

PNAS



Supporting Information for

An inference model gives insights into innate immune adaptation and repertoire diversity

Yawei Qin, Emily M. Mace, John P. Barton

John P. Barton
E-mail: jpbarton@pitt.edu

This PDF file includes:

Supporting text
Figs. S1 to S12
SI References

Supporting Information Text

Methods

A. Description of the inference model.

A.1. Bayesian inference for a Gaussian distribution with unknown mean and precision. We begin with a normal-gamma prior distribution for the Gaussian mean μ and precision λ . After each signal ‘measurement’ x , parameters are updated via Bayes’ rule,

$$P(\mu, \lambda|x) = \frac{P(x|\mu, \lambda)P_{\text{prior}}(\mu, \lambda)}{\int d\mu' \int d\lambda' P(x|\mu', \lambda')P_{\text{prior}}(\mu', \lambda')} . \quad [1]$$

With a normal-gamma prior, the posterior distribution for μ and λ after taking a new measurement x is

$$\begin{aligned} P_{\text{post}}(\mu, \lambda|x) &\propto P(x|\mu, \lambda) \times P_{\text{prior}}(\mu, \lambda|m, \beta, \alpha, \kappa) \\ &= \lambda^\alpha e^{\left[-\beta - \frac{1}{2} \frac{\kappa}{\kappa+1} (x-m)^2 - (\kappa+1) \frac{\left(\mu - \frac{\kappa m + x}{\kappa+1}\right)^2}{2} \right]} . \end{aligned}$$

Recall that

$$P(x|\mu, \lambda) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}} = \sqrt{\frac{\lambda}{2\pi}} e^{-\frac{\lambda(x-\mu)^2}{2}}$$

is the likelihood function, where $\lambda = 1/\sigma^2$ is the precision, and

$$P_{\text{prior}}(\mu, \lambda|m, \beta, \alpha, \kappa) = \frac{\beta^\alpha \sqrt{\kappa}}{\Gamma(\alpha) \sqrt{2\pi}} \lambda^{\alpha-1/2} e^{\lambda \left[-\beta - \kappa \frac{(\mu-m)^2}{2} \right]}$$

is the normal-gamma prior where $\beta > 0$, $\alpha > 0$ and $\kappa > 0$. The posterior thus also follows a normal-gamma distribution with modified parameters m' , β' , α' , and κ' , which are related to those of the prior distribution by

$$\begin{aligned} m' &= \frac{\kappa m + x}{\kappa + 1} & \kappa' &= \kappa + 1 \\ \beta' &= \beta + \frac{\kappa}{\kappa + 1} \frac{(x-m)^2}{2} & \alpha' &= \alpha + \frac{1}{2} . \end{aligned} \quad [2]$$

A.2. Estimating the signal mean and variance. Estimated mean $\hat{\mu}$ and variance $\hat{\sigma}^2$ are the expected values of μ and $\sigma^2 = \frac{1}{\lambda}$ for the distribution $P(\mu, \lambda)$. Prior to the measurement, $\hat{\mu}$ and $\hat{\sigma}^2$ are

$$\begin{aligned} \hat{\mu} &= \iint d\mu d\lambda \mu P(\mu, \lambda|m, \beta, \alpha, \kappa) = m , \\ \hat{\sigma}^2 &= \left\langle \frac{1}{\lambda} \right\rangle = \iint d\mu d\lambda \frac{1}{\lambda} P(\mu, \lambda|m, \beta, \alpha, \kappa) \\ &= \frac{\beta}{\alpha - 1} . \end{aligned} \quad [3]$$

After receiving a signal x , the new estimated mean and variance are

$$\begin{aligned} \hat{\mu}' &= m' = \frac{\kappa m + x}{\kappa + 1} = \frac{\kappa \hat{\mu} + x}{\kappa + 1} , \\ (\hat{\sigma}^2)' &= \frac{\beta'}{\alpha' - 1} = \frac{\beta + \frac{1}{2} \frac{\kappa(x-m)^2}{(\kappa+1)}}{\alpha - \frac{1}{2}} \\ &= \frac{\hat{\sigma}^2 (\alpha - 1) + \frac{1}{2} \frac{\kappa(x-m)^2}{(\kappa+1)}}{\alpha - \frac{1}{2}} . \end{aligned} \quad [4]$$

A.3. Reconstructed signal distribution. In our model, the internal representation of the signal distribution $P_r(x)$ is calculated by integrating Gaussian distributions over mean μ and precision λ , weighted by the probability distribution $P(\mu, \lambda|m, \kappa, \alpha, \beta)$. Performing this integration gives

$$\begin{aligned} P_r(x) &= \iint d\mu d\lambda P(x|\mu, \lambda) P(\mu, \lambda|m, \beta, \alpha, \kappa) \\ &= \int_{-\infty}^{\infty} d\mu \int_0^{\infty} d\lambda \frac{\sqrt{\lambda}}{\sqrt{2\pi}} e^{-\frac{\lambda(x-\mu)^2}{2}} \times \frac{\beta^\alpha}{\Gamma(\alpha)} \left(\frac{\kappa}{2\pi} \right)^{\frac{1}{2}} \lambda^{\alpha-\frac{1}{2}} e^{-\frac{\lambda}{2} [\kappa(\mu-m)^2 + 2\beta]} \\ &= \beta^\alpha \sqrt{\frac{\kappa}{2\pi(\kappa+1)}} \frac{\Gamma(\alpha + \frac{1}{2})}{\Gamma(\alpha)} \times \left[\beta + \frac{1}{2} \frac{\kappa}{\kappa+1} (x-m)^2 \right]^{-\alpha-\frac{1}{2}} . \end{aligned}$$

The distribution above is a scaled, shifted t-distribution

$$P(t) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\pi\nu}\Gamma(\nu/2)} \left(1 + \frac{t^2}{\nu}\right)^{-\frac{\nu+1}{2}},$$

where $t = \sqrt{\frac{\alpha\kappa}{\beta(\kappa+1)}}(x - m)$ and $\nu = 2\alpha$.

A.4. Modified Bayesian inference to adapt to multiple environments. While the standard Bayesian approach described above works well for static signal distributions, it is not suitable for a signal distribution that changes over time. This is because the parameters κ and α are incremented after every measurement in (2), thereby minimizing the contribution of recent data after many measurements have been collected.

To enable adaptation to dynamic environments, we modify the update rules by fixing κ and α . These parameters count the effective number of samples used to estimate the mean and precision, respectively (see (2) and (4)). Fixing α creates a potential problem because β and $(x - m)^2$ are nonnegative, which means that β will increase with each update (see (2)) and thus $\hat{\sigma}^2$ will grow without bound. β' must therefore be shrunk by an α -dependent factor δ_α to keep estimates of the standard deviation finite. We can determine this factor by requiring that the update of the standard deviation follow the same form as in (4), where now $\alpha' = \alpha$,

$$(\hat{\sigma}^2)' = \frac{\beta'}{\alpha - 1} = \delta_\alpha \frac{\beta + \frac{1}{2} \frac{\kappa(x-m)^2}{(\kappa+1)}}{\alpha - 1} = \frac{\beta + \frac{1}{2} \frac{\kappa(x-m)^2}{(\kappa+1)}}{\alpha - \frac{1}{2}}, \quad [5]$$

which yields $\delta_\alpha = (\alpha - 1) / (\alpha - 1/2)$.

Rescaling β' in this way while keeping κ and α fixed,

$$\beta' = \frac{\alpha - 1}{\alpha - \frac{1}{2}} \left[\beta + \frac{\kappa}{\kappa + 1} \frac{(x - m)^2}{2} \right],$$

not only ensures that β is finite, it also gives the same expression for updating the estimated variance as previously,

$$(\hat{\sigma}^2)' = \frac{\beta'}{\alpha' - 1} = \frac{\alpha - 1}{\alpha - \frac{1}{2}} \hat{\sigma}^2 + \frac{\frac{1}{2} \frac{\kappa}{\alpha - \frac{1}{2} \kappa + 1}}{\alpha - \frac{1}{2} \kappa + 1} (x - m)^2.$$

In summary, the adaptive update rules are given by

$$\begin{aligned} m' &= \frac{\kappa m + x}{\kappa + 1} & \kappa' &= \kappa \\ \beta' &= \frac{\alpha - 1}{\alpha - \frac{1}{2}} \left[\beta + \frac{\kappa}{\kappa + 1} \frac{(x - m)^2}{2} \right] & \alpha' &= \alpha. \end{aligned} \quad [6]$$

A.5. Equivalence with exponential weight decay. In standard Bayesian inference, all measurements contribute equally to parameter estimation. Let $\{x_1, x_2, \dots, x_n\}$ represent a set of n signals, where the subscript indicates the order in which each signal was measured. Assuming starting prior parameter values of m_0, β_0, α_0 and κ_0 , and following the standard update rules given in (2), after n signal measurements we have

$$\begin{aligned} m_n &= \frac{\kappa_0}{\kappa_0 + n} m_0 + \frac{1}{\kappa_0 + n} \sum_{i=1}^n x_i, \\ \beta_n &= \beta_0 + \frac{1}{2} \frac{\kappa_0 n}{\kappa_0 + n} \left(\frac{1}{n} \sum_{i=1}^n x_i - m_0 \right)^2 + \frac{1}{2} \sum_{i=1}^n \left(x_i - \frac{1}{n} \sum_{j=1}^n x_j \right)^2. \end{aligned}$$

If we instead follow the adaptive update rules given in (6), the parameter estimates become

$$\begin{aligned} m_n &= \frac{\kappa}{\kappa + 1} m_{n-1} + \frac{1}{\kappa + 1} x_n \\ &= \left(\frac{\kappa}{\kappa + 1} \right)^2 m_{n-2} + \frac{\kappa}{\kappa + 1} \frac{1}{\kappa + 1} x_{n-1} + \frac{1}{\kappa + 1} x_n \\ &= \delta_\kappa^n m_0 + \frac{1}{\kappa + 1} \sum_{i=1}^n \delta_\kappa^{n-i} x_i, \end{aligned}$$

where $\delta_\kappa = \kappa / (\kappa + 1)$, and

$$\begin{aligned} \beta_n &= \frac{\alpha - 1}{\alpha - \frac{1}{2}} \left[\beta_{n-1} + \frac{1}{2} \frac{\kappa}{\kappa + 1} (x_n - m_{n-1})^2 \right] \\ &= \delta_\alpha^n \beta_0 + \frac{1}{2} \delta_\kappa \sum_{i=1}^n \delta_\alpha^{n-i+1} (x_i - m_{i-1})^2, \end{aligned}$$

where $\delta_\alpha = (\alpha - 1) / (\alpha - 1/2)$. Here, each of the n measurements no longer contributes equally to the estimated parameters. New measurements are emphasized more strongly than old ones. With each successive measurement, the effective weight of older signal measurements decreases by a factor of δ_κ for m , and by a factor of δ_α for β .

The argument above facilitates the interpretation of κ and α as parameters controlling the memory length for estimating the signal distribution. Here it is natural to take $\kappa = 2\alpha$, but in principle the memory length for m and β could be separately controlled by adjusting κ and α , respectively. Larger values of κ and α result in longer memory lengths and slower adaptation.

A.6. Inference with signal saturation. In the analysis above, the range of the signal x is in principle unbounded. However, immune cells have a finite number of receptors, limiting the magnitude of activating or inhibitory stimulus that they can receive. To mimic this saturation effect, we considered sigmoidally transforming signals x relative to the current estimated signal mean m . Specifically, we considered a transformed signal x' given by

$$x' = m + r \left[\frac{2}{1 + e^{-(x-m)}} - 1 \right],$$

where r sets the range of signals that can be received, which lies in $[m-r, m+r]$. Because strong activating and inhibitory signals are clipped to a finite range, adaptation to environments that differ substantially from the current one is slower (see **Figs. S5, S6**). In this model, the estimated mean ultimately converges to the true signal mean but the standard deviation converges to a different, smaller value (which we call the transformed standard deviation, σ_{tf}) because the sigmoidal transformation reduces the signal variance.

B. Characterizing the evolution of inferred signal statistics across a population of immune cells.

B.1. Exploring heterogeneity in immune cell populations. As shown in **Fig. 2**, immune cells in our model do not adapt perfectly to the true signal distribution in the environment due to finite memory. Together with the stochastic nature of target cell encounters, this implies that there will be a range of immune cell responses, even among immune cells with identical receptors. **Figure S4** shows an example of this behavior in a population of immune cells transferred between different environments, as shown for a single cell in **Fig. 3**.

To systematically explore this heterogeneity, we sought to characterize the distribution of estimated signal parameters m and β for a population of identical immune cells with finite memory, which inhabit the same environment. As an analytical approach, we developed a continuous approximation of the discrete modified Bayesian update dynamics described in (6). Assuming that parameter updates with each target cell encounter are small, (6) is described by a set of stochastic differential equations. We then derived a Fokker-Planck equation that describes the evolution of the distribution of learned m and β values for a population of immune cells with identical receptors in the same environment. These derivations are described below.

B.2. Derivation of the stochastic differential equations. Here we attempt to characterize the evolution of the estimated signal distribution parameters across a large population of cells with identical receptors. To do this we begin by using the parameter update rules, given in (6), to take the difference between the new the updated parameters,

$$\begin{aligned} \Delta m &= m' - m = \frac{1}{\kappa + 1} (x - m), \\ \Delta \beta &= \beta' - \beta = \frac{1}{2} \left[\delta_\alpha \delta_\kappa (x - m)^2 - \frac{\beta}{\alpha - \frac{1}{2}} \right]. \end{aligned}$$

For an immune cell that is well-adapted to the environment, Δm and $\Delta \beta$ should be small, especially when the memory parameters κ and α are large. Writing the time interval Δt as a single encounter, we can develop stochastic differential equations that describe the evolution of m and β ,

$$\begin{aligned} \frac{\Delta m}{\Delta t} &= \frac{1}{\kappa + 1} (x - m) = \frac{1}{\kappa + 1} (\mu_t - m + \sigma_t \xi), \\ \frac{\Delta \beta}{\Delta t} &= \frac{1}{2} \left[\delta_\alpha \delta_\kappa (x - m)^2 - \frac{\beta}{\alpha - \frac{1}{2}} \right] \\ &= \frac{1}{2} \left\{ \delta_\alpha \delta_\kappa [(\mu_t - m)^2 + 2(\mu_t - m)\sigma_t \xi + \sigma_t^2 \xi^2] - \frac{\beta}{\alpha - \frac{1}{2}} \right\}. \end{aligned}$$

Here we assume a constant environment, where the signal distribution is Gaussian with mean μ_t and standard deviation σ_t . We thus write $x = \mu_t + \sigma_t \xi$, where ξ is a Gaussian white noise.

Collectively, the evolution of $\theta = (m, \beta)$ follows

$$d\theta = \mathbf{A}(\theta, t) dt + \sqrt{\mathbf{D}(\theta, t)} dW, \quad [7]$$

where \mathbf{A} and \mathbf{D} are referred to as the drift vector and the diffusion matrix, respectively. The drift vector describes the expected change in $\boldsymbol{\theta}$ parameters, and the diffusion matrix describes the covariance of changes in $\boldsymbol{\theta}$,

$$\mathbf{A} = \frac{\langle \Delta \boldsymbol{\theta} \rangle}{\Delta t},$$

$$\mathbf{D} = \frac{\langle (\Delta \boldsymbol{\theta})^2 \rangle - \langle \Delta \boldsymbol{\theta} \rangle^2}{\Delta t}.$$

The components of \mathbf{A} and \mathbf{D} can thus be computed by finding the first and second moments of $\Delta \boldsymbol{\theta}$, which are

$$\begin{aligned} \frac{\langle \Delta m \rangle}{\Delta t} &= \frac{1}{\kappa + 1} (\mu_t - m), \\ \frac{\langle \Delta \beta \rangle}{\Delta t} &= \frac{1}{2} \left\{ \delta_\alpha \delta_\kappa [(\mu_t - m)^2 + \sigma_t^2] - \frac{\beta}{\alpha - \frac{1}{2}} \right\}, \\ \frac{\langle (\Delta m)^2 \rangle}{\Delta t} &= \frac{1}{(\kappa + 1)^2} [(\mu_t - m)^2 + \sigma_t^2], \\ \frac{\langle (\Delta \beta)^2 \rangle}{\Delta t} &= \frac{1}{4} \left\{ \left[\delta_\alpha \delta_\kappa [(\mu_t - m)^2 + \sigma_t^2] - \frac{\beta}{\alpha - \frac{1}{2}} \right]^2 + 2\delta_\alpha^2 \delta_\kappa^2 \sigma_t^2 [2(\mu_t - m)^2 + \sigma_t^2] \right\}, \\ \frac{\langle \Delta m \Delta \beta \rangle}{\Delta t} &= \frac{1}{\kappa + 1} (\mu_t - m) \times \frac{1}{2} \left\{ \delta_\alpha \delta_\kappa [(\mu_t - m)^2 + 3\sigma_t^2] - \frac{\beta}{\alpha - \frac{1}{2}} \right\}. \end{aligned}$$

This allows us to write down the cumulants

$$\begin{aligned} \frac{\langle (\Delta m)^2 \rangle - \langle \Delta m \rangle^2}{\Delta t} &= \frac{1}{(\kappa + 1)^2} \sigma_t^2, \\ \frac{\langle (\Delta \beta)^2 \rangle - \langle \Delta \beta \rangle^2}{\Delta t} &= \frac{1}{2} \delta_\alpha^2 \delta_\kappa^2 \sigma_t^2 [2(\mu_t - m)^2 + \sigma_t^2], \\ \frac{\langle \Delta m \Delta \beta \rangle - \langle \Delta m \rangle \langle \Delta \beta \rangle}{\Delta t} &= \frac{1}{\kappa + 1} (\mu_t - m) \times \delta_\alpha \delta_\kappa \sigma_t^2. \end{aligned}$$

Finally we obtain expressions for the vector \mathbf{A} and matrix \mathbf{D} ,

$$\mathbf{A} = \begin{bmatrix} \frac{1}{\kappa + 1} (\mu_t - m) \\ \frac{1}{2} \left\{ \delta_\alpha \delta_\kappa [(\mu_t - m)^2 + \sigma_t^2] - \frac{\beta}{\alpha - \frac{1}{2}} \right\} \end{bmatrix},$$

$$\mathbf{D} = \begin{bmatrix} \frac{1}{(\kappa + 1)^2} \sigma_t^2 & \frac{\mu_t - m}{\kappa + 1} \times \delta_\alpha \delta_\kappa \sigma_t^2 \\ \frac{\mu_t - m}{\kappa + 1} \times \delta_\alpha \delta_\kappa \sigma_t^2 & \frac{1}{2} \delta_\alpha^2 \delta_\kappa^2 \sigma_t^2 [2(\mu_t - m)^2 + \sigma_t^2] \end{bmatrix}.$$

B.3. Fokker-Planck equation. Equation (7) is a stochastic differential equation that describes the evolution of $\boldsymbol{\theta} = (m, \beta)$ for a single immune cell. To understand the distribution of m and β at the population level, we derive the Fokker-Planck equation corresponding to (7). Fokker-Planck equation gives the evolution of the probability distribution of the parameter $\boldsymbol{\theta}$. Following standard methods (1), the form of the Fokker-Planck equation is

$$\begin{aligned} \frac{\partial P(\boldsymbol{\theta}, t)}{\partial t} &= - \left(\frac{\partial}{\partial m} + \frac{\partial}{\partial \beta} \right) \mathbf{A}(\boldsymbol{\theta}, t) P(\boldsymbol{\theta}, t) \\ &\quad + \left(\frac{1}{2} \frac{\partial^2}{\partial m^2} + \frac{1}{2} \frac{\partial^2}{\partial \beta^2} + \frac{\partial^2}{\partial m \partial \beta} \right) \mathbf{D}(\boldsymbol{\theta}, t) P(\boldsymbol{\theta}, t). \end{aligned}$$

Substituting in the expressions for \mathbf{A} and \mathbf{D} derived above, we obtain

$$\begin{aligned} \frac{\partial P(m, \beta, t)}{\partial t} &= \left(\frac{1}{\kappa + 1} + \frac{1}{2\alpha - 1} \right) P - \frac{\mu_t - m}{\kappa + 1} \frac{\partial P}{\partial m} - \left\{ \frac{1}{2} \left[\delta_\alpha \delta_\kappa [(\mu_t - m)^2 + \sigma_t^2] - \frac{\beta}{\alpha - \frac{1}{2}} \right] + \frac{\delta_\alpha \delta_\kappa}{\kappa + 1} \sigma_t^2 \right\} \frac{\partial P}{\partial \beta} \\ &\quad + \frac{1}{2} \frac{1}{(\kappa + 1)^2} \sigma_t^2 \frac{\partial^2 P}{\partial m^2} + \frac{\mu_t - m}{\kappa + 1} \times \delta_\alpha \delta_\kappa \sigma_t^2 \frac{\partial^2 P}{\partial m \partial \beta} + \frac{1}{4} \delta_\alpha^2 \delta_\kappa^2 \sigma_t^2 [2(\mu_t - m)^2 + \sigma_t^2] \frac{\partial^2 P}{\partial \beta^2}. \end{aligned} \tag{8}$$

B.4. Numerical integration of the Fokker-Planck equation. We used the central difference method to solve (8) numerically. Here we simulated adaptation to a true signal distribution with mean $\mu_t = 2$ and standard deviation $\sigma_t = 1$, using the memory parameters $\kappa = 20$ and $\alpha = 10$. We anticipated that the learned signal parameters should be centered around $m = 2$ and $\beta = 9$, which would give $\hat{\mu} = m = 2 = \mu_t$ and $\hat{\sigma}^2 = \frac{\beta}{\alpha-1} = 1 = \sigma_t^2$. Thus we selected a rectangular domain $[1.5, 2.5] \times [5, 15]$ in the (m, β) space to numerically evaluate the equations (see Fig. S2). We set $h = w = 0.01$ as the spatial discretization size of m and β , respectively, dividing the domain into a 100×1000 grid. The discrete first and second derivatives at each grid point can be expressed as

$$\begin{aligned}\frac{\partial P}{\partial m}(m, \beta) &= \frac{P(m+h, \beta) - P(m-h, \beta)}{2h}, \\ \frac{\partial P}{\partial \beta}(m, \beta) &= \frac{P(m, \beta+w) - P(m, \beta-w)}{2w}, \\ \frac{\partial^2 P}{\partial^2 m}(m, \beta) &= \frac{P(m+h, \beta) + P(m-h, \beta) - 2P(m, \beta)}{h^2}, \\ \frac{\partial^2 P}{\partial^2 \beta}(m, \beta) &= \frac{P(m, \beta+w) + P(m, \beta-w) - 2P(m, \beta)}{w^2}, \\ \frac{\partial^2 P}{\partial m \partial \beta}(m, \beta) &= \frac{1}{4hw} [P(m+h, \beta+w) + P(m-h, \beta-w) \\ &\quad - P(m+h, \beta-w) - P(m-h, \beta+w)].\end{aligned}$$

After applying the spatial discretization to the right hand side, (8) can be simplified to

$$\frac{\partial \mathbf{P}}{\partial t} = \mathbf{M} \mathbf{P} \quad [9]$$

where the discrete probability distribution \mathbf{P} is a 10^5 dimensional vector representing values of the function at all grid points. \mathbf{M} is a $10^5 \times 10^5$ matrix with entries determined by (8), using the discrete derivative formulas above. \mathbf{M} does not change over time.

We used the Crank–Nicolson method (2) to solve (9). The Crank–Nicolson discretization is

$$\frac{\mathbf{P}^{n+1} - \mathbf{P}^n}{dt} = \mathbf{M} \frac{\mathbf{P}^{n+1} + \mathbf{P}^n}{2}$$

where n is the discrete time index, proportional to the number of interactions between immune cells and target cells. This is a linear equation with variables \mathbf{P}^{n+1} and \mathbf{P}^n . It can be written in the form of $(\mathbf{I} - \frac{\mathbf{M}}{2} dt) \mathbf{P}^{n+1} = (\mathbf{I} + \frac{\mathbf{M}}{2} dt) \mathbf{P}^n$, where \mathbf{P}^{n+1} is solved for and normalized to 1 at each step. The initial values of the probability distribution $P(m, \beta, 0)$ were set to be uniform among the internal grid points and zero at the boundaries. The time step was set as $dt = 0.001$.

B.5. Interpretation of immune cell heterogeneity. Figure S2A shows one example of the learned distribution of m and β parameters, derived from numerical integration of the Fokker-Planck equation. As expected, the distribution of learned parameters is centered around the true ones. Due to finite memory, some immune cells that have recently interacted with targets that provide unusually low levels of stimulus underestimate the true signal mean in this environment, for example.

Heterogeneous adaptation to the environment results in heterogeneous responses to target cells. Figure S2B shows the probability that an immune cell with particular values of m and β responds to a set of aberrant target cells. Lower values of m and β result in estimated signal distributions $P_r(x)$ that are more concentrated around smaller values of the signal x , making these immune cells more sensitive to aberrant targets (see Equation 3 in the main text). Thus, even for immune cells that express identical sets of receptors, stochastic encounters lead to a range in responsiveness to targets (Fig. S2C). Note that here we describe heterogeneity in average responses across a population of individuals, unlike the stochastic probability of activation for an individual cell shown, for example, in Fig. 3.

Memory length controls the degree of heterogeneity in immune cell adaptation to the environment, and consequently, the degree of heterogeneity in immune cell responses to targets. Figure S3 shows responses of a panel of immune cell populations with different values of the memory parameters κ and α . Shorter memory values result in more heterogeneous responses. Responses against aberrant targets are more reliable for cells with longer memories. But importantly, even fairly short memory lengths result in behavior that is qualitatively similar to immune cells with longer memories and greater precision. Short memory lengths do not necessarily lead to pathological responses such as increased activation against healthy targets or hyposensitivity to strongly activating target cells. And as noted previously, long memories result in slow adaptation to different environments.

C. Comparisons with experimental data. To test the ability of our model to recapitulate experimentally-observed immune cell behaviors, we analyzed three data sets involving both in vitro and in vivo experiments.

First, we considered a recent experiment in which primary human NK cells were exposed to two different types of target cells (3). The target cells were Daudi (a malignant B cell line) cells that were either opsonized with Rituximab, an antibody that can activate NK cells through CD16, or transfected to express MICA, a ligand for the activating NKG2D receptor.

In this experiment, NK cells were presented with two targets in different orders (i.e., Daudi-rituximab/Daudi-rituximab, Daudi-rituximab/Daudi-MICA, Daudi-MICA/Daudi-rituximab, and Daudi-MICA/Daudi-MICA) and it was recorded whether the NK cells killed neither target, the first target only, the second target only, or both target cells. The authors observed that the MICA-expressing target cells appeared to be more strongly activating and that the frequency of NK cell activation depended on past encounters.

We quantified the frequency of each type of interaction by image analysis of Figure 5 of ref. (3). This plot was originally generated using data from 231 pairs of interactions across three independent experiments. We verified that the frequencies that we extracted from the figure matched with numbers given in the text of ref. (3) (e.g., it was noted that 37% of Daudi-rituximab followed by Daudi-MICA interactions resulted in the death of both target cells, in agreement with our measurement from the figure).

To compare our model with the data of ref. (3), we inferred parameters for the signal distribution for each type of target cell (healthy cells, which NK cells encountered prior to exposure to the unhealthy targets, Daudi-rituximab, and Daudi-MICA). For simplicity, we assumed that the means of these signal distributions differed, but that they all had unit standard deviations. We also set the mean of the healthy signal distribution to zero without loss of generality. We then fit the parameters α (the memory length, with κ set to 2α), μ_R (signal mean for Daudi-rituximab), and μ_M (signal mean for Daudi-MICA) *only* to data where NK cells encountered two Daudi-rituximab or two Daudi-MICA targets in a row. This would then allow us to *predict* how NK cells would respond to encounters with mixed targets, Daudi-rituximab/Daudi-MICA or Daudi-MICA/Daudi-rituximab.

Because the expression for the probability of activation in our model is complicated, we used approximate Bayesian computation (ABC) to fit the model parameters to data. This method draws parameter values from a prior distribution and uses them to generate simulated data that is compared with the real data, iteratively moving toward sets of parameters that produce results that are most consistent with the real data. We used a Python implementation of this method called Engine for Likelihood-Free Inference (ELFI) (4). We chose a uniform prior distribution between 1 and 50 for α , and Gamma prior distributions with shape and scale parameters $k = 2$ and $\theta = 2$ for μ_R and μ_M . We used the sum of the squared differences between the simulated and experimental response frequencies as our metric for assessing model fit to data. The best fit parameters were $\alpha = 21.5$, $\mu_R = 2.2$, $\mu_M = 3.2$. Although we only used data from NK cell encounters with two Daudi-rituximab or two Daudi-MICA targets, the model accurately predicts the outcome of Daudi-rituximab/Daudi-MICA and Daudi-MICA/Daudi-rituximab interactions. The total squared error of the best-fit model is 0.08.

As a reference, we also fit a simple ‘null’ model that includes constant probabilities that an NK cell responds to Daudi-rituximab or Daudi-MICA, p_R or p_M , respectively. When fit against *all* the data, we find $p_R = 0.33$ and $p_M = 0.77$. The null model fits the data with a total mean squared error of 0.15. Thus, this null model has a higher error than for our model, even though our model was only fit to a subset of the data.

Second, we analyzed data from an experiment in which IL-2-activated primary NK cells were presented with HeLa-CD48 targets (5). HeLa-CD48 cells are HeLa cells transfected to express CD48, a ligand for the activating receptor 2B4. In this experiment, NK cells encountered up to six HeLa-CD48 targets sequentially, and the fraction of encounters in which NK cells killed their targets was recorded. The authors observed that the fraction of targets killed declined with subsequent encounters.

To quantify data from this experiment, we performed image analysis of Figure 3 of ref. (5). We extracted the total frequency of target cells killed at each encounter number. The number of times that each number of serial encounters was observed was provided in ref. (5) (e.g., 347 encounters with a first target, 250 encounters with a second target, and so forth), allowing us to compute the uncertainty in the fraction of targets that would be killed at each step due to finite sampling.

To test our ability to recover these results, we fit the memory length α and the mean of the signal distribution for the HeLa-CD48 cells, μ_H , to the experimental data. As above, we assumed that the healthy signal distribution was normal with zero mean and unit variance, and we also fixed the variance of the distribution for HeLa-CD48 cells to one. We used the same prior distributions for α and μ as above. Our metric for assessing model fit was the squared difference between the simulated and experimental response frequencies, normalized by the uncertainty in the experimental response frequencies due to finite sampling. Here we normalized these quantities to prevent the results from being strongly biased by response frequency data after a large number of encounters, which were computed from a relatively small number of individual cells. We used $p(1-p)/n$ as the experimental uncertainty in response frequencies, where p is the empirical response frequency and n is the number of individual cells from which this frequency was determined. We found a set of best-fit parameters $\alpha = 25.3$ and $\mu_H = 4.7$, which lead to an excellent fit with the experimental data. Remarkably, we infer almost exactly the same value of the memory length α as we inferred for the first experiment. Here, a ‘null’ model with a constant probability of response will of course be unable to reproduce the observed decline in response frequency as an individual NK cell encounters more target cells.

Finally, we analyzed data from an experiment in which NK cells were transferred between MHC-deficient mice and normal mice (6). After adaptation to the new host, transferred NK cells were then measured for the production of IFN $_{\gamma}$ in response to anti-NK1.1 stimulation. The authors observed that NK cell responses depend on the host mouse: MHC-deficient mice were hyporesponsive compared to normal mice. Mature NK cells transferred from one type of mouse to another adapted over the course of several days to respond like those that had originally developed in the mouse to which they were transferred, demonstrating the plasticity of NK cell responses.

As in the previous two data sets, data from this experiment was extracted through image analysis of Figure S2 of ref. (6). This figure contained data pooled from seven independent experiments.

To compare our model with this data, we fit the mean of the signal distribution in the MHC-deficient mice μ_D and the anti-NK1.1 stimulus μ_A , assuming that the signal distribution in normal mice follows a normal distribution with mean zero. As

in the previous cases, we fixed the standard deviations of the signals in each environment to one. Here, we fit these parameters simply based on the typical probability of response to anti-NK1.1 stimulation for NK cells from normal or MHC-deficient environments, yielding $\mu_A = 1.5$ and $\mu_D = 0.6$. For consistency with the inferred memory parameters in the past two experiments, we set $\alpha = 20$. Assuming roughly one NK-target cell interaction per hour, these parameters fit very well with the experimental data.

D. Extended model with multiple receptors/ligands.

D.1. Model definition. We developed our model with multiple receptors and ligands by assuming that each immune cell with a certain complement of receptors receives a net signal that is the sum of the signal received by each of its individual receptors, described by

$$x = \sum_{i=1}^{n_r} a_i e_i x_i. \quad [10]$$

In this equation, for each individual immune cell, $e_i = 1$ if the i th receptor is expressed by that immune cell and $e_i = 0$ otherwise. We set $a_i = 1$ if the i th receptor is activating and $a_i = -1$ if it is inhibitory. To roughly compare with single cell studies on receptor composition and with the number of distinct receptors that are used by NK cells, we used $n_r = 32$ pairs of receptors and ligands, with 10 activating receptors and 22 inhibitory receptors. We also repeated the simulations described in the main text with different numbers of receptor/ligand pairs, finding qualitatively similar results in all cases where the number of receptor/ligand pairs becomes large ($\gtrsim 10$).

For each individual cell, parameters of the net signal distribution are learned and updated in exactly the same way as for the one-dimensional model we first introduced. We also use the same conditions for activation as in the one-dimensional case.

D.2. Alternative model with real-valued receptor expression. In addition to models with binary receptor expression ($e_i \in \{0, 1\}$), we also considered models in which receptor expression varied continuously. However, consistent with experimental observations, individual immune cells do not always express every receptor at some level. In other words, for some immune cells receptor expression was set exactly to zero, as in the binary case above. For immune cells that do express a particular receptor, the expression level e_i was taken to be a real number chosen from a Gaussian distribution with mean 1 and standard deviation 0.2. We truncated the distribution at zero so that all expression levels e_i were nonnegative. The net signal for immune cells with continuous variation in receptor expression was then computed using (10) as for immune cells with binary receptor expression.

D.3. Simulations of immune cell populations. To check how populations of immune cells would respond to different target cells, we first generated sets of immune cells with different patterns of receptor expression: sparse and dense. In both cases, we called the probability for each immune cell to express a particular receptor p , with $p = 0.2$ in the sparse case and $p = 0.8$ in the dense case. The choice of $p = 0.2$ fits well with the distribution of the number of receptors expressed on individual NK cells in single cell studies (7, 8). For each individual cell, we set $e_i = 1$ if the i th receptor is expressed for the binary model and e_i chosen randomly as described above for the continuous model. If a particular receptor is not expressed, we set $e_i = 0$. For our simulations, we generated 10^4 such individual cells with random receptor expression.

We began our simulations by having each of these immune cells interact with 200 healthy targets. For each healthy target, the level of expression for each ligand was drawn independently from a Gaussian distribution with mean 3 and standard deviation 1. This choice occasionally led to values of ligand expression < 0 for some rare cells. We subsequently repeated our simulations with ligand expression values truncated at zero and found no difference in the results presented below.

After ‘‘pre-training’’ the immune cell population, we considered multiple models for aberrant target cells with different perturbations to surface ligand expression. First, we assumed that changes in ligand expression could either be sparse (concentrated on a single ligand, with normal expression for all other ligands) or dense (perturbation of all ligands, with the amount by which each ligand is modified chosen by generating a random n_r -dimensional vector with Gaussian entries and taking the relative magnitude of each). Second, we considered that perturbations could be either ‘pure’ (i.e., every change increases expression of an activating ligand or decreases expression of an inhibitory ligand) or ‘hidden.’ In the latter case, for sparse perturbations, we selected another ligand at random and increased or decreased its average expression on aberrant target cells such that the average net change in activating - inhibitory ligand expression was zero. For dense hidden perturbations, we first randomly partitioned the group of all ligands into two equally-sized sets. We then chose the sign of perturbations to favor activation for one of these subsets of ligands and inhibition for the other. We then generated 10^4 aberrant target cells by selecting expression levels for each individual ligand from the same distribution as for healthy cells, but with the mean expression values shifted following the perturbations described above.

We tested a variety of total magnitudes for perturbations, ranging from 1 to 9. Unless otherwise noted, figures and results reported in the paper use a perturbation magnitude of 5. We found that this middle range is best able to draw out differences between populations of immune cells with sparse versus dense receptor expression. If perturbations are too small, then the aberrant cells are highly similar to healthy ones and consequently very difficult to distinguish. On the other hand, if perturbations are too large, then the aberrant cells are so different from healthy cells that all immune cells with the proper receptor(s) will always be able to recognize aberrant targets.

We tested the performance of different populations of immune cells by measuring the fraction of the full set of 10^4 aberrant target cells that each individual immune cell recognizes, following (Equation 3 in the main text) with $\theta = 0.01$. We plotted these results for different types of perturbations for populations of immune cells with sparse and dense receptor expression in

Fig. 6 for binary receptor expression and **Fig. S8** for continuous receptor expression. As described in the main text, immune cells with sparse receptor expression are best able to recognize a wide variety of targets.

We further characterized which individual immune cells were best able to recognize particular target cells. To do this, we computed the inner product between the vector of immune receptor expression and the vector of ligand perturbations, multiplied by the sign of the interaction (-1 for inhibitory ligands and $+1$ for activating ligands). As anticipated, we found that immune cells that express a particular complement of receptors that is well-aligned with perturbations in ligand expression on aberrant target cells are best able to distinguish them (**Fig. 6**, **Fig. S9**).

E. Comparison with explicitly time-dependent inference methods. The model that we described above is able to capture signals in changing environments through the use of finite, fixed memory parameters α and κ . However, this approach does not explicitly model a time-varying signal distribution (for example, one characterized by a time-dependent mean μ_t and standard deviation σ_t). One may consider how the approach that we have used would compare with inference methods that attempt to estimate parameters for a time-varying distribution.

E.1. Dynamic Bayesian network estimation. To develop such a model, we used a dynamic Bayesian network (DBN) to estimate time-dependent signal means and precisions from data. In this framework, we assume that data is drawn from a normal distribution with mean μ_t and precision λ_t . For mathematical convenience, we will use a change of variables to work with the natural logarithm of the precision $\theta_t = \ln(\lambda_t)$ instead of λ_t directly. The posterior probability for the parameters μ_t, θ_t conditioned on all past data $x_{1:t} = \{x_1, x_2, \dots, x_t\}$ is then

$$P(\mu_t, \theta_t | x_{1:t}) = P(x_t | \mu_t, \theta_t) \int \int d\mu_{t-1} d\theta_{t-1} \frac{1}{Z} \times P(\mu_t, \theta_t | \mu_{t-1}, \theta_{t-1}) P(\mu_{t-1}, \theta_{t-1} | x_{1:t-1}),$$

where Z is a normalization factor.

We also need to define how the (unobserved) signal mean and log-precision vary over time. Here for simplicity we will assume that both variables are related to each other by independent normal distributions, that is,

$$P(\mu_t, \theta_t | \mu_{t-1}, \theta_{t-1}) \propto e^{-\frac{\kappa_D}{2}(\mu_t - \mu_{t-1})^2} e^{-\frac{\kappa_D}{2}(\theta_t - \theta_{t-1})^2}. \quad [11]$$

Here κ_{DBN} is a parameter that describes how closely-related signal means and precisions from adjacent time points are to one another. In other words, this parameter describes a natural ‘memory length’ for signals in the environment.

E.2. Derivative of posterior derivation. To model the time-varying signal distribution, we used a dynamic Bayesian network (DBN) to estimate the mean $\mu_{1:t}$ and precision $\lambda_{1:t}$ sequences from observed signals sequence $x_{1:t}$. Given the assumptions above, the joint probability for all parameters is

$$P(\mu_{0:t}, \lambda_{0:t}, x_{1:t}) = P(\mu_0, \lambda_0) \prod_{k=1}^t P(\mu_k, \lambda_k | \mu_{k-1}, \lambda_{k-1}) \times \prod_{k=1}^t P(x_k | \mu_k, \lambda_k). \quad [12]$$

$P(\mu_0, \lambda_0)$ is the initial prior distribution for μ and λ at time zero. $P(x_k | \mu_k, \lambda_k)$ describes the how likely x_k is to be observed given μ_k, λ_k , which follows a normal distribution $x_k \sim N(\mu_k, \lambda_k)$. $P(\mu_k, \lambda_k | \mu_{k-1}, \lambda_{k-1})$ describes the relation between parameters in the previous step ($k-1$) and step k .

Without further information, a reasonable restriction is that values of one parameter in adjacent steps should be close to each other. For simplicity, we assumed that the present values satisfy a normal distribution that is centered around the previous value. For the mean and precision, we assume $\mu_k \sim N(\mu_{k-1}, 1/\kappa_D)$ and $\lambda_k \sim N(\lambda_{k-1}, 1/\kappa_D)$, where κ_D is the parameter determines the ‘closeness’ between values in adjacent steps.

The full posterior is

$$P(\mu_{0:t}, \lambda_{0:t} | x_{1:t}) = P(\mu_0, \lambda_0) \prod_{k=1}^t P(\mu_k, \lambda_k | \mu_{k-1}, \lambda_{k-1}, x_{1:k}). \quad [13]$$

The posterior at time t ($t \geq 1$) is

$$\begin{aligned} P(\mu_t, \lambda_t | x_{1:t}) &= \int \int d\mu_{t-1} d\lambda_{t-1} P(\mu_{t-1}, \lambda_{t-1} | x_{1:t-1}) P(\mu_t, \lambda_t | \mu_{t-1}, \lambda_{t-1}, x_{1:t}) \\ &= \int \int d\mu_{t-1} d\lambda_{t-1} P(\mu_{t-1}, \lambda_{t-1} | x_{1:t-1}) P(\mu_t, \lambda_t | \mu_{t-1}, \lambda_{t-1}, x_{1:t-1}) P(x_t | \mu_t, \lambda_t) \\ &= P(x_t | \mu_t, \lambda_t) \int \int d\mu_{t-1} d\lambda_{t-1} P(\mu_{t-1}, \lambda_{t-1} | x_{1:t-1}) P(\mu_t, \lambda_t | \mu_{t-1}, \lambda_{t-1}). \end{aligned} \quad [14]$$

For a given prior $P(\mu_0, \lambda_0)$ and observations $x_{1:t}$, we can numerically solve (14) to calculate the posterior distribution for each step. We use the location of the peak point as the estimated mean and logarithm of precision. As an example, we generated observations $x_{1:t}$ from normal distributions $N(\mu_t, 1/\lambda_t)$ where $\mu_t = 4 \sin \frac{2\pi t}{T}$, $\lambda_t = 1 + \frac{1}{2} \sin \frac{2\pi t}{T}$, with period $T = 100$. With a normal prior distribution, $P(\mu_0) = N(0, 1)$, $P(\lambda_0) = N(2.5, 1)$, we solved (14) in the domain of $\mu \times \lambda = [-10, 10] \times (0, 5]$. The results for different κ_D are shown in **Fig. S11**, where for smaller κ_D (κ_D smaller than the period of true mean and precision), the estimation captures the change of true values. For larger κ_D , the estimation is more stable which is reasonable since this agrees with the definition of κ_D . The step length $dx = 0.1$ which divides the domain into a 200×50 grid.

E.3. Comparison of the update mechanism of our model and the DBN model. Given that the DBN model performs similarly to our simple model, we further analyzed the DBN to understand how parameter updates in this model may be related to our own.

The posterior at time t is given by (14) for the DBN model. This integral is difficult to calculate after the first step ($t \geq 2$), and the posterior cannot be solved analytically. However, since the distribution is convex, we can calculate the maximum of the posterior and explore its connection with the previous distribution. We will use this as an indicator of the DBN model update mechanism. The posterior at time t is

$$P(\mu_t, \lambda_t | x_t) = P_L(x_t | \mu_t, \lambda_t) \iint d\mu_{t-1} d\lambda_{t-1} P_0(\mu_{t-1}, \lambda_{t-1}) P_T(\mu_t, \lambda_t | \mu_{t-1}, \lambda_{t-1}),$$

where

$$P_L(x_t | \mu_t, \lambda_t) = \sqrt{\frac{\lambda_t}{2\pi}} e^{-\frac{1}{2}\lambda_t(x_t - \mu_t)^2}$$

and

$$P_T(\mu_t, \lambda_t | \mu_{t-1}, \lambda_{t-1}) = \sqrt{\frac{\kappa_D}{2\pi}} e^{-\frac{\kappa_D}{2}(\mu_t - \mu_{t-1})^2} \times \sqrt{\frac{\kappa_D}{2\pi}} e^{-\frac{\kappa_D}{2}(\lambda_t - \lambda_{t-1})^2}.$$

At the maximum, $\frac{\partial P(\mu_t, \lambda_t | x_t)}{\partial \mu_t} = 0$, $\frac{\partial P(\mu_t, \lambda_t | x_t)}{\partial \lambda_t} = 0$. Substituting the posterior into the equation, we find

$$\begin{aligned} \mu_t &= \frac{\lambda_t}{\lambda_t + \kappa_D} x_t + \frac{\kappa_D}{\lambda_t + \kappa_D} \frac{\iint \mu_{t-1} P_0 P_T d\mu_{t-1} d\lambda_{t-1}}{\iint P_0 P_T d\mu_{t-1} d\lambda_{t-1}}, \\ \lambda_t &= \frac{1}{2} \left[\langle \lambda_{t-1} \rangle - \frac{(x_t - \mu_t)^2}{2\kappa_D} + \sqrt{\frac{2}{\kappa_D} + \left[\langle \lambda_{t-1} \rangle - \frac{(x_t - \mu_t)^2}{2\kappa_D} \right]^2} \right], \end{aligned}$$

where $\langle \lambda_{t-1} \rangle = \frac{\iint \lambda_{t-1} P_0 P_T d\mu_{t-1} d\lambda_{t-1}}{\iint P_0 P_T d\mu_{t-1} d\lambda_{t-1}}$. Calculation details are presented in the following section. For the update of μ , the maximum of the posterior is determined by the linear combination of the new observation and previous distributions. The update of the precision is complicated, but for large κ_D the last term under the square root can be simplified as

$$\begin{aligned} & \sqrt{\frac{2}{\kappa_D} + \left[\langle \lambda_{t-1} \rangle - \frac{(x_t - \mu_t)^2}{2\kappa_D} \right]^2} \\ &= \sqrt{\langle \lambda_{t-1} \rangle^2 - 2\langle \lambda_{t-1} \rangle \frac{(x_t - \mu_t)^2}{2\kappa_D} + \left(\frac{(x_t - \mu_t)^2}{2\kappa_D} \right)^2 + \frac{2}{\kappa_D}} \\ &\approx \langle \lambda_{t-1} \rangle + \frac{1}{2\langle \lambda_{t-1} \rangle} \left[-2\langle \lambda_{t-1} \rangle \frac{(x_t - \mu_t)^2}{2\kappa_D} + \left(\frac{(x_t - \mu_t)^2}{2\kappa_D} \right)^2 + \frac{2}{\kappa_D} \right] \\ &\approx \langle \lambda_{t-1} \rangle - \frac{(x_t - \mu_t)^2}{2\kappa_D} + \frac{1}{\kappa_D \langle \lambda_{t-1} \rangle} \end{aligned} \tag{15}$$

The second step above used a Taylor expansion, $\sqrt{x+d} \approx \sqrt{x} + \frac{1}{2\sqrt{x}}d$, $d \ll 1$. Thus the update of the precision can be approximated as

$$\lambda_t = \langle \lambda_{t-1} \rangle + \frac{1}{2\kappa_D} \left[\frac{1}{\langle \lambda_{t-1} \rangle} - (x_t - \mu_t)^2 \right]. \tag{16}$$

In summary, the update of the DBN model is

$$\begin{aligned} \mu_t &= \frac{\lambda_t}{\lambda_t + \kappa_D} x_t + \left(1 - \frac{\lambda_t}{\lambda_t + \kappa_D} \right) \frac{\iint \mu_{t-1} P_0 P_T d\mu_{t-1} d\lambda_{t-1}}{\iint P_0 P_T d\mu_{t-1} d\lambda_{t-1}}, \\ \lambda_t &= \langle \lambda_{t-1} \rangle + \frac{1}{2\kappa_D} \left[\frac{1}{\langle \lambda_{t-1} \rangle} - (x_t - \mu_t)^2 \right]. \end{aligned} \tag{17}$$

In our model, μ and λ are used to estimate the mean and precision (3), leading to $\hat{\mu} = m$, $\hat{\sigma}^2 = \beta/(\alpha - 1)$. Since the normal-gamma function (distribution of posterior in our model) is convex, we can show that the maximum of the posterior is close to the mean. The posterior maximum of our model is calculated by $\partial P(\mu, \lambda | m, \beta, \alpha, \kappa) / \partial \mu = 0$, $\partial P(\mu, \lambda | m, \beta, \alpha, \kappa) / \partial \lambda = 0$, which gives $\mu = m$, $\lambda = (\alpha - 1/2) / \beta$. The update of the posterior maximum parameter values is $\mu' = m'$ and $\lambda' = (\alpha - 1/2) / \beta'$. Substituting m' and β' from (6) into this expression, we find $\mu' = \delta_\kappa \mu + (1 - \delta_\kappa) x$ and $1/\lambda' = \delta_\alpha / \lambda + (1 - \delta_\alpha) \delta_\kappa (x - m)^2$, which

is the same as the update of estimated mean and precision since $\sigma^2 = 1/\lambda$, where $\delta_\alpha = (\alpha - 1)/(\alpha - 1/2)$, $\delta_\kappa = \kappa/(\kappa + 1)$. The update of the precision can be approximated by

$$\begin{aligned}\frac{1}{\lambda'} &= \delta_\alpha \frac{1}{\lambda} + (1 - \delta_\alpha) \delta_\kappa (x - m)^2 \\ \lambda' &= \frac{1}{\delta_\alpha \frac{1}{\lambda} + (1 - \delta_\alpha) \delta_\kappa (x - m)^2} \\ &= \frac{\lambda}{\delta_\alpha} \times \frac{1}{1 + \frac{1 - \delta_\alpha}{\delta_\alpha} \delta_\kappa (x - m)^2 \lambda} \\ &\approx \frac{\lambda}{\delta_\alpha} \left[1 - \frac{1 - \delta_\alpha}{\delta_\alpha} \delta_\kappa (x - m)^2 \lambda \right] \\ &\approx \lambda + \frac{1 - \delta_\alpha}{\delta_\alpha} \lambda^2 \left[\frac{1}{\lambda} - \frac{\delta_\kappa}{\delta_\alpha} (x - m)^2 \right].\end{aligned}$$

In summary, the update of our model is

$$\begin{aligned}\mu' &= \delta_\kappa \mu + (1 - \delta_\kappa) x \\ \lambda' &= \lambda + \frac{1 - \delta_\alpha}{\delta_\alpha} \lambda^2 \left[\frac{1}{\lambda} - \frac{\delta_\kappa}{\delta_\alpha} (x - m)^2 \right].\end{aligned}\tag{18}$$

Thus, we see that parameter updates for our model and DBNs, (17) and (18), are quite similar. The main difference is that parameter updates in the DBN model are controlled by κ_D , while in our model these are guided by α, κ . Below, we will explore the relationship between κ_D and α, κ .

E.4. DBN posterior estimates. Here we derive the posterior parameter estimates discussed in the previous section. Partial differentiation of the posterior with respect to μ_t gives

$$\begin{aligned}0 &= \frac{\partial P(\mu_t, \lambda_t | x_t)}{\partial \mu_t} \\ &= \iint d\mu_{t-1} d\lambda_{t-1} P_0(\mu_{t-1}, \lambda_{t-1}) \times \frac{\partial}{\partial \mu_t} [P_L(x_t | \mu_t, \lambda_t) P_T(\mu_t, \lambda_t | \mu_{t-1}, \lambda_{t-1})] \\ &= \iint d\mu_{t-1} d\lambda_{t-1} P_0 \left[\left(\frac{\partial}{\partial \mu_t} P_L \right) P_T + P_L \left(\frac{\partial}{\partial \mu_t} P_T \right) \right] \\ &= \iint d\mu_{t-1} d\lambda_{t-1} P_0 \times [-\lambda_t (\mu_t - x_t) P_L P_T - \kappa_D (\mu_t - \mu_{t-1}) P_L P_T]\end{aligned}$$

We then find

$$(\lambda_t + \kappa_D) \mu_t \iint d\mu_{t-1} d\lambda_{t-1} P_0 P_L P_T = \lambda_t x_t \iint d\mu_{t-1} d\lambda_{t-1} P_0 P_L P_T + \kappa_D \iint d\mu_{t-1} d\lambda_{t-1} \mu_{t-1} P_0 P_L P_T,$$

which can be simplified as

$$\mu_t = \frac{\lambda_t}{\lambda_t + \kappa_D} x_t + \frac{\kappa_D}{\lambda_t + \kappa_D} \frac{\iint \mu_{t-1} P_0 P_T d\mu_{t-1} d\lambda_{t-1}}{\iint P_0 P_T d\mu_{t-1} d\lambda_{t-1}}.$$

For the partial derivative with respect to λ_t , we have

$$\begin{aligned}0 &= \frac{\partial P(\mu_t, \lambda_t | x_t)}{\partial \lambda_t} \\ &= \iint d\mu_{t-1} d\lambda_{t-1} P_0(\mu_{t-1}, \lambda_{t-1}) \times \frac{\partial}{\partial \lambda_t} [P_L(x_t | \mu_t, \lambda_t) P_T(\mu_t, \lambda_t | \mu_{t-1}, \lambda_{t-1})] \\ &= \iint d\mu_{t-1} d\lambda_{t-1} P_0 \left[\left(\frac{\partial}{\partial \lambda_t} P_L \right) P_T + P_L \left(\frac{\partial}{\partial \lambda_t} P_T \right) \right] \\ &= \iint d\mu_{t-1} d\lambda_{t-1} P_0 \left[\frac{1}{2} \lambda_t^{-1} - \frac{1}{2} (\mu_t - x_t)^2 P_L P_T - \kappa_D (\lambda_t - \lambda_{t-1}) P_L P_T \right]\end{aligned}$$

This gives

$$0 = \frac{1}{2} \lambda_t^{-1} - \frac{1}{2} (\mu_t - x_t)^2 - \kappa_D (\lambda_t - \frac{\iint \lambda_{t-1} P_0 P_T d\mu_{t-1} d\lambda_{t-1}}{\iint P_0 P_T d\mu_{t-1} d\lambda_{t-1}}).$$

Considering $\lambda_t > 0$, the solution is

$$\lambda_t = \frac{1}{2} \left[\langle \lambda_{t-1} \rangle - \frac{(x_t - \mu_t)^2}{2\kappa_D} + \sqrt{\frac{2}{\kappa_D} + \left[\langle \lambda_{t-1} \rangle - \frac{(x_t - \mu_t)^2}{2\kappa_D} \right]^2} \right],\tag{19}$$

where $\langle \lambda_{t-1} \rangle = \frac{\iint \lambda_{t-1} P_0 P_T d\mu_{t-1} d\lambda_{t-1}}{\iint P_0 P_T d\mu_{t-1} d\lambda_{t-1}}$.

E.5. Connection between memory parameters in our model and transition factors in dynamic Bayesian networks. Given the similarity between parameter updates in our model and DBNs, we worked to explore the connection between the memory parameters κ and α in our model and the transition factor κ_D in DBNs. In particular, we searched for relationships between κ_D and κ and α such that both sets of models behave similarly.

Since the posterior of DBN model is difficult to calculate, here we used a variational Bayesian (VB) approach to approximate the posterior (9). Our analysis below closely follows work described in ref. (10). VB methods are often used to treat complicated posteriors in Bayesian inference. Instead of calculating the true posterior, VB methods introduce a new distribution Q to approximate the posterior $P(\mu_t, \theta_t | x_{1:t})$. The Kullback–Leibler (KL) divergence, $D_{KL} = \int Q \ln(Q/P)$, describes the ‘closeness’ of P and Q . We used the mean-field form of VB in which Q is factorized into single-variable factors, $Q(Z) = \prod_i q(z_i)$, where $Z = \{z_1, \dots, z_i\}$ are the unknown variables. Every $q(z_i)$ is normalized such that $\int dz_i q(z_i) = 1$. It has been shown that the KL divergence achieves its minimum value when $q(z_i) = \frac{1}{Z} \exp \langle E(z_i, \bar{z}_i, D) \rangle_{Q(\bar{z}_i)}$ for all the variables (9), where D is the observed data, \bar{z}_i means all other variables except z_i , $E(z_i, \bar{z}_i, D) = \ln P(Z, D)$ and $\langle E(z_i, \bar{z}_i, D) \rangle_{Q(\bar{z}_i)}$ means integrating $E(z_i, \bar{z}_i, D)$ and $Q(\bar{z}_i)$ over \bar{z}_i .

In our case, the factorization can be written as $Q(\mu_t, \theta_t) = q(\mu_t)q(\theta_t)$. The optimal Q is achieved when $\ln q(\mu_t) = \int q(\theta_t) \ln P(\mu_t, \theta_t, x_t) d\theta_t + \ln Z_{\mu_t}$ and $\ln q(\theta_t) = \int q(\mu_t) \ln P(\mu_t, \theta_t, x_t) d\mu_t + \ln Z_{\theta_t}$. The Z_{μ_t} and Z_{θ_t} are for normalization. Below, we show that with a normal distribution for $q(\mu_{t-1})$, the optimal $q(\mu_t)$ would be also normal. Thus, if we use a normal distribution as the prior for $P_0(\mu_0)$ in $P_0(\mu_0, \lambda_0) = P_0(\mu_0)P_0(\lambda_0)$, we will obtain normally distributed $q_{t=1, \dots, T}(\mu_t)$.

We assume that $q(\mu_{t-1}) \sim N(m_{t-1}, 1/r_{t-1})$, thus $P(\mu_{t-1}, \lambda_{t-1})$ is replaced with $q(\mu_{t-1})q(\lambda_{t-1})$. The posterior $P(\mu_t, \lambda_t | x_t)$ can be written as

$$\begin{aligned} P(\mu_t, \lambda_t | x_t) &= P(x_t | \mu_t, \lambda_t) \int \int d\mu_{t-1} d\lambda_{t-1} P(\mu_{t-1}, \lambda_{t-1}) P(\mu_t, \lambda_t | \mu_{t-1}, \lambda_{t-1}) \\ &\propto \sqrt{\lambda_t} e^{-\frac{1}{2} \lambda_t (x_t - \mu_t)^2} \int \int d\mu_{t-1} d\lambda_{t-1} e^{-\frac{1}{2} [r_{t-1} (\mu_{t-1} - m_{t-1})^2 + \kappa_D (\mu_t - \mu_{t-1})^2]} \times e^{-\frac{\kappa_D}{2} (\lambda_t - \lambda_{t-1})^2} q(\lambda_{t-1}) \\ &\propto \sqrt{\lambda_t} e^{-\frac{1}{2} \lambda_t (x_t - \mu_t)^2} \times e^{-\frac{1}{2} \frac{\kappa_D r_{t-1}}{\kappa_D + r_{t-1}} (\mu_t - m_{t-1})^2} \times \int d\lambda_{t-1} e^{-\frac{\kappa_D}{2} (\lambda_t - \lambda_{t-1})^2} q(\lambda_{t-1}). \end{aligned} \quad [20]$$

Since $\ln P(\mu_t, \lambda_t, x_t) = \ln P(\mu_t, \lambda_t | x_t) - \ln P(x_t)$ we first calculate the log-posterior

$$\ln(P(\mu_t, \lambda_t | x_t)) \propto \frac{1}{2} \ln(\lambda_t) - \frac{1}{2} \lambda_t (x_t - \mu_t)^2 - \frac{1}{2} \frac{\kappa_D r_{t-1}}{\kappa_D + r_{t-1}} (\mu_t - m_{t-1})^2 + \ln \left(\int d\lambda_{t-1} e^{-\frac{\kappa_D}{2} (\lambda_t - \lambda_{t-1})^2} q(\lambda_{t-1}) \right). \quad [21]$$

For $q(\mu_t)$,

$$\begin{aligned} \ln(q(\mu_t)) &= \int q(\lambda_t) \ln P(\mu_t, \lambda_t, x_t) d\lambda_t + \ln Z_{\mu_t} \\ &\propto -\frac{1}{2} \langle \lambda_t \rangle (x_t - \mu_t)^2 - \frac{1}{2} \frac{\kappa_D r_{t-1}}{\kappa_D + r_{t-1}} (\mu_t - m_{t-1})^2, \end{aligned} \quad [22]$$

where $\langle \lambda_t \rangle = \int \lambda_t q(\lambda_t) d\lambda_t$. Here $q(\mu_t)$ is a normal distribution. We can therefore calculate the maximum (and mean) m_t and precision r_t by calculating the first and second derivatives of (22), $\frac{d \ln q(\mu_t)}{d \mu_t} = 0$ ($\mu_t = m_t$) and $\frac{d^2 \ln q(\mu_t)}{d^2 \mu_t} = -r_t$. Substituting $\ln q(\mu_t)$ into these two equations, we have

$$\begin{aligned} m_t &= \frac{\langle \lambda_t \rangle x_t + \frac{\kappa_D r_{t-1}}{\kappa_D + r_{t-1}} m_{t-1}}{\langle \lambda_t \rangle + \frac{\kappa_D r_{t-1}}{\kappa_D + r_{t-1}}}, \\ r_t &= \langle \lambda_t \rangle + \frac{\kappa_D r_{t-1}}{\kappa_D + r_{t-1}}. \end{aligned} \quad [23]$$

The update for the mean is similar to our model,

$$\langle \mu \rangle = m, \quad m' = \frac{\kappa m + x}{\kappa + 1}.$$

When then find that if κ is set to

$$k = \frac{\kappa_D r_t}{\kappa_D + r_t} \frac{1}{\langle \lambda_t \rangle}, \quad [24]$$

these models should give similar results. We calculated r_t as the inverse of the variance of the mean $r_t = 1/\text{Var}(\int P(\mu_t, \lambda_t | x_t) d\lambda)$, and $\langle \lambda_t \rangle$ as the mean of the precision, $\langle \lambda_t \rangle = \iint d\mu_t d\lambda_t \lambda_t P(\mu_t, \lambda_t | x_t)$. Using the mean of sequence k defined by (24) as

$$\kappa = \frac{1}{T} \sum_{t=1}^T k = \frac{1}{T} \sum_{t=1}^T \frac{\kappa_D r_t}{\kappa_D + r_t} \frac{1}{\langle \lambda_t \rangle},$$

our model would have performance similar to DBNs in estimating the signal distribution in a changing environment (Fig. S12). Thus, within these approximations, it is possible to find a value for κ in our model which leads to performance similar to that of a DBN with a particular value of κ_D .

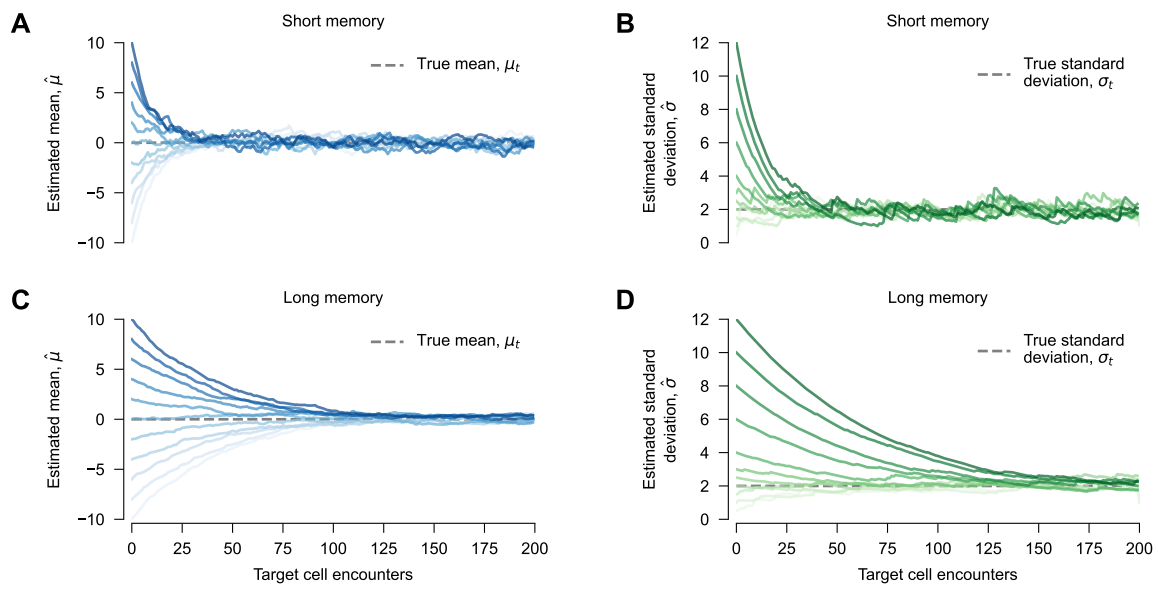


Fig. S1. Immune cells with shorter memory lengths exhibit faster but noisier adaptation. The estimated mean $\hat{\mu}$ and standard deviation $\hat{\sigma}$ approach the true mean ($\mu_t = 0$) and true standard deviation ($\sigma_t = 2$) in a finite number of encounters, regardless of the initial value and memory length. When the memory length is shorter ($\kappa = 10$ and $\alpha = 5$, panels **A** and **B**), estimated values approach the true ones faster than for immune cells with longer memory lengths ($\kappa = 40$ and $\alpha = 20$, panels **C** and **D**). However, adaptation is less precise when the memory length is shorter. **A**, Convergence to true mean ($\mu_t = 0$) from various initial values of m . The initial value of $\beta = (\alpha - 1)$ is the same in all cases. **B**, Convergence to the true standard deviation ($\sigma_t = 2$) from various initial values of β . The initial value of $m = 0$ is the same in all cases. The memory length for both **A** and **B** is set by $\kappa = 10$ and $\alpha = 5$. (**C**, **D**) Display convergence toward the true mean and standard deviation as in (**A**, **B**). Initial parameters are the same as those in (**A**, **B**), except for longer memory lengths $\kappa = 40$ and $\alpha = 20$.

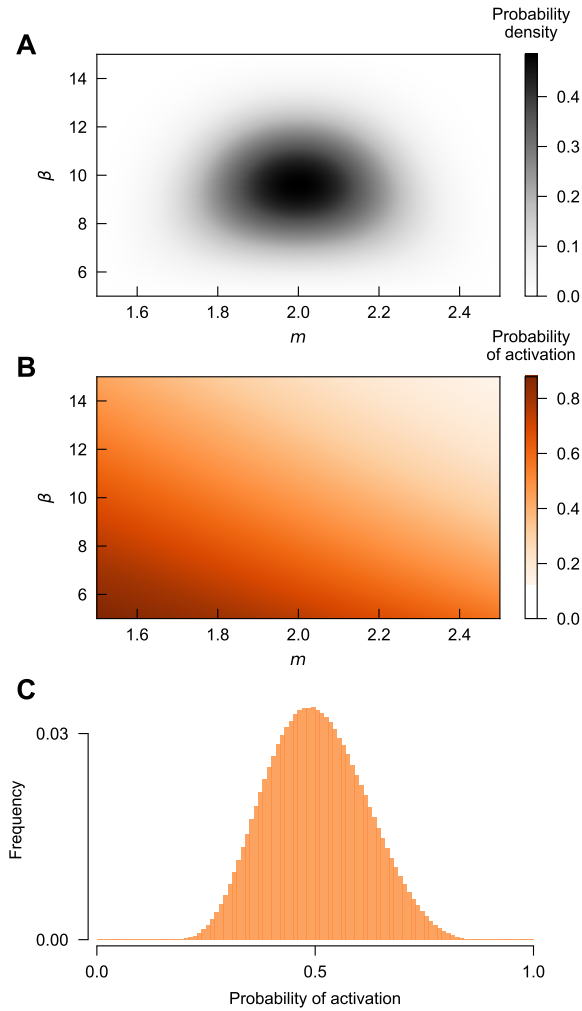


Fig. S2. Steady state distribution of immune cell adaptation and responsiveness due to finite memory. **A**, Steady state joint probability distribution of learned (m, β) parameters, estimated by numerical solution of the Fokker-Planck equation. Here the true signal mean is $\mu_t = 2$, and the true standard deviation is $\sigma_t = 1$. Here we used memory values $\kappa = 20$ and $\alpha = 10$. We observe that the learned parameters are concentrated around the true ones ($\hat{\mu} = m \approx 2$, $\hat{\sigma}^2 = \beta/(\alpha - 1) \approx 1$). **B**, Probability of activation against aberrant targets with signal distribution $\mu_a = 4.5$ and $\sigma_a = 1$ as a function of (m, β) . Immune cells that happen to have lower values of both m and β have higher confidence that target cell signals should be more inhibitory, and thus they are more sensitive to stimulus from aberrant targets. **C**, Net distribution of probabilities of recognizing aberrant targets, obtained by combining the distributions in **A** and **B**. A wide range of responses exists: some immune cells have a high probability of recognizing aberrant targets, while others are relatively unlikely to respond. Heterogeneity can be less pronounced if the signal from aberrant target cells is either much closer or much further away from the healthy cell distribution (see **Fig. S3** for additional examples).

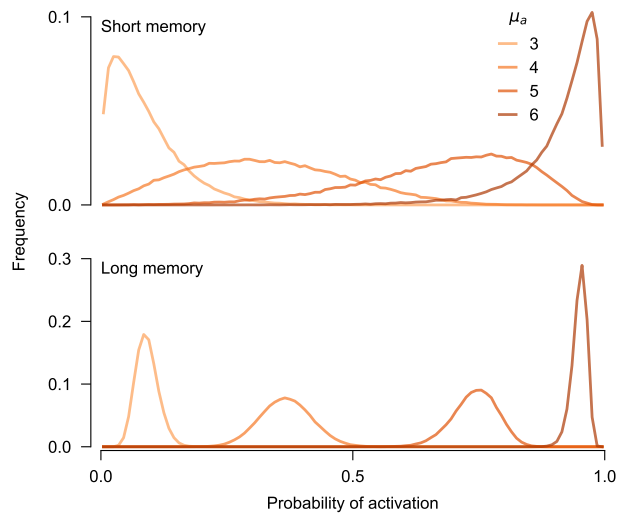


Fig. S3. Immune cell responses are diverse, but follow predictable trends depending on the level of stimulation received from target cells. Distribution of probabilities of activation for populations of immune cells with short ($\kappa = 10$, $\alpha = 5$, *top*) and long ($\kappa = 100$, $\alpha = 50$, *bottom*) memory lengths. Each immune cell population consisted of 10^5 cells with random initial values of m and β uniformly distributed between $[0, 5]$ and $[0, 40]$, respectively. Immune populations were evolved through 10^3 target cell encounters with signal distributions $\mu_t = 2$ and $\sigma_t = 1$. Immune cells were then tested for their probability to activate against aberrant targets that provide different levels of stimulus ($\mu_a = (3, 4, 5, 6)$). Here $\sigma_a = 1$ in all cases. For both shorter and longer memory lengths, the probability of activation is low when μ_a is close to μ_t and high when $\mu_a \gg \mu_t$. However, the spread in activation probabilities is significantly larger for immune cell populations with short memories.

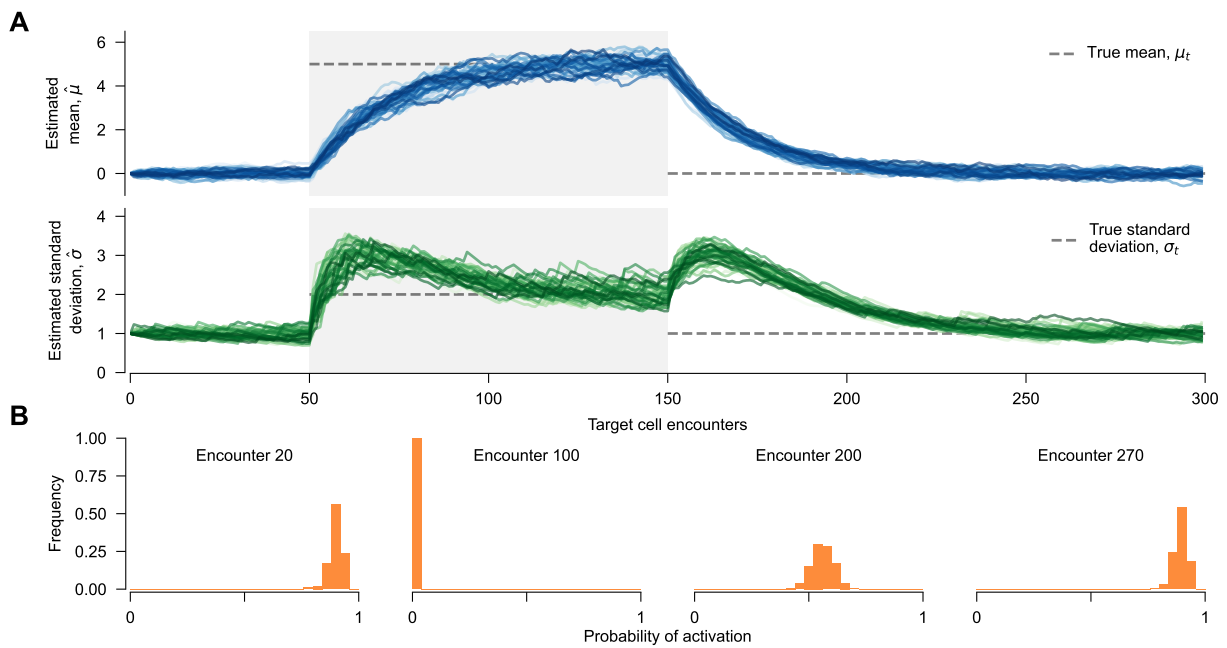


Fig. S4. Finite memory of past interactions with target cells results in heterogeneous immune cell behaviors. **A**, Adaptation of 500 immune cells to changing environments, following the same conventions as in **Fig. 3**. Even though all immune cells start with the same initial values ($m = 0$ and $\beta = 9$, with memory parameters $\kappa = 20$ and $\alpha = 10$), the learned distribution for each immune cell evolves differently over time. This is due to the stochastic nature of signals from target cells, and the finite memory length of past target encounters. **B**, The heterogeneity of learned signal distributions also leads to heterogeneous immune cell responses, characterized by the probability of activation against an aberrant target cell. The signal distribution for aberrant targets is the same as in **Fig. 3**.

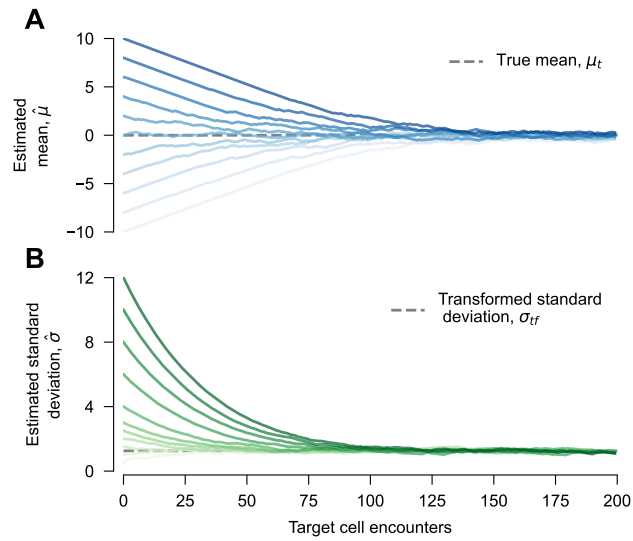


Fig. S5. Immune cell adaptation is more gradual with signal saturation. The estimated mean (A) and standard deviation (B) converge toward the true mean and transformed standard deviation in a model with signal saturation. Parameters and initial conditions are the same as those in Fig. 2. Signal saturation changes the expected standard deviation of the signal when the model is perfectly adapted because large deviations are suppressed (see Methods). We have thus replaced the true standard deviation σ_t with the transformed standard deviation σ_{tf} in B.

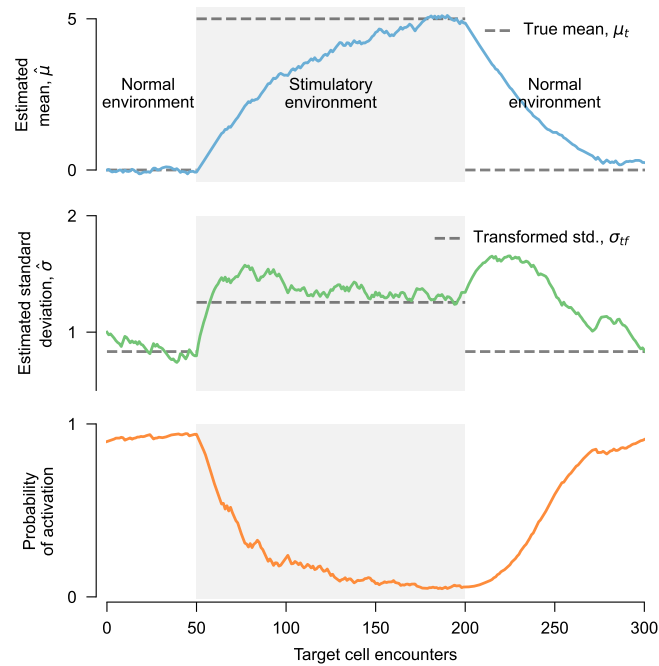


Fig. S6. Immune cell adaptation is more gradual with signal saturation. As in Fig. 3, an immune cell adapts to different signal distributions in 'normal' and 'stimulatory' environments in a model with signal saturation. Adaptation is more gradual due to signal saturation, so we have extended the number of encounters in the stimulatory environment to allow for complete adaptation. As in Fig. 3, the immune cell is able to reliably activate against aberrant targets when it is adapted to a normal environment, but it loses this ability after prolonged exposure to a stimulatory environment. Compared to a model without signal saturation (Fig. 3), adaptation is more gradual and less noisy. In addition, the immune cell retains the ability to activate against aberrant targets for a longer time after being placed in the stimulatory environment.

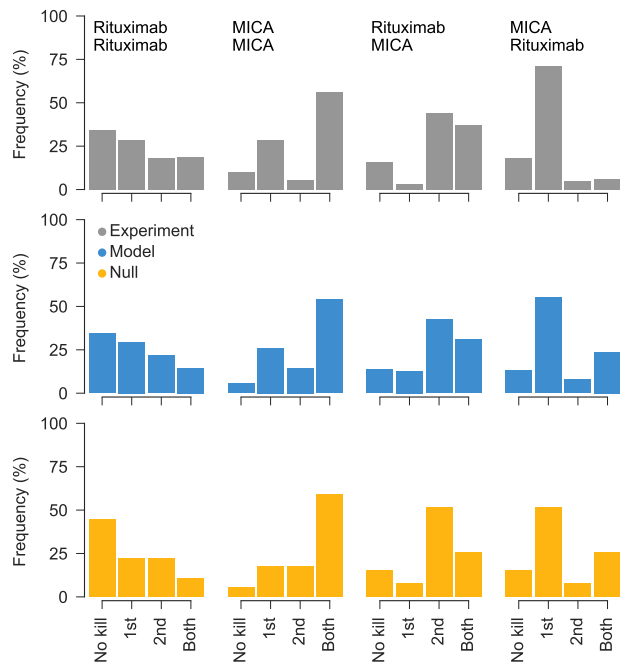


Fig. S7. Our model provides a better fit to data compared to a simple, 'null' model without memory. As in Fig. 4 in the main text, we show experimental data (3) and model predictions for sequential encounters with two types of target cells, labeled Rituximab and MICA, respectively. Here we have also included the best-fitting null model without memory for NK cell killing, in which the probability of killing when encountering each different type of target cell is constant. In experiments where NK cells encounter two identical target cells in sequence, the probability that they kill the second target cell only is lower than the probability that they kill the first target cell only. This phenomenon is naturally recovered in our model with memory, but in the simple null model without memory, these probabilities are always equal. Similarly, the null model is unable to reproduce the observation that the probability of killing both targets is higher if the less stimulatory target (labeled Rituximab here) is encountered first. See Methods for additional details on experimental the data set and model fitting.

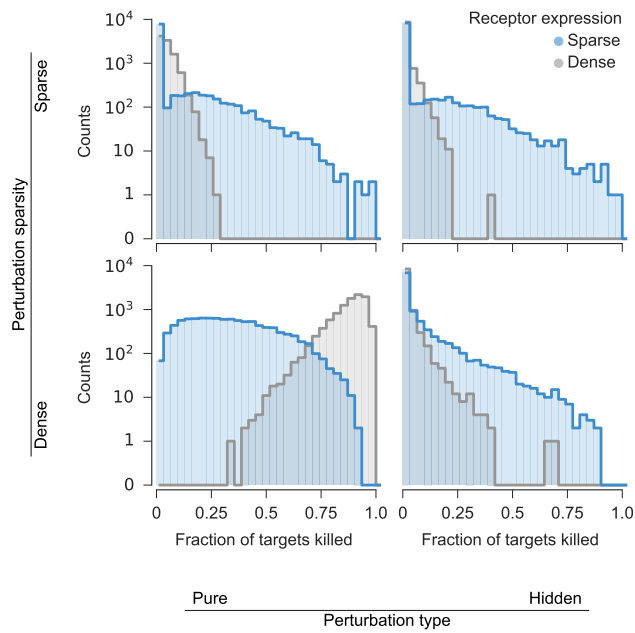


Fig. S8. Immune cells with sparse receptor expression are better able to detect a wide variety of aberrant target cells. This figure is analogous to Fig. 5 in the main text, but with continuous variation in immune cell receptor expression instead of binary expression patterns. As before, for all types of aberrant cells that we considered, there exist some immune cells with sparse receptor expression that are able to kill a large fraction of aberrant cells that they encounter. In contrast, immune cells that express many receptors are unable to efficiently kill aberrant target cells with 'hidden' perturbations, where some changes in ligand expression promote activation and some promote inhibition. Immune cells with dense receptor expression are also unable to efficiently eliminate aberrant targets that only have unusual levels of expression for one or two ligands.

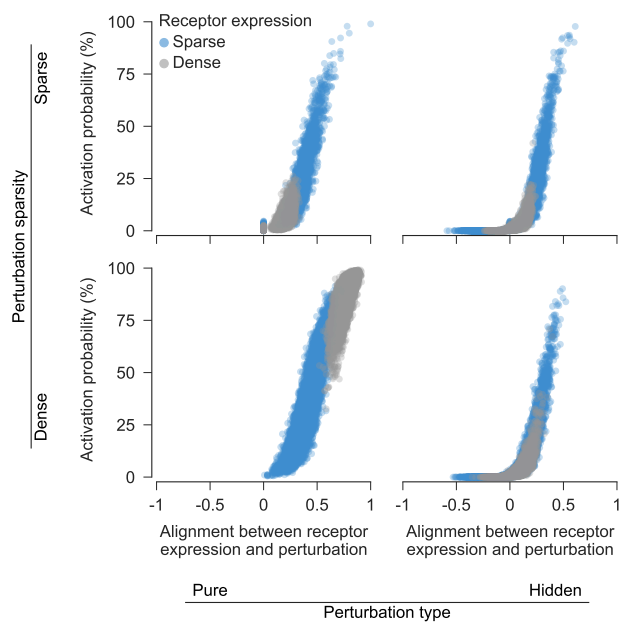


Fig. S9. Immune cells best recognize aberrant targets when they express the specific combination of receptors that recognize ligands with perturbed expression. This figure is analogous to Fig. 6 in the main text, but with continuous variation in immune cell receptor expression instead of binary expression patterns. As shown in the main text, immune cells with sparse receptor expression are more likely to align with changes in ligand expression for a variety of aberrant target cells, allowing them to kill these targets more efficiently. To quantify the alignment between immune cell receptors and perturbations, we treated the perturbations in ligand expression in aberrant cells and receptor expression patterns in immune cells as vectors and computed their normalized inner product.

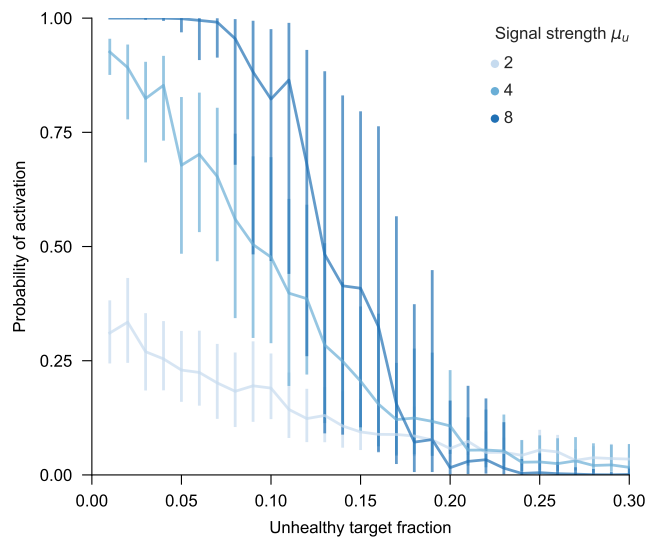


Fig. S10. Immune cells effectively respond to rare unhealthy cells within healthy populations. We considered the ability of immune cells in our model to recognize unhealthy target cells when the population of cells consists of two types: healthy cells, with signal mean $\mu = 0$, and unhealthy cells, for which we considered a range of mean signal strengths μ_{u_i} . For both populations, the signal standard deviation was set to $\sigma = 1$. We allowed immune cells to interact freely with both types of cells, where the probability of encountering each cell type depends on their fraction of the total population (i.e., we assume that the two subpopulations are well-mixed). We find that recognition declines as the fraction of unhealthy cells increases, as measured by the median probability of activation for immune cells. However, due to the stochasticity of target cell encounters, some immune cells in recent contact with healthy cells maintain the ability to kill aberrant targets even when the fraction of unhealthy cells is higher. Together with the median, we also show the interquartile range for immune cell probabilities of activation against unhealthy targets.




Fig. S11. Dynamic Bayesian network (DBN) models estimate the mean and precision of time-varying signals. We tested the ability of DBNs to recover signals with sinusoidally varying mean and precision, $\mu_t = 4 \sin(2\pi t/100)$ and $\lambda_t = 1 + \sin(2\pi t/100)/2$. Here DBNs with values of κ_D (the parameter that controls how strongly correlated estimates of the signal mean and precision should be in time) that are similar or smaller than the period of the true oscillations, $T = 100$, adapt well to moving signals. In general, the precision is more difficult to recover than the mean.




Fig. S12. Similar results for our model and explicitly time-varying DBNs with appropriate parameter choices. Using the conversion between κ_D in the DBN and κ in our model defined in (24), we compared the output from a DBN and our model trained on data from the same model is in Fig. S11. Performance for both methods is similar. When the memory parameters κ_D and κ are close enough to the true time scale of signal fluctuations, both roughly recover underlying parameters. If the memory parameters are too large, then estimates of the mean change too slowly to track the true changes, and the precision is underestimated (correspondingly, the signal variance is overestimated) due to failure to account for the oscillating mean.

References

1. H Risken, *Fokker-planck equation*. (Springer), (1996).
2. J Crank, P Nicolson, A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type in *Mathematical proceedings of the Cambridge philosophical society*. (Cambridge University Press), Vol. 43, pp. 50–67 (1947).
3. K Srpan, et al., Shedding of cd16 disassembles the nk cell immune synapse and boosts serial engagement of target cells. *J. Cell Biol.* **217**, 3267–3283 (2018).
4. J Lintusaari, et al., Elfi: Engine for likelihood-free inference. *J. Mach. Learn. Res.* **19**, 1–7 (2018).
5. I Prager, et al., Nk cells switch from granzyme b to death receptor-mediated cytotoxicity during serial killing. *J. Exp. Medicine* **216**, 2113–2127 (2019).
6. JM Elliott, JA Wahle, WM Yokoyama, Mhc class i-deficient natural killer cells acquire a licensed phenotype after transfer into an mhc class i-sufficient environment. *J. Exp. Medicine* **207**, 2073–2079 (2010).
7. A Horowitz, et al., Genetic and environmental determinants of human nk cell diversity revealed by mass cytometry. *Sci. translational medicine* **5**, 208ra145–208ra145 (2013).
8. DM Strauss-Albee, et al., Human nk cell repertoire diversity reflects immune experience and correlates with viral susceptibility. *Sci. translational medicine* **7**, 297ra115–297ra115 (2015).
9. CW Fox, SJ Roberts, A tutorial on variational bayesian inference. *Artif. intelligence review* **38**, 85–95 (2012).
10. U Beierholm, T Rohe, A Ferrari, O Stegle, U Noppeney, Using the past to estimate sensory uncertainty. *Elife* **9**, e54172 (2020).