Longitudinal linked read sequencing reveals ecological and evolutionary responses of a human gut microbiome during antibiotic treatment

Morteza Roodgar^{1,2*}, Benjamin H. Good^{3*†}, Nandita R. Garud⁴, Stephen Martis⁵, Mohan Avula¹, Wenyu Zhou¹, Samuel Lancaster¹, Hayan Lee¹, Afshin Babveyh¹, Sophia Nesamoney¹, Katherine S. Pollard^{6,7,8†}, Michael P. Snyder^{1†}

¹Department of Genetics, Stanford University, Stanford, California, 94305, ²Institute for Stem Cell Biology and Regenerative Medicine, Stanford University School of Medicine, Stanford, California, 94305, ³Department of Applied Physics, Stanford University, Stanford, California 94305, ⁴Department of Ecology and Evolutionary Biology, University of California Los Angeles, ⁵Department of Physics, University of California, Berkeley, CA 94720 ⁶Gladstone Institutes, San Francisco, CA 94158, ⁷Department of Epidemiology & Biostatistics, University of California, San Francisco, Ca 94158, ⁸Chan Zuckerberg Biohub, San Francisco, CA 94158.

**These authors contributed equally to this work.*

†Correspondence should be addressed to: B.H.G. (bhgood@stanford.edu), K.S.P. (katherine.pollard@gladstone.ucsf.edu), or M.P.S. (mpsnyder@stanford.edu).

Gut microbial communities can respond to antibiotic perturbations by rapidly altering their taxonomic and functional composition. However, little is known about the strain-level processes that drive this collective response. Here we characterize the gut microbiome of a single individual at high temporal and genetic resolution through a period of health, disease, antibiotic treatment, and recovery. We used deep, linked-read metagenomic sequencing to track the longitudinal dynamics of thousands of single nucleotide variants within 36 species, which allowed us to contrast these genetic dynamics with the ecological fluctuations at the species level. We find that antibiotics can drive rapid shifts in the genetic composition of individual species, often involving incomplete genome-wide sweeps of pre-existing variants. Interestingly, genetic changes frequently occur in species without obvious changes in relative species abundance, emphasizing the importance of monitoring diversity below the species level. Our results provide new insights into the population genetic forces that shape individual microbiomes on therapeutically relevant timescales, with potential implications for personalized health and disease.

Introduction

The composition of the gut microbiome varies among human populations and individuals, and it is thought to play a key role in maintaining health and reducing susceptibility to different diseases (1-4). Understanding how this microbial ecosystem changes from week to week—through periods of health, disease and treatment—is important for personalized health management and design of microbiome-aware therapies (5).

Many studies have investigated intra-host dynamics at the species or pathway level (6-16). Among other findings, these studies have shown that oral antibiotics can dramatically influence the composition of the gut microbiome over a period of days, while the community often regains much of its initial composition in the weeks or months after

antibiotics are removed (7-9). This suggests an intriguing hypothesis, in which the long-term composition of a healthy gut community is buffered against brief environmental perturbations.

However, the mechanisms that enable this ecological robustness remain poorly understood. Does species composition recover because external strains are able to recolonize the host? Or do resident strains persist in refugia and expand again once antibiotics are removed? In the latter case, do the resident populations also acquire genetic differences during this time, either due to population bottlenecks or to new selection pressures that are revealed during treatment? To address these questions, it is necessary to map the fine scale genetic diversity below the species or pathway level, and follow how it changes during periods of health, disease, and treatment.

Advances in strain-resolved metagenomics and isolate sequencing (17-19) have made it possible to detect DNA sequence variants within species, and to track how they change within and between hosts. These studies have shown that gut bacteria can acquire genetic differences over time even in healthy human hosts, and that these differences arise from a mixture of external replacement events (18, 20, 21) and the evolution of resident strains (21-23). However, because these studies are based on relatively few timepoints per host, or shallow sampling of their microbiota, the population genetic processes that drive these strain-level dynamics remain poorly characterized. Understanding how the forces of mutation, recombination, selection, and genetic drift operate within hosts is critical for efforts to forecast personalized responses to drugs or other therapies.

To bridge this gap, we used deep metagenomic sequencing to follow the genetic diversity within a single host microbiome at approximately weekly intervals over a period of five months. This longitudinal study included periods of infectious disease and the oral administration of broad-spectrum antibiotics. In contrast to conventional metagenomic studies, we used a linked-read sequencing technique to generate and analyze each of our metagenomic samples. Large molecules of bacterial DNA were isolated in millions of emulsified droplets, digested into shorter fragments, and labelled with a corresponding DNA barcode to follow linked reads from the droplet. Previous work has shown that the linkage information encoded in these barcoded “read clouds” can improve genome assembly (24) and taxonomic assignment (25) from human gut metagenomes. Here, we hypothesized that longitudinal applications of linked read sequencing could also aid the detection and interpretation of genetic changes that occur within individual species.

We developed new statistical methods that leverage these data to simultaneously measure the ecological and evolutionary dynamics across multiple species during the course of antibiotic treatment. We find that natural selection can drive rapid shifts in the genetic composition of individual species, often via incomplete genome-wide sweeps of linked sequence variants. Interestingly, these within-species dynamics can occur even without large changes in the relative abundance of the species, emphasizing the importance of monitoring diversity below the species level. Moreover, we find that many sweeping variants were already segregating in their respective populations before exposure to antibiotics, and quickly revert to their original state once antibiotics are removed, echoing previous observations of robustness at the species level. Together, these results provide new insights into the population genetic forces that shape the gut microbiota of individual hosts, which has important implications for personalized health and disease.

Results

Longitudinal linked read sequencing of a human gut microbiome during disease and treatment

Generation of linked reads requires the preparation of long DNA fragments. We therefore developed an optimized protocol for extracting high-molecular weight DNA from human stool samples (Methods). We used this approach to perform linked read sequencing (10X Genomics) on 19 stool samples collected from a single individual over a period of 5 months (Fig. 1A, B, Table S1). During this time, the individual was diagnosed with Lyme disease and received a two-week course of broad-spectrum oral antibiotics (doxycycline). We generated deep sequencing data for each sample (ranging from ~8-160 Gbp), so that a typical read was present in a “read cloud” (Fig. 1C) containing ~4-30 other read pairs (Fig. 1D, Fig. S1). Consistent with previous studies (25), we observed high rates of read cloud “impurity”, with each read cloud containing fragments from ~5-10 different species (Fig. 1D). To overcome this issue, we used a two-stage approach, which leverages the hybrid nature of the linked read protocol. We first ignored barcodes and used short-read, reference-based methods to track species and sub-species diversity over time (21, 22, 26). We then developed a statistical model for linking genomic regions based on higher than expected rates of barcode sharing given the level of barcode impurity in our data (Methods). Using this hybrid approach, we documented the ecological and evolutionary responses of the gut microbial community before, during, and after antibiotic treatment.

Consistent with previous studies (20, 26-30), we observed a substantial perturbation in species-level composition during and immediately after antibiotic treatment, followed by a return to near baseline values by the end of the sampling interval (Fig. 1B). However, only a few species dramatically declined in relative abundance during this period: of the 48 species that started with a baseline relative abundance greater than 0.1%, only 9 experienced more than a 10-fold reduction in relative abundance by the end of the treatment window. Notable examples include *Alistipes finegoldii* or *Butyrivibrio crossotus* (Fig. 2A). The small number of such examples suggests that a large fraction of the community may have been able to maintain high absolute abundance during treatment, e.g. due to reduced antibiotic sensitivity. Consistent with this hypothesis, we observed a high baseline proportion of doxycycline-related resistance genes among our metagenomic reads (~200 per million mapped), which increased ~2-fold during treatment (Fig. S2). In addition, we found that the *Bacteroides vulgatus* population maintained a high replication origin peak-to-trough ratio (PTR), a proxy for bacterial growth rate, during the antibiotic treatment period (Fig. S3). Since doxycycline is a translation inhibitor, the high PTR values suggest that the *B. vulgatus* population, and by extension, the other species that maintained similar relative abundances during treatment, may have reduced sensitivity to the effects of doxycycline. This is consistent with previous observations of tetracycline resistance in isolates of several *Bacteroides* species (31).

Deep longitudinal sequencing reveals shifts in the genetic composition of 36 species in the same host.

The general pattern of persistence and recovery at the species level is shared by many other classes of antibiotics (28). Yet, the strain-level dynamics that give rise to this long-term stability remain poorly understood. Do the species that persist through disease and treatment remain stable genetically? Or does this general pattern of robustness mask a

larger flux of genetic changes occurring within individual species? Our approach allows us to address these questions by tracking genetic variation within species over time.

We first tracked the genetic composition of each species throughout the time course by aligning our short sequencing reads to a panel of reference genomes and estimating the population frequency of single nucleotide variants (SNVs) at each timepoint (Methods). The high sequencing coverage enabled at least ~10-fold coverage per timepoint for species with abundance >0.3%, and coverages as high as ~500x in some of the most abundant species (Fig. S4). This allowed us to simultaneously monitor SNV dynamics within 36 species that passed our coverage thresholds (Methods), and to contrast these “evolutionary” responses with the “ecological” dynamics observed at the species level (Figs. 2, 3, S5, S6).

This strain-level analysis revealed striking differences in the genetic composition of different species. Consistent with previous work (18, 20, 21, 27), the initial levels of genetic diversity vary widely between species. Some common species, such as *Bacteroides vulgatus* and *Bacteroides uniformis*, have more than ~10,000 SNVs at intermediate frequencies whereas other species, e.g. *Bacteroides coprocola* or *Alistipes* sp, have fewer than ~100 detectable SNVs (Fig. 2B). Of particular interest are those SNVs that undergo large changes in frequency between the initial and later timepoints (e.g. from <20% to >80%, with FDR<0.1, see Methods); these indicate a nearly complete “sweep” within the species of interest. We observe a similarly wide range in the number of SNV differences during and immediately after antibiotic treatment, from more than ~10,000 in some species (e.g. *Eubacterium eligens*) to ~10 (or even 0) in others (Fig. 2C). Of the 36 populations in Fig. 2C, more than half accumulated at least one SNV difference during this period, and more than 80% accumulated SNV differences in at least one portion of the study.

Similar within-host changes have recently been observed in metagenomic analyses from healthy hosts (21-23), though at a significantly lower rate (Fisher’s exact test, $P<0.001$). A major challenge in these earlier studies has been to demonstrate that the temporally variable SNVs are truly linked to their inferred genomic background, and are not simply read mapping artifacts (e.g., from another temporally fluctuating species that shares some parts of the genome). Read cloud sequencing provides an opportunity to address this question. For each SNV difference reported in Fig. 2C, we examined the patterns of barcode sharing with genes in the “core” genomes of our reference genome panel, which provides a proxy for the true genomic background (Methods). This analysis yields positive confirmation for ~80% of the SNV differences in Fig. 2C (where both alleles share read clouds with a core gene in the target species), and negative confirmation for <1% (Fig. S7). We conclude that the majority of these SNVs represent true genetic changes within their respective populations.

The variable genetic responses in different species are not easily explained by their phylogenetic relatedness or their relative abundance trajectories. As an example, Fig. 3 shows the full species abundance and SNV frequency trajectories for six example species, which are chosen to illustrate the range of observed behaviors. This set includes three different members of the *Alistipes* genus that coexist within this particular host. The first two species, *A. sp* and *A. finegoldii*, experience dramatic reductions in relative abundance during treatment, but we observe genetic differences in only one of the populations (*A. finegoldii*) when they recover to their initial levels. In *A. onderdonkii*, by contrast, the relative abundance remains high at the end of the treatment phase, but we observe rapid changes in the frequencies of several SNVs within this species during the same time

period ($P < 0.001$, Methods). These examples show that species abundances alone are not sufficient to predict genetic response within species: relatively constant species abundance trajectories can mask interesting genetic shifts within a species, and vice versa.

Quantifying genetic linkage between SNVs using barcoded read clouds

We next sought to quantify the population genetic processes that could give rise to the SNV changes in Figs. 2 and 3. A key question is the extent of genetic linkage within a species: is recombination sufficiently frequent that genetic drift and natural selection act independently on different SNVs? Or are SNV trajectories tightly correlated because they are linked together on a small number of clonal backgrounds? This question is particularly relevant for species with high levels of SNV diversity like *B. vulgatus* (Fig. S8), where it can mean the difference between $\sim 10^4$ evolutionary trajectories (if SNVs are independent) or possibly only one (if SNVs derive from a mixture of two clonal strains).

Previous analyses of gut bacteria suggest that recombination can efficiently decouple SNVs over long timescales (i.e., millions of bacterial generations) (21, 32), but the extent of genetic linkage within hosts remains unclear. The additional information provided by linked read sequencing now allows us to investigate this question. We developed a statistical approach for detecting linkage between pairs of SNVs (Fig. 4A), which accounts for the substantial coverage variation across different species and read clouds (Methods). Fig. 4B shows how the overall levels of read cloud sharing depend on the coordinate distance between the two SNVs on the reference genome. Consistent with the fragment length estimates from our HMW DNA extraction protocol (Fig. S9), we observed an enrichment of shared read clouds barcodes for SNVs within ~ 10 kb of each other, though the overall fraction of long-range read clouds remains modest ($\sim 10\%$). For the subset of SNVs pairs with significant read cloud sharing, we further quantified genetic linkage by examining the combinations of major and minor alleles that are observed in the same read clouds. In particular, we estimated the number of allelic combinations (or haplotypes) that are observed for each pair of SNVs as a function of their coordinate distance on the reference genome (Fig. 4A). According to the four-gamete test (33), three or fewer haplotypes are consistent with clonal evolution, but the presence of all four haplotypes indicates a possible recombination event between the two SNVs (Methods). Fig. 4D shows that the vast majority of the SNV pairs across species are consistent with clonal evolution: of the ~ 4 million SNV pairs we examined that were separated by more than 2kb (Fig. 4C), only ~ 600 showed significant evidence for all four haplotypes ($q < 0.05$, Methods). Most of these four-haplotype pairs are concentrated in just a few species, with high values of linkage disequilibrium between the two SNVs (Fig. S10). This suggests that to a first approximation, the SNV dynamics within species in this time course reflect a competition between a few clonal haplotypes, rather than independent alleles. This is consistent with previous indirect evidence from the clustering of allele frequencies within hosts (21, 29, 34).

For populations with sufficiently high SNV density (> 1 per kb), the patterns of read cloud sharing can inform efforts to cluster SNVs into smaller numbers of competing haplotypes (Fig. S8). However, many interesting temporal changes occurred in populations with much lower SNV density (< 1 per 10kb, Fig 2). Figure 4 suggests that SNVs will not typically share read clouds in these populations, except in rare cases where they are located in nearby regions of the genome. We therefore used a heuristic approach to infer clusters of perfectly linked SNVs clusters (or multi-SNV haplotypes) based on similarities in their

allele frequency trajectories (Methods). The inferred haplotypes are indicated by the coloring scheme in Fig. 3.

Temporal dynamics of haplotypes reveal cryptic phenotypic differences within species

We next investigated the role of natural selection in driving the genetic changes we observed within species. Many studies assume that intra-host dynamics are dominated by selection pressures that act at the level of species or functional guilds, while genetic variants within a species are mostly interchangeable. In this “neutral” scenario, any shifts in the genetic composition of individual species must be driven by stochastic demographic processes (e.g. *genetic drift*) (35, 36). Other studies have argued that environmental shifts like antibiotics could reveal previously hidden within-species differences, leading to rapid shifts in the frequencies of different genetic variants (37).

The high levels of genetic linkage make it difficult to distinguish between these scenarios using traditional approaches, since selection acts on extended haplotypes rather than individual alleles. For example, the SNV clusters in Fig. 3 contain many synonymous variants (Fig. 5), which are presumably hitchhiking alongside the true causative mutations. These driver mutations may even be missing if they arise from structural variants, mobile elements, or other accessory genes that are not present in our reference genome (23, 24, 38-40).

To overcome these limitations, we focused on the residual information encoded in the shapes of the SNV trajectories in Fig. 3. We developed a new statistical test to determine whether the SNV trajectories in Fig. 3 were consistent with a neutral model with a constant but unknown strength of genetic drift (Methods). This test leverages the fact that, under a constant genetic drift model, the changes in frequency along a single trajectory must be statistically similar to each other, so that a large change in one time-interval is unlikely to be followed by a small change in another interval. The observed trajectories often violate this prediction, and we find significant evidence against the constant genetic drift model for 4 of the 5 species in Fig. 3 (Table S3).

A second possibility is that genetic drift is elevated only during antibiotic treatment, e.g. due to a transient population bottleneck. This could be a particularly plausible hypothesis for the *A. fingoldii* population in Fig. 3A, where the genetic changes coincided with a dramatic reduction in the relative abundance of that species. While it is difficult to rule out similar bottlenecks at unobserved timepoints for the other species in Fig. 3, we still observe significant departures from the constant genetic drift model when the antibiotic timepoints are excluded (Table S3). Closer inspection of the trajectories reveals the likely source of this signal: many of the SNV clusters continue to change in frequency, but in the opposite direction, even after antibiotic treatment has concluded. This behavior, which is recapitulated across the larger set of species in Fig. 2E, cannot be explained by a simple bottleneck during treatment. Instead, we conclude that the initial increases and later reversals are most likely caused by time-dependent selection pressures that act on different haplotypes within these populations.

The high temporal resolution of the SNV clusters yields additional information about the fitness differences between the different haplotypes. For example, the frequency reversals after treatment in Fig. 3 occurred over ~30-40 days, implying a fitness difference about

~10% per day (Methods). The increases in frequency during antibiotics can be even more rapid. In *E. eligens*, the minority haplotype increased from 7% to 90% in just two days, requiring a corresponding fitness difference of at least ~250% per day. We also observe a ~20% increase in the PTR-estimated growth rate between these two timepoints (Fig. S2). This suggests that the dramatic fitness difference arises from a higher growth rate of the sweeping haplotype, rather than an increased death rate of the declining strain.

Statistics of sweeping SNVs reveal strain replacement, evolutionary modification, and selection on standing variation

After demonstrating that genetic changes occur within species, and that these changes are likely driven by selection on linked haplotypes and not necessarily associated with changes in species abundance, we next sought to investigate the origin of these within-host sweeps. A key question is whether the temporally variable SNVs arose *de novo* within the host (*evolutionary modification*) or whether they reflect the invasion of pre-existing strains that diverged for many generations before colonizing the host (*strain replacement*). Following previous work (21), we distinguished between these two scenarios by examining three additional features of the SNV trajectories in Fig. 3: (i) the protein-coding impact of these mutations, (ii) their prevalence across other hosts in a large reference panel (Methods), and (iii) the retention of private marker SNVs (i.e., high-frequency alleles that are unique to the present host). Figure 5 illustrates these quantities for the six example species in Fig. 3, which were chosen to cover the full range of different behaviors we observed.

The *E. eligens* population provides a striking example of strain replacement. The sweeping haplotype in this species contains more than 10,000 SNVs that are widely distributed across the genome (Supplementary Data 1), consistent with the typical genetic differences between *E. eligens* strains in different hosts (18, 20, 21). Few private marker SNVs are retained from the initial timepoint (Fig. 5D), which is again consistent with replacement by a distantly related strain. Similar examples of strain replacement have been observed previously (6, 10, 18, 21, 39) but our densely sampled timecourse provides new information about the dynamics of this process. The SNV frequency trajectories in Fig. 3D show that the distantly related strain was already present at substantial frequencies (~5%) long before its dramatic fitness difference was revealed. Fig. 2D shows that this is also the case for the four other putative replacements in Fig. 2. This indicates that the replacement events we observe here are caused by the sudden increase of previously colonizing strains, rather than the contemporary invasion of new strains from outside the host.

At the opposite extreme, the *Phascolarctobacterium* population in Fig. 3E provides a prototypical illustration of an evolutionary modification. In this case, a cluster of just 6 SNVs (including 5 amino acid changes, all in non-contiguous genes in the reference genome that are unlinked in our read clouds, Supplementary Data 1 and 2) nearly swept to fixation during antibiotic treatment ($f > 99.8\%$, $S > 30\%$ per day), only to decline in frequency later in the study. Unlike the replacement event above, this sweep shared all 42 of the private marker SNVs from the dominant strain at the initial timepoint (Fig. 5E), suggesting that they recently descended from a common ancestor. Interestingly, however, we again observe evidence that the minority haplotype was segregating at substantial frequencies (~1-10%) before treatment, a finding which is recapitulated for several other non-replacement examples in Fig. 2D. This suggests that frequency-dependent selection may have initially driven these mutant lineages to intermediate frequencies – and maintained

them there – before antibiotics or other environmental changes (or subsequent mutations) caused them to sweep through the rest of the population.

In addition to these extreme cases, we also observed a third category of events that seem to bridge the divide between strain replacement and evolutionary modification. For example, in the *Alistipes finegoldii* population in Figs. 3A and 5A, a cluster of ~80 SNVs swept to high frequency when the species recovered from antibiotic treatment, potentially consistent with a population bottleneck. While the high rates of private marker SNVs sharing (52/55) suggest that the sweeping haplotype is a modification of the dominant strain from the initial timepoint, the large fraction of synonymous mutations ($dN/dS=0.16$), many of which are shared across other hosts, is more consistent with a strain replacement event. In contrast to the two examples above, the *A. finegoldii* SNVs fall into a smaller number of contiguous genes in the reference genome and are often linked together in the same read clouds (Supplementary Data 1 and 2). These same SNVs are also frequently co-inherited in “haplotype blocks” among the other hosts in our reference panel (Fig. S11). Taken together, these lines of evidence suggest that the SNVs in Fig. 3A were likely transferred onto their current genetic background through recombination. Similar to the *E. eligens* and *Phascolarctobacterium* examples above, the sweeping haplotype in *A. finegoldii* was already segregating as a minor variant ($f \sim 20\%$) before antibiotic treatment, suggesting that the original recombination event (and its initial rise to observable frequencies) predated the current sampling period.

The *Bacteroides coprocola* population in Figs. 3F and 5F provides another interesting example. In this case, a cluster of 37 SNVs (including reversions of 11 of the 167 private marker SNVs) was already in the process of sweeping through the population before treatment began. In this case, however, the mutations are scattered across many non-contiguous genes in the *B. coprocola* reference genome and are seldom observed in other hosts, so recombination no longer provides a parsimonious explanation. The fraction of synonymous mutations ($dN/dS=0.7$) also lies somewhere between the typical between-host values ($dN/dS \sim 0.1$) and within-host hitchhiking ($dN/dS \geq 1$). This suggests that the lineages may have coexisted with each other for a much longer period of time.

Discussion

The response of gut microbial communities to antibiotics plays a crucial role in their susceptibility to pathogens (7, 41, 42), the spread of antibiotic resistance genes (43, 44), and their long-term stability (8, 45, 46). Numerous studies have documented the resilience of these communities at the taxonomic or pathway level (7-16). Yet the strain-level dynamics that give rise to this ecological robustness remain poorly characterized.

In this study, we sought to characterize these within-species dynamics by following the gut microbiome of a single individual through a period of health, disease, and the oral administration of doxycycline. We used linked read metagenomic sequencing to track the dynamics of single nucleotide variants within 36 different species, and to contrast these within-species dynamics with the broader ecological shifts at the species level. Consistent with our expectations, we found that antibiotic perturbations can lead to widespread shifts in the genetic composition of individual species, and at a higher overall rate than observed in healthy hosts (21-23). However, this genetic response was rarely consistent with the traditional picture of extinction and subsequent recolonization. Instead, we found that genetic responses varied widely across species, with some species accumulating

thousands of SNV differences at the population level, and others accumulating only a handful. These genetic changes were frequently observed in species without large changes in relative abundance in the sampled timepoints, and conversely, large abundance fluctuations were not always accompanied by widespread genetic changes. Furthermore, some of the most dramatic fluctuations at both the species and SNV level occurred in the weeks following the competition of treatment. Together, these findings suggest that the response to antibiotics is not driven by discrete recolonization events, but rather, by the subtler processes of strain-level competition and evolution within the host.

At this population genetic level, our observations revealed qualitative departures from the simplest models of neutral evolution, or the spread of antibiotic resistance phenotypes via classic selective sweeps. The observed genetic responses were much more dynamic: we often observed partial genome-wide sweeps containing multiple linked genetic variants, many of which were segregating at observable frequencies before the onset of treatment. Although their frequencies could increase dramatically on daily or weekly timescales, few of these variants ever fixed in their respective populations. Instead, we observed frequent reversion of sweeps at the single base pair level, consistent with temporally varying selection pressures and strong pleiotropic tradeoffs. These reversions rarely ended in extinction, and more commonly stabilized close to the initial pre-treatment frequency. Together, these dynamics suggest that the sweeping haplotypes may be stably maintained in their respective populations over time, e.g. due to metabolic or spatial niches. This provides a potential explanation for the “oligo-colonization” structure observed in a variety of within-host microbial populations (18, 21, 47). Interestingly, our data show that similar dynamics can occur for mixtures of distantly related strains, as well as for haplotypes that likely evolved *de novo* within the host. This suggests that ongoing ecological diversification could play an important role in shaping the genetic structure of resident populations, echoing a previous finding in *B. fragilis* (23).

There are several important limitations to our study. Since we have focused primarily on single nucleotide variants in well-behaved regions of reference genomes, we are likely missing many of the true targets of selection, particularly in the case of antibiotic resistance where mobile elements (40, 48) and other structural rearrangements (38) are known to play an important role. This makes it difficult to know what fraction of genetic changes are a direct response to antibiotics, as opposed to indirect responses produced by fluctuations in the abundances of other species. It is even possible that nearly all of the mutations that we observe are simply passenger mutations that are hitchhiking alongside the true causative variants. The situation could potentially be improved by combining our approach with *de novo* assembly, similar to previous studies (39, 49, 50). However, given the high levels of genetic linkage we observed, it would be difficult to pinpoint individual selection pressures even with an exhaustive list of mutations, since one can only observe the net effects of selection across entire haplotypes. Our reference-based approach is effectively using this limitation to our advantage, by relying on the dynamics of linked passengers to provide information about the net selection pressures on their corresponding haplotypes.

In addition to these methodological constraints, a second key limitation is our focus on a single host microbiome. While the concentrated resources allowed us to observe a variety of different responses across individual species in the same community, further work will be required to establish the prevalence of these different patterns across larger cohorts, and among different classes of antibiotics. Our high-resolution time course provides a valuable set of templates that can inform these future classification efforts in larger, but lower resolution studies.

In summary, by tracking a host microbiome through periods of disease, treatment, and recovery, we uncovered new evidence that the ecological resilience of microbial communities might extend all the way down to the genetic level. Understanding how this resilience arises from the complex interplay between host genetic, epigenetic, and lifestyle factors, as well its implications for broader evolution of the microbiome, remains an exciting avenue for future work.

FIGURE CAPTIONS

Figure 1: Read cloud sequencing of the gut microbiome of a single individual during disease, antibiotic treatment, and recovery. **a**, Study design. Linked read metagenomic sequencing was performed on 19 fecal samples collected from a single individual over a period of 5 months. During this time, the individual was diagnosed with Lyme disease and received an oral course of doxycycline. **b**, Species-level dynamics over time, estimated from shotgun metagenomic reads (Methods). The “-2yr” sample is taken from a previous study of the same host (27). **c**, Schematic of linked read sequencing with the 10X Genomics platform. High molecular weight metagenomic DNA is partitioned into millions of microfluidic droplets. Amplification and ligation reactions are performed within each droplet, yielding millions of short-read libraries that are tagged with droplet-specific DNA barcodes. The resulting “read clouds” are then pooled together and sequenced on an Illumina instrument. **d**, Observed statistics of read clouds from the first three timepoints. The top panel shows total number of read pairs contributed by read clouds as a function of the number of read pairs they contain. The bottom panel shows a measure of the effective number of species that are detected in each read cloud as a function of the number of read pairs it contains (Methods). Many read clouds contain fragments from several different DNA molecules.

Figure 2: Varied ecological and genetic responses across 36 abundant species in the same host. **a**, Relative abundances of species through time, partitioned according to the epochs defined in Table S1. Each timepoint is indicated by a point, and the timepoints from the same epoch are connected by a vertical line to aid in visualization. For comparison, the grey distribution shows the corresponding values from the Human Microbiome Project (51) cohort (Methods). Species whose relative abundance drops by more than 10-fold between baseline and antibiotic timepoints are indicated with a single star. Only a minority of the most abundant species experience such reductions in relative abundance during treatment. **b**, Within-species nucleotide diversity for each timepoint, as measured by the fraction of core genome sites with intermediate allele frequencies ($0.2 < f < 0.8$, Methods). Points are plotted according to the same scheme as in (a). **c**, The total number of single nucleotide (SNV) differences between a baseline timepoint and each of the later epochs (Methods). The height of the white area indicates the total number of polymorphic SNVs that were tested for temporal variation. Different species display a range of different behaviors, which can be partitioned into putative cases of competition between distantly related strains, and evolution within a dominant resident strain. **d**, Initial frequencies of alleles identified in (c). For species with more than 10 SNV differences, the data are summarized by the median initial frequency (square symbol) and the interquartile range (line). Many alleles have nonzero frequency before the sweep occurs. **e**, Fraction of SNV differences in (c) that are retained at the final timepoint ($f > 0.7$). In many species, only a minority of SNV differences gained during disease or treatment are retained.

Figure 3: Ecological and genetic dynamics in six example species. A subset of the species in Fig. 2 were chosen to illustrate a range of characteristic behaviors (**a-f**). For each of the six species, the top panel shows the relative abundance of that species over time, while the bottom panel shows the frequencies of single nucleotide variants (SNVs) within that species. Colored lines indicate SNVs that underwent a significant shift in frequency over time (Methods), while a subset of non-significant SNVs are shown in light grey for comparison. The colors of temporally varying SNVs are assigned by a hierarchical clustering scheme, which is also used to determine their polarization (Methods).

Figure 4: High levels of linkage disequilibrium (LD) in many resident populations.

a, Schematic of read cloud sharing between two SNVs separated by coordinate distance ℓ on the same reference contig. Three or fewer haplotypes are consistent with clonal evolution, while four haplotypes indicate a possible recombination event. **b**, Observed fraction of shared read clouds as a function of ℓ for SNVs in the six example species in Fig. 3. **c**, Total number of linked SNV pairs (i.e., those with significantly elevated levels of read cloud sharing) for species in Fig. 2 with sufficient coverage (Methods). For each species, the three bars denote SNV pairs with $\ell < 200\text{bp}$, $200\text{bp} < \ell < 2\text{kb}$, and $\ell > 2\text{kb}$, respectively. SNVs are included only if the minor allele has frequency $f > 0.1$. **d**, Observed proportion of SNV pairs in (c) that fall each of the LD categories illustrated in (a). Across species, only a small fraction of SNV pairs provide evidence for recombination.

Figure 5: Signatures of strain replacement and evolutionary modification. a-f,

Statistical properties of temporally varying SNVs from the six example species in Fig. 3. For each species, the bars on the left show the relative proportion of SNVs with different protein coding effects (left) and allele prevalence across other hosts in a larger cohort (right) (Methods). SNVs that are not observed in other hosts are shaded in light red or blue. For each species, the pie chart indicates the relative proportion of private marker SNVs that are preserved or disrupted throughout the sampling interval (Methods). Large fractions of disrupted marker SNVs indicate a strain replacement event.

List of Supplementary Files

Supplementary Information. Supplementary Methods, Supplementary Figures S1-S14, Supplementary Tables S1-S3, Supplementary Data 1-7, and Supplementary references.

ACKNOWLEDGEMENTS

We thank Eitan Yaffe for comments on the manuscript. This work was funded in part by the US National Institutes of Health grants U54DK10255603, R01AT01023202, and 2RM1HG00773506. Sequencing was performed at the Stanford Center for Genomics and Personalized Medicine supported by US National Institutes of Health grant S10OD020141. N.R.G. and K.S.P. acknowledge support from the US National Science Foundation (DMS-1563159), the Chan Zuckerberg Biohub, and the Gladstone Institutes. B.H.G. acknowledges support from the Miller Institute for Basic Research in Science.

AUTHOR CONTRIBUTIONS

M.R. and M.P.S. conceived the study; B.H.G. designed the analysis; M.R. and M.A. developed the HMW DNA extraction protocol and performed the experiments; S.L., H.L., A.B., M.R., S.N., and W.Z. performed sequencing QC and preliminary bioinformatic analyses; B.H.G., N.R.G. and S.M. developed the metagenomic pipeline and analyzed SNV data; B.H.G., S.M., and K.S.P. developed theory and statistical methods; K.S.P. and M.P.S. supervised the study; M.R. and B.H.G. wrote the paper; M.R., B.H.G., N.R.G., K.S.P., and M.P.S. edited the paper.

AUTHOR INFORMATION

The authors declare no competing financial interests. K.S.P. is a consultant for Phylagen and uBiome. Correspondence and requests for materials should be addressed to B.H.G. (bhgood@stanford.edu), K.S.P. (katherine.pollard@gladstone.ucsf.edu), or M.P.S. (mpsnyder@stanford.edu).

REFERENCES CITED

1. G. Sharon, T. R. Sampson, D. H. Geschwind, S. K. Mazmanian, The Central Nervous System and the Gut Microbiome. *Cell* **167**, 915-932 (2016).
2. Q. Feng *et al.*, Gut microbiome development along the colorectal adenoma-carcinoma sequence. *Nat Commun* **6**, 6528 (2015).
3. J. Halfvarson *et al.*, Dynamics of the human gut microbiome in inflammatory bowel disease. *Nat Microbiol* **2**, 17004 (2017).
4. S. R. Gill *et al.*, Metagenomic analysis of the human distal gut microbiome. *Science* **312**, 1355-1359 (2006).
5. P. Spanogiannopoulos, E. N. Bess, R. N. Carmody, P. J. Turnbaugh, The microbial pharmacists within us: a metagenomic view of xenobiotic metabolism. *Nature Reviews Microbiology* **14**, 273 (2016).
6. J. Lloyd-Price *et al.*, Strains, functions and dynamics in the expanded Human Microbiome Project (vol 550, pg 61, 2017). *Nature* **551**, (2017).
7. C. G. Buffie *et al.*, Precision microbiome reconstitution restores bile acid mediated resistance to *Clostridium difficile*. *Nature* **517**, 205 (2015).
8. L. Dethlefsen, D. A. Relman, Incomplete recovery and individualized responses of the human distal gut microbiota to repeated antibiotic perturbation. *Proceedings of the National Academy of Sciences* **108**, 4554-4561 (2011).
9. K. M. Ng *et al.*, Recovery of the Gut Microbiota after Antibiotics Depends on Host Diet, Community Context, and Environmental Reservoirs. *Cell Host Microbe* **26**, 650-665. e654 (2019).
10. M. Yassour *et al.*, Natural history of the infant gut microbiome and impact of antibiotic treatment on bacterial strain diversity and stability. *Sci Transl Med* **8**, 343ra381 (2016).
11. C. Jernberg, S. Löfmark, C. Edlund, J. K. Jansson, Long-term impacts of antibiotic exposure on the human intestinal microbiota. *Microbiology* **156**, 3216-3223 (2010).
12. K. M. Keeney, S. Yurist-Doutsch, M.-C. Arrieta, B. B. Finlay, Effects of antibiotics on human microbiota and subsequent disease. *Annual review of microbiology* **68**, 217-235 (2014).
13. E. Zaura *et al.*, Same Exposure but Two Radically Different Responses to Antibiotics: Resilience of the Salivary Microbiome versus Long-Term Microbial Shifts in Feces. *MBio* **6**, e01693-01615 (2015).
14. F. Raymond *et al.*, The initial state of the human gut microbiome determines its reshaping by antibiotics. *ISME J* **10**, 707-720 (2016).
15. A. Palleja *et al.*, Recovery of gut microbiota of healthy adults following antibiotic exposure. *Nature Microbiology* **3**, 1255-1265 (2018).
16. J. Yin, X.-X. Zhang, B. Wu, Q. Xian, Metagenomic insights into tetracycline effects on microbial community and antibiotic resistance of mouse gut. *Ecotoxicology* **24**, 2125-2132 (2015).
17. M. Scholz *et al.*, Strain-level microbial epidemiology and population genomics from shotgun metagenomics. *Nature methods* **13**, 435 (2016).
18. D. T. Truong, A. Tett, E. Pasolli, C. Huttenhower, N. Segata, Microbial strain-level population structure and genetic diversity from metagenomes. *Genome Res* **27**, 626-638 (2017).

19. D. V. Ward *et al.*, Metagenomic sequencing with strain-level resolution implicates uropathogenic *E. coli* in necrotizing enterocolitis and mortality in preterm infants. *Cell reports* **14**, 2912-2924 (2016).
20. S. Schloissnig *et al.*, Genomic variation landscape of the human gut microbiome. *Nature* **493**, 45-50 (2013).
21. N. R. Garud, B. H. Good, O. Hallatschek, K. S. Pollard, Evolutionary dynamics of bacteria in the gut microbiome within and across hosts. *PLoS biology* **17**, e3000102 (2019).
22. M. Ghalayini *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. *Appl Environ Microb* **84**, (2018).
23. S. Zhao *et al.*, Adaptive evolution within gut microbiomes of healthy people. *Cell Host Microbe* **25**, 656-667. e658 (2019).
24. A. Bishara *et al.*, High-quality genome sequences of uncultured microbes by assembly of read clouds. *Nat Biotechnol*, (2018).
25. D. C. Danko, D. Meleshko, D. Bezdan, C. Mason, I. Hajirasouliha, Minerva: an alignment-and reference-free approach to deconvolve Linked-Reads for metagenomics. *Genome Res* **29**, 116-124 (2019).
26. S. Nayfach, B. Rodriguez-Mueller, N. Garud, K. S. Pollard, An integrated metagenomics pipeline for strain profiling reveals novel patterns of bacterial transmission and biogeography. *Genome Res* **26**, 1612-1625 (2016).
27. V. Kuleshov, M. P. Snyder, S. Batzoglou, Genome assembly from synthetic long read clouds. *Bioinformatics* **32**, i216-i224 (2016).
28. D. A. Relman, M. Lipsitch, Microbiome as a tool and a target in the effort to address antimicrobial resistance. *Proceedings of the National Academy of Sciences* **115**, 12902-12910 (2018).
29. C. Luo *et al.*, ConStrains identifies microbial strains in metagenomic datasets. *Nat Biotechnol* **33**, 1045-1052 (2015).
30. J. D. Lewis *et al.*, Inflammation, Antibiotics, and Diet as Environmental Stressors of the Gut Microbiome in Pediatric Crohn's Disease. *Cell Host Microbe* **18**, 489-500 (2015).
31. B. Rasmussen, K. Bush, F. Tally, Antimicrobial resistance in *Bacteroides*. *Clinical infectious diseases* **16**, S390-S400 (1993).
32. M. Vos, X. Didelot, A comparison of homologous recombination rates in bacteria and archaea. *The ISME journal* **3**, 199 (2009).
33. R. R. Hudson, N. L. Kaplan, Statistical properties of the number of recombination events in the history of a sample of DNA sequences. *Genetics* **111**, 147-164 (1985).
34. C. S. Smillie *et al.*, Strain tracking reveals the determinants of bacterial engraftment in the human gut following fecal microbiota transplantation. *Cell Host Microbe* **23**, 229-240. e225 (2018).
35. M. Kimura, *The neutral theory of molecular evolution*. (Cambridge University Press, 1983).
36. S. P. Hubbell, *The unified neutral theory of biodiversity and biogeography (MPB-32)*. (Princeton University Press, 2001).
37. A. B. Paaby, M. V. Rockman, Cryptic genetic variation: evolution's hidden substrate. *Nature Reviews Genetics* **15**, 247 (2014).

38. X. Jiang *et al.*, Invertible promoters mediate bacterial phase variation, antibiotic resistance, and host adaptation in the gut. *Science* **363**, 181-187 (2019).
39. E. Yaffe, D. A. Relman, Tracking microbial evolution in the human gut using Hi-C. *bioRxiv*, 594903 (2019).
40. A. Bishara *et al.*, Strain-resolved microbiome sequencing reveals mobile elements that drive bacterial competition on a clinical timescale. *BioRxiv*, 125211 (2018).
41. M. Sassone-Corsi, M. Raffatellu, No vacancy: how beneficial microbes cooperate with immunity to provide colonization resistance to pathogens. *The Journal of Immunology* **194**, 4081-4087 (2015).
42. E. A. Cameron, V. Sperandio, Frenemies: signaling and nutritional integration in pathogen-microbiota-host interactions. *Cell Host Microbe* **18**, 275-284 (2015).
43. Y. Hu *et al.*, Metagenome-wide analysis of antibiotic resistance genes in a large cohort of human gut microbiota. *Nature communications* **4**, 2151 (2013).
44. B. Kintses *et al.*, Phylogenetic barriers to horizontal transfer of antimicrobial peptide resistance genes in the human gut microbiota. *Nature microbiology* **4**, 447 (2019).
45. N. Kamada *et al.*, Regulated virulence controls the ability of a pathogen to compete with the gut microbiota. *Science* **336**, 1325-1329 (2012).
46. J. J. Faith *et al.*, The long-term stability of the human gut microbiota. *Science* **341**, 1237439 (2013).
47. M. Yassour *et al.*, Strain-level analysis of mother-to-child bacterial transmission during the first few months of life. *Cell Host Microbe* **24**, 146-154. e144 (2018).
48. S. R. Partridge, G. Tsafnat, Automated annotation of mobile antibiotic resistance in Gram-negative bacteria: the Multiple Antibiotic Resistance Annotator (MARA) and database. *Journal of Antimicrobial Chemotherapy* **73**, 883-890 (2018).
49. A. Crits-Christoph, M. Olm, S. Diamond, K. Bouma-Gregson, J. F. Banfield, Soil bacterial populations are shaped by recombination and gene-specific selection across a meadow. *bioRxiv*, 695478 (2019).
50. C. Quince *et al.*, DESMAN: a new tool for de novo extraction of strains from metagenomes. *Genome biology* **18**, 181 (2017).
51. H. Integrative, The Integrative Human Microbiome Project: dynamic analysis of microbiome-host omics profiles during periods of human health and disease. *Cell Host Microbe* **16**, 276 (2014).

Figures

Figure 1: Read cloud sequencing of the gut microbiome of a single individual during disease, antibiotic treatment, and recovery.

Figure 2: Varied ecological and genetic responses across 36 abundant species in the same host.

Figure 3: Ecological and genetic dynamics in six example species.

Figure 4: High levels of linkage disequilibrium (LD) in many resident populations.

Figure 5: Signatures of strain replacement and evolutionary modification.