A macroecological perspective on genetic diversity in the human gut microbiome

William R. Shoemaker[1,*]

1 Department of Ecology and Evolutionary Biology, University of California, Los Angeles, California, USA

*** Contact:** `williamrshoemaker@gmail.com`

Abstract

While the human gut microbiome has been intensely studied, we have yet to obtain a sufficient understanding of the genetic diversity harbored by this complex system. This difficulty is compounded by the confluence of forces operating within the gut, where the dynamics of a given allele can be driven by evolutionary forces (drift, selection, etc.) or by the ecology of the strain on which said allele is a constituent. To make progress towards this goal, I identified and characterized statistically invariant patterns of genetic diversity in a large cohort of unrelated hosts across 22 prevalent microbial species. The existence of said patterns allows for the set of feasible models to be pared down to those that can reproduce observed patterns without the use of free parameters. Using a feasible distribution of allele frequencies, I evaluated the extent that a model of ecological dynamics was capable of predicting the fraction of hosts harboring a given allele (i.e., prevalence), a statistic that allows for the effect of sampling to be taken into account. I found that the Stochastic Logistic Model (SLM), an ecological model that has been shown to successfully predict the dynamics of strains, succeeded across all species. The accuracy of the predictions was correlated with independent estimates of strain structure, validating the ecological implications of the SLM. This work indicates that a large number of detectable genetic variants in the human gut are subject to ecological, as opposed to evolutionary, dynamics.

Introduction

The human gut microbiome harbors astounding levels of genetic diversity. Hundreds to thousands of species continually reproduce in a typical host, dynamics that culminate in the acquisition of $\sim 10^9$ *de novo* mutations each day [1]. Due to the comparatively brief generation time of microbes in the human gut [2], those mutations that are beneficial can rapidly fix on a timescale of days to months [1,3–8]. Such evolutionary dynamics have the capacity to alter the genetic composition of a species within a given host. However, while all genetic diversity ultimately arises due to mutation, this actuality does not mean that all the genetic variants observed in the human gut are necessarily subject to evolutionary dynamics.

It is becoming increasingly apparent that for many species a large number of genetic variants do not fix or become extinct within a given host. Instead, these variants fluctuate at intermediate frequencies on timescales ranging from months to years [1,3,9–12]. This pattern arises due to the co-occurrence of a few ($\sim \mathcal{O}(1-4)$)

genetically diverged strains belonging to the same species, a process known as oligocolonization [3, 10]. This form of ecology is not simply a descriptive detail, as it has been proposed that strains are the relevant unit of interaction and dynamics in microbial systems [13]. Thus, the dynamics of the genetic variants that comprise a given strain are subject to exogenous and endogenous ecological forces [14–16]. Such dynamics are a clear departure from the dynamics captured by standard origin-fixation models from population genetics, suggesting that measures of genetic diversity estimated within the human gut are shaped by ecology alongside evolution.

This confluence of ecological and evolutionary dynamics in the human gut calls into question the feasibility of attempting to characterize its typical properties. Many studies tackle such complexity by examining individual species [6, 17–19] or by searching for genetic differences between species within the human gut [10, 11, 20–24]. While such approaches are useful for identifying unique species that are reasonable candidates for identifying a cause of specific conditions such as disease or the conference of the ability to metabolize certain resources, it is difficult to leverage isolated observations to identify the general case. This disproportionate focus on individual species and differences between species can lead to idiosyncratic notions about the typical dynamics of a genetic variant in the microbiome.

Instead, it is reasonable to identify genetic patterns that are statistically invariant across species. Such an approach provides the empirical motivation necessary to identify mathematical models that could potentially explain said patterns and aid in the identification of feasible ecological or evolutionary dynamics [9, 25–27]. In recent years, substantial progress has been made towards characterizing the typical microbial evolutionary dynamics across species that operate within the human gut [3, 4, 7, 8, 28–30]. Using this approach it is possible to leverage the richness of species in the human gut microbiome, where each species can be viewed as a draw from an unknown distribution and, as an ensemble, be used to identify patterns that are statistically invariant [31, 32].

To identify such patterns, it is useful to examine attempts that succeed in paring down the complexity of the gut. Notable recent examples come from the discipline of macroecology, which has succeeded at predicting patterns of microbial diversity and abundance at the species level across disparate environments, including the gut microbiome [33–37]. Such an approach emphasizes the benefits of identifying patterns of diversity that are statistically invariant that can then serve as motivation for obtaining quantitative predictions derived from first principles that can account for experimental details (e.g., finite sampling). Recent work suggests that this conjecture may be the case, as macroecological patterns that hold for species also extend to the strain level [37]. This result suggests that the effect of widespread strain structure on genetic diversity may not be as convoluted as a surface level examination would suggest. Namely, models of the ecological dynamics of strains may aid in the search and interpretation of universal patterns of genetic diversity in the human gut microbiome.

In this study, I sought to identify and explain patterns of genetic diversity that are statistically invariant across species in the human gut microbiome. The existence of said patterns provides evidence of universality across species, the impetus necessary to identify a plausible model of allele frequency dynamics. That model is the Stochastic Logistic Model (SLM) of growth, a model that describes the ecological dynamics of microbial communities at the level of species as well as strains [35, 37]. To evaluate the feasibility of the SLM while simultaneously accounting for the effects of sampling, I obtained predictions of the fraction of hosts harboring a given allele (i.e., prevalence) using zero free parameters. To confirm that the success of the SLM was due to the presence of strains, I evaluated the extent that alternative evolutionary models could explain the success of the SLM and determined that the existence of strain structure within a species was correlated with the accuracy of my predictions.

Results

Patterns of genetic diversity are statistically invariant across species

In order to determine whether the task of predicting patterns of genetic diversity across species is even feasible, it is necessary to first investigate the degree of similarity that common measures of genetic diversity share across phylogenetically distant species. Determining the degree that distributions exhibit qualitatively similar forms across species is an appropriate first step. This manner of visualization, known as a data collapse, allows one to assess whether it is reasonable to assume that similar dynamics, regardless of their ultimate causes, are driving the behavior of disparate systems [38–40]. Such an analysis also provides the benefit of allowing for the identification of empirical patterns for subsequent investigation. To determine whether there is evidence that the distributions of measures of genetic diversity have qualitative similar forms across species, I compiled all alleles among all hosts and species that passed the quality control pipeline (Materials and Methods). The total number of processed hosts ranged from $108 - 371$ across species, with a median of 182 (Fig. 1a). The total number of sites ranged from $39 - 37,204$ across species, with a median of 10,269 synonymous and 5,204 nonsynonymous sites (Fig. 1b, S2b).

First, I examined the distribution of within-host allele frequencies across all hosts, as the similarity of this distribution across species informs us as to whether one should expect that the probability distribution of allele frequencies of a given site will have a qualitatively similar form across species. To use population genetics terminology, I examined ensembles of single-site frequency spectra [41, 42]. The macroecological analogue of this object is the distribution of relative abundances of a given species or strain across hosts, the Abundance Fluctuation Distribution [35], an observation that provides insight into how one can methodologically proceed. The distribution of within-host allele frequencies has a similar statistical form across species in the human gut for synonymous (Fig. 1c) and nonsynonymous sites (Fig. S2c), suggesting that a single model may be sufficient. Indeed, the distribution is fairly well-captured by a previously derived form of the gamma distribution [35], suggesting that it is informative to examine models grounded in ecological and evolutionary first principles that predict a gamma distribution of allele frequencies.

Beyond the shape of the distribution of within-host allele frequencies, the statistical moments of within-host frequencies across hosts also exhibit qualitatively similar forms. The distribution of the mean frequency of a given allele across hosts ($\bar{f}$) tends to overlap across species for both synonymous (Fig. 1d) and nonsynonymous sites (Fig. S2d), suggesting that it may be a useful quantity for further inspection. Using this observation, one can examine the relationship between $\bar{f}$ and higher-order central statistical moments. The mean and variance of ensembles of random variables frequently follow linear relationships on a logarithmic scale across biological systems, most notably in patterns of biodiversity in ecological communities [35, 43, 44] but also among population genetic patterns [45–47]. This relationship, known as Taylor's Law [48], implies that the mean and variance of an ensemble of allele frequencies are not independent, reducing the number of parameters necessary to characterize the dynamics of the system [35]. Furthermore, if the variances scales quadratically with the mean then the existence of the relationship implies that the coefficient of variation of f is constant across sites, an observation that can considerably reduce the difficulty of characterizing the dynamics of the system.

By examining the relationship between $\bar{f}$ and the variance of f (σ_f^2), it is clear that the two moments follow a linear relationship on a logarithmic scale for low values of $\bar{f}$ ($0 < \bar{f} \lesssim 0.35$). The exponent of this relationship is ~ 1.96 (bootstrapped 95% CI from

10,000 samples: [1.85, 2.06]) for synonymous sites, a value that is remarkably close to two, implying that the coefficient of variation in f can be viewed as a constant for the range of $\bar{f}$ where the relationship is linear (Fig. 1e). The exponent is slightly reduced for nonsynonymous sites (~ 1.83, 95% CI [1.72, 1.93]; Fig. S2e), suggesting that the variance increases with the mean at a slower rate relative to synonymous sites. Given the evidence that purifying selection is pervasive across species within the human gut [3, 11, 30, 49], it is likely that the typical allele at a nonsynonymous site is deleterious, reducing its variance across hosts for low values of $\bar{f}$. However, the linear relationship does not extend to high values of $\bar{f}$. Given that f is, by definition, a bounded quantity (i.e., $0 \leq f \leq 1$), it is possible that the relationship between $\bar{f}$ and σ_f^2 for values of $\bar{f} \approx 1$ is governed by the upper bound on f. To determine whether this is plausible, I plotted the maximum possible value of σ_f^2 for a given value of $\bar{f}$ constrained on the lower and upper bounds of f. This relationship, known as the Bhatia–Davis inequality, is defined as $\sigma_f^2 \leq (\max(f) - \bar{f})(\bar{f} - \min(f)) = (1 - \bar{f})\bar{f}$ [50]. The empirical relationship between $\bar{f}$ and σ_f^2 follows the Bhatia–Davis inequality across species for $\bar{f} \approx 1$, suggesting that the empirical relationship can be explained solely by the mathematical constraints on f, making the relationship uninformative for the purpose of identifying universal evolutionary or ecological explanations within certain parameter regimes.

It is worth noting that an exponent of two can emerge as the result of conducting an independent and identically distributed sampling process on a sufficiently skewed distribution [51, 52]. However, similar to ecological analyses of the relationship between the mean and variance of species abundances [35, 44, 53], the observation that the observed mean allele frequency varies by close to two orders of magnitude suggests that this patterns reflects a true scaling relationship is and not an artifact of sampling.

Finally, I turned my attention to the fraction of hosts in which a given allele is present. I found that the number of hosts in which a typical allele is present is small for synonymous (Fig. 1f) and nonsynonymous sites (Fig. S2f). Alternatively stated, the fraction of hosts harboring a given allele (i.e., prevalence). Furthermore, the fraction of alleles with a given value of prevalence tends to decay as prevalence increases. This assemblage of results, and their existence across synonymous and nonsynonymous sites provides the empirical basis for quantitative predictions.

Predicting the prevalence of an allele across hosts

The existence of multiple patterns of allelic diversity that are universal across evolutionarily distant bacterial species suggests that comparable dynamics are ultimately responsible. This conjecture presents us with the task of identifying appropriate candidate models of allele frequencies. The shape of the rescaled distribution of allele frequencies suggests that a gamma distribution is a suitable candidate, requiring that any feasible model of ecological or evolutionary dynamics generate said distribution or a distribution of similar form. Disparate approaches can be used to identify such a model. However, given the consistency of the patterns I observed, it is appropriate to focus on models that solely contain parameters that can be measured from empirical data rather than estimating free parameters through statistical inference. Such an approach allows one to make quantitative predictions that can serve as a first step towards the construction of a comprehensive theory [54].

We begin with the assumption that the frequency dynamics of a typical allele are driven by the ecological dynamics of the strain on which said allele resides. Appropriate Langevin equations (i.e., Stochastic Differential Equations) that capture ecological dynamics can be identified, though regardless of the underlying dynamics, the form of the available data constrains the types of solutions that can be evaluated. Given that temporal observational metagenomic data is at-present restricted to a small number of

hosts for the human gut microbiome, I focused on static samples taken at a single timepoint across a large number of unrelated hosts. This detail means that time-dependent solutions of the probability density of f cannot be evaluated across hosts, so I instead focused on stationary solutions and limiting cases where the time-dependence can be captured by parameters that can, in principle, be estimated.

To evaluate the predictive capacity of ecological dynamics, I examined a Langevin describing the dynamics of an allele on the background of a colonizing strain that is subject to self-limiting growth. This equation, known as the Stochastic Logistic Model of growth (SLM), has recently been shown to describe a range of macroecological patterns for microbial communities across disparate environments [35, 55, 56] as well as the temporal dynamics of strains within a single human host [37]. The SLM is defined as

$$\frac{\partial f}{\partial t} = \frac{f}{\tau_i}\left(1 - \frac{f}{K_i}\right) + \sqrt{\frac{\sigma_{\tau_i}}{\tau_i}} f \cdot \eta(t) \tag{1}$$

The terms $\frac{1}{\tau_i}$, K_i, and σ_{τ_i} represent the intrinsic growth rate of the strain, carrying capacity, and the coefficient of variation of growth rate fluctuations of the ith strain. The term $\eta(t)$ is a Brownian noise term where $\langle \eta(t) \rangle = 0$ and $\langle \eta(t)\eta(t') \rangle = \delta(t - t')$ [57]. By obtaining the equivalent partial differential equation of the probability of f (i.e., the Fokker-Planck equation) and deriving the stationary solution (S1 Text), one finds that the SLM predicts that the frequency of a given allele on the background of a strain follows a gamma distribution

$$f_{\text{SLM}} \sim \text{Gamma}\left(2\sigma_{\tau_i}^{-1} - 1, \frac{2}{K_i \sigma_{\tau_i}}\right) \tag{2}$$

This distribution is fully characterized by the mean frequency and the squared inverse of the coefficient of variation ($\beta = \frac{\bar{f}^2}{\sigma_f^2}$) across hosts [35, 58]

$$f \sim \text{Gamma}\left(\beta, \frac{\beta}{\bar{f}}\right) \tag{3}$$

The similarity in the shape of the distribution of $\bar{f}$ across species and the relationship between $\bar{f}$ and σ_f^2 suggests that the mean and variance are appropriate quantities to evaluate the predictive capacity of each model (Fig. 1d,e, S2d,e). However, $\bar{f}$ and σ_f^2 are parameters of the SLM (i.e., inputs), meaning that the SLM cannot be used to generate predictions of these moments. To test the applicability of the SLM, I chose to examine the fraction of hosts harboring a given allele (i.e., prevalence), a quantity that has been used to examine microbial ecology and evolution across systems [35, 59], including the human gut microbiome [3, 7, 35, 60]. Furthermore, this statistic incorporates the sampling effort of each site in each host (i.e., total depth of coverage), allowing for the SLM to be evaluated using experimental details. The predicted value of prevalence is defined as

$$\langle \varrho \rangle = 1 - \frac{1}{M} \sum_{m=1}^{M} \Pr[0 | D_m, \beta, \beta/\bar{f}] \tag{4}$$

where $\Pr[0|D_m, \cdots]$ is the probability of observing zero reads assigned to an allele within the mth host given a total depth of coverage D_m. While Eq. 5 is correct, the pipeline enacted a cutoff for the total depth of coverage, resulting in a coverage cutoff for a given minor allele ($A_{\text{cutoff}} = 10$; Fig. S1). This cutoff effectively truncates the sampling distribution of minor allele read counts, meaning that read counts for an allele less than the specified cutoff are observed as zeros. This issue can be readily accounted

for by summing over the probabilities of observing alternative allele read counts up until the cutoff

$$\langle \varrho \rangle = 1 - \frac{1}{M} \sum_{m=1}^{M} \sum_{A_i=0}^{A_{\text{cutoff}}-1} \Pr[A_i | D_m, \beta, \beta/\bar{f}] \tag{5}$$

The choice of prevalence also allows one to evaluate the extent that the gamma distribution can recapitulate empirical relationships between genetic quantities. One such relationship is that the prevalence of a species (equivalently known as *occupancy* in macroecology) should increase with its mean abundance across communities [61], a pattern that has been found to exhibit statistically similar forms at the species level across microbial systems [62–64] and can be quantitatively explained through the existence of macroecological laws [35].

The Stochastic Logistic Model succeeds at predicting allelic prevalence

By examining the relationship between observed and predicted allele prevalence, I found that the SLM generally succeeded among synonymous and nonsynonymous sites (Figs. 2a, S3, S4). The fraction of sites with relative errors ≤ 0.1 ($\leq 10\%$) ranged from $0.19 - 0.6$ across species, suggesting that a considerable fraction of genetic variants in the human gut are driven by ecological dynamics (Figs. 2b, S5, S6). Furthermore, the SLM generally succeeded in recapitulating the relationship between $\bar{f}$ and prevalence, a population genetic analogue of the abundance-prevalence relationship in ecology (Figs. 2b, S7, S8). While the SLM on its own does not predict Taylor's Law and any value of its exponent can be obtained by any family of distributions under mild conditions, the empirical evidence for the existence of Taylor's Law constrains the parameterization of the SLM and its subsequent interpretation [35, 65]. For the sites that the SLM is able to predict with a high degree of accuracy, the existence of Taylor's Law implies that β is constant across sites ($\sigma_{\tau_i} = \sigma_\tau$; Figs. 1e, S10, S11). Thus, the function for the expected allele frequency reduces to the proportionality $\langle f \rangle \propto K_i$.

Given the success of the SLM, it is worth investigating its degree of accuracy across observed estimates of prevalence. In general, the observed and predicted prevalence values tended to follow a one-to-one relationship across a wide range of observed prevalence values among the species I investigated (Fig. 3a). From this data, can glean a few insights into the accuracy of the gamma distribution predicted by the SLM. First, for most observed prevalence values, when the predicted prevalence differs from the observed value it is generally below the one-to-one line. This pattern does not deviate as the observed prevalence increases, suggesting that the SLM is able to capture the general features of f and its relation to prevalence for the range of observed prevalence values. When SLM is inaccurate it generally underpredicts the true prevalence of an allele. Alternatively stated, the SLM can predict more zero observations ($f = 0$) than there exists in the data. However, there is a large uptick in the predicted prevalence for alleles with high observed values of prevalence, where the predictions are effectively on the one-to-one line. This is the case, as there is a large drop in the error as observed prevalence increases (Fig. 3b), suggesting a negative relationship between the observed prevalence of a SLM and the relative error of the SLM gamma. I found that the correlation between these two quantities was universally negative among all species (Fig. 3c). By permuting the order of these quantities I obtained null distributions of correlation coefficients for all species, where all observed correlation coefficients clearly fall below the bounds of the 95% confidence intervals of the null.

The dependence of the predictive success of the SLM on the observed prevalence of an allele is a curious pattern. Ideally, one expects that the SLM is capable of explaining

the dynamics of alleles if they primarily exist on the genetic background of strains, as the SLM has been found to explain the temporal dynamics of strain frequencies within a single host [37]. The fact that the SLM, when it is inaccurate, tends to underpredict true prevalence allows one to rule out models that predict an excess proportion of values of $f = 0$ than expected from a given distribution (e.g., zero-inflated gamma), as they would further reduce the predicted prevalence. Such models are often representative of competitive exclusion [35]. In the context of this study, competitive exclusion corresponds to the hypothesis that a strain is not found in a given host because it is unable to compete, resulting in alternative stable states. This observation allows one to rule out competitive exclusion as potential major contributor towards the predictive error. While the SLM generally succeeded, its performance was poor among low prevalence alleles, suggesting that one cannot rule out the role of evolution. However, without the availability of additional models predicting distributions of f with zero free parameters it is difficult to parse alleles into "ecological" and "evolutionary" categories. Instead, it is likely that within a given allele with high error there are observed non-zero values of f that would otherwise be zero in the absence of evolution, increasing the observed prevalence beyond what the SLM can predict.

Confirming that the SLM succeeds due to strain structure

While the SLM succeeds at explaining patterns of diversity for an appreciable number of sites across microbial species, in this instance the ability to connect quantitative predictions to empirical distributions ultimately rests on the fact that the SLM predicts a distribution that is fully parameterized by the mean and variance of f (Eq. 3). Alternatively stated, I did not estimate the carrying capacity, I estimated the mean, the expectation of which is a function of the carrying capacity under the SLM, which further reduces to a proportionality when Taylor's law holds [35]. This reliance on statistical moments estimated from observational data suggests that models of alternative dynamics that also predict a gamma distribution are equivalent candidates in the absence of additional evidence. Starting with a Langevin equation, the frequency dynamics of a single allele are governed by forward and backward mutation (μ, υ), selection (s), and random genetic drift for a population of size N.

$$\frac{\partial f}{\partial t} = sf(1-f) + \mu(1-f) + \upsilon f + \sqrt{\frac{f(1-f)}{N}} \cdot \eta(t) \tag{6}$$

To remove the number of free parameters, one can examine Eq. 6 in the low frequency limit ($f \ll 1$) and obtain Langevin equations for evolution under positive ($s > 0$) and purifying selection ($s < 0$)

$$\frac{\partial f}{\partial t} = -sf + \mu + \sqrt{\frac{f}{N}} \cdot \eta(t) \tag{7a}$$

$$\frac{\partial f}{\partial t} = sf + \mu + \sqrt{\frac{f}{N}} \cdot \eta(t) \tag{7b}$$

A gamma-distributed stationary solution can be derived in the case of purifying selection ($s < 0$; [66, 67]), though the stationary solution of Eq. 7ab is trivial. For the $s > 0$ case, a gamma distribution of allele frequencies can be derived where the time-dependence is captured by the maximum frequency that an allele can reach ($f_{max} \equiv \frac{e^{st}-1}{2Ns}$) which, in principle, can be estimated from empirical data (S2 Text for full derivation). This dynamic form of mutation-selection balance is a gamma distribution that is solely parameterized by the maximum obtainable frequency of said allele (f_{max}) and the population scaled mutation rate ($2N\mu$) [68]. Together, these

selection regimes provide two forms of mutation-selection balance that predict the gamma distribution

$$f_{\text{Evo},s<0} \sim \text{Gamma}(2N\mu, 2N|s|) \tag{8a}$$
$$f_{\text{Evo},s>0} \sim \text{Gamma}(2N\mu, f_{max}^{-1}) \tag{8b}$$

Both of these distributions can be parameterized using the mean and variance, where $\langle f \rangle = \frac{\mu}{|s|}$ and $\beta = \frac{1}{2N\mu}$ for $s < 0$ and $\langle f \rangle = 2N\mu f_{max}$ and $\beta = (2N\mu)^2$ for $s > 0$. These distributions provide alternative explanations for the predictive capacity of the gamma distribution, and it is worth investigating their feasibility.

Evidence suggests that purifying selection in the human microbiome is widespread [3, 30], though it is unlikely responsible for the allele frequency fluctuations I was able to observe. After accounting for sequencing error using `MAPGD`, the mean of the lowest inferred allele frequency across all species was ~ 0.02. Give that microbial populations are typically large in size, it is unlikely that the same allele independently reached frequencies ≥ 0.02 in multiple hosts. Assuming that independence among sites holds, the frequencies of individual alleles should not exceed $\frac{1}{N|s|}$ [69], meaning that it is unlikely that a substantial fraction of the alleles with non-zero inferred frequencies were driven by purifying selection.

The dynamic mutation-selection balance parameterization of the gamma distribution also contains forward mutation, a useful feature given that mutation can contribute towards the total observed frequency of a beneficial allele that is increasing in frequency within a given host (Eq. 8ab). Such a feature is appealing, as, when the SLM fails, it does so by under-predicting observed prevalence. However, while the inclusion of forward mutation may increase the predicted value of prevalence, it is unlikely that positive selection substantially contributes towards the typical across-host single-site frequency spectra. Furthermore, the explicit time-dependence in the parameterization of f_{max} implies that a given allele is increasing in frequency in the $f \ll 1$ regime in all hosts where non-zero frequencies were observed. This dynamic is highly useful for describing the dynamics of an ensemble of populations where the initial state is known, but it is likely inapplicable to single-timepoint samples across unrelated human hosts. To the extent that this parameterization holds, it likely does so for alleles that are observed at low frequencies, restricting the model's applicability to low prevalence alleles. Because the SLM tends to fail for this prevalence regime, it is possible that a distribution derived from a model of evolutionary dynamics, gamma or otherwise, is necessary to predict allelic prevalence.

Beyond the consideration of evolutionary models, it is ultimately necessary to determine whether the success of the gamma is due to the existence of strain structure. This is a difficult task given that there is no guarantee that a given allele is present in multiple hosts because it is on the background of genetically identical strains. This limitation arises due to the difficulty inherent in determining whether a given host harbors a given strain from a single static metagenomic sample. Such difficulty persists in part due to technical and statistical limitations, but also due to the lack of practical strain definitions [70–73]. It is difficult to phase strains from short-read sequencing data where physical linkage between variants cannot be established, making it necessary to instead identify a single haplotype [3, 10]. So while it is currently possible to identify the prevalence of dominant lineages across hosts, there is no straightforward means to assign an observed allele to a given strain.

Given these methodological constraints, I instead identified the existence of strain-level structure for a given species within a given host using `StrainFinder`, a program that infers the number of strains using the shape of the site-frequency spectrum (Materials and Methods; [74]). I found that the fraction of hosts harboring

strain structure ranged from $0.089 - 0.64$ among the species I examined, with a median of ~ 0.27 (Fig. 4a). If alleles with higher prevalence were driven by the presence of strain structure, then one would expect a positive correlation between the accuracy of the predicted allele prevalence and the fraction of hosts containing strain structure across species, translating to a negative correlation between relative error and the fraction of hosts containing strain structure. This prediction held, as the correlation between these two variables was typically negative within a given range of prevalence values, but tended to be more negative when only alleles with high prevalence estimates were included (Fig. 4b). To resolve this relationship between the relative error and strain prevalence, I examined the relationship between the two across a wide range of observed prevalence thresholds (Fig. 4c). The correlation tended to increase with prevalence, though there was a substantial decrease once a prevalence threshold of ~ 0.1 was reached. The inferred estimates of strain prevalence can provide context. These estimates place empirical bounds on the range of allele prevalences that can be affected by strain structure, as the dynamics of alleles with prevalence values much smaller than the lowest observed strain prevalence are unlikely to be driven by strain structure. Given that the lowest observed strain prevalence value was close to this descent, it is likely that this value is where the ecological dynamics of strains began to dominate the trajectories of segregating alleles.

Discussion

Identifying and explaining universal patterns is a difficult task in biology. Though once they are identified, the extent that one can use mathematical models grounded in first principles to generate quantitative predictions reflects the state of one's understanding of the system. In this study, I identified patterns of genetic diversity in the human gut microbiome that were statistically invariant across evolutionarily distant species. Motivated by these results, I identified a prospective model of ecological dynamics (i.e., the Stochastic Logistic Model [35]) that could potentially explain said patterns through the sole use of measurable parameters. By generating quantitative predictions, I was able to predict the fraction of hosts that harbored a given allele using the SLM for a considerable fraction of sites across all species. I hypothesize that this success was driven by the existence of strain structure, which was tested by confirming a relationship between the accuracy of the SLM and the existence of strain structure.

The success of the SLM for common alleles suggests that true strain absences within a host are sufficiently rare such that they do not exert a detectable influence on the prevalence of an allele across hosts, implying that competitive exclusion is rare at the strain level [35]. This conclusion seemingly contrasts with the observation that strains exhibit varying frequencies across hosts [3], though fluctuations are not demonstrative of competitive exclusion. Fluctuation are expected under the SLM and some number of absences will inevitably arise due to the finite nature of sampling. However, it is possible that the carrying capacity of a given strain could vary from host-to-host, widening the distribution of frequency fluctuations. Such a detail can readily be incorporated into the SLM if one assumes that the carrying capacity in a given host is an independently drawn random variable from some unknown distribution, where the expected frequency of an allele on the background of a strain with constant carrying capacity $\langle f \rangle \propto K_i$ is replaced by the mean carrying capacity $\langle f \rangle \propto \frac{1}{M} \sum_{m=1}^{M} K_{i,m}$. Thus, given certain assumptions, the success of the SLM is reconcilable with the view that the carrying capacity of a strain is host-dependent.

While the SLM successfully predicted the prevalence of common alleles across hosts and has been shown to describe the temporal dynamics of strains within a host [37], all models have limitations, and it is useful to briefly discuss those that are applicable in

this case. A fundamental limitation of the SLM is that it is phenomenological in nature. It is unclear how microscopic details that are plainly relevant to the ecological dynamics of microbial communities, namely, consumer-resource dynamics [31, 75, 76], map onto the SLM. Models that incorporate resource consumption dynamics are capable of recapitulating the temporal dynamics of microbial communities at the species level [77], suggesting that models with such details may sufficiently describe the dynamics of strains. However, as a contrast, the phenomenological nature of the SLM could be viewed as an asset when one wants to capture and predict multiple empirical patterns using an analytic solution without the use of fitted parameters.

Finally, it is worth noting that the success of a prediction can be contingent on one's observations. Using a maximum likelihood approach to account for sequencing error (i.e., `MAPGD` [78]), the minimum inferred allele frequencies for sites that met the quality-control thresholds ranged from $0.006 - 0.06$ across species with an average of ~ 0.02. However, a straightforward calculation suggests that this range is higher than the true minimum frequency by several orders of magnitude. The mean relative abundances across hosts of species in this study had a range of $0.02 - 0.1$. Previously established order-of-magnitude estimates of the typical number of cells in the human gut range from $\sim 10^{13} - 10^{14}$, from which one can use the mean relative abundance distribution across species to calculate a first-pass range of empirical abundances of $\sim 10^{11} - 10^{13}$ [1, 79]. This range suggests that the minimum allele frequency (i.e., $1/N$) is at least eight orders of magnitude lower than the minimum allele frequency that I could infer after accounting for sequencing error. It is clear from this calculation that one is only able to examine a narrow range of the possible frequencies of segregating variants within the microbiome.

The implications of such a relatively narrow observational window are that the success of these predictions, and the feasibility of the single-locus models of evolution I examined, are likely contingent on present limitations on the depth and error rates of shotgun metagenomic sequencing. As advances in sequencing technology and statistical inference continue to permit lower observational thresholds and provide information about the physical linkage between variants [5, 80–83], one expects for an increasingly higher fraction of observable alleles to be subject to evolutionary dynamics rather than be the constituents of strains, affecting the aptness of the SLM as a general model. Succinctly stated, the existence of strain structure suggests that the dynamics of allele frequencies are likely dependent on the frequency scale over which observations can be made. By expanding the range of frequencies over which observations can be made, it may be feasible to identify a frequency threshold where the alleles one observes are primarily driven by evolutionary dynamics rather than the ecological dynamics of strains, providing the means to test quantitative predictions of recent developments in population genetic theory [68, 69]. Recognition of the possibility of such scale-dependence has the potential to shape future studies and rigorously assess the purported universality of empirical patterns of genetic diversity in the human gut microbiome.

Materials and Methods

Data acquisition and processing

To investigate patterns of genetic diversity within the human gut microbiome, I used shotgun metagenomic data from 468 healthy North American individuals sequenced by the Human Microbiome Project [16, 84]. I first processed the data using a previously developed analysis pipeline to identify the set of sites in core genes [3]. This pipeline uses a standard reference-based approach (`MIDAS` v1.2.2 [20]) to map reads from each

metagenomic sample to reference genes across a panel of prevalent species and filter reads based on quality scores and read mapping criteria. The relative abundance of each species in a given host was inferred using `merge_midas.py species`. Then, the command `merge_midas.py genes` was run with the following flags: `--sample_depth 10`, `--min_samples 1`, and `--max_species 150`. Using this processed gene output, the command `merge_midas.py snps` was run with the following flags: `--sample_depth 5`, `--site_depth 3`, `--min_samples 1`, `--max_species 150`, and `--site_prev 0.0`. I only processed samples from unrelated hosts and removed all temporal samples, leaving us with a set of samples that correspond to observations from unique hosts. I retained alleles in sites that were present in a core gene (i.e., a gene that was present in $\geq 90\%$ of hosts) with a minimum total within-host depth of coverage of 20 ($D \geq 20$) in at least 20 hosts.

After identifying the set of sites in core genes that passed the quality control thresholds, I obtained appropriate BAM files so that allele frequencies could ultimately be estimated. First, using the reference genomes used by `MIDAS` and the BAM file generated by `MIDAS` for all species in a single host, I split the BAM file containing all species into separate BAM files for each species using the command `samtools view` with default settings [85]. Each BAM file was sorted using the `samtools sort` command with default settings, from which a `.header` file was using `samtools view`.

Sequencing errors often obfuscate estimates of low frequency alleles ($f \ll 1$). To account for such errors, I elected to use the full maximum likelihood estimator `MAPGD` v0.4.40 [78]. This estimator was chosen due its comparative high sensitivity to low frequency alleles without the tendency to sacrifice its false discovery rate [86]. Using BAM and `.header` files processed from `MIDAS` output, `MAPGD` was run with a log-likelihood ratio polymorphism cutoff of 20 (`-a 20`). Samples with insufficient coverage for `MAPGD` to be run were removed from all downstream analyses. I then polarized alleles based on the major allele across hosts.

As an independent estimate of strain structure, I estimated strain frequencies for all species using all sites with ≥ 20 fold coverage within each host by applying `StrainFinder` v1.0 [74] to the frequency spectra obtained from the upstream pipeline [3]. The program `StrainFinder` was run on each sample for each species using 10 initial conditions using local convergence criteria with the following flags: `--dtol 1`, `--ntol 2`, `--max_time 20000`, `--converge`. The program was run for strain numbers ranging from one to four the estimates with the top five log-likelihoods were retained. I then selected strain frequencies with the lowest Bayesian Information Criterion for each species in each sample. Joint density plots (e.g., Fig. 2) were made using functions from `macroecotools` v0.4.0 [87].

Evaluating the success of prevalence predictions

Full derivations of the predicted prevalence of each model can be found in the Supporting Information. The predicted values of prevalence were compared to the following estimate of observed prevalence.

$$\hat{\varrho} = \frac{1}{M}\sum_{m=1}^{M}(1-\delta_{f_m,0}) = 1 - \frac{1}{M}\sum_{m=1}^{M}\delta_{f_m,0} \tag{9}$$

where f_m is the frequency of the alternative allele in the mth host and the Kronecker delta $\delta_{f_m,0}$ is equal to 1 if $f_m = 0$ and zero otherwise. To evaluate the success of the predictions I calculated the relative error of a given prediction

$$\varepsilon = \left|\frac{\text{Obs.} - \text{Pred.}}{\text{Obs.}}\right| \tag{10}$$

We performed permutation tests to determine whether the SLM had higher success among alleles with higher prevalence. By permuting all values of ε for a given species and calculating the Pearson correlation coefficient between it and the observed prevalence, I obtained a null distribution of correlation coefficients from which I calculated 95% intervals.

To determine whether there was a relationship between the error of the SLM and the fraction of hosts with strain structure across species, I implemented a permutation approach. First, for each species, I calculated the number of alleles in a given prevalence threshold (T total thresholds). I then permuted all values of ε and calculated the mean ε using the number of alleles that were found in each prevalence threshold, $(\bar{\varepsilon}_1, \cdots, \bar{\varepsilon}_T)$. The correlation coefficient was then calculated between $\bar{\varepsilon}_t$ and the fraction of hosts containing strains among all species for each prevalence threshold, allowing us to obtain a null distribution of correlation coefficients for all values of t. I only retained a prevalence threshold for a given species if there were at least 10 sites within the threshold.

Data and code availability

All raw public data used in this study was previously described [3]. All code is publicly available on GitHub: `https://github.com/garudlab/snv_prevalence`.

Acknowledgments

I thank S. Bald and R.W. Wolff for their assistance with `StrainFinder` and both N.R. Garud and R.W. Wolff for their insights at various stages of the project. I would also like to thank S. Bubnovich, D. Reyes-González and J. Grilli for their invaluable feedback on the manuscript. Thanks to B.H. Good for pivotal discussions and for making their lecture notes available to the public. Finally, thanks to M.S. Ackerman for their assistance with `MAPGD`. This work was supported by the NSF Postdoctoral Research Fellowships in Biology Program under Grant No. 2010885 (W.R.S.).

Author contributions

W.R.S. conceptualized the project, performed all derivations and analyses, and wrote the manuscript.

References

1. Zhao, S. *et al.* Adaptive Evolution within Gut Microbiomes of Healthy People. *Cell Host & Microbe* **25**, 656–667.e8 (2019).

2. Ghosh, O. M. & Good, B. H. Emergent evolutionary forces in spatial models of luminal growth in the human gut microbiota. Tech. Rep. (2021). URL `https://www.biorxiv.org/content/10.1101/2021.07.15.452569v1`. Company: Cold Spring Harbor Laboratory Distributor: Cold Spring Harbor Laboratory Label: Cold Spring Harbor Laboratory Section: New Results Type: article.

3. Garud, N. R., Good, B. H., Hallatschek, O. & Pollard, K. S. Evolutionary dynamics of bacteria in the gut microbiome within and across hosts. *PLOS Biology* **17**, e3000102 (2019). URL `https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.3000102`. Publisher: Public Library of Science.

4. Yaffe, E. & Relman, D. A. Tracking microbial evolution in the human gut using Hi-C reveals extensive horizontal gene transfer, persistence and adaptation. *Nature Microbiology* **5**, 343–353 (2020).

5. Roodgar, M. *et al.* Longitudinal linked-read sequencing reveals ecological and evolutionary responses of a human gut microbiome during antibiotic treatment. *Genome Research* **31**, 1433–1446 (2021). URL `https://genome.cshlp.org/content/31/8/1433`. Company: Cold Spring Harbor Laboratory Press Distributor: Cold Spring Harbor Laboratory Press Institution: Cold Spring Harbor Laboratory Press Label: Cold Spring Harbor Laboratory Press Publisher: Cold Spring Harbor Lab.

6. Ghalayini, M. *et al.* Evolution of a Dominant Natural Isolate of Escherichia coli in the Human Gut over the Course of a Year Suggests a Neutral Evolution with Reduced Effective Population Size. *Applied and Environmental Microbiology* **84**, e02377–17 (2018). URL `https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5835743/`.

7. Chen, D. W. & Garud, N. R. Rapid evolution and strain turnover in the infant gut microbiome. Tech. Rep. (2021). URL `https://www.biorxiv.org/content/10.1101/2021.09.26.461856v1`. Company: Cold Spring Harbor Laboratory Distributor: Cold Spring Harbor Laboratory Label: Cold Spring Harbor Laboratory Section: New Results Type: article.

8. Groussin, M. *et al.* Elevated rates of horizontal gene transfer in the industrialized human microbiome. *Cell* **184**, 2053–2067.e18 (2021). URL `https://www.sciencedirect.com/science/article/pii/S0092867421002415`.

9. Good, B. H. & Hallatschek, O. Effective models and the search for quantitative principles in microbial evolution. *Current Opinion in Microbiology* **45**, 203–212 (2018).

10. Truong, D. T., Tett, A., Pasolli, E., Huttenhower, C. & Segata, N. Microbial strain-level population structure and genetic diversity from metagenomes. *Genome Research* **27**, 626–638 (2017).

11. Schloissnig, S. *et al.* Genomic variation landscape of the human gut microbiome. *Nature* **493**, 45–50 (2013). URL https://www.nature.com/articles/nature11711. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 7430 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Genetic variation;Metagenomics;Microbiota Subject_term_id: genetic-variation;metagenomics;microbiota.

12. Faith, J. J. *et al.* The long-term stability of the human gut microbiota. *Science (New York, N.Y.)* **341**, 1237439 (2013).

13. Goyal, A., Bittleston, L. S., Leventhal, G. E., Lu, L. & Cordero, O. X. Interactions between strains govern the eco-evolutionary dynamics of microbial communities. *eLife* **11**, e74987 (2022). URL https://doi.org/10.7554/eLife.74987. Publisher: eLife Sciences Publications, Ltd.

14. Arumugam, M. *et al.* Enterotypes of the human gut microbiome. *Nature* **473**, 174–180 (2011). URL https://www.nature.com/articles/nature09944. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 7346 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Metagenomics;Microbiota Subject_term_id: metagenomics;microbiota.

15. Ley, R. E., Lozupone, C. A., Hamady, M., Knight, R. & Gordon, J. I. Worlds within worlds: evolution of the vertebrate gut microbiota. *Nature Reviews Microbiology* **6**, 776–788 (2008). URL https://www.nature.com/articles/nrmicro1978. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 10 Primary_atype: Research Publisher: Nature Publishing Group.

16. Lloyd-Price, J. *et al.* Strains, functions and dynamics in the expanded Human Microbiome Project. *Nature* **550**, 61–66 (2017).

17. Baumgartner, M., Bayer, F., Pfrunder-Cardozo, K. R., Buckling, A. & Hall, A. R. Resident microbial communities inhibit growth and antibiotic-resistance evolution of Escherichia coli in human gut microbiome samples. *PLOS Biology* **18**, e3000465 (2020). URL https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.3000465. Publisher: Public Library of Science.

18. Tett, A., Pasolli, E., Masetti, G., Ercolini, D. & Segata, N. Prevotella diversity, niches and interactions with the human host. *Nature Reviews Microbiology* **19**, 585–599 (2021). URL https://www.nature.com/articles/s41579-021-00559-y. Number: 9 Publisher: Nature Publishing Group.

19. Karcher, N. *et al.* Analysis of 1321 Eubacterium rectale genomes from metagenomes uncovers complex phylogeographic population structure and subspecies functional adaptations. *Genome Biology* **21**, 138 (2020). URL https://doi.org/10.1186/s13059-020-02042-y.

20. Nayfach, S., Rodriguez-Mueller, B., Garud, N. & Pollard, K. S. An integrated metagenomics pipeline for strain profiling reveals novel patterns of bacterial transmission and biogeography. *Genome Research* **26**, 1612–1625 (2016).

21. Niu, J. *et al.* Evolution of the Gut Microbiome in Early Childhood: A Cross-Sectional Study of Chinese Children. *Frontiers in Microbiology* **11** (2020). URL https://www.frontiersin.org/article/10.3389/fmicb.2020.00439.

22. Kundu, P., Blacher, E., Elinav, E. & Pettersson, S. Our Gut Microbiome: The Evolving Inner Self. *Cell* **171**, 1481–1493 (2017). URL https://www.sciencedirect.com/science/article/pii/S0092867417313715.

23. Yatsunenko, T. *et al.* Human gut microbiome viewed across age and geography. *Nature* **486**, 222–227 (2012). URL https://www.nature.com/articles/nature11053. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 7402 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Epidemiology;Genomics;Microbial ecology;Microbial genetics Subject_term_id: epidemiology;genomics;microbial-ecology;microbial-genetics.

24. Tierney, B. T. *et al.* The Landscape of Genetic Content in the Gut and Oral Human Microbiome. *Cell Host & Microbe* **26**, 283–295.e8 (2019). URL https://www.sciencedirect.com/science/article/pii/S193131281930352X.

25. Prosser, J. I. *et al.* The role of ecological theory in microbial ecology. *Nature Reviews Microbiology* **5**, 384–392 (2007). URL https://www.nature.com/articles/nrmicro1643. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 5 Primary_atype: Comments & Opinion Publisher: Nature Publishing Group.

26. Marquet, P. A. *et al.* On Theory in Ecology. *BioScience* **64**, 701–710 (2014). URL https://doi.org/10.1093/biosci/biu098.

27. Good, B. H. & Rosenfeld, L. B. Eco-evolutionary feedbacks in the human gut microbiome. Tech. Rep., bioRxiv (2022). URL https://www.biorxiv.org/content/10.1101/2022.01.26.477953v1. Section: New Results Type: article.

28. Garud, N. R. & Pollard, K. S. Population Genetics in the Human Microbiome. *Trends in Genetics* **36**, 53–67 (2020). URL https://www.sciencedirect.com/science/article/pii/S0168952519302215.

29. Poyet, M. *et al.* A library of human gut bacterial isolates paired with longitudinal multiomics data enables mechanistic microbiome research. *Nature Medicine* **25**, 1442–1452 (2019).

30. Shoemaker, W. R., Chen, D. & Garud, N. R. Comparative Population Genetics in the Human Gut Microbiome. *Genome Biology and Evolution* (2021). URL https://doi.org/10.1093/gbe/evab116.

31. Cui, W., Marsland, R. & Mehta, P. Diverse communities behave like typical random ecosystems. *Physical Review E* **104**, 034416 (2021). URL https://link.aps.org/doi/10.1103/PhysRevE.104.034416. Publisher: American Physical Society.

32. Advani, M., Bunin, G. & Mehta, P. Statistical physics of community ecology: a cavity solution to MacArthur's consumer resource model. *Journal of Statistical Mechanics (Online)* **2018**, 033406 (2018).

33. Descheemaeker, L. & de Buyl, S. Stochastic logistic models reproduce experimental time series of microbial communities. *eLife* **9**, e55650 (2020). URL https://doi.org/10.7554/eLife.55650. Publisher: eLife Sciences Publications, Ltd.

34. Shoemaker, W. R., Locey, K. J. & Lennon, J. T. A macroecological theory of microbial biodiversity. *Nature Ecology & Evolution* **1**, 1–6 (2017). URL https://www.nature.com/articles/s41559-017-0107. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 5 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Biodiversity;Community ecology;Macroecology;Microbial ecology Subject_term_id: biodiversity;community-ecology;macroecology;microbial-ecology.

35. Grilli, J. Macroecological laws describe variation and diversity in microbial communities. *Nature Communications* **11**, 4743 (2020). URL https://www.nature.com/articles/s41467-020-18529-y. Bandiera_abtest: a Cc_license_type: cc_by Cg_type: Nature Research Journals Number: 1 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Biodiversity;Community ecology;Ecological modelling;Macroecology;Microbial ecology Subject_term_id: biodiversity;community-ecology;ecological-modelling;macroecology;microbial-ecology.

36. Ji, B. W., Sheth, R. U., Dixit, P. D., Tchourine, K. & Vitkup, D. Macroecological dynamics of gut microbiota. *Nature Microbiology* **5**, 768–775 (2020). URL https://www.nature.com/articles/s41564-020-0685-1. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 5 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Microbial communities;Ecological modelling;Microbial ecology Subject_term_id: communities;ecological-modelling;microbial-ecology.

37. Wolff, R., Shoemaker, W. & Garud, N. Ecological Stability Emerges at the Level of Strains in the Human Gut Microbiome. Tech. Rep. (2021). URL https://www.biorxiv.org/content/10.1101/2021.09.30.462616v1. Company: Cold Spring Harbor Laboratory Distributor: Cold Spring Harbor Laboratory Label: Cold Spring Harbor Laboratory Section: New Results Type: article.

38. Phillips, R. Theory in Biology: Figure 1 or Figure 7? *Trends in Cell Biology* **25**, 723–729 (2015). URL https://www.cell.com/trends/cell-biology/abstract/S0962-8924(15)00194-4. Publisher: Elsevier.

39. Bhattacharjee, S. M. & Seno, F. A measure of data collapse for scaling. *Journal of Physics A: Mathematical and General* **34**, 6375–6380 (2001). URL https://doi.org/10.1088/0305-4470/34/33/302. Publisher: IOP Publishing.

40. Stanley, H. E. Scaling, universality, and renormalization: Three pillars of modern critical phenomena. *Reviews of Modern Physics* **71**, S358–S366 (1999). URL https://link.aps.org/doi/10.1103/RevModPhys.71.S358. Publisher: American Physical Society.

41. Theys, K. *et al.* Within-patient mutation frequencies reveal fitness costs of CpG dinucleotides and drastic amino acid changes in HIV. *PLOS Genetics* **14**, e1007420 (2018). URL https://journals.plos.org/plosgenetics/article?id=10.1371/journal.pgen.1007420. Publisher: Public Library of Science.

42. Vogl, C. & Bergman, J. Computation of the Likelihood of Joint Site Frequency Spectra Using Orthogonal Polynomials. *Computation* **4**, 6 (2016). URL https://www.mdpi.com/2079-3197/4/1/6. Number: 1 Publisher: Multidisciplinary Digital Publishing Institute.

43. Marquet, P. A. *et al.* Scaling and power-laws in ecological systems. *Journal of Experimental Biology* **208**, 1749–1769 (2005). URL https://journals.biologists.com/jeb/article/208/9/1749/9378/Scaling-and-power-laws-in-ecological-systems.

44. Ramsayer, J., Fellous, S., Cohen, J. E. & Hochberg, M. E. Taylor's Law holds in experimental bacterial populations but competition does not influence the slope. *Biology Letters* **8**, 316–319 (2012). URL https://royalsocietypublishing.org/doi/10.1098/rsbl.2011.0895. Publisher: Royal Society.

45. Kendal, W. S. & Jørgensen, B. Taylor's power law and fluctuation scaling explained by a central-limit-like convergence. *Physical Review E* **83**, 066115 (2011). URL https://link.aps.org/doi/10.1103/PhysRevE.83.066115.

46. Kendal, W. S. A scale invariant clustering of genes on human chromosome 7. *BMC Evolutionary Biology* 10 (2004).

47. Kendal, W. S. An Exponential Dispersion Model for the Distribution of Human Single Nucleotide Polymorphisms. *Molecular Biology and Evolution* **20**, 579–590 (2003). URL https://academic.oup.com/mbe/article-lookup/doi/10.1093/molbev/msg057.

48. Taylor, L. R. Aggregation, Variance and the Mean. *Nature* **189**, 732–735 (1961). URL https://www.nature.com/articles/189732a0. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 4766 Primary_atype: Research Publisher: Nature Publishing Group.

49. He, M. *et al.* Evolutionary dynamics of Clostridium difficile over short and long time scales. *Proceedings of the National Academy of Sciences of the United States of America* **107**, 7527–7532 (2010).

50. Bhatia, R. & Davis, C. A Better Bound on the Variance. *The American Mathematical Monthly* **107**, 353–357 (2000). URL https://doi.org/10.1080/00029890.2000.12005203. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/00029890.2000.12005203.

51. Cohen, J. E. & Xu, M. Random sampling of skewed distributions implies Taylor's power law of fluctuation scaling. *Proceedings of the National Academy of Sciences* **112**, 7749–7754 (2015). URL https://www.pnas.org/content/112/25/7749. Publisher: National Academy of Sciences Section: Biological Sciences.

52. Giometto, A., Formentin, M., Rinaldo, A., Cohen, J. E. & Maritan, A. Sample and population exponents of generalized Taylor's law. *Proceedings of the National Academy of Sciences* **112**, 7755–7760 (2015). URL https://www.pnas.org/content/112/25/7755. Publisher: National Academy of Sciences Section: Biological Sciences.

53. Xiao, X., Locey, K. J. & White, E. P. A Process-Independent Explanation for the General Form of Taylor's Law. *The American Naturalist* **186**, E51–E60 (2015). URL https://www.journals.uchicago.edu/doi/full/10.1086/682050. Publisher: The University of Chicago Press.

54. Transtrum, M. K. *et al.* Perspective: Sloppiness and emergent theories in physics, biology, and beyond. *The Journal of Chemical Physics* **143**, 010901 (2015). URL https://aip.scitation.org/doi/10.1063/1.4923066. Publisher: American Institute of Physics.

55. Zaoli, S. & Grilli, J. The stochastic logistic model with correlated carrying capacities reproduces beta-diversity metrics of microbial communities. Tech. Rep. (2021). URL https://www.biorxiv.org/content/10.1101/2021.11.16.468765v1. Company: Cold Spring Harbor Laboratory Distributor: Cold Spring Harbor Laboratory Label: Cold Spring Harbor Laboratory Section: New Results Type: article.

56. Zaoli, S. & Grilli, J. A macroecological description of alternative stable states reproduces intra- and inter-host variability of gut microbiome. *Science Advances* **7**, eabj2882 (2021). URL https://www.science.org/doi/10.1126/sciadv.abj2882. Publisher: American Association for the Advancement of Science.

57. Gardiner, C. W. *Stochastic methods: a handbook for the natural and social sciences*. No. 13 in Springer series in synergetics (Springer, Berlin Heidelberg, 2009), 4th ed edn.

58. Engen, S. & Lande, R. Population Dynamic Models Generating Species Abundance Distributions of the Gamma Type. *Journal of Theoretical Biology* **178**, 325–331 (1996). URL https://www.sciencedirect.com/science/article/pii/S0022519396900284.

59. Shoemaker, W. R. & Lennon, J. T. Predicting Parallelism and Quantifying Divergence in Microbial Evolution Experiments. *mSphere* (2022). URL https://journals.asm.org/doi/abs/10.1128/msphere.00672-21. Publisher: American Society for Microbiology 1752 N St., N.W., Washington, DC.

60. Madi, N., Chen, D., Wolff, R., Shapiro, B. J. & Garud, N. R. Community diversity is associated with intra-species genetic diversity and gene loss in the human gut microbiome. Tech. Rep., bioRxiv (2022). URL https://www.biorxiv.org/content/10.1101/2022.03.08.483496v1. Section: New Results Type: article.

61. Gaston, K. J. *et al.* Abundance–occupancy relationships. *Journal of Applied Ecology* **37**, 39–59 (2000). URL https://onlinelibrary.wiley.com/doi/abs/10.1046/j.1365-2664.2000.00485.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1046/j.1365-2664.2000.00485.x.

62. Shade, A. *et al.* Macroecology to Unite All Life, Large and Small. *Trends in Ecology & Evolution* **33**, 731–744 (2018).

63. Sloan, W. T., Woodcock, S., Lunn, M., Head, I. M. & Curtis, T. P. Modeling Taxa-Abundance Distributions in Microbial Communities using Environmental Sequence Data. *Microbial Ecology* **53**, 443–455 (2007). URL https://doi.org/10.1007/s00248-006-9141-x.

64. Burns, A. R. *et al.* Contribution of neutral processes to the assembly of gut microbial communities in the zebrafish over host development. *The ISME Journal* **10**, 655–664 (2016). URL https://www.nature.com/articles/ismej2015142. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 3 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Microbial ecology Subject_term_id: microbial-ecology.

65. Cohen, J. E. Every variance function, including Taylor's power law of fluctuation scaling, can be produced by any location-scale family of distributions with positive mean and variance. *Theoretical Ecology* **13**, 1–5 (2020). URL https://doi.org/10.1007/s12080-019-00445-7.

66. Nei, M. The frequency distribution of lethal chromosomes in finite populations. *Proceedings of the National Academy of Sciences of the United States of America* **60**, 517–524 (1968). URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC225078/.

67. Ewens, W. J. *Mathematical population genetics* (Springer, 2010). OCLC: 778838329.

68. Good, B. H. Linkage disequilibrium between rare mutations. *bioRxiv* 2020.12.10.420042 (2020). URL https://www.biorxiv.org/content/10.1101/2020.12.10.420042v1. Publisher: Cold Spring Harbor Laboratory Section: New Results.

69. Cvijović, I., Good, B. H. & Desai, M. M. The Effect of Strong Purifying Selection on Genetic Diversity. *Genetics* **209**, 1235–1278 (2018). URL https://doi.org/10.1534/genetics.118.301058.

70. Yan, Y., Nguyen, L. H., Franzosa, E. A. & Huttenhower, C. Strain-level epidemiology of microbial communities and the human microbiome. *Genome Medicine* **12**, 71 (2020). URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7427293/.

71. Dijkshoorn, L., Ursing, B. M. & Ursing, J. B. Strain, clone and species: comments on three basic concepts of bacteriology. *Journal of Medical Microbiology* **49**, 397–401 (2000).

72. Zhu, A., Sunagawa, S., Mende, D. R. & Bork, P. Inter-individual differences in the gene content of human gut bacterial species. *Genome Biology* **16**, 82 (2015). URL https://doi.org/10.1186/s13059-015-0646-9.

73. Shapiro, B. J. & Polz, M. F. Ordering microbial diversity into ecologically and genetically cohesive units. *Trends in Microbiology* **22**, 235–247 (2014).

74. Smillie, C. S. *et al.* Strain Tracking Reveals the Determinants of Bacterial Engraftment in the Human Gut Following Fecal Microbiota Transplantation. *Cell Host & Microbe* **23**, 229–240.e5 (2018). URL https://www.cell.com/cell-host-microbe/abstract/S1931-3128(18)30038-6. Publisher: Elsevier.

75. Chesson, P. MacArthur's consumer-resource model. *Theoretical Population Biology* **37**, 26–38 (1990). URL https://www.sciencedirect.com/science/article/pii/004058099090025Q.

76. Good, B. H., Martis, S. & Hallatschek, O. Adaptation limits ecological diversification and promotes ecological tinkering during the competition for substitutable resources. *Proceedings of the National Academy of Sciences* **115**, E10407–E10416 (2018). URL https://www.pnas.org/content/115/44/E10407. Publisher: National Academy of Sciences Section: PNAS Plus.

77. Ho, P.-Y., Good, B. & Huang, K. C. Competition for fluctuating resources reproduces statistics of species abundance over time across wide-ranging microbiotas. Tech. Rep., bioRxiv (2021). URL

https://www.biorxiv.org/content/10.1101/2021.05.13.444061v1. Section: New Results Type: article.

78. Lynch, M., Bost, D., Wilson, S., Maruki, T. & Harrison, S. Population-Genetic Inference from Pooled-Sequencing Data. *Genome Biology and Evolution* **6**, 1210–1218 (2014). URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4040993/.

79. Sender, R., Fuchs, S. & Milo, R. Revised Estimates for the Number of Human and Bacteria Cells in the Body. *PLOS Biology* **14**, e1002533 (2016). URL https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.1002533. Publisher: Public Library of Science.

80. Kuleshov, V. *et al.* Synthetic long-read sequencing reveals intraspecies diversity in the human microbiome. *Nature Biotechnology* **34**, 64–69 (2016). URL https://www.nature.com/articles/nbt.3416. Bandiera_abtest: a Cg_type: Nature Research Journals Number: 1 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Metagenomics Subject_term_id: metagenomics.

81. Zlitni, S. *et al.* Strain-resolved microbiome sequencing reveals mobile elements that drive bacterial competition on a clinical timescale. *Genome Medicine* **12**, 50 (2020). URL https://doi.org/10.1186/s13073-020-00747-0.

82. DeMaere, M. Z. & Darling, A. E. bin3C: exploiting Hi-C sequencing data to accurately resolve metagenome-assembled genomes. *Genome Biology* **20**, 46 (2019). URL https://doi.org/10.1186/s13059-019-1643-1.

83. Press, M. O. *et al.* Hi-C deconvolution of a human gut microbiome yields high-quality draft genomes and reveals plasmid-genome interactions. Tech. Rep. (2017). URL https://www.biorxiv.org/content/10.1101/198713v1. Company: Cold Spring Harbor Laboratory Distributor: Cold Spring Harbor Laboratory Label: Cold Spring Harbor Laboratory Section: New Results Type: article.

84. Methé, B. A. *et al.* A framework for human microbiome research. *Nature* **486**, 215–221 (2012). URL https://www.nature.com/articles/nature11209. Bandiera_abtest: a Cc_license_type: cc_y Cg_type: Nature Research Journals Number: 7402 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Medical research;Metagenomics;Microbiome Subject_term_id: medical-research;metagenomics;microbiome.

85. Danecek, P. *et al.* Twelve years of SAMtools and BCFtools. *GigaScience* **10** (2021). URL https://doi.org/10.1093/gigascience/giab008. Giab008, https://academic.oup.com/gigascience/article-pdf/10/2/giab008/36332246/giab008.pdf.

86. Guirao-Rico, S. & González, J. Benchmarking the performance of Pool-seq SNP callers using simulated and real sequencing data. *Molecular Ecology Resources* **21**, 1216–1229 (2021). URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1755-0998.13343. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1755-0998.13343.

87. Xiao, X., Thibault, K., Harris, D. J., Baldridge, E. & White, E. weecology/macroecotools: v0.4.0 (2016). URL https://zenodo.org/record/166721.

Figures

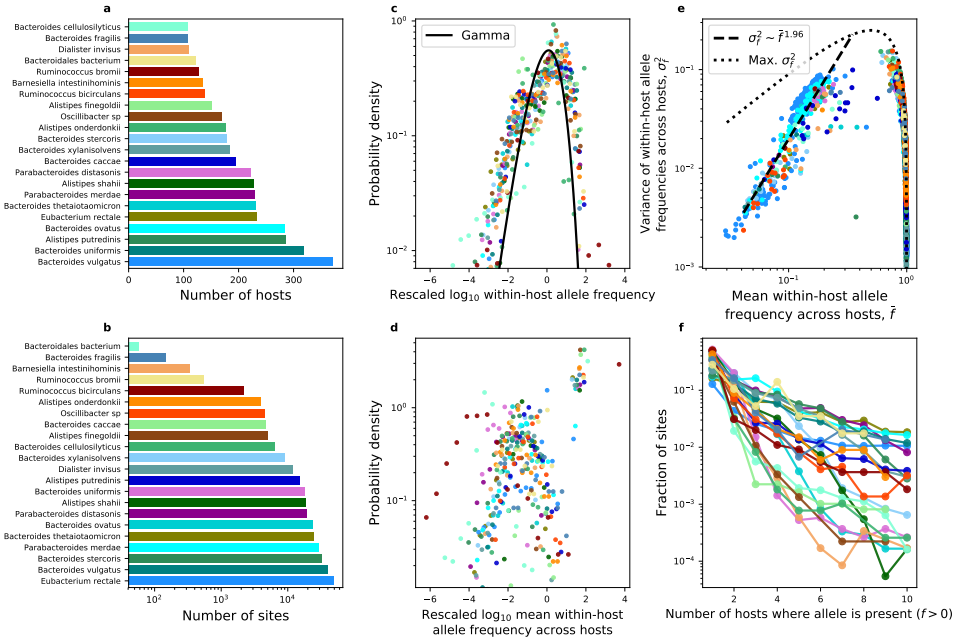

Figure 1. Distributions of genetic diversity exhibit similar statistical forms across phylogenetically distant species in the human gut. a,b) Using the set of hosts and alleles that passed the quality control criteria, I was able to evaluate the degree of similarity in patterns of genetic diversity for the 22 most prevalent bacterial species. **c)** The distribution of within-host frequencies for all alleles for all hosts as well as **(d)** the distribution of mean allele frequencies are qualitatively similar across species, where the logarithm of all values for each species has been rescaled using the standard score to facilitate visual comparison. The similarity in the form of each distribution across species suggests that it is reasonable to identify a single model of allelic dynamics. **e)** The relationship between statistical moments of within-host allele frequencies is consistent across species, as there is a strong linear relationship between the mean frequency of an allele and its the variance on a log-log scale (i.e., Taylor's Law). To reduce the contribution of an excess number of zeros towards estimates of $\bar{f}$ and σ_f^2, I only examined alleles with non-zero values of f in $\geq 35\%$ of hosts. **f)** Finally, the number of hosts where an allele is present can contain information about its ecological or evolutionary dynamics. I found that the typical allele is present in a low number of hosts and that the fraction of alleles in a given number of hosts tends to decay as the number of hosts increases across species. All sites in this analysis are synonymous, identical analyses were performed on alleles at nonsynonymous sites (Fig. S2). Species within the same genus were assigned the same primary/secondary color with varying degrees of saturation.

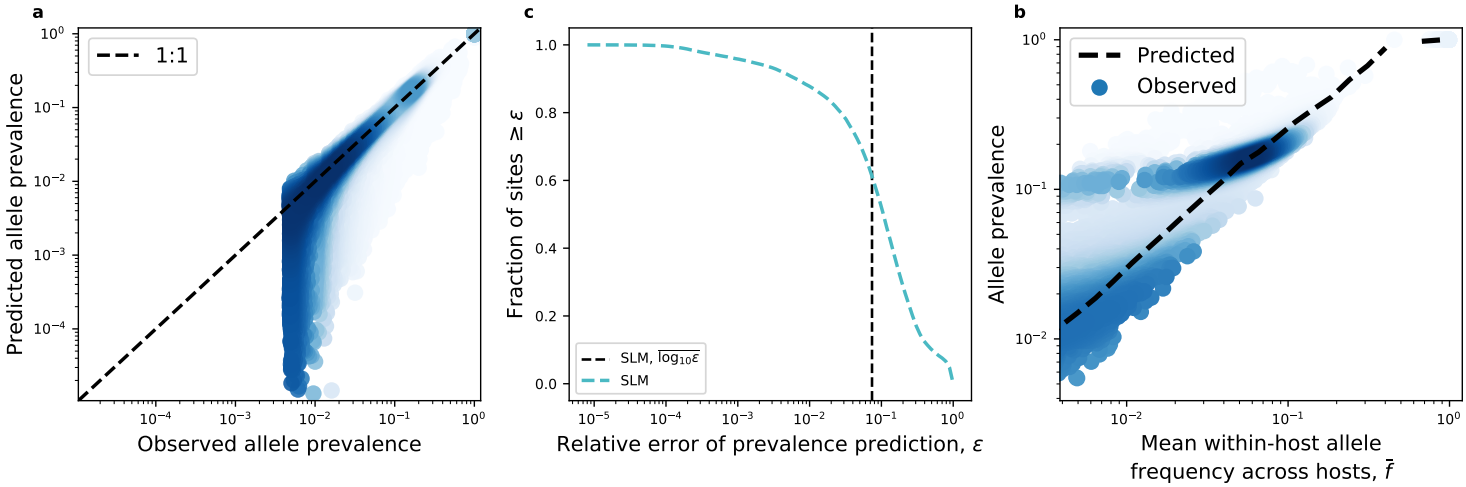

Figure 2. The SLM generally succeeds in predicting patterns of genetic diversity. To demonstrate a particular instance of this pattern, I examined the degree of success for each model on a common commensal gut species, *B. vulgatus*. **a)** The SLM generally succeeds in predicting the prevalence of alleles. This assessment is particularly true for alleles with high observed estimates of prevalence, all of which fall on a one-to-one line indicating perfect accuracy. **b)** The visual evidence of the predictive success of the SLM is confirmed by quantifying the relative error of the prevalence predictions. The mean of the logarithm of the relative error of the SLM ($\overline{\log_{10}\varepsilon}$, dashed black line) is ~ 0.1. **c)** This success is also present in the relationship between patterns of within-host vs. among-host allelic diversity, as the SLM generally succeeded in predicting the relationship between the mean frequency of a allele ($\bar{f}$) and its prevalence. All analyses here were performed on alleles at synonymous sites. Visualizations of the predictions in this plot for all species for nonsynonymous and synonymous sites can be found in the Supplemental Text.

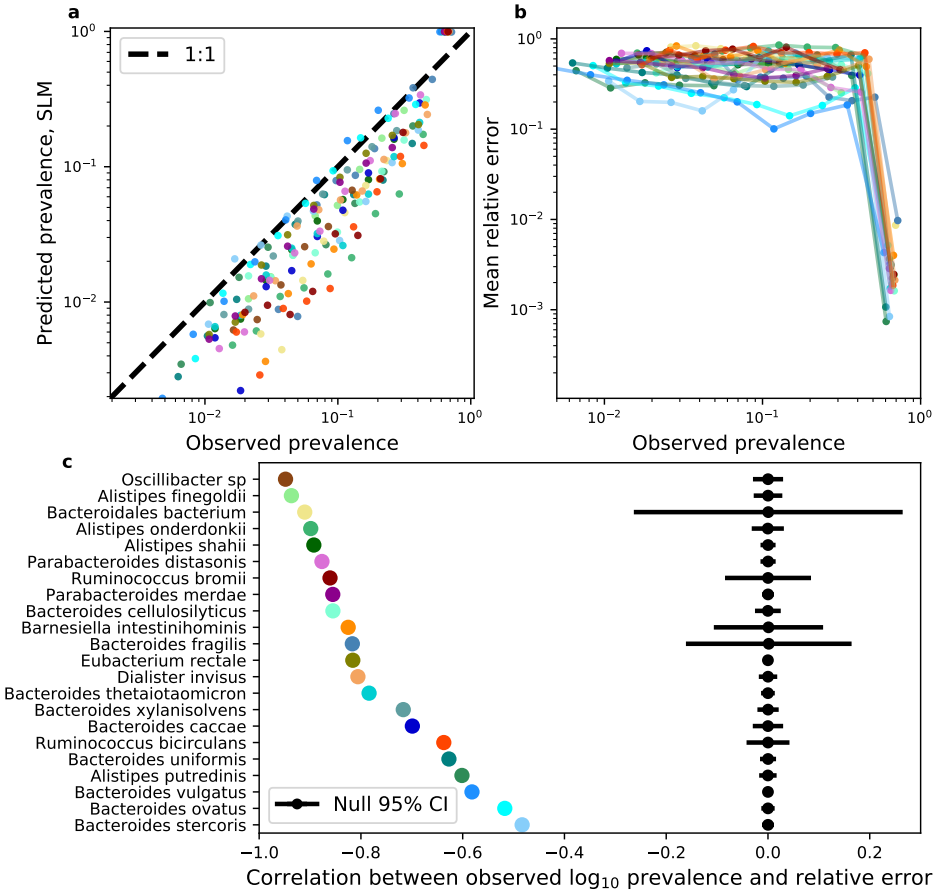

Figure 3. The SLM succeeds and fails in a consistent manner across phylogenetically distant species. a) By binning alleles by their observed prevalence, it is clear that the observed vs. predicted allele prevalence values tend to follow a one-to-one relationship across all species. **b)** This assessment holds with the relative error is examined, as it tends to remain flat across a wide range of prevalence values until the observed prevalence approaches high values. **c)** A permutation test confirms that there is a negative relationship between the accuracy the SLM and the observed prevalence of an allele, as the correlation between the two quantities is far more negative than expected by chance. All analyses in this plot were performed using synonymous sites. Identical analyses and equivalent results for nonsynonymous sites can be found in the Supplemental Text (Fig. S9).

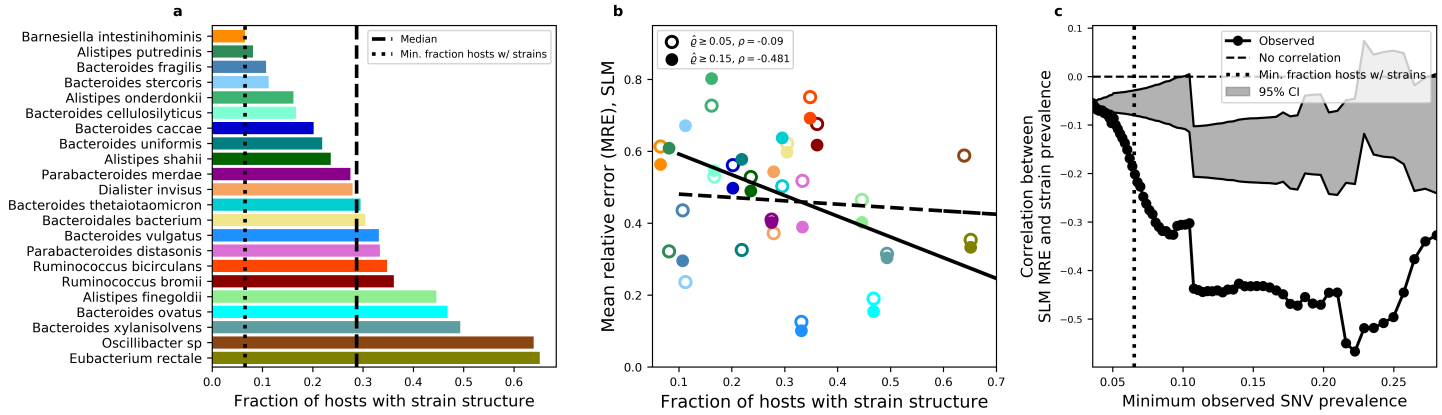

Figure 4. The presence of strain structure is correlated with the accuracy of the SLM. a) The presence of absence of strain structure can be inferred for a given species within a given host, providing the means to determine whether the degree of deviation from perfect accuracy (i.e., relative error) of the SLM is correlated with the fraction of hosts that harbor strain structure. **b)** By examining the relationship between error and strain structure across species within a low ($\hat{\varrho} \geq 0.05$: dashed line) and high ($\hat{\varrho} \geq 0.15$: solid line) threshold value of observed allele prevalence, I found that the correlation is higher for alleles with a high prevalence threshold. **c)** One can resolve this dependence using a range of observed prevalence bins and an appropriate permutation test, where the correlation between the error and the fraction of hosts with strain structure tends to become increasingly negative (Materials and Methods). Notably, the correlation drops off right after the minimum fraction of hosts where strain structure was inferred (> 1 strain) and continues as allele prevalence increases. In order to maintain a high number of species across values where the existence of strain structure could feasibly contribute, I restricted my analysis to prevalence values less than the median strain prevalence across species.