

Optimal Maturation Protocols for High-Affinity Antibody Targets: A Path-Integral Approach

Marian Huot^{1,*}, Marco Molari^{1,*}, Rémi Monasson^{1,†,‡} and Simona Cocco^{1,†,§}

Laboratory of Physics of the Ecole Normale Supérieure, PSL Research & CNRS UMR8023, Sorbonne Université, Paris, France



(Received 19 December 2025; accepted 14 August 2026; published 29 September 2026)

Affinity maturation is a stochastic evolutionary process allowing the adaptive immune system to produce B cells capable of recognizing antigenic molecules. One of the main factors influencing the quality of the maturation outcome, quantified by the affinity of the produced antibodies to the antigen, is the time course of antigen availability. In this paper, we develop a stochastic model for affinity maturation, calibrated against *in vivo* lineage dynamics and deep mutational scanning data, and address the following question: What is the best antigen-concentration protocol maximizing the probability of reaching a target affinity value at the end of the process? We introduce a path-integral formalism to identify the maturation trajectories of the rare clones ending at the desired target affinity, estimate their probabilities, and maximize them over the antigen-concentration temporal profile. The theoretical optimal concentration protocol is approximated by a discrete three-injection schedule; we show that such temporal modulation of selection pressure outperforms constant-dosage regimes.

DOI: [10.1103/79gx-ctyk](https://doi.org/10.1103/79gx-ctyk)

I. INTRODUCTION

Vaccination remains one of the most effective interventions in medical history, yet the development of potent immunogens against highly mutable pathogens, such as HIV and influenza, continues to defy conventional empirical approaches. The efficacy of a vaccine relies fundamentally on its ability to trigger affinity maturation (AM), a rapid Darwinian evolutionary process occurring within germinal centers (GCs). In these microanatomical structures, B cells undergo iterative rounds of somatic hypermutation and selection, progressively evolving B-cell receptors (BCRs) capable of binding antigen with high affinity [1–4]. While the biological mechanisms governing this process—clonal expansion, mutation, and competitive selection for survival signals—are well characterized, the quantitative principles that determine the optimal evolutionary trajectory of a B-cell lineage remain an active area of theoretical investigation.

A central variable governing the efficiency of this evolutionary search is the availability of antigen. Antigen concentration acts as a tunable control parameter for selection pressure [5–7]: it determines the energetic threshold a B cell must overcome to capture antigen, present it to T follicular helper cells, and receive survival signals. Previous computational and experimental studies have highlighted a fundamental tradeoff regarding antigen dosage [8–10].

High concentrations relax selection pressure, allowing low-affinity clones to survive and dominate the population, thereby stalling the maturation process. In contrast, strict selection mediated by low antigen concentrations accelerates affinity gains but increases the risk of population extinction before high-affinity variants can emerge [11]. It is therefore important to identify the best antigen administration protocols, maximizing the likelihood of reaching a high-affinity final state.

We address the problem of optimal immunization by leveraging established computational models of affinity maturation [10–14]. Despite their idealized nature, these frameworks effectively encapsulate the stochastic fluctuations and complex selection pressures that define the process. Specifically, we investigate the following: What is the probability that a B-cell lineage achieves a target affinity, and how can the temporal profile of antigen be tuned to maximize this likelihood? Our objective is to identify the most probable evolutionary trajectory leading from a wild-type state to a specific affinity threshold, along with its associated probability of success.

To do so we introduce a mathematical framework that treats affinity maturation as a stochastic trajectory optimization problem on a fitness landscape. We map the discrete, agent-based dynamics of B-cell evolution onto a continuous path-integral formalism to identify the “least-action” trajectories—the most probable evolutionary histories leading to high-affinity states. We also take into account the fluctuations around these least-action trajectories. This formalism allows us to analytically derive the time-dependent antigen concentration that maximizes the probability of evolving high-affinity clones while respecting biological constraints such as population viability and carrying capacity.

We calibrate our model using time-resolved lineage data and deep mutational scanning profiles from *in vivo* experiments, ensuring that our theoretical predictions are grounded in realistic biophysical parameters. Contrary to the common intuition that stronger selection always yields better

*These authors contributed equally to this work.

†Co-last authors

‡Contact author: remi.monasson@phys.ens.fr

§Contact author: simona.cocco@phys.ens.fr




FIG. 1. Schematic of the affinity maturation model. The simulation proceeds through iterative rounds of evolution. (a) Duplication and somatic hypermutation: B cells proliferate and mutate, with outcomes classified as silent, lethal, or affinity affecting. (b) Selection and regulation: Cells undergo a two-step selection process involving antigen binding and competition for T-cell help. The population is subsequently regulated via differentiation and random removal to enforce carrying capacity. (c) Stochastic affinity trajectories of B cells over GC rounds. The red curves highlight the subset of target trajectories that successfully reach high-affinity states, which are the focus of this study.

antibodies, we show that the optimal protocol requires a dynamic modulation of selection pressure that keeps the population near a critical “tipping point” of extinction. We further demonstrate how this theoretically optimal continuous solution can be discretized into a practical multi-injection vaccination schedule, providing a rational basis for the design of immunization protocols that maximize evolutionary gain.

II. MODEL AND COMPARISON WITH EXPERIMENTAL DATA

A. Affinity maturation model

We describe a stochastic model for antibody affinity maturation within a GC. The variable component of the B-cell receptor affinity is quantified by the affinity parameter ϵ , measured in units of $k_B T$, where k_B is the Boltzmann constant and T is the organism’s temperature. This choice of units allows one to simply express Boltzmann factors as e^ϵ . This affinity is related to the dissociation constant, K_D , between the B-cell receptor and the antigen through $\epsilon = -\log(K_D) + \epsilon_a$, where ϵ_a is the wild-type log dissociation constant. Here, K_D

is expressed in nanomolar (nM) units; choosing different units would shift the values of ϵ by a constant amount.

To capture site-specific granularity, we define the B-cell affinity as the sum of these individual site contributions:

$$\epsilon = \sum_{i=1}^L \epsilon_i. \quad (1)$$

In particular, the wild-type (unmutated) B cell has $\epsilon_i = 0$ for all sites. The maturation process proceeds through iterative rounds of duplication, mutation, and selection (Fig. 1) [10,12] that modify the values of the ϵ_i ’s. For the sake of mathematical tractability, our model assumes that mutational effects are independent of the background, see Eq. (1), and that the evolution of B cells can be modeled as a succession of affinity-independent duplications and mutations, while selection is driven by affinity only.

Cells first duplicate in the GC dark zone and then mutate one site with probability μ for each sequence [Fig. 1(a)]. Mutation effects can be silent (with probability q_{sil}), lethal (q_{let} , resulting in the immediate removal of the cell), or affinity affecting (q_{aa}) [10,12,14] (see Table I). In the latter case, we

TABLE I. Parameter values used for the simulations.

Symbol	Reference	Description
$L = 40$	[15]	Number of CDR residues
$T = 140$	[15]	Number of GC rounds
$\epsilon_a = \log(40)$	[15]	$\log K_d(\text{WT})$ (nM)
$\Delta\epsilon_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$	[15]	Effect of site mutations
$\mu = 0.001 \times L \times 3$	[12]	Mutation probability per sequence of size L
$q_{\text{sil}} = 0.3$	[10,12,14]	Probability of silent mutation
$q_{\text{let}} = 0.5$	[10,12,14]	Probability of lethal mutation
$q_{\text{aa}} = 0.2$	[10,12,14]	Probability of affinity-affecting mutation
$N_{\text{max}} = 2500$	[10]	GC max population size
C	Free parameter	Antigen concentration (nM)

update a random site i by a random additive contribution $\Delta\epsilon_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$. We make these contributions site specific, with the intention of setting them from experimental measurements.

As a consequence, each daughter cell keeps the same affinity with probability $p_{\text{sil}} = (1 - \mu) + \mu q_{\text{sil}}$, dies with probability $p_{\text{let}} = \mu q_{\text{let}}$, or modifies affinity with probability $p_{\text{aa}} = \mu q_{\text{aa}}$.

Following mutation, cells migrate to the Light Zone to undergo selection [Fig. 1(b)]. In the first selection step, B cells attempt to bind the antigen (Ag) exposed on Follicular Dendritic Cells. The probability of successfully binding and surviving this step depends on the antigen concentration $C(t)$, expressed in nMol, and the total binding energy. Cells that fail to bind undergo apoptosis. The probability of antigen binding is

$$P_{\text{Ag}}(\epsilon, t) = \frac{C(t) \exp[\epsilon - \epsilon_a]}{1 + C(t) \exp[\epsilon - \epsilon_a]}. \quad (2)$$

Surviving cells compete for survival signals provided by T cells. This introduces a competitive element where B cells with higher affinity are preferentially selected. The probability of surviving this selection step is

$$P_T(\epsilon, t) = \frac{C(t) \exp[\epsilon]}{C(t) \exp[\epsilon] + \bar{\Phi}(t)}, \quad (3)$$

where $\bar{\Phi}(t)$ represents the mean fitness of the current population, acting as a competitive threshold. It is defined as the population average of the affinity term:

$$\bar{\Phi}(t) = \frac{1}{N(t)} \int \rho(t, \epsilon) \exp[\epsilon] d\epsilon, \quad (4)$$

where the population density $\rho(t, \epsilon)$ represents the expected number of cells with affinity ϵ at time t . The total population size is given by the integral

$$N(t) = \int d\epsilon \rho(t, \epsilon). \quad (5)$$

The $\bar{\Phi}$ -dependent term at the denominator of Eq. (3) ensures that the selection pressure adapts dynamically to the improving affinity of the population.

The surviving cells differentiate with uniform probability $p_{\text{diff}} = 10\%$. Cells that differentiate are removed from the population.

Finally, if the population exceeds the maximum carrying capacity N_{max} , cells in excess are randomly removed, irrespective of their affinity [10]. Iterating these GC cycles yields stochastic trajectories of affinity maturation, as illustrated in Fig. 1(c).

B. Comparison with experimental affinity maturation trajectories

We calibrate the model parameters in this stochastic algorithm using experimental data from DeWitt *et al.* [15], who investigated the reproducibility of evolutionary trajectories by “replaying” germinal center reactions starting from a monoclonal B-cell population. In their study, B cells specific for the chicken IgY antigen were transferred into GC-deficient mice that were subsequently immunized with IgY to initiate the affinity maturation process. Following immunization, the authors reconstructed the phylogenetic histories and affinity trajectories of B cells from over 100 individual GCs. Furthermore, they utilized Deep Mutational Scanning (DMS) to quantify the effects of virtually all possible single-amino acid substitutions on the receptor’s binding affinity.

We directly incorporate these empirical measurements into our model parameters. Specifically, the activation energy ϵ_a is derived from the measured wild-type affinity to antigen, and the stochastic update factor $\Delta\epsilon_i$ is sampled from a site-specific normal distribution $\mathcal{N}(\mu_i, \sigma_i^2)$ fitted on the amino acid mutational effects determined by their DMS experiments (Table I). We restrict our analysis to mutations that are one nucleotide away from wild type. We focus our analysis on the $L = 40$ CDR residues, as they exhibit significantly higher mutability in both light and heavy chains compared to other residues [15].

Then, we validate our simulation by reproducing the population-level affinity maturation trajectories observed in their time-course sequencing data. We first compare the mean affinity evolution of our stochastic simulation (averaged over ten independent realizations) against the experimental measurements [Fig. 2(a)]. By adjusting a single free parameter—a constant antigen concentration $C(t) = C_{\text{exp}}$ —our model accurately captures the dynamics of affinity gain, reproducing the characteristic rapid initial rise followed by a slower improvement observed *in vivo*. We emphasize that assuming a constant concentration here is a temporary simplification to establish a baseline against experimental data; the core of our subsequent theoretical analysis is dedicated to deriving and optimizing dynamic, time-dependent concentration protocols.

To investigate the lineage dynamics underlying these trajectories, we reconstruct the phylogenetic history of a representative simulated founder cell [Fig. 2(b)]. The resulting tree topology reveals a distinct distribution of fitness states: while the backbone of the lineage is dominated by high-affinity clones, B cells with lower affinities compared to the population mean are predominantly observed only at the final time steps (the leaves of the tree). This observation is consistent with the “pull of the present” effect described by DeWitt *et al.*, where recent deleterious mutations are sampled simply




FIG. 2. Comparison with experimental data from DeWitt *et al.* (a) Comparison of simulated affinity trajectories (mean of ten independent stochastic realizations) against experimentally measured affinity maturation in DeWitt *et al.* Only one free parameter (antigen concentration C) was adjusted to match the experimental results. (b) Evolutionary tree of a single founder cell, with daughter cells observed at the final time step. B cells with lower affinities relative to the population average are predominantly observed only at the sampling endpoint. (c) Comparison of simulated and experimental site-specific affinities at different time points of the maturation process. Pearson correlation coefficients with p value are indicated.

because selection has not yet had time to purge them from the population. Conversely, the historical nodes show a survivorship bias (“push of the past”) where only fitness-increasing mutations tend to fixate and leave descendants.

Furthermore, we demonstrate that our model preserves the statistical properties of site-specific affinity contributions. To do so, we compare the effects of single mutations observed in experiments with the site-specific affinity contributions derived from our simulations [Fig. 2(c)]. We observe strong

agreement across different time points, with the Pearson correlation coefficient reaching $\rho = 0.5$ (p value = 0.0011) at the end of simulation. Although we assume that mutational effects are independent of the current site affinity (similar to Wang *et al.* [12])—a simplification that neglects the biological reality of diminishing returns [16], where advantageous mutations become scarcer as affinity grows—we find that the simulated site-specific affinities remain consistent with the experimental range.

III. BACKTRACKING AFFINITY MATURATION TRAJECTORIES

A. Continuous description

We now derive a continuous description of the evolutionary dynamics. To make the problem computationally tractable, we reduce the high-dimensional shape space to a single effective binding-energy coordinate, denoted hereafter by the scalar ϵ . The mutational effects on affinity of this single site are obtained from a normal distribution $\mathcal{N}(\mu_M, \sigma_M^2)$ fitted on all sites mutational effects for single nucleotide away mutations. Note that the affinity updates drawn from this normal distribution share the same first and second moments as the pooled updates in the stochastic simulation, which are sampled from the Gaussian mixture $\frac{1}{L} \sum_{i=1}^L \mathcal{N}(\mu_i, \sigma_i^2)$.

In the limit of large population sizes and frequent, weak-effect mutations, the evolution of the binding-energy distribution $\rho(t, \epsilon)$ can be approximated by a continuity equation describing an advection-diffusion process.

We first decompose the deterministic growth and decay dynamics. The baseline growth rate, λ , aggregates the contributions of cell duplication, lethal mutations (occurring with probability p_{let}), and differentiation (probability p_{diff}) during a single evolution round:

$$\lambda = \log 2 + \log(1 - p_{\text{let}}) + \log(1 - p_{\text{diff}}). \quad (6)$$

Selection pressures modulate this baseline rate based on the cell's affinity. The net fitness landscape is shaped by the probability of antigen binding, $P_{\text{Ag}}(\epsilon, t)$, and the probability of receiving T-cell help, $P_T(\epsilon, t)$. The total energy-dependent growth rate is therefore

$$\Gamma(t, \epsilon) = \lambda + \log P_{\text{Ag}}(\epsilon, t) + \log P_T(\epsilon, t) - \Omega(t). \quad (7)$$

To strictly enforce the carrying capacity N_{max} , we introduce a pressure term $\Omega(t)$. This term prevents unbounded growth by penalizing the population growth rate when the global population size $N(t)$ exceeds the threshold:

$$\Omega(t) = \begin{cases} 0 & \text{if } N(t) < N_{\text{max}} \\ \max[0, \lambda + \langle \log P_{\text{Ag}} + \log P_T \rangle_{\text{pop}}] & \text{otherwise.} \end{cases} \quad (8)$$

Next, we characterize the stochastic effects of affinity-affecting mutations. While these mutations do not alter the instantaneous population size, they redistribute cells along the binding-energy axis. In the continuum limit, this jump process is approximated by a drift-diffusion operator defined by drift velocity v and diffusion coefficient D (see Appendix A). Physically, since the majority of random mutations are deleterious, we have a negative drift ($v < 0$) that acts as an entropic force opposing selection.

Combining these contributions, the time evolution of the population density $\rho(t, \epsilon)$ is governed by the following continuity equation:

$$\partial_t \rho(t, \epsilon) = \left[\Gamma(t, \epsilon) - v \partial_\epsilon + \frac{D}{2} \partial_\epsilon^2 \right] \rho(t, \epsilon). \quad (9)$$

This formulation is similar to a standard Fokker-Planck equation, but it is nonconservative due to the presence of the growth and decay term Γ .

The continuum (diffusion) limit requires a nearly continuous supply of mutations to neglect demographic stochasticity, meaning $N\mu \gg 1$ (in our simulations, $N\mu \approx 300$).

To systematically investigate the impact of antigen availability on evolutionary outcomes, we numerically solved the continuity equation over 140 GC cycles under varying antigen concentrations (Fig. 3). The dynamics reveal a critical dependency on the strength of selection pressure. Higher concentrations create a permissive environment where selection is weak [Figs. 3(a) and 3(b)]. Under these conditions the population maintains high diversity but fails to achieve significant affinity improvement relative to the wild type. On the other hand, at insufficient antigen concentrations, the stringent selection barrier exceeds the population's adaptive capacity, leading to extinction [Fig. 3(c)]. We benchmark this continuous approach against the discrete stochastic model introduced in a previous section [Fig. 3(c)]. Interestingly, we find that the population density in the continuous approximation tends to overestimate the speed of the affinity wave, resulting in final affinity distributions that are slightly shifted toward higher affinities compared to stochastic realizations, notably for smaller population sizes N_{max} [Fig. 3(d)]. This is a well-documented phenomenon in statistical physics and evolutionary dynamics [17,18]. In a continuous diffusion equation, the leading edge of the fitness wave contains infinitesimal, nonzero continuous densities (effectively, "fractions of a cell") that deterministically multiply and pull the wave forward. In contrast, affinity improvement in a discrete simulation is strictly bounded by the requirement of having at least one individual; densities below 1 act as a natural cutoff that slows the wave speed. While this can be corrected by introducing a population density truncation, we opt to maintain the smooth formalism to preserve consistency with our incoming least-action framework.

The optimal evolutionary outcome is achieved at intermediate concentrations; in this regime, selection pressure is sufficiently strong to drive the population toward high-affinity states yet permissive enough to sustain population viability, resulting in the most robust maturation [10].

B. Path-integral representation and backward trajectories

We next try to provide the shape of the most likely trajectory leading to any final energy state. The infinitesimal generator of the dynamics defined by the continuity equation allows us to construct the path integral for the trajectory $\epsilon(t)$ ending at the affinity ϵ_f at the final time t_f . After introducing the auxiliary conjugate field to enforce local constraints and integrating it out (see Appendix B), we obtain the path-integral representation [Fig. 4(a)]:

$$\rho(t_f, \epsilon_f) = \int d\epsilon_{\text{init}} \rho(t_0, \epsilon_{\text{init}}) \int_{\epsilon(t_0)=\epsilon_{\text{init}}}^{\epsilon(t_f)=\epsilon_f} \mathcal{D}[\epsilon] \exp[-S[\epsilon]]. \quad (10)$$

Here, the effective action S of a trajectory $\epsilon = \{\epsilon(t), 0 \leq t \leq t_f\}$ is given by

$$S(\epsilon) = - \int_{t_0}^{t_f} dt \left[\Gamma(t, \epsilon(t)) - \frac{1}{2D} (\dot{\epsilon}(t) - v)^2 \right]. \quad (11)$$




FIG. 3. Impact of antigen concentration on affinity maturation outcomes. (a), (b) Temporal evolution of the B-cell population density, obtained by numerically solving the continuity equation over 140 GC cycles under varying antigen concentrations. Color indicates the population density at different time steps. Excessively high concentration results in relaxed selection pressure and a lower final affinity of the B-cell population. (c) Comparison of affinity maturation outcomes between a stochastic simulation and the continuous approximation across varying antigen concentrations. (d) Impact of population size N_{max} on mean affinity in stochastic simulation (computed over ten runs) and continuous approximation. Error bars denote the standard deviation of the affinity distributions.

The path integral is dominated by the saddle-point trajectory that minimizes this action. For a population starting with a given affinity ϵ_{init} , the population density at the final time is approximated by the contribution of this single “most likely” evolutionary path:

$$\rho(\epsilon_f) \propto \exp[-S_{min}(\epsilon_{init}, \epsilon_f)]. \quad (12)$$

We identify the effective time-dependent Lagrangian,

$$\mathcal{L} = -\Gamma(t, \epsilon) + \frac{1}{2D}(\dot{\epsilon} - v)^2. \quad (13)$$

Minimizing the action via the Euler-Lagrange equation,

$$\frac{d}{dt} \partial_{\dot{\epsilon}} \mathcal{L} - \partial_{\epsilon} \mathcal{L} = 0, \quad (14)$$




FIG. 4. Least-action analysis of affinity maturation trajectories obtained with fitted experimental concentration $C_{\text{exp}} = 170$. (a) Schematic representation of the path-integral formulation, where a continuous trajectory is discretized into steps of size δt . (b) Reconstruction of the most likely evolutionary history. The dashed line represents the theoretical least-action trajectory, computed by minimizing the action $S[\varepsilon]$ conditioned on a fixed final affinity of $\varepsilon_f = 10$. $\varepsilon_f = 10$ corresponds to highest affinity reached in experiment [15]. (c) Comparison of theoretical least-action path and stochastic lineages that successfully evolved to the target affinity $\varepsilon_f \approx 10$. (d) Comparison of theoretical least-action path and mean of stochastic lineages that successfully evolved to the target affinity ε_f . (e) Validation of the least-action approximation. The minimum action $S_{\min}$ corresponding to a specific final state scales linearly with the logarithm of the final population density $\rho(\varepsilon_f)$, confirming the approximation $\rho \propto \exp(-S_{\min})$. S_0 is a constant fixed to the average difference between $S_{\min}$ and $\log \rho$. (f) Comparison of the entropic correction $-\Delta S$ and the error made by the least-action approximation, compared to the population density in continuous approximation. ΔS_0 is a constant offset introduced to ensure both quantities are plotted on a comparable range.

yields the equation of motion for the optimal trajectory (see Appendix C):

$$\partial_\epsilon \Gamma(t, \epsilon(t)) + \frac{1}{D} \ddot{\epsilon}(t) = 0. \quad (15)$$

This second-order ordinary differential equation describes a “particle” moving through the fitness landscape Γ under the influence of an effective inertial term proportional to the inverse diffusion coefficient.

By imposing boundary conditions fixed at the wild-type state $\epsilon(t_0) = 0$ and a target $\epsilon(t_f) = \epsilon_f$, we numerically solve the corresponding Euler-Lagrange equation [Fig. 4(b)]. The resulting least-action trajectory accurately traces the ridge of the population density obtained from the direct integration of the continuity equation, identifying the most probable route through the dynamic fitness landscape. This saddle-point approximation relies on the weak-noise limit, requiring the diffusion coefficient $D \ll 1$ (in our simulations, $D \approx 0.08$). Moreover, we test the predictive power of this theoretical framework against the microscopic dynamics of the discrete model. We isolate the subset of stochastic realizations that successfully reach the target affinity ϵ_f by the final time point [Figs. 4(c) and 4(d)]. Notably, we observe in both the stochastic simulations and the least-action paths that the divergence between high- and low-affinity trajectories becomes particularly pronounced after one month of maturation [Fig. 4(d)]. While individual lineages fluctuate due to the stochastic nature of mutation and selection, their ensemble average is remarkably well described by the deterministic least-action solution. This agreement demonstrates that the Euler-Lagrange equation captures the mean evolutionary dynamics of successful lineages.

C. Probabilities of best and nearby trajectories

We validate the least-action approximation by systematically comparing the computed action against the numerical solution of the population density in the continuous approximation. In this limit, the population density $\rho(\epsilon_f) \equiv \rho(t_f, \epsilon_f)$ is expected to be proportional to $\exp[-S_{\min}(\epsilon_{\text{init}}, \epsilon_f)]$. In Fig. 4(e) we plot the negative minimized action against the logarithm of the population density for various final states. The strict linear relationship confirms that the system operates in a regime where the path integral is indeed dominated by the single most likely evolutionary path.

The approximation of Eq. (12) does not consider the contribution to the total probability of nearby trajectories, which are less likely than the least-action one but more numerous. To estimate the entropic correction to the action computation due to this bundle of trajectories, we start by linearly expanding the log-argument of the integral—the action $S[\epsilon]$ —around the least-action trajectory $\hat{\epsilon}$ connecting ϵ_{init} at time $t = 0$ to ϵ_f at time $t = t_f$. We define the Hessian matrix of the action through the second-order functional derivatives around this optimal trajectory:

$$\mathbf{H}(t, t') = \frac{\delta^2}{\delta \epsilon(t) \delta \epsilon(t')} S[\epsilon] \Big|_{\hat{\epsilon}}. \quad (16)$$

This Hessian matrix describes the local curvature of the trajectory landscape around the optimal path. Integrating over

the fluctuations of ϵ leads to a Gaussian integral, resulting in an additive contribution to the least action:

$$\Delta S = \frac{1}{2} \ln \det \mathbf{H}. \quad (17)$$

Physically, this correction is equal to (minus) the entropy of the bundle of trajectories surrounding the optimal path. A small determinant corresponds to a shallow valley in the action landscape (high entropy), allowing for significant fluctuations, whereas a large determinant indicates a steep valley where trajectories are tightly confined. Details about the calculation of the Hessian matrix and its determinant are found in Appendix D. We validate this higher-order approximation in Fig. 4(f). In particular, we find that the small discrepancies observed in the zeroth-order approximation is perfectly explained by the entropic correction.

Finally, we further validate the robustness of this framework by performing the least-action analysis across a wide range of antigen concentrations [Figs. 5(a) and 5(b)]. We find that the entropic correction is minor and that the least action alone reliably represents the log population density.

IV. OPTIMAL PROTOCOLS

A. Optimizing growth of high-affinity trajectories

We seek the optimal antigen administration protocol $C(t)$ that maximizes the likelihood of a specific evolutionary trajectory. In practice, we minimize the least action $S[\hat{\epsilon}]$ with respect to the control parameter $c(t) = \log C(t)$. We obtain two distinct regimes (see Appendix E):

$$\begin{aligned} \frac{\delta S}{\delta c(t)} &= -\frac{\partial \Gamma(t, \epsilon(t))}{\partial c(t)} \\ &= \begin{cases} P_{\text{Ag}} + P_T - 2, & \text{if } N < N_{\text{max}} \text{ or } \Gamma_{\text{pop}} < 0, \\ (P_{\text{Ag}} + P_T) - \langle P_{\text{Ag}} + P_T \rangle_{\text{pop}}, & \text{otherwise.} \end{cases} \end{aligned} \quad (18)$$

Here, $\Gamma_{\text{pop}} \equiv \langle \log P_{\text{Ag}} + \log P_T \rangle_{\text{pop}} + \lambda$ represents the net population growth rate. Consequently, minimizing the action with respect to antigen concentration is equivalent to maximizing the instantaneous growth rate along the evolutionary trajectory. We define the *critical concentration*, C_0 , as the threshold where the net selection pressure exactly counterbalances the baseline rates of growth, mutation, and differentiation:

$$\Gamma_{\text{pop}}(C_0) = (\log P_{\text{Ag}} + \log P_T)_{\text{pop}} + \lambda = 0. \quad (19)$$

Below this threshold ($C < C_0$), the system operates in an unconstrained growth regime where the population pressure term $\Omega(C)$ vanishes [Fig. 6(a)]. Since the Ag-binding and T-cell selection probabilities are bounded by unity, the derivative of the action remains negative ($P_{\text{Ag}} + P_T - 2 < 0$). Thus, the effective growth rate Γ increases monotonically with concentration up to C_0 , implying that the optimal protocol must maintain $C(t) \geq C_0$ to keep the population at carrying capacity N_{max} [Fig. 6(b)]. Conversely, above C_0 , the pressure term $\Omega(C)$ activates to enforce the population limit. In this regime, low-affinity B cells continue to benefit from higher antigen availability, but high-affinity clones experience a decrease in growth rate as the global population penalty Ω outweighs their saturated binding probabilities. Figure 6(b) illustrates this by comparing the final-step growth rate of a high-affinity target




FIG. 5. Validation of the least-action approximation across varying antigen concentrations. (a) Comparison of least-action and log population density. (b) Comparison of least action with entropic correction and log population density.

cell ($\epsilon = 10$, matching the experimental maximum) against a lower-affinity cell ($\epsilon = 6$), demonstrating that the optimal growth for high-affinity targets is strictly achieved at C_0 .

Since C_0 depends on $\langle \log P_{Ag} + \log P_T \rangle_{pop}$, it varies as a function of time. Furthermore, because modifying the antigen concentration alters the dynamic fitness landscape and consequently shifts the optimal least-action trajectory $\epsilon(t)$, determining the true optimal protocol maximizing growth rate of high-affinity B cells requires an iterative procedure. We strictly alternate between solving the continuity equation to recompute the population affinity and updating the optimal $C_0(t)$ based on the evolving population mean until convergence is reached. As population affinity gets higher with GC rounds, the resulting optimal concentration $C_0(t)$ decreases [Fig. 6(c)].

DeWitt *et al.* [15] observed that clonal bursts were not a main driver of affinity maturation and that cells initiating these bursts were not affinity outliers. However, our model can be extended to account for other biological evidence suggesting that higher-affinity B cells undergo more divisions [19,20]. In this framework, mutation is “silenced” during rapid expansion phases, which serves to protect high-quality clones from losing their affinity due to random mutational degradation [20,21]. Precisely, building on studies observing that cell division is proportional to the amount of antigen captured

and presented, we can model the division rate component as $\lambda(\epsilon) = \lambda_0 + kP_{Ag}(\epsilon)$ (see Appendix F). In this scenario, the optimal concentration C_0 is shifted to slightly lower values; because higher-affinity clones benefit from an increased proliferation rate, less antigen is required to sustain the population at carrying capacity.

B. Optimal immunization scheme with discrete injections

The case of continuous antigen control is instructive but lacks realism. In reality, antigen is usually administered in a small number of injections and decays exponentially over time. For an immunization scheme consisting of N_{inj} injections, we can consider the antigen concentration to evolve as

$$C_{vaccine}(t) = \sum_{i=1}^{N_{inj}} \theta(t - t_i) C_i e^{-(t-t_i)/\tau}, \quad (20)$$

where $\tau = 20$ GC cycles is the antigen decay time. For every injection $i = 1, \dots, N_{inj}$, t_i corresponds to the injection time and C_i is the amount of antigen administered (related to the logarithmic amount by $C_i = e^{c_i}$). The function $\theta(\Delta t)$ is the Heaviside step function.

To find the optimal injection schedule, we employ a two-stage optimization procedure. First, we minimize the




FIG. 6. Optimal vaccination protocols. (a) Carrying capacity pressure $\Omega(C)$ as a function of antigen concentration. The pressure term activates above a threshold C_0 , mediating the random removal of cells independently of affinity, and converges to the baseline growth rate λ in the limit of infinite concentration. (b) Antigen concentration that maximizes the B-cell growth rate. We show the growth rate at the final step of the continuous approximation for a target-affinity B cell ($\epsilon = 10$, matching the highest experimental affinities) and a lower-affinity cell ($\epsilon = 6$). (c) Optimal concentration protocol maximizing the probability of high-affinity trajectories, alongside its best approximation using tunable vaccine injections, using antigen decay time $\tau = 20$. (d) Final B-cell affinity distributions comparing effect of different vaccination schedules using antigen decay time $\tau = 20$. (e) Effect of the antigen decay time τ on the final mean affinity, computed via the continuity equation across different vaccination schedules. (f) Validation of vaccine schedule efficiency using stochastic simulations using antigen decay time $\tau = 20$. The final mean affinity (averaged over ten runs, with extinct populations assigned zero affinity) is compared against a constant-dosage baseline and a standard three-dose regimen (one injection every four weeks with same concentration as in first dose of optimized three-dose regimen).

deviation between the discrete parametrization $C_{\text{vaccine}}(t)$ and the theoretically optimal continuous antigen control $C_0(t)$ derived in the previous section using the Mean Squared Error (MSE) loss function $\mathcal{L}_{\text{MSE}}$:

$$\mathcal{L}_{\text{MSE}}(\{t_i, C_i\}) = \frac{1}{N_\tau} \sum_{k=0}^{N_\tau} [C_{\text{vaccine}}(t_k) - C_0(t_k)]^2. \quad (21)$$

Using the results of the MSE minimization as an initial guess, we then perform a secondary optimization where we minimize the objective function $\mathcal{L}_{\text{total}} = S(\epsilon_f) + \mathcal{L}_{\text{extinction}}$ with gradient descent. The extinction penalty, defined as $\mathcal{L}_{\text{extinction}} = \int_{t_0}^{t_f} \max(0, \frac{1}{N(t)} - 1)^2 dt$, introduces an asymptotic barrier that activates exclusively when the population drops below the single-cell threshold ($N(t) < 1$). This barrier prevents infinitesimal populations from surviving tight bottlenecks in the continuous description. Here, the target $\epsilon_f = 30$ corresponds to the maximal affinity reached in all simulations. The gradients of $\mathcal{L}_{\text{total}}$ with respect to the $2N_{\text{inj}}$ injection parameters are computed numerically via finite differences by simulating the continuous population dynamics for each parameter perturbation and evaluating the action on the resulting optimal trajectory.

The results of this optimization are presented in Fig. 6(c). We observe that a simplified protocol consisting of $N_{\text{inj}} = 3$ tunable injections is sufficient to maintain antigen levels close to the least-action solution, preventing premature population extinction while sustaining selection pressure. To validate the efficacy of this protocol, we simulate the continuity equation under various immunization regimes [Fig. 6(d)], finding that the three-injection schedule yields the highest affinity.

To further assess the robustness of these schedules, we investigate the influence of the antigen exponential decay time, τ , on the maturation outcome [Fig. 6(e)]. While the three-injection protocol remains superior across most regimes, we find that larger values of τ (slower decay) make a single injection more competitive. Lastly, we cross-validate our theoretically optimized protocols using the discrete stochastic simulation [Fig. 6(f)]. Although the protocols are derived using the continuous least-action formalism, they remain highly effective in the stochastic regime, consistently outperforming constant-dosage regimes. This confirms that the principle of tracking the critical concentration C_0 provides a robust heuristic for vaccine design, even in the presence of the demographic noise inherent to germinal center reactions.

V. DISCUSSION

In this work, we have combined stochastic population dynamics with analytical optimization to define immunization protocols that maximize evolutionary gain. While previous studies [8–11] have established that antigen dosage is a critical determinant of affinity maturation outcomes, they primarily relied on the forward simulation of static dosages or predefined heuristic schedules. A rigorous derivation of the optimal temporal immunization profile has remained elusive. In particular, while using similar modeling assumptions, we advance beyond the descriptive modeling framework of Molari *et al.* [10] by introducing a path-integral formulation that transforms the analysis from a forward simulation of population

dynamics into an inverse optimal control problem. A direct practical prediction of this least-action framework, which was unavailable in prior literature, is the analytical derivation of a continuous, time-dependent antigen-concentration protocol that actively maintains the germinal center at a critical selection threshold to maximize adaptation without causing extinction.

Our results demonstrate that the most effective protocols rely on a dynamic modulation of selection pressure that guides the population along a “least-action” evolutionary trajectory. The main innovation of our approach is the expression of complex agent-based dynamics as a path integral over all possible maturation trajectories. We then estimate this path integral based on the least-action path, similar to the WKB approximation, a method increasingly utilized in evolutionary biology to solve problems where extinction [22] or adaptation [23] is driven by rare, large fluctuations. We also compute the Gaussian fluctuations around the least-action trajectory that maximizes the probability of reaching a final high-affinity state. We validate this framework by directly incorporating empirical measurements into our model parameters, specifically, the activation energy and mutational effects derived from Deep Mutational Scanning (DMS) experiments by DeWitt *et al.* [15]. The remarkable agreement between our analytical predictions and the ensemble average of successful stochastic lineages confirms that, despite the inherent randomness of mutation, the mean evolutionary path is deterministic and predictable under strong selection.

Our findings corroborate and quantify the long-standing observation that affinity maturation is not monotonically improved by higher antigen dosages. Consistent with prior literature [2,7,8,10,11], we observe that excessive antigen concentrations relax selection pressure, allowing low-affinity clones to persist and dominate the population through sheer numbers rather than fitness. Conversely, insufficient antigen leads to overly stringent bottlenecks that extinguish the population before beneficial mutations can fixate. Our model formalizes this tradeoff by identifying a critical, time-dependent concentration that keeps the system effectively at the “tipping point” of extinction. Beyond schedule timing, the availability of time-resolved sequence data now allows for the reconstruction of B-cell phylogenetic trees within individual subjects [24–27]. Tracking individual phylogeny could further improve our understanding of how vaccination protocols impact maturation by comparing these real histories to our theoretical “least-action” paths [28]. This comparison would help identify the specific selection pressures required to drive lineages toward high-affinity states, bridging the gap between clinical observation and deterministic evolutionary theory.

A key biological insight is that the optimal injection schedule maintains high selectivity over time. While our model does not incorporate all complexities of the immune response—such as the distinct recruitment kinetics of memory B cells [10,29,30]—it positions the injection protocol as a key factor, here supporting delayed second and third injections that effectively maintain a nonzero antigen concentration. This approach aligns with “slow delivery” immunization strategies, which have been shown to enhance germinal center responses in the context of HIV [31]. This principle is further supported by extensive research into SARS-CoV-2 strategies,

where delayed dosing has consistently improved outcomes. In particular, extending the interval between doses has been associated with higher neutralizing antibody titers and stronger long-term protection against infection compared to standard intervals [32,33].

The fundamental principle of the optimal schedule—dynamically maintaining the population at the critical selection threshold—can be adapted across plausible variations in key biological parameters. For instance, the carrying capacity $N_{\max}$, estimated in the range of a few thousand [34], determines the system’s resilience to selection pressure; a larger $N_{\max}$ reduces demographic stochasticity and provides a larger population buffer, allowing the optimal protocol to tolerate more aggressive selection via lower antigen concentrations $C(t)$. This larger population size also accelerates affinity maturation by increasing the total number of mutational trials occurring simultaneously. Similarly, variations in the mutation rate—typically estimated between 0.7 and 1.1×10^{-3} per base pair per division [35]—directly modulate the effective diffusion coefficient D in our continuous approximation. An increased mutation rate accelerates the exploration of the binding-energy landscape, leading to a faster affinity maturation. Consequently, the optimal continuous protocol should be adjusted with a marginally lower antigen concentration to maintain the population at carrying capacity. Furthermore, the initial affinity distribution dictates the starting coordinate of the evolutionary trajectory. A founder population seeded with a higher average affinity or greater standing variance requires a shorter evolutionary path. This correspondingly reduces both the required duration and the total cumulative dosage of the optimal schedule.

To achieve mathematical tractability, we introduce necessary simplifications. In our model the fitness advantage is captured qualitatively through the selection probabilities P_{Ag} and P_r . In addition, while individual follicles in our model risk extinction at population bottlenecks, inter-GC migration allows the system to share successful clones across follicles, ensuring that local failures do not result in a total population collapse [36]. The framework can be enriched by incorporating GC initiation and explicit antigen depletion dynamics. *In vivo*, antibody-mediated feedback creates “antigen masking,” where secreted antibodies block epitopes on Follicular Dendritic Cells, raising the selection threshold for competing B cells [37]. Beyond masking, high-affinity clones further tighten stringency through active antigen depletion—physically internalizing the resource and accelerating the extinction of lower-affinity competitors [12]. Qualitatively, this feedback suggests that maintaining the population along an optimal evolutionary trajectory would require higher cumulative dosages and more frequent or escalating injection schedules than our current model predicts, as they must compensate for the accelerated loss of available antigen caused by both antibody masking and active B-cell consumption.

Finally, while our evolutionary target is high affinity to a single antigen, effective intervention against highly mutable pathogens like HIV [12–14] and SARS-CoV-2 [25] requires targeting breadth—the capacity to bind a diverse spectrum of viral variants. This objective was recently modeled by Mahdisoltani *et al.* [38], who similarly applied a path-integral framework to identify optimal vaccination

protocols for broadly neutralizing antibodies. However, our study differs fundamentally in modeling choices and approximation schemes. Biologically, they model affinity based on the distance between B cells and vaccine antigens within a theoretical, multidimensional space. In contrast, our model bypasses this spatial abstraction by tracking a single binding-energy coordinate that is directly calibrated using empirical Deep Mutational Scanning data.

Ultimately, translating these theoretical insights into clinical applications will require rigorous experimental validation. Concrete *in vivo* tests, such as comparing our proposed three-injection schedule against standard immunization regimens in mouse GC models, would be highly valuable. By utilizing lineage sequencing to measure the probability of reaching specific affinity thresholds and comparing the observed clonal phylogenies against our predicted least-action lineages, future studies could directly test the predictive power of this path-integral framework.

ACKNOWLEDGMENTS

We are grateful to A. Chakraborty for useful discussions. We acknowledge the reviewer for suggesting an improved vaccine parameter optimization procedure.

M.H., M.M., S.C., and R.M. designed the research, analyzed the results, and wrote the paper. M.H. and M.M. carried out numerical analysis and simulations.

DATA AVAILABILITY

The source code and data that support the findings of this article are publicly available [39].

APPENDIX A: DRIFT AND DIFFUSION COEFFICIENTS

If we exclude lethal mutations (their effect is already accounted for in the growth term λ), then the probability for a cell to undergo a binding-energy change $\epsilon \rightarrow \epsilon + \Delta\epsilon$ in one round is given by the kernel K :

$$K(\Delta\epsilon) = f_{\text{aa}} K_{\text{aa}}(\Delta\epsilon) + f_{\text{sil}} \delta(\Delta\epsilon). \quad (\text{A1})$$

Here, $f_{\text{aa}} = \frac{p_{\text{aa}}}{p_{\text{aa}} + p_{\text{sil}}}$ and $f_{\text{sil}} = 1 - f_{\text{aa}}$ are the fractions of affinity-affecting and silent mutations, respectively. The kernel K_{aa} is a normal distribution with mean μ_M and variance σ_M^2 , while δ is the Dirac delta function centered at zero.

In the limit of large population size, affinity-affecting mutations can be described by a diffusion process. The effective drift and diffusion coefficients are then

$$v = \int d\Delta\epsilon \Delta\epsilon K(\Delta\epsilon) = f_{\text{aa}} \mu_M, \quad (\text{A2})$$

$$D = \int d\Delta\epsilon (\Delta\epsilon - v)^2 K(\Delta\epsilon) = f_{\text{aa}} (\sigma_M^2 + f_{\text{sil}} \mu_M^2). \quad (\text{A3})$$

For the numerical simulation of the continuity equation, we employed a time step of $\delta t = 0.001$ and a spatial discretization of $\delta\epsilon = 0.05$.

APPENDIX B: PATH-INTEGRAL FORMULATION

For a short time increment $\delta t \ll 1$, we formally write the evolution of the population density as

$$\rho(t + \delta t, \epsilon) \simeq \exp \left\{ \delta t \left[\Gamma(t, \epsilon) - v \partial_\epsilon + \frac{D}{2} \partial_\epsilon^2 \right] \right\} \rho(t, \epsilon). \quad (\text{B1})$$

We insert the Dirac delta in Fourier representation,

$$\delta(\epsilon - \epsilon') = \int \frac{d\tilde{\epsilon}}{2\pi} \exp[-i\tilde{\epsilon}(\epsilon - \epsilon')], \quad (\text{B2})$$

and evaluate (B1) at an initial position ϵ_{init} and a final position $\epsilon_f \equiv \epsilon_{\text{init}} + \Delta\epsilon$. We obtain

$$\rho(t + \delta t, \epsilon_f) = \int d\epsilon_{\text{init}} \int \frac{d\tilde{\epsilon}}{2\pi} \exp \left\{ \delta t \left[\Gamma(t, \epsilon_{\text{init}}) + i\tilde{\epsilon}v - \frac{D}{2} \tilde{\epsilon}^2 - i\tilde{\epsilon} \frac{\Delta\epsilon}{\delta t} \right] \right\} \rho(t, \epsilon_{\text{init}}). \quad (\text{B3})$$

Equation (B3) is the *one-step propagator* between two consecutive time slices of duration δt .

We partition the total interval $[t_0, t_f]$ into $T = t_f - t_0 = N_\tau \delta t$ steps labeled by an integer $\tau = 0, \dots, N_\tau - 1$. Repeatedly inserting (B3) yields the joint integral over all intermediate positions $\{\epsilon_\tau\}$ and their conjugate Fourier modes $\{\tilde{\epsilon}_\tau\}$:

$$\rho(t_f, \epsilon_f) = \int d\epsilon_{\text{init}} \rho(t_0, \epsilon_{\text{init}}) \int \left[\prod_{\tau=0}^{N_\tau-1} \frac{d\tilde{\epsilon}_\tau}{2\pi} \prod_{\tau=1}^{N_\tau-1} d\epsilon_\tau \right] \times \exp \left\{ \sum_{\tau=0}^{N_\tau-1} \delta t \mathcal{L}_\tau \right\}, \quad (\text{B4})$$

where we introduced the *discrete Lagrangian*:

$$\mathcal{L}_\tau = \Gamma_\tau + i\tilde{\epsilon}_\tau v - \frac{D}{2} \tilde{\epsilon}_\tau^2 - i\tilde{\epsilon}_\tau \frac{\epsilon_{\tau+1} - \epsilon_\tau}{\delta t}, \quad \Gamma_\tau \equiv \Gamma(t_\tau, \epsilon_\tau). \quad (\text{B5})$$

Equation (B4) is *quadratic* in $\tilde{\epsilon}$, which can be integrated out analytically using the Gaussian identity:

$$\int d\tilde{\epsilon} \exp \left[-\frac{a}{2} \tilde{\epsilon}^2 + ib\tilde{\epsilon} \right] = \sqrt{\frac{2\pi}{a}} \exp \left[-\frac{b^2}{2a} \right], \quad \Re a > 0.$$

Setting $a = D\delta t$ and $b = \delta t v - (\epsilon_{\tau+1} - \epsilon_\tau)$, and performing all N integrations gives the purely *real* sum:

$$\rho(t_f, \epsilon_f) = \int d\epsilon_{\text{init}} \rho(t_0, \epsilon_{\text{init}}) \frac{1}{\sqrt{2\pi\delta t D}} \left[\prod_{\tau=1}^{N_\tau-1} \frac{d\epsilon_\tau}{\sqrt{2\pi\delta t D}} \right] \times \exp \left\{ \sum_{\tau=0}^{N_\tau-1} \delta t \left[\Gamma_\tau - \frac{1}{2D} \left(v - \frac{\epsilon_{\tau+1} - \epsilon_\tau}{\delta t} \right)^2 \right] \right\}. \quad (\text{B6})$$

Taking the continuum limit $\delta t \rightarrow 0$ ($N_\tau \rightarrow \infty$ at fixed T), we identify $\epsilon_\tau \rightarrow \epsilon(t)$ and the exponent becomes a time integral. We thus obtain the path-integral representation:

$$\rho(t_f, \epsilon_f) = \int d\epsilon_{\text{init}} \rho(t_0, \epsilon_{\text{init}}) \int_{\epsilon(t_0)=\epsilon_{\text{init}}}^{\epsilon(t_f)=\epsilon_f} \mathcal{D}[\{\epsilon\}] \times \exp[-S[\{\epsilon\}]], \quad (\text{B7})$$

with the real action

$$S[\{\epsilon\}] = - \int_{t_0}^{t_f} dt \left[\Gamma(t, \epsilon(t)) - \frac{1}{2D} (\dot{\epsilon}(t) - v)^2 \right]. \quad (\text{B8})$$

For the numerical computation of action, we considered $N_\tau = 1000$ time steps.

APPENDIX C: LEAST ACTION

The dominant contribution to (B7) in the weak-noise limit arises from the trajectory that minimizes $S[\epsilon]$:

$$\rho(t_f, \epsilon_f) \approx N_0 \exp[-S_{\text{min}}(\epsilon_{\text{init}}, \epsilon_f)],$$

if $\rho(t_0, \epsilon) = N_0 \delta(\epsilon - \epsilon_{\text{init}})$.

We identify the integrand of the action $S[\epsilon] = \int dt \mathcal{L}$ as the effective Lagrangian:

$$\mathcal{L}(\epsilon, \dot{\epsilon}, t) = -\Gamma(t, \epsilon) + \frac{1}{2D} (\dot{\epsilon} - v)^2. \quad (\text{C1})$$

Extremizing the action yields the Euler-Lagrange equation:

$$\frac{d}{dt} \left[\frac{\partial \mathcal{L}}{\partial \dot{\epsilon}} \right] - \frac{\partial \mathcal{L}}{\partial \epsilon} = 0. \quad (\text{C2})$$

Computing the partial derivatives,

$$\frac{\partial \mathcal{L}}{\partial \dot{\epsilon}} = \frac{1}{D} (\dot{\epsilon} - v), \quad \frac{\partial \mathcal{L}}{\partial \epsilon} = -\partial_\epsilon \Gamma,$$

and substituting them back gives

$$\frac{d}{dt} \left[\frac{1}{D} (\dot{\epsilon} - v) \right] - (-\partial_\epsilon \Gamma) = 0.$$

This simplifies to the second-order ordinary differential equation:

$$\partial_\epsilon \Gamma(t, \epsilon(t)) + \frac{1}{D} \ddot{\epsilon}(t) = 0. \quad (\text{C3})$$

In our numerical experiment, the trajectory must satisfy boundary conditions:

$$\epsilon(t_0) = 0, \quad \epsilon(t_f) = \epsilon_f. \quad (\text{C4})$$

The Euler-Lagrange equations are solved using `scipy.integrate.solve_bvp` with its default relative tolerance of 0.001.

APPENDIX D: ENTROPIC CORRECTIONS

We evaluate the population density $\rho(t_f, \epsilon_f)$, when starting from the wild type $\epsilon_{\text{init}} = 0$, by integrating over all possible trajectories ϵ . While the path integral is dominated by the least-action trajectory $\hat{\epsilon}$, fluctuations around this optimal path contribute to the total probability (the entropic contribution).

We decompose a general trajectory ϵ into the optimal path $\hat{\epsilon}$ and a fluctuation vector $\Delta\epsilon$:

$$\epsilon = \hat{\epsilon} + \Delta\epsilon. \quad (\text{D1})$$

We expand the action $S[\epsilon]$ to the second order around $\hat{\epsilon}$ using a functional Taylor expansion:

$$S[\hat{\epsilon} + \Delta\epsilon] \approx S[\hat{\epsilon}] + \underbrace{\int_0^{t_f} dt \frac{\delta S}{\delta \epsilon(t)} \Big|_{\hat{\epsilon}} \Delta\epsilon(t)}_{=0} + \frac{1}{2} \int_0^{t_f} dt \int_0^{t_f} dt' \Delta\epsilon(t) H(t, t') \Delta\epsilon(t'). \quad (\text{D2})$$

The first-order term vanishes because $\hat{\epsilon}$ is a stationary point of the action (satisfying the Euler-Lagrange equations). The second-order term is governed by the Hessian matrix $\mathbf{H}$, which describes the local curvature of the fitness landscape.

The elements of the Hessian $\mathbf{H}$ matrix, see Eq. (16), can be expressed in a discrete approximation with step δt , with $N_\tau = T/\delta t$ time steps between the initial and final times. This yields a symmetric tridiagonal matrix where the diagonal components for $1 \leq \tau \leq N_\tau - 1$ are

$$\mathbf{H}_{(\tau, \tau)} = -\delta t \frac{\partial^2}{\partial \epsilon^2} \Gamma(\tau \delta t, \epsilon_\tau) + \frac{2}{D \delta t}, \quad (\text{D3})$$

and the off-diagonal components for $1 \leq \tau \leq N_\tau - 2$ are

$$\mathbf{H}_{(\tau, \tau+1)} = \mathbf{H}_{(\tau+1, \tau)} = -\frac{1}{D \delta t}. \quad (\text{D4})$$

To have a well-defined limit when $\delta t \rightarrow 0$, it is appropriate to rescale the Hessian matrix as $\mathbf{H} \rightarrow D \cdot \delta t \cdot \mathbf{H}$.

We therefore obtain

$$\rho(t_f, \epsilon_f) \propto \exp\{-S[\hat{\epsilon}]\} \times \frac{1}{\sqrt{\det \mathbf{H}}} \quad (\text{D5})$$

up to an irrelevant multiplicative factor.

In practice, this involves computing the determinant of a Hessian of size $(N_\tau - 1) \times (N_\tau - 1)$, where we choose $N_\tau = 1000$. We compute the entropic correction by building the Hessian with PyTorch (torch.autograd.functional.hessian) and extracting the log determinant using numpy.linalg.slogdet.

APPENDIX E: CONCENTRATION OPTIMIZATION

We control the time-dependent antigen concentration through its logarithm

$$c_t = \log C(t).$$

The action minimized by the least-action trajectory $\epsilon(t)$ is

$$S[\epsilon] = - \int_0^T \left[\Gamma(t, \epsilon(t)) - \frac{1}{2D} (\dot{\epsilon} - v)^2 \right] dt,$$

with growth term

$$\Gamma(t, \epsilon) = \lambda + \log P_{\text{Ag}}(\epsilon, t) + \log P_T(\epsilon, t) - \Omega(t).$$

Because c_t enters only through Γ ,

$$\frac{\delta S}{\delta c_t} = -\frac{\partial \Gamma}{\partial c_t}. \quad (\text{E1})$$

Differentiating with respect to c_t ,

$$\frac{\partial}{\partial c_t} \log P_{\text{Ag}} = 1 - P_{\text{Ag}}, \quad (\text{E2})$$

$$\frac{\partial}{\partial c_t} \log P_T = 1 - P_T, \quad (\text{E3})$$

so that

$$\frac{\partial}{\partial c_t} [\log P_{\text{Ag}} + \log P_T] = 2 - P_{\text{Ag}} - P_T. \quad (\text{E4})$$

The size constraint is enforced by

$$\Omega(t) = \begin{cases} 0, & N(t) < N_{\text{max}}, \\ \lambda + \langle \log P_{\text{Ag}} + \log P_T \rangle_{\text{pop}}, & N(t) = N_{\text{max}}. \end{cases}$$

Hence

$$\frac{\partial \Omega}{\partial c_t} = \begin{cases} 0, & N < N_{\text{max}} \text{ or} \\ \langle \log P_{\text{Ag}} + \log P_T \rangle_{\text{pop}} \leq -\lambda, & \\ 2 - \langle P_{\text{Ag}} + P_T \rangle_{\text{pop}}, & \text{otherwise.} \end{cases} \quad (\text{E5})$$

Combining (E1), (E4), and (E5),

$$\frac{\delta S}{\delta c_t} = -(2 - P_{\text{Ag}} - P_T) + \frac{\partial \Omega}{\partial c_t}, \quad (\text{E6})$$

which becomes, explicitly,

$$\frac{\delta S}{\delta c_t} = \begin{cases} P_{\text{Ag}} + P_T - 2, & N < N_{\text{max}} \text{ or} \\ \langle \log P_{\text{Ag}} + \log P_T \rangle_{\text{pop}} \leq -\lambda, & \\ P_{\text{Ag}} + P_T - \langle P_{\text{Ag}} + P_T \rangle_{\text{pop}}, & \text{otherwise.} \end{cases} \quad (\text{E7})$$

APPENDIX F: AFFINITY-DEPENDENT CELLULAR DIVISION

We consider an affinity-dependent cellular division, where the number of divisions depends on antigen captured:

$$\lambda(\epsilon) = \lambda_0 + k P_{\text{Ag}}(\epsilon), \quad (\text{F1})$$

leading to growth term

$$\Gamma(t, \epsilon) = \lambda_0 + k P_{\text{Ag}}(\epsilon, t) + \log P_{\text{Ag}}(\epsilon, t) + \log P_T(\epsilon, t) - \Omega(t).$$

Differentiating with respect to c_t , we obtain

$$\begin{aligned} \frac{\partial}{\partial c_t} [k P_{\text{Ag}} + \log P_{\text{Ag}} + \log P_T] \\ = k P_{\text{Ag}}(1 - P_{\text{Ag}}) + 2 - P_{\text{Ag}} - P_T. \end{aligned} \quad (\text{F2})$$

The size constraint is enforced by

$$\Omega(t) = \begin{cases} 0, & N(t) < N_{\text{max}}, \\ \lambda_0 + k \langle P_{\text{Ag}} \rangle_{\text{pop}} + \langle \log P_{\text{Ag}} + \log P_T \rangle_{\text{pop}}, & N(t) = N_{\text{max}}. \end{cases}$$

Hence

$$\frac{\partial \Omega}{\partial c_t} = \begin{cases} 0, & N < N_{\max} \text{ or} \\ \lambda_0 + k\langle P_{\text{Ag}} \rangle_{\text{pop}} + \langle \log P_{\text{Ag}} + \log P_T \rangle_{\text{pop}} \leq 0, & \\ k\langle P_{\text{Ag}}(1 - P_{\text{Ag}}) \rangle_{\text{pop}} + 2 - \langle P_{\text{Ag}} + P_T \rangle_{\text{pop}}, & \text{otherwise.} \end{cases} \quad (\text{F3})$$

Combining (E1), (F2), and (F3),

$$\frac{\delta S}{\delta c_t} = -(kP_{\text{Ag}}(1 - P_{\text{Ag}}) + 2 - P_{\text{Ag}} - P_T) + \frac{\partial \Omega}{\partial c_t}, \quad (\text{F4})$$

which becomes, explicitly,

$$\frac{\delta S}{\delta c_t} = \begin{cases} P_{\text{Ag}} + P_T - 2 - kP_{\text{Ag}}(1 - P_{\text{Ag}}), & N < N_{\max} \text{ or} \\ \lambda_0 + k\langle P_{\text{Ag}} \rangle_{\text{pop}} + \langle \log P_{\text{Ag}} + \log P_T \rangle_{\text{pop}} \leq 0, & \\ (P_{\text{Ag}} + P_T - kP_{\text{Ag}}(1 - P_{\text{Ag}})) & \\ - \langle P_{\text{Ag}} + P_T - kP_{\text{Ag}}(1 - P_{\text{Ag}}) \rangle_{\text{pop}}, & \text{otherwise.} \end{cases} \quad (\text{F5})$$

Similarly to before adding the affinity-dependent cell division, we define the *critical concentration*, C_0 , as the threshold where net selection and dynamic proliferation exactly counterbalance the baseline death and differentiation rates:

$$\Gamma_{\text{pop}}(C_0) = \lambda_0 + k\langle P_{\text{Ag}} \rangle_{\text{pop}} + \langle \log P_{\text{Ag}} + \log P_T \rangle_{\text{pop}} = 0. \quad (\text{F6})$$

Below this threshold ($C < C_0$), the effective growth rate Γ increases monotonically with concentration up to C_0 . Thus, the optimal protocol must maintain $C(t) \geq C_0$ to keep the population at carrying capacity $N_{\max}$. Conversely, above C_0 the pressure term $\Omega(C)$ activates to enforce the population limit.

In this regime, low-affinity B cells continue to benefit from higher antigen availability, but high-affinity clones experience a decrease in growth rate as the global population penalty Ω outweighs their saturated binding and capped proliferation rates.

-
- [1] N. S. De Silva and U. Klein, Dynamics of B cells in germinal centres, *Nat. Rev. Immunol.* **15**, 137 (2015).
- [2] G. D. Victora and M. C. Nussenzweig, Germinal centers, *Annu. Rev. Immunol.* **30**, 429 (2012).
- [3] L. Mesin, J. Ersching, and G. Victora, Germinal center B cell dynamics, *Immunity* **45**, 471 (2016).
- [4] O. Bannard and J. G. Cyster, Germinal centers: Programmed for affinity maturation and antibody diversification, *Curr. Opin. Immunol.* **45**, 21 (2017).
- [5] H. N. Eisen, Affinity enhancement of antibodies: How low-affinity antibodies produced early in immune responses are followed by high-affinity antibodies later and in memory B-cell responses, *Cancer Immunol. Res.* **2**, 381 (2014).
- [6] J. Foote and H. N. Eisen, Kinetic and affinity limits on antibodies produced during immune responses, *Proc. Natl. Acad. Sci. USA* **92**, 1254 (1995).
- [7] M. Kang, T. J. Eisen, E. A. Eisen, A. K. Chakraborty, and H. N. Eisen, Affinity inequality among serum antibodies that originate in lymphoid germinal centers, *PLoS One* **10**, e0139222 (2015).
- [8] E. A. Goidl, W. E. Paul, G. W. Siskind, and B. Benacerraf, The effect of antigen dose and time after immunization on the amount and affinity of anti-hapten antibody, *J. Immunol.* **100**, 371 (1968).
- [9] S. J. Rhodes, G. M. Knight, D. E. Kirschner, R. G. White, and T. G. Evans, Dose finding for new vaccines: The role for immunostimulation/immunodynamic modelling, *J. Theor. Biol.* **465**, 51 (2019).
- [10] M. Molari, K. Eyer, J. Baudry, S. Cocco, and R. Monasson, Quantitative modeling of the effect of antigen dosage on B-cell affinity distributions in maturing germinal centers, *eLife* **9**, e55678 (2020).
- [11] M. Molari, R. Monasson, and S. Cocco, Survival probability and size of lineages in antibody affinity maturation, *Phys. Rev. E* **103**, 052413 (2021).
- [12] S. Wang, J. Mata-Fink, B. Kriegsman, M. Hanson, D. Irvine, H. Eisen, D. Burton, K. Wittrup, M. Kardar, and A. Chakraborty, Manipulating the selection forces during affinity maturation to generate cross-reactive HIV antibodies, *Cell* **160**, 785 (2015).
- [13] J. S. Shaffer, P. L. Moore, M. Kardar, and A. K. Chakraborty, Optimal immunization cocktails can promote induction of broadly neutralizing Abs against highly mutable pathogens, *Proc. Natl. Acad. Sci. USA* **113**, E7039 (2016).
- [14] K. G. Sprenger, J. E. Louveau, P. M. Murugan, and A. K. Chakraborty, Optimizing immunization protocols to elicit broadly neutralizing antibodies, *Proc. Natl. Acad. Sci. USA* **117**, 20077 (2020).
- [15] W. S. DeWitt, A. A. Vora, T. Araki, J. G. Galloway, T. Alkutkar, J. Bortolatto, T. B. Castro, W. Dumm, C. Jennings-Shaffer, T. Jia, L. Mesin, G. Ozorowski, J. Pae, D. K. Ralph, J. D. Bloom, A. Nourmohammad, Y. S. Song, A. B. Ward, T. N. Starr, F. A. Matsen, *et al.*, Replaying germinal center evolution on a quantified affinity landscape, *Cell* **189**, 4980 (2025).
- [16] E. R. Jerison and M. M. Desai, Genomic investigations of evolutionary dynamics and epistasis in microbial evolution experiments, *Curr. Opin. Genet. Dev.* **35**, 33 (2015).
- [17] L. S. Tsimring, H. Levine, and D. A. Kessler, RNA virus evolution via a fitness-space model, *Phys. Rev. Lett.* **76**, 4440 (1996).

- [18] E. Brunet and B. Derrida, Shift in the velocity of a front due to a cutoff, *Phys. Rev. E* **56**, 2597 (1997).
- [19] A. D. Gitlin, Z. Shulman, and M. C. Nussenzweig, Clonal selection in the germinal centre by regulated proliferation and hypermutation, *Nature (London)* **509**, 637 (2014).
- [20] J. Merkenschlager, A. G. T. Pyo, G. S. Silva Santos, D. Schaefer-Babajew, M. Cipolla, H. Hartweger, A. D. Gitlin, N. S. Wingreen, and M. C. Nussenzweig, Regulated somatic hypermutation enhances antibody affinity maturation, *Nature (London)* **641**, 495 (2025).
- [21] J. Pae, N. Schwan, B. Ottino-Löffler, W. S. DeWitt, A. Garg, J. Bortolatto, A. A. Vora, J.-J. Shen, A. Hobbs, T. B. R. Castro, L. Mesin, F. A. Matsen, M. Meyer-Hermann, and G. D. Victora, Transient silencing of hypermutation preserves B cell affinity during clonal bursting, *Nature (London)* **641**, 486 (2025).
- [22] N. R. Smith and B. Meerson, Extinction of oscillating populations, *Phys. Rev. E* **93**, 032109 (2016).
- [23] P. Ashcroft, F. Michor, and T. Galla, Stochastic tunneling and metastable states during the somatic evolution of cancer, *Genetics* **199**, 1213 (2015).
- [24] I. Setliff, W. J. McDonnell, N. Raju, R. G. Bombardi, A. A. Murji, C. Scheepers, R. Ziki, C. Mynhardt, B. E. Shepherd, A. A. Mamchak, N. Garrett, S. A. Karim, S. A. Mallal, J. E. Crowe, L. Morris, and I. S. Georgiev, Multi-donor longitudinal antibody repertoire sequencing reveals the existence of public antibody clonotypes in HIV-1 infection, *Cell Host Microbe* **23**, 845 (2018).
- [25] I. Paciello, G. Pierleoni, E. Pantano, G. Antonelli, P. Pileri, G. Maccari, D. Cardamone, G. Realini, F. Perrone, M. M. Neto, S. Pozzessere, M. Fabbiani, F. Panza, I. Rancan, M. Tumbarello, F. Montagnani, D. Medini, P. Maes, N. Temperton, E. Simon-Loriere, *et al.*, Antigenic sin and multiple breakthrough infections drive converging evolution of COVID-19 neutralizing responses, *Cell Rep.* **43**, 114645 (2024).
- [26] W. Kim, J. Q. Zhou, S. C. Horvath, A. J. Schmitz, A. J. Sturtz, T. Lei, Z. Liu, E. Kalaidina, M. Thapa, W. B. Alsoussi, A. Haile, M. K. Klebert, T. Suessen, L. Parra-Rodriguez, P. A. Mudd, S. P. J. Whelan, W. D. Middleton, S. A. Teefey, I. Pusic, J. A. O'Halloran, *et al.*, Germinal centre-driven maturation of B cell response to mRNA vaccination, *Nature (London)* **604**, 141 (2022).
- [27] A. H. Ellebedy, K. J. L. Jackson, H. T. Kissick, H. I. Nakaya, C. W. Davis, K. M. Roskin, A. K. McElroy, C. M. Oshansky, R. Elbein, S. Thomas, G. M. Lyon, C. F. Spiropoulou, A. K. Mehta, P. G. Thomas, S. D. Boyd, and R. Ahmed, Defining antigen-specific plasmablast and memory B cell subsets in human blood after viral infection or vaccination, *Nat. Immunol.* **17**, 1226 (2016).
- [28] N. Spisak, G. Athènes, T. Dupic, T. Mora, and A. M. Walczak, Combining mutation and recombination statistics to infer clonal families in antibody repertoires, *eLife* **13**, e86181 (2024).
- [29] M. Z. Syeda, T. Hong, C. Huang, W. Huang, and Q. Mu, B cell memory: From generation to reactivation: A multipronged defense wall against pathogens, *Cell Death Discovery* **10**, 117 (2024).
- [30] A. Mayer, V. Balasubramanian, A. M. Walczak, and T. Mora, How a well-adapting immune system remembers, *Proc. Natl. Acad. Sci. USA* **116**, 8815 (2019).
- [31] K. M. Cirelli, D. G. Carnathan, B. Nogal, J. T. Martin, O. L. Rodriguez, A. A. Upadhyay, C. A. Enemuo, E. H. Gebru, Y. Choe, F. Viviano, C. Nakao, M. G. Pauthner, S. Reiss, C. A. Cottrell, M. L. Smith, R. Bastidas, W. Gibson, A. N. Wolabaugh, M. B. Melo, B. Cossette, *et al.*, Slow delivery immunization enhances HIV neutralizing antibody and germinal center responses via modulation of immunodominance, *Cell* **177**, 1153 (2019).
- [32] V. G. Hall, V. H. Ferreira, H. Wood, M. Ierullo, B. Majchrzak-Kita, K. Manguiat, A. Robinson, V. Kulasingam, A. Humar, and D. Kumar, Delayed-interval BNT162b2 mRNA COVID-19 vaccination enhances humoral immunity and induces robust T cell responses, *Nat. Immunol.* **23**, 380 (2022).
- [33] K. Shioda, A. Breskin, P. Harati, A. T. Chamberlain, T. Komura, B. A. Lopman, and E. T. Rogawski McQuade, Comparative effectiveness of alternative intervals between first and second doses of the mRNA COVID-19 vaccines, *Nat. Commun.* **15**, 1214 (2024).
- [34] J. M. J. Tas, L. Mesin, G. Pasqual, S. Targ, J. T. Jacobsen, Y. M. Mano, C. S. Chen, J.-C. Weill, C.-A. Reynaud, E. P. Browne, M. Meyer-Hermann, and G. D. Victora, Visualizing antibody affinity maturation in germinal centers, *Science* **351**, 1048 (2016).
- [35] S. H. Kleinstein, Y. Louzoun, and M. J. Shlomchik, Estimating hypermutation rates from clonal tree data, *J. Immunol.* **171**, 4639 (2003).
- [36] J. Zhang and E. I. Shakhnovich, Optimality of mutation and selection in germinal centers, *PLoS Comput. Biol.* **6**, e1000800 (2010).
- [37] Y. Zhang, M. Meyer-Hermann, L. A. George, M. T. Figge, M. Khan, M. Goodall, S. P. Young, A. Reynolds, F. Falciani, A. Waisman, C. A. Notley, M. R. Ehrenstein, M. Kosco-Vilbois, and K.-M. Toellner, Germinal center B cells govern their own fate via antibody feedback, *J. Exp. Med.* **210**, 457 (2013).
- [38] S. Mahdisoltani, P. Murugan, A. K. Chakraborty, and M. Kardar, Minimal framework for optimizing vaccination protocols targeting highly mutable pathogens, *Phys. Rev. E* **110**, 064137 (2024).
- [39] M. Huot, M. Molari, R. Monasson, and S. Cocco, m-huot/AFFINITY_MAT: PRX life, version v0.0, Zenodo, 2026, https://github.com/m-huot/AFFINITY_MAT.