

UC San Diego

UC San Diego Previously Published Works

Title

Kinetic coevolutionary models predict the temporal emergence of HIV-1 resistance mutations under drug selection pressure.

Permalink

<https://escholarship.org/uc/item/23d6z6sv>

Journal

Proceedings of the National Academy of Sciences, 121(15)

Authors

Biswas, Avik

Choudhuri, Indrani

Arnold, Eddy

et al.

Publication Date

2024-04-09

DOI

10.1073/pnas.2316662121

Peer reviewed



Kinetic coevolutionary models predict the temporal emergence of HIV-1 resistance mutations under drug selection pressure

Avik Biswas^{a,b,c,1} , Indrani Choudhuri^{a,d,1}, Eddy Arnold^e , Dmitry Lyumkis^{b,f} , Allan Haldane^{a,g,2}, and Ronald M. Levy^{a,d,2}

Edited by Peter Wolynes, Rice University, Houston, TX; received September 25, 2023; accepted February 23, 2024

Drug resistance in HIV type 1 (HIV-1) is a pervasive problem that affects the lives of millions of people worldwide. Although records of drug-resistant mutations (DRMs) have been extensively tabulated within public repositories, our understanding of the evolutionary kinetics of DRMs and how they evolve together remains limited. Epistasis, the interaction between a DRM and other residues in HIV-1 protein sequences, is key to the temporal evolution of drug resistance. We use a Potts sequence-covariation statistical-energy model of HIV-1 protein fitness under drug selection pressure, which captures epistatic interactions between all positions, combined with kinetic Monte-Carlo simulations of sequence evolutionary trajectories, to explore the acquisition of DRMs as they arise in an ensemble of drug-naïve patient protein sequences. We follow the time course of 52 DRMs in the enzymes protease, RT, and integrase, the primary targets of antiretroviral therapy. The rates at which DRMs emerge are highly correlated with their observed acquisition rates reported in the literature when drug pressure is applied. This result highlights the central role of epistasis in determining the kinetics governing DRM emergence. Whereas rapidly acquired DRMs begin to accumulate as soon as drug pressure is applied, slowly acquired DRMs are contingent on accessory mutations that appear only after prolonged drug pressure. We provide a foundation for using computational methods to determine the temporal evolution of drug resistance using Potts statistical potentials, which can be used to gain mechanistic insights into drug resistance pathways in HIV-1 and other infectious agents.

HIV | epistasis | drug-resistance mutation (DRM) | kinetic Monte-Carlo (KMC) | timeline of resistance

The HIV type 1 (HIV-1) currently infects ~40 million people worldwide. In the absence of a cure, antiretroviral therapy (ART) presents the primary treatment option (1). However, all antiretroviral drugs, including those from newer drug classes, are at risk of becoming partially or fully inactive due to the emergence of drug-resistant mutations (DRMs) (2–5). The rapid mutation rate of HIV-1 plays a major role in the failure of ARTs among infected patients, leading to DRMs occurring in response to drug selection pressure (6, 7).

The viral enzymes protease (PR), RT, and Integrase (IN), which are encoded by the *Pol* gene of HIV-1, have been the major focus of ART over the past several decades (8–14). The fitness landscape of these enzymes is determined by the combined effects of the host immune response and selection pressure from ART, and how these interplay with the proteins' structure, function, thermodynamics, and kinetics (15–20). As a result, complex mutation profiles often arise in these proteins located both near and distal from the active site (17, 21, 22). These profiles can be observed in HIV-1 patient protein sequences that are available on large public databases such as the Stanford HIV-1 drug resistance database (HIVDB) (23) and the Los Alamos HIVDB (24, 25), from which we can derive specific patterns and relationships.

Primary DRMs generally occur with a fitness penalty in viral sequences found in drug-naïve patients. The effect of a mutation, however, is dependent on the entire genetic background in which it occurs, a phenomenon known as “epistasis”. Due to epistatic interactions with the sequence background, primary DRMs can become favorable (26, 27) in sequence backgrounds in which accessory mutations accumulate, such that there is a fitness penalty for reversion to the wild-type (WT) residue that leads to evolutionary trapping or “entrenchment” of the primary mutation (28–32). There is also feedback between the appearance of primary and background mutations, with the accumulation of accessory mutations increasing the likelihood that the primary mutation also arises (“contingency”), and vice versa. In the presence of drug pressure, this leads to a complex interplay between the functions of the primary and accessory mutations (33, 34). Theoretical considerations suggest that epistasis can slow the rate of evolution by creating a rugged fitness landscape with low-fitness intermediate states forming barriers between local fitness optima (35, 36). Here, we study such phenomena using empirical data.

Significance

HIV-1 affects the lives of millions of people worldwide; cases of panresistant HIV are emerging. We use a kinetic Monte-Carlo method to simulate the evolution of drug resistance based on HIV-1 patient-derived sequence data available on public databases. Our simulations capture the reported time to acquire drug-resistance mutations (DRMs) across the major HIV-1 drug-target enzymes: Protease, RT, and Integrase. The network of epistatic interactions with a primary DRM determines its acquisition rate, which is not explained by the overall fitness of the DRM or features of the genetic code, but instead by an “epistatic barrier”. This work provides a framework for the development of computational methods that forecast the time-course over which resistance to antivirals develops in patients.

Author contributions: I.C., E.A., A.H., and R.M.L. designed research; A.B. and I.C. performed research; A.B. and I.C. analyzed data; E.A. provide helpful suggestions to improve the paper; and A.B., I.C., D.L., A.H., and R.M.L. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2024 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹A.B. and I.C. contributed equally to this work.

²To whom correspondence may be addressed. Email: allan.haldane@temple.edu or ronlevy@temple.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2316662121/-/DCSupplemental>.

Published April 1, 2024.

Studies have illustrated the effects of epistasis on the fitness landscape of HIV-1 proteins (29, 34, 37–42), but it is unclear why some DRMs are acquired rapidly while others are acquired much more slowly, and how this is influenced by the epistatic network. Recently, we introduced a kinetic model that describes the evolution of HIV-1 sequences on a fitness landscape constructed using a Potts sequence-covariation statistical energy model (43) inferred from drug-experienced patient protein sequences in the Stanford HIVDB (7, 23). This kinetic evolution model has the key feature that it models epistasis between all positions in the sequences. We have previously shown that this feature is critical for making the model numerically consistent with the observed between-host sequence mutation covariation statistics obtained from public repositories, such as the Stanford HIVDB. When simulating many long evolutionary trajectories in parallel by our method, starting from a drug-naïve sequence ensemble and collecting the final sequences, the mutational statistics (frequencies of single point and higher-order combinations of mutations) of the generated sequences match those of the drug-experienced dataset that was used to train the fitness model. To establish a baseline for understanding how epistasis affects kinetics associated with the development of drug resistance, we previously used this model to follow the kinetics of a DRM within the drug-experienced ensemble of patient protein sequences and concluded that epistasis has a strong effect on evolutionary dynamics (37). However, the emergence of drug resistance is best understood in the context of a changing environment as the virus is newly exposed to drug treatment. In the current work, we focus on the kinetics of the emergence of drug-resistance mutations (DRMs) in an ensemble of drug-naïve HIV-1 patient sequences evolving under the influence of newly applied drug pressure.

Our goal is to use the kinetic model to probe the relative times at which DRMs arise in HIV-1 under drug selection pressure. We focus on 52 DRMs in the three HIV-1 drug-target enzymes encoded by the *Pol* gene (23). These 52 DRMs are chosen based on two primary criteria: They are classified in the Stanford HIVDB as “primary” mutations that generally affect drug susceptibility ≥ 5 to 10 fold on their own (23), and they are observed in drug-experienced patients with frequencies $\geq 1\%$. The principal result of our study is that the rates at which DRMs are acquired in initially drug-naïve strains predicted using our model are highly correlated with the corresponding observed acquisition rates that are reported in the literature. This suggests that the rates at which DRMs are acquired depend strongly on epistatic interactions between the focal mutation and the other residues in the sequence background. The acquisition rates cannot be explained by the equilibrium frequency of a DRM, which is a proxy for its fitness in the drug-experienced patient population after averaging over sequence backgrounds or by features of the genetic code such as nucleotide changes (Δnuc) at the codon level, or transitions (Ti) vs. transversions (Tv). We propose that some DRMs are acquired more slowly because they face an “epistatic barrier”, and outline how this arises. The Potts model parameterized on drug-experienced HIV-1 patients, combined with kinetic Monte-Carlo (KMC) techniques, is a powerful predictor of the relative DRM acquisition times leading to drug resistance.

Results

Relative Rates at Which HIV-1 DRMs Are Acquired with a KMC Evolution Model Match the Literature. We follow the temporal evolution of a set of primary DRMs in HIV-1 PR, RT, and IN using a KMC method to evolve, in parallel, an ensemble

of initially drug-naïve consensus sequences. Our KMC model employs a coarse-grained representation of the intrahost HIV-1 evolutionary process, as described in more detail in *Computational Methods*. During HIV-1 infection, many viral variants are present in a host, but the sequences of these within-host variants are generally $\sim 99\%$ identical (44). This stands in contrast to the sequence identity between evolving viral populations in different hosts (between-hosts), which can be much lower (e.g. 90% for PR) (45). These observations suggest that a host’s viral population can be represented at coarse grain by its population consensus sequence. Indeed, in many HIV-1 sequence datasets, such as derived from the Stanford HIVDB, a population consensus sequence is effectively sampled from each host by averaging over the viral diversity within the host. We model the evolution of an ensemble of consensus sequences representing multiple host populations, approximating this process as a series of point-mutation events occurring at a constant rate in each consensus sequence, consistent with observations (46), and the mutations are either fixed or lost according to the fitness landscape inferred based on between-host sequence data. Our KMC scheme models this as a series of point-mutation proposals for each sequence, which are either accepted or rejected according to the Metropolis Criterion, analogous to fixation or loss of the mutant strain in a host viral population. While this model coarse-grains several details of the intrahost HIV-1 evolutionary dynamics, for instance, clonal competition, recombination, immune-pathogen coevolution, spatial and temporal drug heterogeneity, and nucleotide-level biases, it faithfully reproduces the observed pairwise and higher-order mutation patterns of the data.

We model the scenario in which the initially drug-naïve sequences are first exposed to drugs upon starting ART. These drug-naïve sequences have previously evolved in a drug-naïve fitness landscape, which we do not explicitly model here. Upon initiation of ART, the new environment imposes a drug-experienced fitness landscape. We model this by initializing the trajectories starting with the sequence ensemble taken from the drug-naïve sequences collected in the Stanford HIVDB and follow their evolution forward in time under the drug-experienced Potts fitness model, which was inferred using the drug-experienced sequence dataset collected in the Stanford HIVDB.

The drug-experienced Potts statistical energy model captures the sequence covariation due to selection pressure in the presence of drugs. The Potts “statistical” energy $E(S)$ of a sequence S (*Computational Methods*) predicts how sequences will appear in the dataset with probability $P(S) \propto e^{-E(S)}$ such that sequences with favorable statistical energies are more prevalent in the multiple sequence alignment (MSA), and thus more fit. This overall fitness predicted by the Potts model, which is inferred based on observed mutation prevalence, is the net result of many phenotypic contributions, including both the “replicative fitness” and “transmission fitness” commonly measured by virological assays (47–49). A key feature of the Potts model of fitness is that the effect of a mutation on $E(S)$ is “background-dependent”. This means that a mutation at one position will affect the likelihood of mutations at all other positions, both directly and indirectly, through chains of epistatic interactions involving one or many intermediate residues.

Fig. 1 illustrates how drug exposure drives the appearance of DRMs in general, as demonstrated with the enzyme RT. The distribution of the total number of mutations per sequence measured with respect to the WT HIV-1 subtype-B consensus sequence (hereafter referred to more simply as the WT) evolves from a narrow initial distribution in drug-naïve sequences with a peak at

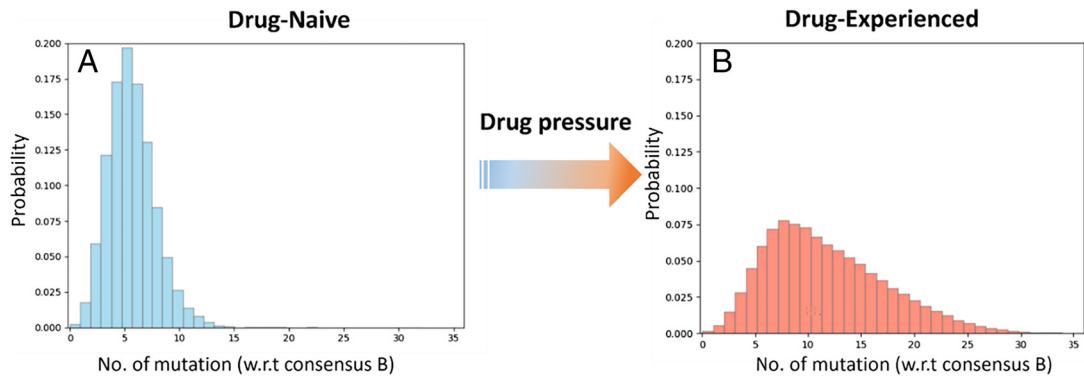


Fig. 1. The distribution of the total number of mutations in HIV-1 RT in (A) the drug-naïve MSA, and (B) the drug-experienced MSA available from the Stanford HIVDB (22, 24).

~six mutations and a maximum of ~20 mutations to a much broader distribution in drug-experienced sequences with a peak at ~12 mutations and a maximum of more than 30 mutations. Corresponding plots for PR and IN are shown in *SI Appendix, Figs. S1 and S2*. We note that IN is more conserved under drug pressure than either PR or RT.

We compute a characteristic acquisition time for each primary DRM in PR, RT, and IN from the KMC simulations using an exponential fit for the change in the frequency of the DRM starting from the initial drug-naïve ensemble and ending in the final drug-experienced state (*SI Appendix, Fig. S3*). We follow 52 primary DRMs with a wide range of acquisition times in the clinical literature: 13 protease inhibitor (PI) DRMs, 14 nucleoside-analog RT inhibitor (NRTI) DRMs, 11 nonnucleoside RT inhibitor (NNRTI) DRMs, and 14 integrase strand-transfer inhibitor (INSTI) DRMs. Fig. 2 shows the correspondence between the acquisition times to acquire these 52 major DRMs estimated from the KMC simulations, alongside the corresponding timelines reported in the literature (*SI Appendix, Fig. S3 and Tables S1–S4*).

The Spearman correlation coefficients between the KMC simulated and observed DRM acquisition times are listed in Table 1, and correlation plots are shown in Fig. 3. The average Spearman correlation coefficient for all 52 DRMs in the three HIV-1 viral enzymes is $\rho = 0.85$, $P < 0.001$ (for individual enzymes, $\rho = 0.75$, 0.92, and 0.90 for PR, RT, and IN, respectively).

The very strong correlation between the predicted and observed acquisition times for primary DRMs is noteworthy. The acquisition times for each DRM in PR, RT (NRTI-selected), RT (NNRTI-selected), and IN are shown in *SI Appendix, Tables S1–S4*). The time span to acquire DRMs is large, and the fastest DRMs are acquired ~20 times more rapidly than the slowest DRMs. To illustrate the temporal evolution of DRMs with a contrasting timeline of resistance, we divide the primary DRMs into three categories based on their acquisition times in the literature: fast, between 0 to 3 mo; intermediate, between 4 to 5 mo; and slow, more than 6 mo. The acquisition times (τ) from our simulations are correspondingly classified for each category, between ~1 to 10 for fast, ~10 to 24 for intermediate, and >24 for slow.

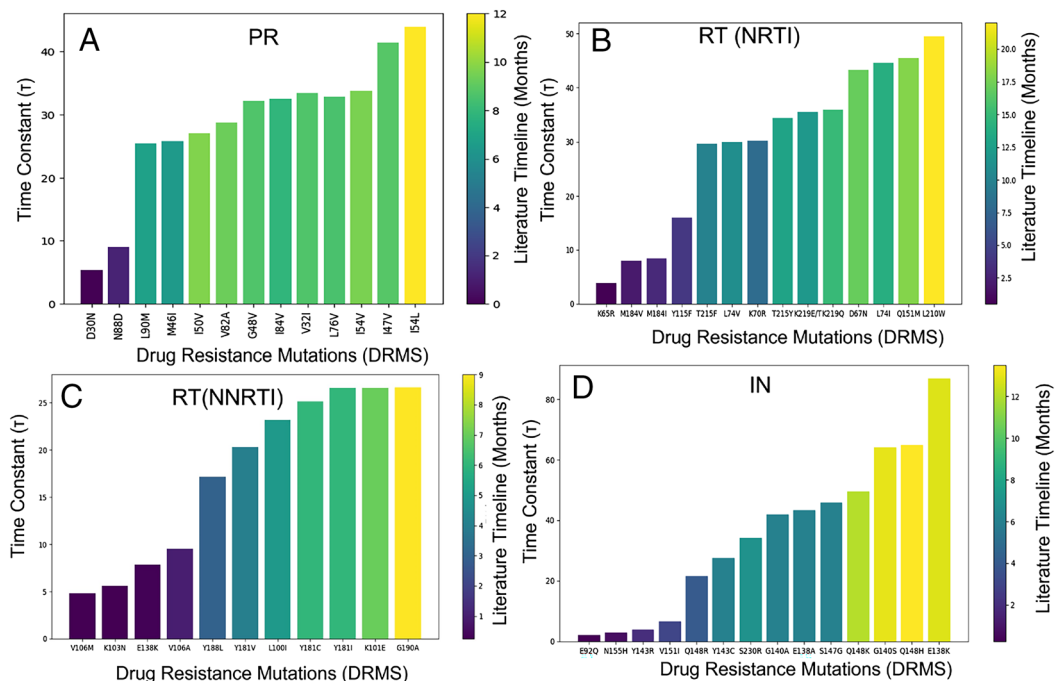


Fig. 2. The literature correspondence with the KMC acquisition time of the major (A) PR DRMs in response to PIs, (B) RT DRMs in response to NRTI, (C) RT DRMs in response to NNRTI, and (D) IN DRMs in response to INSTIs. The height of each bar represents the time constant (τ) calculated from KMC simulations. The color bar represents the time required for each DRM to emerge in the patient's population under drug selection pressure, which is collected from the literature.

Table 1. Spearman's rank correlation coefficient (ρ) test

Models	Targets	ρ	P-value
KMC-predicted time constant (τ)	All DRMs	0.85	1×10^{-22}
	PR	0.75	1×10^{-05}
	RT	0.92	1×10^{-16}
	IN	0.90	1×10^{-08}
Mutant prevalence (drug experienced frequency)	All DRMs	0.26	0.05
	PR	0.30	0.29
	RT	0.47	0.02
	IN	0.02	0.93
Genetic code (Ti vs. Tv)	All DRMs	0.13	0.33
	PR	0.18	0.54
	RT	0.23	0.25
	IN	0.15	0.50
Genetic code (Δ nuc)	All DRMs	0.14	0.34
	PR	0.23	0.45
	RT	0.11	0.57
	IN	0.29	0.15

Among the NRTI selected primary DRMs, 3 are acquired rapidly, while 10 are acquired slowly. The DRMs K65R, M184V, and M184I are acquired within ~ 1.5 to 3 mo after initiation of therapy with KMC time constants ranging between ~ 4 to 8 KMC time units, whereas T215F/Y, L74V/I, K70R, K219E/Q, D67N, and L210W are acquired between 9 to 20 mo after initiation of therapy

with KMC time constants between 30 to 50 KMC time units (*SI Appendix, Tables S2 and S3*). The relative times required for acquisition of the major drug resistance mutations in PR and IN, (*SI Appendix, Tables S1 and S4*) are also recapitulated by the KMC simulations. We conclude from the results summarized in *SI Appendix, Fig. S3 and Tables S1–S4* that the KMC model of evolution of drug resistance mutations in RT captures fundamental kinetic features of the mutational landscape for this enzyme evolving under drug pressure, which distinguishes DRMs that are acquired slowly from those acquired rapidly, as reported in the literature. The KMC simulations correspond to a fitness-based coarse-grained epistatic model of the HIV-1 evolutionary process under drug selection pressure. We have found that despite this coarse-graining, the model is an excellent predictor of DRM acquisition times.

Slowly Acquired DRMs Have an Epistatic Barrier to Resistance.

We next investigate the reasons why some DRMs arise fast whereas others arise slow under our model. We find that the acquisition time is best explained by an epistatic barrier effect. Specifically, DRMs that are acquired slowly are contingent on the acquisition of accessory mutations that must arise first as the drug-naïve ensemble evolves under newly applied drug selection pressure, while DRMs that are acquired rapidly have higher fitness after drug pressure is applied even before accessory mutations accumulate. Regardless of whether the DRMs are acquired rapidly or slowly, accessory mutations are eventually acquired, and they function to trap or entrench the DRMs (*SI Appendix, Fig. S5*) in

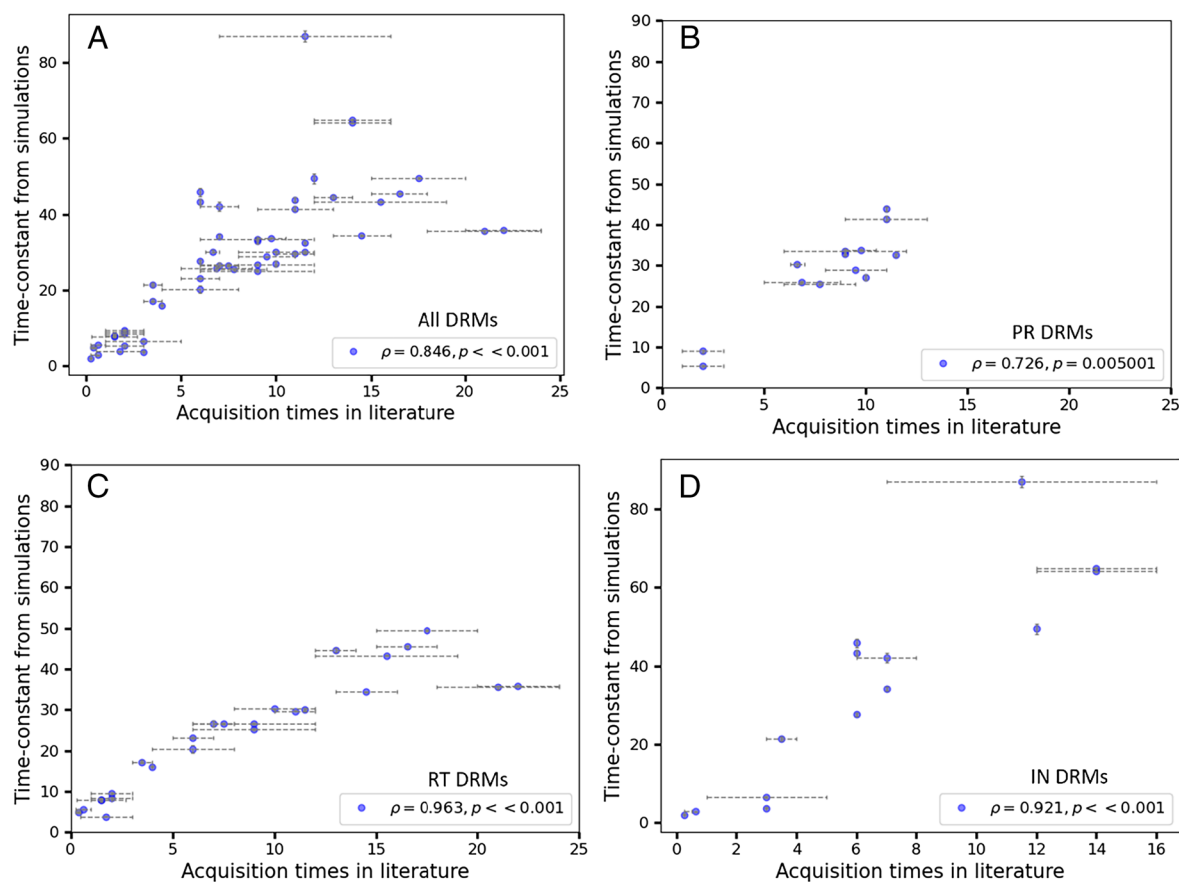


Fig. 3. Spearman rank-order correlations between acquisition times reported in the literature and our KMC simulations for DRMs. (A) All mutations, (B) PR (C) RT (RT/NRTI+NNRTI) and (D) IN. The gray dashed lines along the abscissa axis indicate the ranges in acquisition times reported the literature when multiple literature reports are available; along the ordinate axis represent the errors in the obtained computational time constants from the best fit curves in *SI Appendix, Fig. S3*. Sampling alternate sets of acquisition times uniformly from within the ranges gives largely similar correlation coefficients (mean Spearman's rank-order correlation coefficient = 0.83, 0.72, 0.95, 0.93 for panels A–D, respectively from 1,000 samples).

a subset of the drug-experienced sequence ensemble wherein they appear with high probability (34).

In Table 1, we first show that acquisition time is not explained by various simpler hypotheses. We carried out tests to determine whether acquisition times are affected by other features that have been suggested to affect these times but are not explicitly accounted for in our model, including the drug-experienced frequency of the DRM, biases due to kinetic effects associated with nucleotide transition vs. transversion rates, or differences in the rates of codon changes due to single vs. double-nucleotide mutations. From the low correlation between these properties and the times to acquire corresponding DRMs reported in the literature (Table 1), we conclude that these are not major determinants of the rates at which drug resistance is acquired, in contrast to the KMC-predicted time constant, which is highly correlated. Notably, even if a DRM has high fitness in the sense that it ultimately rises to high frequency after drug exposure, this does not imply it is a fast DRM, because the drug-experienced “mutation prevalence” is poorly correlated with reported acquisition time. This suggests that more complex epistatic effects control DRM acquisition times.

To measure how epistatic constraints change as the sequences evolve, we define the “adaptive-frequency”, $\theta_\alpha^i(S)$, which reflects the likelihood of a DRM α to occur at position i in a specific sequence S , and differs for each sequence due to epistasis. In other words, if residue positions other than i were to be held fixed in sequence S , the adaptive-frequency $\theta_\alpha^i(S)$ measures the equilibrium frequency at which position i will have mutated to residue α in the drug-experienced fitness environment. It can be thought of as a proxy for the viral fitness of the DRM in that sequence background due to a combination of selection pressures including drug-pressure and epistatic constraint. We define it as:

$$\theta_\alpha^i(S) = \frac{e^{-E(S_\alpha^i)}}{\sum_\beta e^{-E(S_\beta^i)}}, \quad [1]$$

where S_α^i is the sequence S with position i mutated to character α , $E(S)$ is the Potts statistical energy of sequence S parameterized on the drug-experienced patient sequence data, which includes both position-dependent “field” and epistatic “coupling” terms between the focal position and all other positions; the lower sum runs over all possible residues at position i .

Eq. 1 gives $\theta_\alpha^i(S)$ for a specific background S , and we also compute the mean adaptive-frequency $F_\alpha^i(t) = \langle \theta_\alpha^i \rangle_t$ over all sequences in the ensemble at time t . The mean adaptive-frequency of a DRM only changes over time due to changing epistatic constraints caused by other mutations arising in the sequence backgrounds. As a limiting case at long times after all the residue positions equilibrate, the mean adaptive-frequency must equal the equilibrium frequency of the mutation observed in the Stanford HIVDB of drug-experienced patient sequences, that is: $F_\alpha^i(\infty) = f_\alpha^i(\infty)$ where $f_\alpha^i(t)$ is the frequency of residue α at position i in the ensemble at time t , i.e. the average over the indicator function $f_\alpha^i(t) = \langle \delta_\alpha^i \rangle_t$, which at $t = 0$ equals the frequency in the drug-naive dataset and at $t = \infty$ that in the drug-experienced dataset. Note that at $t = 0$, $F_\alpha^i(0) \neq f_\alpha^i(0)$.

In Fig. 4, we track how the adaptive-frequency of various DRMs changes over time due to changing epistatic constraints, averaged over an ensemble of trajectories evolving under the drug-experienced Potts statistical potential from the drug-naive initial state. We simultaneously track each DRM’s frequency over the same ensemble, computed by averaging an indicator function

that is 1 when the DRM is present and 0 otherwise. We empirically observe that the DRM frequency initially quickly changes to track its adaptive-frequency, which itself changes more slowly as mutations accumulate in the sequence background, such that at long times $F_\alpha^i(\infty) = f_\alpha^i(\infty)$, as expected. The fact that the DRM frequency depends on the slowly changing epistatic constraint indicates the latter a major controlling factor of DRM acquisition.

To illustrate this, we contrast a “slow” DRM and a “fast” DRM for each target protein. These include D30N (fast) and V32I (slow) in HIV-1 PR in response to PIs, M184V (fast) and D67N (slow) in RT in response to NRTIs, K103N (fast) and Y181C (slow) in RT in response to NNRTIs, and N155H (fast) and Q148H (slow) in IN in response to INSTIs. The fast/slow pairs were selected such that their equilibrium frequencies in the initial (drug-naive) and final (drug-experienced) states observed in the Stanford HIVDB are approximately the same. We refer the reader to *SI Appendix, Tables S1–S4* for a list of the 52 DRMs that we modeled, together with their drug-naive and drug-experienced frequencies, and their adaptive-frequency following the application of drugs but before any additional mutations can accumulate.

We observe that the fast DRMs have high initial adaptive-frequency, close to their final equilibrium frequency, in contrast to slow DRMs whose adaptive-frequency is initially relatively low but gradually rises as the epistatic constraint is relaxed and the sequence backgrounds evolve. We refer to this phenomenon as an epistatic barrier which we find is a central feature of all the DRMs that are acquired slowly. For instance, the fast PR D30N mutation is initially present at less than 1% frequency in the drug-naive ensemble at $t = 0$ but has a high initial mean adaptive-frequency of $F_N^{30}(0) = 6\%$, close to its final equilibrium drug-experienced frequency of 8% (Fig. 4A). Its frequency, i.e., its indicator function average, quickly increases to match its adaptive-frequency, and both subsequently increase by a small amount to their final value of $F_N^{30}(\infty) = f_N^{30}(\infty) = \sim 8\%$ as additional mutations arise in the sequence backgrounds. In contrast, the slow PR V32I mutation initially has a low mean adaptive-frequency of $F_I^{32}(0) < 1\%$ after drug pressure is applied, relative to its final equilibrium frequency of $\sim 5\%$, and both its frequency and adaptive-frequency slowly rise to this much larger final value. This suggests the slow DRMs are contingent on the appearance of other mutations in the sequence background following the application of drug pressure. The same behavior is observed for other protein targets in Fig. 4 B–D.

In *SI Appendix*, we show that the ensemble average behavior of the adaptive-frequency can also be understood in terms of Potts energy values $E(S)$, by examining how the distribution of DRM fitness effects ΔE (*SI Appendix, Fig. S4*) of the sequence ensemble changes over time. Here ΔE is the Potts energy difference between a sequence with the DRM and the same sequence with the WT residue at the DRM position. A concise summary of the change in ΔE distributions under the drug experienced Hamiltonian for the fast DRM M184V is tabulated in *SI Appendix, Table S5*, corresponding to adaptive-frequency values shown in Fig. 4B.

The analysis above leads us to define a “constraint ratio”, which measures the fraction of the total change in frequency of a DRM that is explained by the relaxation of epistatic constraint. We expect DRMs with a high constraint ratio will be slow to acquire. The results in Fig. 4 show how each DRM’s overall change in frequency upon drug exposure can be understood as a sum of two effects: 1) The new selective pressures of the drug exposed environment cause an immediate initial increase in DRM adaptive-frequency

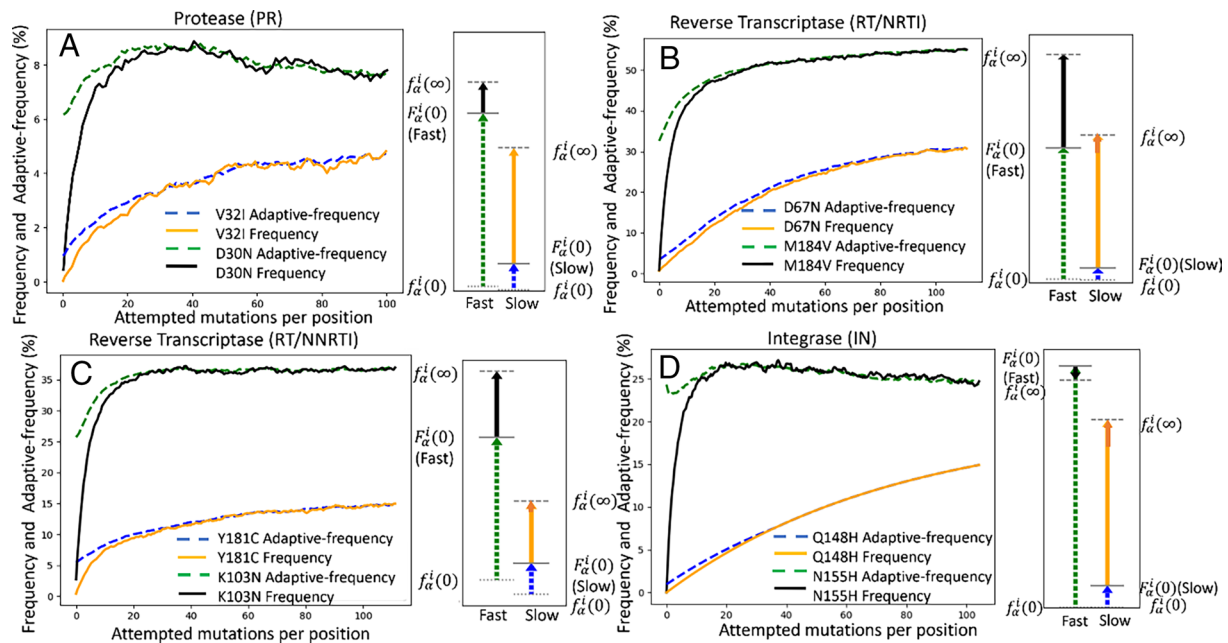


Fig. 4. Temporal evolution of the mean adaptive-frequency $F_{\alpha}^i(t)$ and frequency $f_{\alpha}^i(t)$ of a mutation as a function of the number of attempted mutations per position for two differently evolving DRMs. The panels refer to (A) D30N (fast) and V32I (slow) DRMs in PR, (B) M184V (fast) and D67N (slow) DRMs in RT in response to NRTIs, (C) K103N (fast) and Y181C (slow) DRMs in RT in response to NNRTIs, and (D) N155H (fast) and Q148H (slow) DRMs in IN in response to INSTIs. The arrows in the *Right* side panel of each plot illustrate schematically the initial adaptive-frequency (Initial increase) relative to initial drug-naïve frequency $f_{\alpha}^i(0)$ immediately after exposure to drugs at time 0 (dashed green/blue arrow for fast/slow) and the subsequent increase (black/orange solid arrow for fast/slow) from this value to the final value ($f_{\alpha}^i(\infty)$) after the background has had time to equilibrate. The initial adaptive-frequency $F_{\alpha}^i(0)$ for the faster-evolving DRMs is much closer to the final drug-experienced frequency (dashed green arrows are larger than the dashed blue arrows). The darker orange arrow at the tip of the orange line shows the additional increase in frequency for the slow DRMs after 100 attempted mutations per position. Further discussion about the position-specific contributions of the additional mutations necessary for a focal DRM to occur are discussed in Fig. 5.

compared to its frequency in the drug-naïve environment. This is represented by the dashed green/blue arrows in Fig. 4, equal to $F_{\alpha}^i(0) - f_{\alpha}^i(0)$. This indicates how, before any other mutations occur in the sequence background, many DRMs would increase in frequency upon exposure to drugs, for instance by providing resistance to the drug directly, even if the stability of the target protein is adversely affected. 2) Subsequently, there is a further increase in adaptive-frequency or fitness caused by the relaxation of epistatic constraint through mutations evolving in the sequence background. This is represented by the solid black/orange arrows in Fig. 4, equal to $f_{\alpha}^i(\infty) - F_{\alpha}^i(0)$. A key distinguishing feature of the fast vs. slow DRMs is that for slow DRMs a large fraction of the overall change in frequency (dashed arrow + solid arrow) at long times is accounted for by the relaxation of epistatic constraint (solid arrow). Indeed, we find this constraint ratio, $(f_{\alpha}^i(\infty) - F_{\alpha}^i(0)) / (f_{\alpha}^i(\infty) - f_{\alpha}^i(0))$, strongly correlates (SI Appendix, Fig. S6) with the time constant τ computed from the kinetic simulations. The fact that this ratio, which measures the effect of epistatic constraints induced through the coupling terms of the Potts model, correlates well with literature-reported DRM acquisition time supports our hypothesis that an epistatic barrier is the major determinant of DRM acquisition times.

We also tested whether simpler measures of DRM fitness are correlated with DRM acquisition time, which are presented in SI Appendix, Table S6. However, we did not find strong correlations between any other measures and the DRM acquisition time. These include the final mutation prevalence of the DRM, $f_{\alpha}^i(\infty)$, or equivalently its final adaptive-frequency $F_{\alpha}^i(\infty)$, with correlation 0.26, the overall increase in frequency of the DRM reflecting its net change in fitness, $f_{\alpha}^i(\infty) - f_{\alpha}^i(0)$, with correlation 0.24, or its initial adaptive-frequency $F_{\alpha}^i(0)$ with correlation 0.36. Unlike the

constraint ratio, these values do not measure how much of the DRM's increase in frequency is due to relaxation of epistatic constraint, which appears to be the major determinant of whether a DRM is fast or slow. Fig. 4 suggests why this is the case. The DRM frequency evolves to quickly track the adaptive-frequency, and we find the adaptive-frequency only changes gradually as mutations in the sequence background reduce the epistatic constraint, so DRMs where this subsequent gradual change is most important will be slower to acquire.

Position-Specific Contributions to the DRM Fitness. We next demonstrate how the Potts kinetic model can be analyzed to determine which patterns of mutations cause the epistatic barrier. Taking advantage of its simple and interpretable form, we define a score to estimate the contribution of each background position to the epistatic barrier for a focal DRM. For a given sequence, we compare the adaptive-frequency for the DRM at position i in that sequence to the expected adaptive-frequency if background position j were mutated to other residues in proportion to which they appear in the equilibrated (drug-experienced) sequence ensemble. This score is defined as

$$\Delta\theta_{\alpha}^{ij}(S) = \left[\sum_{\beta} f_{\beta}^j(\infty) \theta_{\alpha}^i(S_{\beta}^j) \right] - \theta_{\alpha}^i(S), \quad [2]$$

where i is the DRM position, α the DRM mutation, j the background position, and f_{β}^j is the frequency of residue β at position j in the drug-experienced ensemble.

This measures the average change in the sequence-specific adaptive-frequency at position i from its initial value caused by letting only position j evolve to the drug-experienced ensemble

frequencies, keeping all other residues in the sequence fixed. In this way, only positions j , which are both epistatically coupled to position i and have mutations β that arise during evolution to the drug experienced ensemble, will have significant $\Delta\theta(S)$ values. Further, we average $\Delta\theta_{\alpha}^{ij}(S)$ over the drug-naïve sequence ensemble to give a “change in fitness” score $\langle\Delta\theta_{\alpha}^{ij}\rangle_{naive}$. The sum of this score over all positions j , $\sum_j \langle\Delta\theta_{\alpha}^{ij}\rangle_{naive}$ will approximate the net change in adaptive-frequency of the DRMs at position i due to its direct coupling to all positions j as the ensemble evolves under drug pressure to its drug-experienced values (solid arrows).

The coupled mutations identified with Eq. 2 are shown in Fig. 5 for pairs of fast and slow DRMs from each protein target. The additional coupled mutations are largely consistent with the literature, in that numerous studies identify them as being associated with each focal DRM. This observation shows that the Potts predicted epistatic barrier can be rationally decomposed and additionally implies that such analyses can be used to quantitatively identify and study novel couplings that were previously missed. *SI Appendix, section S1* contains a detailed discussion about the literature survey pertaining to every pair of fast and slow mutations

featured in Fig. 5, along with an additional discussion on the RT/NRTI Q151M complex (*SI Appendix, Fig. S7*). This decomposition is not only broadly consistent with the literature but also quantifies the relative strength of the coupled interactions for each DRM and specifically identifies “directly” coupled positions as opposed to mutations indirectly coupled through an epistatic network.

Structural Underpinnings for the Rate of Emergence of DRMs.

Our study highlights the central role epistasis plays in kinetics leading to major DRMs in HIV, which can span a large time range. However, these predictions do not explain the mechanistic origins for different rates of emergence of DRMs. Here we explore a structural rationale for the differential fitness of resistance mutations in a drug-naïve background under drug selection pressure. We include examples of fast and slow mutations from each of the four drug classes. The analyses suggest a general principle whereby faster mutations induce changes that are less disruptive and can be more readily compensated, whereas slower mutations typically induce more disruptive changes and/or lead to drug excision through indirect mechanisms.

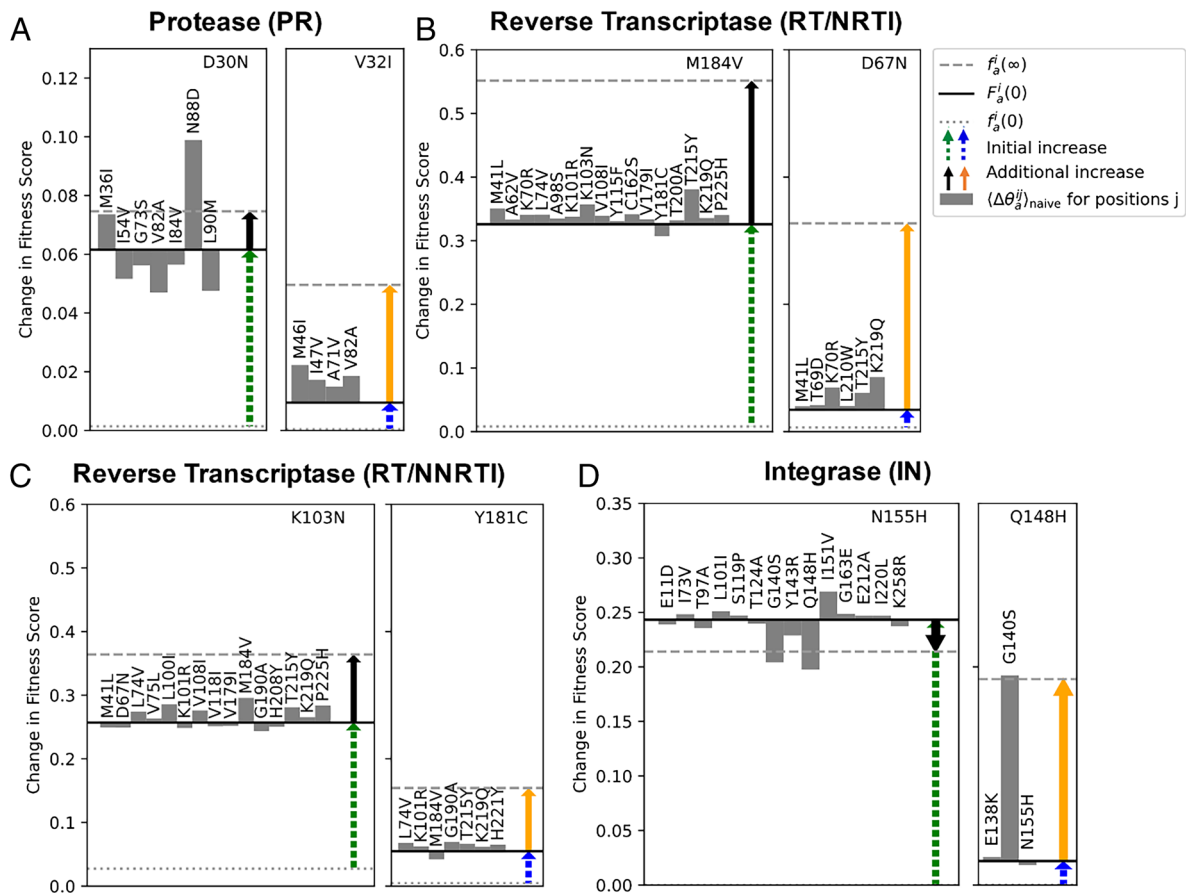


Fig. 5. Comparison of the additional mutations necessary for a focal DRM to occur when comparing DRMs with fast and slow acquisition times. The DRMs shown are (A) PI DRMs in PR, for D30N (fast) and V32I (slow), (B) NRTI DRMs in RT for M184V (fast) and D67N (slow), (C) NNRTI DRMs in RT for K103N (fast) and Y181C (slow), and (D) INSTI DRMs in IN for N155H (fast) and Q148H (slow). In each plot, the focal DRM is listed in the *Upper Right*. The dotted gray line reflects the DRM’s frequency in the drug-naïve ensemble, the solid black line reflects the DRM’s average fitness under the drug-experienced Hamiltonian in the drug-naïve ensemble, and the dashed gray line reflects the DRM’s frequency in the drug-experienced ensemble. The dashed green/blue arrow then represents the initial increase in the DRM’s frequency upon exposure to drugs starting from the drug-naïve state if the background were held fixed, identical to the dashed green/blue arrows in Fig. 4, and the solid black/orange arrow represents the additional increase in DRM frequency once the sequence backgrounds are allowed to vary and accumulate coupled mutations, identical to the solid black/orange arrow in Fig. 4. The gray bars represent the fitness measured using Eq. 2, averaged over the drug-naïve ensemble, giving the first-order contribution to the increase in the fitness of the DRM at i due to the evolution of j to the drug-experienced state. The sum of the gray bar magnitudes approximates the length of the solid black/orange arrow. For the fast DRMs, the effect of the initial increase in fitness (dashed green arrow) is larger than the additional increase in fitness due to coupled mutations (dashed blue arrow).

Fig. 6A shows the changes in the PR active site evolving in response to PIs. The fast-arising mutation D30N leads to the loss of a crucial hydrogen bond formed between the inhibitor and PR, which immediately reduces binding affinity and engenders drug resistance. Although the removal of negative charge also results in altered electrostatic interaction with the native substrate, and reduced proteolytic activity, this is easily overcome by the compensatory charge swap mutation N88D, which also arises quickly (50). D30N, in conjunction with N88D, rapidly emerges and mutually entrench each other. In contrast to the fast-arising D30N (and N88D), the slow-arising mutation V32I works through an indirect mechanism. V32 resides at the periphery of the PR active site and makes an important network of hydrophobic interactions with residues I47 and V82 that frequently arise in concert with V32I, all of which are required for PR catalytic activity (51). The mutation V32I causes an extensive pattern of rearrangements that ultimately results in repositioning of the inhibitor. PR variants containing V32I mutations frequently display comparatively larger dynamic fluctuations, which propagate throughout the enzyme, in comparison to other PR variants. For example, as many as 12 hydrogen bonds can change significantly in comparison to the WT enzyme (52). Thus, the fast-arising mutation D30N directly affects ligand binding, whereas V32I works through an indirect mechanism.

Fig. 6B shows the changes in the RT active site evolving in response to NRTIs. The fast-arising M184V DRM in RT results in fewer disruptive changes to the enzyme and acts directly to

displace bound ligand. All approved NRTIs lack a 3-OH and, when incorporated into the nascent DNA primer strand by RT, act as chain terminators. The M184V mutation replaces a flexible side chain near the polymerase active site with a branched amino acid that selectively discriminates against NRTIs, while still allowing for the incorporation of dNTPs with normal deoxyribose rings. Thus, M184V directly displaces the NRTI, but has minimal effect on normal enzyme activity. The slow-arising D67N mutation, however, resides in a different location—at the tip of the flexible 3 to 4 loop with its side chain facing the adenosine triphosphate (ATP)—and typically arises in combination with other mutations. D67N retains a similar size but eliminates the negatively charged environment imparted by the original Asp67. To compensate for this change in the electronic environment, RT must acquire additional subtle and interconnected background mutations that would remain conducive to ATP binding while allowing the mutant enzyme to excise a broad array of NRTIs (54). Thus, D67N arises more slowly, due to the requirement of developing compensatory background mutations that increases the fitness substantially above the drug-naïve value with the D67N mutation.

Fig. 6C shows the changes in the RT active site evolving in response to NNRTIs. The fast-arising K103N DRM leads to a novel hydrogen bond between N103 and Y188, which is otherwise absent in WT RT. The protein interaction network surrounding the newly formed hydrogen bond stabilizes the closed-pocket conformation of the enzyme, thus impeding NNRTI access to the binding pocket (55, 56). Notably, K103N induces minor changes

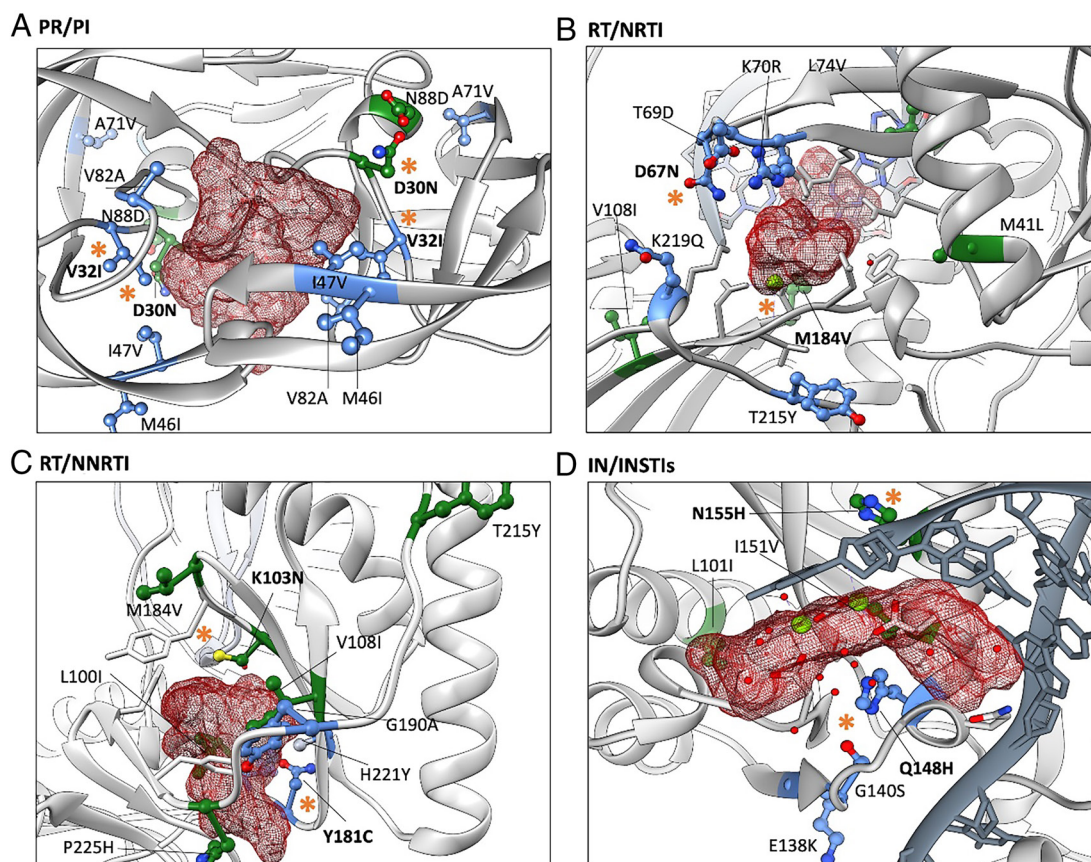


Fig. 6. The structural underpinnings for the rate of emergence of DRMs are shown for all four classes of drugs. Fast (green) and slow (blue) mutations are shown for (A) HIV-1 PR against PIs (PDB: 7DPQ, 4Q1X), (B) RT against NRTIs (PDB: 6UJY), (C) RT against NNRTIs (PDB: 3BGR), and (D) IN against INSTIs (PDB: 8FNP). The extent of the known molecular envelope for different bound drugs is shown as a red mesh (more details in the *Computational Method* section). The slow and fast mutations being followed are indicated in bold and orange asterisks (*). The other mutations involved in pathways with the slow or fast mutations are shown in the same color code (green and blue). Where mutant structures are not available, mutations are introduced through the UCSF Chimera (53) Structure Editing module and the highest probable rotameric configurations of the mutation side chains, based on previously determined atomic structures, are shown.

to the pocket residues compared to WT RT. However, resistance due to the slow-arising Y181C is more disruptive to the binding pocket. Y181C abrogates the π - π stacking interactions between two aromatic rings of residues on RT (Y181, Y188) and an aromatic ring of bound NNRTIs (55, 57, 58). Changes associated with the binding pocket in RT are more extensive in response to Y181C mutations, and thus the enzyme benefits from the addition of specific additional background mutations. Again, the slow development of Y181C can be explained by the requirement to develop additional compensatory changes.

Fig. 6D shows the changes in the IN active site evolving in response to INSTIs. The fast-arising mutation N155H leads to a salt-bridge interaction with the vDNA phosphate, which was hypothesized to affect the kinetics of INSTI binding (59). Enzyme activity is minimally affected by the N155H mutation. In contrast, the slow-arising mutation Q148H is well-known for its detrimental effect on enzyme activity. The mechanism of resistance for Q148H can be explained by the introduction of an electropositive moiety underneath the two Mg^{2+} metal ions, weakening metal chelation and leading to INSTI displacement (53, 60). Importantly, Q148H also significantly compromises enzyme activity, because the Mg^{2+} ions are also directly involved in catalysis. Since Q148H leads to a more extensive modulation to the structure by itself, this DRM also leads to a more substantial fitness cost in the drug-naive viral population than N155H. To account for the greater drop-in fitness associated with Q148H, the key compensatory G140S mutation must evolve to restore replicative capacity, while other mutations frequently accumulate with this G140S/Q148H pair. Despite the very different timescales for DRM emergence, both N155H and Q148H are the two most frequently encountered IN mutations in the Stanford drug resistance database (61), indicating that the rate of emergence of the DRM is not necessarily correlated to its final frequency in the population; the final frequencies depend heavily on the background.

To generalize the analyses from structural biology, there are two scenarios that discriminate fast vs. slow mutations in the context of drug binding, which may affect either [1] ligand binding or [2] enzymatic activity. In the first scenario, assuming enzymatic activity remains constant, fast mutations will affect the ligand binding directly, whereas slow mutations may work through indirect mechanisms and must accumulate in conjunction with other changes that eventually displace the ligand. In the second scenario, assuming that their effect on ligand binding remains constant, fast mutations will generally have a smaller effect on enzyme fitness, whereas slow mutations lead to more profound detrimental changes that affect the natural function of the enzyme and must therefore be compensated by additional background changes. We note that these scenarios are not mutually exclusive, and most cases are likely to be explained by a combination of these effects.

Conclusion

The evolution of HIV-1 under drug pressure and internal epistatic constraints induces correlated mutations that change the frequencies at which DRMs appear in the population over time. Literature surveys show that the timeline of emergence of DRMs from the drug-naive patient population varies from a few weeks to a year or more when these patients receive ART. In the present study, we modeled the kinetics of the emergence of DRMs using KMC simulations on a fitness-landscape described by an epistatic Potts statistical potential parameterized on drug-experienced sequence data. We propagate an initial sequence ensemble that matches the patterns observed in the drug-naive population as it evolves to a

final ensemble that matches the patterns observed in the drug-experienced ensemble.

We selected 52 DRMs from three different protein targets (PR, RT [NRTI and NNRTI], IN) to study their kinetics by calculating the acquisition times (τ) and compared them with the timeline of emergence reported in the literature. The times to acquire DRMs predicted by the KMC model are highly correlated with the acquisition times reported in the literature ($\rho = \sim 0.85$, $P < 0.001$). Qualitatively, for DRMs that are reported in the literature as acquired rapidly (emergence time range between 0.5 to 3 mo), the predicted KMC time constants are $\tau < 10$; while for DRMs reported in the literature as acquired slowly (~ 8 to 20 mo), the KMC time constants are $\tau > 24$. These results provide strong evidence in support of the role of epistasis—the couplings between DRMs—as the determining factor, which distinguishes DRMs that are acquired rapidly from those that are acquired slowly.

We introduced a score, the sequence-dependent adaptive-frequency $\theta_{\alpha}^i(S)$ as a proxy for fitness, which is a measure of the likelihood of a DRM in a fixed background S under the drug-experienced Potts statistical energy model, and the corresponding ensemble averaged adaptive-frequency $F_{\alpha}^i(t)$. The most important feature that distinguishes DRMs that are acquired slowly from those that are rapidly acquired is the initial fitness of the DRM in the sequence when drug pressure is applied before additional mutations have time to accumulate. For DRMs acquired rapidly, the frequency begins to increase within the drug-naive ensemble as soon as drug pressure is applied, without the need to first accumulate additional mutations. In contrast, for DRMs acquired slowly, the frequency increases in tandem with the accumulation of additional mutations. We interpret this as the existence of an epistatic barrier for the slow DRMs, consistent with results suggesting epistasis reduces evolutionary rates (35, 36). In contrast, two other factors that have been suggested to have a large effect on the rates at which DRMs are acquired, including within host effects associated with the genetic code [transitions (Ti) vs. transversions (Tv), and the number of nucleotide changes (Δ nuc) required per codon change] and the overall fitness of the DRM as estimated by its prevalence in the drug-experienced population, are not well correlated with the times to acquire drug resistance reported in the literature. Rather the distinction between fast DRMs and slow DRMs arises from the relative contributions of the two effects: a fast initial increase in the DRM frequency and a (observed) significantly slower relaxation of the sequence background. If a large fraction of the overall frequency change of the DRM is explained by the gradual change in the epistatic constraints as measured by the constraint ratio, then the DRM is slow. In contrast, if the gradual change in epistatic constraints only amounts to a small fraction of the total change in frequency of the DRM, then we expect the DRM will be fast to acquire, with a lack of epistatic barrier. Specifically, in the fast case, the sequence backgrounds are already more fit for acquisition of the DRM, but in the slow case the DRM cannot rise sufficiently to approach close to its final frequency until the sequence backgrounds change to relax the epistatic constraint.

This work provides a framework for the development and application of computational methods to forecast the time course and the pathways over which drug resistance to antivirals develops in patients. We envision that a more informed picture could be developed using Potts models, to classify drugs in silico by how difficult it will be for HIV to evolve resistance, through the identification and analysis of specific pathways required for the acquisition of compensatory sets of mutations. Currently, the Potts model

cannot be used to predict frequencies of mutation(s) that have never been observed in any dataset under drug selection pressure, their acquisition times, or the corresponding mutational pathways. However, Potts models based on in silico computed mutational effects can be utilized to address this shortcoming. Furthermore, signatures of novel (evolving) mutations may already be present in the form of the appearance of sets of accessory mutations that aid in the evolution and entrenchment of a primary DRM, that can be used to predict the likelihood of evolving resistance based on identifying pathway constraints. Such analysis may also be extensible to broadly neutralizing antibody design for an HIV vaccine. While the KMC simulations for each of the three HIV-1 target enzymes (PR, RT, and IN) consist of tens of thousands of trajectories that propagate the initial drug-naive ensemble to the final drug-experienced ensemble, both of which are observed in the Stanford HIVDB, we expect that there are far fewer pathways that actually predominate and cause drug resistance in patients. Identifying discrete clusters of pathways and defining the constraints that distinguish them is the subject of ongoing research.

Computational Methods. In this section, we present the Potts Hamiltonian model and the motivation behind the model. We also describe the KMC methods which are used to study the temporal evolution of HIV-1 under drug selection pressure (34, 62). Here, we measure time in these kinetic simulations in the units of the number of mutations attempted per position, i.e., the number of attempted mutations throughout the sequence divided by the sequence length. The length of PR, RT, and IN are 99, 188 and 263 respectively (33); and with this normalization scheme of the KMC algorithm ensures that, a time unit has the same meaning across the three HIV-1 enzymes, PR, RT, and IN, which have differential lengths. The details of the data processing, mutation classification, alphabet reduction, and length of amino acids are discussed in our previous paper (34).

The University of California, San Francisco (UCSF) Chimera (63) Structure Editing module is used to prepare Fig. 5. The extent of the known molecular envelope for different bound drugs is shown as a red mesh; The molecular envelope is created using UCSF Chimera (58) and using the PDBs: PIs: DRV (PDB: 4Q1X) SQV (PDB: 3S56), NFV (PDB: 3EKX), DRV (PDB: 3D20), LPV (PDB: 2QHC), TPV (PDB: 4NJU); NRTIs: FTC (PDB: 6UIR), AZT (PDB: 3KLG); NNRTIs: RPV (PDB: 3BGR), DOR (PDB: 7Z2G), EFV (PDB: 1IKW), ETR (PDB: 3MEC), NVP (PDB: 4PWD); INSTIs: DTG (8FNZ), BIC (6PUW).

Potts Hamiltonian Model. We use the Potts model which is a probabilistic model designed to describe the probabilities of observing specific states of a system that is constructed to be as unbiased as possible except to agree with the average first- and second-order observable (marginals) from the sequence data (64–70). The Potts model has a long history of use in statistical physics and analysis of protein sequence. In a set of protein sequences, the single and pair amino acid frequencies are average quantities that can be estimated from the finite samples using the data. The details of the models are described in our previous work (34, 37).

KMC Simulations. The KMC simulation is a Monte Carlo method which is intended to simulate the time evolution processes with known transition rates between states.

The Metropolis algorithm (71) is used to evaluate the metropolis acceptance probability of a mutation such as W (WT)→M

(Mutant) at a randomly chosen position i in a given sequence background at every simulation step given by $f_{W \rightarrow M}^{MFT}$

$$f_{W \rightarrow M}^{MET} = \min \left\{ 1, \frac{P_M}{P_W} \right\} \quad [3]$$

$$= \min \left\{ 1, e^{-\Delta E_{W \rightarrow M}} \right\},$$

where $\Delta E_{W \rightarrow M} = E_M - E_W$ is the change in Potts energy in going from residue W to M in the given background.

At the beginning of the simulation process a set of seed sequences (drug-naive sequences) are taken and a random position i and random mutant residue α (Reduced alphabet of four letters are used) are chosen the amino acid character at the chosen position i is either preserved or mutated based at the chosen position i is either preserved or mutated based on based on the Metropolis acceptance rate $f_{W \rightarrow M}^{MFT}$. For example, a mutation V32I from V (valine) → I (isoleucine) at position 32 in HIV-1 PR (99 residues long) has a probability $\left(\frac{1}{99}\right) e^{-\Delta E_{L \rightarrow M}^{32}}$ associated with the mutation at each KMC step. The algorithm used here allowed self-mutations during the simulation.

When evaluating fitnesses using a Potts model, it is possible to use alternative “selection temperatures” other than the inferred value of $T = 1$, giving a modified sequence probability distribution $P(S) \propto e^{-E(S)/T}$. However, we use $T = 1$ in our simulations as this ensures a “generative” model that accurately replicates the observed mutation frequencies in the HIV sequence dataset affected by drugs. Modifying the temperature would alter the mutant frequencies and higher-order mutation patterns, resulting either an excess or dearth of mutations relative to the observed HIV sequence data, inaccurately mimicking drug pressure in the observed drug-experienced ensemble. Our kinetic Potts model is a coarse-grained representation of the within-host evolutionary process. Evolution within a host is driven by various selective forces due to host immune response or to maintain viral viability, nonselective forces such as genetic drift, and is affected by other well-known aspects of HIV-1 infection such as retroviral dormancy, compartmentalization, and high recombination rates. In the chronic phase of HIV-1 infection, a large number of viral variants may be present at any point in time; however, this population is typically measured to have high sequence identity of close to 99% between pairs of viral genomes in a single host. This is much larger than the typical sequence identity of consensus sequences from different hosts of 90% for our datasets and justifies summarizing a host’s viral population by a single “consensus” sequence. Additionally, the host consensus sequence is observed to accumulate substitutions in a relatively clock-like manner, suggestive of sequential selective sweeps. Therefore, instead of modeling the detailed “microscopic” evolutionary forces, we use a coarse-graining which only tracks the consensus sequence of a host viral population over time as it accumulates substitutions due to these underlying forces.

In this way, one interpretation of our coarse-grained kinetics is that it models a series of point-mutation events in a viral population which occur according to a Poissonian mutational process, and these mutations are either fixed or lost from the population according to a fitness landscape inferred based on between-host sequence data. We coarse-grain a number of aspects of evolutionary dynamics, for instance we model amino-acid sequences instead of nucleotide sequences, assuming all amino acids can mutate to all others as is commonly done in phylogenetic analyses, for

instance in the Whelan and Goldman (WAG) and Jones-Taylor-Thornton (JTT) models. While this coarse-grained model is necessarily a simplification of HIV-1 viral dynamics, there are key properties of its construction which support conclusions drawn from it. First, an inferred “Potts” model prevalence landscape will implicitly capture many of the averaged effects of various microscopic evolutionary forces because it is fitted to HIV-1 sequence datasets which arose under the microscopic dynamics. For instance, it will capture mutational biases as these causes an increase in the inferred prevalence of the biased amino acids. Second, this model is numerically consistent with the observed between-host sequence variation data: If we use this kinetic model to simulate parallel trajectories (representing evolution in different hosts) and collect the final sequences, then by construction the mutational statistics of the generated sequences (frequencies of amino acids and k-mers of amino acids) match those of the between-host sequence datasets used to train the fitness model. We use a particular inference technique which we have confirmed gives a generative model which very closely reproduces the natural patterns of HIV-1 sequence variation for high order k-mers in generated sequences.

We assume an underlying Poisson mutational process, such that mutation arises at a rate μ . We implemented this by assigning each step in the Metropolis Algorithm a time drawn from an exponential distribution which is the waiting time for a Poisson process.

$$P(t|\mu) = \mu e^{-\mu t}. \quad [4]$$

With this overall KMC scheme our simulations match empirical data in two ways. First, the value of μ can be calibrated so that the simulated evolutionary trajectories accumulated substitutions at the same rate as observed experimentally. Second, when using this scheme to run many parallel trajectories until equilibrium, the bivariate residues of the resulting MSA match those observed in the HIV-1 sequence database used to train the Potts model.

Potential Coarse Graining Approximations. Here we discuss possible coarse-graining errors we have investigated. One aspect which our model has coarse-grained is drug-specific selection pressure and host-specific immune pressure. The ranges for literature acquisition times listed in Table 1 reflect, in part, the fact that some mutations arise in response to multiple drugs, but with some difference in acquisition time for each drug, this effect cannot be captured by our fitness model which reflects an averaged selection pressure of all drugs. But the overall correspondence of our model and literature suggests the major determinant of acquisition time is the epistatic interaction of the primary DRM with accessory mutations, which is captured by our fitness model and is independent of specific drug. It may be possible to explicitly model such drug-specific modulation of selection strengths through extensions of our fitness model, which we intend to investigate in the future.

Our model also coarse-grains the mutational process, and for instance does not explicitly distinguish between transition (Ti)

and transversion (Tv) mutations, which occur at different rates, and does not explicitly distinguish between mutations in the amino-acid sequence which correspond to single-nucleotide and double-nucleotide mutations at the nucleotide level. Because our fitness landscape is inferred from HIV-1 sequences which evolved in vivo under the influence of these mutational biases, the model implicitly captures their effect to some degree. To further investigate whether these mutational biases significantly affect DRM acquisition time, we investigated their effect for the mutants listed in Table 1.

The alterations in the genetic code corresponding to each DRMs are listed in the second column of Table 1 and *SI Appendix, Tables S1–S4* to assess the impact of genetic codes on the timeline of drug resistance evolution. The majority of DRMs, regardless of the protein type (PR, RT, or IN), are linked to single nucleotide changes (approximately 90%). Consequently, for these mutants, the number of genetic code alterations does not directly affect the time required to develop drug resistance in HIV. The DNA substitution mutations can be categorized as transition (Ti) or transversion (Tv). Transitions involve the exchange of two-ring purines (A → G) or one-ring pyrimidines (C → T), representing bases with similar sizes. On the other hand, transversions involve the interchange of purine for pyrimidine bases, which entails the exchange between one-ring and two-ring structures. We have included the nature of nucleotide changes (Δ_{nuc}) for each DRM of all proteins (PR, RT, and IN) in the third column of Table 1 and *SI Appendix, Tables S1–S4*. Our model does not provide information about the DNA level of the sequences; therefore, silent substitutions due to wobble base pair effects are not considered when determining the nature of nucleotide changes. In summary, we conclude that mutational biases have a negligible or relatively minor influence on acquisition time, whereas the epistatic interactions captured by our fitness model have a much more significant effect.

Data, Materials, and Software Availability. All study data are included in the article and/or *SI Appendix*. The code to perform the KMC simulations is available on Github (72).

ACKNOWLEDGMENTS. This work has been supported by the NIH through grants awarded to R.M.L. (R01 AI178849, U01 AI136680 S10OD020095). The NSF also provided funding through a grant awarded to A.H. and R.M.L. (1934848). E.A. was supported by NIH grants U54-AI150472 and R01 AI027690. D.L. was supported by NIH grants U54-AI150472, U01 AI136680, R01 AI146017, the Margaret T. Morris and the Hearst Foundations. A.B. was also supported by the Eric and Wendy Schmidt AI in Science Postdoctoral Fellowship, a program of Schmidt Sciences. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author affiliations: ^aCenter for Biophysics and Computational Biology, College of Science and Technology, Temple University, Philadelphia, PA 19122; ^bLaboratory of Genetics, The Salk Institute for Biological Studies, La Jolla, CA 92037; ^cDepartment of Physics, University of California San Diego, La Jolla, CA 92093; ^dDepartment of Chemistry, Temple University, Philadelphia, PA 19122; ^eDepartment of Chemistry and Chemical Biology, Center for Advanced Biotechnology and Medicine, Rutgers University, Piscataway, NJ 08854; ^fGraduate School of Biological Sciences, Department of Molecular Biology, University of California San Diego, La Jolla, CA 92093; and ^gDepartment of Physics, Temple University, Philadelphia, PA 19122

- UNAIDS, The path that ends AIDS: UNAIDS Global AIDS Update 2023 (Joint United Nations Programme on HIV/AIDS, Geneva, 2023). <https://www.unaids.org/en/resources/fact-sheet> (UNAIDS, 2022). Accessed 15 April 2022.
- J. M. Coffin, The virology of AIDS: 1990. *AIDS* **4**, S9 (1990).
- J. M. Coffin, HIV population dynamics in vivo: Implications for genetic variation, pathogenesis, and therapy. *Science* **267**, 483–489 (1995).
- A. S. Perelson, A. U. Neumann, M. Markowitz, J. M. Leonard, D. D. Ho, HIV-1 dynamics in vivo: Virion clearance rate, infected cell life-span, and viral generation time. *Science* **271**, 1582–1586 (1996).
- S.-Y. Rhee *et al.*, Public availability of HIV-1 drug resistance sequence and treatment data: A systematic review. *The Lancet Microbe* **3**, e392–e398 (2022).
- J. M. Cuevas, R. Geller, R. Garijo, J. López-Aldeguer, R. Sanjuán, Extremely high mutation rate of HIV-1 in vivo. *PLoS Biol.* **13**, e1002251 (2015).
- J. P. Barton *et al.*, Relative rate and location of intra-host HIV evolution to evade cellular immunity are predictable. *Nat. Commun.* **7**, 11660 (2016).
- P. L. Boyer, S. G. Sarafianos, E. Arnold, S. H. Hughes, Selective excision of AZTMP by drug-resistant human immunodeficiency virus reverse transcriptase. *J. Virol.* **75**, 4832–4842 (2001).
- S. G. Sarafianos *et al.*, Structure and function of HIV-1 reverse transcriptase: Molecular mechanisms of polymerization and inhibition. *J. Mol. Biol.* **385**, 693–713 (2009).
- D. O. Passos *et al.*, Structural basis for strand-transfer inhibitor binding to HIV intasomes. *Science* **367**, 810–814 (2020).

11. M. F. Kearney *et al.*, Lack of detectable HIV-1 molecular evolution during suppressive antiretroviral therapy. *PLoS Pathog.* **10**, e1004010 (2014).
12. J. I. Boucher *et al.*, Constrained mutational sampling of amino acids in HIV-1 protease evolution. *Mol. Biol. Evol.* **36**, 798–810 (2019).
13. E. Spielvogel *et al.*, Selection of HIV-1 for resistance to fifth-generation protease inhibitors reveals two independent pathways to high-level resistance. *eLife* **12**, e80328 (2023).
14. T. C. Butler, J. P. Barton, M. Kardar, A. K. Chakraborty, Identification of drug resistance mutations in HIV from constraints on natural evolution. *Phys. Rev. E* **93**, 022412 (2016).
15. S. W. Lockless, R. Ranganathan, Evolutionarily conserved pathways of energetic connectivity in protein families. *Science* **286**, 295–299 (1999).
16. J. D. Bloom, L. I. Gong, D. Baltimore, Permissive secondary mutations enable the evolution of influenza oseltamivir resistance. *Science* **328**, 1272–1275 (2010).
17. O. Haq, M. Andrec, A. V. Morozov, R. M. Levy, Correlated electrostatic mutations provide a reservoir of stability in HIV protease. *PLoS Comput. Biol.* **8**, e1002675 (2012).
18. R. M. Levy, A. Haldane, W. F. Flynn, Potts Hamiltonian models of protein co-variation, free energy landscapes, and evolutionary fitness. *Curr. Opin. Struct. Biol.* **43**, 55–62 (2017).
19. P. Tian, J. M. Louis, J. L. Baber, A. Aniana, R. B. Best, Co-evolutionary fitness landscapes for sequence design. *Angew. Chem. Int. Ed. Engl.* **57**, 5674–5678 (2018).
20. R. A. Neher, B. I. Shraiman, Competition between recombination and epistasis can cause a transition from allele to genotype selection. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 6866–6871 (2009).
21. M. W. Chang, B. E. Torbett, Accessory mutations maintain stability in drug-resistant HIV-1 protease. *J. Mol. Biol.* **410**, 756–760 (2011).
22. W. F. Flynn *et al.*, Deep sequencing of protease inhibitor resistant HIV patient isolates reveals patterns of correlated mutations in gag and protease. *PLoS Comput. Biol.* **11**, e1004249 (2015).
23. R. W. Shafer, Rationale and uses of a public HIV drug-resistance database. *J. Infect. Dis.* **194**, S51–S58 (2006).
24. C. Apetrei *et al.*, HIV Sequence Compendium 2021 (LA-UR-23-22840, Theoretical Biology and Biophysics Group, Los Alamos National Laboratory, NM, 2021). <http://www.hiv.lanl.gov/>.
25. A. M. Wensing *et al.*, 2022 update of the drug resistance mutations in HIV-1. *Top Antivir. Med.* **30**, 559–574 (2022).
26. A. L. Ferguson *et al.*, Translating HIV sequences into quantitative fitness landscapes predicts viral vulnerabilities for rational immunogen design. *Immunity* **38**, 606–617 (2013).
27. J. P. Barton, M. Kardar, A. K. Chakraborty, Scaling laws describe memories of host-pathogen riposte in the HIV population. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 1965–1970 (2015).
28. D. D. Pollock, G. Thiltgen, R. A. Goldstein, Amino acid coevolution induces an evolutionary Stokes shift. *Proc. Natl. Acad. Sci. U.S.A.* **109**, E1352–E1359 (2012).
29. P. Shah, D. M. McCandlish, J. B. Plotkin, Contingency and entrenchment in protein evolution under purifying selection. *Proc. Natl. Acad. Sci. U.S.A.* **112**, E3226–E3235 (2015).
30. D. M. McCandlish, P. Shah, J. B. Plotkin, Epistasis and the dynamics of reversion in molecular evolution. *Genetics* **203**, 1335–1351 (2016).
31. M. A. Depriost, D. M. Weinreich, D. L. Hartl, Missense meanderings in sequence space: A biophysical view of protein evolution. *Nat. Rev. Genet.* **6**, 678–687 (2005).
32. M. J. Harms, J. W. Thornton, Historical contingency and its biophysical basis in glucocorticoid receptor evolution. *Nature* **512**, 203–207 (2014).
33. N. K. Yilmaz, C. A. Schiffer, *Drug Resistance to HIV-1 Protease Inhibitors: Molecular Mechanisms and Substrate Coevolution* (Springer International Publishing, 2017), pp. 535–544.
34. A. Biswas, A. Haldane, E. Arnold, R. M. Levy, Epistasis and entrenchment of drug resistance in HIV-1 subtype B. *eLife* **8**, e50524 (2019).
35. Y. Iwasa, F. Michor, M. A. Nowak, Stochastic tunnels in evolutionary dynamics. *Genetics* **166**, 1571–1579 (2004).
36. Y. Guo, M. Vucelja, A. Amir, Stochastic tunneling across fitness valleys can give rise to a logarithmic long-term fitness trajectory. *Sci. Adv.* **5**, eaav3842 (2019).
37. I. Choudhuri, A. Biswas, A. Haldane, R. M. Levy, Contingency and entrenchment of drug-resistance mutations in HIV viral proteins. *J. Phys. Chem. B* **126**, 10622–10636 (2022).
38. A. Biswas, A. Haldane, R. M. Levy, Limits to detecting epistasis in the fitness landscape of HIV. *PLoS One* **17**, e0262314 (2022).
39. J. A. de la Paz, C. M. Nartey, M. Yuvaraj, F. Morcos, Epistatic contributions promote the unification of incompatible models of neutral molecular evolution. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 5873–5882 (2020).
40. J. Gizzio, A. Thakur, A. Haldane, R. M. Levy, Evolutionary divergence in the conformational landscapes of tyrosine vs serine/threonine kinases. *eLife* **11**, e83368 (2022).
41. J. Rodriguez-Rivas, G. Croce, M. Muscat, M. Weigt, Epistatic models predict mutable sites in SARS-CoV-2 proteins and epitopes. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2113118119 (2022).
42. M. Bisardi, J. Rodriguez-Rivas, F. Zamponi, M. Weigt, Modeling sequence-space exploration and emergence of epistatic signals in protein evolution. *Mol. Biol. Evol.* **39**, msab321 (2022).
43. A. Haldane, R. M. Levy, Mi3-GPU: MCMC-based inverse Ising inference on GPUs for protein covariation analysis. *Comput. Phys. Commun.* **260**, 107312 (2021).
44. A. Chaillon *et al.*, HIV persists throughout deep tissues with repopulation from multiple anatomical sources. *J. Clin. Invest.* **130**, 1699–1712 (2020).
45. K. Theys *et al.*, The impact of HIV-1 within-host evolution on transmission dynamics. *Curr. Opin. Virol.* **28**, 92–101 (2018).
46. A. F. Feder, K. N. Harper, C. J. Brumme, P. S. Pennings, Understanding patterns of HIV multi-drug resistance through models of temporal and spatial drug heterogeneity. *eLife* **10**, e69032 (2021).
47. A. R. Wargo, G. Kurath, Viral fitness: Definitions, measurement, and current insights. *Curr. Opin. Virol.* **2**, 538–545 (2012).
48. E. Domingo, J. J. Holland, RNA virus mutations and fitness for survival. *Annu. Rev. Microbiol.* **51**, 151–178 (1997).
49. G. Sella, A. E. Hirsh, The application of statistical physics to evolutionary biology. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 9541–9546 (2005).
50. S. C. Bihani, G. D. Gupta, M. V. Hosur, Molecular basis for reduced cleavage activity and drug resistance in D30N HIV-1 protease. *J. Biomol. Struct. Dyn.* **40**, 13127–13135 (2022).
51. S. Pawar *et al.*, Structural studies of antiviral inhibitor with HIV-1 protease bearing drug resistant substitutions of V32I, I47V and V82I. *Biochem. Biophys. Res. Commun.* **514**, 974–978 (2019).
52. D. A. Ragland *et al.*, Drug resistance conferred by mutations outside the active site through alterations in the dynamic and structural ensemble of HIV-1 protease. *J. Am. Chem. Soc.* **136**, 11956–11963 (2014).
53. N. J. Cook *et al.*, Structural basis of second-generation HIV integrase inhibitor action and viral resistance. *Science* **367**, 806–810 (2020).
54. X. Tu *et al.*, Structural basis of HIV-1 resistance to AZT by excision. *Nat. Struct. Mol. Biol.* **17**, 1202–1209 (2010).
55. M. T. Lai *et al.*, Mechanistic study of common non-nucleoside reverse transcriptase inhibitor-resistant mutations with K103N and Y181C substitutions. *Viruses* **8**, 263 (2016).
56. Y. Hsiou *et al.*, The Lys103Asn mutation of HIV-1 RT: A novel mechanism of drug resistance. *J. Mol. Biol.* **309**, 437–445 (2001).
57. K. Das *et al.*, Crystal structures of 8-Cl and 9-Cl TIBO complexed with wild-type HIV-1 RT and 8-Cl TIBO complexed with the Tyr181Cys HIV-1 RT drug-resistant mutant. *J. Mol. Biol.* **264**, 1085–1100 (1996).
58. J. Ren *et al.*, Structural mechanisms of drug resistance for mutations at codons 181 and 188 in HIV-1 reverse transcriptase and the improved resilience of second generation non-nucleoside inhibitors. *J. Mol. Biol.* **312**, 795–805 (2001).
59. S. Hare *et al.*, Molecular mechanisms of retroviral integrase inhibition and the evolution of viral resistance. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 20057–20062 (2010).
60. M. Li *et al.*, Mechanisms of HIV-1 integrase resistance to dolutegravir and potent inhibition of drug-resistant variants. *Sci. Adv.* **9**, eadg5953 (2023).
61. P. L. Tzou *et al.*, Integrase strand transfer inhibitor (INSTI)-resistance mutations for the surveillance of transmitted HIV-1 drug resistance. *J. Antimicrob. Chemother.* **75**, 170–182 (2020).
62. W. F. Flynn, A. Haldane, B. E. Torbett, R. M. Levy, Inference of epistatic effects leading to entrenchment and drug resistance in HIV-1 protease. *Mol. Biol. Evol.* **34**, 1291–1306 (2017).
63. E. F. Pettersen *et al.*, UCSF Chimera—a visualization system for exploratory research and analysis. *J. Comput. Chem.* **25**, 1605–1612 (2004).
64. S. Cocco, C. Feinauer, M. Figliuzzi, R. Monasson, M. Weigt, Inverse statistical physics of protein sequences: A key issues review. *Rep. Prog. Phys.* **81**, 032601 (2018).
65. M. Mézard, T. Mora, Constraint satisfaction problems and neural networks: A statistical physics perspective. *J. Physiol. Paris* **103**, 107–113 (2009).
66. M. Weigt, R. A. White, H. Szurmant, J. A. Hoch, T. Hwa, Identification of direct residue contacts in protein-protein interaction by message passing. *Proc. Natl. Acad. Sci. U.S.A.* **106**, 67–72 (2009).
67. J. I. Sulikowska, F. Morcos, M. Weigt, T. Hwa, J. N. Onuchic, Genomics-aided structure prediction. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 10340–10345 (2012).
68. R. Nassar, G. L. Dignon, R. M. Razban, K. A. Dill, The protein folding problem: The role of theory. *J. Mol. Biol.* **433**, 167126 (2021).
69. F. Morcos, N. P. Schafer, R. R. Cheng, J. N. Onuchic, P. G. Wolynes, Coevolutionary information, protein folding landscapes, and the thermodynamics of natural selection. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 12408–12413 (2014).
70. N. Bhattacharya *et al.*, Interpreting potts and transformer protein models through the lens of simplified attention. *Pac. Symp. Biocomput.* **27**, 34–45 (2022).
71. W. K. Hastings, Monte carlo sampling methods using markov chains and their applications. *Biometrika* **57**, 97 (1970).
72. A. Biswas, I. Choudhuri, A. Haldane, Code from "Kinetic coevolutionary models predict the temporal emergence of HIV-1 resistance mutations under drug selection pressure". GitHub. <https://github.com/avikbiswas/KMCsequence>. Deposited 11 March 2023.