Constrained Evolutionary Funnels Shape Viral Immune Escape

Marian Huot[1,2], Dianzhuo Wang[2,3], Eugene Shakhnovich[2,*], Rémi Monasson[1], and Simona Cocco[1,*]

[1]Laboratory of Physics of the École Normale Supérieure, CNRS UMR 8023 and PSL Research, Sorbonne Université, 24 rue Lhomond, Paris, France
[2]Department of Chemistry and Chemical Biology, Harvard University, Cambridge, MA
[3]John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA
[*]Correspondence: shakhnovich@chemistry.harvard.edu, simona.cocco@phys.ens.fr

ABSTRACT

Understanding how viral proteins adapt under immune pressure while preserving structural viability is crucial for anticipating the emergence of antibody-resistant variants. Here, we present a probabilistic framework that predicts the evolutionary trajectories of viral escape, revealing immune evasion is funneled through a remarkably small number of viable paths compared to total mutational space. These escape funnels arise from the combined constraints of protein viability and escape from antibodies, which we model using a generative model trained on structural homologs and deep mutational scanning data. We derive a mean-field approximation of evolutionary path ensembles, enabling us to quantify both the fitness and entropy of escape routes. Applied to the SARS-CoV-2 receptor binding domain, our framework reveals convergent evolution patterns, accurately predicts mutation sites in emerged variants of concern, and explains the differential effectiveness of antibody cocktails. In particular, we show that combinations of antibodies with de-correlated escape profiles slow viral adaptation by increasing the mutational effort and viability cost required for escape.

SIGNIFICANCE

Viruses must continually mutate to evade our immune defenses, yet they cannot mutate freely. Like navigating a minefield, each step toward immune escape comes at the potential cost of structural stability. This study shows that despite the astronomical number of potential mutations, viruses are funneled into narrow, predictable evolutionary paths to escape antibodies while preserving structural integrity. Using a statistical physics inspired model grounded in experimental antibody and SARS-CoV-2 epidemiology data, we model these “escape funnels” and show how they predict convergent evolution observed during the past pandemic. We also reveal why certain antibody cocktails are more resistant to viral escape: they constrain evolution to longer, less viable mutational routes.

KEYWORDS

Viral adaptation, Mutational pathways, Restricted Boltzmann machines, SARS-CoV-2, Antibody escape, Protein evolution

INTRODUCTION

Anticipating the evolutionary trajectories of SARS-CoV-2 has become pivotal to pandemic preparedness, driven by the virus's continuous adaptation under immune pressure[1,2]. Central to these adaptations is the receptor-binding domain (RBD) of the viral spike protein, critical for host cell entry and a primary target of neutralizing antibodies[3]. Immune responses elicited by prior infections and widespread vaccinations have imposed selective pressures that significantly shape viral evolution, resulting in the emergence of variants characterized by mutations that enhance transmissibility and enable antibody evasion[4–6].

Several studies[7–9] documented striking cases of convergent evolution in SARS-CoV-2, showing that distinct viral lineages can independently acquire similar mutations in the RBD. These convergent mutations, which cluster at key residues, help the virus escape neutralizing antibodies while maintaining ACE2 binding. Cao et al.[7] further demonstrated that this pattern is especially evident in recent Omicron subvariants, where immune imprinting from earlier exposures narrows the diversity of neutralizing antibodies, intensifying selective pressures and driving parallel evolutionary pathways.

Significant progress has been made in demonstrating the effects of individual viral mutations, largely through powerful experimental approaches such as deep mutational scanning (DMS)[7,10–18]. These techniques provide invaluable prospective data on how single amino acid changes in the SARS-CoV-2 spike protein or its RBD influence key viral traits, including receptor-binding affinity, viral entry, and antibody neutralization. Such insights have proven critical for viral surveillance and have even enabled the forecasting of evolutionary success for specific viral clades[19–22]. Yet, quantifying how these mutations combine and interact across entire evolutionary trajectories remains a substantially more complex and less explored challenge. While many computational approaches can now predict the fitness effects of individual mutations[23–25] or low-order combinations[26], they typically overlook the collective, high-dimensional epistatic constraints that ultimately govern which mutational paths are viable. In realistic protein landscapes, epistatic interactions emerge progressively as mutations accumulate, dynamically reshaping the accessibility of subsequent mutations and giving rise to long, emergent evolutionary timescales[27].

Recent theoretical advances, including transition path sampling[28,29] and study of navigability in high-dimensional genotype–phenotype maps[30], have begun to reveal how epistasis and evolutionary constraint jointly shape accessible evolutionary paths. Yet, a comprehensive, quantitative framework for predicting how functional and immune selection co-constrain the ensemble of escape trajectories, validated against real-world data, remains lacking.

Here, we introduce a probabilistic framework to characterize immune escape as a constrained dynamical process through sequence space. By modeling both protein viability and antibody evasion, we show that viral adaptation is funneled through a small number of viable mutational trajectories, which we term escape funnels. To uncover and quantify these escape funnels, we define escape trajectories as sequences of single-residue mutations and model their probability under joint structural and immune constraints. Structural viability is captured by Restricted Boltzmann Machines (RBMs) trained on homologous sequences, while immune escape is modeled using antibody-specific binding scores derived from deep mutational scanning. We benchmark our approach on a solvable lattice protein model and then apply it to SARS-CoV-2 RBD using experimental DMS data. Crucially, we derive a mean-field approximation of the path ensemble, enabling tractable computation of the free energy, entropy, and continuity of escape paths. This approach reveals a sharply reduced set of viable escape paths, explains convergent evolution in real-world variants, and quantifies how antibody combinations reshape the mutational landscape.

RESULTS

Escape paths

To investigate how proteins can escape immune pressure while maintaining viability, we developed a framework to sample evolutionary paths[28,29] under immune pressure (Figure 1). Each path is defined as a sequence of amino acid variants $\{v_t\}_{t=0}^{T}$, starting from the wildtype v_0, and constructed through single-residue mutations per step, such that the Hamming distance between consecutive variants satisfies $d_H(v_t, v_{t+1}) \leq 1$. The overall path probability is factorized as

$$P_{\text{path}}\big(\{v_t\}_{t=0}^{T}\big) \propto \prod_{t=0}^{T} P(v_t) \prod_{t=0}^{T-1} \mathbf{1}\{d_H(v_t, v_{t+1}) \leq 1\}\,, \tag{1}$$

where each $P(v_t)$ balances two competing pressures: the requirement to maintain protein viability (which includes stability and capacity to bind to ACE2 receptor) and the need to escape antibody binding.

To model protein viability, we learn the probability of each protein sequence $v = (v_1, \ldots, v_N)$ using a Restricted Boltzmann Machine (RBM)[31,32]. The RBM is trained on a multiple sequence alignment of protein variants and serves as a generative model for viable sequences. It captures both local constraints on individual sites (through site-specific fields g_i) and higher-order epistatic interactions (through a hidden layer). In this model, each visible unit i represents the amino-acid identity v_i at site i in the sequence, while each hidden (latent) unit μ summarizes a collective interaction pattern across sites via the input $I_\mu(v) = \sum_i w_{i\mu}(v_i)$, where $w_{i\mu}$ encodes the contribution of residue v_i at site i to hidden unit μ. The probability of observing a sequence is then:

$$P_{\text{RBM}}(v) \propto \exp\left(\sum_{i=1}^{N} g_i(v_i) + \sum_{\mu=1}^{M} \Gamma_\mu\big(I_\mu(v)\big)\right), \tag{2}$$

where the non-linear function Γ_μ is the cumulant generating function associated with hidden unit μ.

To model immune escape, we define

$$P_{\text{escape}}(v) = \prod_{\eta} \left(1 - P_{\text{bind}}^{(\eta)}(v)\right), \tag{3}$$

where $P_{\text{bind}}^{(\eta)}(v)$ denotes the binding probability of sequence v to antibody η, and the product runs over a pool of antibodies.

The full fitness model

$$P(v) \propto P_{\text{RBM}}(v) \cdot P_{\text{escape}}(v), \tag{4}$$

combines these two contributions and determines the probability of visiting each sequence along the evolutionary path.

To explore such trajectories, we extended a Monte Carlo sampling algorithm of transition paths (MCMC) previously introduced[28] (see Methods). Each path is constructed under adjacency constraints (single-residue mutations per step), ensuring evolutionary continuity while respecting the underlying sequence probability in (4).

Lattice protein model

We first benchmark this approach in a controlled setting, and apply it to lattice proteins (LP)[33,34], a simplified yet physically grounded model of protein folding (Figure 2a). LP sequences consist

of 27 amino acids folding into compact, self-avoiding walks on a $3 \times 3 \times 3$ cubic lattice, which defines the possible structures S. In this model, the ground-truth viability of a sequence V is the probability $P_{\text{nat}}(S_0|\mathbf{v})$ that it folds in the native structure S_0 rather than in one of the many competing structures (see Methods). To mimic the situation in which we do not know the ground-truth fitness, we use the RBM trained on a set of high-P_{nat} LP sequences, as a surrogate model for P_{nat}[35].

To represent the immune pressure exerted by an abstract antibody, we introduce an escape potential that penalizes similarity to the wildtype within a targeted epitope region:

$$P_{\text{bind}}(v) = e^{-(\epsilon + d_H(v, v_0))}, \tag{5}$$

where $d_H(v, v_0)$ denotes the Hamming distance between the candidate sequence v and wildtype v_0, restricted to epitope sites, and $\epsilon << 1$ enables non infinite log escape probability, see Eq. (3), even when the epitope is not mutated.

We then sample escape paths of length $T = 6$ using our MCMC algorithm, starting from an initial sequence with high P_{nat} (Figure S1). Choosing $T = 6$ allows mutating up to $\sim 20\%$ of the protein, roughly twice the mutation fraction observed in the SARS-CoV-2 RBD and comparable to the cumulative sequence divergence of influenza hemagglutinin since 1968[36]. Despite immune pressure, the sampled paths maintain high folding probabilities throughout the trajectory (Figure 2b): while individual paths exhibit some variability, the mean folding probability remains as high as wildtpe folding probability, demonstrating that the RBM constraint successfully preserves structural integrity under immune-driven evolution.

To assess mutation dynamics, we compare the per-site probability of being non-wildtype along escape paths. When sampling paths, viability is either approximated with RBM, or with the ground-truth LP model, used as a reference. At each path step, we compute for every site the probability of carrying a mutation relative to the wildtype. As shown in Figure S1, the dynamics obtained from RBM-constrained paths closely track those of the LP model, whereas paths generated under a site-independent model are less predictive for which site will mutate. This improved agreement highlights the RBM's ability to capture epistatic interactions that underlie structural dependencies between sites.

Simultaneously, escape progresses along these paths, as seen in Figure 2c: the Hamming distance from the wildtype increases steadily on the epitope, indicating mutations accumulate progressively and steadily in the epitope region, reflecting adaptive escape from antibody pressure. Interestingly, although many mutational routes are theoretically possible—given that 9 epitope residues are mutable—the final escape variants converge into a small number of distinct sequence clusters. Finally, Figure 2d shows different paths lead to similar functional outcomes in epitope region, shaped by the joint pressure to evade antibodies while maintaining protein stability.

Mean field for path sampling

A key advantage of using an RBM for sequence modeling is that its latent-unit structure admits a tractable mean-field (MF) approximation of the path ensemble, providing theoretical insight and scaling beyond sampled trajectories[28]. In the MF framework of statistical physics, each path is described via three sets of order parameters: (1) the sequence dependent inputs to each of the RBM hidden units $m_t^\mu = \frac{1}{L}\sum_i w_{i\mu}(v_{i,t})$, (2) the binding energy per antibody $\gamma_t^\eta = \frac{1}{L}\sum_i \alpha_i^\eta(v_{i,t})$ and (3) a continuity parameter along the path $q_t = \frac{1}{L}\sum_i \delta_{v_{i,t}, v_{i,t+1}}$ encoding the overlap between successive sequences of length L (see Methods).

These parameters summarize the trajectory at each time step t, and the probability of sam-

pling paths with such parameters is approximated as a Boltzmann weight

$$P_{\text{path}}(\{m_t^\mu, \gamma_t^\eta, q_t\}) \propto \exp\big[-L\beta\, f_{\text{path}}(\{m_t^\mu, \gamma_t^\eta, q_t\})\big], \tag{6}$$

where the path free energy f_{path} can be exactly computed (see Methods). Here, β plays the role of an inverse temperature, allowing us to select paths of variable quality.

To validate the framework, we first apply the MF formulation to our lattice model, with $\gamma_t = \frac{1}{L}\sum_{i\in epitope} 1_{v_{i,t}\neq v_{i,t_0}}$. Optimizing the MF path under this drive recovers the key mutated epitope sites (Figure 2e). Moreover, projecting MCMC trajectories and the MF path onto the order-parameter space (m_t^μ, γ_t) (Figure S2) shows close agreement: MF reproduces the geometry and progression of individual trajectories.

Application to SARS-CoV-2 receptor binding domain

We next apply our framework to a more realistic setting by modeling SARS-CoV-2 receptor binding domain (RBD) sequences under functional and immune constraints. Viability and binding to ACE2 are modeled using an RBM trained on 1000 synthetic homologs generated with ESM-Inverse Folding, conditioned on the RBD–ACE2 complex structure. Antibody escape is incorporated using experimental deep mutational scanning (DMS) data, inspired by Greaney et al.[37], for a fixed set of 29 antibodies identified early in the pandemic (see Methods)[16]. Although additional experimental DMS datasets became available later[7], this panel spans all four antibody classes and offers high-quality coverage with few missing single mutants. The objective is to determine whether an escape path exists: a rapid route with few mutational steps that circumvents the specified antibody set while maintaining protein viability. Unlike previous approaches[20,38], we do not explicitly model antibody concentration and apply a uniform immune pressure to all antibodies (see Methods).

We first validate the RBM's learned landscape by comparing its single-site mutational effect against experimental phenotypes. For each site i, we computed the RBM log-likelihood score for single-site mutations in the wild-type background:

$$\Delta\text{LL}_i^{\text{RBM}} = \langle \log P_{\text{RBM}}(v_i, \{wt_{j\neq i}\})\rangle_{v_i} - \log P_{\text{RBM}}(wt_i, \{wt_{j\neq i}\}), \tag{7}$$

where $\langle\cdot\rangle_{v_i}$ denotes the average over all amino-acid substitutions v_i at position i. This score reflects the expected effect of mutations at site i on protein viability, with positive values indicating mutations more favorable than the wild-type residue[39]. As shown in Figure S3, these RBM-derived scores correlate well with DMS measurements of both protein expression and ACE2 binding, indicating that the model accurately captures the viability landscape of the RBD. The model achieves a Spearman correlation of 0.60 for expression data—comparable to the DCA model of Rodriguez-Rivas *et al.*[26] and the RBM of Huot *et al.*[38] (both trained on SARS-CoV homologous sequences). It further reaches a Spearman correlation of 0.69 for ACE2 binding, a performance not reported in these earlier studies. This increase in predictive power is likely due to the use of synthetic homologs generated with ESM-Inverse Folding, which are structurally conditioned and more likely to retain ACE2 binding, in contrast to far SARS-CoV-2 homologs which were likely not subject to the ACE2 binding constraint. The RBM retains most of the predictive power for expression, compared to the original ESM-Inverse Folding model, while remaining equally effective at predicting binding.

We then consider escape paths of length $T = 20$, corresponding to at most 20 mutations along a trajectory, consistent with the mutational load observed in circulating variants. A key advantage of MF is that it enables the quantification of the number of viable immune escape paths through the computation of a path entropy (see Methods). This entropy reflects the diversity

of trajectories consistent with both structural and immune constraints, and is reduced not only by limiting the number of mutations per step, but also by the functional constraints encoded in the RBM, which restrict the set of viable trajectories. In practice, the number of escape paths shrinks from $Q^{LT} \sim 10^{4631}$ in the unconstrained model to $e^{LS_{\text{path}}} \sim 10^{21}$. Notably, the continuity constraint of one mutation per step reduces entropy far more strongly than the RBM fields alone (site-independent model).

Constraints	none	site-independent field	1 mutation per step	full model
Nb of paths	10^{4631}	10^{1985}	10^{71}	10^{21}

Table 1: Number of paths under different model constraints.

A small number of viable escape paths suggests that only a limited subset of mutations can enable immune evasion without compromising protein foldability. To explore the size of this subset of mutations, we project MF-derived site entropies (see Methods) onto the 3D structure of the RBD–ACE2 complex (Figure 3a), revealing that highly mutable positions (in red) are surface-exposed and often overlap with residues mutated in known variants of concern (VOCs). Notably, top predictions include known escape hotspots such as T376, N440, N460 and Q498 mutated in most of VOCs, further supporting the biological relevance of the model.

Altogether, these results demonstrate that our MF approximation identifies key mutational targets and offers an interpretable and efficient way to characterize dominant escape trajectories. We then projected the MCMC-sampled trajectories into antigenic space, defined by the top two principal components (PCs) of antibody binding scores $\{-\sum_i \alpha_i^\eta(v_i)\}_\eta$ of observed sequences in pandemic (Figure 3b). Observed SARS-CoV-2 variants were colored by date of emergence. This allows us to directly compare model-predicted antigenic drift with real-world viral evolution. Notably, MCMC trajectories align with the direction of antigenic evolution observed in circulating variants, progressing from wildtype toward major variants of concern (VOCs) such as BA.5, BQ.1.1, and XBB. Later variants appear further along the model-predicted escape path, highlighting the model's ability to predict the temporal and directional structure of antigenic drift.

We then compare computed antibody binding scores across all 29 antibodies at first and last steps along the trajectories (Figure 3c and Figure S5). At each time point, the MF predictions show strong correlation with those from MCMC, with agreement improving as escape progresses. This indicates that the MF approximation successfully captures the global escape dynamics. In particular, the MF is highly consistent with stochastic sampling at later steps, where immune pressure drives more pronounced divergence from the wildtype. Notably, class 1 and 2 antibodies tend to lose binding earliest, reflecting their greater vulnerability to initial escape mutations.

To test whether these predicted trajectories correspond to actual antibody escape, we computed amino acid probabilities at the final-step to deduce antibody binding scores for 438 antibodies from Cao et al.[7] (Figure 3d) and compared them with experimental IC50 measurements. We find that the model reliably identifies ineffective antibodies, with area under the ROC curve (AUC) values of 0.58 for BA.5, 0.67 for BQ.1.1, and 0.73 for XBB—demonstrating strong predictive power, especially for the more immune-evasive variants. These results show that our framework not only reproduces past antigenic evolution but also has the potential to prospectively forecast antibody failure, without requiring any a priori knowledge of which mutations will arise. Finally, our framework is not limited to escape paths originating from the wildtype. When initialized from BA.1, the MF trajectory also reaches an antigenic profile consistent with XBB, showing that convergent escape routes can also emerge from already mutated variants (Figure 3e).

Effect of cocktails

Our MF framework enables us to quantify how different antibody combinations shape viral escape dynamics and to understand why some cocktails offer greater short-term protection (evaluated here over a trajectory of T = 9 steps). We examined the evolution of the escape probability P_{escape} along the MF trajectory under four different conditions: selection by individual antibodies (COV2-2196 or REGN10987) in Figures 4a–b and selection by cocktails pairing each with REGN10933 in Figures 4c–d. For each case, we identified the first point along the trajectory where the fraction of sequences with $P_{\text{escape}} \geq 50\%$ is reached, providing a quantitative measure of escape speed.

The results show that the REGN10933 + REGN10987 cocktail substantially delays escape compared to REGN10987 alone, while the COV2-2196 + REGN10933 cocktail provides only marginal additional protection over COV2-2196. This suggests that not all cocktails are equally effective at slowing escape.

To understand these differences, we examined the correlation between escape profiles of REGN10933 with other antibodies discovered early in pandemic[16], measured as the covariance of their site-wise binding effects $\alpha_i^\eta(v_i)$ weighted by P_{RBM} covariation[38]. This restricts the analysis to mutations that preserve protein viability. Antibody pairs with positively correlated escape profiles tend to be evaded by the same mutations, whereas uncorrelated profiles require distinct and often compensatory mutations. As shown in Figure 4e, the number of additional steps required to reach 50% escape when REGN10933 is added to a cocktail increases when covariance between the two antibodies decreases (Pearson r = -0.75). Notably, combining multiple correlated antibodies in a cocktail only modestly increased the number of steps required for escape (Figure S6). Finally, these results generalize beyond cocktails involving REGN10933, suggesting broader principles for effective antibody combinations. Notably, fast escape primarily occurs when combining Class 1 and Class 2 antibodies (Figure S7). These findings agree with a previous study[38] that demonstrated how decorrelated cocktails lead to a lower number of viable escape mutants. Our work extends this by showing that these decorrelated cocktails are harder to escape because they force the virus to follow longer, more constrained evolutionary paths through sequence space.

We next translate this increased “trajectory length” into explicit fitness costs, as evaluated by the RBM score. Figure 4f shows the additional energy ΔE_{RBM} along each trajectory. Notably, individual antibodies REGN10933, COV2-2196 and REGN10987 alone require few extra loss of viability because they target variable sites where naturally occurring mutations are often sufficient to reduce binding by 50%. In contrast, antibody combinations require each component antibody to reach a higher escape level before the RBD escapes the cocktail, leading to a mutational cost greater than the sum for the individual antibodies. This effect is especially pronounced for decorrelated antibodies, where about two full additional mutations (relative to natural evolution) are needed to achieve 50% escape. Together, these results illustrate how decorrelated cocktails not only extend the evolutionary path but also force the virus to accumulate changes that compromise structural stability.

DISCUSSION

Understanding how viruses adapt under immune pressure while preserving protein function is critical for anticipating antigenic evolution and designing therapies that remain effective over time. Here, we present a framework that captures the constrained nature of viral escape from neutralizing antibodies, by modeling evolution as a path through sequence space shaped by both viability constraints and immune selection. Compared to the astronomical combinatorial size of protein sequence space, we find that viable escape is funneled through a remarkably narrow set of mutational trajectories that we term "escape funnels". These trajectories are not only constrained by viability, but also exhibit convergence consistent with real-world viral evolution, such as that observed in SARS-CoV-2 variants of concern[7].

A core advance of our approach lies in combining generative sequence modeling with a statistical physics treatment of evolutionary paths[28]. We leverage RBMs to capture complex epistatic patterns learned from protein sequences, including both local and higher-order dependencies that are critical for maintaining structural integrity. The RBM accurately recovers folding properties in benchmark lattice models[35] and enables distillation of sequence ensembles generated by powerful models like ESM-Inverse Folding[40] into a simpler, tractable representation. We then apply a mean-field approximation to the ensemble of evolutionary trajectories, allowing us to analytically characterize dominant paths while quantifying key properties such as path entropy, and fitness cost.

This framework yields several key biological insights. First, we show that viability and immune constraints together lead to strongly convergent evolution, where escape trajectories rapidly collapse onto a limited number of mutational solutions. This finding is validated in both lattice protein models and the SARS-CoV-2 RBD, where we observe that mean-field generated trajectories recapitulate mutational patterns found in natural sequences. Notably, sites predicted to exhibit high mutational entropy under immune pressure correspond closely to those that have repeatedly mutated in real-world variants such as BA.5 and XBB.

Second, our model enables prospective forecasting of antibody escape from any variant and, crucially, incorporates evolutionary time, unlike previous approaches[22,41] focused on viral variants generation that tested antibody effectiveness on mutants far away from wildtype. By projecting trajectories into antigenic space, we find that model-predicted paths align with the direction of observed antigenic drift in SARS-CoV-2 evolution. Furthermore, the predicted antibody binding loss correlates with experimental measurements of antibody resistance across a panel of 438 antibodies, demonstrating the utility of the model for identifying at-risk therapeutics in advance.

Third, we provide a mechanistic explanation for the effectiveness of antibody cocktails. We show that cocktails composed of antibodies with anti-correlated escape profiles force the virus to traverse longer and more structurally costly paths to achieve immune evasion. These findings are not only consistent with previous empirical results[42–44], but also yield a quantitative framework for optimizing cocktail combinations. In particular, we find that the REGN10933 + REGN10987 cocktail imposes both the greatest mutational burden and the largest stability penalty, providing a biophysical explanation for its superior durability relative to other combinations.

While our framework captures structural and immune constraints that define escape funnels, it deliberately omits the temporal dynamics of immune responses and treats the antibody landscape as fixed. This abstraction is efficient for forecasting which mutations enable fastest and viable escape, given a known set of antibodies, but it cannot account for shifting repertoires or the delayed emergence of certain antibody classes after boosting[45].

More broadly, our framework is not virus-specific. The use of generative models like RBMs and the statistical mechanics of constrained path ensembles make this approach generalizable to other rapidly evolving systems. As a matter of fact, RBMs and similar probabilistic models

have already been used to successfully capture epistatic interactions in HIV[46] and influenza hemagglutinin[47] evolution.

METHODS

Fitness model of lattice sequences

The folding probability of a sequence $\mathbf{v}$ into a given structure S is determined by the energetic contributions of residue-residue interactions within the folded configuration[33,48]. These interactions occur between amino acid pairs that are spatially adjacent in the structure but not sequentially contiguous. The geometry of the structure is encoded by a contact map c_{ij}^S, where $c_{ij}^S = 1$ indicates that residues i and j are in contact in structure S, and $c_{ij}^S = 0$ otherwise.

The total energy of a sequence $\mathbf{v}$ in a particular structure S is computed as:

$$\mathcal{E}_{\text{LP}}(\mathbf{v}|S) = \sum_{i<j} c_{ij}^S \, E_{\text{MJ}}(v_i, v_j) \tag{8}$$

where $E_{\text{MJ}}(v_i, v_j)$ denotes the Miyazawa-Jernigan interaction energy[49] between amino acids v_i and v_j, reflecting their physico-chemical compatibility.

The probability that sequence $\mathbf{v}$ adopts structure S is then given by a Boltzmann distribution:

$$P_{\text{nat}}(S|\mathbf{v}) = \frac{\exp\left[-\mathcal{E}_{\text{LP}}(\mathbf{v}|S)\right]}{\sum_{S'} \exp\left[-\mathcal{E}_{\text{LP}}(\mathbf{v}|S')\right]} \tag{9}$$

where the denominator sums over all possible compact, self-avoiding conformations S' in the lattice. In this formulation, P_{nat} represents the ground-truth thermodynamic probability that sequence $\mathbf{v}$ adopts a given structure, based on its folding energy $\mathcal{E}_{\text{LP}}$. However, computing this distribution exactly requires summing over all compact, self-avoiding conformations—a combinatorially expensive task for realistic protein spaces.

In our work, we therefore approximate P_{nat} with P_{RBM} coming from an RBM trained on sequences known to fold into the desired structure, following the procedure of Jacquin et al.[35]. Precisely, these sequences were sampled from a low temperature MC sampling using $-\beta \log P_{nat}(|S_{nat})$ with $\beta = 10^3$ as effective energy.

To account for antibody escape, penalize low Hamming distance d_{H} to wildtype v_0 over a fixed epitope subset:

$$d_{\text{H}}(\mathbf{v}, \mathbf{v}_0) = \sum_{i \in \mathcal{E}_{\text{epitope}}} \mathbf{1}_{\{v_i \neq v_{0,i}\}}, \quad \text{where} \quad \mathcal{E}_{\text{epitope}} = \{9, 10, 11, 12, 13, 16, 17, 25, 26\}$$

To combine structural stability with immune escape, we define the fitness of a sequence $\mathbf{v}$ as the product of its native folding probability and the antibody escape factor:$P(\mathbf{v}) = P_{\text{RBM}}(\mathbf{v}) \cdot P_{\text{escape}}(\mathbf{v})$

ESMIF distillation

For the SARS-CoV-2 RBD, we generated an artificial multiple sequence alignment (MSA) of 1,000 sequences using ESM-Inverse Folding at temperature 1, conditioned on the full RBD–ACE2 complex structure from PDB entry 6M0J. Conditioning on ACE2 binding ensures that the sampled sequences are compatible with receptor interaction, in contrast to prepandemic coronaviridae or animal SARS-related sequences that may diverge functionally. The RBM was then trained using the persistent contrastive divergence algorithm to approximate this distribution, thereby capturing structural constraints relevant for ACE2 recognition.

Antibody binding

To model effect of single mutations on binding to RBD, we use deep mutational scans : data provided in bloomlab repository, that integrates data from several previous studies[10–16], as well as data taken from work of Cao et al.[7].

We then define the binding energy contribution of amino acid v_i at site i to antibody η : $\alpha_i^\eta := -\log(1-\epsilon_i^\eta)$, where

$$\epsilon_i^\eta = \begin{cases} 0.001 \text{ if } v_i = wt_i \text{ or } v_i \text{ not in DMS} \\ \text{escape ratio of mutant } v_i \text{ otherwise} \end{cases} .$$

We assume that all wildtype (wt) amino acids, as well as single mutants missing from DMS, do not provide escape ($\epsilon_i^\eta = 0.001$). Value of 0.001 was taken as it was close to minimum measurable escape ratio among all single variants in DMS. Note that since $0 \leq \epsilon < 1$, binding energy contributions $\alpha_i^\eta \geq 0$.

Following Greaney et al.[37], we then define

$$P_{\text{bind}}^{(\eta)}(\mathbf{v}) = \exp\left(-\sum_i \alpha_i^\eta(v_i)\right) .$$

In other words, antibody η binds to RBD only if all mutations maintain binding.

In previous studies[20,38], binding probability to an antibody with concentration C and dissociation constant K_d was typically modeled as $P_{\text{bind}}(\mathbf{v}) = C/(C + K_d(\mathbf{v}))$. In the low-concentration regime ($C \ll K_d(\mathbf{v})$), $P_{\text{bind}}(\mathbf{v}) \approx C/K_d(\mathbf{v})$, making the binding energy $\sum_i \alpha_i^\eta(v_i)$ equal to $\log K_d(\mathbf{v})$ used in previous approaches (up to an additive constant).

Pandemic sequences

All 15,371,428 spike sequences on GISAID[50] as of 14-April-2023 were downloaded and aligned, following approach in the work of Starr et al.[15]. Sequences from non-human origins and with lengths outside $[1260, 1276]$ were removed. They were then aligned via mafft[51] and sequences containing unicode errors, gap or ambiguous characters were removed. To avoid regional bias, we filtered to only keep sequences observed in the USA. Overall, we retained 3,712,556 submissions represented by 10,453 unique RBD sequences. RBD amino-acid mutations were enumerated compared to the reference Wuhan-Hu-1 SARS-CoV-2 RBD sequence (Genbank MN908947, residues N331-T531).

To align with the MSA data, the model was finally applied to sequences of length 178, spanning residues S349 to G526, leading to 8,167 unique RBD sequences.

Path sampling with MCMC

We sample mutation paths under the target distribution P_{path}. Paths start from a fixed wildtype v_0 and are updated by single–mutation Metropolis–Hastings moves. Proposed updates depend on the Hamming distance between neighboring states: if neighbors coincide, a new mutation is drawn randomly; if they differ by one residue, the update is restricted to that site; and if they differ by two residues, the update flips to the only compatible intermediate. Each proposal is accepted with probability ensuring detailed balance, so that the algorithm converges to the target distribution P_{path} (see SI, MCMC algorithm for path sampling).

Mean field

We extend the framework of Mauri et al.[28,29], which formulated mean–field dynamics of mutational paths under folding constraints, in two key ways. First, we incorporate antibody escape by introducing binding–score order parameters that act in concurrence with folding pressure, thereby capturing the trade–off between structural stability and immune evasion. Second, we free the end of the trajectories: unlike transition–path formulations that assume a fixed evolutionary destination, our paths are open-ended and evolve solely under the joint action of folding and immune constraints.

Sequence distribution is given by the RBM with fields g_i and hidden unit cumulant generating function Γ_μ (see SI, RBM-based sequence modeling):

$$P_{\text{RBM}}(v) = \exp\left[\sum_{i=1}^{N} g_i\left(v_i\right) + \sum_{\mu=1}^{M} \Gamma_\mu\left(Lm_\mu\right)\right] \tag{10}$$

Antibody escape is given by the product of the probability to escape each antibody:

$$P_{\text{escape}}(v) = \prod_{\eta}\left(1 - \exp(-L\gamma_\eta)\right) \tag{11}$$

Sequence probability is a combination of stability and antibody escape:

$$P(v) \propto \exp\left[\sum_{i=1}^{N} g_i\left(v_i\right) - L\zeta(m_\mu, \gamma_\eta)\right] \tag{12}$$

with the potential

$$\zeta(m_\mu, \gamma_\eta) = -\frac{1}{L}\left[\sum_{\mu} \Gamma_\mu\left(Lm_\mu\right) + \sum_{\eta} \log\left(1 - \exp(-L\gamma_\eta)\right)\right] \tag{13}$$

The potential Φ should forbid large jumps along the paths. We thus consider a hard-wall repulsive potential,

$$\Phi(q) = \begin{cases} \dfrac{1}{L} \cdot \dfrac{1/T^2}{q - q_c} & \text{if } q_c < q \leq 1, \\ +\infty & \text{otherwise} \end{cases} \tag{14}$$

The location of the hard wall, $q_c = 1 - 1/L$, allows the path to explore at most $T \times L(1 - q_c) = T$ mutations in T steps.

The free energy is given by:

$$f_{\text{path}} = \sum_{t} \zeta(m_t, \gamma_t) + \sum_{t} \Phi(q_t) - \frac{1}{\beta} S_\beta \tag{15}$$

The entropy term S_β captures the diversity of sequence realizations with the same order parameters. The mean-field trajectory is obtained by minimizing f_{path} with respect to the order parameters, subject to initial conditions fixed at one given sequence, such as the wildtype (see SI, Mean field).

Site entropy

To quantify the mutational variability at each site i of the RBD, we compute the Shannon entropy:

$$S_i = -\sum_a P_i(a) \log P_i(a), \tag{16}$$

where $P_i(a)$ is the marginal probability of observing amino acid a at position i. In the mean–field approximation, these marginals follow from the derivative of the partition function with respect to the local fields (see SI, Amino acid probability in mean field), yielding

$$P_i(a) = \left\langle \delta_{v_{i,t},a} \right\rangle_{\text{path}} = \frac{\partial \ln Z}{\partial \beta g_i(a)}. \tag{17}$$

High entropy indicates that the site tolerates diverse mutations, while low entropy reflects evolutionary constraints or functional importance.

Path Entropy

To build intuition, we consider limiting cases for the path entropy normalized by the protein length L. When paths are constrained to a single mutation per step and no external fields are applied, the entropy is given by $S_{\text{path}} = \frac{T}{L} \log[L(Q-1)]$. In the case where continuity is removed and no field is present, the entropy becomes $S_{\text{path}} = T \log Q$. If the field is site-independent and continuity is still removed, the entropy takes the form $S_{\text{path}} = T \cdot \langle S_{\text{site}} \rangle$, where $\langle S_{\text{site}} \rangle$ denotes the average entropy across sites induced by the field. Using the full model, path entropy (see SI, Path Entropy) can be computed as:

$$S_{\text{path}} = -\frac{df_{\text{path}}}{d(1/\beta)} = S_\beta - \beta \frac{\partial S_\beta}{\partial \beta} \tag{18}$$

Correlation of antibody escape profiles

To quantify the similarity between antibody escape profiles, we computed the covariance between the binding scores of antibodies identified early in the pandemic[16]. For a given antibody η, the binding score, defined as

$$-L\gamma^\eta(v) = -\sum_i \alpha_i^\eta(v_i), \tag{19}$$

was evaluated over sequences v sampled from P_{RBM}, thereby biasing the analysis toward structurally viable variants.

To quantify antibody synergy, we compute the covariance between dissociation constants associated with antibodies 1 and 2[38]:

$$\text{Cov}(Ab_1, Ab_2) = \sum_{\{v_i, v_j\}} \alpha^1(v_i)\ \alpha^2(v_j) \text{Cov}(v_i, v_j), \tag{20}$$

where Cov represents the covariance of amino-acid occurrences when sampling sequences from the distribution P_{RBM}. Sampling sequences from the RBM distribution ensures that escape profile correlations capture mutational accessibility under stability constraints, rather than relying solely on epitope comparison as done in previous studies[14,37].

RESOURCE AVAILABILITY

Lead contact

Requests for further information and resources should be directed to and will be fulfilled by the lead contact, Simona Cocco (simona.cocco@phys.ens.fr).

Data and code availability

The code used in this study will be available at `https://github.com/m-huot/ESCAPE_PATHS` after manuscript acceptance.

ACKNOWLEDGMENTS

SC is grateful to G. Parisi for having suggested the application of transition paths to SARS-Cov2 evolution. We acknowledge funding from the Agence Nationale de la Recherche (ANR ProDiGen, AAP2024 CE45 to S.C. and R.M.). This work is supported by NIH R35GM139571. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

We gratefully acknowledge all data contributors, i.e., the Authors and their Originating laboratories responsible for obtaining the specimens, and their Submitting laboratories for generating the genetic sequence and metadata and sharing via the GISAID Initiative, on which this research is based.

AUTHOR CONTRIBUTIONS

M.H., D.W., E.S., R.M., and S.C. conceived the project, supervised the research, and wrote the manuscript. M.H. coded the computational model and performed data analysis.

DECLARATION OF INTERESTS

The authors declare no competing interests.

SUPPLEMENTAL INFORMATION INDEX

Supplementary methods

Supplementary figures and legends, included after main figures.

Main Figures and Tables

481

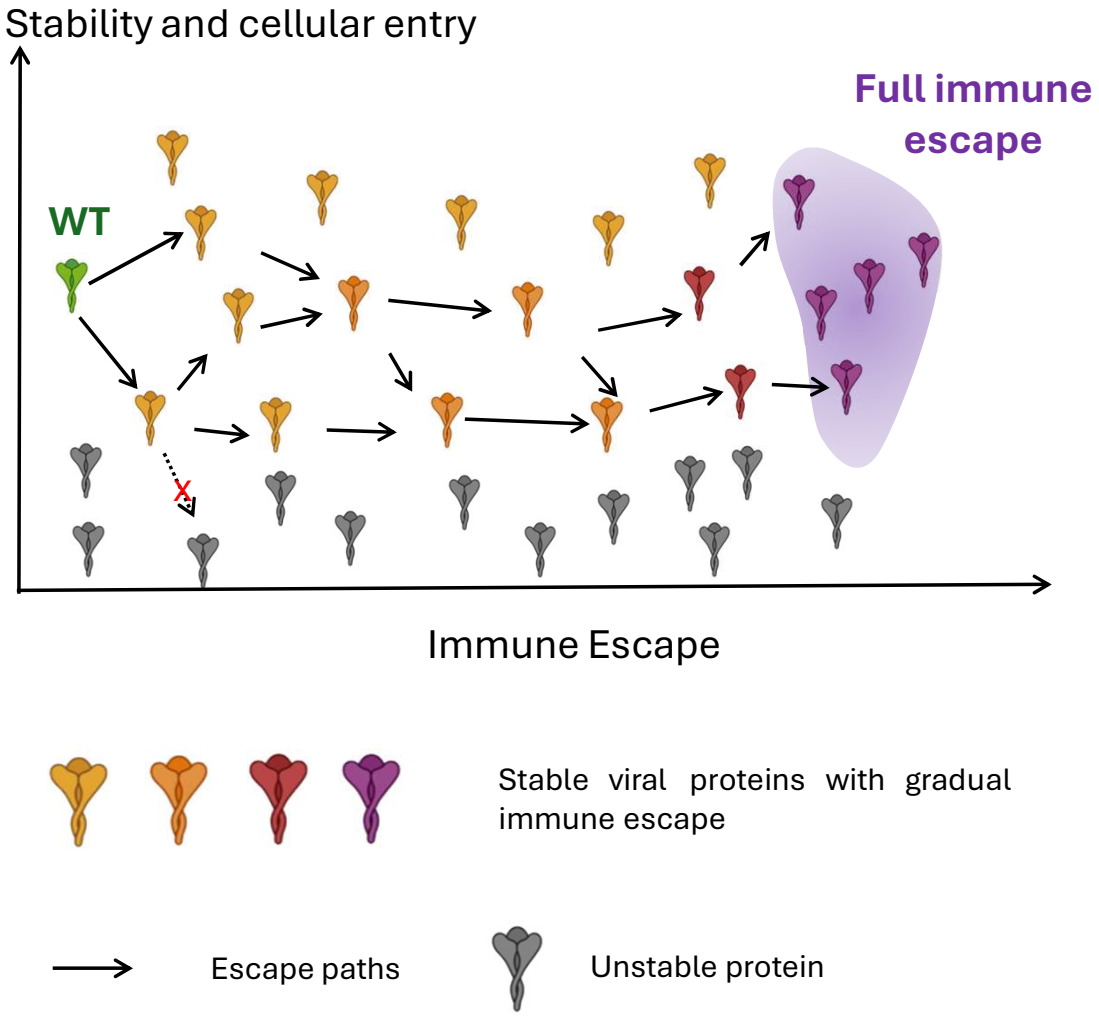

Fig. 1: Evolutionary paths toward antibody escape. Schematic representation of plausible evolutionary trajectories from the wildtype to antibody-escaping variants. Paths progress through a constrained sequence space, funneling through a small set of viable intermediates that gradually accumulate immune escape mutations.

Figure 1

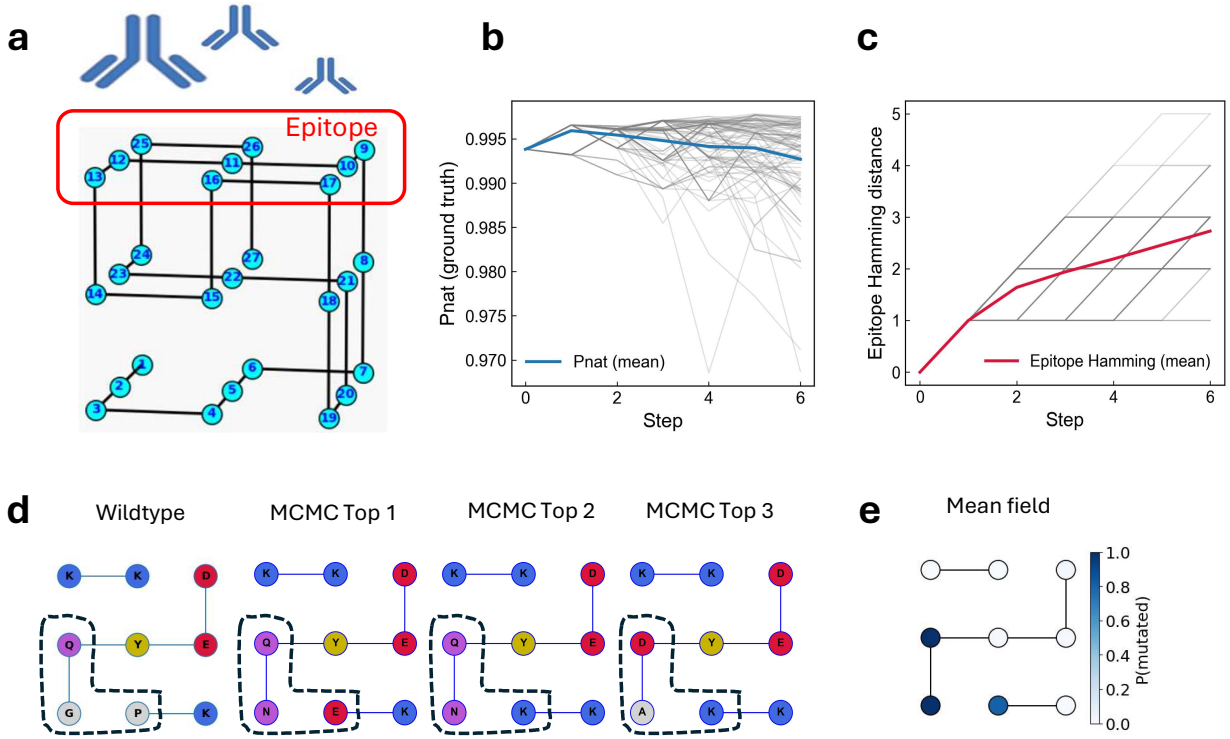

Fig. 2: Folding constraints shape convergent escape routes under immune pressure on lattice protein model. (a) Native fold structure S_{nat}, with the antibody-targeted epitope (upper face) highlighted in red. During escape, the protein must accumulate epitope mutations while preserving high fold probability. **(b)** Evolution of folding probability P_{nat} along sampled escape paths. The blue line indicates the mean across paths; grey lines show individual trajectories. **(c)** Accumulation of epitope mutations along escape paths, measured as Hamming distance from the wildtype. Each grey line represents a single path. Mutations accumulate steadily across the epitope, indicating progressive immune escape. **(d)** Wild-type epitope and three representative epitope variants at the terminal step of MCMC-sampled escape paths; amino-acid colors denote biochemical classes. Escape routes consistently converge on a small set of mutations with similar biochemical properties. **(e)** Site-wise probability of being mutated at the final step of the mean-field trajectory.

Figure 2

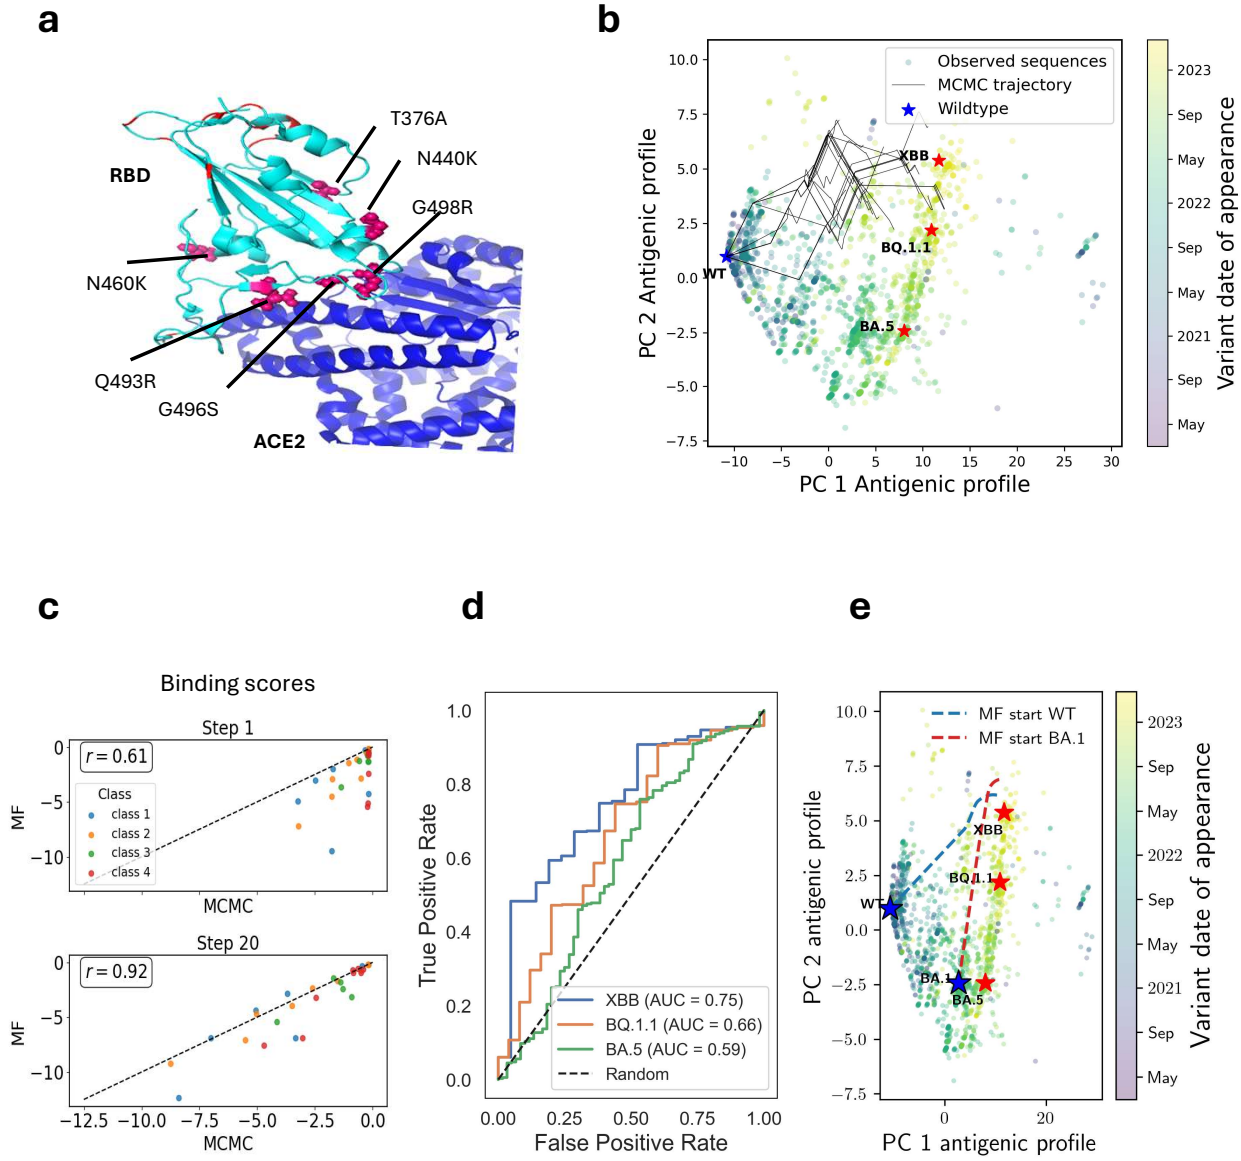

Fig. 3: Escape paths capture mutational and antigenic features of SARS-CoV-2 variants. (a) 3D structure of the RBD (light blue) bound to ACE2 (dark blue). Residues in red are in the top 10% of predicted entropy in the final step of the mean field trajectory. Pink spheres mark the top 10% predicted sites that were also mutated in variants of concern (VOC), labeled with the VOC-observed mutation. **(b)** MCMC trajectories projected onto the top two principal components of antibody binding coefficients (antigenic space). SARS-CoV-2 variants observed during the pandemic are plotted and colored by their date of appearance. Key variants of concern (VOCs) such as Wildtype (WT), BA.5, BQ.1.1, and XBB are annotated. **(c)** Comparison of antibody binding scores (across 29 antibodies, colored by class) between MCMC-sampled paths and the mean-field (MF) trajectory. The MF trajectory progressively aligns with MCMC results as escape accumulate with class 1 and 2 antibodies losing binding earliest. **(d)** ROC curves showing the predictive performance of the mean-field trajectory in identifying antibodies escaped by each of three major SARS-CoV-2 variants of concern: BA.5, BQ.1.1, and XBB. Predictions are based on the final-step values of antibody binding order parameters computed along the MF path. Experimental escape data—measured as IC50 values for 438 antibodies from Cao et al.—serve as the ground truth. **(e)** MF trajectories starting from WT and BA.1 variants, projected onto the top two principal components of antibody binding coefficients (antigenic space). SARS-CoV-2 variants observed during the pandemic are plotted and colored by their date of appearance.

Figure 3

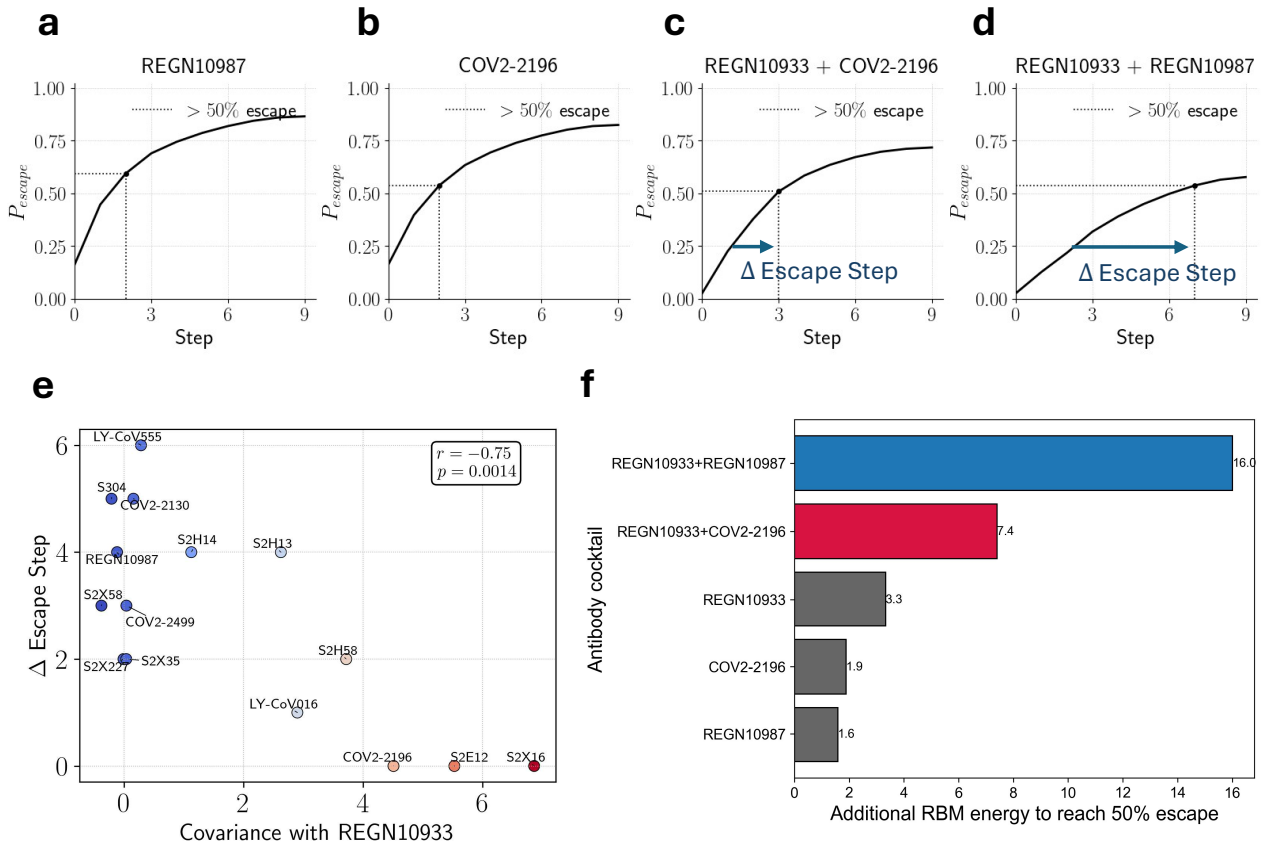

Fig. 4: Mean field trajectories reveal resistance of antibody cocktails to viral escape. **(a–d)** Mean field trajectories showing antibody escape under pressure from individual antibodies (COV2-2196, REGN10987) and cocktails with REGN10933. The vertical dotted line indicates the step at which at least 50% escape is achieved. **(e)** Relationship between antibody binding covariance with REGN10933 and additional steps required to reach 50% escape when the antibody is used in a cocktail with REGN10933. **(f)** Additional RBM energy, compared to evolution without immune pressure, to reach 50% escape level averaged over mean field path, for different antibody combinations. Additional RBM energy can be interpreted as protein loss of viability.

Figure 4

References

1. Raharinirina, N.A., Gubela, N., Börnigen, D., Smith, M.R., Oh, D.Y., Budt, M., Mache, C., Schillings, C., Fuchs, S., Dürrwald, R., Wolff, T., Hölzer, M., Paraskevopoulou, S., and Von Kleist, M. (2025). SARS-CoV-2 evolution on a dynamic immune landscape. Nature *639*, 196–204. URL: https://www.nature.com/articles/s41586-024-08477-8. doi: 10.1038/s41586-024-08477-8.

2. Meijers, M., Ruchnewitz, D., Eberhardt, J., Łuksza, M., and Lässig, M. (2023). Population immunity predicts evolutionary trajectories of SARS-CoV-2. Cell *186*, 5151–5164.e13. URL: https://linkinghub.elsevier.com/retrieve/pii/S0092867423010760. doi: 10.1016/j.cell.2023.09.022.

3. Jackson, C.B., Farzan, M., Chen, B., and Choe, H. (2022). Mechanisms of SARS-CoV-2 entry into cells. Nature Reviews Molecular Cell Biology *23*, 3–20. URL: https://www.nature.com/articles/s41580-021-00418-x. doi: 10.1038/s41580-021-00418-x.

4. Carabelli, A.M., Peacock, T.P., Thorne, L.G., Harvey, W.T., Hughes, J., COVID-19 Genomics UK Consortium, De Silva, T.I., Peacock, S.J., Barclay, W.S., De Silva, T.I., Towers, G.J., and Robertson, D.L. (2023). SARS-CoV-2 variant biology: immune escape, transmission and fitness. Nature Reviews Microbiology. URL: https://www.nature.com/articles/s41579-022-00841-7. doi: 10.1038/s41579-022-00841-7.

5. Nabel, K.G., Clark, S.A., Shankar, S., Pan, J., Clark, L.E., Yang, P., Coscia, A., McKay, L.G.A., Varnum, H.H., Brusic, V., Tolan, N.V., Zhou, G., Desjardins, M., Turbett, S.E., Kanjilal, S., Sherman, A.C., Dighe, A., LaRocque, R.C., Ryan, E.T., Tylek, C., Cohen-Solal, J.F., Darcy, A.T., Tavella, D., Clabbers, A., Fan, Y., Griffiths, A., Correia, I.R., Seagal, J., Baden, L.R., Charles, R.C., and Abraham, J. (2022). Structural basis for continued antibody evasion by the SARS-CoV-2 receptor binding domain. Science *375*, eabl6251. URL: https://www.science.org/doi/10.1126/science.abl6251. doi: 10.1126/science.abl6251.

6. Wang, Q., Iketani, S., Li, Z., Liu, L., Guo, Y., Huang, Y., Bowen, A.D., Liu, M., Wang, M., Yu, J., Valdez, R., Lauring, A.S., Sheng, Z., Wang, H.H., Gordon, A., Liu, L., and Ho, D.D. (2023). Alarming antibody evasion properties of rising SARS-CoV-2 BQ and XBB subvariants. Cell *186*, 279–286.e8. URL: https://linkinghub.elsevier.com/retrieve/pii/S0092867422015318. doi: 10.1016/j.cell.2022.12.018.

7. Cao, Y., Jian, F., Wang, J., Yu, Y., Song, W., Yisimayi, A., Wang, J., An, R., Chen, X., Zhang, N., Wang, Y., Wang, P., Zhao, L., Sun, H., Yu, L., Yang, S., Niu, X., Xiao, T., Gu, Q., Shao, F., Hao, X., Xu, Y., Jin, R., Shen, Z., Wang, Y., and Xie, X.S. (2022). Imprinted SARS-CoV-2 humoral immunity induces convergent Omicron RBD evolution. Nature. URL: https://www.nature.com/articles/s41586-022-05644-7. doi: 10.1038/s41586-022-05644-7.

8. Feng, S., Reid, G.E., Clark, N.M., Harrington, A., Uprichard, S.L., and Baker, S.C. (2024). Evidence of SARS-CoV-2 convergent evolution in immunosuppressed patients treated with antiviral therapies. Virology Journal *21*, 105. URL: https://virologyj.biomedcentral.com/articles/10.1186/s12985-024-02378-y. doi: 10.1186/s12985-024-02378-y.

9. Jian, F., Feng, L., Yang, S., Yu, Y., Wang, L., Song, W., Yisimayi, A., Chen, X., Xu, Y., Wang, P., Yu, L., Wang, J., Liu, L., Niu, X., Wang, J., Xiao, T., An, R., Wang, Y., Gu, Q., Shao, F., Jin, R., Shen, Z., Wang, Y., Wang, X., and Cao, Y. (2023). Convergent evolution of SARS-CoV-2 XBB lineages on receptor-binding domain 455–456 synergistically enhances

antibody evasion and ACE2 binding. PLOS Pathogens *19*, e1011868. URL: https://dx.plos.org/10.1371/journal.ppat.1011868. doi: 10.1371/journal.ppat.1011868.

10. Greaney, A.J., Starr, T.N., Gilchuk, P., Zost, S.J., Binshtein, E., Loes, A.N., Hilton, S.K., Huddleston, J., Eguia, R., Crawford, K.H., Dingens, A.S., Nargi, R.S., Sutton, R.E., Suryadevara, N., Rothlauf, P.W., Liu, Z., Whelan, S.P., Carnahan, R.H., Crowe, J.E., and Bloom, J.D. (2021). Complete Mapping of Mutations to the SARS-CoV-2 Spike Receptor-Binding Domain that Escape Antibody Recognition. Cell Host & Microbe *29*, 44–57.e9. URL: https://linkinghub.elsevier.com/retrieve/pii/S1931312820306247. doi: 10.1016/j.chom.2020.11.007.

11. Greaney, A.J., Starr, T.N., Barnes, C.O., Weisblum, Y., Schmidt, F., Caskey, M., Gaebler, C., Cho, A., Agudelo, M., Finkin, S., Wang, Z., Poston, D., Muecksch, F., Hatziioannou, T., Bieniasz, P.D., Robbiani, D.F., Nussenzweig, M.C., Bjorkman, P.J., and Bloom, J.D. (2021). Mapping mutations to the SARS-CoV-2 RBD that escape binding by different classes of antibodies. Nature Communications *12*, 4196. URL: https://www.nature.com/articles/s41467-021-24435-8. doi: 10.1038/s41467-021-24435-8.

12. Greaney, A.J., Loes, A.N., Crawford, K.H., Starr, T.N., Malone, K.D., Chu, H.Y., and Bloom, J.D. (2021). Comprehensive mapping of mutations in the SARS-CoV-2 receptor-binding domain that affect recognition by polyclonal human plasma antibodies. Cell Host & Microbe *29*, 463–476.e6. URL: https://linkinghub.elsevier.com/retrieve/pii/S1931312821000822. doi: 10.1016/j.chom.2021.02.003.

13. Tortorici, M.A., Czudnochowski, N., Starr, T.N., Marzi, R., Walls, A.C., Zatta, F., Bowen, J.E., Jaconi, S., Di Iulio, J., Wang, Z., De Marco, A., Zepeda, S.K., Pinto, D., Liu, Z., Beltramello, M., Bartha, I., Housley, M.P., Lempp, F.A., Rosen, L.E., Dellota, E., Kaiser, H., Montiel-Ruiz, M., Zhou, J., Addetia, A., Guarino, B., Culap, K., Sprugasci, N., Saliba, C., Vetti, E., Giacchetto-Sasselli, I., Fregni, C.S., Abdelnabi, R., Foo, S.Y.C., Havenar-Daughton, C., Schmid, M.A., Benigni, F., Cameroni, E., Neyts, J., Telenti, A., Virgin, H.W., Whelan, S.P.J., Snell, G., Bloom, J.D., Corti, D., Veesler, D., and Pizzuto, M.S. (2021). Broad sarbecovirus neutralization by a human monoclonal antibody. Nature *597*, 103–108. URL: https://www.nature.com/articles/s41586-021-03817-4. doi: 10.1038/s41586-021-03817-4.

14. Starr, T.N., Greaney, A.J., Dingens, A.S., and Bloom, J.D. (2021). Complete map of SARS-CoV-2 RBD mutations that escape the monoclonal antibody LY-CoV555 and its cocktail with LY-CoV016. Cell Reports Medicine *2*, 100255. URL: https://linkinghub.elsevier.com/retrieve/pii/S2666379121000719. doi: 10.1016/j.xcrm.2021.100255.

15. Starr, T.N., Greaney, A.J., Addetia, A., Hannon, W.W., Choudhary, M.C., Dingens, A.S., Li, J.Z., and Bloom, J.D. (2021). Prospective mapping of viral mutations that escape antibodies used to treat COVID-19. Science *371*, 850–854. URL: https://www.science.org/doi/10.1126/science.abf9302. doi: 10.1126/science.abf9302.

16. Starr, T.N., Czudnochowski, N., Liu, Z., Zatta, F., Park, Y.J., Addetia, A., Pinto, D., Beltramello, M., Hernandez, P., Greaney, A.J., Marzi, R., Glass, W.G., Zhang, I., Dingens, A.S., Bowen, J.E., Tortorici, M.A., Walls, A.C., Wojcechowskyj, J.A., De Marco, A., Rosen, L.E., Zhou, J., Montiel-Ruiz, M., Kaiser, H., Dillen, J.R., Tucker, H., Bassi, J., Silacci-Fregni, C., Housley, M.P., Di Iulio, J., Lombardo, G., Agostini, M., Sprugasci, N., Culap, K., Jaconi, S., Meury, M., Dellota Jr, E., Abdelnabi, R., Foo, S.Y.C., Cameroni, E., Stumpf, S., Croll, T.I., Nix, J.C., Havenar-Daughton, C., Piccoli, L., Benigni, F., Neyts, J., Telenti, A., Lempp, F.A., Pizzuto, M.S., Chodera, J.D., Hebner, C.M., Virgin, H.W., Whelan,

S.P.J., Veesler, D., Corti, D., Bloom, J.D., and Snell, G. (2021). SARS-CoV-2 RBD antibodies that maximize breadth and resistance to escape. Nature *597*, 97–102. URL: https://www.nature.com/articles/s41586-021-03807-6. doi: 10.1038/s41586-021-03807-6.

17. Dadonaite, B., Brown, J., McMahon, T.E., Farrell, A.G., Figgins, M.D., Asarnow, D., Stewart, C., Lee, J., Logue, J., Bedford, T., Murrell, B., Chu, H.Y., Veesler, D., and Bloom, J.D. (2024). Spike deep mutational scanning helps predict success of SARS-CoV-2 clades. Nature *631*, 617–626. URL: https://www.nature.com/articles/s41586-024-07636-1. doi: 10.1038/s41586-024-07636-1.

18. Yisimayi, A., Song, W., Wang, J., Jian, F., Yu, Y., Chen, X., Xu, Y., Yang, S., Niu, X., Xiao, T., Wang, J., Zhao, L., Sun, H., An, R., Zhang, N., Wang, Y., Wang, P., Yu, L., Lv, Z., Gu, Q., Shao, F., Jin, R., Shen, Z., Xie, X.S., Wang, Y., and Cao, Y. (2024). Repeated Omicron exposures override ancestral SARS-CoV-2 immune imprinting. Nature *625*, 148–156. URL: https://www.nature.com/articles/s41586-023-06753-7. doi: 10.1038/s41586-023-06753-7.

19. Ito, J., Strange, A., Liu, W., Joas, G., Lytras, S., The Genotype to Phenotype Japan (G2P-Japan) Consortium, Matsuno, K., Nao, N., Sawa, H., Mizuma, K., Kojima, I., Li, J., Tsubo, T., Tanaka, S., Tsuda, M., Wang, L., Oda, Y., Ferdous, Z., Shishido, K., Fukuhara, T., Tamura, T., Suzuki, R., Suzuki, S., Tsujino, S., Ito, H., Kaku, Y., Misawa, N., Plianchaisuk, A., Guo, Z., Hinay, A.A., Usui, K., Saikruang, W., Uriu, K., Kosugi, Y., Fujita, S., M.Tolentino, J.E., Chen, L., Pan, L., Li, W., Suganami, M., Chiba, M., Yoshimura, R., Yasuda, K., Iida, K., Ohsumi, N., Tanaka, S., Okumura, K., Yoshimura, K., Sadamas, K., Nagashima, M., Asakura, H., Yoshida, I., Nakagawa, S., Takaori-Kondo, A., Shirakawa, K., Nagata, K., Nomura, R., Horisawa, Y., Tashiro, Y., Kawai, Y., Takayama, K., Hashimoto, R., Deguchi, S., Watanabe, Y., Nakata, Y., Futatsusako, H., Sakamoto, A., Yasuhara, N., Hashiguchi, T., Suzuki, T., Kimura, K., Sasaki, J., Nakajima, Y., Yajima, H., Irie, T., Kawabata, R., Sasaki-Tabata, K., Ikeda, T., Nasse, H., Shimizu, R., Begum, M.M., Jonathan, M., Mugita, Y., Leong, S., Takahashi, O., Ichihara, K., Ueno, T., Motozono, C., Toyoda, M., Saito, A., Shofa, M., Shibatani, Y., Nishiuchi, T., Zahradni, J., Andrikopoulos, P., Padilla-Blanco, M., Konar, A., and Sato, K. (2025). A protein language model for exploring viral fitness landscapes. Nature Communications *16*, 4236. URL: https://www.nature.com/articles/s41467-025-59422-w. doi: 10.1038/s41467-025-59422-w.

20. Wang, D., Huot, M., Mohanty, V., and Shakhnovich, E.I. (2024). Biophysical principles predict fitness of sars-cov-2 variants. Proceedings of the National Academy of Sciences *121*, e2314518121. URL: https://www.pnas.org/doi/abs/10.1073/pnas.2314518121. doi: 10.1073/pnas.2314518121. arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.2314518121.

21. Maher, M.C., Bartha, I., Weaver, S., Di Iulio, J., Ferri, E., Soriaga, L., Lempp, F.A., Hie, B.L., Bryson, B., Berger, B., Robertson, D.L., Snell, G., Corti, D., Virgin, H.W., Kosakovsky Pond, S.L., and Telenti, A. (2022). Predicting the mutational drivers of future SARS-CoV-2 variants of concern. Science Translational Medicine *14*, eabk3445. URL: https://www.science.org/doi/10.1126/scitranslmed.abk3445. doi: 10.1126/scitranslmed.abk3445.

22. Huot, M., Wang, D., Liu, J., and Shakhnovich, E.I. (2025). Predicting high-fitness viral protein variants with bayesian active learning and biophysics. Proceedings of the National Academy of Sciences *122*, e2503742122. URL: https://www.pnas.org/doi/abs/10.1073/pnas.2503742122. doi: 10.1073/pnas.2503742122. arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.2503742122.

23. Thadani, N.N., Gurev, S., Notin, P., Youssef, N., Rollins, N.J., Ritter, D., Sander, C., Gal, Y., and Marks, D.S. (2023). Learning from prepandemic data to forecast viral escape. Nature *622*, 818–825. URL: https://www.nature.com/articles/s41586-023-06617-0. doi: 10.1038/s41586-023-06617-0.

24. Wang, G., Liu, X., Wang, K., Gao, Y., Li, G., Baptista-Hon, D.T., Yang, X.H., Xue, K., Tai, W.H., Jiang, Z., Cheng, L., Fok, M., Lau, J.Y.N., Yang, S., Lu, L., Zhang, P., and Zhang, K. (2023). Deep-learning-enabled protein–protein interaction analysis for prediction of SARS-CoV-2 infectivity and variant evolution. Nature Medicine *29*, 2007–2018. URL: https://www.nature.com/articles/s41591-023-02483-5. doi: 10.1038/s41591-023-02483-5.

25. Dianzhuo Wang, Huot, M., Zechen Zhang, Kaiyi Jiang, Shakhnovich, E.I., and Esvelt, K.M. (2025). Without Safeguards, AI-Biology Integration Risks Accelerating Future Pandemics. URL: https://rgdoi.net/10.13140/RG.2.2.29765.15849. doi: 10.13140/RG.2.2.29765.15849. Publisher: Unpublished.

26. Rodriguez-Rivas, J., Croce, G., Muscat, M., and Weigt, M. (2022). Epistatic models predict mutable sites in SARS-CoV-2 proteins and epitopes. Proceedings of the National Academy of Sciences *119*, e2113118119. URL: https://pnas.org/doi/full/10.1073/pnas.2113118119. doi: 10.1073/pnas.2113118119.

27. Bari, L.D., Bisardi, M., Cotogno, S., Weigt, M., and Zamponi, F. (2024). Emergent time scales of epistasis in protein evolution. Proceedings of the National Academy of Sciences *121*, e2406807121. URL: https://www.pnas.org/doi/abs/10.1073/pnas.2406807121. doi: 10.1073/pnas.2406807121. arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.2406807121.

28. Mauri, E., Cocco, S., and Monasson, R. (2023). Mutational Paths with Sequence-Based Models of Proteins: From Sampling to Mean-Field Characterization. Physical Review Letters *130*, 158402. URL: https://link.aps.org/doi/10.1103/PhysRevLett.130.158402. doi: 10.1103/PhysRevLett.130.158402.

29. Mauri, E., Cocco, S., and Monasson, R. (2023). Transition paths in Potts-like energy landscapes: General properties and application to protein sequence models. Physical Review E *108*, 024141. URL: https://link.aps.org/doi/10.1103/PhysRevE.108.024141. doi: 10.1103/PhysRevE.108.024141.

30. Greenbury, S.F., Louis, A.A., and Ahnert, S.E. (2022). The structure of genotype-phenotype maps makes fitness landscapes navigable. Nature Ecology & Evolution *6*, 1742–1752. URL: https://www.nature.com/articles/s41559-022-01867-z. doi: 10.1038/s41559-022-01867-z.

31. Fischer, A., and Igel, C. (2014). Training restricted Boltzmann machines: An introduction. Pattern Recognition *47*, 25–39. URL: https://linkinghub.elsevier.com/retrieve/pii/S0031320313002495. doi: 10.1016/j.patcog.2013.05.025.

32. Tubiana, J., Cocco, S., and Monasson, R. (2019). Learning protein constitutive motifs from sequence data. eLife *8*, e39397. URL: https://elifesciences.org/articles/39397. doi: 10.7554/eLife.39397.

33. Lau, K.F., and Dill, K.A. (1989). A lattice statistical mechanics model of the conformational and sequence spaces of proteins. Macromolecules *22*, 3986–3997. URL: https://pubs.acs.org/doi/abs/10.1021/ma00200a030. doi: 10.1021/ma00200a030.

34. Loffredo, E., Vesconi, E., Razban, R., Peleg, O., Shakhnovich, E., Cocco, S., and Monasson, R. (2023). Evolutionary dynamics of a lattice dimer: a toy model for stability vs. affinity trade-offs in proteins. Journal of Physics A: Mathematical and Theoretical *56*, 455002. URL: https://iopscience.iop.org/article/10.1088/1751-8121/acfddc. doi: 10.1088/1751-8121/acfddc.

35. Jacquin, H., Gilson, A., Shakhnovich, E., Cocco, S., and Monasson, R. (2016). Benchmarking Inverse Statistical Approaches for Protein Structure and Design with Exactly Solvable Models. PLOS Computational Biology *12*, e1004889. URL: https://dx.plos.org/10.1371/journal.pcbi.1004889. doi: 10.1371/journal.pcbi.1004889.

36. Yang, H., Carney, P.J., Chang, J.C., Guo, Z., Villanueva, J.M., and Stevens, J. (2015). Structure and receptor binding preferences of recombinant human A(H3N2) virus hemagglutinins. Virology *477*, 18–31. URL: https://linkinghub.elsevier.com/retrieve/pii/S004268221400573X. doi: 10.1016/j.virol.2014.12.024.

37. Greaney, A.J., Starr, T.N., and Bloom, J.D. (2022). An antibody-escape estimator for mutations to the sars-cov-2 receptor-binding domain. Virus Evolution *8*, veac021. URL: https://doi.org/10.1093/ve/veac021. doi: 10.1093/ve/veac021. arXiv:https://academic.oup.com/ve/article-pdf/8/1/veac021/43671163/veac021.pdf.

38. Huot, M., Rosenbaum, P., Planchais, C., Mouquet, H., Monasson, R., and Cocco, S. (2025). Generative model of sars-cov-2 variants under functional and immune pressure unveils viral escape potential and antibody resilience. bioRxiv. URL: https://www.biorxiv.org/content/early/2025/05/13/2025.05.12.653592. doi: 10.1101/2025.05.12.653592. arXiv:https://www.biorxiv.org/content/early/2025/05/13/2025.05.12.653592.full.pdf.

39. Cocco, S., Feinauer, C., Figliuzzi, M., Monasson, R., and Weigt, M. (2018). Inverse statistical physics of protein sequences: a key issues review. Reports on Progress in Physics *81*, 032601. URL: https://iopscience.iop.org/article/10.1088/1361-6633/aa9965. doi: 10.1088/1361-6633/aa9965.

40. Hsu, C., Verkuil, R., Liu, J., Lin, Z., Hie, B., Sercu, T., Lerer, A., and Rives, A. (2022). Learning inverse folding from millions of predicted structures. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, eds. Proceedings of the 39th International Conference on Machine Learning vol. 162 of *Proceedings of Machine Learning Research*. PMLR pp. 8946–8970. URL: https://proceedings.mlr.press/v162/hsu22a.html.

41. Youssef, N., Gurev, S., Ghantous, F., Brock, K.P., Jaimes, J.A., Thadani, N.N., Dauphin, A., Sherman, A.C., Yurkovetskiy, L., Soto, D., Estanboulieh, R., Kotzen, B., Notin, P., Kollasch, A.W., Cohen, A.A., Dross, S.E., Erasmus, J., Fuller, D.H., Bjorkman, P.J., Lemieux, J.E., Luban, J., Seaman, M.S., and Marks, D.S. (2025). Computationally designed proteins mimic antibody immune evasion in viral evolution. Immunity *58*, 1411–1421.e6. URL: https://linkinghub.elsevier.com/retrieve/pii/S1074761325001785. doi: 10.1016/j.immuni.2025.04.015.

42. Baum, A., Fulton, B.O., Wloga, E., Copin, R., Pascal, K.E., Russo, V., Giordano, S., Lanza, K., Negron, N., Ni, M., Wei, Y., Atwal, G.S., Murphy, A.J., Stahl, N., Yancopoulos, G.D., and Kyratsous, C.A. (2020). Antibody cocktail to sars-cov-2 spike protein prevents rapid mutational escape seen with individual antibodies. Science *369*, 1014–1018. URL: https://www.science.org/doi/abs/10.1126/science.abd0831. doi: 10.1126/science.abd0831. arXiv:https://www.science.org/doi/pdf/10.1126/science.abd0831.

43. Ku, Z., Xie, X., Davidson, E., Ye, X., Su, H., Menachery, V.D., Li, Y., Yuan, Z., Zhang, X., Muruato, A.E., I Escuer, A.G., Tyrell, B., Doolan, K., Doranz, B.J., Wrapp, D., Bates, P.F., McLellan, J.S., Weiss, S.R., Zhang, N., Shi, P.Y., and An, Z. (2021). Molecular determinants and mechanism for antibody cocktail preventing SARS-CoV-2 escape. Nature Communications *12*, 469. URL: https://www.nature.com/articles/s41467-020-20789-7. doi: 10.1038/s41467-020-20789-7.

44. Yu, T.C., Thornton, Z.T., Hannon, W.W., DeWitt, W.S., Radford, C.E., Matsen, F.A., and Bloom, J.D. (2022). A biophysical model of viral escape from polyclonal antibodies. Virus Evolution *8*, veac110. URL: https://academic.oup.com/ve/article/doi/10.1093/ve/veac110/6889254. doi: 10.1093/ve/veac110.

45. Andreano, E., Paciello, I., Pierleoni, G., Piccini, G., Abbiento, V., Antonelli, G., Pileri, P., Manganaro, N., Pantano, E., Maccari, G., Marchese, S., Donnici, L., Benincasa, L., Giglioli, G., Leonardi, M., De Santi, C., Fabbiani, M., Rancan, I., Tumbarello, M., Montagnani, F., Sala, C., Medini, D., De Francesco, R., Montomoli, E., and Rappuoli, R. (2023). B cell analyses after SARS-CoV-2 mRNA third vaccination reveals a hybrid immunity like antibody response. Nature Communications *14*, 53. URL: https://www.nature.com/articles/s41467-022-35781-6. doi: 10.1038/s41467-022-35781-6.

46. Shekhar, K., Ruberman, C.F., Ferguson, A.L., Barton, J.P., Kardar, M., and Chakraborty, A.K. (2013). Spin models inferred from patient-derived viral sequence data faithfully describe HIV fitness landscapes. Physical Review E *88*, 062705. URL: https://link.aps.org/doi/10.1103/PhysRevE.88.062705. doi: 10.1103/PhysRevE.88.062705.

47. Doelger, J., Kardar, M., and Chakraborty, A.K. (2022). Inferring the intrinsic mutational fitness landscape of influenzalike evolving antigens from temporally ordered sequence data. Physical Review E *105*, 024401. URL: https://link.aps.org/doi/10.1103/PhysRevE.105.024401. doi: 10.1103/PhysRevE.105.024401.

48. Shakhnovich, E.I., and Gutin, A.M. (1993). Engineering of stable and fast-folding sequences of model proteins. Proceedings of the National Academy of Sciences *90*, 7195–7199. URL: https://pnas.org/doi/full/10.1073/pnas.90.15.7195. doi: 10.1073/pnas.90.15.7195.

49. Miyazawa, S., and Jernigan, R.L. (1996). Residue – Residue Potentials with a Favorable Contact Pair Term and an Unfavorable High Packing Density Term, for Simulation and Threading. Journal of Molecular Biology *256*, 623–644. URL: https://linkinghub.elsevier.com/retrieve/pii/S002228369690114X. doi: 10.1006/jmbi.1996.0114.

50. Elbe, S., and Buckland-Merrett, G. (2017). Data, disease and diplomacy: GISAID's innovative contribution to global health: Data, Disease and Diplomacy. Global Challenges *1*, 33–46. URL: https://onlinelibrary.wiley.com/doi/10.1002/gch2.1018. doi: 10.1002/gch2.1018.

51. Katoh, K., and Standley, D.M. (2013). MAFFT Multiple Sequence Alignment Software Version 7: Improvements in Performance and Usability. Molecular Biology and Evolution *30*, 772–780. URL: https://academic.oup.com/mbe/article-lookup/doi/10.1093/molbev/mst010. doi: 10.1093/molbev/mst010.

Supplementary Materials for: Constrained Evolutionary Funnels Shape Viral Immune Escape

Huot et al.

Corresponding author: simona.cocco@phys.ens.fr

RBM-based sequence modeling

We model the probability distribution over receptor-binding domain (RBD) sequences using *Restricted Boltzmann Machines* (RBMs), using the (PGM package)[32]. This model is trained on homologous sequences to capture essential evolutionary, structural, and functional constraints on the RBD.

An RBM defines a joint probability distribution between protein sequences and latent variables using a bipartite graphical model. The visible layer corresponds to the amino acid sequence $v = (v_1, v_2, \ldots, v_N)$, with each v_i taking one of 21 values (20 amino acids plus a gap). The hidden layer comprises real-valued units $h = (h_1, h_2, \ldots, h_M)$ that represent learned features across the sequence.

The joint distribution over visible and hidden units is given by:

$$P_{\text{RBM}}(v, h) \propto \exp\left(\sum_{i=1}^{N} g_i(v_i) - \sum_{\mu=1}^{M} \mathcal{U}_\mu(h_\mu) + \sum_{i,\mu} h_\mu w_{i\mu}(v_i)\right), \tag{21}$$

where $g_i(v_i)$ and $\mathcal{U}_\mu(h_\mu)$ are local potentials (biases) on the visible and hidden units, and $w_{i\mu}(v_i)$ are interaction weights between them.

The hidden unit potential $\mathcal{U}_\mu(h)$ follows a double rectified linear unit (dReLU) form:

$$\mathcal{U}_\mu(h) = \frac{1}{2}\gamma_{\mu,+}h_+^2 + \frac{1}{2}\gamma_{\mu,-}h_-^2 + \theta_{\mu,+}h_+ + \theta_{\mu,-}h_-, \quad \text{with} \quad h_+ = \max(h, 0), \quad h_- = \min(h, 0). \tag{22}$$

Model parameters $(\gamma_{\mu,+}, \gamma_{\mu,-}, \theta_{\mu,+}, \theta_{\mu,-})$ are estimated by maximizing the likelihood of sequences in the multiple sequence alignment (MSA), i.e., optimizing the marginal distribution:

$$P_{\text{RBM}}(v) \propto \int \mathrm{d}h\, P_{\text{RBM}}(v, h) = \exp\left(\sum_{i=1}^{N} g_i(v_i) + \sum_{\mu=1}^{M} \Gamma_\mu(I_\mu(v))\right), \tag{23}$$

where $I_\mu(v) = \sum_i w_{i\mu}(v_i)$ is the input to hidden unit μ, and the cumulant generating function Γ_μ is defined as:

$$\Gamma_\mu(I) = \log \int \mathrm{d}h_\mu \, \exp\left(I h_\mu - \mathcal{U}_\mu(h_\mu)\right). \tag{24}$$

and computed following the method described in Tubiana et al.[32].

MCMC algorithm for path sampling

Algorithm 1 Metropolis–Hastings sweep over a mutation path of length $T+1$

Input: path $\{v_t\}_{t=0}^{T}$, weight $P(\cdot)$, alphabet $\mathcal{A}$ of size Q
for $k = 1$ **to** T **do**
  Draw $t \sim \mathrm{Unif}\{1, \ldots, T\}$
  Set $v_- \leftarrow v_{t-1}, \quad v \leftarrow v_t, \quad [\, v_+ \leftarrow v_{t+1}$ if $t < T\,]$

(A) Proposal v'
  if $t = T$ **or** $d_H(v_-, v_+) = 0$ **then**
    Choose $i \sim \mathrm{Unif}\{1, \ldots, L\}$
    Choose $z \sim \mathrm{Unif}(\mathcal{A})$
    $v' \leftarrow M(v_-,\, i,\, z)$
  else if $d_H(v_-, v_+) = 1$ **then**
    Let i be the unique site with $v_{-,i} \neq v_{+,i}$
    Choose $z \sim \mathrm{Unif}(\mathcal{A})$
    $v' \leftarrow M(v_-,\, i,\, z)$
  else if $d_H(v_-, v_+) = 2$ **then**
    Identify the two single–mutation candidates compatible with both neighbours
    $v' \leftarrow$ the one different from v
  end if

(B) Acceptance probability:

$$p_{\mathrm{acc}} = \min\Big\{1,\ \frac{P(v')}{P(v)}\,\frac{\delta(v)}{\delta(v')}\Big\}, \quad \delta(u) = \big|\{(i,z)\colon M(v_-, i, z) = u\}\big|.$$

  With probability p_{acc} set $v_t \leftarrow v'$
end for
Output: updated path $\{v_t\}$

Let

$$\delta(u) = \big|\{(i,z)\colon M(v_-, i, z) = u\}\big|.$$

If $d_H(v_-, v_+) = 0$, then

$$\delta(u) = \begin{cases} L, & \text{if } u = v_-, \\ 1, & \text{if } d_H(u, v_-) = 1, \\ 0, & \text{otherwise}, \end{cases}$$

and therefore the acceptance probability

$$p_{\mathrm{acc}} = \min\Big\{1,\ \frac{P(v')}{P(v)}\,\frac{\delta(v)}{\delta(v')}\Big\}$$

specializes to

$$p_{\mathrm{acc}} = \begin{cases} 1, & v' = v, \\ \min\Big\{1,\ \dfrac{P(v')}{P(v)}\Big\}, & v, v' \neq v_-, \\ \min\Big\{1,\ \dfrac{P(v_-)}{L\,P(v)}\Big\}, & v' = v_-,\ v \neq v_-, \\ \min\Big\{1,\ L\,\dfrac{P(v')}{P(v)}\Big\}, & v = v_-,\ v' \neq v_-. \end{cases}$$

If $d_H(v_-, v_+) \neq 0$, then every proposal and its reverse are generated in exactly one way, so $\delta(v) = \delta(v') = 1$, and hence

$$p_{\mathrm{acc}} = \min\Big\{1,\ \frac{P(v')}{P(v)}\Big\}.$$

Proof of detailed balance. Let

$$q(v \to v') = \frac{\delta(v')}{L\,Q}, \quad p_{\mathrm{acc}}(v \to v') = \min\Big\{1,\ \frac{P(v')\,\delta(v)}{P(v)\,\delta(v')}\Big\}.$$

Then

$$P(v)\,q(v \to v')\,p_{\mathrm{acc}}(v \to v') = P(v)\,\frac{\delta(v')}{LQ}\,\min\Big\{1,\ \frac{P(v')\,\delta(v)}{P(v)\,\delta(v')}\Big\} = \frac{\min\{P(v)\,\delta(v'),\ P(v')\,\delta(v)\}}{LQ}.$$

By symmetry of the minimum,

$$\frac{\min\{P(v)\,\delta(v'),\ P(v')\,\delta(v)\}}{LQ} = P(v')\,\frac{\delta(v)}{LQ}\,\min\Big\{1,\ \frac{P(v)\,\delta(v')}{P(v')\,\delta(v)}\Big\} = P(v')\,q(v' \to v)\,p_{\mathrm{acc}}(v' \to v).$$

Moreover,

$$\frac{P(v')}{P(v)} = \frac{P_{\mathrm{path}}(\dots, v', \dots)}{P_{\mathrm{path}}(\dots, v, \dots)}.$$

Hence $P(v)\,q(v \to v')\,p_{\mathrm{acc}}(v \to v') = P(v')\,q(v' \to v)\,p_{\mathrm{acc}}(v' \to v)$ implies detailed balance for P_{path} as well.

Path sampling in nucleotide space

Path sampling in nucleotide space follows the same Metropolis–Hastings framework described above, with the key distinction that each v_t represents a full nucleotide sequence and each transition $(v_{t-1} \to v_t)$ introduces at most one nucleotide mutation. The path distribution is proportional to $P_{\mathrm{path}}\big(\{v_t\}\big) \propto \prod_{t=0}^{T} P(v_t)$ where

$$P(v) \propto P_{\mathrm{RBM}}\big(\mathrm{translate}(v)\big) \times \prod_{n \in v} f(n|\mathrm{aa}(n)), \tag{25}$$

and $f(n|\mathrm{aa}) = 1/\#\mathrm{codons}(\mathrm{aa})$ is a uniform codon-bias term, following Di-Bari et al.[27]. Proposals introducing stop codons are automatically rejected.

Mean field

We extend the framework developed by Mauri et al.[28,29] for transition paths, to allow for paths that are unconstrained at their endpoints and evolve under external immune pressure. To describe the evolution of such paths, we introduce the following order parameters: the binding score for antibody η, $\gamma_\eta = \frac{1}{L}\sum_i \alpha_i^\eta(v_i)$; the activation of hidden unit μ, $m_\mu = \frac{I_\mu}{L}$; and the average site overlap between successive sequences, $q_t = \frac{1}{L}\sum_i \delta_{v_{i,t},v_{i,t+1}}$.

Sequence stability is given by the RBM:

$$P_{\text{RBM}}(v) = \exp\left[\sum_{i=1}^{N} g_i\left(v_i\right) + \sum_{\mu=1}^{M} \Gamma_\mu\left(Lm_\mu\right)\right] \tag{26}$$

Antibody escape is given by the product of the probability to escape each antibody:

$$P_{\text{escape}}(v) = \prod_\eta \left(1 - \exp(-L\gamma_\eta)\right) \tag{27}$$

Sequence probability is a combination of stability and antibody escape:

$$P(v) = \exp\left[\sum_{i=1}^{N} g_i\left(v_i\right) - L\zeta(m_\mu, \gamma_\eta)\right] \tag{28}$$

with the potential

$$\zeta(m_\mu, \gamma_\eta) = -\frac{1}{L}\left[\sum_\mu \Gamma_\mu\left(Lm_\mu\right) + \sum_\eta \log\left(1 - \exp(-L\gamma_\eta)\right)\right] \tag{29}$$

The potential Φ should forbid large jumps along the paths. We thus consider a hard-wall repulsive potential,

$$\Phi(q) = \begin{cases} \dfrac{1}{L}\cdot\dfrac{1/T^2}{q - q_c} & \text{if } q_c < q \leq 1, \\ +\infty & \text{otherwise} \end{cases} \tag{30}$$

The location of the hard wall, $q_c = 1 - 1/L$, allows the path to explore at most $T \times L(1 - q_c) = T$ mutations in T steps.

$$P_{path}(\mathbf{m_t}, \gamma_\mathbf{t}, q_t) \propto \sum_{\{\mathbf{v}_t\}} \prod_{t,\mu} \delta\left(\frac{1}{L}\sum_i w_{i\mu}(v_{i,t}) - m_t^\mu\right) \prod_{t,\eta} \delta\left(\frac{1}{L}\sum_i \alpha_{i\eta}(v_{i,t}) - \gamma_t^\eta\right) \prod_t \delta\left(\frac{1}{L}\sum_i \delta_{v_{i,t},v_{i,t+1}} - q_t\right)$$
$$\exp\left(\beta\sum_i g_i(v_{i,t}) + L\left(-\sum_t \zeta(\mathbf{m_t}, \gamma_\mathbf{t}) - \sum_t \Phi(q_t)\right)\right)$$

The probability of a path is given by a Boltzmann weight involving the path free energy f_{path}:

$$P_{path}(\mathbf{m_t}, \gamma_\mathbf{t}, q_t) \propto \exp\left[-L\beta\, f_{\text{path}}\left(\{\mathbf{m}_t, \gamma_t, q_t\}\right)\right], \tag{31}$$

$$f_{\text{path}}(\{\mathbf{m}_t, \gamma_t, q_t\}) = \sum_t \zeta(\mathbf{m}_t, \gamma_t) + \sum_t \Phi(q_t) - \frac{1}{\beta} S_\beta(\{\mathbf{m}_t, \gamma_t, q_t\}_{t=1}^T), \tag{32}$$

$$S_\beta(\{\mathbf{m}_t, \gamma_t, q_t\}_{t=1}^T) = \min_{\hat{m},\hat{\gamma},\hat{q}} \frac{1}{L}\sum_i \log Z_i(\{\hat{m}_t, \hat{\gamma}_t, \hat{q}_t\}) - \sum_{t,\mu} m_t^\mu \hat{m}_t^\mu - \sum_{t,\eta} \gamma_t^\eta \hat{\gamma}_t^\eta - \sum_t q_t \hat{q}_t, \tag{33}$$

$$Z_i(\{\hat{m}_t, \hat{\gamma}_t, \hat{q}_t\}) = \sum_{\{v_t\}} \exp\left[\beta \sum_t g_i(v_t) + \sum_{t,\mu} \hat{m}_t^\mu w_{i\mu}(v_t) + \sum_{t,\eta} \hat{\gamma}_t^\eta \alpha_{i\eta}(v_t) + \sum_t \hat{q}_t \delta_{v_{i,t}, v_{i,t+1}}\right]. \quad (34)$$

Minimizing the path free–energy density:

$$\frac{\partial f_{\text{path}}}{\partial m_t^\mu} = 0, \quad \frac{\partial f_{\text{path}}}{\partial \gamma_t^\eta} = 0, \quad \frac{\partial f_{\text{path}}}{\partial q_t} = 0, \quad (35)$$

leads to the following saddle point equations for the conjugate fields:

$$\frac{\hat{m}_t^\mu}{\beta} = -\frac{\partial \zeta(\mathbf{m}_t, \boldsymbol{\gamma}_t)}{\partial m_t^\mu}, \quad \frac{\hat{\gamma}_t^\eta}{\beta} = -\frac{\partial \zeta(\mathbf{m}_t, \boldsymbol{\gamma}_t)}{\partial \gamma_t^\eta}, \quad \frac{\hat{q}_t}{\beta} = -\frac{\partial \Phi(q_t)}{\partial q_t}, \quad (36)$$

In addition, the minimization of the entropy term S yields the self-consistency conditions:

$$\begin{aligned} m_t^\mu &= \frac{\partial}{\partial \hat{m}_t^\mu}\left(\frac{1}{L}\sum_i \log Z_i(\{\hat{m}_t, \hat{\gamma}_t, \hat{q}_t\})\right), \\ \gamma_t^\eta &= \frac{\partial}{\partial \hat{\gamma}_t^\eta}\left(\frac{1}{L}\sum_i \log Z_i(\{\hat{m}_t, \hat{\gamma}_t, \hat{q}_t\})\right), \\ q_t &= \frac{\partial}{\partial \hat{q}_t}\left(\frac{1}{L}\sum_i \log Z_i(\{\hat{m}_t, \hat{\gamma}_t, \hat{q}_t\})\right), \end{aligned} \quad (37)$$

See SI for free energy minimisation and computation of optimal $(\mathbf{m_t}, \gamma_\mathbf{t}, q_t)$.

Free energy optimisation in mean field

We first initialize $(\mathbf{m}_t, \boldsymbol{\gamma}_t, q_t)$ to $(\mathbf{m}_0, \boldsymbol{\gamma}_0, 1), \forall t$ and apply algorithm 2.

Algorithm 2 Mean-Field Iteration to Compute $\mathbf{m}_t, \boldsymbol{\gamma}_t, q_t$

1: Compute conjugate fields:

$$\hat{m}_t^\mu \leftarrow -\frac{\partial \zeta(\mathbf{m}_t, \boldsymbol{\gamma}_t)}{\partial m_t^\mu},$$
$$\hat{\gamma}_t^\eta \leftarrow -\frac{\partial \zeta(\mathbf{m}_t, \boldsymbol{\gamma}_t)}{\partial \gamma_t^\eta},$$
$$\hat{q}_t \leftarrow -\frac{\partial \Phi(q_t)}{\partial q_t}.$$

2: Solve the following equations for $dm_t^\mu, d\gamma_t^\eta, dq_t$:

$$dm_t^\mu \leftarrow \frac{\partial}{\partial \hat{m}_t^\mu}\left(\frac{1}{L}\sum_i \log Z_i(\{\hat{m}_t, \hat{\gamma}_t, \hat{q}_t\})\right),$$
$$d\gamma_t^\eta \leftarrow \frac{\partial}{\partial \hat{\gamma}_t^\eta}\left(\frac{1}{L}\sum_i \log Z_i(\{\hat{m}_t, \hat{\gamma}_t, \hat{q}_t\})\right),$$
$$dq_t \leftarrow \frac{\partial}{\partial \hat{q}_t}\left(\frac{1}{L}\sum_i \log Z_i(\{\hat{m}_t, \hat{\gamma}_t, \hat{q}_t\})\right).$$

3: Perform gradient updates:

- $m_t^\mu \leftarrow m_t^\mu + \epsilon \cdot (dm_t^\mu - m_t^\mu)$
- $\gamma_t^\eta \leftarrow \gamma_t^\eta + \epsilon \cdot (d\gamma_t^\eta - \gamma_t^\eta)$
- $q_t \leftarrow q_t + \epsilon \cdot (dq_t - q_t)$

After each update of the conjugate fields, the computation of $\log Z_i$ is carried out using a transfer matrix method. This exploits the Markovian structure of the temporal sequence v_t to recursively build the partition function. .

Gradients of $\log Z_i$ with respect to the conjugate fields are obtained via automatic differentiation, using the `PyTorch` library.

Amino acid probability in mean field

We introduce a time–resolved local field $H_{i,t}(a) = \beta\, g_i(a)$.

$$\left\langle \delta_{v_{i,t},a} \right\rangle_{\text{path}} = \frac{\partial \ln Z}{\partial H_{i,t}(a)} = \frac{\partial \ln Z_i}{\partial H_{i,t}(a)} = \frac{1}{\beta}\frac{\partial \ln Z_i}{\partial g_{i,t}(a)}. \tag{38}$$

Path Entropy

To build intuition, we first consider limiting cases for the path entropy per site S_{path}.

When paths are constrained to a single mutation per step and no external fields are applied, the entropy is given by

$$S_{\text{path}} = \frac{T}{L}\log[L(Q-1)]. \tag{39}$$

In the case where continuity is removed and no field is present, the entropy becomes

$$S_{\text{path}} = T \log Q. \tag{40}$$

If the field is site-independent and continuity is still removed, the entropy takes the form

$$S_{\text{path}} = T \cdot \langle S_{\text{site}} \rangle \,, \tag{41}$$

where $\langle S_{\text{site}} \rangle$ denotes the average entropy across sites induced by the field.

To compute the entropy in the full mean-field model, we begin with the expression for the path free energy:

$$f_{\text{path}}(\beta; \{\boldsymbol{m}_t, \boldsymbol{\gamma}_t, q_t\}) = \sum_t \zeta(\boldsymbol{m}_t, \boldsymbol{\gamma}_t) + \sum_t \Phi(q_t) - \frac{1}{\beta} S_\beta(\{\boldsymbol{m}_t, \boldsymbol{\gamma}_t, q_t\}). \tag{42}$$

where

$$S_\beta = \frac{1}{L} \sum_i \log Z_i(\{\hat{m}_t, \hat{\gamma}_t, \hat{q}_t\}) - \sum_{t,\mu} m_t^\mu \hat{m}_t^\mu - \sum_{t,\eta} \gamma_t^\eta \hat{\gamma}_t^\eta - \sum_t q_t \hat{q}_t, \tag{43}$$

and S_β derivative with respect to β follows from the explicit dependence of Z_i:

$$\frac{\partial S_\beta}{\partial \beta} = \frac{1}{L} \sum_i \frac{\partial}{\partial \beta} \log Z_i = \frac{1}{L} \sum_{t,i} \langle g_i(v_t) \rangle_{Z_i} \,. \tag{44}$$

Combining all expressions, we obtain the final form for the path entropy following:

$$S_{path} = -\frac{df_{\text{path}}}{d(1/\beta)} = S_\beta - \beta \frac{\partial S_\beta}{\partial \beta}$$

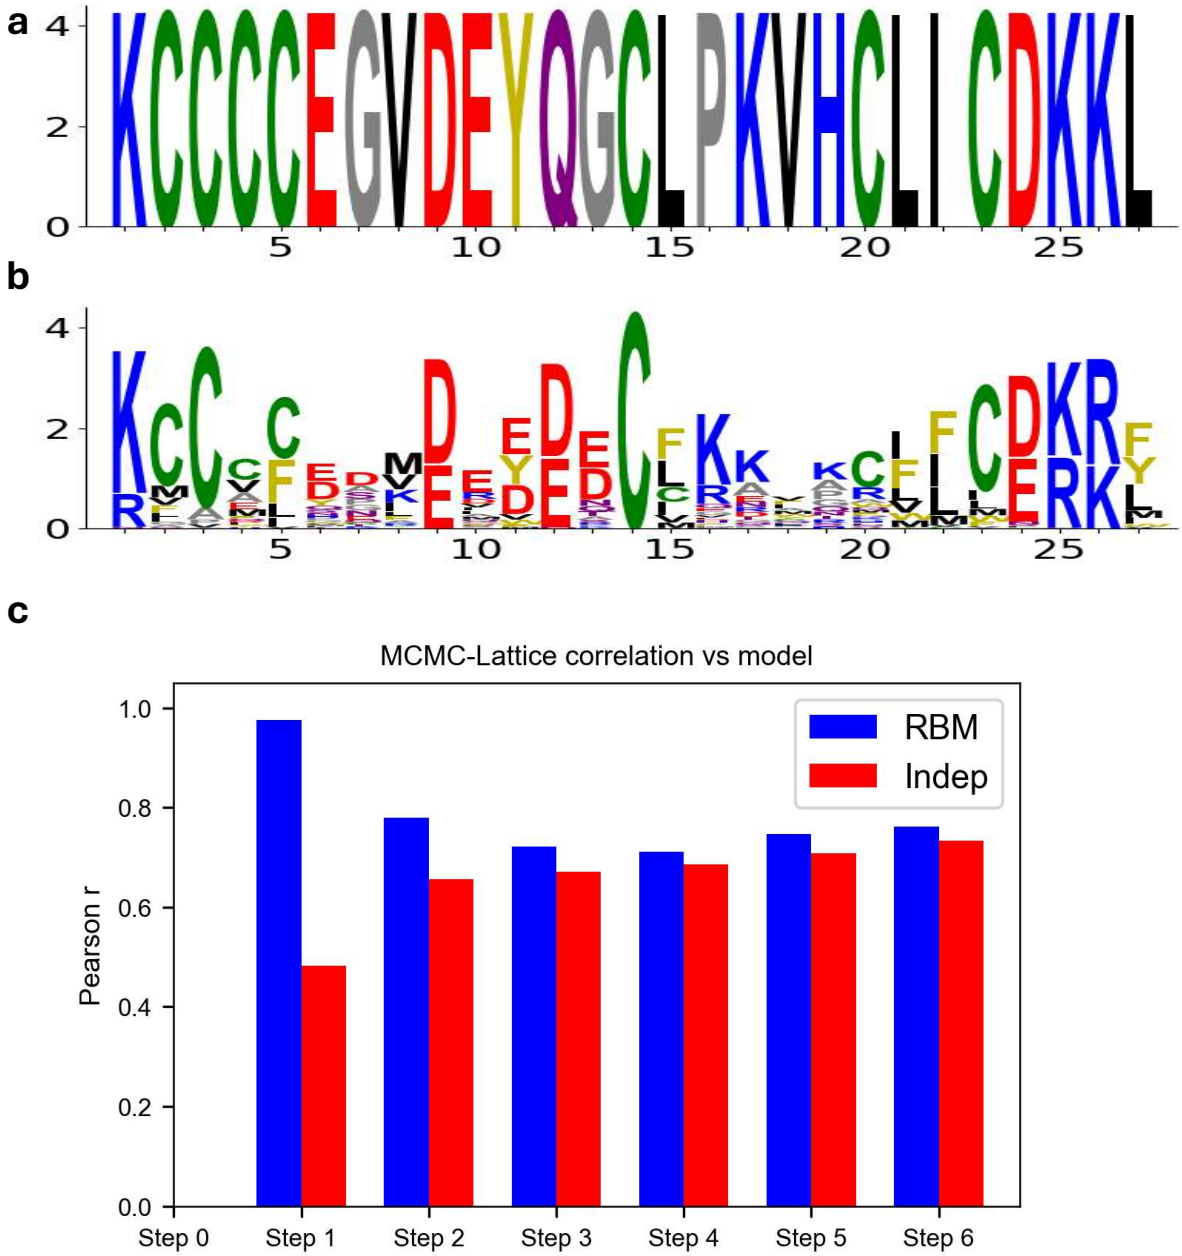

Fig. S1: Escape paths in protein lattice model.
(a) Wildtype sequence used as reference in the lattice model. **(b)** Sequence logo from an MSA of high-folding-probability sequences generated using the lattice model ($\beta = 1000$). Larger letters indicate higher probability of observing a given amino acid at that site. **(c)** Pearson correlation between per-site mutation probabilities from RBM-based MCMC paths and ground truth lattice model (blue). Comparison with an independent-site model (red).

Figure S1

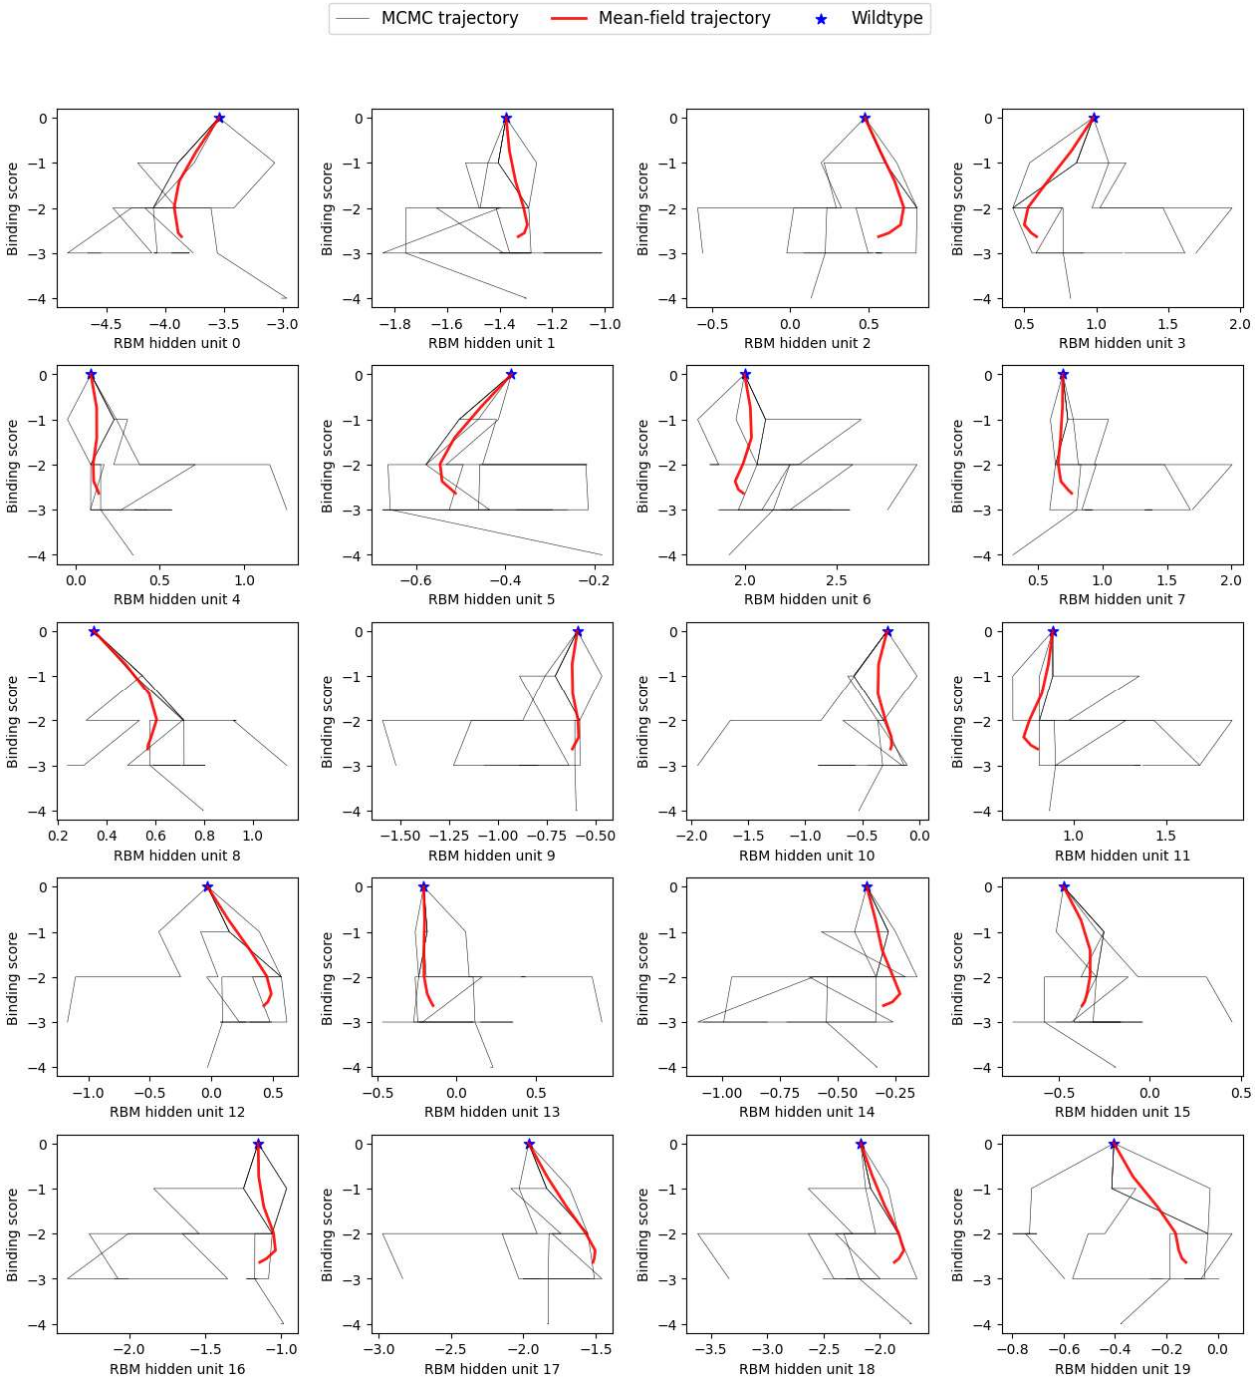

Fig. S2: Mean field and MCMC trajectories in lattice model. Projection of mean field and MCMC trajectories in space defined by order parameters corresponding to antibody binding score and one RBM hidden unit.

Figure S2

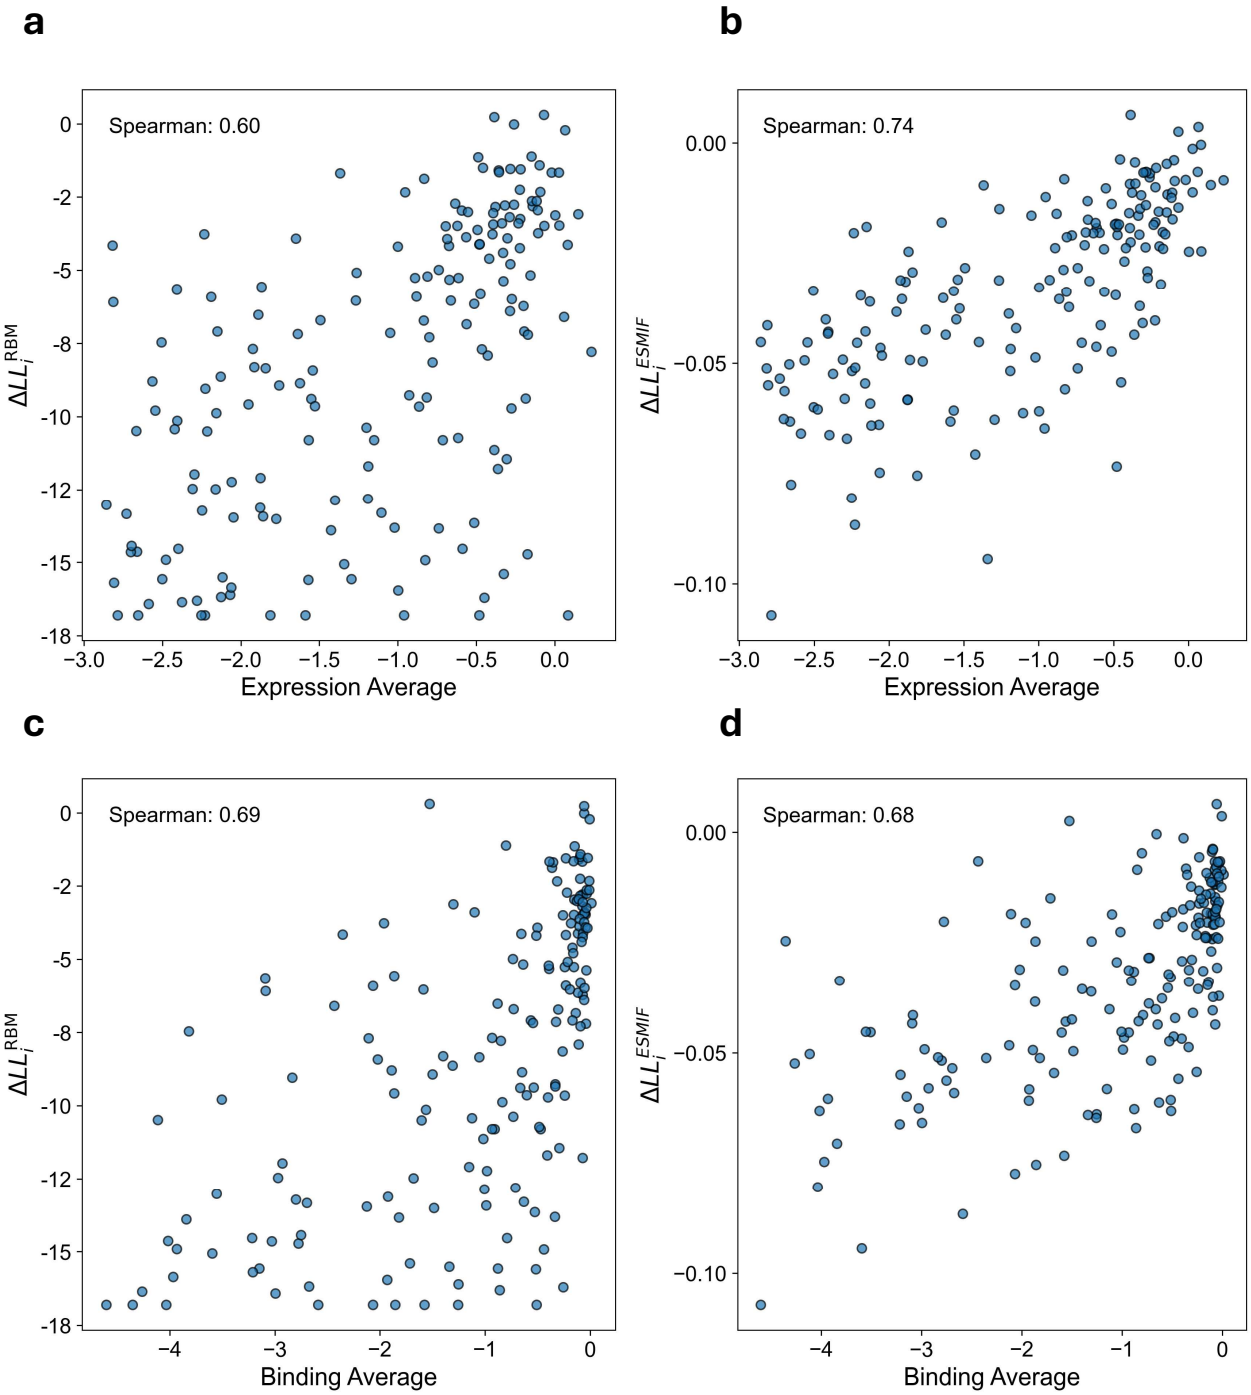

Fig. S3: Distillation of ESMIF with RBM.
(a-b) Comparison between experimentally measured changes in expression (Starr *et al.*), averaged per site for single mutants in the Wuhan background, and RBM/ESM-Inverse Folding–derived mutational scores ($\Delta \mathrm{LL}_i$). **(c-d)** Same as (a-b) with experimentally measured changes in ACE2 binding (Starr *et al.*).

Figure S3

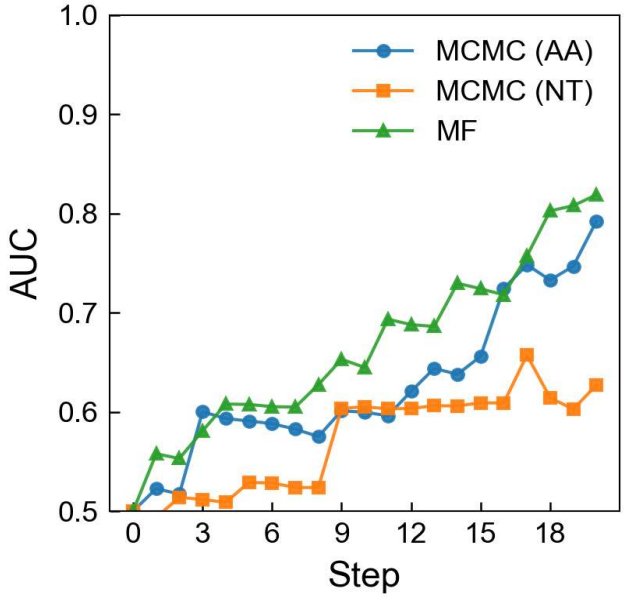

Fig. S4: Stepwise comparison of mean field and MCMC predictions of antibody binding. Area Under the Curve (AUC) for predicting highly mutable sites in GISAID (defined as sites with ≥ 2 representative amino acids). Site entropy from MF trajectory is 510 used as a predictor and compared with site entropy from MCMC samples in amino acid (AA) or nucleotide (NT) space. AUC of 0.5 represents random predictions.

Figure S4

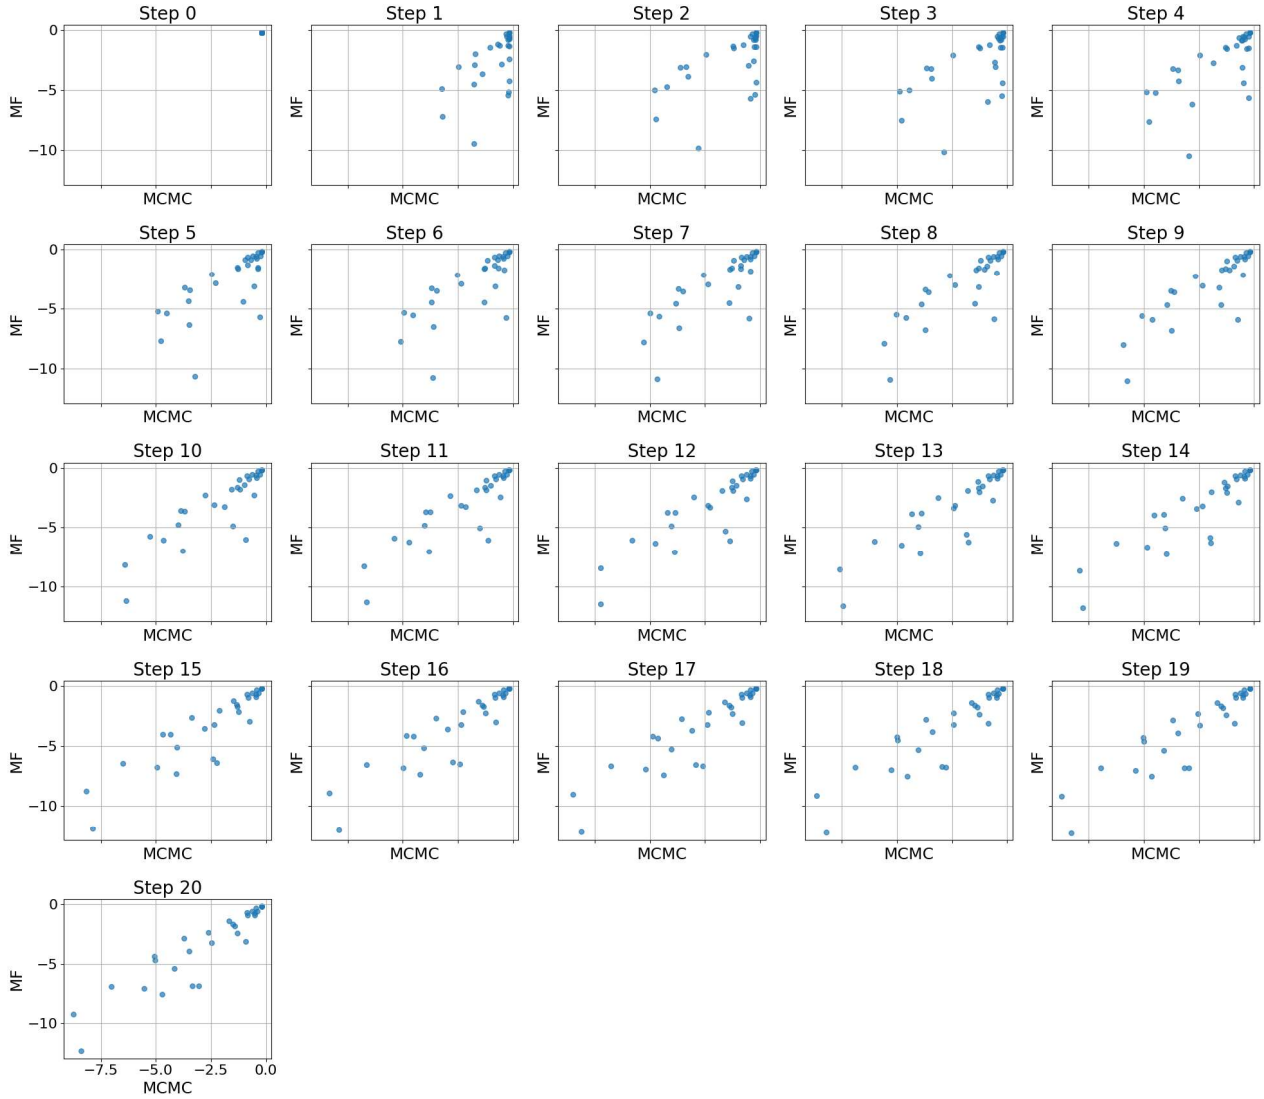

Fig. S5 Stepwise comparison of mean field and MCMC predictions of antibody binding. Correlation between antibody binding scores computed from MCMC-sampled sequences and mean field (MF) trajectory, shown for all 20 path steps. Each dot represents one antibody.

Figure S5

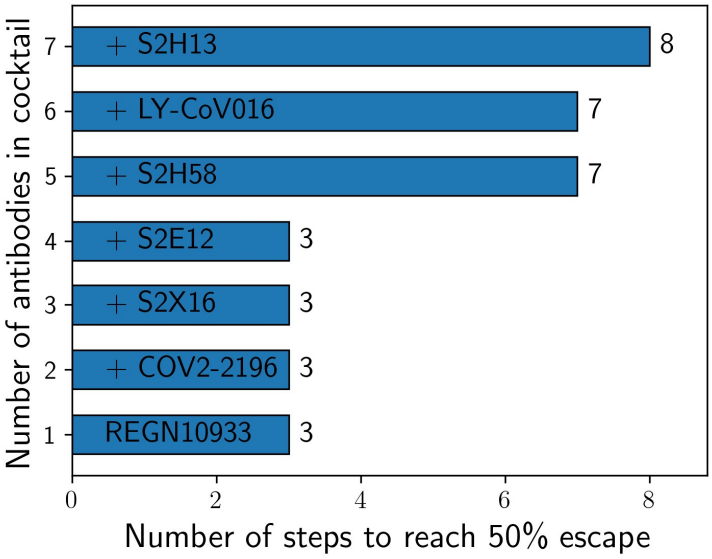

Fig. S6 : Impact of correlated antibodies in a cocktail.
Number of steps required to reach 50% escape as a function of the number of additional antibodies in a REGN10933-based cocktail. Since antibodies show high correlation, viral escape is not delayed.

Figure S6

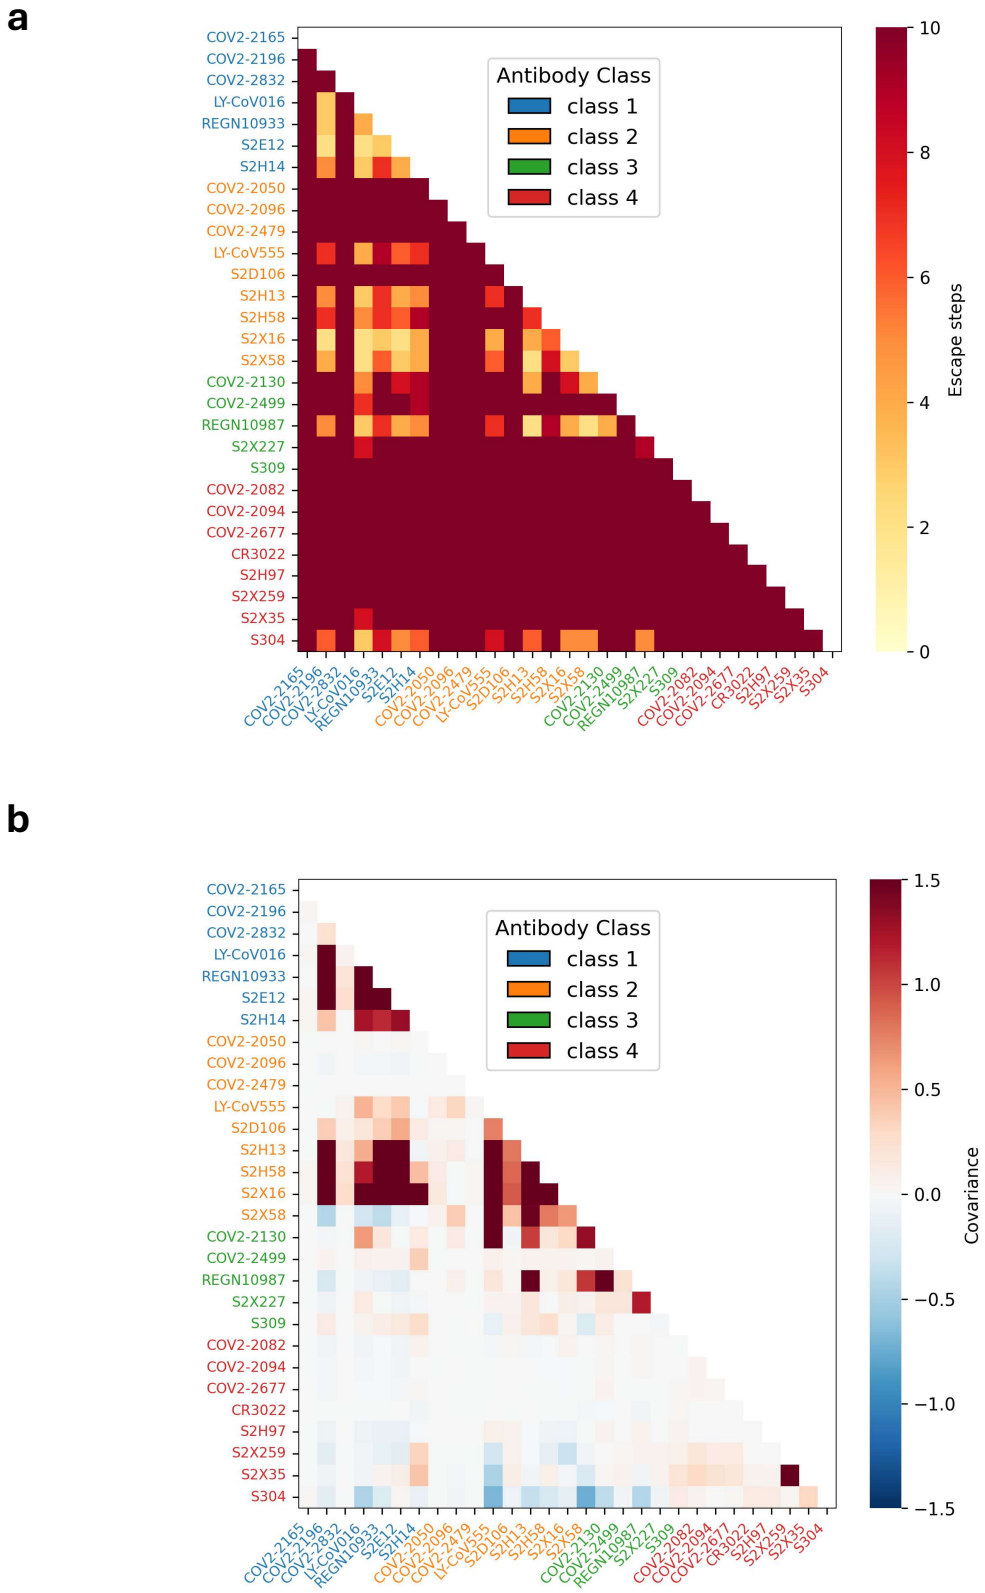

Fig. S7: Cocktails of two antibodies.
(a) Number of steps required for escape in cocktails of two antibodies. **(b)** Covariance between pairs of antibodies.

Figure S7

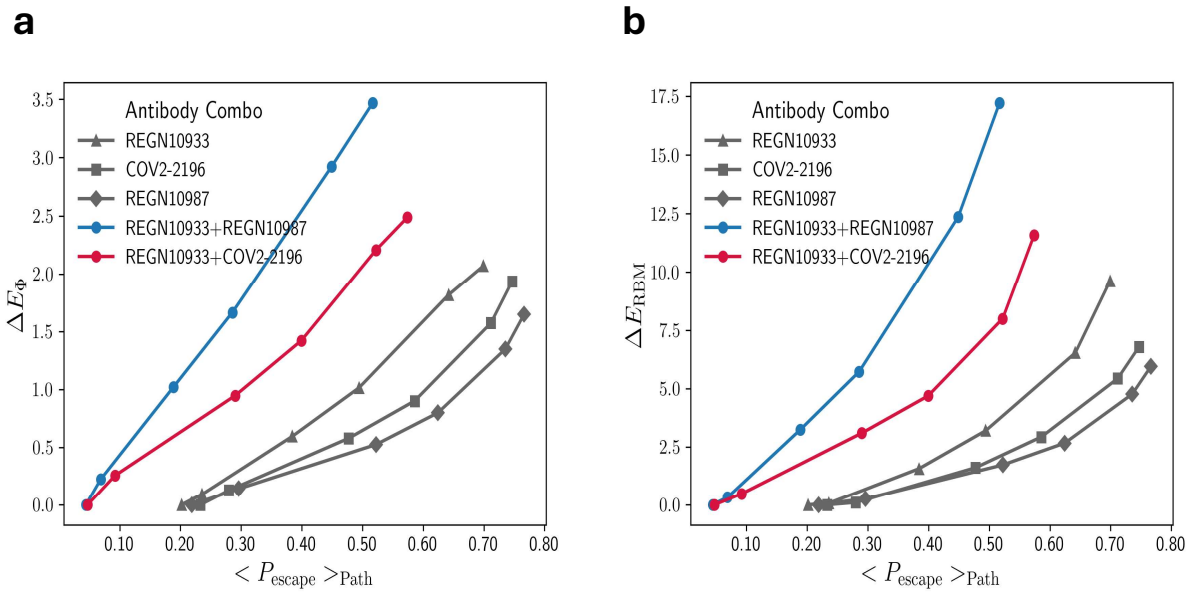

Fig. S8: Additional energy required to reach different escape levels across mean field paths.
(a) Cumulative continuity penalty (interpreted as mutational effort) required to reach increasing levels of escape, averaged over the trajectory, for each antibody combination. **(b)** Cumulative RBM energy (interpreted as folding or stability cost) required to reach corresponding escape levels.

Figure S8