction rate (*Methods*). We varied the two timescales of the model, τ and θ , from months to years, while keeping $\alpha = 2\theta/\tau$ within the observed experimental range 1.1 to 1.25 (35).




Fig. 3. Inferring dynamical parameters from data. (A) Naive inference. The empirical clonotype frequencies in the RepSeq samples at the two time points, \hat{f}_1 and \hat{f}_2 , are treated as the true ones, f_1 and f_2 . We estimate the two parameters τ and θ from the average and SD over all observed TCR clones of the log-fold change in frequency between two time points, $\log(\hat{f}_2/\hat{f}_1)$, which the model predicts is distributed normally. The distribution of \hat{s} is shown in light red and the Gaussian fit in solid red. (B) Full inference. The empirical frequencies \hat{f} are modeled as noisy read-outs of the true ones f , through a probabilistic noise model. First, the noise model $P(\hat{f}|f)$ is inferred from replicate experiments such as shown in Fig. 1C. The inference procedure also learns the distribution of frequencies $\rho(f)$, assumed to follow a power law with adjustable exponent α and minimal frequency f_{\min} . Second, using the noise model, the parameters of the dynamical propagator $P(f_2|f_1, \tau, \theta)$ are inferred from two time points, where f_1 and f_2 are treated as latent variables and \hat{f}_1, \hat{f}_2 as observables, using a Maximum likelihood estimator. (C and D) Validation of naive and full inference models on synthetic data. Model parameters: $t_2 - t_1 = 2$ y, $\tau = 1, 2, 5, 10$ y, $\alpha = 1.11, 1.18, 1.25$, with all 12 combinations tested; number of cells $N_c = 10^{10}$; initial clone size $n_0 = 40$; the other parameters (number of clones, thymic output) are deduced assuming steady state (Methods). Sampling model: number of sampled reads $N_r = 10^6$; noise model parameters $a = 0.7$ and $b = 1.1$. Error bars are standard deviations over 10 simulations.

We model experimental sampling using a negative binomial distribution with variance parameters $a = 0.7$ and $b = 1.1$. Sequencing depth was set to $N_r = 10^6$ reads at both time points (we checked that asymmetric numbers of reads at each time point did not affect the results, SI Appendix, Fig. S2), resulting in $\sim 10^5$ sampled distinct clonotypes. For each set of parameters, we generated 10 longitudinal datasets to assess errors. We then performed the naive and full inference methods on these datasets, restricted to clones with $\hat{f}_1 \geq \hat{f}_{\text{th}}$ and $\hat{f}_2 > 0$, and compared the inferred values of τ and θ to the true ones (Fig. 3 C and D and SI Appendix, Table S2).

While the full inference (blue points) works for all values of the parameters and frequency threshold, the naive inference (red points) performs poorly for large values of the parameters. Increasing the cutoff frequency to $\hat{f}_{\text{th}} = 10^{-4}$ improves the naive inference by limiting the effect of the sampling noise, which is relatively smaller in large clones. For lower values of the threshold, the more numerous small clones dominate the

inference, yielding an erroneous estimate, underestimating large values and overestimating low values of the parameters. However, since the naive inference does not require replicates or a noise model, and is faster to implement, it provides a practical solution for learning τ and θ for large clones.

E. Analysis of Repertoires. We applied the full inference to longitudinal data sets of healthy individual TCR repertoires presented in Fig. 1 for which replicates were available, focusing on large enough clones ($\hat{f}_1 \geq \hat{f}_{\text{th}} = 10^{-5}$). With this cutoff, we limit experimental noise and focus mainly on memory clones since large clones are more likely to have arisen from expansion and belong to the memory pool (14).

For all individuals, the inferred values of τ and θ broadly follow the line defined by $\tau \approx 2\theta$ (Fig. 4A) corresponding to a predicted exponent of $\alpha \approx 2\theta/\tau = 1$ in the power law of the clone size distribution. This result is in agreement with the empirical observations of Fig. 1D. A more refined comparison

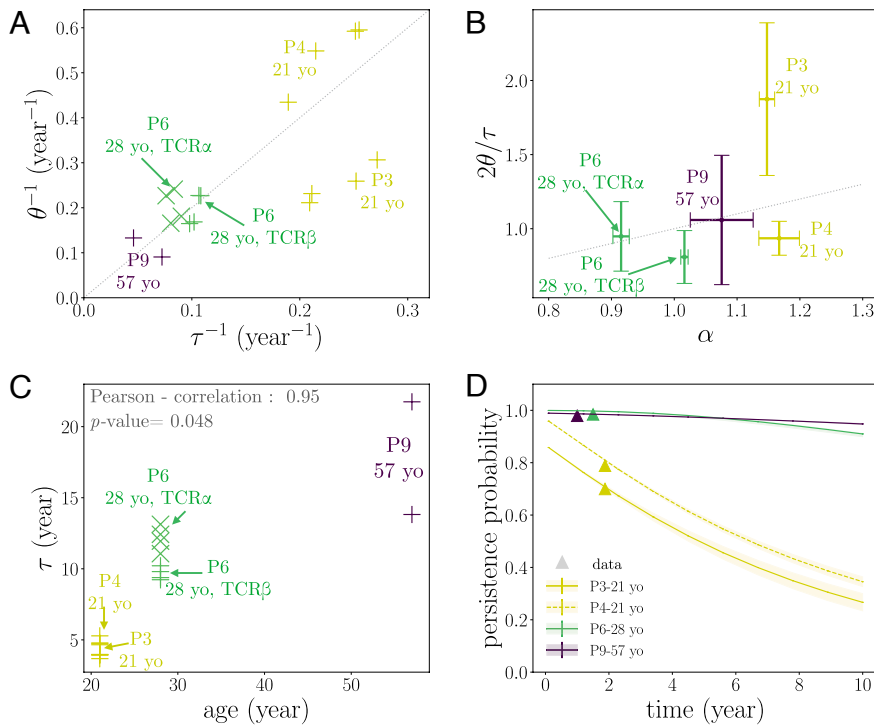


Fig. 4. Dynamical parameters of healthy TCR repertoires. (A) Typical decay rate τ^{-1} , and inverse fluctuation amplitude θ^{-1} for the five donors for whom replicates were available, as obtained using the full inference procedure with $f_{th} = 10^{-5}$. All donors but one are consistent with the relation $2\theta = \tau$, corresponding to $\alpha = 1$. Error bars are standard deviations over all combinations of the replicates at each time point. (B) Direct test of the prediction $\alpha = 2\theta/\tau$. Most values of α fall close to one, allowing for only a narrow range of tested values. Error bars on α show standard deviations across time points. (C) Turnover parameter τ as a function of donor age. (D) Probability for a clone detected at some time point with frequency $\hat{f} \geq 10^{-5}$ to be detected again at a later time point (with the experimental dataset size). Symbols are empirical estimates. The model predictions show excellent agreement. Error bars in (B–D) are propagated from (A)

of the predicted exponent, $2\theta/\tau$, with the one directly inferred from the distribution of clone sizes, α , gives consistent but noisy results (Fig. 4B), primarily because of the narrow range of values of α (0.9 to 1.2) and the small number of individuals.

We emphasize that the two inferred values of α use completely independent pieces of information. On the one hand, the exponent α describing the TCR frequency power-law distribution was learned using NoisET (34) from biological replicates taken at the same time point. On the other hand, the timescales forming the ratio $2\theta/\tau$ were learned using changes in individual clonal frequencies of TCR between RepSeq samples taken at two different time points, without assumptions about their distribution (except to calibrate the noise model, which has only a minor effect on the full inference and none on the naive inference).

Since our approach is probabilistic, it provides as a byproduct the posterior distribution of the fold change of individual clones (Methods). The average of this posterior over clones gives a direct visualization of the distribution of fold-changes in clonal frequencies, and agrees very well with the model propagator Eq. 4 (SI Appendix, Fig. S3), validating its consistency with the data.

The decay time τ increases with age, from a few years at age 21 to ~ 50 y at age 57 (Fig. 4C, Pearson's $\rho = 0.95$, $p = 0.047$, on the mean τ vs age of each individual). Since the ratio $2\theta/\tau$ is constrained to be ≈ 1 , this implies that the amplitude of the stochastic stimulations, θ^{-1} , decreases with age. The TCR repertoire is more dynamic, with faster clonal decay and thus turnover, for young individuals, who also have a larger rate of introduction of new TCR clones from the thymus than older individuals. At the same time, a decay time of ~ 20 years

at the age of 40 suggests that the repertoires of adults remain dynamic despite greatly reduced thymic output. Despite good concordance between results obtained across replicates, as well as between the α and β chain for P6 and between twins P3 and P4, these results rely only on four individuals and should be confirmed with a larger cohort.

For individual P6, both TCR α and TCR β RepSeq samples were available. We recover very similar dynamic parameters for both receptor chains (Fig. 4A and B). This justifies our hypothesis that the bulk sequencing of single chains captures well the dynamics of $\alpha\beta$ clonotypes.

For comparison, we also applied the naive inference procedure, which allows us to include all nine patients even when replicates are not available. This inference generally gave much larger rates τ^{-1} and θ^{-1} (SI Appendix, Fig. S4A), suggesting confounding effects of the noise on both parameters (reversion to the mean for τ^{-1} and larger variance for θ^{-1}). Indeed, results obtained for a larger value of the frequency threshold ($f_{th} = 10^{-4}$, SI Appendix, Fig. S4B) gave smaller values, and in better agreement with the age dependence.

To ask whether the clonal dynamics depended on the cell type, we separately analyzed the longitudinally sampled CD4 and CD8 repertoires of P6, the only individual for which such data were available. The clone size distribution of CD4 falls off with a larger exponent than that of CD8, meaning that its largest expanded clones are relatively smaller (SI Appendix, Fig. S5A), as already noted (28). We then applied the naive inference procedure with $f_{th} = 10^{-4}$ (since we did not have replicates for the CD4 and CD8 repertoires). The inference (SI Appendix, Fig. S5B) reveals that CD4 clones turn over more slowly than CD8 cells

(smaller τ^{-1}) but also have much smaller fluctuations in their sizes (smaller θ^{-1}). This result is consistent with a shorter tail of large clones and a larger α in CD4 than in CD8 (*SI Appendix, Fig. S5C*). The proximity of the bulk and CD8 results also suggests that the analysis of the full repertoire is dominated by the CD8 subpopulation thanks to its longer large-clone tail, although this should be confirmed on more individuals.

The inference results can be used to predict the persistence of clones, whose turnover has been discussed in the context of aging (28, 29, 35). For a given individual, we define persistence as the probability that a clone initially observed at frequency $\geq \hat{f}_{th} = 10^{-5}$ is resampled at a later time. This probability strongly depends on the rate of change of the TCR repertoire, and therefore on the age of the individual, as well as on the time interval between the two samples (Fig. 4D). We can estimate this persistence probability directly from data and compare it to the predictions of our inferred dynamics, showing excellent agreement. This analysis shows that even moderately large clones persist for many years and even decades in older individuals.

Our model assumes that clones have unique trajectories but that the statistical properties of these trajectories are uniform. However, because of their distinct histories and phenotypical compositions, clones may differ in those dynamical properties. To investigate that possibility, we asked whether the inferred time scales τ and θ depended on the value of the clone size threshold \hat{f}_{th} . Low values of \hat{f}_{th} mean that all clones are taken into account in the inference, while high values mean that we focus on the largest clones only. We found that the values of both τ^{-1} and θ^{-1} increase with the threshold (Fig. 5 A and B) for the younger individuals P3 and P4, suggesting that their large clones tend to be more dynamic, with a faster turnover. Therefore, while the overall trends reported in Fig. 4 are still correct, these results

imply that the model should be revisited to allow for frequency-dependent dynamics.

To measure the frequency dependence of the dynamic parameters more finely, we separately inferred τ and θ for clones sorted into contiguous intervals according to their initial count $\hat{n}(t_1)$. Both time scales showed an approximately linear dependency on the logarithm of the initial frequency (Fig. 5 C and D), although this dependency is visible only in the younger individuals P3 and P4. This confirms the observation that their large clones tend to have faster dynamics than small ones. Longitudinal samples from more donors would be needed to assess the individual heterogeneity and age dependence of these observations. The two time scales τ and θ vary in concert, so that their ratio remains approximately constant across frequencies (*SI Appendix, Fig. S6*). As we have argued before, this ratio is linked to the power-law exponent of the distribution of clone sizes at steady state. This exponent can be read off as the slope of that distribution on a double logarithmic scale, which is consistently observed to be constant in the data (Fig. 1D). Finally, we applied the naive inference procedure to learn the frequency-dependent dynamics of clones in all nine individuals. As expected, this inference yielded more noisy and less stable results than the full inference, especially at low frequencies for which noise is largest (*SI Appendix, Fig. S7*). The dependence of the inferred parameters on clonal frequency varies across individuals but confirms a picture in which older individuals have more stable large clones.

3. Discussion

The sizes of T cell clones change constantly throughout the lifetime of an individual, not only due to specific stimulation. Large clones have been reported in aging individuals (39) as well as in laboratory mice with no active exposure to pathogens (27).

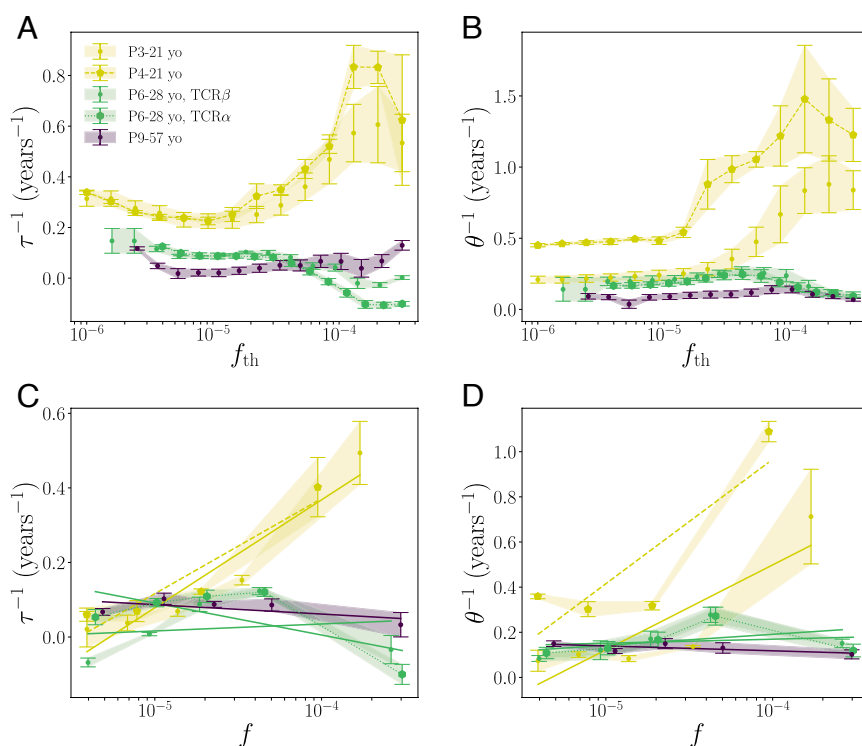


Fig. 5. Clonal dynamics are frequency dependent. (A and B) Results of the full inference as a function of the minimal frequency threshold \hat{f}_{th} for τ^{-1} and θ^{-1} . (C and D) Dynamical parameters as a function of clone frequency. The inference was performed on separate subsets of clones sorted by their frequency in intervals $n_{min} < \hat{n} \leq n_{max}$, with $n_{min,max}$ consecutive numbers in (2, 5, 10, 20, 100, ∞). Error bars are estimated as in Fig. 4.

We used data sampled on timescales of the order of a year from individuals who did not undergo any strong identified antigenic stimulations to learn the repertoire turnover dynamics. These dynamics include both random unstimulated T cell proliferation and death as well as asymptomatic or weakly symptomatic antigenic stimulation. We showed that a geometric Brownian motion correctly captures the clone dynamics. This model imposes strict relations that link the exponent of the TCR clone size distribution at steady state, α , with the parameters of the dynamics, τ , and θ . We showed that, for all individuals, we were able to predict the measured exponent $\alpha \approx 1$ from the inferred dynamical parameters, suggesting that geometric Brownian motion is a good description of the process. Although the actual timescales vary between individuals, with younger individuals having faster clone turnover dynamics than older individuals, their ratio is fixed. Indeed, as already noted on a larger cohort of individuals in ref. 35, the exponent of the power-law distribution of clone sizes does not depend on age. A limitation of this study is that the dependence between age and dynamics was shown on only four individuals. The naive inference, which does not correct for noise and is therefore less reliable, but allows for including all nine individuals, shows only a weaker correlation. In addition, some features appear to be individual-specific, such as the way the abundance of a clone affects its dynamics (Fig. 5 C and D).

The faster turnover in younger individuals is uncovered but not explained by our analysis. It can be linked to a larger thymic output rate (1), imposing a faster turnover. It could also be linked to more a rapid formation of new immune memories at a young age. We did not attempt to separately learn the dynamics of memory and naive pools since we did not have sorted longitudinal data for which abundance information could be trusted. While it is sometimes assumed that larger clones have a memory phenotype because they must have expanded, a recent study in mice has shown that naive clones can be large as well (14). It will be interesting to perform a separate analysis of carefully sorted naive and memory repertoires in the future using the method described here, especially for individuals of different ages. Our mathematical model assumes proliferation as a result of antigenic stimulations, which is the hallmark of the memory pool. Since naive cells are not supposed to proliferate in response to stimuli, we may expect its repertoire to be described by neutral ecological dynamics, where differences between clones are due to stochastic divisions and deaths (5, 14, 38). A difficulty for future analyses is that these dynamics lead to exponentially decaying clone size distributions, with very small clonal frequencies, which cannot be easily measured in human blood samples.

More generally, we expect clonal dynamics to be linked to the cellular phenotype, as our preliminary analysis showed for CD4 and CD8 cells. Phenotypes can be characterized with increasing resolution using single-cell expression data (40), which also provides paired TCR information (41). Future work combining longitudinal sampling with single-cell techniques could help explore the relationship between neutral clonal dynamics and cell type. Additionally, we know that TCRs with similar sequences form clusters that often respond to similar stimulants (17, 42), and methods are being developed to annotate repertoire with cluster membership (43) or specificity (44–48). As these annotations become comprehensive, one will be able to study the dynamics of specificity clusters and to assess the persistence of specific immune memories across different immune challenges.

Our current model is based on two effective parameters that describe the timescale for clone decay, τ , and the timescale of

random changes, θ . Two major assumptions underlie this model. First, it assumes that antigenic stimulation happens repeatedly on short time scales, so that its cumulated effects on longer time scales look like random fluctuations of the net growth rate. Testing this assumption would require longer time traces of the clonal dynamics, to look for memory effects in the clonal growth rates. Second, it assumes that dynamical properties do not depend on the clone size. As observed in Fig. 5, this assumption is only partially verified, with clear violations for 2 of the youngest donors, in which the larger clones display much faster dynamics than the smaller ones. While it might suggest that those large, dynamic clones are involved in some immune response, repertoire studies of vaccinations or infections have shown that responsive clones are typically not among the largest ones (31, 49). In addition, we checked that the top 100 β -chain clones of donors P3, P4, P6, and P9 had little overlap in their CDR3 with the VDJdb database of pathogen-associated TCRs (50) (*Methods*), with 15 matches in total (compared to 12.0 ± 3.4 for a random choice of 100 clones with at least five reads from each donor). The longitudinal analysis of larger cohorts with a broad age distribution would be required to investigate this effect in detail.

The decay time scales we infer range from a few years to 50 y, depending on the age of the individual. It has been shown that even sparsely sampled T cell repertoires can provide a fingerprint that uniquely identifies individuals (51). The stability of this immune fingerprint is guaranteed for tens of years, provided that the decay rate is of the order of years or more, as we showed here.

Direct measurements of T cell lifetimes using heavy water (10) give lifetimes of months for memory cells, to a few years for naive cells. These estimates are consistent with our findings: Our time scale τ is linked to the inverse of the net growth rate of the clone, which results from the balance between cell proliferation and death, while experiments based on heavy water measure the turnover of individual cells. For instance, memory cells are short lived but also divide rapidly to compensate for death, so that the size of memory clones remains stable. One may also want to compare our estimate with the previously reported persistence time of clonotypes believed to be of fetal origin, ≈ 37 y (52). This persistence time is not directly comparable to τ , which is the decay rate of the abundance of each clone, but it is similar to the characteristic decay of the persistence probability (Fig. 4D), which may be slower. Another caveat is that fetal clonotypes are also primarily naive and take up only a few percent of the repertoire, so that they may not be representative of the overall properties of the clonal dynamics.

Deciphering the baseline dynamics is an important step toward interpreting TCR dynamics in the presence of antigenic stimuli. The expectation based on the model can be useful in identifying clones that have unusual dynamics, quantified as being significantly different from the TCR bulk repertoire. More precisely, for any given clonal trajectory, we can use the model to compute a P -value corresponding to the posterior probability that the frequency fold change s falls outside a certain number of standard deviations relative to the rest of the repertoire ($|s - t/\tau|/\sqrt{t/\theta}$). This can allow us to identify candidate clones associated with a response to an infection or an immune challenge.

Our work was possible because we were able to calibrate the noise using replicate samples. However, replicates are not always available. In this case, the dynamics can still be learned for large clones: We showed using simulations that above a certain frequency threshold, the sampling error becomes small and we can use empirical observations to learn TCR repertoire dynamics

directly from read counts. This allows us to correctly estimate the dynamics of large clones without a noise model if the clone sizes are large at both time points. However, since the repertoire is described by a power law distribution, the role of small clones is far from negligible. An alternative to replicates may be to use close-by time points (relative to the time scales of the dynamics) as surrogate replicates. While we had such time points separated by 1 mo for P1, P2, and P7, we did not attempt a full inference on these samples: We did not manage to learn a reliable noise model for these donors because we lacked both the raw sequencing reads and details about the processing procedure (PCR amplification, error correction, etc). In particular, unlike uniquely barcoded cDNA sequencing, PCR amplification of gDNA used for these donors inflates rare clonotypes (as suggested by the low-frequency plateau in the clone size distributions, Fig. 1D), potentially confounding the analysis.

One of the main conclusions of our work is that repertoires are very dynamic systems, with clone frequencies changing by orders of magnitude on timescales of years, even in the absence of strong known stimulation. This observation challenges our ability to identify responding clonotypes to direct immune stimulation, such as vaccination or diseases. This work builds the ground for inference procedures that are not only correct for experimental and biological noise but also for the natural repertoire dynamics. The methods we designed are general and can be used on larger cohorts of individuals presenting different health statuses, ages, and immunodeficiency features. They provide a promising tool to better understand the maintenance and efficiency of T cells, enabling to quantify immunosenescence (53), which plays an important role in vaccine performance and cancer research.

Methods

A. Longitudinal Data. The datasets analyzed in this study are summarized in *SI Appendix, Table S1*, along with accession number and links to databases.

Data were collected from four different studies, which use two different techniques for repertoire sequencing. Data from refs. 28, 30–32 were generated by sequencing TCR mRNA of PBMCs from healthy individuals, while data from ref. 29 were obtained by directly sequencing genomic DNA (gDNA), as described in detail in each original study.

Briefly, mRNA sequencing was done through cDNA synthesis with template switch allowing for the addition of a unique molecular identifier (UMI), followed by 2-step PCR amplification of the TCR loci (alpha and/or beta), multiplexing, sequencing on an Illumina platform, and processing using the MixCR software package (54) to obtain lists of clonotypes (V and J segments and Complementarity Determining Region 3 nucleotide sequence) corrected for UMI multiplicity and sequencing errors. gDNA sequencing was done by extracting genomic DNA and performing multiplex PCR to amplify the TCR beta gene before sequencing on an Illumina HiSeq system. Raw data processing was performed using closed software. Since the raw data are not available, we used the processed data provided on the ImmuneAccess platform.

B. Naive Inference. The naive inference method directly uses the observed TCR clonal frequencies to learn τ and θ parameters, assuming that they represent exactly the true frequencies: $f = \hat{f} = \hat{n}/N_r$. We aim here at maximizing directly the log-likelihood $\mathcal{L}(\tau, \theta) = \log \mathbb{P}(\{(f_i(t_1), f_i(t_2))\} | \tau, \theta)$, which can be expressed by integrating Eq. 1:

$$\mathbb{P}(\{(\ln f_i(t_1), \ln f_i(t_2) \}) | \tau, \theta) = \prod_i^{N_c} G(\ln f_i(t_2) | \ln f_i(t_1); \tau, \theta) P(\ln f_i(t_1)) \quad [4]$$

where

$$G(x|y; \tau, \theta) = \sqrt{\frac{\theta}{2\pi\Delta t}} \exp - \frac{\theta(x - y - \Delta t\tau^{-1})^2}{2\Delta t} \quad [5]$$

is the propagator of the Brownian motion, $\Delta t = t_2 - t_1$ the time interval between the two time points, and where we have assumed that N_{cell} is a constant of time. Maximizing the log-likelihood with respect to τ and θ is equivalent to performing a linear regression of $\ln f(t_2) - \ln f(t_1)$ against Δt .

C. Full Inference. Using same-day replicates at time t_j , we jointly learn the parameters (α, f_{\min}, a, b) of the clone-size distribution $\rho(f) = Cf^{-1-\alpha}$ (for $f_{\min} \leq f \leq 1$) and the noise model $P(\hat{n}|f) = \text{NegBin}(\hat{n}; N_r f, N_r f + a(N_r f)^b)$ using NoiSET software (34), where $\text{NegBin}(n; x, \sigma)$ is a negative binomial of mean x and variance σ . The learned parameters are reported in *SI Appendix, Fig. S1*.

We then learn the parameters of the dynamics by maximizing the likelihood of samples taken at two different time points, using the noise model to account for the discrepancy between true frequencies and sequence counts: For one clone, the full model likelihood reads

$$\mathbb{P}(\hat{n}_i(t_1) = \hat{n}_1, \hat{n}_i(t_2) = \hat{n}_2 | \tau, \theta) = \int_{[f_{\min}, 1]^2} df_1 \rho(f_1) \frac{df_2}{f_2} G(\ln f_2 | \ln f_1; \tau, \theta) P(\hat{n}_1 | f_1) P(\hat{n}_2 | f_2) \quad [6]$$

where the noise models are specific to each time point.

The maximum likelihood estimator is given by:

$$(\tau^*, \theta^*) = \underset{(\tau, \theta)}{\text{argmax}} \prod_{i=1}^{N_c} \frac{\mathbb{P}(\hat{n}_i(t_1), \hat{n}_i(t_2) | \tau, \theta)}{\mathbb{P}(\hat{n}_i(t_1) \geq N_r f_{\text{th}}, \hat{n}_i(t_2) > 0 | \tau, \theta)} \quad [7]$$

where the denominator accounts for the condition that the clone be included in the analysis: $\hat{f}_i(t) \geq f_{\text{th}}$ and $\hat{n}_2 > 0$. The choice to impose a threshold on the first time point is justified by the fact we are learning the forward propagator of the dynamics, which is conditioned on the initial frequency. Typically around 50 to 70% of clones above the threshold on the first time point remain above the threshold in the second time point. Likewise, similar percentages (40 to 80%) of clones above the threshold in the second time point were also seen in the first time point. This loss is expected since the frequencies follow stochastic trajectories, many of which are likely to cross the threshold between the two time points.

The persistence probability of Fig. 4D is linked to that normalization and is computed as follows:

$$P_{\text{pers}}(\tau, \theta) = \frac{\mathbb{P}(\hat{n}_i(t_1) \geq N_r f_{\text{th}}, \hat{n}_i(t_2) > 0 | \tau, \theta)}{\mathbb{P}(\hat{n}_i(t_1) \geq N_r f_{\text{th}} | \tau, \theta)} \quad [8]$$

Once the model is learned, the posterior distribution of fold changes $s_i \equiv \ln f_i(t_2) - \ln f_i(t_1)$ of each clone i is computed through

$$\mathbb{P}(s_i = s | \hat{n}_1, \hat{n}_2, \tau^*, \theta^*) = \frac{\int_{f_{\min}}^1 df_1 \rho(f_1) G(\ln f_1 + s | \ln f_1; \tau^*, \theta^*) P(\hat{n}_1 | f_1) P(\hat{n}_2 | f_1 e^s)}{\mathbb{P}(\hat{n}_1, \hat{n}_2 | \tau^*, \theta^*)} \quad [9]$$

The overall posterior distribution over all clones (solid lines in *SI Appendix, Fig. S3*) is then given by $\mathbb{P}_{\text{post}}(s) = (1/N_c) \sum_{i=1}^{N_c} \mathbb{P}(s_i = s | \hat{n}_1, \hat{n}_2, \tau^*, \theta^*)$. The prior distribution (dashed line), by contrast, is directly given by $G(\ln f_1 + s, \ln f_1 | \tau^*, \theta^*)$, which is independent of f_1 .

When performing inference in each frequency bin, the product in Eq. 7 runs over clones that fall in the bin, and the normalization in the denominator is replaced by the probability to observe $\hat{n}_i(t_1)$ in the bin of interest, and

$\hat{n}_i(t_2) > 0$. The maximization is performed using the `minimize` function from the Scipy package, with the Sequential Least Squares Programming (SLSQP) method (55) with parameters `tol=1e-8` and `maxiter=300` and initial condition $\tau = 2, \theta = .5$ and constraint $\theta^{-1} > 10^{-3}$.

D. Synthetic Data. Synthetic data were generated by simulating Eq. 1 with a source term producing new clones with rate S at initial size $n = n_0 = 40$ and an absorbing boundary condition at $n = 1$. We work with the $x = \ln n$ variable for convenience. The simulation is initialized at steady state, which can be computed analytically (5, 9). The analytical solution gives us the expected number of cells and clones as a function of the model parameters: $N_{\text{cell}} = S(n_0 - 1)/(\tau^{-1} - \theta^{-1}/2)$ and $N_c = S\tau \ln n_0$. Fixing the number of cells to $N_{\text{cell}} = 10^{10}$, we then compute the number of clones necessary to achieve that size, $N_c = N_{\text{cell}}(1 - \tau/2\theta) \ln n_0 / (n_0 - 1)$. We then draw the size $n_i(t_1) = e^{x_i(t_1)}$ of each clone $i = 1, \dots, N_c$ from the continuous steady-state distribution (5):

$$\rho_x(x) = \begin{cases} S\tau(1 - e^{-\alpha x}) & \text{if } x \leq x_0 \equiv \ln n_0 \\ S\tau e^{-\alpha x} (e^{\alpha x_0} - 1) & \text{if } x > x_0, \end{cases} \quad [10]$$

with $\alpha = 2\theta/\tau$.

Then, the evolution of each clone from time t_1 to $t_2 = t_1 + \Delta t$ is determined by the modified propagator with the absorbing boundary condition at $x = 0$:

$$G_{\text{abs}}(x|y) = G(x|y) - e^{-\alpha y} G(x| - y), \quad [11]$$

where $G(x|y)$ is defined in Eq. 5. In practice, we kill clone i with probability $1 - P_{\text{surv}}(x_i(t_1)) \equiv 1 - \int_0^\infty dx G_{\text{abs}}(x|x_i(t_1))$, which can be expressed in terms of error functions. Otherwise, its new log-size $x_i(t_2)$ is drawn from the distribution $G_{\text{abs}}(x|x_i(t_1))/P_{\text{surv}}(x_i(t_1))$.

In addition, new clones are introduced during Δt . We draw their number from a Poisson distribution of mean $S\Delta t$ and their introduction times t from a uniform distribution in the interval $[t_1, t_2]$. Then, their dynamics are drawn in the same way as for the initial clones, but with $\Delta t = t_2 - t$ instead of $t_2 - t_1$.

Once the abundances ($n_i(t_1) = e^{x_i(t_1)}, n_i(t_2) = e^{x_i(t_2)}$) have been determined, the number of reads $\hat{n}_i(t_1), \hat{n}_i(t_2)$ from each time point is drawn from a negative binomial distribution of mean $\langle \hat{n}_i(t_1) \rangle = N_r n_i(t_1)/N_{\text{cell}}$ and variance $\langle \hat{n}_i(t_1) \rangle + a \langle \hat{n}_i(t_1) \rangle^b$, and likewise for $\hat{n}_i(t_2)$, with $N_r = 10^6, a = 0.7$ and $b = 1.1$.

E. Comparison to the VDjdb Database. We downloaded the 2022-03-30 release of VDjdb (50) at <https://github.com/antigenomics/vdjdb-db/releases/download/2022-03-30/vdjdb-2022-03-30.zip> and restricted our search to CDR3s associated with antigens from the following species: CMV, InfluenzaA, EBV, SARS-CoV-2, HIV-1, HCV, YFV, HTLV-1, DENV1, DENV3/4, HIV, HSV-2, M.tuberculosis, DENV2, HCoV-HKU1, HPV, MCPyV, StreptomycesKanamyceticus, E.Coli, HIV1, HHV, PseudomonasAeruginosa, PseudomonasFluorescens, SaccharomycesCerevisiae, SelaginellaMoellendorffii, totaling 65616 CDR3 amino acid sequences.

Data, Materials, and Software Availability. All scripts to produce the figures can be found at https://github.com/statbiophys/inferring_TCR_repertoire_dynamics/. Previously published data were used for this work. O. V. Britanova, et al., Dynamics of individual T cell repertoires: From cord blood to centenarians. *J. Immunol.* 196, 5005–5013 (2016). N. D. Chu, et al., Longitudinal immunosequencing in healthy people reveals persistent T cell receptors rich in highly public receptors. *BMC Immunol.* 20, 19 (2019). A. L. Sycheva, et al., Quantitative profiling reveals minor changes of T cell receptor repertoire in response to subunit inactivated influenza vaccine. *Vaccine* 36, 1599–1605 (2018). M. V. Pogorelyy, et al., Precise tracking of vaccine-responding T cell clones reveals convergent and personalized response in identical twins. *Proc. Natl. Acad. Sci. U.S.A.* 115, 12704–12709 (2018). A. A. Minervina, et al., Primary and secondary anti-viral response captured by the dynamics and phenotype of individual T cell clones. *eLife* 9, e53704 (2020).

ACKNOWLEDGMENTS. This work was partially supported by the European Research Council Consolidator Grant no. 724208 and ANR-19-CE45-0018 “RESP-REP” from the Agence Nationale de la Recherche.

- A. J. Yates, Theories and quantification of thymic selection. *Front. Immunol.* 5, 13 (2014).
- I. Bains, R. Thiébaud, A. J. Yates, R. Callard, Quantifying thymic export: Combining models of naive T cell proliferation and TCR excision circle dynamics gives an explicit measure of thymic output. *J. Immunol.* 183, 4329–4336 (2009).
- S. C. Jameson, Maintaining the norm: T-cell homeostasis. *Nat. Rev. Immunol.* 2, 547–556 (2002).
- M. R. Dowling, P. D. Hodgkin, Modelling naive T-cell homeostasis: Consequences of heritable cellular lifespan during ageing. *Immunol. Cell Biol.* 87, 445–456 (2009).
- J. Desponds, T. Mora, A. M. Walczak, Fluctuating fitness shapes the clone-size distribution of immune repertoires. *Proc. Natl. Acad. Sci. U.S.A.* 113, 274–279 (2016).
- T. Mora, A. M. Walczak, *Quantifying Lymphocyte Receptor Diversity in Systems Immunology* (CRC Press, 2019).
- Q. Qi et al., Diversity and clonal selection in the human T-cell repertoire. *Proc. Natl. Acad. Sci. U.S.A.* 111, 13139–13144 (2014).
- G. Lythe, R. E. Callard, R. L. Hoare, C. Molina-Paris, How many TCR clonotypes does a body maintain? *J. Theor. Biol.* 389, 214–224 (2016).
- T. Mora, A. M. Walczak, How many different clonotypes do immune repertoires contain? *Curr. Opin. Syst. Biol.* 18, 104–110 (2019).
- R. J. De Boer, D. Homann, A. S. Perelson, Different dynamics of CD4+ and CD8+ T cell responses during and after acute lymphocytic choriomeningitis virus infection. *J. Immunol.* 171, 3928–3935 (2003).
- K. Kedzierska, S. A. Valkenburg, P. C. Doherty, M. P. Davenport, V. Venturi, Use it or lose it: Establishment and persistence of T cell memory. *Front. Immunol.* 3 (2012).
- A. Mayer, Y. Zhang, A. S. Perelson, N. S. Wingreen, Regulation of T cell expansion by antigen presentation dynamics. *Proc. Natl. Acad. Sci. U.S.A.* 116, 5914–5919 (2019).
- I. Bains, R. Antia, R. Callard, A. J. Yates, Quantifying the development of the peripheral naive CD4+ T-cell pool in humans. *Blood* 113, 5480–5487 (2009).
- P. C. de Greef et al., The naive T-cell receptor repertoire has an extremely broad distribution of clone sizes. *eLife* 9, e49900 (2020).
- R. Dessalles et al., How naive T-cell clone counts are shaped by heterogeneous thymic output and homeostatic proliferation. *Front. Immunol.* 12 (2022).
- R. J. De Boer, A. S. Perelson, T cell repertoires and competitive exclusion. *J. Theor. Biol.* 169, 375–390 (1994).
- P. Dash et al., Quantifiable predictive features define epitope-specific T cell receptor repertoires. *Nature* 547, 89–93 (2017).
- A. Mayer, V. Balasubramanian, A. M. Walczak, T. Mora, How a well-adapting immune system remembers. *Proc. Natl. Acad. Sci. U.S.A.* 116, 8815–8823 (2019).
- P. L. F. Johnson, A. J. Yates, J. J. Goronzy, R. Antia, Peripheral selection rather than thymic involution explains sudden contraction in naive CD4 T-cell diversity with age. *Proc. Natl. Acad. Sci. U.S.A.* 109, 21432–21437 (2012).
- R. J. De Boer, A. S. Perelson, Quantifying T lymphocyte turnover. *J. Theor. Biol.* 327, 45–87 (2013).
- P. Lindau, H. S. Robins, Advances and applications of immune receptor sequencing in systems immunology. *Curr. Opin. Syst. Biol.* 1, 62–68 (2017).
- M. M. Davis, S. D. Boyd, Recent progress in the analysis of $\alpha\beta$ T cell and B cell receptor repertoires. *Curr. Opin. Immunol.* 59, 109–114 (2019).
- A. Minervina, M. Pogorelyy, I. Mamedov, T-cell receptor and B-cell receptor repertoire profiling in adaptive immunity. *Transplant Int.* 32, 1111–1123 (2019).
- J. D. Burgos, P. Moreno-Tovar, Zipf-scaling behavior in the immune system. *Biosystems* 39, 227–232 (1996).
- H. Koch, D. Starenki, S. J. Cooper, R. M. Myers, Q. Li, powerTCR: A model-based approach to comparative analysis of the clone size distribution of the T cell receptor repertoire. *PLoS Comput. Biol.* 14, e1006571 (2018).
- Y. N. Naumov, E. N. Naumova, K. T. Hogan, L. K. Selin, J. Gorski, A fractal clonotype distribution in the CD8+ memory T cell repertoire could optimize potential for immune responses. *J. Immunol.* 170, 3994–4001 (2003).
- V. I. Zarnitsyna, B. D. Evavold, L. N. Schoettle, J. N. Blattman, R. Antia, Estimating the diversity, completeness, and cross-reactivity of the T cell repertoire. *Front. Immunol.* 4 (2013).
- O. V. Britanova et al., Dynamics of individual T cell repertoires: From cord blood to centenarians. *J. Immunol.* 196, 5005–5013 (2016).
- N. D. Chu et al., Longitudinal immunosequencing in healthy people reveals persistent T cell receptors rich in highly public receptors. *BMC Immunol.* 20, 19 (2019).
- A. L. Sycheva et al., Quantitative profiling reveals minor changes of T cell receptor repertoire in response to subunit inactivated influenza vaccine. *Vaccine* 36, 1599–1605 (2018).
- M. V. Pogorelyy et al., Precise tracking of vaccine-responding T cell clones reveals convergent and personalized response in identical twins. *Proc. Natl. Acad. Sci. U.S.A.* 115, 12704–12709 (2018).
- A. A. Minervina et al., Primary and secondary anti-viral response captured by the dynamics and phenotype of individual T cell clones. *eLife* 9, e53704 (2020).
- M. Puelma Touzel, A. M. Walczak, T. Mora, Inferring the immune response from repertoire sequencing. *PLoS Comput. Biol.* 16, e1007873 (2020).
- M. Bensouda Koraiichi, M. P. Touzel, T. Mora, A. M. Walczak, NoisET: Noise learning and expansion detection of T-cell receptors with Python. [q-bio]. arXiv [Preprint] (2021). <http://arxiv.org/abs/2102.03568>.
- M. U. Gaimann, M. Nguyen, J. Desponds, A. Mayer, Early life imprints the hierarchy of T cell clone sizes. *eLife* 9, e61639 (2020).
- T. Dupic, Q. Marcou, A. M. Walczak, T. Mora, Genesis of the alpha beta T-cell receptor. *PLoS Comput. Biol.* 15, e1006874 (2019).
- J. M. Heather, M. Ismail, T. Oakes, B. Chain, High-throughput sequencing of the T-cell receptor repertoire: Pitfalls and opportunities. *Briefings Bioinf.* bbw138 (2017).

38. G. Altan-Bonnet, T. Mora, A. M. Walczak, Quantitative immunology for physicists. *Phys. Rep.* **849**, 1–83 (2020).
39. O. V. Britanova *et al.*, Age-related decrease in TCR repertoire diversity measured with deep and normalized sequence profiling. *J. Immunol.* **192**, 2689–2698 (2014).
40. J. A. Pai, A. T. Satpathy, High-throughput and single-cell T cell receptor sequencing technologies. *Nat. Methods* **18**, 881–892 (2021).
41. S. Valkiers *et al.*, Recent advances in T-cell receptor repertoire analysis: Bridging the gap with multimodal single-cell RNA sequencing. *Immunoinformatics* **5**, 100009 (2022).
42. J. Glanville *et al.*, Identifying specificity groups in the T cell receptor repertoire. *Nature* **547**, 94–98 (2017).
43. K. Mayer-Blackwell *et al.*, TCR meta-clonotypes for biomarker discovery with Tcrdist3 enabled identification of public, HLA-restricted clusters of SARS-CoV-2 TCRs. *eLife* **10**, e68605 (2021).
44. S. Gielis *et al.*, Detection of enriched T cell epitope specificity in full T cell receptor sequence repertoires. *Front. Immunol.* **10** (2019).
45. A. Montemurro *et al.*, NetTCR-2.0 enables accurate prediction of TCR-peptide binding by using paired TCR α and β sequence data. *Commun. Biol.* **4**, 1–13 (2021).
46. J. W. Sidhom, H. B. Larman, D. M. Pardoll, A. S. Baras, DeepTCR is a deep learning framework for revealing sequence concepts within T-cell repertoires. *Nat. Commun.* **12**, 1605 (2021).
47. I. Springer, N. Tickotsky, Y. Louzoun, Contribution of T cell receptor alpha and beta CDR3, MHC Typing, V and J genes to peptide binding prediction. *Front. Immunol.* **12** (2021).
48. W. Zhang *et al.*, A framework for highly multiplexed dextramer mapping and prediction of T cell receptor sequences to antigen specificity. *Sci. Adv.* **7**, eabf5835 (2021).
49. A. A. Minervina *et al.*, Longitudinal high-throughput TCR repertoire profiling reveals the dynamics of T-cell memory formation after mild COVID-19 infection. *eLife* **10**, e63502 (2021).
50. D. V. Bagaev *et al.*, VDJdb in 2019: Database extension, new analysis infrastructure and a T-cell receptor motif compendium. *Nucleic Acids Res.* **48**, D1057–D1062 (2020).
51. T. Dupic *et al.*, Immune fingerprinting through repertoire similarity. *PLOS Genet.* **17**, e1009301 (2021).
52. M. V. Pogorelyy *et al.*, Persisting fetal clonotypes influence the structure and overlap of adult human T cell receptor repertoires. *PLOS Comput. Biol.* **13**, e1005572 (2017).
53. H. Zhang, C. M. Weyand, J. J. Goronzy, Hallmarks of the aging T-cell system. *FEBS J.* **288**, 7123–7142 (2021).
54. D. A. Bolotin *et al.*, MixCR: Software for comprehensive adaptive immunity profiling. *Nat. Methods* **12**, 380–381 (2015).
55. P. Virtanen *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nat. Methods* **17**, 261–272 (2020).