

Fitness effects of mutations to SARS-CoV-2 proteins

Jesse D. Bloom^{1,2,3,†} and Richard A. Neher^{4,5,†}

¹Basic Sciences and Computational Biology, Fred Hutchinson Cancer Center, 1100 Fairview Ave N, Seattle, WA 98109, USA, ²Department of Genome Sciences, University of Washington, 3720 15th Ave NE, Seattle, WA 98195, USA, ³Howard Hughes Medical Institute, 1100 Fairview Ave N, Seattle, WA 98109, USA,

⁴Biozentrum, University of Basel, Spitalstrasse 41, Basel 4056, Switzerland and ⁵Swiss Institute of Bioinformatics, Lausanne 1015, Switzerland

[†]<https://orcid.org/0000-0003-1267-3408>

[†]<https://orcid.org/0000-0003-2525-1407>

*Corresponding author: E-mail: jbloom@fredhutch.org

Abstract

Knowledge of the fitness effects of mutations to SARS-CoV-2 can inform assessment of new variants, design of therapeutics resistant to escape, and understanding of the functions of viral proteins. However, experimentally measuring effects of mutations is challenging: we lack tractable lab assays for many SARS-CoV-2 proteins, and comprehensive deep mutational scanning has been applied to only two SARS-CoV-2 proteins. Here, we develop an approach that leverages millions of publicly available SARS-CoV-2 sequences to estimate effects of mutations. We first calculate how many independent occurrences of each mutation are expected to be observed along the SARS-CoV-2 phylogeny in the absence of selection. We then compare these expected observations to the actual observations to estimate the effect of each mutation. These estimates correlate well with deep mutational scanning measurements. For most genes, synonymous mutations are nearly neutral, stop-codon mutations are deleterious, and amino acid mutations have a range of effects. However, some viral accessory proteins are under little to no selection. We provide interactive visualizations of effects of mutations to all SARS-CoV-2 proteins (<https://jbloomlab.github.io/SARS2-mut-fitness/>). The framework we describe is applicable to any virus for which the number of available sequences is sufficiently large that many independent occurrences of each neutral mutation are observed.

Keywords: COVID-19; dN/dS; mutation rate; deep mutational scanning; fitness; UShER.

1. Introduction

The rapid evolution of SARS-CoV-2 has led to the emergence of viral variants with enhanced transmissibility, escape from therapeutics, or reduced recognition by immunity (Harvey et al. 2021; Abdool Karim and de Oliveira 2021). To anticipate and mitigate this evolution, the scientific community has launched efforts to assess the risk of new viral variants (DeGrace et al. 2022) and create therapeutics that target constrained regions of the virus where resistance is less likely to evolve (Moghadas et al. 2023; Iketani et al. 2022b; Hiscox et al. 2021). Both efforts require determining how specific mutations affect viral fitness.

Unfortunately, experimentally measuring the effects of mutations is challenging for most SARS-CoV-2 proteins. For spike, tractable lab assays have identified key functional and antigenic mutations (Harvey et al. 2021; Weisblum et al. 2020), and enabled deep mutational scanning measurements of how most mutations affect receptor binding, cellular infection, and antibody recognition (Starr et al. 2020; Dadonaite et al. 2023; Greaney et al. 2021; Cao et al. 2022a). These experimental data are valuable for assessing new spike variants (DeGrace et al. 2022; Greaney, Starr and Bloom 2022; Tzou et al. 2022) and designing antibody therapeutics with greater resistance to escape (Starr et al. 2021; Rappazzo et al. 2021; Cao et al. 2022b). But most SARS-CoV-2 proteins lack

tractable lab assays, despite contributing to viral fitness (Thorne et al. 2022; Syed et al. 2021; McGrath et al. 2022) and being targets of efforts to develop anti-viral drugs (Tao et al. 2021). The only non-spike SARS-CoV-2 protein with large-scale experimental measurements of mutation effects is Mpro (Flynn et al. 2022; Iketani et al. 2022a).


An alternative to experiments is to estimate effects of mutations by analyzing natural viral sequences. The amount of data available for such analyses has increased dramatically over the last few years with the sequencing of SARS-CoV-2 from millions of human infections. So far analyses of these sequences have focused on analyzing expanding viral clades to identify mutations that mediate immune escape or increase transmissibility (Obermeyer et al. 2022; Lee et al. 2022; Maher et al. 2022). The basic idea is that mutations that repeatedly appear near the base of clades that increase in relative frequency are likely beneficial to the virus. However, only a small minority of all possible mutations are beneficial, with most being nearly neutral or deleterious. For purposes such as identifying constrained drug targets or understanding the function of viral proteins, it is important to estimate the effects of neutral or deleterious mutations as well as beneficial ones. Other studies have analyzed broader alignments of coronaviruses substantially diverged from SARS-CoV-2 (Rodriguez-Rivas et al. 2022;

A


Expected mutation counts from four-fold degenerate sites.

The third nucleotide in codon `acc` is four-fold degenerate since `aca`, `acc`, `acg`, and `act` all encode threonine.

We count the number of *mutations* on branches of the tree, not the final alignment: in tree at right there are three `c→g` mutations although more than three sequences have the mutation.

**Actual counts of mutations at each site.**

In tree at right, there is one count of the indicated `c→g` mutation, which induces an alanine to glycine mutation. The fact that the actual counts are less than the expected counts suggests this mutation is deleterious.



B




Figure 1. Expected versus actual counts of mutations. (A) The number of expected counts of each type of nucleotide mutation is computed from four-fold degenerate sites, and then compared the actual counts of each mutation. (B) Expected versus actual counts for each nucleotide mutation type aggregated across all viral clades and averaged across all sites where the mutation is four-fold degenerate, synonymous (including four-fold degenerate), nonsynonymous, or introduces a stop codon. See https://jbloomlab.github.io/SARS2-mut-fitness/avg_counts.html for an interactive version of panel B that enables mouseovers to read off specific values.

Thadani et al. 2022), but the resulting estimates are limited by sparse sampling and possible changes in the impacts of some mutations across divergent viruses.

Here, we develop a new approach that uses natural sequences to estimate the effects of mutations. Our basic insight is that there are now so many SARS-CoV-2 sequences that all non-deleterious single-nucleotide mutations are expected to independently occur many times along the observed phylogenetic tree. We therefore first calculate the number of expected observations of independent occurrences of each mutation based on the neutral mutation rate of SARS-CoV-2. We then compare these expected observations to the actual observations in the SARS-CoV-2 tree to estimate the effect of each mutation. The resulting estimates correlate well with existing deep mutational scanning data. Most viral proteins have regions under strong selective constraints. However, some accessory proteins show only weak selection against amino acid and even stop-codon mutations. Overall, our work demonstrates a new approach to determine the effects of mutations and provides detailed maps of mutational effects across the SARS-CoV-2 proteome.

2. Results

2.1 Mutation effects from actual versus expected counts

To determine how many times each mutation is expected to be observed, we used the pre-built UShER tree (McBroome et al. 2021; Turakhia et al. 2021; Lanfear 2020) of ~7-million public SARS-CoV-2 sequences to count nucleotide mutations at four-fold degenerate sites (Fig. 1A; Bloom et al. 2023). Because mutations at such sites never alter the amino acid sequence, these counts reflect the mutation process in the absence of protein-level selection (see below for caveats about nucleotide-level selection). The expected counts of a mutation from nucleotide x to y is simply the average count of this type of mutation across all four-fold degenerate sites with parental identity x . Importantly, we count independent occurrences of each mutation along the branches of the tree, not the sequences with the mutation in the final alignment (Fig. 1A). We also compute expected counts separately for each SARS-CoV-2 clade to account for shifts in mutation spectrum (Bloom et al. 2023; Ruis et al. 2022), and apply quality-control steps to remove spurious mutations (see Methods).

The expected counts per mutation vary with mutation type, ranging from ~565 for $C \rightarrow T$ to only ~9 for $T \rightarrow G$ mutations (Fig. 1B). This variation is because the SARS-CoV-2 mutation spectrum is highly biased toward specific mutation types (Bloom et al. 2023; Ruis et al. 2022; De Maio et al. 2021; Neher 2022). However, because there are so many SARS-CoV-2 sequences we are able to estimate the rate of even the rarest mutation types with high accuracy (Bloom et al. (2023)). For instance, there are $\sim 1.9 \times 10^4$ observed occurrences of $T \rightarrow G$ mutations across all ~2,100 four-fold degenerate sites with a parental identity of T , which is enough to estimate the $T \rightarrow G$ mutation rate (and therefore the expected counts of each mutation) with high accuracy.

We compared the expected counts to the actual observed counts of mutations averaged across sites (Fig. 1). For synonymous mutations, the expected and actual counts are similar. But for nonsynonymous and especially stop-codon mutations, the actual counts are substantially lower than the expected counts, reflecting purifying selection for protein function.

The ratio of actual to expected counts for each mutation is related to its effect on viral fitness. The intuition is straightforward: mutations arise at all sites, but viruses with deleterious mutations are less likely to transmit and be observed in sequencing of human SARS-CoV-2. Therefore, the ratio of actual to expected counts will be one for neutral mutations, and less than one for deleterious mutations. In the Methods and Appendix, we show that under plausible assumptions about SARS-CoV-2 evolution and sampling intensity (fraction of viruses sequenced), the fitness cost of a deleterious mutation scales roughly inversely with the ratio of actual to expected counts for mutations with costs greater than a few percent. A key result is the dependence on sampling intensity: if all human SARS-CoV-2 were sequenced even deleterious mutations would have a high chance of being sampled and we would need to study the subsequent spread of the mutations to assess their fitness. But the actual sampling intensity is ~0.1%, since there are ~7-million publicly available SARS-CoV-2 sequences and the total number of human infections is now probably roughly on par with the total global population of ~8-billion people. At this sampling intensity, the number of times a mutation is observed reflects more subtle reductions in transmission efficiency. We quantify the effect of each mutation as the logarithm of the ratio of actual to expected counts after summing counts for all nucleotides that encode the relevant amino acid. The

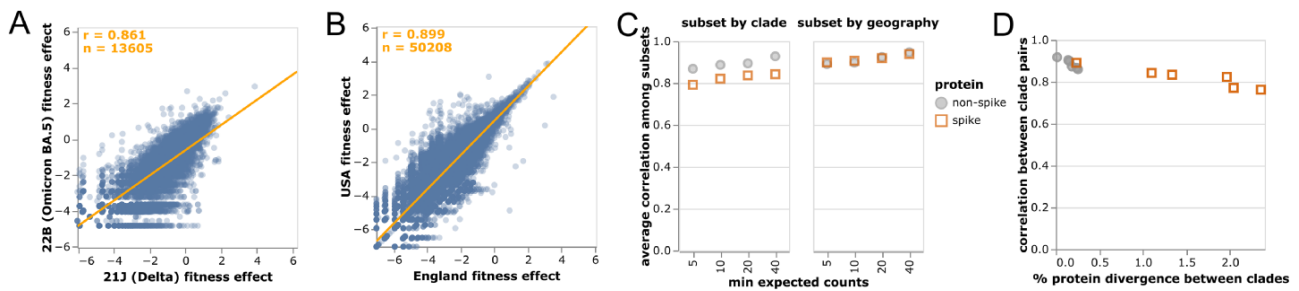


Figure 2. Correlations of mutation fitness effect estimates made using subsets of natural sequences. Correlations between estimates made (A) just using sequences from the Delta or Omicron BA.5 clades or (B) just from the USA or England. Each point is an amino acid mutation, the orange line is a least-squares regression, and orange text at upper left shows the number of mutations and Pearson's correlation coefficient. Only mutations with at least 10 expected counts are shown, which is why panels have different numbers of mutations shown (sequence subsets vary in size). Different subset size are also the reason why the regression line in (A) deviates from the identity $x = y$. (C) Correlations between clade or geography subsets become higher with an increasingly large threshold for minimum expected counts. Spike mutations have a worse correlation when subsetting by viral clade (plot shows average correlation over all pairwise combinations of Delta, BA.1, BA.2, and BA.5), but not when subsetting by geography (USA or England). (D) Correlations in estimated mutation-effects decline for clades with higher protein divergence, with the effect most noticeable for spike since spike is more diverged among SARS-CoV-2 clades than other viral proteins. See https://jbloomlab.github.io/SARS2-mut-fitness/clade_corr_chart.html and https://jbloomlab.github.io/SARS2-mut-fitness/subset_corr_chart.html for versions of A and B that include all viral clades with at least 500,000 total expected counts (summed across all mutations) and have other interactive options.

statistical noise is greater for mutations with fewer expected counts: the figures in this paper show mutations with ≥ 10 expected counts unless otherwise noted, with legends linking to interactive plots that enable adjustment of this threshold.

2.2 Mutation-effect estimates are robust to subsampling, with some evidence of epistasis in spike

We computed the correlations among mutation-effect estimates made using subsets of SARS-CoV-2 sequences from different viral clades or geographic locations. These estimates were well correlated, with some modest variation in estimates across sequence subsets (Fig. 2A, B).

The modest variation in estimates from different sequence subsets could have two causes: statistical noise due to finite mutation counts, or real shifts in mutation effects during SARS-CoV-2 evolution (Starr et al. 2022a; Moulana et al. 2022). To test for statistical noise, we computed correlations with different thresholds for how many expected counts are required before making an estimate for a mutation (Fig. 2C). Correlations increased with this count threshold, consistent with reduced statistical noise for larger mutation counts. But the correlation for spike mutations was consistently lower for cross-clade but not cross-geography comparisons (Fig. 2C). The lower cross-clade correlation for spike appears due to epistatic shifts in mutation effects (Starr et al. 2022a; Moulana et al. 2022; Pollock, Thiltgen and Goldstein 2012; Shah, McCandlish and Plotkin 2015; Lee et al. 2018) or changes in the selective landscape (Sun et al. 2023) between SARS-CoV-2 clades, since the correlation is lower between clades with higher spike divergence (Fig. 2D). In particular, the interactive version of Fig. 2A shows that mutations that are more beneficial in Omicron BA.5 than Delta are often antibody-escape mutations (e.g., K444N or G446S in spike (Greaney, Starr and Bloom 2022))—a result that makes sense, since newer variants like Omicron BA.5 are evolving under increased immune selection compared to earlier variants like Delta that circulated in a more immunologically naive population.

Despite evidence for some shifts in mutation effects in spike, for the rest of this paper we aggregate counts across viral clades to make a single estimate for each amino acid mutation. The

reason is that the accuracy of the estimates increases with the number of counts (Fig. 2C), and several mutation types only have enough counts for reasonable estimates when aggregating across clades (Fig. 1B). For the purposes of this paper, we deemed it preferable to have more accurate and comprehensive pan-SARS-CoV-2 estimates than noisier clade-specific estimates for fewer mutations. However, the interactive version of Fig. 2A linked in the legend enables exploration of mutations with disparate estimates among clades.

An important question is whether the mutation fitness estimates are affected by noise from limited statistical sampling of mutations or whether sequencing errors and bioinformatic artifacts distort the estimates. To assess if this is the case, we repeated the entire fitness estimation using an even larger pre-built USHER mutation-annotated tree of all ~ 14 -million SARS-CoV-2 sequences in GISAID Shu and McCauley (2017) as of March-29-2023. There is an extremely high correlation between fitness estimates made using the ~ 7 -million publicly available sequences and the larger GISAID tree (Supplementary Fig. S1). This concordance indicates that the set of ~ 7 -million public sequences is large enough that doubling the data does not appreciably shift the estimates, and so throughout this paper we use that sequence set due to our preference for publicly available data. Furthermore, fitness estimates using sequences from specific countries (USA and England, Fig. 2B) are also highly concordant, suggesting that sequencing and bioinformatic workflows are not driving the signal. Lastly, positions known to be under strong constraint *a priori* (e.g., start codons and the ribosomal slippage site) typically have no or only few mutations, suggesting that sequencing errors in consensus sequences are rare.

2.3 Structural and non-structural proteins are under strong purifying selection, but most accessory proteins are not

The distributions of mutation effects concur with biological intuition about how different classes of mutations impact protein function. Most synonymous mutations are nearly neutral, most stop codons are highly deleterious, and amino acid mutations range from slightly beneficial to highly deleterious (Fig. 3A). The handful of synonymous or noncoding mutations with highly




Figure 3. Distribution of effects of different classes of mutations. (A) Histograms of effects of synonymous, nonsynonymous, and stop-codon mutations across all viral genes. Neutral mutations have effects of zero (dashed gray vertical lines), and deleterious mutations have negative effects. (B) Effects of each class of mutation for each viral gene. Dark squares indicate the median effect, and the lighter rectangles span the interquartile range. Mutation types are color-coded as in panel (A). The apparent constraint on synonymous mutations in ORF9b is probably because this gene is encoded in an overlapping reading frame with N (Jungreis et al. 2021). See https://jbloomlab.github.io/SARS2-mut-fitness/effects_histogram.html and https://jbloomlab.github.io/SARS2-mut-fitness/effects_dist.html for plots that allow adjustment of the expected-count cutoff and other interactive options (such as separate histograms for each gene). See [Supplementary Fig. S3](#) for a version of panel B with genes ordered by genomic position rather than constraint on nonsynonymous mutations.

deleterious effects are in either regions of known non-coding constraint (e.g., the ORF1ab ribosomal slippage site (Bhatt et al. 2021)) or two regions in the center of E and the end of M (Supplementary Fig. S2).

To investigate differences in functional constraint among viral proteins, we computed the distributions of mutation effects separately for each gene (Figs. 3B and S3). SARS-CoV-2 proteins are grouped into three categories: nonstructural (or nsp) proteins, structural proteins (spike, M, N, and E), and accessory proteins (names prefixed with ‘ORF’) (V’kovski et al. 2021). The nonstructural and structural proteins are essential, and these proteins show strong selection against stop codons and clear although variable purifying selection against amino acid mutations (Figs. 3B and S3; e.g., nsp13 is under stronger protein-level constraint than nsp1).

However, most accessory proteins are under little constraint (Figs. 3B and S3). Stop-codon and amino acid mutations to ORF7a and ORF8 are not more deleterious than synonymous mutations (although recall that our estimates are only sensitive to fitness costs greater than a few percent). The lack of deleterious mutations to ORF8 is consistent with the fact that viruses with deletions in this gene have spread in humans (Su et al. 2020) and that major variants had stop codons early in ORF8. Indeed, the loss of accessory proteins such as ORF8 appears to occur with some regularity during the early evolution of non-human viruses in humans (Rochman, Wolf and Koonin 2022). The only accessory protein under strong purifying selection against stop codons is ORF3a (Fig. 3B), for which stop codons in the first 240 residues

are clearly deleterious (Supplementary Fig. S4). These observations concur with experiments showing SARS-CoV-2 is attenuated by deletion of ORF3a but there is little effect of deleting ORF6, ORF7a, or ORF8 (McGrath et al. 2022; Silvas et al. 2021; Liu et al. 2022). However, ORF3a’s function must be relatively insensitive to its protein sequence, since other than selection against stop codons there is only amino acid level constraint at a few sites like 135 and 138 (Supplementary Fig. S4). Observations such as these could help guide experimental studies to better understand protein function.

The accessory proteins are encoded in the last 5kb of the SARS-CoV-2 genome, meaning that our results show that this portion of the genome (with the exception of the genes coding for nucleoprotein, E, and M) is under relatively weak purifying selection (Supplementary Fig. S3). Note that some accessory proteins also vary in presence/absence across coronaviruses (V’kovski et al. 2021; Llanes et al. 2020), and are mostly dispensable in cell culture (McGrath et al. 2022; Silvas et al. 2021; Liu et al. 2022; Liu et al. 2014). Decreased purifying selection in genes encoded near the end of the genome is also observed in more distantly related orders of large single-stranded RNA viruses such as closteroviruses (Wang et al. 2011; Dawson and Folimonova 2013).

2.4 Mutation-effect estimates correlate with experiments

We examined how the mutation effects estimated using our approach compare with prior high-throughput deep mutational

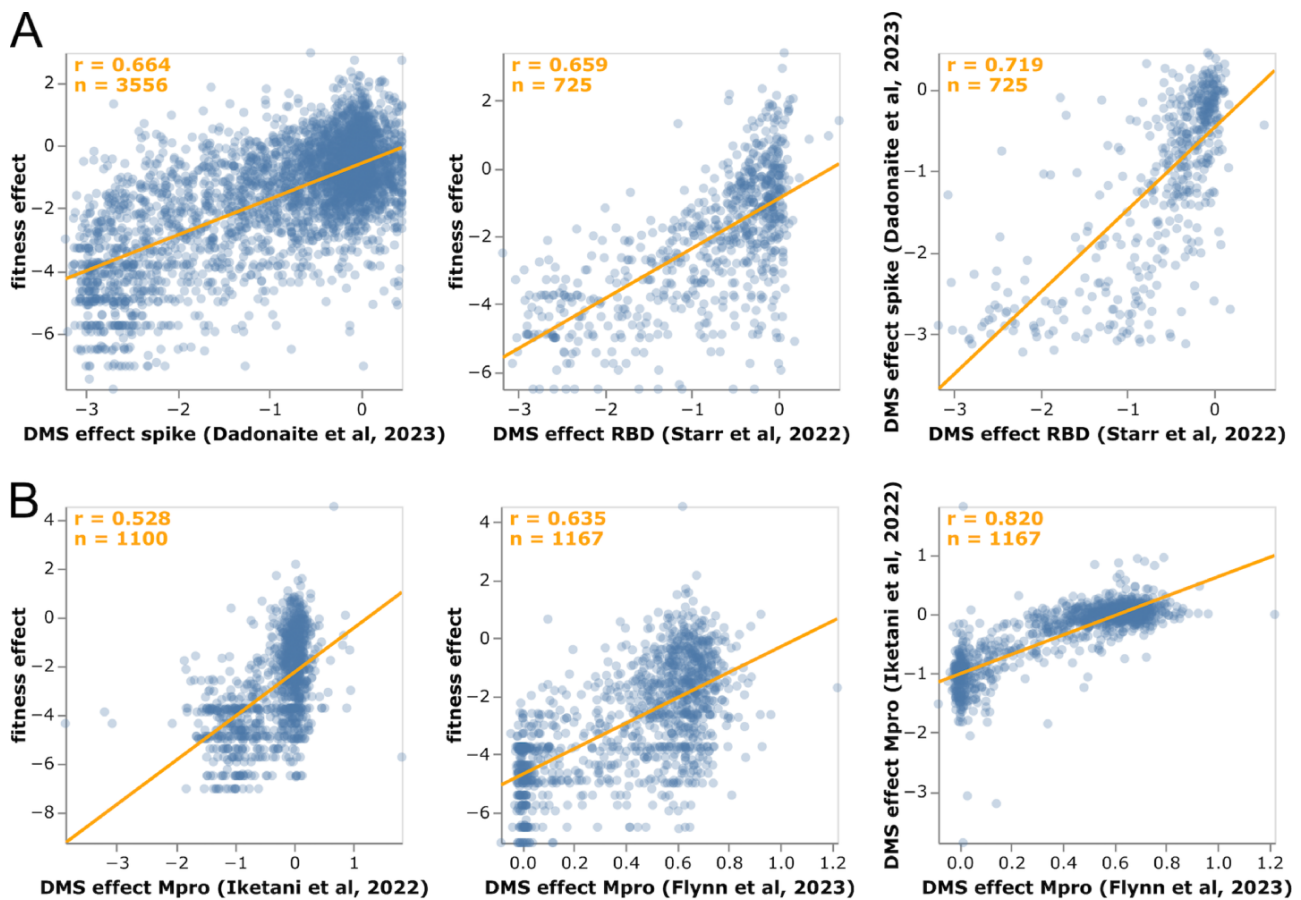


Figure 4. Correlation of mutation-effect estimates with experimental deep mutational scanning measurements for (A) the full spike (Dadonaite et al. 2023) or its RBD (Starr et al. 2022b), and (B) Mpro (Flynn et al. 2023; Iketani et al. 2022a). Each point is an amino acid mutation, the orange line is a least-squares regression, and orange text in the upper left shows the number of mutations and Pearson’s correlation coefficient. Each subpanel shows a different set of mutations (depending on which mutations were measured in that experiment). See https://jbloomlab.github.io/SARS2-mut-fitness/dms_S_corr.html and https://jbloomlab.github.io/SARS2-mut-fitness/dms_nsp5_corr.html for plots that also show the Mpro dataset from (Flynn et al. 2022) and have various interactive options. The plots in this figure show the average of the multiple phenotypes measured in the deep mutational scanning of Starr et al. (2022b); see https://jbloomlab.github.io/SARS2-mut-fitness/dms_S_all_corr.html for each phenotype separately. This figure only shows mutations with at least 20 expected counts, which is higher than the threshold of 10 used in most of the rest of this paper (this threshold can be adjusted in the interactive plots).

scanning measurements. For spike, two distinct experimental methodologies have been used to characterize large numbers of mutations: yeast display of the receptor-binding domain (RBD) (Starr et al. 2020; Starr et al. 2022b) and spike pseudotyped lentiviruses (Dadonaite et al. 2023). For Mpro (also known as nsp5 or 3CLpro), two different labs have performed deep mutational scanning using the same basic methodology of assaying protease cleavage in yeast (Flynn et al. 2022; Flynn et al. 2023; Iketani et al. 2022a).

For spike, our estimates from natural sequences correlate with the experiments almost as well as the two experimental methodologies correlate with each other (Fig. 4A), with Pearson’s correlations of 0.66 between the estimates and experiments versus 0.72 between the two experiments. Neither experiment fully captures how mutations affect viral fitness, since both RBD yeast display and lentiviral pseudotyping are imperfect proxies for spike function during actual human infections. Therefore, it is unclear how much the differences between the mutation-effect estimates and experiments are due to noise in the estimates versus limitations of the experiments. However, the fact that the estimates correlate with the experiments almost as well as the experiments correlate

with each (Fig. 4A) suggests the estimates are of comparable quality to experimental measurements. At least some of the mutations with the greatest divergence between our estimates and the deep mutational scanning likely represent experimental artifacts. For instance, P527L, which is favorable in the RBD deep mutational scan but deleterious in the sequence-based estimates and full-spike scan, is at the C-terminus of the yeast-displayed RBD (Starr et al. 2020) where it may adopt a non-native conformation.

The sequence-based estimates for Mpro also correlate with the deep mutational scans for that protein, although in this case the experiments correlate substantially better with each other than with our estimates (Fig. 4B). However, the Mpro experiments all use a similar yeast-based methodology (Flynn et al. 2022; Flynn et al. 2023; Iketani et al. 2022a) that fails to capture significant aspects of Mpro’s function during human infections. For instance, a stop codon at Q306 is well tolerated in the deep mutational scans but extremely unfavorable in our sequence-based estimates, and such a mutation would clearly be highly deleterious to actual virus as it would truncate the polyprotein. Similarly, K61N is well tolerated in the deep mutational scans but extremely unfavorable in our estimates, probably because in the full viral polyprotein

this residue mediates important interactions between Mpro and nsp7-10 (Yadav et al. 2022).

2.5 Mutation-effect estimates better capture functional constraint than dN/dS ratios or predictions from other methods

A longstanding approach for analyzing protein constraint is to compare rates of nonsynonymous (dN) and synonymous (dS) substitutions at each site (Nielsen and Yang 1998; Kosakovsky Pond and Frost 2005). These dN/dS ratios can be calculated by counting mutations or using phylogenetic substitution models (Kosakovsky Pond and Frost 2005; Yang and Nielsen 2000). A limitation of dN/dS ratios is they cannot be interpreted in terms of the fitness effects, since they simply represent the relative rate of amino acid substitution at a site rather than the effects of specific mutations (Spielman and Wilke 2015; Kryazhimskiy and Plotkin 2008). Nonetheless, we can compare dN/dS ratios to our mutation-effect estimates as measures of the average constraint at each site. The mutation-effect estimates greatly outperform dN/dS ratios as a measure of site-level constraint as assessed by correlation with deep mutational scanning experiments (Supplementary Fig. S5). The reason is in part because some aspects of functional constraint cannot be captured by a dN/dS ratio. For instance, the ACE2-affinity enhancing spike mutation N501Y arose in several SARS-CoV-2 variants early in the pandemic, and has since remained fixed due its importance for receptor binding (Starr et al. 2022b). Our mutation-effect estimates correctly reflect that site 501 is strongly prefers tyrosine, but the site has a high dN/dS ratio due to the early convergent evolution of this site to that preferred amino acid.

Our mutation-effect estimates also correlate better with deep mutational scanning experiments than predictions from two algorithms trained to learn epistatic models of mutation effects from phylogenetically broader but more sparsely sampled sequence data (Thadani et al. 2022; Rodriguez-Rivas et al. 2022) (Supplementary Fig. S6). Our estimates also correlate better with experiments than predictions by a machine-learning algorithm that integrates sequence and epidemiological data (Maher et al. 2022) (Supplementary Fig. S6). These results suggest that our straightforward approach of directly reading out the effects of mutations from their actual versus expected counts can outperform much more complex models when millions of sequences are available.

2.6 Fixed mutations tend to have beneficial or neutral effects

Amino acid mutations that have fixed in at least one viral clade are estimated to mostly have neutral or beneficial effects, whereas most other mutations are deleterious (Supplementary Fig. S7). This fact is unsurprising: viral lineages that expand into new clades do so because they have acquired beneficial mutations while avoiding deleterious ones (Luksza and Lässig 2014; Koelle and Rasmussen 2015; Huddleston et al. 2020). But the fact that the beneficial effects of fixed mutations are correctly estimated by our approach, which simply counts mutation occurrences and does not incorporate information on lineage size, demonstrates such mutations occur independently in many viral lineages that are more successful than average.

Most fixed mutations are estimated to be beneficial regardless of whether estimates are made using all viral clades, or just clades that did not fix the mutation (Supplementary Fig. S8). However, a few beneficial fixed mutations show epistatic entrenchment (Shah, McCandlish and Plotkin 2015; Starr et al. 2018) in

the sense that they are not particularly beneficial in clades in which they did not fix (Supplementary Fig. S8). The most striking example is S373P in spike, which has experimentally been shown to be neutral or slightly deleterious in pre-Omicron clades, but strongly beneficial in the Omicron clades in which it fixed (Starr et al. 2022b; Moulana et al. 2022).

2.7 Interactive exploration of amino acid fitnesses

To enable easy access to the mutation-effect estimates, we created interactive plots to enable exploration of the data for each protein. A static view of one of these plots is in Fig. 5; see <https://jbloomlab.github.io/SARS2-mut-fitness> for interactive versions for all proteins. These plots enable both high-level inspection of functional constraint across each protein, and detailed interrogation of the effects of specific mutations.

3. Discussion

Enough SARS-CoV-2 viruses have now been sequenced that many independent occurrences of every tolerated single-nucleotide mutation have been observed along the viral phylogeny. Here, we have described a new approach that leverages this fact to estimate the effects of these mutations. In essence, we treat natural evolution as a deep mutational scan, with the millions of publicly available SARS-CoV-2 sequences providing a readout of this experiment. The key is simply to calculate how many times each mutation has been ‘tested’ along the history of sampled viral sequences, and compare that expectation to the actual observations of the mutation among viruses sufficiently fit to have been sequenced in actual human infections.

The resulting estimates of mutational effects are robust to subsetting on specific viral clades or geographies, and correlate well with experimental measurements. In broad strokes, the mutation effects illuminate patterns of constraint: for instance, there is strong selection on structural and non-structural proteins, but only limited purifying selection on the accessory proteins.

However, the real value of our approach is in the detailed maps of effects of specific mutations to all viral proteins, including proteins with poorly understood functions not easily characterized in the lab. These maps will be of value for designing drugs that target constrained sites, interpreting the consequences of mutations observed during viral surveillance, and guiding experiments to mechanistically characterize protein function.

There are several caveats to our approach. First, because the number of observations of any given mutation is small compared to the millions of SARS-CoV-2 sequences being analyzed, our approach requires careful quality control to remove sequencing errors. Second, we assume the rate of each type of nucleotide mutation is uniform across the viral genome, and neglect higher-order context that may influence mutation rate (Sadykov et al. 2021; Beale et al. 2004). Likewise, we neglect constraint on nucleotide identity beyond the encoded protein sequence (Huston et al. 2021; Kuo and Masters 2013)—although this probably has only a minor effect, since our analyses show just a handful of synonymous sites are under strong selection (Supplementary Fig. S2). Third, the exact relationship between the statistics we calculate and viral fitness depend on the fraction of all infections that are sequenced (sampling intensity) and viral population dynamics. Although we derive this relationship, we do not adjust for sampling intensity and population dynamics when estimating mutation effects. Fourth, we make a single estimate for each mutation across all SARS-CoV-2, neglecting the




Figure 5. Effects of amino acid mutations to E protein. The area plot at top shows the average effects of mutations at each site, and the heatmap shows the effects of specific amino acids, with **x** denoting the amino acid identity in the Wuhan-Hu-1 strain. See <https://jbloomlab.github.io/SARS2-mut-fitness/E.html> for an interactive version of this plot that enables zooming, mouseovers, adjustment of the minimum expected count threshold, and layering of stop codon effects on the site plot. See <https://jbloomlab.github.io/SARS2-mut-fitness> for comparable interactive plots for all SARS-CoV-2 proteins.

epistasis that can affect some mutations (Starr et al. 2022a; Moulana et al. 2022). Finally there are a few technical caveats to how we count mutations that are discussed in the Methods section.

Conceptually, our approach differs from prior methods that aim to identify beneficial SARS-CoV-2 mutations associated with viral clades that increase in frequency (Obermeyer et al. 2022; Lee et al. 2022; Maher et al. 2022). Those methods draw information primarily from what happens downstream of a mutation. In contrast, we treat all mutations equivalently regardless of whether they are on a tip node or at the base of a large clade. Our approach is better for estimating effects of deleterious or nearly neutral mutations, but clade-growth methods may be better for beneficial mutations. In particular, clade size carries information beyond that contained in mutation counts alone (Supplementary Fig. S9). Hopefully, future work can combine mutation-counting and clade-growth methods for even better estimates of SARS-CoV-2 mutation effects. Note our approach is conceptually similar to estimating fitness costs of HIV or polio mutations from mutation-selection balance in deep sequencing of intra-population viral quasispecies (Zanini et al. 2017; Acevedo, Brodsky and Andino 2014), except we analyze mutation occurrences rather than frequencies to account for the phylogenetic structure and genetic hitchhiking that characterize global SARS-CoV-2 evolution.

The power of the approach we have described will increase with more viral sequencing. SARS-CoV-2 is the first virus with enough sequences that every tolerated mutation is observed multiple independent times. As costs drop, it is easy to imagine a future with even more viral sequences. As this occurs, viral genomic sequencing—which has traditionally been used primarily to track evolution and spread—will also become an increasingly precise tool to determine the effects of specific mutations.

4. Methods

4.1 Code and data availability

See the GitHub repository at <https://github.com/jbloomlab/SARS2-mut-fitness> for the computer code and processed data (eg, fitness estimates and mutation counts). That repository contains a README with links to specific data files as well as a description of the computational pipeline. See https://github.com/jbloomlab/SARS2-mut-fitness/blob/main/results/aa_fitness/aa_fitness.csv for final estimates of amino acid fitnesses across all clades; other intermediate data files are also provided in the GitHub repository. The specific version of the repository used for this paper is tagged as ‘bioRxiv-v2’ on GitHub (<https://github.com/jbloomlab/SARS2-mut-fitness/tree/bioRxiv-v2>). The pipeline is fully reproducible and is run using `snakemake` (Mölder et al. 2021) with interactive plots rendered using `altair` (VanderPlas et al. 2018).

The interactive plots are rendered at <https://jbloomlab.github.io/SARS2-mut-fitness> via GitHub pages.

4.2 Versioning of analyses of different sequence sets

The figures in this manuscript show analyses of the set of all publicly available sequences as of 11 May 2023. However, the pipeline can be run on different sequence sets. The sequence sets on which the analysis is currently run include all sets listed under the ‘mat_trees’ key in <https://github.com/jbloomlab/SARS2-mut-fitness/blob/main/config.yaml>; these include the sets of all public sequences from several earlier dates (such as those available for the first version of this analysis), as well as the set of all sequences in GISAID as of 29 March 2023. A version of the results for each sequence set is provided in the GitHub repository (<https://github.com/jbloomlab/SARS2-mut-fitness>) in subdirectories with names like ‘results_public_2023-05-11’, and the index

page for the interactive plots (<https://jbloomlab.github.io/SARS2-mut-fitness>) links at the bottom to plots for each sequence set. The 'current_mat' key in <https://github.com/jbloomlab/SARS2-mut-fitness/blob/main/config.yaml> specifies which sequence set is used to generate the main results that in the 'results' sub-directory in the GitHub repository and are shown by default in the interactive plots; we anticipate periodically updating this to newer sequence sets as more sequences become available. See https://jbloomlab.github.io/SARS2-mut-fitness/mat_aa_fitness_correlations.html for the correlations among mutation effects estimated from the different sequence sets.

For the GISAID sequence set, we acknowledge the submitters of the sequences listed at the following URLs: <https://doi.org/10.55876/gis8.230403ab>, <https://doi.org/10.55876/gis8.230403hg>, <https://doi.org/10.55876/gis8.230403ht>, and <https://doi.org/10.55876/gis8.230403tg>.

4.3 Counting mutations along the phylogenetic tree

We counted occurrences of each mutation in each viral clade using the USHER pre-built mutation-annotated tree (McBroome et al. 2021; Turakhia et al. 2021; Lanfear 2020) from May-11-2023 (http://vhgdownload.soe.ucsc.edu/goldenPath/wuhCor1/USHER_SARS-CoV-2/2023/05/11/public-2023-05-11.all.masked.nextclade.pangolin.pb.gz), which contains all ~7-million SARS-CoV-2 sequences that are available in public databases. To make these counts at a per-clade level, we first subsetted the mutation-annotated tree on all sequences for each Nexstrain clade (Aksamentov et al. 2021), retained only clades with at least 10^4 sequences, and then used the matUtils program distributed with USHER to extract the nucleotide mutations on every branch of the each clade-subsetted mutation-annotated tree. For the analyses by geographic location (Fig. 2), we subsetted on all sequences that began with 'USA' or 'England' as these were the two locations with the most publicly available sequences.

We then performed quality control by ignoring any branch that met any of the following criteria:

- it had more than four nucleotide mutations;
- it contained more than one nucleotide mutation that was a reversion to the Wuhan-Hu-1 reference sequence;
- it contained more than one nucleotide mutation that was a reversion to the founder sequence for that clade as provided at https://raw.githubusercontent.com/neherlab/SC2_variant_rates/7e738194a8c6592082f1caa9a6ca70cb68289790/data/clade_gts.json by Neher (2022);
- it contained more than one nucleotide mutation to the same codon.

The rationale for the first exclusion is that highly mutated branches are often indicative of sequencing errors or viral evolution in chronically infected humans, neither of which correspond to the pattern of typical SARS-CoV-2 transmission in acute infections. Because the virus's evolution is very densely sampled, only a small fraction of branches have more than four mutations (Supplementary Fig. S10). The rationale for the second and third exclusions is that excess reversions can arise from base-calling pipelines that erroneously call low-coverage sites as reference. We ignore branches with multiple nucleotide mutations to the same codon (this is very rare) because as detailed below our method is only designed to make estimates for mutations that represent single-nucleotide changes from the clade founder. Note also that

the mutation-annotated tree does not include insertion or deletion mutations, and so we only consider (and make estimates for) point mutations.

We then specified for exclusion certain mutations and sites that are prone to sequencing or base-calling errors. Specifically, we excluded

- the sites specified in Table S1 of Turakhia et al. (2020) as being error prone;
- sites 5629, 6851, 7328, 28095, and 29362 since they had very high error rates in some clades;
- the problematic sites listed at https://github.com/W-L/ProblematicSites_SARS-CoV2, which are masked in the pre-built mutation-annotated tree;
- for each clade, the clade-specific sites listed in https://github.com/jbloomlab/SARS2-mut-fitness/blob/main/data/usher_masked_sites.yaml, which are masked in the pre-built mutation-annotated tree;
- for each clade, any mutation that was a reversion from the clade founder to the Wuhan-Hu-1 reference, and the reverse complements of these mutations.

The last exclusion criteria are because some bioinformatics pipelines called low-coverage sites as reference.

See https://github.com/jbloomlab/SARS2-mut-fitness/blob/main/results/mutation_counts/agggregated.csv for the final counts of each nucleotide mutation in each clade; note that this file also contains excluded mutations.

4.4 Calculation of expected counts

To calculate the expected counts for each nucleotide mutation, we analyzed just the four-fold degenerate sites in each clade in an approach paralleling that of Bloom et al. (2023). Specifically, we identify all non-excluded four-fold degenerate sites in each clade founder. We then count nucleotide mutations just at those sites in each clade, and calculate the expected per-site number of mutations from nucleotide x to y as the total number of x to y mutations at four-fold degenerate sites divided by the number of four-fold degenerate sites with x as the parental identity. This analysis is done at the clade level for two reasons: referencing mutations to the clade founder (rather than the Wuhan-Hu-1 reference) limits problem with the approach that would arise at sites that substitute multiple times in the history of a sequence (since each clade is a relatively high-identity group multiple mutations at the same site within a clade are very rare), and because it is known that SARS-CoV-2 mutation rates vary somewhat among clades (Bloom et al. 2023; Ruis et al. 2022). We only retain clades with at least 5000 mutations at four-fold degenerate sites in order to avoid inaccurate estimates of expected counts due to low sampling of mutations.

4.5 Mutational effects from actual versus expected counts

To estimate the effects of mutations, we simply compare the expectation to the actual counts in the pre-built mutation-annotated tree. See https://github.com/jbloomlab/SARS2-mut-fitness/blob/main/results/expected_vs_actual_mut_counts/expected_vs_actual_mut_counts.csv for these expected versus actual counts on a per-clade basis; note that this file also includes counts at excluded sites.

To estimate the effects of mutations, we first sum the counts of all non-excluded nucleotide mutations that encode each amino acid mutation to convert the nucleotide counts to amino acid

counts. In doing this, we exclude any mutations that are not from the clade-founder codon identity: in other words, we ignore sequences with histories that involve multiple mutations at the same codon in the same clade (this is a caveat of the approach, although because each clade is relatively high identity it does not have a major effect). For the overall estimates reported in this paper, we also sum these counts across all retained clades; for the analyses in Fig. 2 we also make estimates without summing across clades and only for counts from sequences from specific geographic locations. We then compute the estimated fitness Δf of each mutation as simply the natural logarithm of the ratio of actual to expected counts after adding a pseudocount of $P - 0.5$ to each count, namely $\Delta f = \log\left(\frac{n_{\text{actual}+P}}{n_{\text{expected}+P}}\right)$.

Note that these mutation-effect estimates will have more statistical noise the smaller the value of the expected counts for each mutation. Therefore, we also track the expected counts alongside the estimates. In this paper, we only show estimates for mutations with expected counts of at least 10 unless otherwise noted. However, the figures link to interactive legends that allow adjustment of this threshold: larger values (e.g., 20 or more) will lead to slightly more accurate estimates but drop some mutations, lower values can be used if you need a noisier estimate for a mutation that has less than 10 expected counts.

See https://github.com/jbloomlab/SARS2-mut-fitness/blob/main/results/aa_fitness/aamut_fitness_all.csv for the estimates of amino acid mutation effects across all clades, and see https://github.com/jbloomlab/SARS2-mut-fitness/blob/main/results/aa_fitness/aamut_fitness_by_clade.csv for the clade-specific estimates. The all-clade estimates of mutation effects are what are shown in Fig. 3.

For the clade correlations plotting in Fig. 2, we only include clades with at least 5×10^5 expected counts across all sites, as only these clades have enough counts for reasonable per-clade estimates.

4.6 Mutation effects to amino acid fitnesses

For the final estimates of amino acid fitnesses shown in the heatmaps such as in Fig. 5, we need a single estimate for each amino acid. This is straightforward for sites that have the same amino acid identity in all clade founders: the ‘wildtype’ residue shared across all clades has a fitness of zero, and all other amino acids have fitnesses equal to the effect of mutating from the ‘wildtype’ to that amino acid. However, for sites that change amino acid identity between clade founders, things are more complicated and we need to take the extra step below.

For each clade we have estimated the change in fitness Δf_{xy} caused by mutating a site from amino acid x to y , where x is the amino acid in the clade founder sequence. For each such mutation, we also have n_{xy} which is the number of expected mutations from the clade founder amino acid x to y . These n_{xy} values are important because they give some estimate of our ‘confidence’ in the Δf_{xy} values: if a mutation has high expected counts (large n_{xy}) then we can estimate the change in fitness caused by the mutation more accurately, and if n_{xy} is small then the estimate will be much noisier.

However, we would like to aggregate the data across multiple clades to estimate amino acid fitness values at a site under the assumption that these are constant across clades. Things get complicated if not all clade founders have the same amino acid identity at a site. For instance, let us say at our site of interest, the clade founder amino acid is x in one clade and z in another clade. For each clade we then have a set of Δf_{xy} and n_{xy} values for the first clade (where y ranges over the 20 amino acids, including stop

codon, that are not x), and another set of up to 20 Δf_{zy} and n_{zy} values for the second clade (where y ranges over the 20 amino acids that are not z).

From these sets of mutation fitness changes, we would like to estimate the fitness f_x of each amino acid x , where the f_x values satisfy $\Delta f_{xy} = f_y - f_x$ (in other words, a higher f_x means higher fitness of that amino acid). When there are multiple clades with different founder amino acids at the site, there is no guarantee that we can find f_x values that precisely satisfy the above equation since there are more Δf_{xy} values than f_x values and the Δf_{xy} values may have noise (and in some cases even real shifts among clades due to epistasis). Nonetheless, we can try to find the f_x values that come closest to satisfying the above equation.

First, we choose one amino acid to have a fitness value of zero, since the scale of the f_x values is arbitrary and there are really only 20 unique parameters among the 21 f_x values (there are 21 amino acids since we consider stops, but we only measure differences among them, not absolute values). Typically if there was just one clade, we would set the wildtype value of $f_x = 0$ and then for mutations to all other amino acids y we would simply have $f_y = \Delta f_{xy}$. However, when there are multiple clades with different founder amino acids, there is no longer a well-defined ‘wildtype’. So we choose the most common non-stop parental amino acid for the observed mutations and set that to zero. In other words, we find x that maximizes $\sum_y n_{xy}$ and set that f_x value to zero.

Next, we choose the f_x values that most closely match the measured mutation effects, weighting more strongly mutation effects with higher expected counts (since these should be more accurate). Specifically, we define a loss function as:

$$L = \sum_x \sum_{y \neq x} n_{xy} (\Delta f_{xy} - [f_y - f_x])^2,$$

where we ignore effects of synonymous mutations (the $x \neq y$ term in second summand) because we are only examining protein-level effects. We then use numerical optimization to find the f_x values that minimize that loss L .

Finally, we would still like to report an equivalent of the n_{xy} values for the Δf_{xy} values that give us some sense of how accurately we have estimated the fitness f_x of each amino acid. To do that, we tabulate $N_x = \sum_y (n_{xy} + n_{yx})$ as the total number of mutations either from or to amino acid x as the ‘count’ for the amino acid. Amino acids with larger values of N_x should have more accurate estimates of f_x .

See https://github.com/jbloomlab/SARS2-mut-fitness/blob/main/results/aa_fitness/aa_fitness.csv for these overall amino acid fitness estimates.

4.7 Site numbering and protein naming

All sites are numbered according to the sequential Wuhan-Hu-1 reference numbering scheme, using the reference sequence at <http://hgdownload.soe.ucsc.edu/goldenPath/wuhCor1/bigZips/wuhCor1.fa.gz>. The protein annotations are taken from the associated GTF at <http://hgdownload.soe.ucsc.edu/goldenPath/wuhCor1/bigZips/genes/ncbiGenes.gtf.gz>. Those protein annotations refer to the polyproteins encoding the non-structural proteins as ORF1a and ORF1ab. To convert to from ORF1ab numbering/naming to the nsp-based naming (e.g., nsp1, nsp2, etc.) we use the conversions specified under ‘orf1ab_to_nsp’ in <https://github.com/jbloomlab/SARS2-mut-fitness/blob/main/config.yaml>, which are in turn taken from Theo Sanderson’s annotations at <https://github.com/theosanderson/Codon2Nucleotide/blob/main/src/App.js>.

4.8 Comparison to deep mutational scanning

Deep mutational scanning data were taken from published studies (Dadonaite et al. 2023; Starr et al. 2022b; Flynn et al. 2022; Flynn et al. 2023; Iketani et al. 2022a), using the data at the links specified under the 'dms_datasets' key in <https://github.com/jbloomlab/SARS2-mut-fitness/blob/main/config.yaml>. For the spike deep mutational scanning (Dadonaite et al. 2023) we only included mutations with 'times seen' values of at least three in the deep mutational scanning. The RBD data (Starr et al. 2022b) include measurements for two phenotypes (ACE2 affinity and RBD expression), and one of the Mpro studies (Flynn et al. 2022) includes measurements for three different phenotypes in yeast (growth, FRET, and transcription factor activity). Figs. 4, S5, and S6 show the effect averaged across all phenotypes measured by each of these studies. For plots that break the correlations out by phenotype, see https://jbloomlab.github.io/SARS2-mut-fitness/dms_S_all_corr.html and https://jbloomlab.github.io/SARS2-mut-fitness/dms_nsp5_all_corr.html.

4.9 Comparison to dN/dS and other mutation-effect prediction algorithms

For the comparison to the dN/dS approaches shown in Supplementary Fig. S5, we used the dN/dS values available at <https://github.com/spond/SARS-CoV-2-variation> (Martin et al. 2021) for all SARS-CoV-2 sequences, which were calculated using the FEL approach (Kosakovsky Pond and Frost 2005). Specifically, the file <https://raw.githubusercontent.com/spond/SARS-CoV-2-variation/master/windowed-sites-fel-all.csv> contains dN/dS estimates made using large-scale sequence sets from GISAID in 3-month windows starting with the earliest sequences from December 2019 and continuing up until 31 January 2022. We averaged the dN/dS values for each site over all three month windows, and analyzed those time-window averaged dN/dS values. In order to ensure comparability in the sequence sets used, for the comparisons in Supplementary Fig. S5, we used fitness estimates from our approach made using only sequences in the mutation-annotated tree as of 31 January 2022. We restricted our analysis to this timeframe because the large-scale dN/dS analyses at <https://github.com/spond/SARS-CoV-2-variation> (which use state-of-the-art methods) were only available for that range of dates.

The dN/dS ratios only provide a single number for each site, which cannot be directly compared to either the mutation-effect estimates or the deep mutational scanning, which estimate the effects of individual amino acid mutations. We therefore computed site-summary metrics of the mutation-effect estimates and the deep mutational scanning as the average effect of all measured amino acid mutations at each site, excluding stop codons. The correlations in Supplementary Fig. S5 are with those site-summary metrics.

We also compared both our mutation-effect estimates and the spike deep mutational scanning measurements (Dadonaite et al. 2023) to predictions from three other algorithms:

- the EpiScores reported by Maher et al. (2022),
- the DCA mutability scores reported by Rodriguez-Rivas et al. (2022), and
- the EVE scores reported by Thadani et al. (2022).

These comparisons are shown in Supplementary Fig. S6. The Maher et al. (2022) and Thadani et al. (2023) studies report mutation-level predictions and so are compared directly to

the deep mutational scanning our mutation-effect estimates; Rodriguez-Rivas et al. (2023) report only site-level metrics and so are compared to site-summary metrics as for the dN/dS analysis.

4.10 Derivation of relationship between actual to expected count ratio and viral fitness

The ratio of actual to expected counts that we calculate in this paper is related to the probability that we observe a viral lineage containing an occurrence of a specific mutation among sequenced human SARS-CoV-2. This probability depends on three factors: the fitness effect of the mutation, the fraction of all SARS-CoV-2 viruses that are sequenced (sampling intensity), and the growth dynamics of the viral population. In the supplementary appendix, we derive the approximate relationship between this probability as a function of the fitness cost s and sampling intensity ϵ for deleterious mutations for both a constant and exponentially growing viral population.

We show that for a constant viral population size, the probability of observing a lineage containing a deleterious mutation with cost s is roughly $\frac{\epsilon}{s+\epsilon}$ when $s^2\epsilon\ll 1$, and more weakly dependent on s for smaller fitness costs (when $s^2\epsilon\gg 1$). The intuitive explanation is that the average size of a mutant lineage with fitness cost s is $1/s$ and we basically ask whether we sample the lineage before it disappears. If we sample more intensely (larger ϵ), whether a lineage gets sampled depends primarily on the stochastic dynamics and little on the fitness effect. With a typical sampling intensity for SARS-CoV-2 between 1/1000 and 1/100, this means our approach is sensitive to fitness effects larger than a few percent per serial interval; mutations with fitness costs smaller than that will not show an appreciable difference from neutral mutations in their ratio of actual to expected accounts.

In an exponentially growing population, the probability of observing a mutant lineage with fitness cost s again scales as $\sim \frac{\epsilon}{\epsilon+s}$ if $sT > 1$, where T is the time over which the variant has expanded. If T is \sim months, that is 20 generations, which again corresponds to s of at least a few percent for $sT > 1$. For mutations with smaller fitness costs, the dependence scales more as $\sim \epsilon(1-sT)$.

Overall, these calculations indicate that for multiple different growth dynamics of the viral population, the ratio of expected to actual counts will scale inversely with the fitness cost of deleterious mutations for mutations with costs that exceed a few percent. Note that the approach we use in this paper does not account for variation in sampling intensity across space or time, does not attempt to adjust for changes in viral growth dynamics over time, uses the heuristic formula of calculating the effect as the log ratio of counts, and applies this same formula to all mutations regardless of whether they are deleterious, neutral, or beneficial. A more complete derivation might try to calculate the fitness effects from the full distribution of lineage sizes more rigorously and incorporate information about the sampling intensity and viral growth dynamics. However, such a derivation (if possible at all) is beyond the scope of this study, and we also note that good empirical data is generally lacking to precisely account for sampling intensity and viral growth dynamics over the full span of time and space from which the sequences we analyze are drawn. The key point of the derivations for our current study is simply that our approach should be sensitive to detecting the effects of mutations with fitness costs greater than a few percent.

Supplementary data

Supplementary data are available at Virus Evolution online.

Acknowledgements

We thank Angie Hinrichs for providing the pre-built mutation-annotated trees and thank the UShER team for promptly answering and addressing GitHub issues related to use of this package and its pre-built trees. We thank the sequence submitters to GISAID, who are listed in the tables cited in the Methods section. The work of JDB was supported in part by the NIH/NIAID under Grant Nos. U19AI171399 and R01AI141707 and Contract No. 75N93021C00015. JDB is an Investigator of the Howard Hughes Medical Institute.

Conflict of interest statement: J.D.B. consults Apriori Bio, Aerium Therapeutics, Invivyd, the Vaccine Company, GSK, and Pfizer on topics related to viral evolution. J.D.B. receives royalty payments as an inventor on Fred Hutch licensed patents related to deep mutational scanning of viral proteins.

References

- Abdool Karim, S. S. and T., de Oliveira (2021) 'New SARS-CoV-2 variants—clinical, public health, and vaccine implications', *New England Journal of Medicine*, **384**: 1866–1868.
- Acevedo, A., L., Brodsky and R., Andino (2014) 'Mutational and fitness landscapes of an RNA virus revealed through population sequencing', *Nature*, **505**: 686–690.
- Aksamentov, I. et al. (2021) 'Nextclade: clade assignment, mutation calling and quality control for viral genomes', *Journal of Open Source Software*, **6**: 3773.
- Beale, R. C. et al. (2004) 'Comparison of the differential context-dependence of DNA deamination by APOBEC enzymes: correlation with mutation spectra in vivo', *Journal of Molecular Biology*, **337**: 585–596.
- Bhatt, P. R. et al. (2021) 'Structural basis of ribosomal frameshifting during translation of the SARS-CoV-2 RNA genome', *Science*, **372**: 1306–1313.
- Bloom, J. D. et al. (2023) 'Evolution of the SARS-CoV-2 mutational spectrum', *Molecular Biology and Evolution*, **40**: msad085.
- Cao, Y. et al. (2022a) 'Imprinted SARS-CoV-2 humoral immunity induces convergent Omicron RBD evolution', *Nature*, **614**: 521–529.
- Cao, Y. et al. (2022b) 'Rational identification of potent and broad sarbecovirus-neutralizing antibody cocktails from SARS convalescents', *Cell reports*, **41**: 111845.
- Dadonaite, B. et al. (2023) 'A pseudovirus system enables deep mutational scanning of the full SARS-CoV-2 spike', *Cell*, **186**: 1263–1278.
- Dawson, W. O. and S. Y., Folimonova (2013) 'Virus-based transient expression vectors for woody crops: a new frontier for vector design and use', *Annual Review of Phytopathology*, **51**: 321–337.
- DeGrace, M. M. et al. (2022) 'Defining the risk of SARS-CoV-2 variants on immune protection', *Nature*, **605**: 640–652.
- De Maio, N. et al. (2021) 'Mutation rates and selection on synonymous mutations in SARS-CoV-2', *Genome Biology and Evolution*, **13**: evab087.
- Flynn, J. M. et al. (2022) 'Comprehensive fitness landscape of SARS-CoV-2 Mpro reveals insights into viral resistance mechanisms', *eLife*, **11**: e77433.
- Flynn, J. M. et al. (2023) 'Systematic analyses of the resistance potential of drugs targeting SARS-CoV-2 main protease', *ACS Infectious Diseases*, **9**: 1372–1386.
- Greaney, A. J. et al. (2021) 'Complete mapping of mutations to the SARS-CoV-2 spike receptor-binding domain that escape antibody recognition', *Cell Host & Microbe*, **29**: 44–57.
- Greaney, A. J., T. N., Starr and J. D., Bloom (2022) 'An antibody-escape estimator for mutations to the SARS-CoV-2 receptor-binding domain', *Virus Evolution*, **8**: veac021.
- Harvey, W. T. et al. (2021) 'SARS-CoV-2 variants, spike mutations and immune escape', *Nature Reviews Microbiology*, **19**: 409–424.
- Hiscox, J. A. et al. (2021) 'Shutting the gate before the horse has bolted: is it time for a conversation about SARS-CoV-2 and antiviral drug resistance?', *Journal of Antimicrobial Chemotherapy*, **76**: 2230–2233.
- Huddleston, J. et al. (2020) 'Integrating genotypes and phenotypes improves long-term forecasts of seasonal influenza A/H3N2 evolution', *Elife*, **9**: e60067.
- Huston, N. C. et al. (2021) 'Comprehensive in vivo secondary structure of the SARS-CoV-2 genome reveals novel regulatory motifs and mechanisms', *Molecular Cell*, **81**: 584–598.
- Iketani, S. et al. (2022a) 'Functional map of SARS-CoV-2 3CL protease reveals tolerant and immutable sites', *Cell Host & Microbe*, **30**: 1354–1362.
- Iketani, S. et al. (2022b) 'Multiple pathways for SARS-CoV-2 resistance to nirmatrelvir', *Nature*, **613**: 558–564.
- Jungreis, I. et al. (2021) 'Conflicting and ambiguous names of overlapping ORFs in the SARS-CoV-2 genome: a homology-based resolution', *Virology*, **558**: 145–151.
- Koelle, K. and D. A., Rasmussen (2015) 'The effects of a deleterious mutation load on patterns of influenza A/H3N2's antigenic evolution in humans', *Elife*, **4**: e07361.
- Kosakovsky Pond, S. L. and S. D., Frost (2005) 'Not so different after all: a comparison of methods for detecting amino acid sites under selection', *Molecular Biology and evolution*, **22**: 1208–1222.
- Kryazhimskiy, S. and J. B., Plotkin (2008) 'The population genetics of dN/dS', *PLoS Genetics*, **4**: e1000304.
- Kuo, L. and P. S., Masters (2013) 'Functional analysis of the murine coronavirus genomic RNA packaging signal', *Journal of Virology*, **87**: 5182–5192.
- Lanfear, R. (2020) 'A global phylogeny of SARS-CoV-2 sequences from GISAID', *Zenodo*.
- Lee, J. M. et al. (2018) 'Deep mutational scanning of hemagglutinin helps predict evolutionary fates of human H3N2 influenza variants', *Proceedings of the National Academy of Sciences*, **115**: E8276–E8285.
- Lee, B. et al. (2022) 'Inferring effects of mutations on SARS-CoV-2 transmission from genomic surveillance data', *MedRxiv*, 2021–12.
- Liu, D. X. et al. (2014) 'Accessory proteins of SARS-CoV and other coronaviruses', *Antiviral Research*, **109**: 97–109.
- Liu, Y. et al. (2022) 'A live-attenuated SARS-CoV-2 vaccine candidate with accessory protein deletions', *Nature Communications*, **13**: 1–14.
- Llanes, A. et al. (2020) 'Betacoronavirus genomes: how genomic information has been used to deal with past outbreaks and the COVID-19 pandemic', *International Journal of Molecular Sciences*, **21**: 4546.
- Luksza, M. and M., Lässig (2014) 'A predictive fitness model for influenza', *Nature*, **507**: 57–61.
- Maher, M. C. et al. (2022) 'Predicting the mutational drivers of future SARS-CoV-2 variants of concern', *Science Translational Medicine*, **14**: eabk3445.
- Martin, D. P. et al. (2021) 'The emergence and ongoing convergent evolution of the SARS-CoV-2 N501Y lineages', *Cell*, **184**: 5189–5200.
- McBroome, J. et al. (2021) 'A daily-updated database and tools for comprehensive SARS-CoV-2 mutation-annotated trees', *Molecular Biology and Evolution*, **38**: 5819–5824.

- McGrath, M. et al. (2022) 'SARS-CoV-2 variant spike and accessory gene mutations alter pathogenesis', *Proceedings National Academy of Sciences USA*, **119**: e2204717119.
- Moghadasi, S. A. et al. (2023) 'Transmissible SARS-CoV-2 variants with resistance to clinical protease inhibitors', *Science Advances*, **9**: eade8778.
- Mölder, F. et al. (2021) 'Sustainable data analysis with snakemake', *F1000Research*, **10**: 33.
- Moulana, A. et al. (2022) 'Compensatory epistasis maintains ACE2 affinity in SARS-CoV-2 Omicron BA. 1', *Nature Communications*, **13**: 7011.
- Neher, R. A. (2022) 'Contributions of adaptation and purifying selection to SARS-CoV-2 evolution', *Virus Evolution*, **8**: veac113.
- Nielsen, R. and Z., Yang (1998) 'Likelihood models for detecting positively selected amino acid sites and applications to the HIV-1 envelope gene', *Genetics*, **148**: 929–936.
- Obermeyer, F. et al. (2022) 'Analysis of 6.4 million SARS-CoV-2 genomes identifies mutations associated with fitness', *Science*, **376**: 1327–1332.
- Pollock, D. D., G. Thiltgen, and R. A. Goldstein, 2012 Amino acid coevolution induces an evolutionary Stokes shift. *Proceedings of the National Academy of Sciences* **109**: E1352–E1359.
- Rappazzo, C. G. et al. (2021) 'Broad and potent activity against SARS-like viruses by an engineered human monoclonal antibody', *Science*, **371**: 823–829.
- Rochman, N. D., Y.I., Wolf and E. V., Koonin (2022) 'Molecular adaptations during viral epidemics', *EMBO reports*, **23**: e55393.
- Rodriguez-Rivas, J. et al. (2022) 'Epistatic models predict mutable sites in SARS-CoV-2 proteins and epitopes', *Proceedings of the National Academy of Sciences*, **119**: e2113118119.
- Ruis, C. et al. (2022) 'Mutational spectra distinguish SARS-CoV-2 replication niches', *BioRxiv*, 2022–09.
- Sadykov, M. et al. (2021) 'Short sequence motif dynamics in the SARS-CoV-2 genome suggest a role for cytosine deamination in CpG reduction', *Journal of Molecular Cell Biology*, **13**: 225–227.
- Shah, P., D. M., McCandlish and J. B., Plotkin (2015) 'Contingency and entrenchment in protein evolution under purifying selection', *Proceedings of the National Academy of Sciences*, **112**: E3226–E3235.
- Shu, Y. and J., McCauley (2017) 'GISAID: Global initiative on sharing all influenza data—from vision to reality', *Eurosurveillance*, **22**: 30494.
- Silvas, J. A. et al. (2021) 'Contribution of SARS-CoV-2 accessory proteins to viral pathogenicity in K18 human ACE2 transgenic mice', *Journal of Virology*, **95**: e00402–21.
- Spielman, S. J. and C. O., Wilke (2015) 'The relationship between dN/dS and scaled selection coefficients', *Molecular Biology and Evolution*, **32**: 1097–1108.
- Starr, T. N. et al. (2018) 'Pervasive contingency and entrenchment in a billion years of Hsp90 evolution', *Proceedings of the National Academy of Sciences*, **115**: 4453–4458.
- Starr, T. N. et al. (2020) 'Deep mutational scanning of SARS-CoV-2 receptor binding domain reveals constraints on folding and ACE2 binding', *Cell*, **182**: 1295–1310.
- Starr, T. N. et al. (2021) 'SARS-CoV-2 RBD antibodies that maximize breadth and resistance to escape', *Nature*, **597**: 97–102.
- Starr, T. N. et al. (2022a) 'Shifting mutational constraints in the SARS-CoV-2 receptor-binding domain during viral evolution', *Science*, **377**: 420–424.
- Starr, T. N. et al. (2022b) 'Deep mutational scans for ACE2 binding, RBD expression, and antibody escape in the SARS-CoV-2 Omicron BA. 1 and BA. 2 receptor-binding domains', *PLoS Pathogens*, **18**: e1010951.
- Su, Y. C. et al. (2020) 'Discovery and genomic characterization of a 382-nucleotide deletion in ORF7b and ORF8 during the early evolution of SARS-CoV-2', *mBio*, **11**: e01610–20.
- Sun, K. et al. (2023) 'Rapidly shifting immunologic landscape and severity of SARS-CoV-2 in the Omicron era in South Africa', *Nature Communications*, **14**: 246.
- Syed, A. M. et al. (2021) 'Rapid assessment of SARS-CoV-2-evolved variants using virus-like particles', *Science*, **374**: 1626–1632.
- Tao, K. et al. (2021) 'SARS-CoV-2 antiviral therapy', *Clinical Microbiology reviews*, **34**: e00109–21.
- Thadani, N. N. et al. (2022) 'Learning from pre-pandemic data to forecast viral antibody escape', *BioRxiv*, 2022–07.
- Thorne, L. G. et al. (2022) 'Evolution of enhanced innate immune evasion by SARS-CoV-2', *Nature*, **602**: 487–495.
- Turakhia, Y. et al. (2020) 'Stability of SARS-CoV-2 phylogenies', *PLoS Genetics*, **16**: e1009175.
- Turakhia, Y. et al. (2021) 'Ultrafast Sample placement on Existing tRees (USHER) enables real-time phylogenetics for the SARS-CoV-2 pandemic', *Nature Genetics*, **53**: 809–816.
- Tzou, P. L. et al. (2022) 'Coronavirus Resistance Database (CoV-RDB): SARS-CoV-2 susceptibility to monoclonal antibodies, convalescent plasma, and plasma from vaccinated persons', *Plos one*, **17**: e0261045.
- VanderPlas, J. et al. (2018) 'Altair: interactive statistical visualizations for Python', *Journal of Open Source Software*, **3**: 1057.
- V'kovski, P. et al. (2021) 'Coronavirus biology and replication: implications for SARS-CoV-2', *Nature Reviews Microbiology*, **19**: 155–170.
- Wang, J. et al. (2011) 'Genetic diversity in the 3' terminal 4.7-kb region of grapevine leafroll-associated virus 3', *Phytopathology*, **101**: 445–450.
- Weisblum, Y. et al. (2020) 'Escape from neutralizing antibodies by SARS-CoV-2 spike protein variants', *eLife*, **9**: e61312.
- Yadav, R. et al. (2022) 'Biochemical and structural insights into SARS-CoV-2 polyprotein processing by Mpro', *Science Advances*, **8**: eadd2191.
- Yang, Z. and R., Nielsen (2000) 'Estimating synonymous and nonsynonymous substitution rates under realistic evolutionary models', *Molecular Biology and evolution*, **17**: 32–43.
- Zanini, F. et al. (2017) 'In vivo mutation rates and the landscape of fitness costs of HIV-1', *Virus Evolution*, **3**: vex003.