Transient Kinetic Proofreading

G. Iyengar[1*] **and M. Perry**[2]

***For correspondence:**
garud@ieor.columbia.edu (GI)

[1]IEOR Department, Columbia University; [2]IEOR Department, Columbia University

Abstract We propose a new stochastic model for understanding the transient kinetic proofreading mechanism in a T-cell. Our model indicates that a stochastic version of absolute ligand discrimination is a consequence of the finite number of receptors on the cell surface; thus, pointing to receptor number control as being critical to T-cell activation. We propose four different metrics to characterize the performance of kinetic proofreading mechanisms. We explore the numerical experiments that explore the trade-offs between speed, specificity, sensitivity, and robustness of T-cell activation as a function of the model parameters. We also consider the impact of receptor clustering on these trade-offs.

1 Introduction

Kinetic proofreading is a mechanism that explains the extremely low probability of error observed in a number of biological processes, e.g. microtubule growth, DNA replication, and the activation of T-cells by antigen presenting cells (***Hopfield, 1974***; ***Ninio, 1975***). Kinetic proofreading exploits the fact that cycling through intermediate complexes on the path to a final product can ensure that the path to undesired products has a significantly lower probability than the path to the desired final product. In this work, we are concerned with kinetic proofreading in switching systems where the cell takes one of two actions based on environmental signals. In particular, we will focus on T-cell activation.

T-cells are a component of the adaptive arm of the immune system (***Perelson and Weisbuch, 1997***). In contrast to the innate immune system, which attacks a broad range of pathogens based on common cell-surface markers, each individual cell in the adaptive immune system is specially tailored to sense and react to antigens associated with a specific type of pathogen. Among other things, a T-cell is responsible for sensing and killing pathogens expressing the T-cell's target antigen on their cell surface. T-cells are able to sense very small differences in the binding energy between a specific "non-self" antigen and T-cell receptors (TCRs) and that of a "self" antigen and the TCR to activate only in the presence of the "non-self" antigen. Kinetic proofreading has been proposed as a mechanism that allows T-cells to discriminate with much better accuracy than that implied by the difference in binding energy alone.

McKeithan (***1995***) proposed the first kinetic proofreading-based model for TCR activation. This model is described by the following set of ODEs ***François and Altan-Bonnet*** (***2016***).

$$
\begin{aligned}
\frac{dC_0}{dt} &= -\left(k_u + k_p\right) C_0(t) + k_b A R, \\
\frac{dC_n}{dt} &= -\left(k_u + k_p\right) C_n(t) + k_p C_{n-1}(t), \quad n = 1, 2, \ldots, N-1, \\
\frac{dC_N}{dt} &= -k_u C_N + k_p C_{N-1},
\end{aligned}
$$

where A denotes the density of the antigen interacting with the TCR, R denotes the density of the TCR, $C_0(t)$ is the density of a TCR-antigen complex at time t, $C_n(t)$ is the density of a TCR-antigen complex with n phosphorylations at time t, k_b is the rate at which antigens bind to TCRs and is assumed to be the same for all antigens, k_u is the rate at which antigens unbind from TCRs and varies between different antigens, and k_p is the rate at which TCR-antigen bindings are phosphorylated. The affinity τ of an antigen is defined as the reciprocal of the unbinding rate k_u, i.e. $\tau = 1/k_u$. The model assumes that a TCR-antigen complex with N phosphorylations is required for T-cell activation.

The steady state value of $C_N \propto L/(k_u)^N$. Suppose a self antigen with affinity $\tau_u^{(s)} = 1/k_u^{(s)}$ and a non-self antigen with affinity $\tau_u^{(o)} = 1/k_u^{(o)}$ are present with the same concentration R. Then the ratio of the steady-state concentration $C_N^{(s)}$ of the self-antigen-TCR complex and the steady-state concentration $C_N^{(n)}$ of the non-self-antigen-TCR complex is $C_N^{(s)}/C_N^{(n)} = (\tau_u^{(s)}/\tau_u^{(o)})^N$. Thus, by repeatedly dissociating the receptor-antigen bindings en route to a final product, the proofreading network ensures that the rate of error decreases exponentially with the size of the network.

The analysis above assumes that the concentration of self and non-self antigens are equal. This is almost never the case. There is evidence of densities as low as ten non-self antigens per cell leading to T-cell activation (***Irvine et al., 2002***), whereas self antigens are far more abundant. T-cells activate even to non-self antigen at very low concentrations, and do not activate even when self antigen is present in very high concentrations. The performance of a kinetic proofreading mechanism can be summarized by its *activation set* $\mathcal{A}$ which is defined as the set of pairs (L, τ) for which the mechanism activates.

The ability of the T-cells to distinguish between self and non-self antigens despite very large discrepancies in concentration is known as *absolute ligand discrimination* (***François and Altan-Bonnet, 2016***). ***François and Altan-Bonnet*** (***2016***) (see, also (***Altan-Bonnet and Germain, 2005***; ***Lalanne and François, 2013***)) proposed and analyzed deterministic mean-field models for absolute ligand discrimination that were based on a feedback loop that regulates the phosphorylation rate of antigen-bound TCRs via a kinase whose activity is itself regulated by the quantity of phosphorylated antigen-bound TCRs. This feedback loop, called *adaptive sorting*, decouples the behavior of the proofreading mechanism from antigen density, thereby achieving absolute ligand discrimination. Although the impact of this feedback loop on T-cell ligand discrimination in vivo is questionable (***Paster et al., 2015***), decoupling the behavior of the proofreading network from antigen affinity is key for absolute discrimination. The activation set $\mathcal{A}^I$ corresponding to *ideal* absolute ligand discrimination is given by $\mathcal{A}^I = \{(L, \tau) : L \geq L_{\min}, \tau \geq \tau_{min}\}$, i.e. the lower boundary of the set $\mathcal{A}$ is given by an asymptote in the L-direction located at $\tau_{\min}$ and an asymptote in the τ-direction located at $L_{\min}$. See Figure 1.

The existing models of absolute ligand distribution are both mean field and steady state. Therefore, these models are not able to represent the trade-off between the delay and other performance measures. Furthermore, these approaches are also not able to model the impact of stochastic noise. In this work, we propose a model that accounts for both transience and stochasticity. We assume the kinetic proofreading mechanism needs to make the activation decision within a specified delay D, and therefore the decision inherently involves stochastic analysis. The performance metric of interest is now the probability $\alpha \in (0, 1)$ by which the mechanism activates by delay D. Consequently, in the stochastic setting, the activation set $\mathcal{A}_\alpha$ is now indexed by a probability α. Following standard practice, we set $\alpha = 0.95$. See Figure 1.

The focus of the existing work on kinetic proofreading is to develop mechanisms that yield error-correction, and absolute ligand discrimination models have focused on the existence of asymptotes in the L and τ directions. We propose that these mechanisms must be evaluated on several other dimensions. In particular, we propose the following four metrics:

(a) *Sensitivity*: The level L_α of the asymptote in the τ_u-dimension, i.e. the lowest density at which

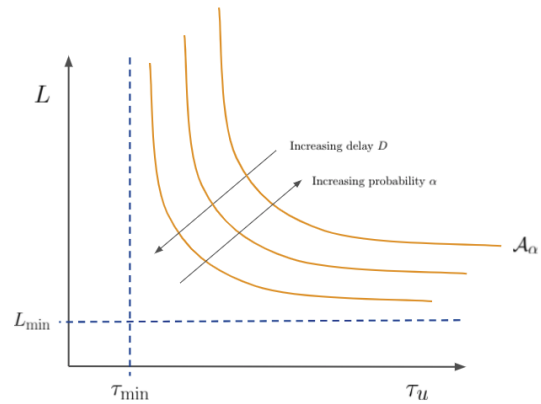

Figure 1. The activation set $\mathcal{A}_\alpha$ bounded by the vertical asymptote at $\tau_{\min}$ and the horizontal asymptote at $L_{\min}$ defining the ideal absolute ligand discrimination activation set $\mathcal{A}^I$. The delay D can be fixed and the probability α of activation increased to move $\mathcal{A}_\alpha$ up and to the right, and the probability of activation can be fixed and the delay can be increased to move $\mathcal{A}_\alpha$ down and to the left.

the cell activates within delay D with probability at least α for any value of the affinity τ_u.

(b) *Specificity*: The gap $\tau_{0.95} - \tau_{0.05}$ between the asymptotes along the L-direction for $\alpha = 0.05$ and $\alpha = 0.95$, i.e. the difference between the affinity values for which the cell almost never activates (i.e. $\alpha = 0.05$) and almost always activates (i.e. $\alpha = 0.95$) in the limit of large density. The affinity of the non-self antigen should be greater than $\tau_{0.95}$, and the affinity of the self-antigen should be less than $\tau_{0.05}$ to guarantee that the mechanism makes very few errors. Therefore, the smaller the gap, the more specific is the response of the mechanism.

(c) *Fidelity*: This metric quantifies the distance of the activation set $\mathcal{A}_\alpha$ associated with a mechanism, and that associated with ideal absolute ligand discrimination, i.e. the distance of the curve from the two asymptotes. We provide one possible definition of this metric in Section 2.3.

(d) *Robustness*: This metric measures how the activation set changes when the parameters defining the mechanism are perturbed. Robustness is an important metric since noise becomes very important at low copy numbers.

Although versions of these metrics apply to any proofreading mechanism involved in switching, we trade-off between these desirable properties in the context of a particular model for absolute ligand discrimination.

We propose a new stochastic model where a T-cell activates when the total number of antigen-TCR-downstream complexes exceeds a certain threshold Z. We show how to efficiently compute the probability that the T-cell activates with a certain delay D, allowing us to investigate the impact of delay on each performance metric.

In our model absolute ligand discrimination emerges as a direct consequence of the fact that the number of TCRs on the cell surface is finite. This connection between absolute ligand discrimination and receptor number control (***Sorkin and Von Zastrow, 2009***) opens up the possibility that the properties of a single T-cell can be modulated by controlling the number of receptors.

T-cell activation in our model is controlled by the number of TCRs N, the threshold Z on the number of antigen-TCR-downstream molecule complexes, the delay D, and the probability α. We show how these parameters trade-off sensitivity, specificity, fidelity, and robustness. We also show how to use the model to explore the impact of receptor clustering, and show that such clustering improves fidelity and robustness to biochemical errors.

The rest of this paper is organized as follows. In Section 2.1 we introduce the model, and in Section 2.2 show that the model achieves absolute ligand discrimination because the number of re-

ceptors on the cell surface is finite. In Section 2.3 we discuss the results of numerical experiments investigating the trade-offs between sensitivity, specificity, fidelity, speed, and robustness. In Section 2.4 we discuss the impact of clustering on ligand discrimination, and in Section 3 we conclude with a discussion of possible experiments to test the different behaviors seen in the model.

2 Results

2.1 Model

Consider a single T-cell exposed to a single antigen with a uniform density L on the cell surface. The cell has N receptors spread uniformly across its surface that bind to the antigen. We work with the following approximation for the T-cell activation dynamics. An antigen-bound receptor can recruit one more molecule that we refer to as Lck [1], and the total number of antigen-TCR-Lck complexes needs to be larger than a threshold for the T-cell to activate. It will be clear from the discussion below, one could add more downstream molecules to the model without changing the analysis or the computation.

Following ***Husain*** (***2018***); ***Taylor et al.*** (***2017***), we model the response of a T-cell to an antigen by a two-dimensional stochastic model. The events in the model are as follows:

1. A single T-cell has a finite number N of TCRs available on its surface. The assumption that N is finite is a significant departure from ***Husain*** (***2018***); ***Taylor et al.*** (***2017***), and we show that this is directly leads to absolute ligand discrimination.
2. An unbound TCR binds to an antigen with rate $k_b L$. We implicitly assume that the TCR only reacts to a single antigen, and are ignoring the effects of ligand antagonism (***François et al., 2016***).
3. Antigen-bound TCRs recruit a single Lck molecule at rate k_{on}.
4. Antigens unbind from TCRs at rate $k_u = \tau_u^{-1}$. We assume that the antigen-TCR-Lck complex loses the Lck molecule if the antigen unbinds. In line with ***McKeithan*** (***1995***), we assume that Lck's do not unbind from antigen-TCR complexes unless the complex itself breaks apart. This assumption does not change the qualitative behavior of the model.
5. The cell activates once $Z(\geq 2)$ TCRs have recruited Lck molecules [2].

The state of a T-cell in our model is described by two variables: the number of antigen-TCR complexes $n \leq N$ and the number of antigen-TCR-Lck complexes $z \leq Z$. We let (n_t, z_t) denote the state of the T-cell at any time $t \geq 0$. Since $n_t, z_t \geq 0$, $z_t \leq n_t$, and the cell activates when $z_t = Z$, it follows that the state space is

$$(n_t, z_t) \in \mathcal{S} = \{(n, z) : 0 \leq z \leq Z, 0 \leq n \leq N, z \leq n\},$$

and the dynamics of $\{(n_t, z_t) : t \geq 0\}$ is described by a two-dimensional continuous time Markov chain with an absorbing boundary at

$$\mathcal{B} = \{(n, z) \in \mathcal{S} : z = Z, Z \leq n \leq N\}, \tag{1}$$

i.e. $\mathcal{B}$ refers to all cell states corresponding to an activated cell. See Figure 2 for a schematic of the Markov chain transitions and Appendix A for more details on the model. In our numerical simulations, we use $k_{on}^{-1} = 51s$ and $k_b^{-1} = 2250s/\mu m^2$ (***Husain, 2018***; ***Taylor et al., 2017***).

We assume that the T-cell must make a decision within a specified delay D from the first presentation of the antigen, i.e. if the cell does not have Z Lck's recruited by time D, the cell has missed its opportunity to sense the antigen in time. This is a departure from previous kinetic proofreading models, in which activation is a function of the steady-state concentrations and information theo-

[1] The details of the second molecule, i.e. whether it is phosphorus, Lck, ZAP, etc., are not relevant to our model.

[2] For $Z = 1$, the dynamics for $((1+\epsilon)L, \tau_u/(1+\epsilon))$ are the same as that for (L, τ_u). Therefore, absolute ligand discrimination is impossible.

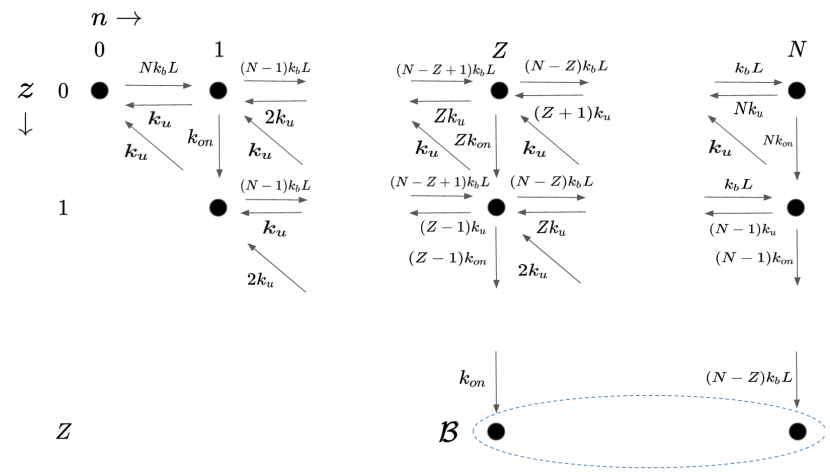

Figure 2. State space and transition rates for model of T-cell activation. Horizontal movements correspond to changes in the number of antigen-bound TCRs n, while vertical movements correspond to changes in the number of antigen-TCR-Lck complexes z. $\mathcal{B}$ corresponds to states for which the cell activates.

retic tools can be used to analyze the system (***Cover, 1999***; ***Ganti et al., 2020***; ***Das et al., 2018***). We want to characterize the set $\mathcal{A}_\alpha$ of density-affinity pairs (L, τ_u) for which the T-cell activates before D with probability $\alpha \in (0,1)$. In order to characterize the operating characteristics of the T-cell, we define the first passage time $T_{(i,j)}$ from the state $(i,j) \in \mathcal{S}\backslash\mathcal{B}$ to a state $(n, Z) \in \mathcal{B}$ as follows:

$$T_{(i,j)}(L, \tau_u) = \min\left\{t \geq 0 : (n_t, z_t) \in \mathcal{B} \text{ s.t. } (n_0, z_0) = (i,j)\right\} \tag{2}$$

i.e. $T_{(i,j)}$ is the first time at which a cell, initially with $i \geq 0$ antigen-bound TCRs and $0 \leq j \leq i$ antigen-TCR-Lck complexes, has Z antigen-TCR-Lck complexes. We emphasize that $T_{(i,j)}$ is a function of the density-affinity pair (L, τ_u) since the rate matrix Q (see Appendix for details) depends on (L, τ_u). Since the random time to activation after the initial presentation of the antigen is given by $T_{(0,0)}(L, \tau_u)$, it follows that

$$\mathcal{A}_\alpha = \left\{(L, \tau_u) : \mathbb{P}(T_{(0,0)}(L, \tau_u) \leq D) \geq \alpha\right\}. \tag{3}$$

In order to characterize $\mathcal{A}_\alpha$, we need to characterize the entire distribution of $T_{(0,0)}(L, \tau_u)$ as opposed to just the mean first passage time (MPFT), as was done in earlier works (***Husain, 2018***; ***Taylor et al., 2017***). We are interested in characterizing sets $\mathcal{A}_\alpha$ for a range of α in order to understand the impact of the probability threshold α.

Since an analytical expression for the distribution of the first passage time $T_{(0,0)}$ is unknown, we compute and invert the Laplace transform $\mathcal{L}_{(0,0)}(s) = \mathbb{E}[e^{-sT_{(0,0)}}]$ of $T_{(0,0)}$ to compute the entire cumulative distribution function (CDF) $\mathbb{P}(T_{(0,0)} \leq t)$ for all $t \geq 0$. We show that the complexity of computing the entire distribution is comparable to computing the MPFT. Furthermore, since we compute the entire CDF, we can compute the expected value of other performance measures. See Appendix A for details.

2.2 Finite number of receptor leads to absolute ligand discrimination

Let

$$\mathcal{L}_\alpha(\tau_u) = \inf\left\{L : (L, \tau_u) \in \mathcal{A}_\alpha\right\} = \inf\left\{L : \mathbb{P}(T_{(0,0)}(L, \tau_u) \leq D) \geq \alpha\right\}, \tag{4}$$

denote the lowest density for an antigen with affinity τ_u that ensures that the cell activates within a delay D. In Figure 3 we plot $\mathcal{L}_\alpha(\tau_u)$ vs τ for $\alpha \in \{0.05, 0.5, 0.95\}$. The set $\mathcal{A}_\alpha$ is above and to the right of the corresponding $\mathcal{L}_\alpha$-curve. We also plot the minimum density required for the mean

first passage time (MFPT) to be below the delay D. The plots in this figure also clearly show the additional information in the collection of plots $\mathcal{L}_\alpha$ for different values of α. For example, the plots show that the set $\mathcal{A}_\alpha$ is much more sensitive to affinity τ_u as compared to the density L. This information is not available in the single plot corresponding to the MFPT.

The $\mathcal{L}_\alpha$ curve is very closely approximated by two asymptotes: an asymptote in the L-direction at the affinity $\tau_{\min}$, and an asymptote in the τ_u-direction at the density $L_{\min}$. (See Figure 3 for the asymptotes corresponding to the $\mathcal{L}_{0.05}(\tau_u)$ curve.) The thresholds $(L_{\min}, \tau_{\min})$ are functions of the parameters of the mechanism (N, Z), the allowed delay D, and the probability threshold α. This asymptotic behavior directly leads to a probabilistic version of the absolute ligand discrimination described in ***Lalanne and François*** (***2013***). For antigens with sufficiently large affinity, the cell activates with very high probability by the cutoff time D, and it does so for all antigen density L above some threshold $L_{\min}$. For a fixed density above this lower threshold, one can decrease the affinity until a critical affinity $\tau_{\min}$ is reached, at which point the probability of activation starts to dramatically decrease as the affinity is further decreased. Finally, this critical affinity is the same for almost all densities above the density threshold.

We show that the asymptotes are an immediate consequence of the fact that N is finite, and do *not* exist if N is allowed to be infinite. First consider the asymptote in the L-direction. For large enough densities L, all of the N TCRs are very quickly bound to an antigen, and therefore the Lck dynamics, and hence the activation probability, are independent of density L. We rigorously prove this in section C in the appendix.

The dynamics that lead to the asymptote in the τ_u-direction are more subtle. Justifying the existence of a strictly positive density $L_{\min}$ below which the T-cell never activates for any value of the affinity τ_u is equivalent to showing that $L_{\min} = \lim_{\tau_u \to \infty} \mathcal{L}_\alpha(\tau_u) > 0$. Since increasing the affinity only makes it easier for the TCR-antigen-Lck complex to recruit and hold on to Lck, it is clear that $\mathcal{L}_\alpha(\tau_u)$ is a decreasing function of τ_u. So $L_{\min} \geq \mathcal{L}_\alpha(\infty)$. When $\tau_u = \infty$, antigens never never unbind, and so both n_t and z_t are non-decreasing over time. Since $z_t \leq n_t$, it follows that $\mathbb{P}(z_t \geq Z) \leq \mathbb{P}(n_t \geq Z)$ i.e. the likelihood of the cell being in an activated state by time t must be less than the likelihood of the cell having Z antigen-bound TCRs at time t. Furthermore, since the rate of recruiting one more antigen-bound TCR is bounded above by Nk_bL, it follows that the probability that there are Z antigen-bound TCRs by time D is less than $\mathbb{P}\left(\sum_{i=1}^{Z-1} X_i \leq D\right)$, where each X_i is independent, exponentially distributed with rate Nk_bL. Thus, it follows that $L_{\min} \geq \mathcal{L}_\alpha(\infty) \geq \inf\{L : \mathbb{P}(\sum_{i=1}^{Z-1} X_i \leq D) \geq \alpha\} > 0$ for sufficiently large D (see section B in the appendix). Note that the last term is strictly positive only if $N < \infty$ i.e. we lose the existence of the non-zero horizontal asymptote when the number of receptors is infinite. We rigorously establish this result in section D in the appendix.

We conclude by noting that absolute ligand discrimination holds even in the presence of additional downstream molecules. Suppose that each antigen-bound TCR must recruit $M > 1$ additional molecules, each at rate k_{on}, and that the cell activates once it has Z_M TCRs that have recruited M molecules. The $\mathcal{L}_\alpha$ curves for this model will still have a vertical asymptote i.e. some minimum affinity $\tau_{\min}$ below which no antigen will activate the cell in time with sufficient probability, even if the antigen has extremely high density. This occurs for the same reason mentioned before - for sufficiently high antigen density L, all N TCRs are bound to an antigen with very high probability for all points in time. Therefore, the dynamics of the proofreading network in this limit is effectively independent of the affinity τ_u. The $\mathcal{L}_\alpha$ curves will also have a horizontal asymptote i.e. a density $L_{\min}$ below which no antigen will activate the cell in time even with extremely high affinity, and this holds due to the exact same argument above that the time until activation is lower-bounded by the time to recruit Z_M antigens. These arguments, however, do not explain how properties of the kinetic proofreading network scale as the dimension of the model is increased. We leave this discussion for future work.

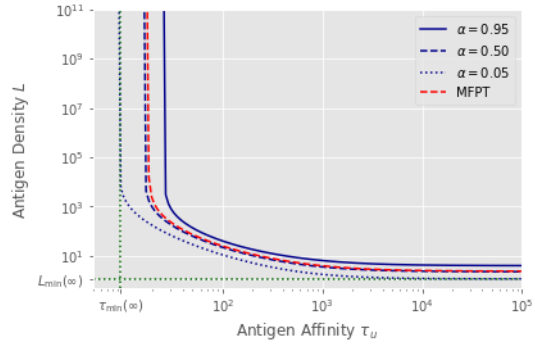

Figure 3. Let $\mathcal{L}_\alpha(\tau_u)$ denote the minimum density required to activate by delay D with probability α (see (4)). Plots of $\mathcal{L}_\alpha(\tau_u)$ for $Z = 8$ $N = 12$ and delay $D = 10^3$. For each affinity τ_u, the minimum density L was found using bisection over densities ranging from 10^{-4} to 10^{16}. The delay $D = 10^3$ seconds was chosen because it represents a realistic amount of time for human immune cells to operate over (10^3 seconds ≈ 15 minutes). Horizontal and vertical asymptotes $L_{\min}$ and $\tau_{\min}$ displayed for $\alpha = 0.05$. Also included is the minimum density L needed for a cell exposed to antigen with affinity τ_u to have a MFPT less than D. Time to compute $\alpha = 0.50$ curve was 30 seconds, time to compute MFPT curve was 2 seconds. This factor of 15 difference comes from using 15 summands to approximate the Bromwich integral. This same ratio holds when testing larger N and Z values (see section A in the appendix).

2.3 Sensitivity, specificity, fidelity, and robustness

In the Introduction, we introduced four metrics that we argued effectively characterize the performance of a kinetic proofreading mechanism: specificity, sensitivity, fidelity and robustness. In this section, we define these metrics for our proposed mechanism. In order to define these metrics, we need to define the following function:

$$\tau_\alpha(L) = \inf\{\tau_u : (L, \tau_u) \in \mathcal{A}_\alpha\} = \inf\{\tau_u : \mathbb{P}(T_{(0,0)}(L, \tau_u) \leq D) \geq \alpha\}, \tag{5}$$

i.e. $\tau_\alpha(L)$ is the minimum affinity at which the proofreading network activates with probability α by delay D when the antigen density is L. Note that $\tau_{\min} = \tau_\alpha(\infty)$ is the x-axis location of the asymptote in the L-direction. Next, we define the four metrics.

(a) *Specificity*: The level of the horizontal asymptote $\mathcal{L}_{0.95}(\infty)$ for $\alpha = 0.95$. Clearly, the smaller the sensitivity, the better the network is at responding to small amounts of antigen.

(b) *Specificity*: The relative gap $(\tau_{0.95}(\infty) - \tau_{0.05}(\infty))/\tau_{0.95}(\infty)$ between the x-axis location of the asymptote in the L-direction for $\alpha = 0.95$ and $\alpha = 0.05$. High specificity implies that the likelihood of activation decreases quickly as the antigen affinity is decreased. Therefore, networks with higher specificity are better at discriminating between two affinity values.

(c) *Fidelity*: The first two metrics refer to the asymptotes. For the ideal proofreading network, the $\mathcal{L}_\alpha$ curve is the L-shaped curve defined by the two asymptotes. Fidelity is a metric that measures how far $\mathcal{L}_\alpha$ is from the ideal L-shape. We need to define some intermediate quantities in order to define fidelity.

Fix $\epsilon > 0$. We say that the curve $\mathcal{L}_\alpha(\tau)$ reaches within a factor $(1+\epsilon)$ of the asymptote in the τ_u-direction at $\tau_\epsilon = \tau_\alpha((1+\epsilon)L_{\min})$, and reaches within a factor $(1+\epsilon)$ of the asymptote in the L-direction at $L_\epsilon = \mathcal{L}_\alpha((1+\epsilon)\tau_{\min})$. See Figure 4. Monotonicity properties of the model imply that the $\mathcal{L}_\alpha(\tau_u)$ curve lies in the following two L-shaped regions:

$$\mathcal{C}_\epsilon^{\tau_u} = \left\{(L, \tau_u) : L_{\min} \leq L \leq L_\epsilon \text{ or } \tau_{\min} \leq \tau_u \leq (1+\epsilon)\tau_{\min}\right\},$$
$$\mathcal{C}_\epsilon^{L} = \left\{(L, \tau_u) : L_{\min} \leq L \leq (1+\epsilon)L_{\min} \text{ or } \tau_{\min} \leq \tau_u \leq \tau_\epsilon\right\}.$$

The set $\mathcal{C}_\epsilon^{\tau_u}$ would represent an ideal proof-reading curve provided the maximum of the x-width $\epsilon\tau_{min}$ and y-width $(L_\epsilon - L_{\min})$ is close to zero, i.e. $d_\epsilon^{\tau_u} = \max\left\{\epsilon\tau_{\min}, L_\epsilon - L_{\min}\right\} \approx 0$. Similarly, define $d_\epsilon^L = \max\left\{\tau_\epsilon - \tau_{\min}, \epsilon L_{\min}\right\})$, and $d_\epsilon = \min\{d_\epsilon^{\tau_u}, d_\epsilon^L\}$ to be the minimum of the dis-

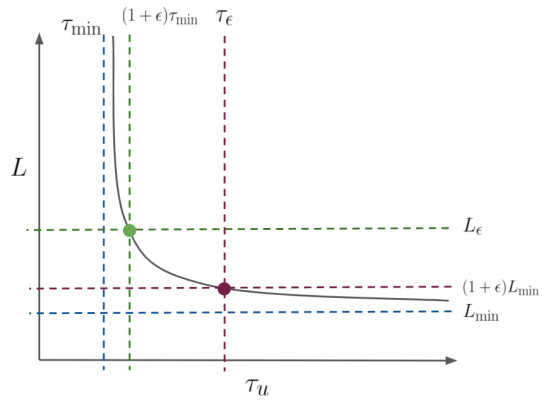

Figure 4. Examples of covering $(\tau_u, \mathcal{L}_\alpha(\tau_u))$ curve with $C_\epsilon^{\tau_u}$ (in green) and C_ϵ^{L} (in red).

tance to ideal performance implied by either of the two covering sets. We define the *fidelity* $F(N, Z, D, \alpha)$ as the inverse of the minimum "distance" over all possible choices for ϵ, i.e.

$$\begin{aligned} F(N, Z, D, \alpha) &= \frac{1}{\min_{\epsilon>0}\{d_\epsilon\}}, \\ &= \frac{1}{\min_{\epsilon>0}\left\{\min\left\{\max\left\{\epsilon\tau_{\min}, L_\epsilon - L_{\min}\right\}, \max\left\{\tau_\epsilon - \tau_{\min}, \epsilon L_{\min}\right\}\right\}\right\}} \end{aligned} \tag{6}$$

Note that high fidelity implies that the transition is closer to ideal behavior.

(d) *Robustness*: Since the actual number of TCRs $N + \delta N$ on the cell surface may differ from N because of transport noise, the corresponding $\mathcal{L}_\alpha$ curve may no longer be able to discriminate between the self and non-self antigens. We define the robustness $r_N(N, Z, D)$ of the kinetic proofreading network with parameters (N, Z, D) to errors in N as follows:

$$r_N(N, Z, D) = \frac{1}{N}\max\left\{\delta N : \tau_{\text{self}} \le \tau_{0.05}(N \pm \delta N, Z, T), \quad \tau_{\text{non-self}} \ge \tau_{0.95}(N \pm \delta N, Z, T)\right\}, \tag{7}$$

i.e. $r_N(N, Z, D, \alpha)$ is the maximum allowed relative change in N before the kinetic proofreading network is unable to differentiate between the self and non-self antigens.

The robustness in Z is defined as

$$r_Z(N, Z, D) = \frac{1}{Z}\max\left\{\delta Z : \tau_{\text{self}} \le \tau_{0.05}(N, Z \pm \delta Z, T), \quad \tau_{\text{non-self}} \ge \tau_{0.95}(N, Z \pm \delta Z, T)\right\} \tag{8}$$

It is possible that the robustness $r_N(N, Z, D, \alpha)$ is very sensitive to Z. In order to study this effect, we define robustness $r_{NZ}(N, Z, D, \alpha)$ to simultaneous perturbations in N and Z as follows:

$$r_{NZ}(N, Z, D) = \min_{|\delta Z| \le m}\left\{r_N(N, Z + \delta Z, D)\right\} \tag{9}$$

To summarize, a kinetic proofreading network that exhibits perfect absolute ligand discrimination has zero specificity, zero sensitivity, and infinite fidelity. Next, we investigate the possible values of sensitivity, specificity, and fidelity achievable by the proposed Markov chain model as a function of the parmeters (N, Z, D).

We simulated the performance of the kinetic proofreading system for

$$\begin{aligned} (\tau_{\text{self}}, \tau_{\text{non-self}}) &\in \{(1, 10), (3, 10)\} \text{ seconds}, \\ N &\in \{2, \ldots, 100\}, \\ Z &\in \{2, \ldots, N\}, \\ D &\in \left[1.1D^*(N, Z, \alpha = 0.95), 100D^*(N, Z, \alpha = 0.95)\right]. \end{aligned} \tag{10}$$

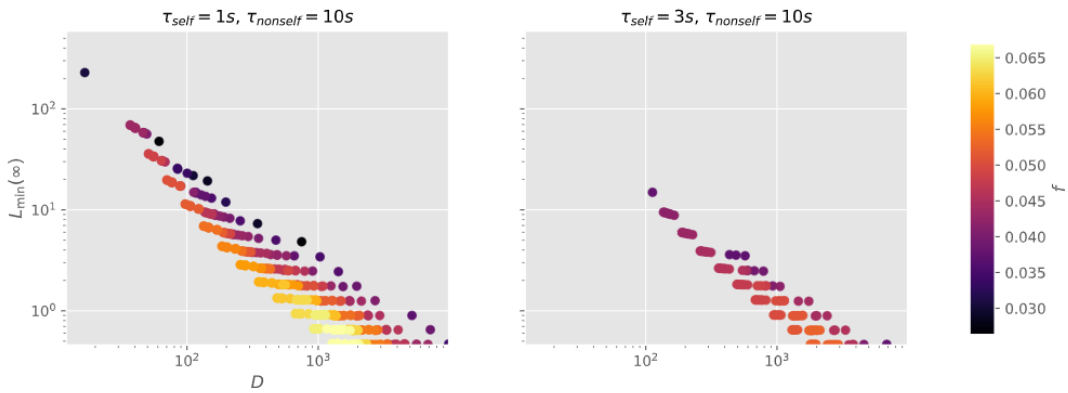

Figure 5. Achievable values of D and $L_{\min}$ for fixed τ_{self} and $\tau_{\text{non-self}}$. Points are colored based on the value of $F(L, D)$.

We call (N, Z, D) *admissible* for the pair of self and non-self antigen affinities $(\tau_{\text{self}}, \tau_{\text{non-self}})$ if the $\alpha = 0.95$ vertical asymptote $\tau_{0.95}(N, Z, D) \leq \tau_{\text{non-self}}$, and the $\alpha = 0.05$ vertical asymptote $\tau_{0.05}(N, Z, D) \geq \tau_{\text{self}}$. This filtering step fixes the specificity of the proofreading network.

In Figure 5 we plot sensitivity $L_{\min}(N, Z, D, \alpha = 0.95)$ vs the delay D for all admissible values of (N, Z, D) – each point corresponds to an admissible triple (N, Z, D). The lower envelop of the set of points corresponds to values of (N, Z, D) that are Pareto optimal with respect to sensitivity and delay, i.e. these points represent the optimal trade-off between these two metrics. Note that on the Pareto front, the sensitivity $L_{\min}$ is a decreasing function of the delay D, i.e. the proofreading network can be designed to be more sensitive if a higher delay is tolerable.

Next, we want to visualize the (L, D, F) Pareto front. There are often several close together (L, D) values that result in very different values for fidelity. In order to account for this, we take the maximum value of fidelity in a neighborhood of (L, D), and define

$$F(L, D) = \max\left\{F(N', Z', D') : ||(L_{\min}(N', Z', D'), D') - (L, D)||_2 / ||(L, D)||_2 \leq \delta = 0.05\right\}.$$

The color of the point (L, D) in figure 5 denotes the value of $F(L, D)$. On the Pareto front, the fidelity and sensitivity both increase as the allowed delay D increases, and begins to level off between $D = 100$ and $D = 1000$ seconds, a biologically realistic range of activation times. In the right-hand plot in Figure 5 we see that, as one would expect, increasing the required specificity by increasing τ_{self} limits the delay, sensitivity, and fidelity with which the cell can operate.

In Figure 6 the points are colored by robustness. In the first row the points are colored by robustness to perturbations in N, i.e. r_N defined in (7), in the second row the points are colored by robustness to perturbations in Z, i.e. r_Z defined in (8), and the third row the points are colored by the joint robustness r_{NZ} defined in (9). We see that the values of r_N and r_Z on the Pareto front increase as D increases. The trends in the N and Z-robustness r_{NZ} are slightly more complicated – on the Pareto front, r_{NZ} is maximized for intermediate values of D. Increasing specificity by only a small amount (e.g. a 30% increase in τ_{self}) can significantly reduce the achievable level of robustness in the kinetic proofreading model. Increasing specificity had the same effect on fidelity, although this change is not as drastic as that for robustness.

2.4 Clustering

Husain (***2018***); ***Taylor et al.*** (***2017***) proposed the model that we refine in this work to understand the role of receptor clustering in the activation of T-cells. Spatial heterogeneity and clustering of cell surface receptors is of general interest as well (***Burroughs and Van Der Merwe, 2007***; ***Iyengar and Rao, 2014***). ***Taylor et al.*** (***2017***) found that clustering of antigen-bound TCRs is crucial for ligand

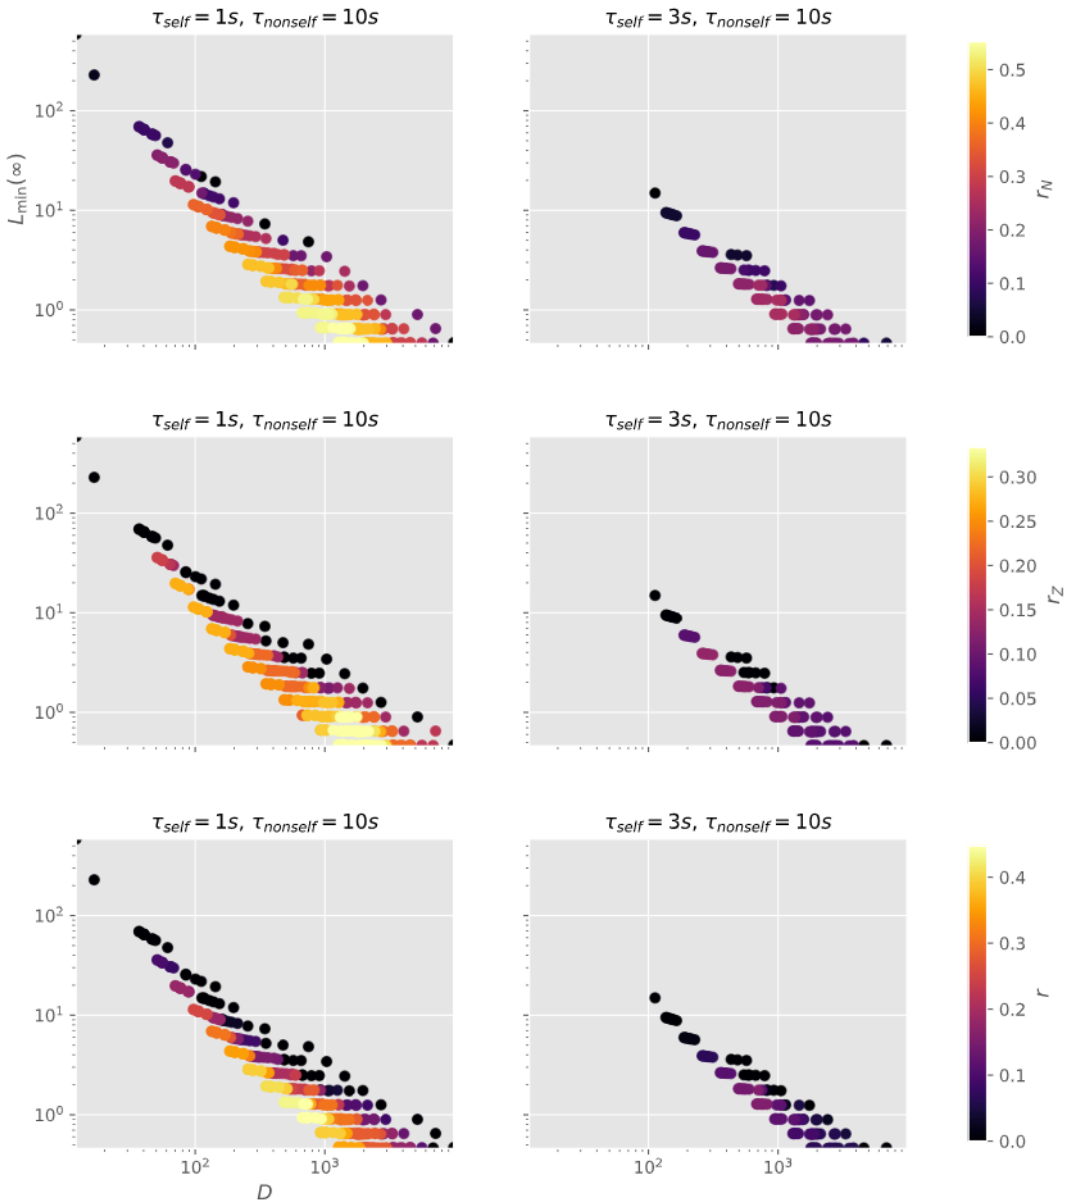

Figure 6. Achievable values of D and $L_{\min}$ for fixed τ_{self} and $\tau_{\text{non-self}}$. Points in the top two plots are colored based on r_N. Points in the middle two plots are colored based on r_Z. Points in the bottom two plots are colored based on r_{NZ}. In all three rows, highest robustness among nearby points is used to color the points.

discrimination in T-cells. We consider here what a cell potentially gains by incorporating clustering into its proofreading mechanism.

To approximate the impact of clustering we modify the original model as follows. We assume that the cell surface has two compartments, each initially having N TCRs. The dynamics of the two compartments are assumed to be independent. We let L_i denote the antigen density on compartment $i = 1, 2$, and say that the cell activates once one of the compartments has recruited Z Lck molecules. We introduce volatility in the antigen density by setting $L_1 = (1 - \epsilon)L$ and $L_2 = (1 + \epsilon)L$ for some $\epsilon \sim \text{uniform}(-1, 1)$, so that the average antigen density over the whole surface is still L. We assume that the cell is able to sense density and places $(1 + \delta)N$ receptors in the compartment where the density is higher, and $(1 - \delta)N$ receptors in compartment with the lower density.

Relaxing the assumptions of this simple model i.e. that the antigen density on each compartment is fixed over time, the cell surface only has two compartments, the cell can pre-emptively move δN receptors from compartment 1 to compartment 2, and that the number of receptors in each compartment stays fixed over time, is left for future analysis.

Figure 7 plots admissible (L_{min}, D) values at a specified level of specificity for models without clustering or variation in antigen density, without clustering but with variation in density, with clustering but without variation in density, and with both both clustering and variation in density. Points are colored based on fidelity. Figure 8 is the same as figure 7 but with points colored by robustness to simultaneous errors in N and Z.

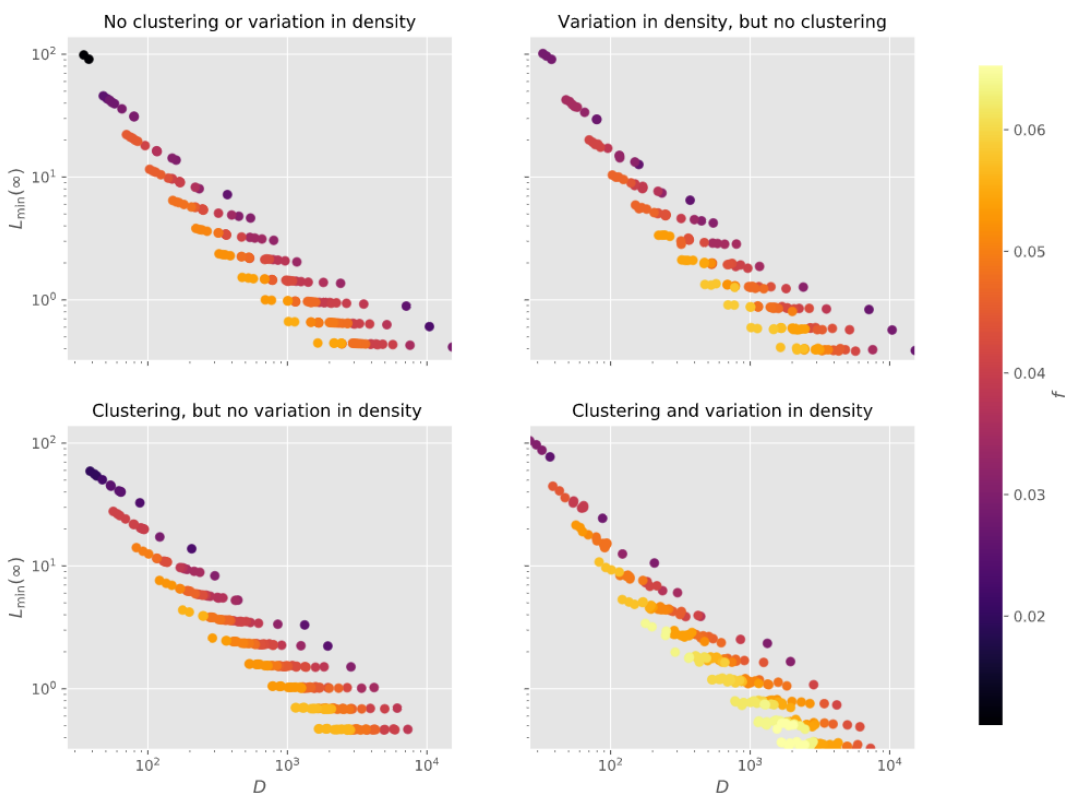

Figure 7. Achievable values of D and L_{min} for $\tau_{\text{self}} = 1$s and $\tau_{\text{nonself}} = 10$s. Points colored based on fidelity. Plots in top row have no clustering ($\delta = 0.0$), plots in bottom row have clustering with $\delta = 0.50$. Plots in left column have no variation in antigen density ($\epsilon = 0.0$), whereas plots in right column have variation in density with $\epsilon \sim \text{uniform}(-1, 1)$. For plots with variation in density, probability of activation was calculated with 20 samples of ϵ.

From figure 7 we see that clustering enhances the fidelity of the proofreading network given the existence of variation in the antigen density, with no loss in sensitivity or speed at the given level

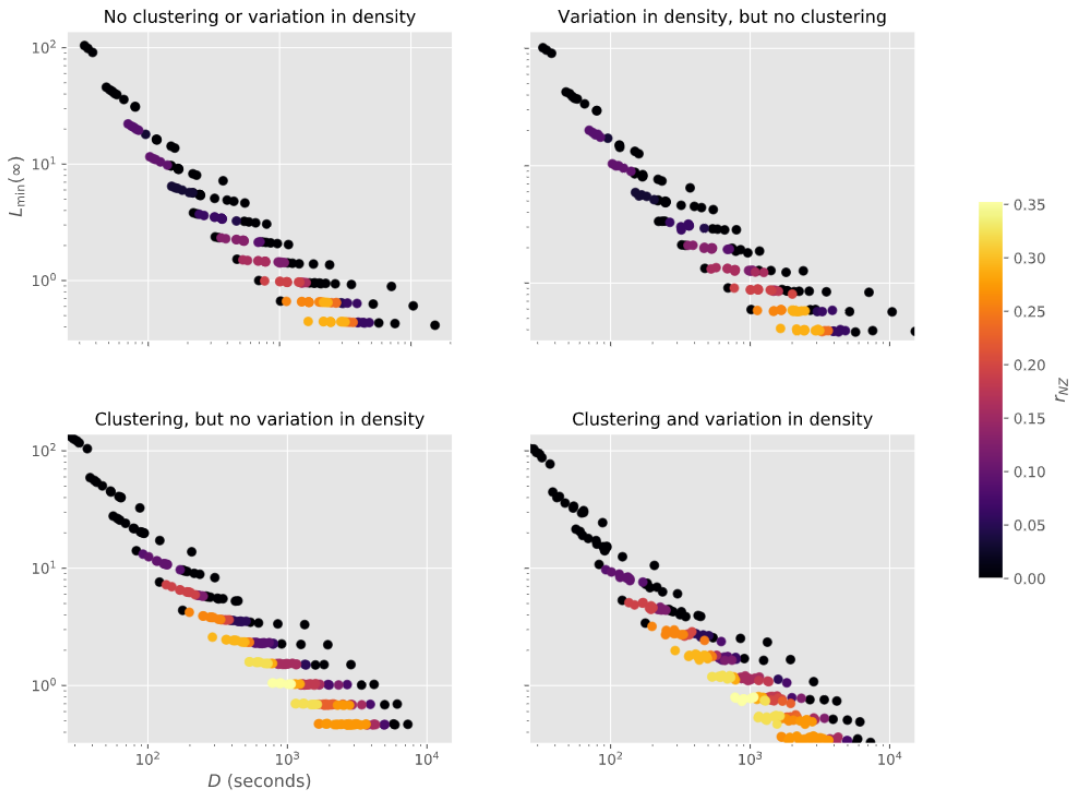

Figure 8. Same layout and parameters as figure 7, but with points colored based on robustness to simultaneous errors in N and Z.

of specificity. A similar result is seen in figure 8, with clustering improving the robustness to simultaneous errors in the number of surface receptors N and activation threshold Z, albeit regardless of whether or not there's variation in the antigen density. The impact of clustering on robustness to N and Z separately is similar, and figures for these quantities are omitted. These figures point towards clustering as a potential way of improving the fidelity and robustness of proofreading.

3 Discussion

We propose a kinetic proofreading model that exhibits a stochastic version of absolute ligand discrimination as a direct consequence of controlling the number of receptors. We propose that the performance of a kinetic proofreading mechanism can be characterized by sensitivity, specificity, fidelity, robustness to biochemical errors and delay. We explored the trade-off between these different metrics within our proposed model. We also examined the impact of receptor clustering on these trade-offs.

In Figure 9 we plot the predictions of our model as the number of receptors N is changed. Our model predicts a decrease in the level $L_{\min}$ of the horizontal asymptote as N increases; furthermore, the rate of decrease in $L_{\min}$ is decreasing in N, i.e. successive doubling of N does not result in equal amounts of decrease in $L_{\min}$. A similar prediction holds for $\tau_{\min}$ as well. Moreover, the model predicts that the transition from the horizontal asymptote to the vertical one happens over a larger interval of affinity values as N increases. These predictions can be experimentally validated by inhibiting or enhancing receptor endocytosis (***Sorkin and Von Zastrow, 2009***) and observing the impact on antigen sensing.

We also focused on the trade-offs between various metrics in the proposed kinetic proofreading model. Given the many different functions performed by cells and the various environments they

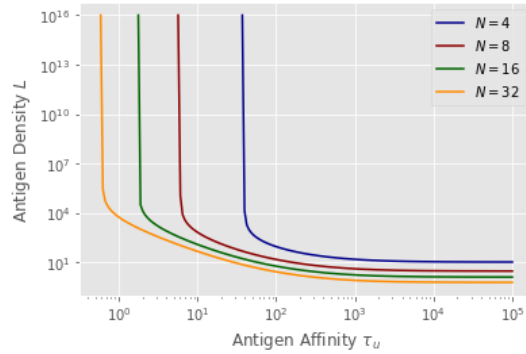

Figure 9. $\mathcal{L}_\alpha$ plots for various values of N. $Z = 4$, $\alpha = 0.95$, $D = 10^3$ seconds.

must operate in, one would expect there to be differences in how cells prioritize sensitivity, specificity, speed, and robustness. Controlling the number of surface receptors, the threshold Z on the downstream molecule, and the delay D, may be a way for cells to trade-off these metrics. Given different priorities among cells, it could be worth surveying different types of T-cells, as well as B-cells ***Tsourkas et al.*** (***2012***), and seeing whether the quantity of receptors displayed on their surface matches up with what they prioritize.

The model also showed that clustering may serve the purpose of increasing the fidelity and robustness of proofreading, particularly when the antigen density varies over the cell surface. It may be interesting to attempt an experiment where the extracellular fluid around a collection of T-cells is agitated so as to induce heterogeneities in antigen density. The model here would predict more clustering the more the fluid is perturbed.

As a final thought, our the model can be modified to study the effect of cell heterogeneity. Biological systems are subject to large amounts of noise that can both help and hurt the system. Developing this model further to study how cell heterogeneity can help a population of T-cells be more robust to changing environmental conditions ***Kussell and Leibler*** (***2005***) and antigen variability ***Mayer et al.*** (***2015***), and how the noisy responses from heterogeneous cells can be aggregated in an effective manner ***Butler et al.*** (***2013***), could be promising extensions.

References

Abate J, Choudhury GL, Whitt W. An introduction to numerical transform inversion and its application to probability models. In: *Computational probability* Springer; 2000.p. 257–323.

Altan-Bonnet G, Germain RN. Modeling T cell antigen discrimination based on feedback control of digital ERK responses. PLoS Biol. 2005; 3(11):e356.

Burroughs NJ, Van Der Merwe PA. Stochasticity and spatial heterogeneity in T-cell activation. Immunological reviews. 2007; 216(1):69–80.

Butler TC, Kardar M, Chakraborty AK. Quorum sensing allows T cells to discriminate between self and nonself. Proceedings of the National Academy of Sciences. 2013; 110(29):11833–11838.

Cover TM. Elements of information theory. John Wiley & Sons; 1999.

Das SG, Rao M, Iyengar G. Cascade of transitions in molecular information theory. Journal of Statistical Mechanics: Theory and Experiment. 2018; 2018(9):093402.

François P, Altan-Bonnet G. The case for absolute ligand discrimination: modeling information processing and decision by immune T cells. Journal of Statistical Physics. 2016; 162(5):1130–1152.

François P, Hemery M, Johnson KA, Saunders LN. Phenotypic spandrel: absolute discrimination and ligand antagonism. Physical Biology. 2016; 13(6):066011.

Ganti RS, Lo WL, McAffee DB, Groves JT, Weiss A, Chakraborty AK. How the T cell signaling network processes information to discriminate between self and agonist ligands. Proceedings of the National Academy of Sciences. 2020; 117(42):26020–26030.

Hopfield JJ. Kinetic proofreading: a new mechanism for reducing errors in biosynthetic processes requiring high specificity. Proceedings of the National Academy of Sciences. 1974; 71(10):4135–4139.

Husain K. The Mechanics of Cell Surface Clusters and Their Role in Signalling. PhD thesis, National Center for Biological Sciences; 2018.

Irvine DJ, Purbhoo MA, Krogsgaard M, Davis MM. Direct observation of ligand recognition by T cells. Nature. 2002; 419(6909):845–849.

Iyengar G, Rao M. A cellular solution to an information-processing problem. Proceedings of the National Academy of Sciences. 2014; 111(34):12402–12407.

Kussell E, Leibler S. Phenotypic diversity, population growth, and information in fluctuating environments. Science. 2005; 309(5743):2075–2078.

Lalanne JB, François P. Principles of adaptive sorting revealed by in silico evolution. Physical review letters. 2013; 110(21):218102.

Mayer A, Balasubramanian V, Mora T, Walczak AM. How a well-adapted immune system is organized. Proceedings of the National Academy of Sciences. 2015; 112(19):5950–5955.

McKeithan TW. Kinetic proofreading in T-cell receptor signal transduction. Proceedings of the national academy of sciences. 1995; 92(11):5042–5046.

Ninio J. Kinetic amplification of enzyme discrimination. Biochimie. 1975; 57(5):587–595.

Paster W, Bruger AM, Katsch K, Grégoire C, Roncagalli R, Fu G, Gascoigne NR, Nika K, Cohnen A, Feller SM, et al. A THEMIS: SHP 1 complex promotes T-cell survival. The EMBO journal. 2015; 34(3):393–409.

Perelson AS, Weisbuch G. Immunology for physicists. Reviews of modern physics. 1997; 69(4):1219.

Privault N. Understanding Markov Chains. Examples and Applications, Publisher Springer-Verlag Singapore. 2013; .

Sigman K. IEOR 6711: Continuous-Time Markov Chains; 2009, http://www.columbia.edu/~ks20/stochastic-I/stochastic-I-CTMC.pdf.

Sorkin A, Von Zastrow M. Endocytosis and signalling: intertwining molecular networks. Nature reviews Molecular cell biology. 2009; 10(9):609–622.

Taylor MJ, Husain K, Gartner ZJ, Mayor S, Vale RD. A DNA-based T cell receptor reveals a role for receptor clustering in ligand discrimination. Cell. 2017; 169(1):108–119.

Tsourkas PK, Liu W, Das SC, Pierce SK, Raychaudhuri S. Discrimination of membrane antigen affinity by B cells requires dominance of kinetic proofreading over serial engagement. Cellular & molecular immunology. 2012; 9(1):62–74.

A Model details

The full rate matrix Q for the model described in section 2.1 is defined as follows (***Sigman, 2009***).

$$
\begin{aligned}
Q_{(n,z),(n+1,z)} &= k_b L \mathbb{1}\,(0 \le n < N, 0 \le z < Z) \\
Q_{(n,z),(n-1,z)} &= (n-z) k_u \mathbb{1}\,(z < n \le N, 0 \le z < Z) \\
Q_{(n,z),(n,z+1)} &= (n-z) k_{on} \mathbb{1}\,(z < n \le N, 0 \le z < Z) \\
Q_{(n,z),(n-1,z-1)} &= z k_u \mathbb{1}\,(z \le n \le N, 0 < z < Z) \\
Q_{(n,z),(n,z)} &= -\left(Q_{(n,z),(n+1,z)} + Q_{(n,z),(n-1,z)} + Q_{(n,z),(n,z+1)} + Q_{(n,z),(n-1,z-1)}\right) \\
Q_{(n,z),(\eta,\zeta)} &= 0,\ (\eta,\zeta) \notin \{(n+1,z),(n-1,z),(n,z+1),(n-1,z-1),(n,z)\}
\end{aligned}
\tag{11}
$$

These rates follow from the assumptions described in the text. For example, if the T-cell is in state (n, z) and it has not yet activated, i.e. $z < Z$, the antigen binds to an available TCR, i.e. if $n < N$, with rate $k_b L$; thus, $Q_{(n,z),(n+1,z)} = k_b L$ if $n < N$ and $z < Z$. All the other rates similarly follow from the assumptions above. Note that $Q_{(n,z),(n,z)}$ is set to ensure that $\sum_{(n',z')\in S} Q_{(n,z),(n',z')} = 0$.

The first step in computing the distribution of $T_{(0,0)}$ is to compute the Laplace transform $\mathcal{L}_{(0,0)}(s)$ of $T_{(0,0)}$.

Lemma A.1 (Laplace Transform)**.** *Let $\mathcal{L}_{(i,j)} = \mathbb{E}\left[e^{-sT_{(i,j)}}\right]$ denote the Laplace transform of the density function for the first passage time $T_{(i,j)}$ from state (i, j) to the absorbing boundary $\mathcal{B}$. Then*

$$
\mathcal{L}_{(i,j)}(s) = \begin{cases} 1 & (i,j) \in B, \\ \frac{1}{s - Q_{(i,j),(i,j)}} \sum_{(k,l)\neq(i,j)} Q_{(i,j),(k,l)} \mathcal{L}_{(k,l)}(s) & (i,j) \notin B. \end{cases}
\tag{12}
$$

Therefore, $\mathcal{L}_{(0,0)}(s)$ can be computed by solving the above set of equations. The complexity of computing $\mathcal{L}_{(0,0)}$ is $\mathcal{O}(N^2 Z^2)$.

Proof. Since $\mathcal{B}$ is absorbing, it follows that $T_{(i,j)} = 0$ for all $(i.j) \in \mathcal{B}$. Therefore, $\mathcal{L}_{(i,j)}(s) = 1$ for all s.

Fix $(i, j) \notin \mathcal{B}$. Then the first-step analysis approach implies that ***Privault*** (***2013***)

$$
\begin{aligned}
\mathcal{L}_{(i,j)}(s) &= \sum_{(k,l)\neq(i,j)} P_{(i,j),(k,l)} \mathbb{E}\left[e^{-sT_{(i,j)}} \mid \text{jump from state } (i,j) \text{ to state } (k,l)\right] \\
&= \sum_{(k,l)\neq(i,j)} P_{(i,j),(k,l)} \mathbb{E}\left[e^{-s\left(H_{(i,j)}+T_{(k,l)}\right)}\right] = \sum_{(k,l)\neq(i,j)} P_{(i,j),(k,l)} \mathbb{E}\left[e^{-sH_{(i,j)}}\right]\left[e^{-sT_{(k,l)}}\right] \\
&= \frac{1}{s - Q_{(i,j),(i,j)}} \sum_{(k,l)\neq(i,j)} Q_{(i,j),(k,l)} \mathcal{L}_{(k,l)}(s)
\end{aligned}
$$

where $P_{(i,j),(k,l)} = \frac{Q_{(i,j),(k,l)}}{\sum_{(k,l)\neq(i,j)} Q_{(i,j),(k,l)}} = -\frac{Q_{(i,j),(k,l)}}{Q_{(i,j),(i,j)}}$ is the probability of transitioning to state $(k, l) \neq (i, j)$ upon leaving (i, j), and $H_{(i,j)}$ is the holding time for state (i, j) and is exponentially distributed with parameter $-Q_{(i,j),(i,j)}$. □

The cumulative density function (CDF) is given by

$$
\mathbb{P}\left(T_{(0,0)} \le t\right) = \mathcal{L}^{-1}\left\{\frac{\mathcal{L}_{(0,0)}(s)}{s}\right\}(t) = \lim_{\sigma\to\infty} \frac{1}{2\pi i} \int_{b-i\sigma}^{b+i\sigma} e^{st} \frac{\mathcal{L}_{(0,0)}(s)}{s} ds,
\tag{13}
$$

where b is any real number to the right of all singularities of $\mathcal{L}_{(0,0)}(s)/s$. Let $s = c + id$, for $c, d \in \mathbb{R}$. Then

$$
\left|\frac{\mathcal{L}_{(0,0)}(s)}{s}\right| = \frac{1}{\sqrt{c^2+d^2}} \left|\mathbb{E}\left[e^{-cT_{(0,0)}} e^{-idT_{(0,0)}}\right]\right| \le \frac{1}{\sqrt{c^2+d^2}} \mathbb{E}\left[e^{-cT_{(0,0)}}\right].
$$

Thus, it follows that any singularity of $\mathcal{L}_{(0,0)}(s)/s$ has $\Re(s) \le 0$, and therefore any $b > 0$ is sufficient. We use $b = 1$ in our numerical study. Next, using the trapezoidal rule with step size h, we approximate

the Bromwich intergral in (13) with an infinite series ***Abate et al.*** (***2000***)

$$\lim_{\sigma\to\infty}\frac{1}{2\pi i}\int_{b-i\sigma}^{b+i\sigma} e^{st}\frac{\mathcal{L}_{(0,0)}(s)}{s}ds \approx \frac{he^{bt}}{\pi}+\frac{2he^{bt}}{\pi}\sum_{k=1}^{+\infty}\mathrm{Re}(\mathcal{L}_{(0,0)}(b+ikh))\cos(kht),$$

and appropriately truncate the series to approximate the integral. For h and the number of terms in the approximation, we use the specification in ***Abate et al.*** (***2000***). We found that the number of terms was independent of the parameters (N, Z, D), i.e. only a constant number of Laplace transform evaluations are needed for any parameter values. Thus, it follows that the computational complexity of computing the CDF in terms of (N, Z) is the same as that of computing the mean first passage time (MFPT), since computing the MFPT reduces to solving a system of equations very similar to those used to compute $\mathcal{L}_{(0,0)}$.

It is also clear that this method of computing the first passage time until activation will still work even if the number of downstream molecules that can bind to each TCR is increased to $M > 1$, albeit the system of equations to solve for the Laplace transforms $\mathcal{L}_{(i,j_1,\ldots,j_M)}(s)$ of the first passage times $T_{(i,j_1,\ldots,j_M)}$ until activation (which in this case corresponds to the cell reaching some critical threshold Z_M antigen-bound TCRs with M additional molecules recruited) will take $\mathcal{O}(N^2 Z_1^2 \ldots Z_M^2)$ time.

B Minimum possible activation time

Throughout the text we implicitly assume that the delay D is chosen to be large enough so that it's possible for the cell to activate quickly enough with sufficiently high probability. In this section, we compute a lower bound for the delay D such that $\mathbb{P}(T_{(0,0)}(L,\tau_u) \le D) \ge \alpha$ for some (L, τ_u).

Since the forward transitions, i.e. those that increase (n, z), are increasing in L and the backward transitions, i.e. those that decrease (n, z), are decreasing in τ_u, it follows that the first passage time $T_{(0,0)}$ is monotonically decreasing in both L and τ_u. Thus, in order to ensure there exist (L, τ_u) for which $\mathbb{P}(T_{(0,0)}(L,\tau_u) \le D) \ge \alpha$, it is sufficient to show that $\lim_{\tau_u\to\infty}\lim_{L\to\infty}\mathbb{P}(T_{(0,0)}(L,\tau_u) \le D) \ge \alpha$. In Section C, we establish that the first passage time $T_{(0,0)}(\tau_u, L)$ in the limit $L \to \infty$ is approximated by the first passage time $T_0^{(B)}$ from state 0 to state Z in a birth-death process with birth rates $\lambda_z = (N-z)k_{on}$ and death rate $\mu_z = z/\tau_u$. Next, taking the limit $\tau_u \to \infty$, the limit

$$\begin{aligned}\lim_{\tau_u\to\infty} T_0^{(B)} &= \mathrm{Exp}(Nk_{on}) + \mathrm{Exp}((N-1)k_{on}) + \cdots + \mathrm{Exp}((N-Z+1)k_{on}),\\ &= \mathrm{Hypo}\left(Nk_{on}, \ldots, (N-Z+1)k_{on}\right),\end{aligned}$$

where $\mathrm{Exp}(\lambda)$ denotes an exponential random variable with rate λ and $\mathrm{Hypo}(\lambda_1, \ldots, \lambda_k)$ is a hypo-exponential distribution defined by the given rates. The same limit results from taking the limit $L \to \infty$ in the Markov chain $M^{(f)}$ that we discuss in Section D. Thus, in order to ensure that the network reaches $\mathcal{B}$ with probability at least α, we need

$$D \ge D^*(N, Z, \alpha) \triangleq \inf\left\{T : \mathbb{P}\left(\mathrm{Hypo}\left(Nk_{on}, \ldots, (N-Z+1)k_{on}\right) \le T\right) \ge \alpha\right\}. \tag{14}$$

We ensure that $D \ge D^*(N, Z, \alpha)$ in all of our analyses so that $\tau_{\min}(\infty)$ and $L_{\min}(\infty)$ are well-defined.

C Vertical asymptote

Fix $D \ge D^*$ (see (14)), and define

$$\tau_{\min}(L; N, Z, D, \alpha) \triangleq \inf\left\{\tau_u > 0 : \mathbb{P}\left(T_{(0,0)}(\tau_u, L; N, Z) \le D\right) \ge \alpha\right\}. \tag{15}$$

where $T_{(0,0)}(\tau_u, L; N, Z)$ is the first passage time from state $(0, 0)$ to the absorbing boundary $\mathcal{B}$ for an antigen with affinity τ_u and density L when the cell has N TCRs and activates when Z antigen-TCR-Lck complexes form. Informally, $\tau_{\min}(L; N, Z, D, \alpha)$ denotes the minimum affinity τ_u needed to

ensure that the cell activates with probability at least α by time D with antigen density L, when the cell has N TCRs on the cell surface, and Z antigen-TCR-Lck complexes are needed for activation. For brevity, let $\tau_{\min}(L)$ denote $\tau_{\min}(L; N, Z, D, \alpha)$.

To study the behavior of $\tau_{\min}(L)$ as $L \to \infty$, we define

$$\tau_{\min}(\infty) \triangleq \tau_{\min}\left(\lim_{L\to\infty} L\right) = \inf\left\{\tau_u > 0 : \lim_{L\to\infty} \mathbb{P}\left(T_{(0,0)}(\tau_u, L) \le D\right) \ge \alpha\right\} \tag{16}$$

where the last equality holds by Fatou's lemma and the fact that $\lim_{L\to\infty}\{T_{(0,0)}(\tau_u, L) \le D\}$ exists and is equal to $\{T_{(0,0)}(\tau_u, \infty) \le D\}$, so we can exchange the limit and probability.

Note that as $L \to \infty$, all the N TCRs instantly convert into TCR-antigen complexes, and the bottleneck event that determines activation is the recruitment of Lck, i.e. activation is controlled by the Lck dynamics. We formalize this in Theorem C.1. Let $B = \{b_t : t \ge 0\}$ denote a birth-death process defined on $\{z : 0 \le z \le Z\}$ with birth rates $\lambda_z = (N - z)k_{on}$ and death rates $\mu_z = z/\tau_u$, and an absorbing boundary at Z. Let $T_0^{(B)}(\tau_u; N, Z)$ denote the first passage time from state 0 to state Z in a birth-death process B. Then, we have the following characterization for $\tau_{\min}(\infty)$.

Theorem C.1 (Vertical Asymptote)**.** *The minimum affinity $\tau_{\min}(\infty)$ required to activate in the limit $L \to \infty$ is given by*

$$\tau_{\min}(\infty) = \tau_{\min}^{(B)}(N, Z, D, \alpha) \triangleq \inf\left\{\tau_u > 0 : \mathbb{P}\left(T_0^{(B)}(\tau_u; N, Z) \le D\right) \ge \alpha\right\}.$$

Thus, Theorem C.1 implies that, in order to characterize $\tau_{\min}(\infty)$ as a function of (N, Z, T, α), we only need to compute the distribution of the first passage time $T_0^{(B)}(\tau_u)$ for the birth-death process B. The Laplace transforms $\mathcal{L}_i(s)$ for the first passage time from state i to state Z for B can be computed in the same way as the Laplace transforms $\mathcal{L}_{(i,j)}(s)$ for the original model. Computing $\mathcal{L}_i(s)$ for $0 \le i \le Z$ reduces to solving a tri-diagonal system of equations,with a computational complexity of $\mathcal{O}(Z)$ as opposed to $\mathcal{O}(N^2Z^2)$. Figure D.1 confirms that $\tau_{\min}(\infty) = \tau_{\min}^{(B)}$ and that $\tau_{\min}^{(B)}$ can be computed much faster than $\tau_{\min}(\infty)$ (~10-100 times faster for moderate-sized N and Z).

Proof. This result follows from showing that $\lim_{L\to\infty} \mathbb{P}\left(T_{(0,0)}(\tau_u, L) \le D\right) = \mathbb{P}\left(T_0^{(B)}(\tau_u) \le D\right)$. Let (n_t, z_t) be the state of the Markov chain at time t. Since the boundary $\mathcal{B}$ is absorbing,

$$\mathbb{P}\left(T_{(0,0)}(\tau_u, L) \le D\right) = \mathbb{P}_{(0,0)}\left((n_D, z_D) \in \mathcal{B}\right) = \mathbb{P}_{(0,0)}\left(z_D = Z\right)$$

with $\mathbb{P}_{(n,z)}$ denoting the distribution of (n_t, z_t) given $(n_0, z_0) = (n, z)$. Therefore, it suffices to show $\lim_{L\to\infty} \mathbb{P}_{(0,0)}\left((n_T, z_T) \in \mathcal{B}\right) = \mathbb{P}\left(T_0^{(B)}(\tau_u) \le D\right)$. We establish this by showing

$$\lim_{L\to\infty} \mathbb{P}_{(0,0)}\left(z_t = z\right) = \lim_{L\to\infty} \mathbb{P}_{(N,0)}\left(z_t = z\right) = \mathbb{P}_0\left(B_t = z\right)$$

The first equality is established in Lemma C.2, and the second equality is established in Lemma C.3. □

The intuition here is that in the limit of $L \to \infty$, all N TCRs are instantaneously bound to antigens, and the activation is determined by the dynamics of the antigen bound TCRs acquiring Lck molecules. In Lemma C.2 we rigorously establish the first result.

Lemma C.2. *For all $t > 0$, $\lim_{L\to\infty} \mathbb{P}_{(0,0)}(z_t = z) = \lim_{L\to\infty} \mathbb{P}_{(N,0)}(z_t = z)$.*

Proof. The rate out of each state (n, z) in the chain is upper bounded by $k_b L + \frac{N}{\tau_u} + \frac{Z}{\tau_u} + Nk_{on}$. We

define the following four independent Poisson processes.

$$P_t^{(E)} \sim Poisson(k_b L t), \qquad P_t^{(N)} \sim Poisson(N k_{on} t)$$
$$P_t^{(W)} \sim Poisson(\frac{N}{\tau_u} t), \qquad P_t^{(SW)} \sim Poisson(\frac{Z}{\tau_u} t)$$

We think of the arrivals for the process $P_t^{(E)}$ as "eastward" movements from (n, z) to $(n, z+1)$, arrivals for $P_t^{(W)}$ as "westward" movements from (n, z) to $(n-1, z)$, arrivals for $P_t^{(N)}$ as "northward" movements from (n, z) to $(n, z+1)$, and arrivals for $P_t^{(SW)}$ as "southwestward" movements from (n, z) to $(n-1, z-1)$. We can see that

$$\begin{aligned}
\mathbb{P}_{(0,0)}\left(\substack{n_{K(\epsilon)/L}=N, z_{K(\epsilon)/L}=0, \\ n_t=n, z_t=z}\right) &= \mathbb{P}_{(0,0)}\left(n_{K(\epsilon)/L} = N, z_{K(\epsilon)/L} = 0\right) \mathbb{P}_{(N,0)}\left(n_{t-K(\epsilon)/L} = n, z_{t-K(\epsilon)/L} = m\right) \\
&\leq \mathbb{P}_{(0,0)}\left(n_t = n, z_t = z\right) \\
&= \mathbb{P}_{(0,0)}\left(\substack{n_{K(\epsilon)/L}=N, z_{K(\epsilon/L)}=0, \\ n_t=n, z_t=z}\right) + \mathbb{P}_{(0,0)}\left(\substack{(n_{K(\epsilon)/L}, z_{K(\epsilon/L)})\neq(N,0), \\ n_t=n, z_t=z}\right) \\
&\leq \mathbb{P}_{(0,0)}\left(\substack{n_{K/L}=N, z_{K/L}=0, \\ n_t=n, z_t=z}\right) + \mathbb{P}_{(0,0)}\left(\left(n_{K/L}, z_{K/L}\right) \neq (N, 0)\right)
\end{aligned}$$

Next, we have

$$\begin{aligned}
\mathbb{P}_{(0,0)}\left(n_{K/L} = N, z_{K/L} = 0\right) &\geq \mathbb{P}\left(P_{K/L}^{(E)} \geq N, P_{K/L}^{(W)} = P_{K/L}^{(N)} = P_{K/L}^{(SW)} = 0\right) \\
&= \mathbb{P}(P_{K/L}^{(E)} \geq N) e^{-\frac{K}{L}\left(\frac{N}{\tau_u} + \frac{Z}{\tau_u} + N k_{on}\right)}
\end{aligned}$$

Since $e^{-x}/x^j \to 0$ as $x \to \infty$ for any finite j, it follows that $\mathbb{P}\left(P_{K/L}^{(E)} \geq N\right) = 1 - e^{-K k_b} \sum_{j=0}^{N-1} \frac{(K k_b)^j}{j!} \to 1$ as $K \to \infty$, and $e^{-\frac{K}{L}\left(\frac{N}{\tau_u} + \frac{Z}{\tau_u} + N k_z\right)} \to 1$ as $L \to \infty$ for any fixed K. Thus, it follows that for all $\epsilon > 0$, there exists some $K = K(\epsilon)$ such that for all $L \geq L(\epsilon)$, $\mathbb{P}\left(P_{K/L}^{(E)} \geq N\right) e^{-\frac{K}{L}\left(\frac{N}{\tau_u} + \frac{Z}{\tau_u} + N k_z\right)} \geq 1 - \epsilon$. Hence,

$$\begin{aligned}
(1-\epsilon)\mathbb{P}_{(N,0)}\left(n_{t-K(\epsilon)/L} = n, z_{t-K(\epsilon)/L} = z\right) &\leq \mathbb{P}_{(0,0)}\left(n_{K(\epsilon)/L} = N, z_{K(\epsilon)/L} = 0\right) \mathbb{P}_{(N,0)}\left(\substack{n_t=z, \\ z_t=z} \,\middle|\, \substack{n_{K(\epsilon)/L}=N, \\ z_{K(\epsilon)/L}=0}\right) \\
&= \mathbb{P}_{(0,0)}\left(n_t = n, z_t = z\right) \\
&= \mathbb{P}_{(0,0)}\left(\substack{n_t=n, z_t=z, \\ n_{K(\epsilon)/L}=N, z_{K(\epsilon)/L}=0}\right) + \mathbb{P}_{(0,0)}\left(\substack{n_t=n, z_t=z, \\ (n_{K(\epsilon)/L}, z_{K(\epsilon)/L})\neq(N,0)}\right) \\
&\leq \mathbb{P}_{(0,0)}\left(\substack{n_t=n, z_t=z, \\ n_{K(\epsilon)/L}=N, z_{K(\epsilon)/L}=0}\right)) + \mathbb{P}_{(0,0)}((n_{K(\epsilon)/L}, z_{K(\epsilon)/L}) \neq (n, 0)) \\
&\leq \mathbb{P}_{(N,0)}\left(n_{t-K(\epsilon)/L} = n, z_{t-K(\epsilon)/L} = z\right) + \epsilon
\end{aligned}$$

Now, choose any sequence of densities $\{L_k\}$ with $L_k \to \infty$, as well as a sequence $\{\epsilon_k\}$ such that $\epsilon_k \to 0$ and $K(\epsilon_k)/L_k \to 0$. Taking the limit along (L_k, ϵ_k) gives

$$\lim_{k\to\infty} \mathbb{P}_{(N,0)}\left(n_{t-K(\epsilon_k)/L_k} = n, z_{t-K(\epsilon)/L_k} = z\right) = \lim_{k\to\infty} \mathbb{P}_{(0,0)}\left(n_t = n, z_t = z\right)$$

In the limit $k \to \infty$, $K(\epsilon_k)/L_k \to 0$, and the state distribution of a continuous time Markov chain is continuous in time, so

$$\lim_{k\to\infty} \mathbb{P}_{(N,0)}\left(n_{t-K(\epsilon_k)/L_k} = n, z_{t-K(\epsilon_k)/L_k} = z\right) = \lim_{k\to\infty} \mathbb{P}_{(N,0)}\left(n_t = n, z_t = z\right)$$

and we can conclude that $\lim_{L\to\infty} \mathbb{P}_{(0,0)}\left(n_t = n, z_t = z\right) = \lim_{L\to\infty} \mathbb{P}_{(N,0)}\left(n_t = n, z_t = z\right)$. □

Next, we establish that dynamics of the 2D Markov chain in the limit of $L \to \infty$ is given by the 1D birth-death process in the Lck dimension.

Lemma C.3. $\lim_{L\to\infty} \mathbb{P}_{(N,0)}(z_t = z) = \mathbb{P}_0(B_t = z)$.

Proof. Instead of tracking the aggregate behavior (n_t, z_t) of the chain, we keep track of individual

arrivals and departures of TCRs and Lcks, and show $\lim_{L\to\infty} \mathbb{P}_{(N,0)}(z_t \le z) = \mathbb{P}_0(B_t \le z)$ using the following coupling

1. Let (a_t, d_t) denote the birth and death processes associated with the birth-death process B_t i.e. $B_t = a_t - d_t$.
2. Suppose there is an arrival at time s. Then, increase z_s by 1 with probability n_s/N and keep it the same with probability $1 - n_s/N$.
3. Suppose there is a departure at time s, decrease z_s by 1 only if the associated arrival had led to an increase in z_s.

By construction, z_t defined above has the same distribution as the number of Lcks z_t in the Markov chain. Also by construction, $z_t \le B_t$ on a sample-path basis, and therefore $\mathbb{P}_{(N,0)}(z_t \le z) \ge \mathbb{P}_0(B_t \le z)$, i.e. $\lim_{L\to\infty} \mathbb{P}_{(N,0)}(z_t \le z) \ge \mathbb{P}_0(B_t \le z)$.

Additionally, $B_t - z_t \le A(T_N(t))$ on a sample-path basis, with $T_N(t)$ being the total amount of time the Markov chain spends in states with $n_s < N$ for $s \in (0,t]$, and $A(T_N(t))$ being the number of arrivals in B_t during the periods of time in $(0,t]$ when $n_s < N$. Thus, we have that

$$\begin{aligned}
\mathbb{P}_{(N,0)}\left(z_t \le z\right) &\le \mathbb{P}_0\left(B_t - A\left(T_N(t)\right) \le z\right) \\
&\le \mathbb{P}_0\left(B_t \le A\left(T_N(t)\right) + z \middle| T_N(t) \le u\right) \mathbb{P}_{(N,0)}\left(T_N(t) \le u\right) + \mathbb{P}_{(N,0)}\left(T_N(t) \ge u\right) \quad \forall u > 0 \\
&\le \mathbb{P}_0\left(B_t \le A(u) + z\right) \mathbb{P}_{(N,0)}\left(T_N(t) \le u\right) + \mathbb{P}_{(N,0)}\left(T_N(t) > u\right) \\
&\le \mathbb{P}_0\left(B_t \le A(u) + z\right) + \mathbb{P}_{(N,0)}\left(T_N(t) > u\right)
\end{aligned}$$

where the second-to-last inequality uses the fact that $T_N(t)$ is independent of B_t. Taking limits, we see

$$\begin{aligned}
\lim_{L\to\infty} \mathbb{P}_{(N,0)}\left(z_t \le z\right) &\le \mathbb{P}_0\left(B_t \le A(u) + z\right) + \lim_{L\to\infty} \mathbb{P}_{(N,0)}\left(T_N(t) > u\right), && \forall u > 0 \\
&\le \mathbb{P}_0\left(B_t \le A(u) + z\right), && \forall u > 0 \\
&\le \inf_{u>0} \mathbb{P}_0(B_t \le A(u) + z) \\
&= \mathbb{P}_0\left(B_t \le z\right)
\end{aligned}$$

The first line uses the fact that B_t and $A(u)$ are independent of L. The second inequality comes from the fact $\limsup_{L\to\infty} \mathbb{P}_{(N,\cdot)}\left(T_N(t) > u\right) = 0$ for all $u > 0$, with $\mathbb{P}_{(n,\cdot)}$ being the conditional probability when the chain starts with n TCRs and arbitrary $0 \le z \le n$ at time 0. This result is established in Lemma C.4. The third inequality comes from the fact that, in general, the CDF of a random variable is monotonically increasing in its argument, and $A(u)$ is monotonically decreasing in a sample-path sense, so $\mathbb{P}_0\left(B_t \le A(u) + z\right)$ is minimized by taking $u \to 0$. The last equality uses the fact that CDFs by definition are right-continuous, and $\lim_{u\to\infty} A(u) = 0$ in a sample-path sense. We conclude that $\lim_{L\to\infty} \mathbb{P}_{(N,0)}\left(z_t \le z\right) = \mathbb{P}_0\left(B_t \le z\right)$. □

Lemma C.4. $\limsup_{L\to\infty} \mathbb{P}_{(N,\cdot)}(T_N(t) \ge u) = 0$.

Proof. By Markov's inequality, we know

$$\limsup_{L\to\infty} \mathbb{P}_{(N,\cdot)}\left(T_N(t) \ge u\right) \le \limsup_{L\to\infty} \frac{\mathbb{E}_{(N,\cdot)}\left[T_N(t)\right]}{u}.$$

Next, let $D_N(t)$ be the number of transitions from a state with N TCRs to a state with $N-1$ TCRs in $(0,t]$ and S be the sojourn time from a state with $N-1$ TCRs to a state with N TCRs. By definition, $T_N(t) = \sum_{k=1}^{D_N(t)} \min\{S_k, t\}$ with the S_k's being iid copies of S, so by Wald's equation we have

$\mathbb{E}_{(N,\cdot)}\left[T_N(t)\right] = \mathbb{E}_{(N,\cdot)}\left[D_N(t)\right]\mathbb{E}_{(N-1,\cdot)}\left[\min\{S,t\}\right]$. We see that

$$\begin{aligned}
\mathbb{E}_{(N-1,\cdot)}\left[\min\{S,t\}\right] &= \mathbb{E}_{(N-1,\cdot)}\left[\min\{S,t\}\middle|\text{first transition to } (N,z)\right]\mathbb{P}_{(N-1,\cdot)}\left(\text{first transition to } (N,z)\right) \\
&\quad + \mathbb{E}_{(N-1,\cdot)}\left[\min\{S,t\}\middle|\text{first transition not to } (N,z)\right]\mathbb{P}_{(N-1,\cdot)}\left(\text{first transition not to } (N,z)\right) \\
&\le \mathbb{E}_{(N-1,\cdot)}\left[\min\{S,t\}\middle|\text{first transition to } (N,z)\right] + t\,\mathbb{P}_{(N-1,\cdot)}\left(\text{first transition not to } (N,z)\right) \\
&\le \mathbb{E}_{(N-1,\cdot)}\left[S\middle|S\le t, \text{first transition to } (N,z)\right] + t\frac{Nk_{on}+(N+Z)/\tau_u}{k_bL+Nk_{on}+(N+Z)/\tau_u} \\
&\le \frac{1}{k_bL} + t\frac{Nk_{on}+(N+Z)/\tau_u}{k_bL+Nk_{on}+(N+Z)/\tau_u}
\end{aligned}$$

Since the rates for the Poisson processes $P_t^{(W)}$, $P_t^{(SW)}$, and $P_t^{(N)}$ are all finite for all $\tau_u > 0$ and independent of L, $\mathbb{E}_{(N,\cdot)}\left[D_N(t)\right]$ is finite for all L. Therefore, we conclude that $\limsup_{L\to\infty}\mathbb{E}_{(N,\cdot)}\left[T_N(t)\right] = 0$, and $\limsup_{L\to\infty}\mathbb{P}_{(N,\cdot)}\left(T_N(t)\ge u\right) = 0$. □

D Horizontal asymptote

Fix $D \ge D^*$ (see (14)), and let

$$L_{\min}(\tau_u; N, Z, D, \alpha) = \inf\left\{L > 0 : \mathbb{P}\left(T_{(0,0)}(\tau_u, L; N, Z) \le D\right) \ge \alpha\right\}$$

denote the minimum density L needed to ensure that the T-cell activates with delay D with probability at least α. For brevity, let $L_{\min}(\tau_u)$ denote $L_{\min}(\tau_u; N, Z, D, \alpha)$.

To study the behavior of $L_{\min}(\tau_u)$ as $\tau_u \to \infty$, we define

$$L_{\min}(\infty) \triangleq L_{\min}\left(\lim_{\tau_u\to\infty}\tau_u\right) = \inf\left\{L > 0 : \lim_{\tau_u\to\infty}\mathbb{P}\left(T_{(0,0)}(\tau_u, L) \le D\right) \ge \alpha\right\}$$

where the last equality holds because the limit and probability can be exchanged for the same reason as in section C. In the limit of $\tau_u \to \infty$, the unbinding rate $k_u \to 0$, and therefore the rates $Q_{(n,z),(n-1,z)} = Q_{(n,z),(n-1,z-1)} \to 0$, i.e. the activation time should be approximated by a Markov chain where the state (n,z) is always monotonically increasing. To this end, define the Markov chain $M^{(f)}$ with only forward transitions i.e. the rate matrix $Q^{(f)}$ is given by

$$\begin{aligned}
Q^{(f)}_{(n,z),(n+1,z)} &= (N-n)k_bL\mathbb{1}\left(n<N, z<Z\right) \\
Q^{(f)}_{(n,z),(n,z+1)} &= (n-z)k_z^{on}\mathbb{1}\left(z<\min\{n,Z\}\right) \\
Q^{(f)}_{(n,z),(n,z)} &= -\left(Q^{(f)}_{(n,z),(n+1,z)} + Q^{(f)}_{(n,z),(n,z+1)}\right) \\
Q^{(f)}_{(n,Z),(\cdot,\cdot)} &= 0
\end{aligned}$$

and all other rates not specified are 0. The only difference between the original rate matrix Q, and the rate matrix $Q^{(f)}$ is that $Q^{(f)}_{(n,z),(n-1,z)} = Q^{(f)}_{(n,z),(n-1,z-1)} = 0$. Let $T^{(f)}_{(0,0)}$ denote the first passage time from the state $(0,0)$ to the boundary $\mathcal{B}$ for the chain $M^{(f)}$. Then we have the following result.

Theorem D.1 (Horizontal Asymptote)**.** *The minimum density $L_{\min}(\infty)$ required to activate in the limit $\tau_u \to \infty$ is given by*

$$L_{\min}(\infty) = L^{(f)}_{\min}(N, Z, D, \alpha) \triangleq \inf\left\{L > 0 : \mathbb{P}\left(T^{(f)}_{(0,0)}(L; N, Z) \le D\right) \ge \alpha\right\}.$$

Theorem D.1 implies that we can compute $L_{\min}(\infty)$ and understand its dependence on (N, Z, D, α) by computing the first passage time distribution for the simpler Markov chain $M^{(f)}$. Computing the Laplace transforms $\mathcal{L}_{(i,j)}(s)$ for $M^{(f)}$ reduces to solving a sparse upper triangular set of equations. These equations can be solved via back-substitution resulting in $\mathcal{O}(NZ)$ complexity as compared to $\mathcal{O}(N^2Z^2)$ for computing the Laplace transforms for the first passage times for M. Figure D.1

confirms that $L_{\min}(\infty) = L_{\min}^{(f)}$ and that $L_{\min}^{(f)}$ can be computed much faster than $L_{\min}(\infty)$ (~10-100 times faster for moderate-sized N and Z). Computing the horizontal asymptote is faster for smaller N and Z because it only requires one pass through the problem data (of size $\mathcal{O}(NZ)$), whereas computing the vertical asymptote requires two passes through the problem data (of size $\mathcal{O}(Z)$). The difference in complexity starts to show up once N and/or Z become larger.

Proof. We equivalently show that

$$\lim_{\tau_u \to \infty} \mathbb{P}\left(T_{(0,0)}(\tau_u, L) \leq D\right) = \mathbb{P}\left(T_{(0,0)}^{(f)} \leq D\right)$$

As in the analysis of the vertical asymptote, it suffices to show

$$\lim_{\tau_u \to \infty} \mathbb{P}_{(0,0)}\left(n_t = n, z_t = z\right) = \mathbb{P}_{(0,0)}^{(f)}\left(n_t = n, z_t = z\right)$$

for all $(n, z) \in \mathcal{S}$ and $t > 0$, with $\mathbb{P}_{(0,0)}^{(f)}$ being the distribution of the modified Markov chain $M^{(f)}$ given $(n_0, z_0) = (0, 0)$. We have

$$\mathbb{P}_{(0,0)}\left(n_t = n, z_t = z\right) = \mathbb{P}_{(0,0)}\left(\begin{smallmatrix} n_t = n, z_t = z, \text{ no backward} \\ \text{transitions before time } t \end{smallmatrix}\right) + \mathbb{P}_{(0,0)}\left(\begin{smallmatrix} n_t = n, z_t = z., \text{ at least one} \\ \text{backward transition before } t \end{smallmatrix}\right)$$

We also see that

$$\mathbb{P}_{(0,0)}\left(\text{no backward transitions before time } t\right) = e^{-\frac{t}{\tau_u}(N+Z)} \to 1,$$

as $\tau_u \to \infty$, and

$$\mathbb{P}_{(0,0)}\left(n_t = n, z_t = z | \text{no backward transitions by time } t\right) = \mathbb{P}_{(0,0)}^{(f)}\left(n_t = n, z_t = z\right)$$

Putting these two together tells us that

$$\lim_{\tau_u \to \infty} \mathbb{P}_{(0,0)}\left(\begin{smallmatrix} n_t = n, z_t = z, \\ \text{no backward transitions before time } t \end{smallmatrix}\right) = \mathbb{P}_{(0,0)}^{(f)}\left(n_t = n, z_t = z\right).$$

A similar argument establishes that $\lim_{\tau_u \to \infty} \mathbb{P}_{(0,0)}\left(\begin{smallmatrix} n_t = n, z_t = z, \\ \text{at least one backward transition before } t \end{smallmatrix}\right) = 0$. Therefore, $\lim_{\tau_u \to \infty} \mathbb{P}_{(0,0)}\left(n_t = n, z_t = z\right) = \mathbb{P}_{(0,0)}^{(f)}\left(n_t = n, z_t = z\right)$. □

E 1-D approximation for finite L, τ_u

In Sections C and D we established that the vertical and horizontal asymptotes of the $L_\alpha(\tau_u)$ curve can be characterized in terms of simpler Markov chains whose first passage time distributions can be computed much more quickly. In this section, we study the shape of the $\mathcal{L}_\alpha$ curve for finite values of L and τ_u. We do this because being able to efficiently calculate $\tau_\alpha(L)$ and $\mathcal{L}_\alpha(\tau)$ is essential for computing the fidelity metric discussed in section 2.3.

In order to efficiently simulate the entire $\mathcal{L}_\alpha$ curve and investigate its "distance" from ideal absolute ligand distribution behavior, we develop a 1-dimensional approximation of the full model M where the n- and z–dynamics are decoupled. We show that this decoupling significantly reduces the complexity of computing $(\tau_\epsilon, L_\epsilon)$.

Consider a birth-death process $B = \{b(t) : t \geq 0\}$ on the state space $\mathcal{Z} = \{0, 1, ..., Z\}$ with an absorbing boundary at Z. The transition rates are

$$\begin{aligned} \mu_{z,z-1} &= z/\tau_u, \\ \mu_{z,z+1} &= \frac{1}{\mathbb{E}\left[\gamma_z(L, \tau_u)\right]}, \end{aligned} \tag{17}$$

with $\gamma_z(L, \tau_u)$ being the absorption time in a z-dependent Markov chain defined below. See Figure E.1.

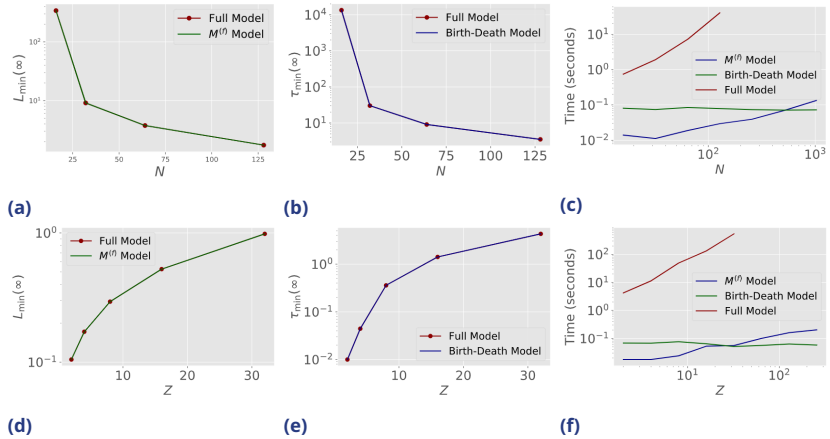

Figure D.1. Comparison of location of horizontal and vertical asymptotes computed using the full model vs. using the simplified models, as well as the speed to compute them. (a) $L_{\min}(\infty)$ computed using full model (red dots) and simplified model $M^{(f)}$ (green line), as a function of $N \in \{16, 32, 64, 128\}$, $Z = 16$, $\alpha = 0.95$, $D = 300$s. (b) $\tau_{\min}(\infty)$ computed using full model and birth-death process (blue line). N, Z, α, D same as (a). (c) Speed to compute $L_{\min}(\infty)$ using $M^{(f)}$, speed to compute $\tau_{\min}(\infty)$ using birth-death process, and speed to compute $\tau_{\min}(\infty)$ using full model, as a function of N (speed to compute $\tau_{\min}(\infty)$ and speed to compute $L_{\min}(\infty)$ using full model are the same). For $L_{\min}(\infty)$ and $\tau_{\min}(\infty)$, N values extended to $N \in \{16, 32, 64, 128, 256, 512, 1024\}$. (d)-(f) Same as (a)-(c) but as a function of Z. $N = 256$, $\alpha = 0.95$, $D = 450$s. $Z \in \{2, 4, 8, 16, 32\}$ for (a), (b), and full model in (c); $Z \in \{2, 4, 8, 16, 32, 64, 128, 256\}$ in (c) for simplified models.

Let $M^{(z)} = \{n^{(z)}(t) : t \geq 0\}$ denote a Markov chain on $\{z, z+1, ..., N\} \cup \{\mathcal{A}\}$, where the state $\mathcal{A}$ is an absorbing state. For state $i \in \{z, z+1, ..., N\}$ the transition rates are given by

$$\begin{aligned}
\mu^{(z)}_{i,i-1} &= (i-z)/\tau_u, \quad i > z, \\
\mu^{(z)}_{i,i+1} &= (N-i)k_b L, \quad i < N, \\
\mu^{(z)}_{i,\mathcal{A}} &= (i-z)k_{on}.
\end{aligned}$$

The Markov chain $M^{(z)}$ starts from the initial state $n^{(z)}(0)$ defined below. Before describing the procedure for selecting $n^{(z)}(0)$, we describe the rationale for this decoupling of the (n, z) dynamics.

The T-cell activates when $z = Z$. We would be able to efficiently estimate the quantities of interest if one could describe a Markov chain where the transition rates are only a function of the state z. The rate from state z to state $z-1$ in the 2-D Markov chain M is given by $\frac{z}{\tau_u}$, i.e. it only depends on z, and so we keep this rate as it is. However, the rate from state z to $z+1$ depends on the number of TCR-antigen complexes n, and therefore cannot be written as a function of only the state z. We know that $n \geq z$, and so define a Markov chain $M^{(z)}$ with the state space consisting of $\{z, \ldots, N\}$ and an absorbing state $\mathcal{A}$ that represents the event $z \to z+1$. Note that in the 2-D Markov chain, it's possible that n is much larger than z when the chain makes a transition from $z \to z+1$, and so setting $n^{(z)}(0) = z$ could significantly overestimate the first passage time. To correct this, we set the initial states $n^{(z)}(0)$ in the following recursive manner:

$$n^{(z)}(0) = \inf\left\{ i \geq n^{(z-1)}(0) : \mu^{(z-1)}_{i,i+1} < \mu^{(z-1)}_{i,\mathcal{A}} \right\},$$

i.e. $n^{(z)}(0)$ is set to the value of $i \geq n^{(z-1)}(0)$ where the rate of transition to the stopping state $\mathcal{A}$ is larger than the rate to the state $i+1$. Setting $n^{z}(0)$ in this manner allows the approximation to mimic the behaviour of the 2-D chain M both for when $L\tau_u$ is very large and the transitions in M are likely to be close to $(0,0) \to (1,0) \to (2,0) \cdots \to (N,0) \to (N,1) \cdots \to (N,Z)$, and when $L\tau_u$ is very small and the the transitions in M are likely to be close to $(0,0) \to (1,0) \to (1,1) \to (2,1) \to \cdots \to (N, Z-1) \to (N, Z)$. This setting also allows for more complicated transitions associated with moderate values of $L\tau_u$. Finally, we replace the first passage time $\gamma_z(L, \tau_u)$ from $n_z(0)$ to $\mathcal{A}$ by the maximum entropy non-negative random variable with the same mean, i.e. an exponential random variable with rate

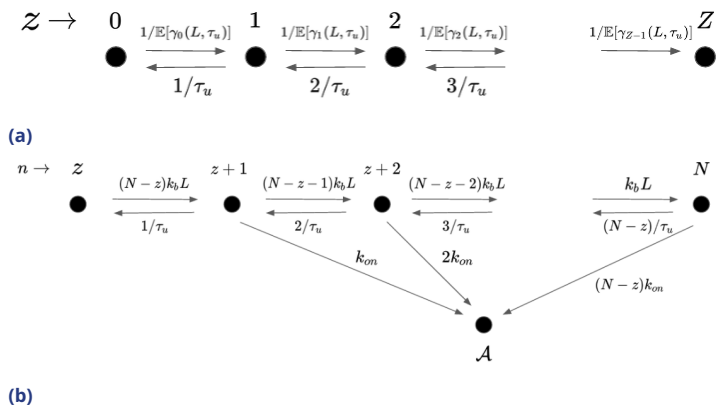

Figure E.1. Proposed model used to approximate the convergence of $\tau_{\min}(L) \rightarrow \tau_{\min}(\infty)$ and $L_{\min}(\tau_u) \rightarrow L_{\min}(\infty)$. The first passage time to the boundary $\mathcal{B}$ in M is approximated with the first passage time to Z in the birth-death process shown in (a). $\gamma_z(L,\tau_u)$ denotes the first passage time to $\mathcal{A}$ starting value $n^{(z)}(0)$ in the chain displayed in (b). We replace $\gamma_z(L,\tau_u)$ by the maximum entropy distribution with mean $\mathbb{E}[\gamma_z(L,\tau_u)]$, i.e. an exponential with rate $1/\mathbb{E}[\gamma_z(L,\tau_u)]$. Once the model in subfigure (a) transitions to a new value of z (so either to $z-1$ or $z+1$), the model here changes to reflect the change in z and restarts.

$\frac{1}{\mathbb{E}[\gamma_z(L,\tau_u)]}$. This explains the forward rate in (17).

Inverting the Laplace transform for the birth-death process B allows us to approximate τ_ϵ and L_ϵ. Since calculating the mean first passage time to $\mathcal{A}$ for each of the $M^{(z)}$ processes involves solving a tri-diagonal system of equations, this step takes at most $\mathcal{O}(N)$ time, resulting in a total complexity of $\mathcal{O}(NZ)$. The Laplace transform for the birth-death process can be computed in $\mathcal{O}(Z)$ time. Thus, the approximate model results in an $\mathcal{O}(NZ)$ speedup compared to using the model M.

Figures E.2 and E.3 compare the values of $\tau_\epsilon(N,Z,D,\alpha)$ and $L_\epsilon(N,Z,D,\alpha)$ computed using the full model compared to the values computed using the approximation detailed in this section. We see that the true and approximated values of τ_ϵ and L_ϵ are typically within a factor of 2 from one another, and the time to approximate these two values is $10-100$ times faster for the values of N and Z tested.

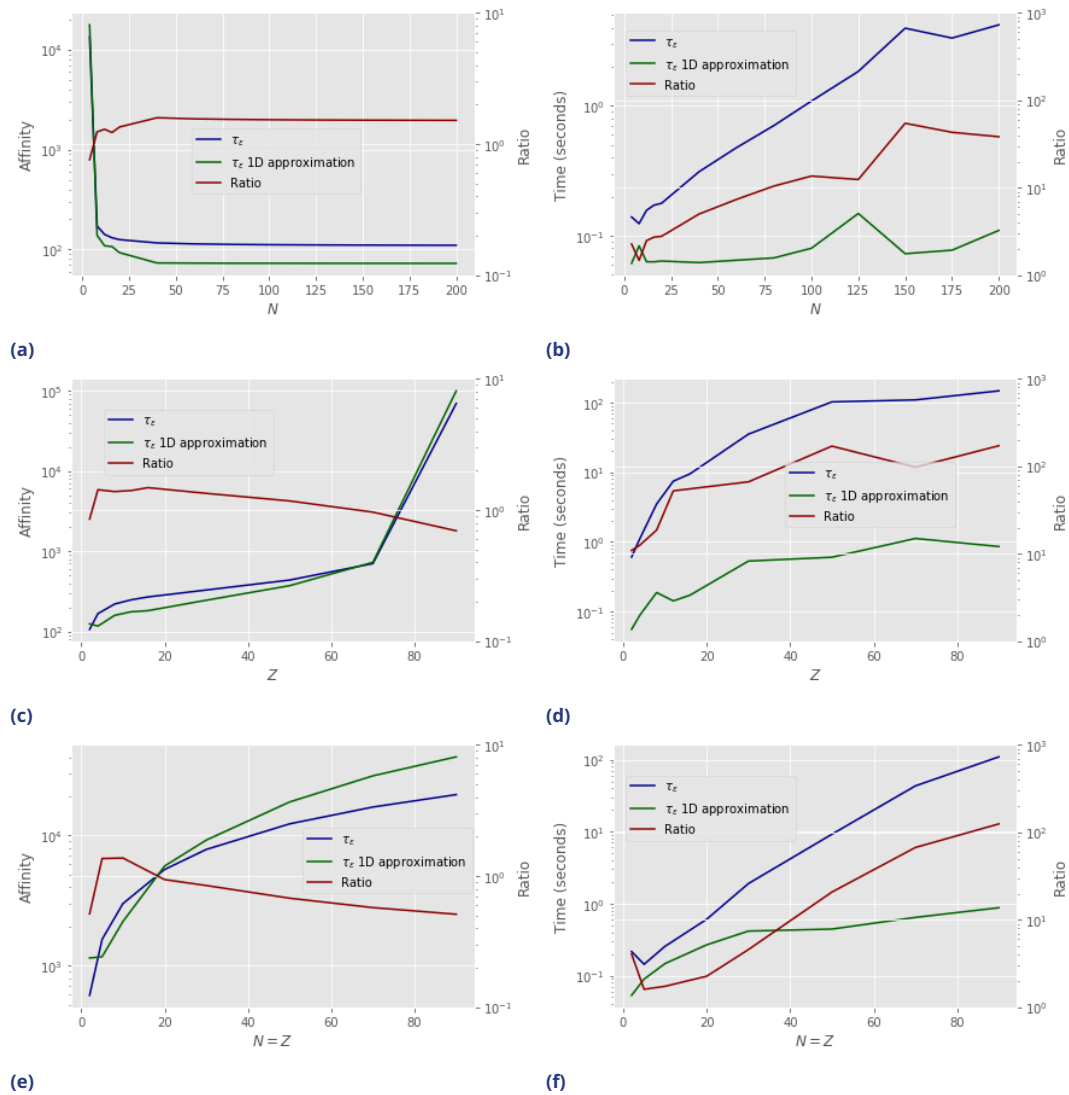

Figure E.2. Comparison of τ_ϵ using the full model M and the approximation in Figure E.1. $\epsilon = 1$ and $\alpha = 0.95$ for all figures. (a) τ_ϵ as computed by the full model M (blue line) and the approximation (green line), as well as the ratio of the two (red line), as a function of $N = 4, ..., 200$. $Z = 4$, $D = 225$. (b) Time to compute τ_ϵ using original model M (blue line) and approximation (blue line), as well as ratio between the two, as a function of N. Same parameters as (a). (c) Same as (a) but as a function of $Z = 2, ..., 90$. $N = 90$, $D = 400$. (d) Same as (b), but as a function of Z. Same parameters as (c). (e) Same as (a), but as a function of $N = Z = 2, ..., 90$, with D set to $1.25D^*(N, Z, \alpha)$. (f) Same as (b), but as a function of $N = Z$. Same parameters as (e).

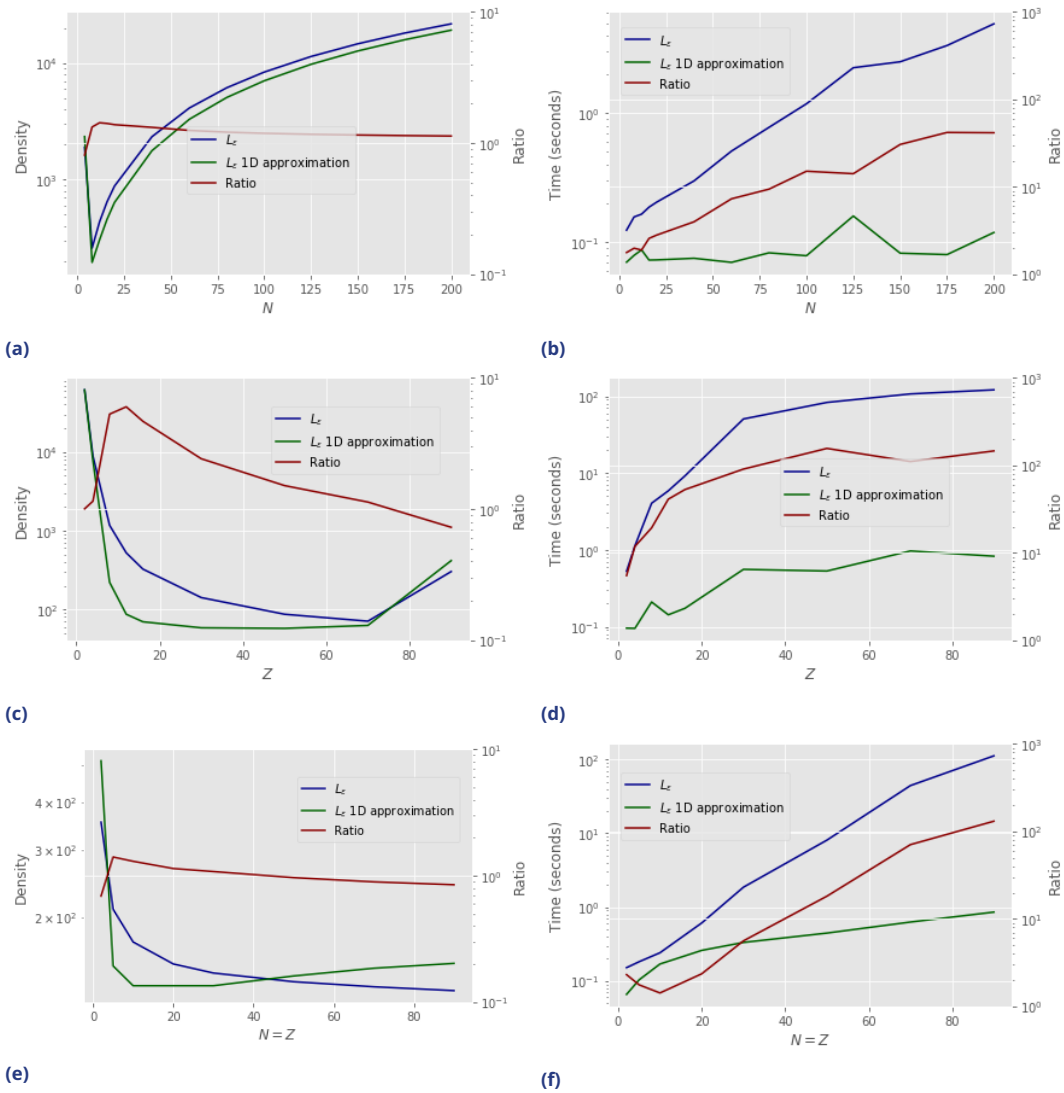

Figure E.3. Comparison of L_ϵ using original model M and approximation from Figure E.1. Same layout as Figure E.2.