Strain-level bacterial reconstruction and inference in patients receiving FMT for UC

Microbiome Ignite
Byron J. Smith
2020-03-20

Thanks for the introduction and thanks for joining me for my presentation.
It's great to be here, virtually, today to share some preliminary results from an exciting in-progress study.

I’d like to start by acknowledging the vital support I’ve had both from Katie and my lab

As well as a number of collaborators, institutions, and funding sources.

Ulcerative Colitis

Nearly 1M Americans have been diagnosed with UC, trending upwards

Characterized by chronic inflammation of the colon

Patients stratified by disease severity and what drugs they respond to

Associated with the microbiome (maybe causally) but mechanism is unknown

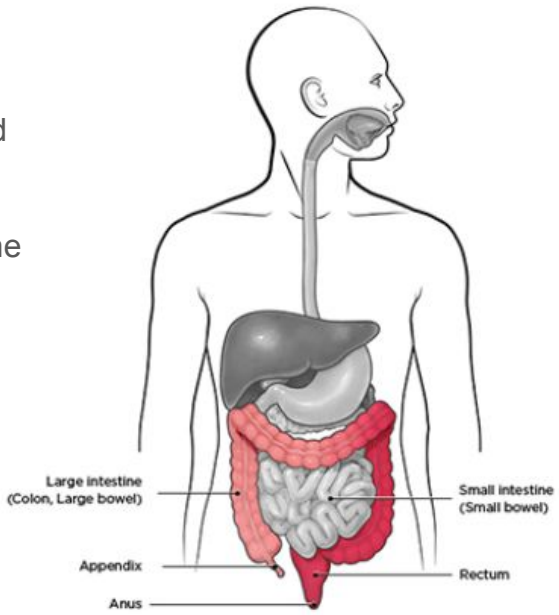

Ulcerative colitis is a one of the two types of IBD
affects nearly 1M americans and the prevalence is increasing.
The disease is characterized by chronic inflammation of the colon,

We do know that it's associated with the microbiome,
and there's fairly strong evidence that this can be a causal relationship,
but the exact mechanism remains unknown.

FMT treats UC in some patients

Approximately 36% remission after FMT, with favorable results in randomized, controlled trials

No established mechanism in successful treatment, and **inconsistent associations with microbiome characteristics**

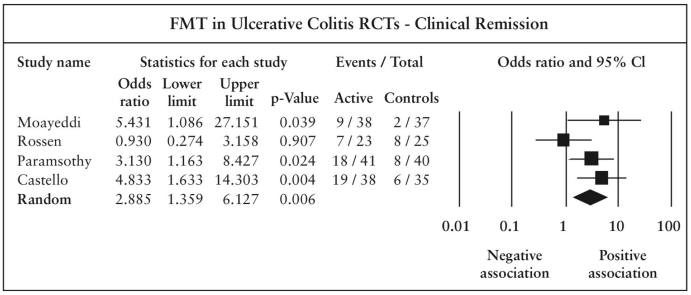

Paramsothy et al. 2017

FMT has been tested as a treatment with favorable results, including in randomized controlled trials.
This is a figure from a meta-analysis of the handful of such studies showing greater improvement in FMT recipients than in controls.

However, it certainly does not have the success rate of the flagship FMT use-case, C. difficile infection,
and the mechanism of successful treatment is very unclear.

How do we explain variability of clinical outcomes?

Could enable:

- Improved understanding of mechanisms of disease/recovery
- Better screening for patient or donor suitability
- More precise therapeutics

Taxonomic resolution in past studies may have been *insufficient* to identify changes in microbial **function**

*Hypothesis: The relevant microbial functions will be found at the level of **strain** rather than species.*

Current research should focus on explaining the variability in clinical outcomes from FMT
Understanding this variability will enable:
() Basic scientific understanding of the mechanisms of disease and recovery
() Greater targeting of therapies to the patients they are most likely to help
() The design of therapies that are more precise than simple FMT,
perhaps individual bacteria or metabolites or drugs that target relevant pathways

In the work that I'm sharing today, I'm exploring the possibility that
past studies have not been able to link microbiome characteristics of the recipient or donor stool
because the taxonomic resolution of microbiome analysis has been insufficient.

Specifically, I'm hypothesizing that relevant microbial functions are *strain specific*, and their presence or absence is not reflected in traditional taxonomic analysis.

The methodological approach that I am taking to explore this hypothesis
can be broadly described as **strain reconstruction** using metagenomic data.

FMT for UC study design

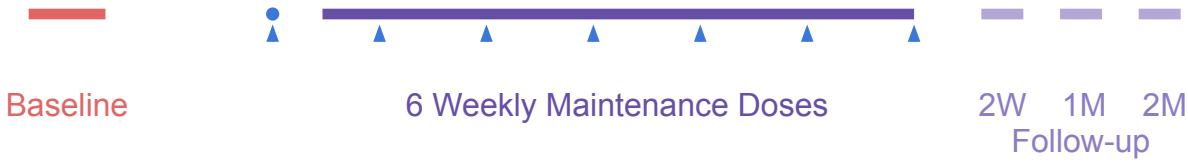

The data that I'll be sharing today comes from an FMT trial being carried out here at UCSF

Donor material is applied to patients repeatedly for six weeks, starting with direct application during a colonoscopy, followed by weekly maintenance doses.

FMT for UC study design

Patient response is assessed via colonoscopy,
comparing the baseline to a second scope at the 2-week followup.

FMT for UC study design

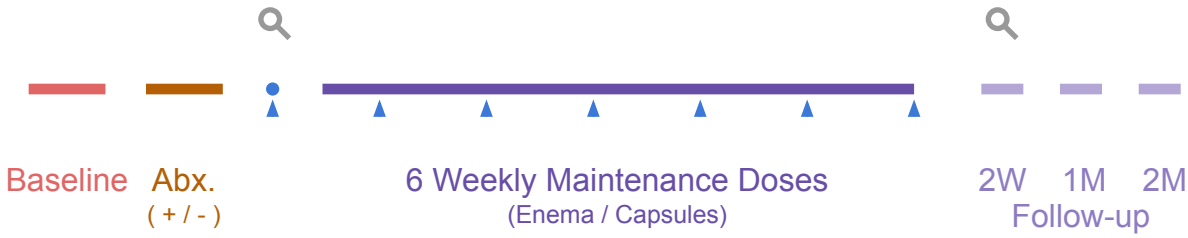

Patients are randomly assigned to arms determining
Whether or not they receive antibiotic pretreatment prior to FMT application as well as the mode of delivery of the maintenance doses, either via enema or capsules.

FMT for UC study design

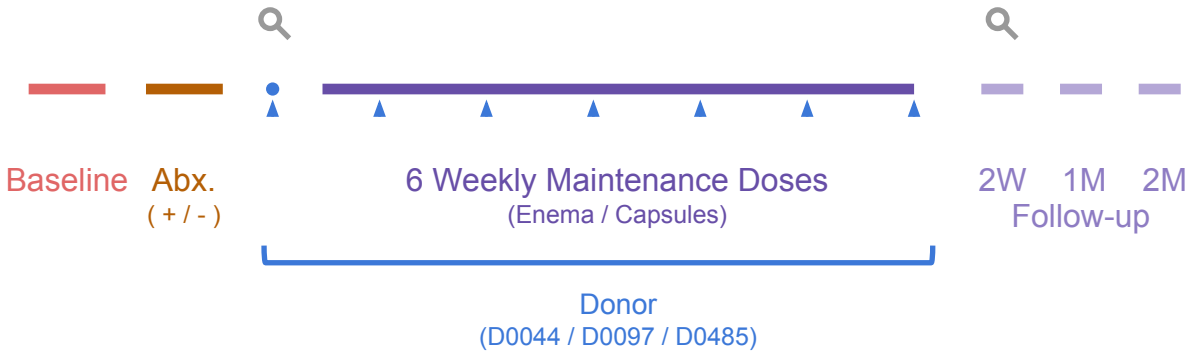

The study is designed so that all transplants come from the same donor
With three different donors in the sample set I'll describe today.

FMT for UC study design

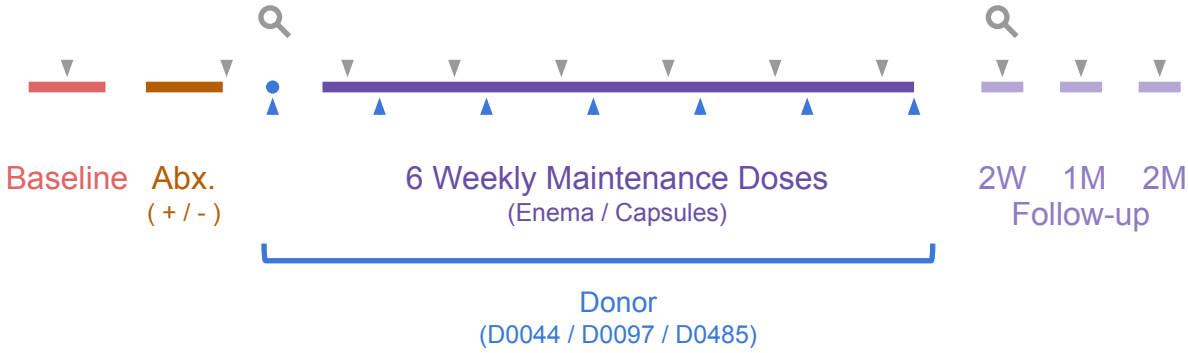

Fecal samples were collected at numerous time points, including a single baseline sample,
Six samples during the maintenance period, and three followup samples.

FMT for UC study design

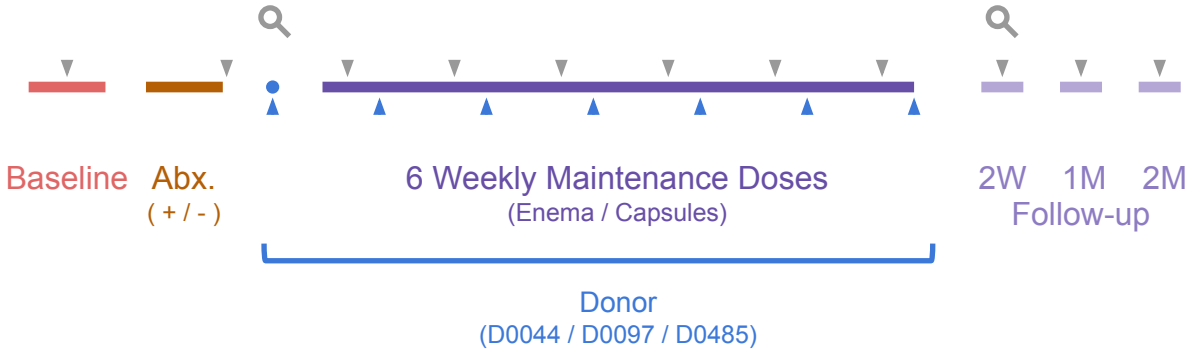

Currently n ≈ 12 with 16S, Metagenomics, Metabolomics

I'll describe data from twelve subjects, for whom we have 16S, metagenomics, and metabolomics data.

My part of this work

Track engraftment from donors to recipients

Analyze functional potential of the microbiome in metagenomic time-series

Question: Can strain-reconstruction provide novel insights into this system?

This study has required a diverse set of experts

And my role has been to track the engraftment of the microbiota from donors to patients
And to use primarily metagenomic data to understand temporal dynamics in the taxonomic composition and functional potential of the microbiome

I'd like to know if there's something to be learned from strain-level inference in these data.
"Can strain-reconstruction provide novel insights into this system?"

"*Strain*" reconstruction

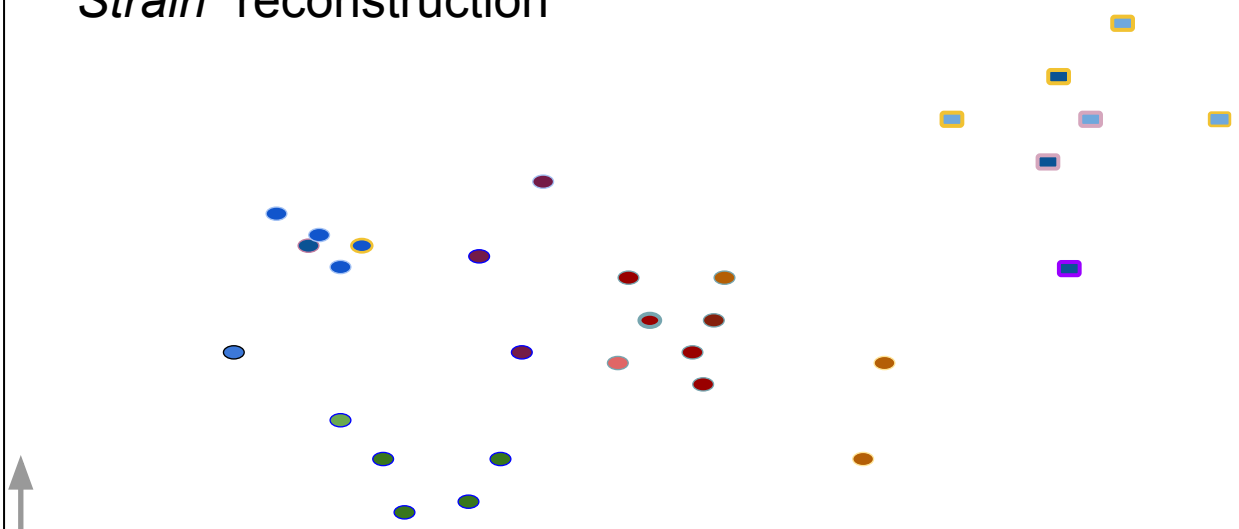

To better explain what I mean by "strain reconstruction", first I'd like to tell you what I mean by "strain".

Here I've drawn a cartoon version of some small amount of bacterial diversity space, with individual genotypes (colored ovals) spread throughout

Classically, we think of a strain of bacteria as a single type, just one of these ovals. With genetic information we usually mean a single genotype, although mutation in the lab often means that
an isolated strain still has very limited amounts of diversity.

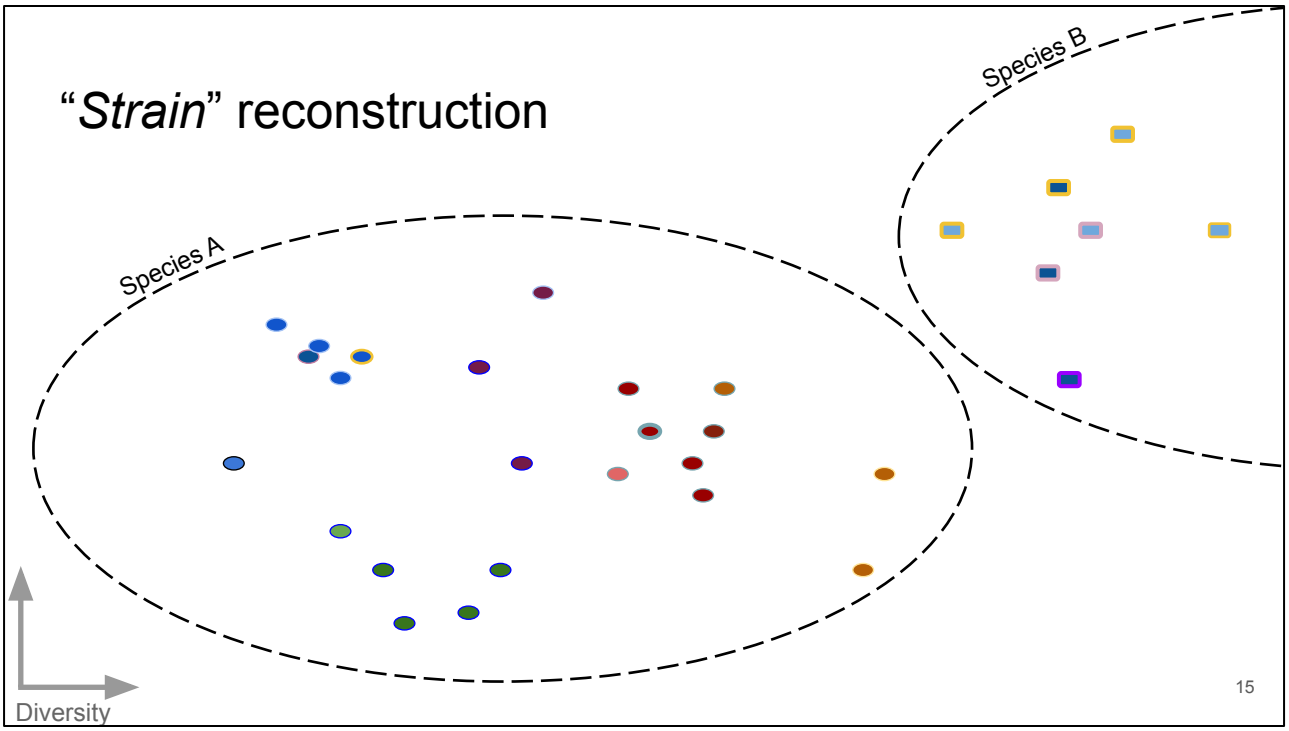

Of course this is much more precise than what we currently achieve using SOPs like 16S.
Instead, we're lucky to resolve individual species of bacteria, each of which may encompass a broad
swath of this diversity space.

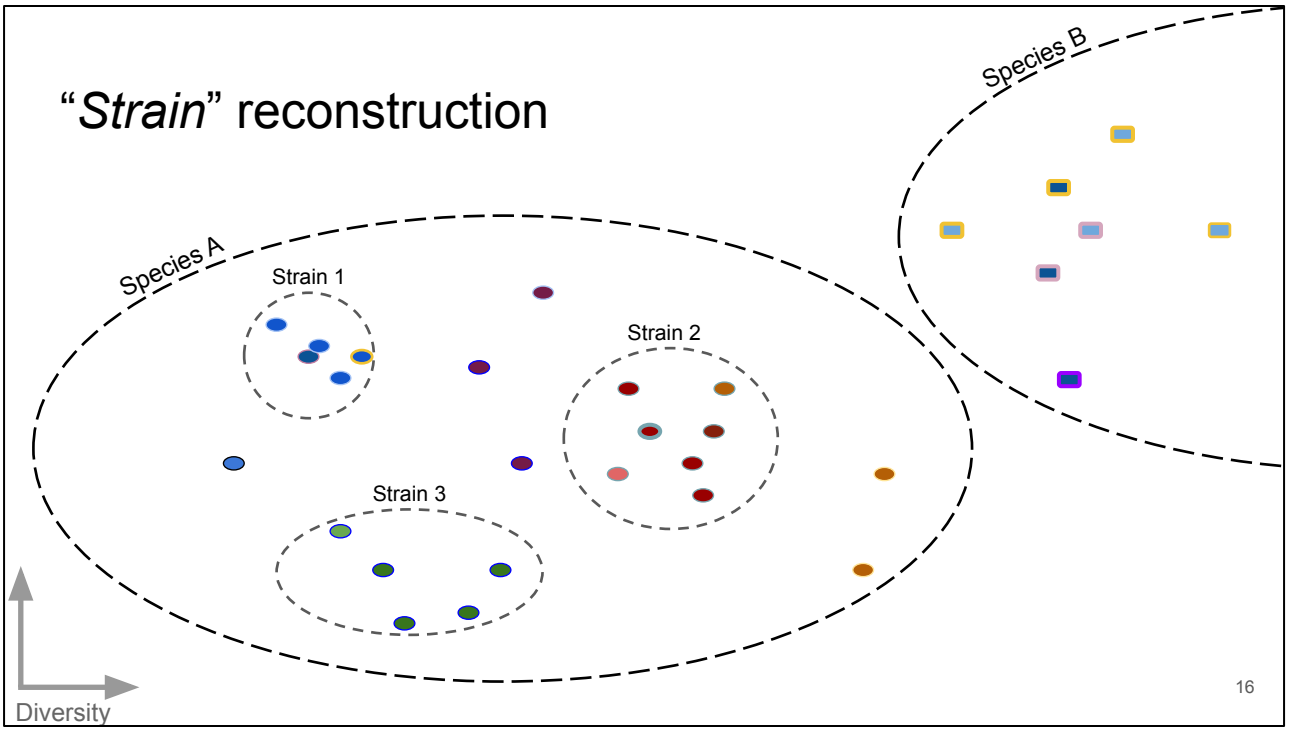

While I'd love to be able to resolve individual genotypes,
This is unrealistic given the limitations of sequencing depth and technical variability.

Instead, I use the term "strain" loosely here to denote genetic variants that can be distinguished **given** the observations we have available.

Strain "*reconstruction*"

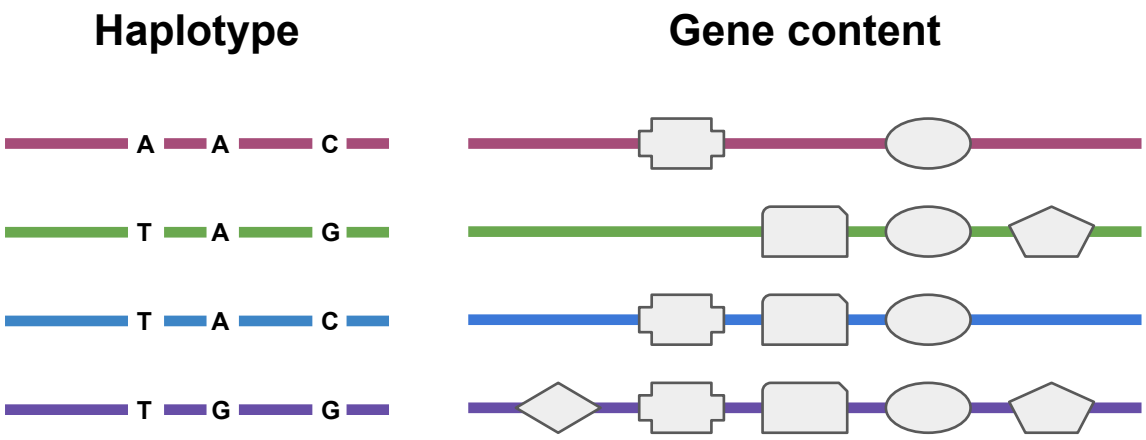

17

Next, I'd like to explain what I mean by "reconstruction".

Really, I mean two different things:

(1) Identifying the patterns of alleles at variable positions in the genome (the haplotype)

(2) Determining which genes are present in the genomes of individual strains.

The former can inform our understanding of the evolutionary relationships between strains

While the latter lends itself to making predictions about their functional potential

Strain haplotypes from metagenomes

Strain A

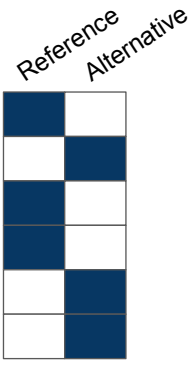

So, how do I go about reconstructing haplotypes?

A single strain is defined by a pattern of bases at variable positions
here depicted as colored boxes for either the reference or the alternative allele.

Strain haplotypes from metagenomes

Strain A Strain B

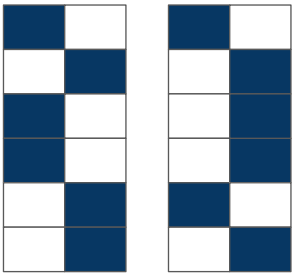

Different strains may differ in a subset of these positions
While sharing alleles in other positions.

Strain haplotypes from metagenomes

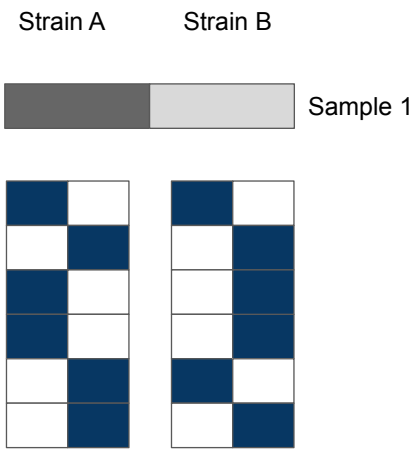

A single sample is composed of a mixture of one or more of these strains.

Strain haplotypes from metagenomes

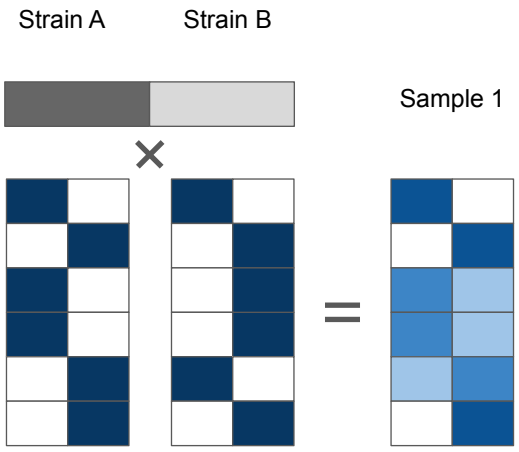

What we're able to observe through metagenomics are the frequencies of each base at each position in each sample:
the sample's genotype

These represent the product of the latent haplotypes and their (also unmeasured) relative abundances.

Strain haplotypes from metagenomes

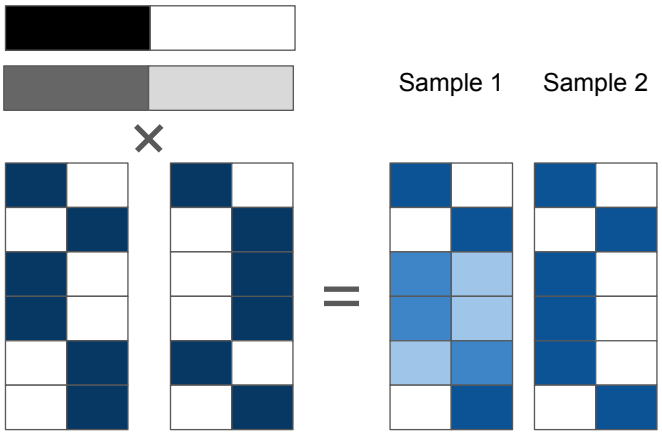

[Next Slide]

Strain haplotypes from metagenomes

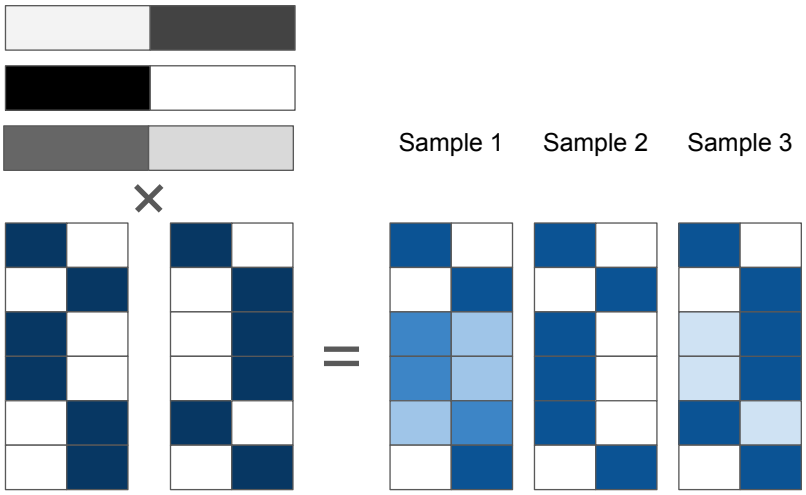

As the relative frequencies of strains vary across samples, the observations reflect this fact.

Strain haplotypes from metagenomes

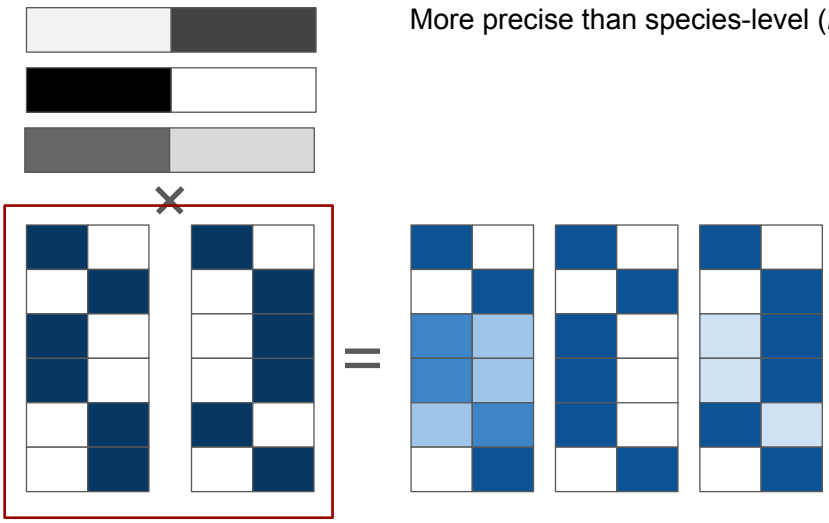

More precise than species-level (*much* better than 16S)

This model-based inference reveals the existence of strains defined at the level of haplotypes.
Since these variable positions are common in quickly evolving genes, the taxonomic resolution is much higher than with other methods.

Strain haplotypes from metagenomes

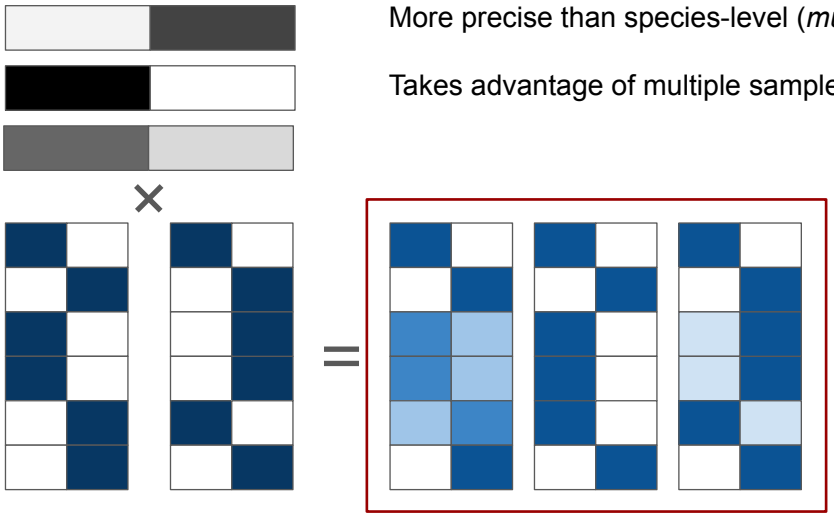

More precise than species-level (*much* better than 16S)

Takes advantage of multiple samples

The inference is particularly accurate in datasets for which we have many samples and
An expectation of shared strains across several of these samples.

Strain haplotypes from metagenomes

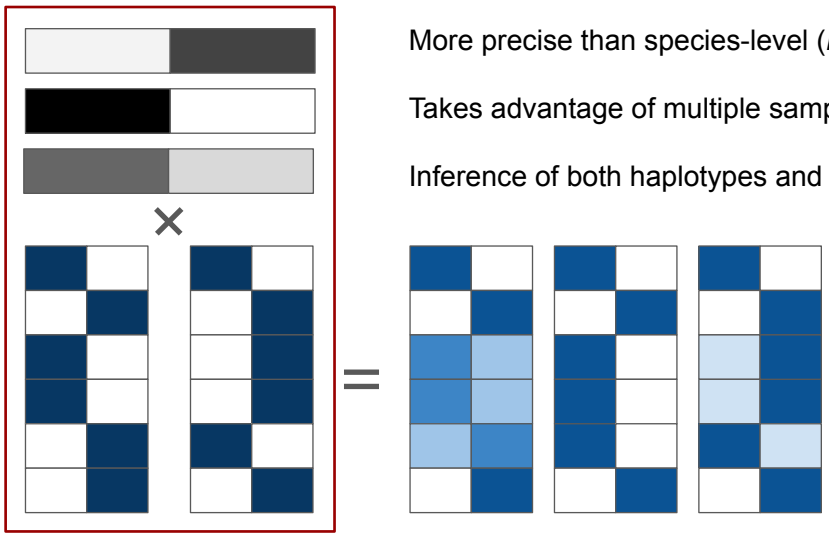

More precise than species-level (*much* better than 16S)

Takes advantage of multiple samples

Inference of both haplotypes and abundance

By fitting this probabilistic model on real data
We can infer both these haplotypes and their frequencies.

Strain haplotypes from metagenomes

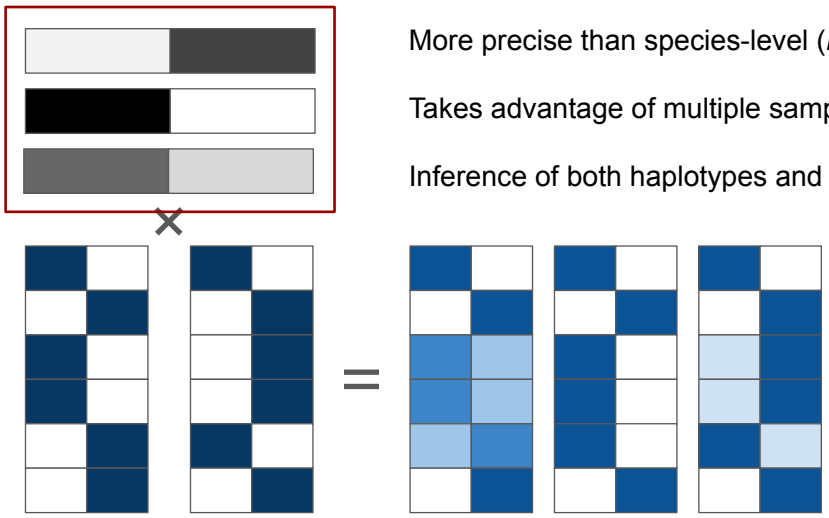

More precise than species-level (*much* better than 16S)

Takes advantage of multiple samples

Inference of both haplotypes and **abundance**

27

While the haplotypes themselves will be very exciting to study,
for the remainder of this talk, I'm going to focus on the inferred frequencies,

which give us a powerful view into engraftment and fluctuations at a much more precise taxonomic level.

Species tracking in longitudinal samples

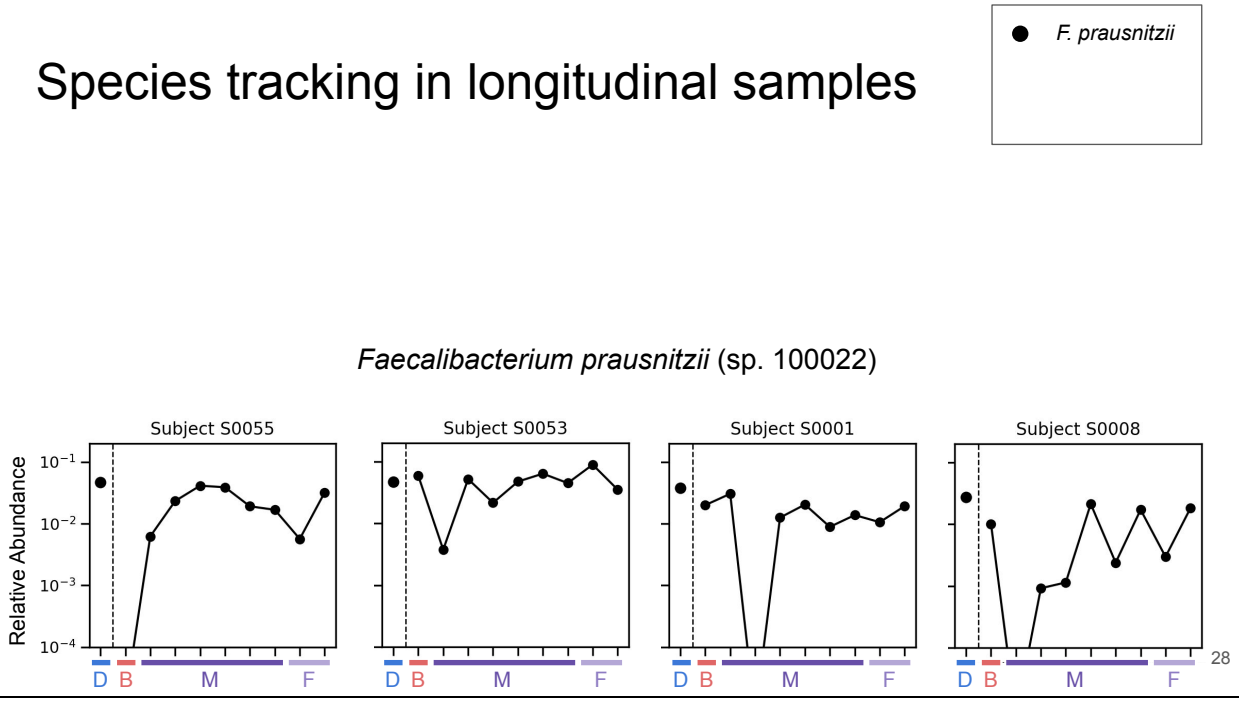

So, before I show these strain data,
Let's start by taking a look at the species level time-series.

Here I'm plotting the relative abundance of *F. prausnitzii* in four different study subjects.
I've indicated the samples along the X-axes,
B: for baseline, M: for samples taken prior to each of the six maintenance doses, and F: for the 2 week and 1 month followup samples.

D, separated from the other points by the dashed line,
indicates the relative abundance of F. prausnitzii in the donor,
which differs among these four subjects.

In these results, you can see that subject 55 (on the left) did not have Fp
In their baseline sample, but that it colonized during the FMT and was sustained
Into followups.

The other subjects all had Fp and, despite some fluctuation, this did
Not change throughout the time-series.

Strain tracking in longitudinal samples

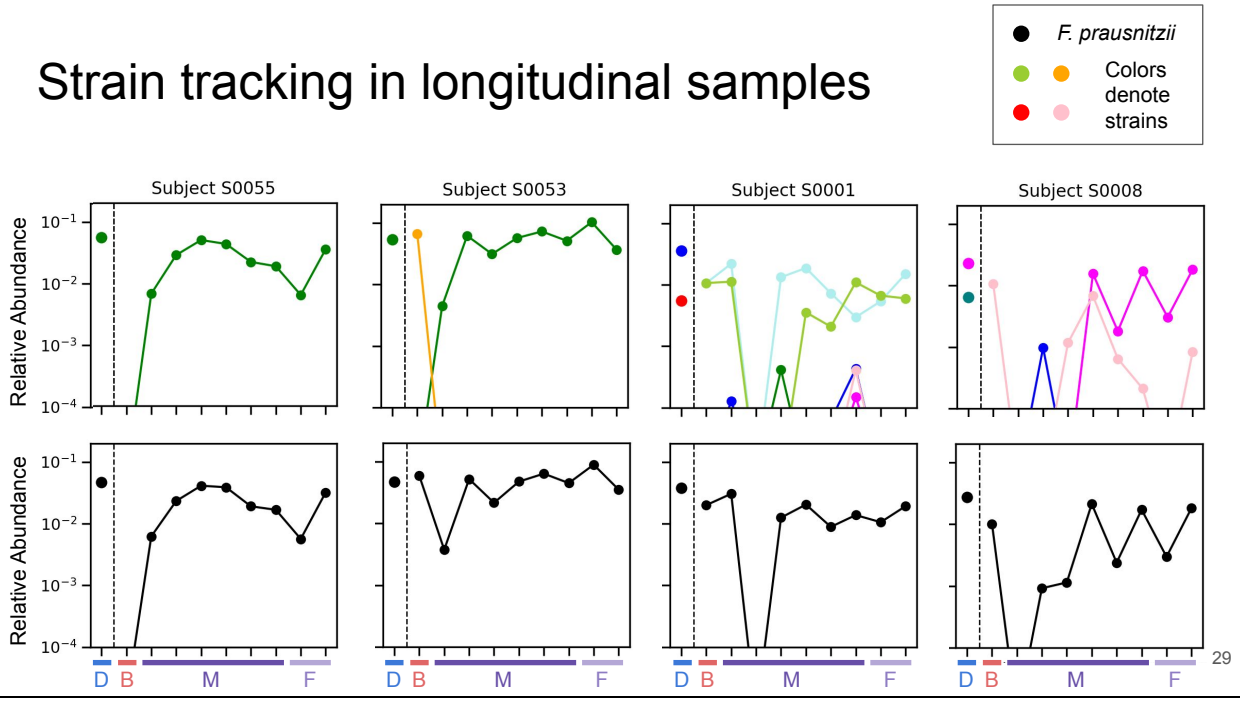

I think the strain results are much more exciting.
Here strains are differentiated by colors.

Now you can see a lot more, for instance:

- In subject 53 the donor strain has actually fully replaced the strain that was present at baseline,
- In subject 1, two different strains co-exist throughout the time-series

Strains reveal diverse outcomes of FMT

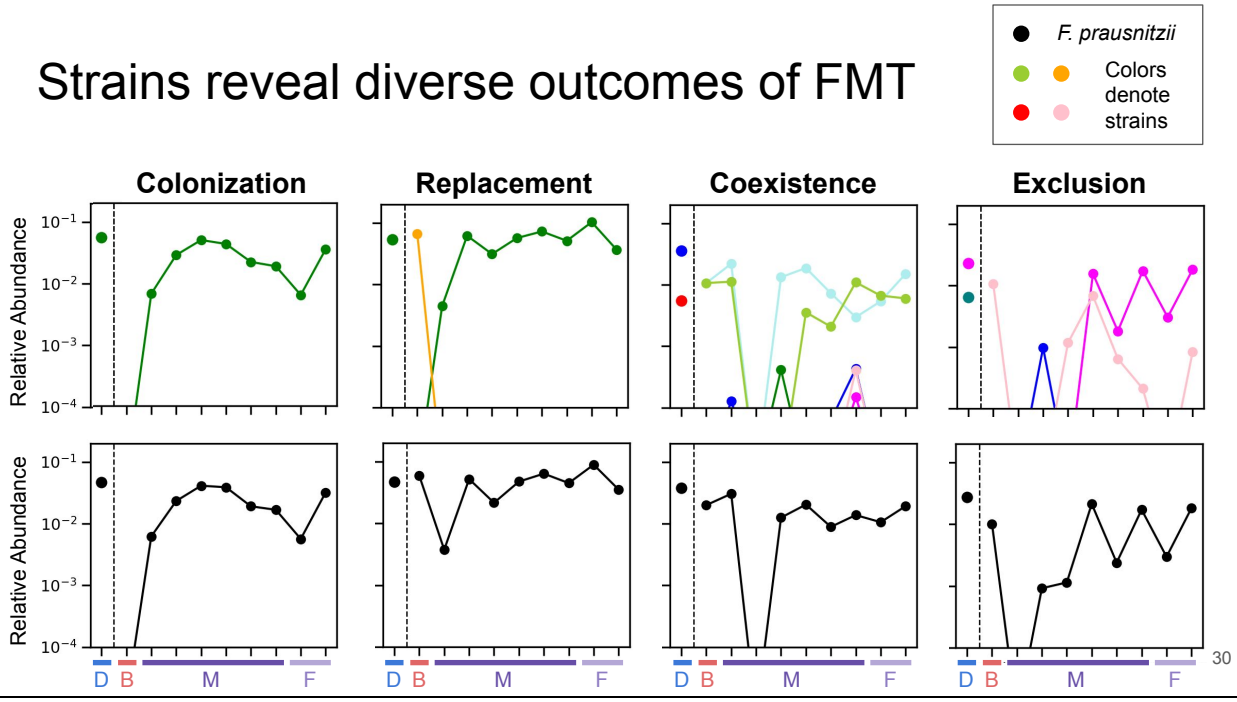

In just this one species and these four subjects we can see examples of colonization, replacement, coexistence, and exclusion between donor and recipient strains.

Dynamics that we were almost entirely blind to using standard, species-level methods.

FMT protocol influences microbiota transfer

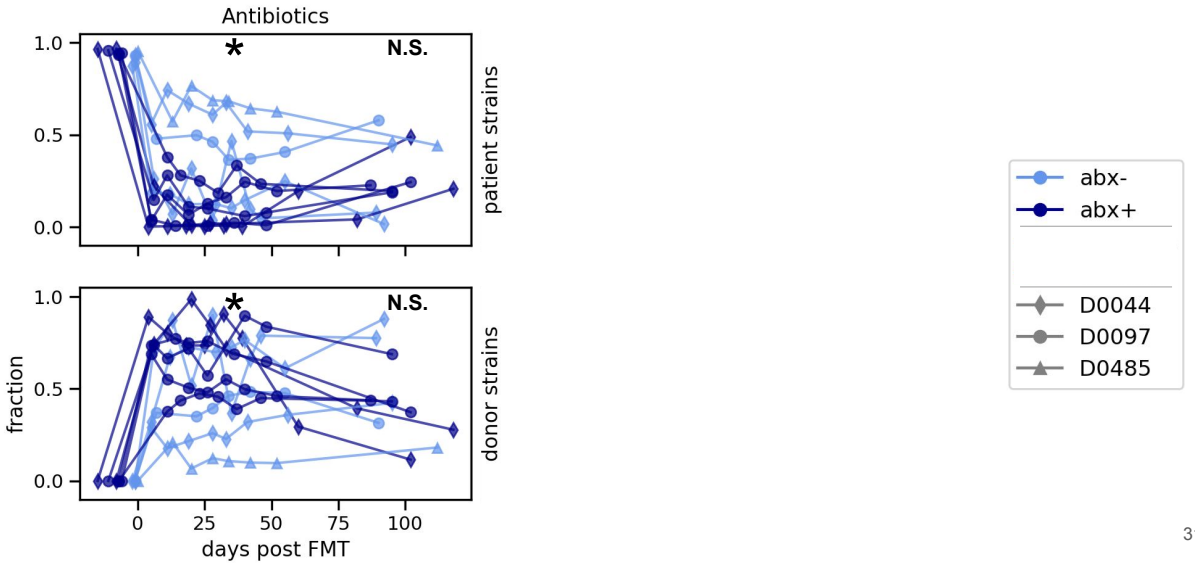

This enables several very interesting analyses.

Here lines in each panel represent different subjects.
I'm plotting the total relative abundance of strains in each subject
Partitioned into those strains that were:

- found in the patient's baseline sample, but not in the donor (top panel)
- Or found in the donor but not the patient's baseline (bottom panel)

Lines are colored by whether the subject received antibiotic pre-treatment before FMT.

You can see, and statistics confirm, that

- Patient communities shift substantially to resemble donor communities during and
  After FMT
- And antibiotic pre-treatment both decreased the abundance of patient strains
  And increased the number of donor strains.
  During maintenance dosing (but not during followup)

FMT protocol influences microbiota transfer

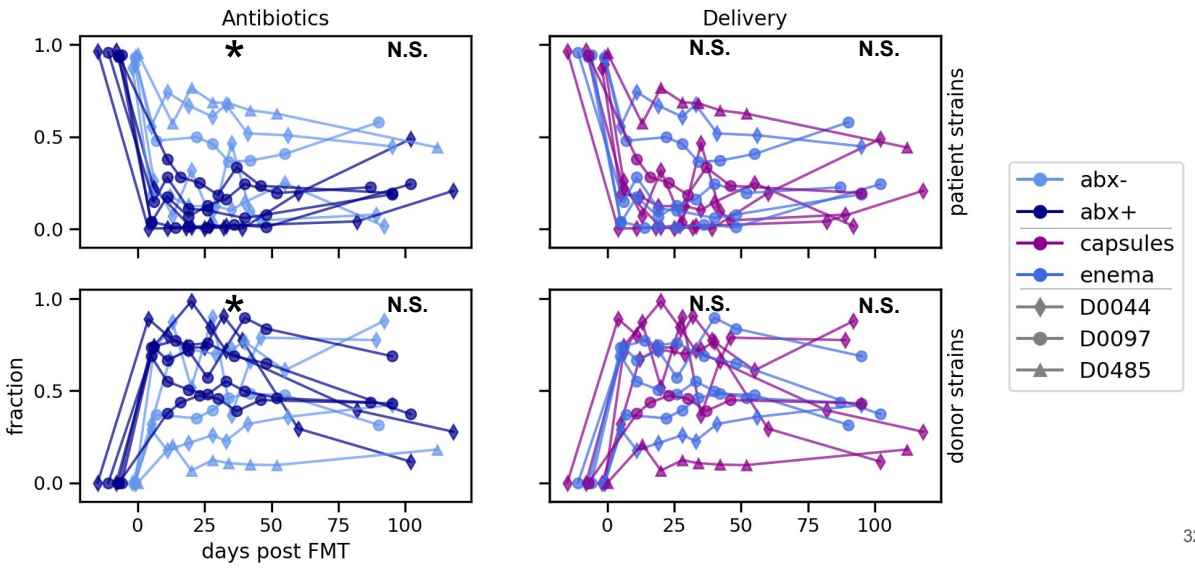

An analogous analysis stratifying by the delivery for the maintenance doses (capsules or enema)
Shows no effect during either maintenance or followup.

How do we link strains to functions?

Level of strain may be *too* precise

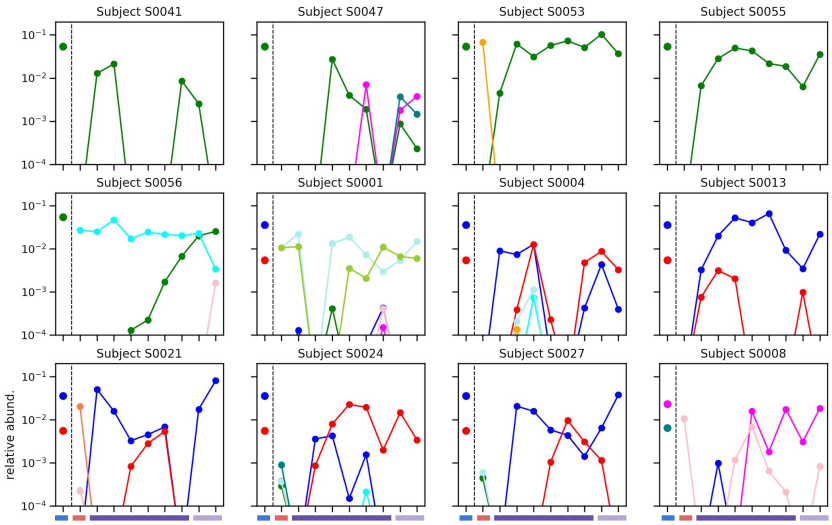

This kind of global engraftment analysis is a powerful way to compare FMT protocols
But it doesn't get us any closer to understanding the functional impacts of this engraftment.

In fact, looking at this highly specific strain-level may lose sight of the functions shared
By all members of a species.

We don't know if the green strain and the pink strain differ in their impacts on the host.

How do we link strains to functions?

Level of strain may be *too* precise

Gene content reconstruction, eventually

Phenotypic data can guide us: metabolomics

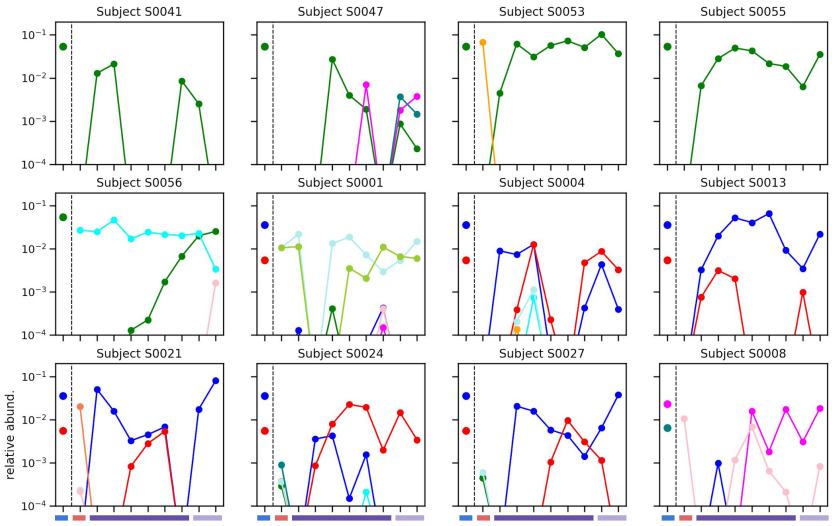

Gene content reconstruction is part of the answer to this dilemma,
But I'd like to suggest that integration of phenotypic data from another -Omics —metabolomics—can inform our search for differences that might drive Variability in FMT efficacy.

Bile acids

Produced in the liver and secreted (conjugated) into the **small intestine** to assist in the absorption of **hydrophobic compounds** (e.g. fats)

95% re-absorbed, but relevant amounts reach the colon and are transformed into a diverse set of **secondary bile acids**

Both metabolized by and affect the microbiota with evidence of **strain-specific** bacterial physiology

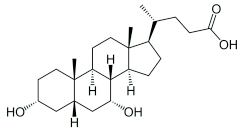

Cholate

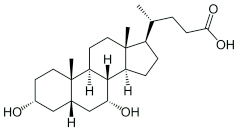

Chenodeoxycholate

Let's introduce the metabolites that I'm going to talk about today.

Bile acids are important components of the digestive system, contributing to the breakdown
and absorption of hydrophobic compounds (e.g. fats) in the small intestine.

Bile acids

Produced in the liver and secreted (conjugated) into the **small intestine** to assist in the absorption of **hydrophobic compounds** (e.g. fats)

95% re-absorbed, but relevant amounts reach the colon and are transformed into a diverse set of **secondary bile acids**

Both metabolized by and affect the microbiota with evidence of **strain-specific** bacterial physiology

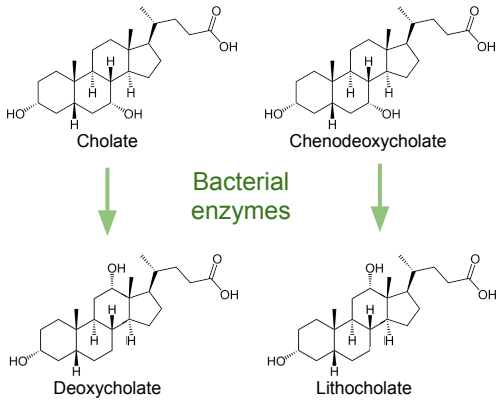

Most bile acids are reabsorbed in the distal small intestine, but a significant amount Enters the large intestine where it is acted upon by the microbiota to produce secondary bile acids.

Besides the two most abundant forms, deoxycholate and lithocholate,
Secondary bile acids include a diverse array of compounds
Which can have differential effects on both host and microbial physiology.

Strain-level variability in this metabolism has been documented.

Bile acids

Produced in the liver and secreted (conjugated) into the **small intestine** to assist in the absorption of **hydrophobic compounds** (e.g. fats)

95% re-absorbed, but relevant amounts reach the colon and are transformed into a diverse set of **secondary bile acids**

Both metabolized by and affect the microbiota with evidence of **strain-specific** bacterial physiology

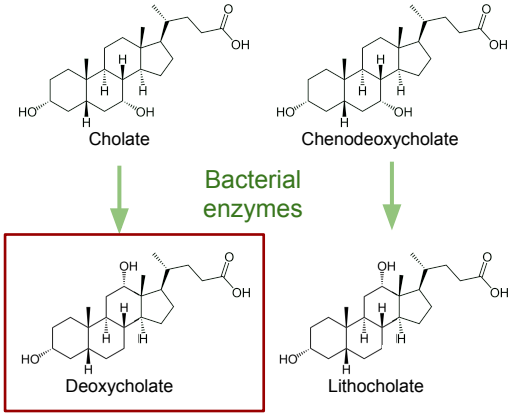

I'm going to focus on just one of these compounds, deoxycholate, which is Usually abundant, but highly variable in stool.

Deoxycholate may be associated with UC

Known to induce inflammation and experimentally linked to several cancers

Powerful agonist of several receptors with downstream effects on metabolism and immunomodulation: TGR5 and FXRα

Along with lithocholate, decreased in IBD associated dysbiosis

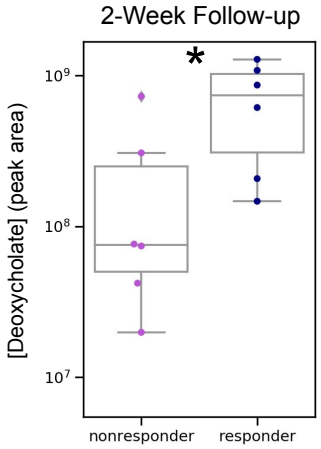

I'm choosing to focus on deoxycholate because it has been extensively linked to gut health, with
Connections to increased inflammation and experimental evidence that it can cause cancer.

Several affectors of deoxycholate are known, including the bile acid receptors TGR5 and FXR-alpha.

Perhaps contradictorily, previous studies have demonstrated that deoxycholate concentrations are decreased in patients with
IBD, perhaps due to decreased abundance of the relevant bacterial taxa.

Data from our study reproduces that pattern when comparing deoxycholate concentrations in stool at the two week followup in subject that were
Found to have responded to FMT treatment.

Predictive model of deoxycholate (MWAS)

$$[\text{Deoxycholate}] \sim \text{SubjectEffects} + \text{SpeciesEffects} + \text{StrainEffects}$$

Given these prior results, I wanted to ask whether strain-level resolution could improve our ability to predict deoxycholate concentrations in feces.

To answer that question, I fit probabilistic model with parameters for subject effects, as well as effects at both the species level and the strain level.

By analogy to similar models used for genome-wide association studies, I might call this a microbiome-wide association study (MWAS).

Predictive model of deoxycholate (MWAS)

$$[\text{Deoxycholate}] \sim \text{SubjectEffects} + \text{SpeciesEffects} + \text{StrainEffects}$$

However:

It is unclear if this prediction is accurate out-of-sample

Large number of predictors relative to samples

Coefficients do not have the same causal interpretation as in GWAS

It is worth pointing out some caveats, however.

First, we have a relative small sample size and a relatively large number of parameters, so
over-fitting is a real problem and
It is worth being skeptical of out-of-sample relevance.

At the same time, because the microbial community also responds to bile acid concentrations,
we cannot say that the predictors precede the response, so we lose the causal interpretation that is
sometimes enjoyed by GWAS.

Model comparison

Accurate predictions of deoxycholate concentration

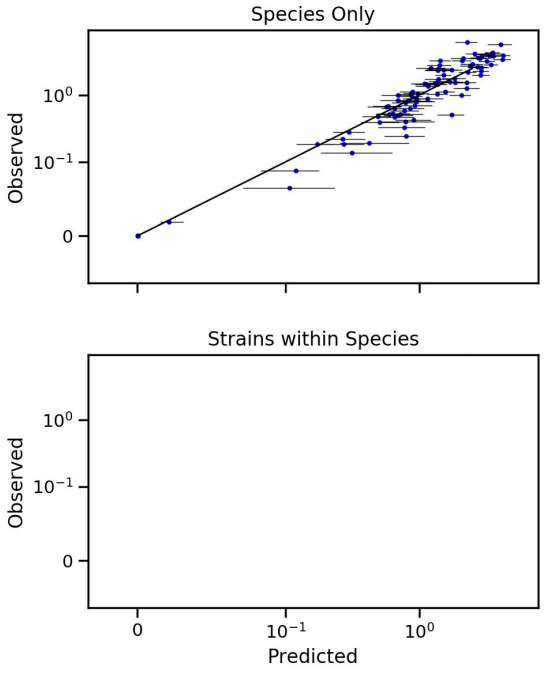

Nonetheless, our model does a very good job of predicting deoxycholate concentrations.
Here I'm showing the observed deoxycholate concentrations plotted against concentrations predicted when using a reduced model
that does not include strain effects.

While the fit is pretty striking...

Model comparison

Accurate predictions of deoxycholate concentration

WAIC indicates superior performance of strains model (dWAIC >> 0),

Suggests that strain-level variation is important

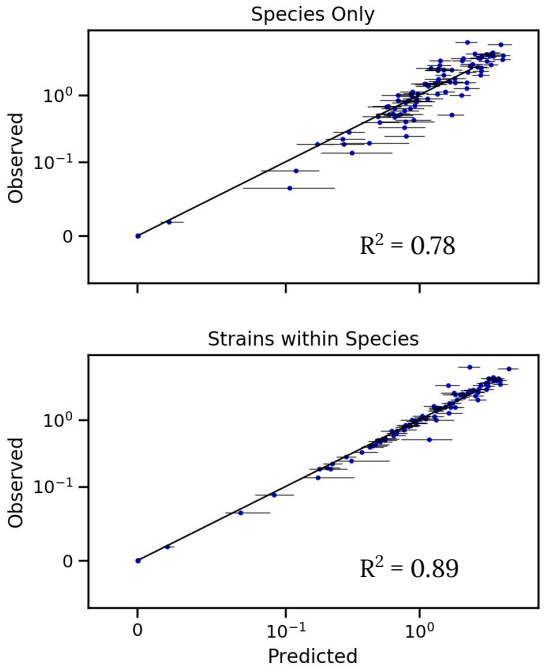

It is not nearly as good as a model that also includes strain-level effects.
The r-squared increases substantially,
and comparing the widely-applicable information criterion lends strong support to the inference that
Strain-level effects are important for predicting deoxycholate concentrations.

Strains and species associated with deoxycholate

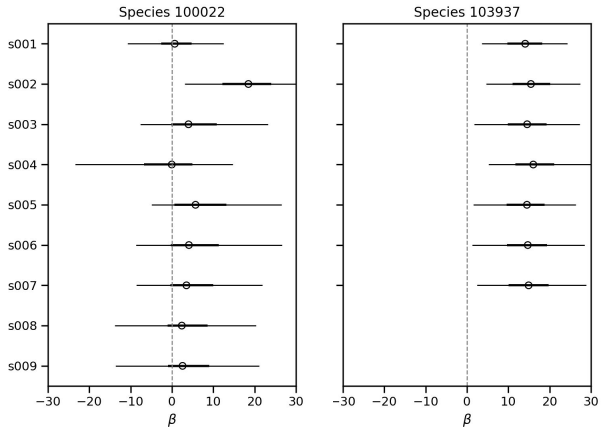

Fitted parameters indicate particular taxa associated with deoxycholate, evidence of strain-specific associations.

Suggest the presence of relevant metabolic pathways.

Despite the limitations, we can go beyond prediction and ask about the fitted parameters themselves,
Using these values to make inferences about which taxa are associated with DCA concentrations.

Here I'm showing parameter estimates for strains from two different species.
You can see that Species 103937 is positively associated with DCA irrespective Of strain,
While for 100022 just one strain is significantly linked.

An associated *F. prausnitzii* strain is found in donor D0044

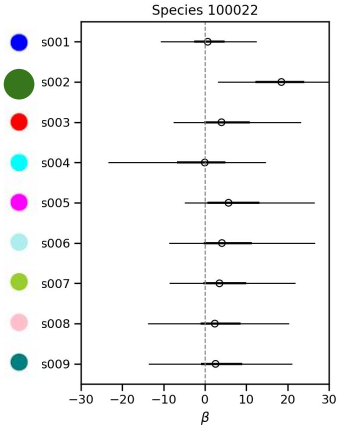

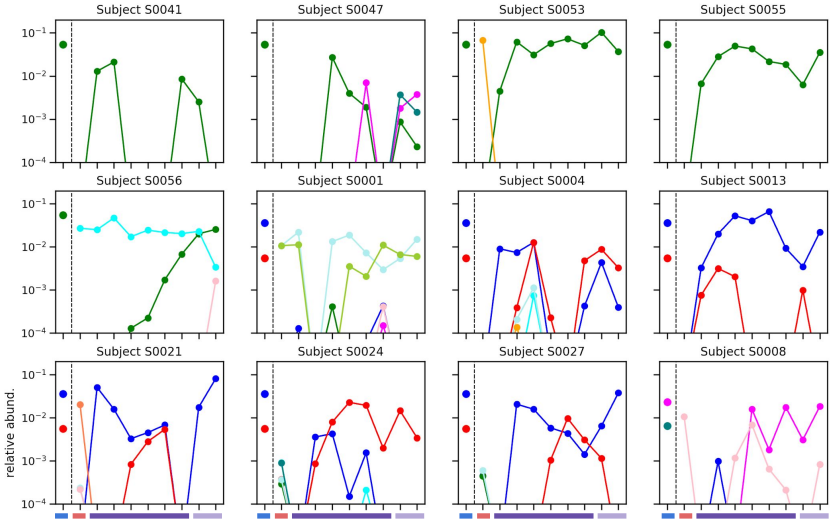

This strain-specific association is particularly interesting given because Strain 2 is Found in only Donor 44...

An associated *F. prausnitzii* strain is found in donor D0044

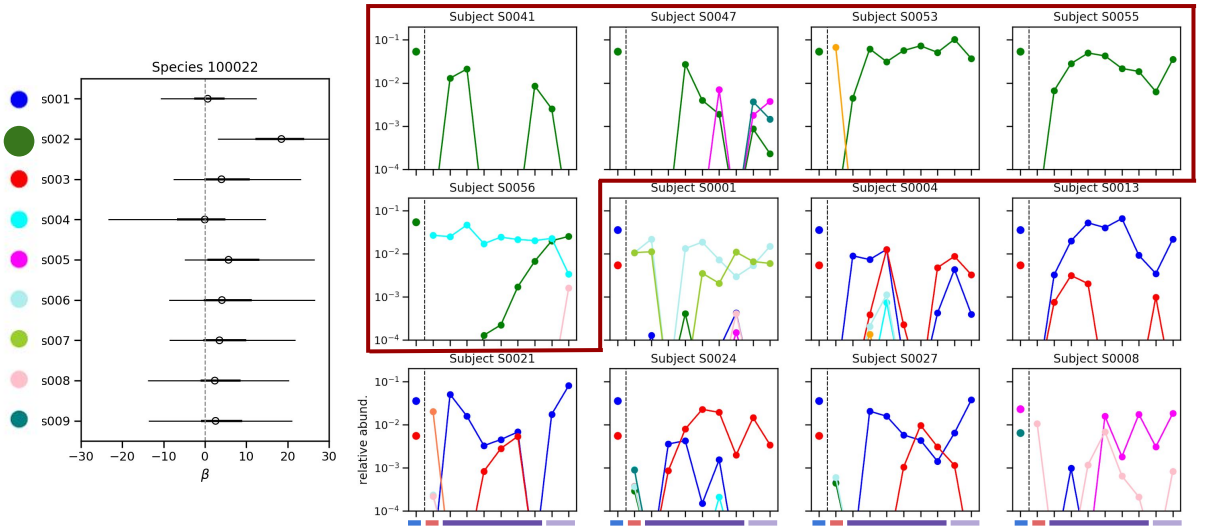

...and is often transferred to the patients receiving material from this donor.

In the few moments I have left, I'd like to describe some very preliminary results from An attempt to reconstruct genome content for this strain in particular, and compare It's functional potential to other strains.

Strain-informed genome reconstruction

Briefly:

Identify genes with correlated coverage in samples with *only* the focal strain.

Very briefly, we reconstructed genome content by finding genes whose metagenomic read coverage
Was correlated with species abundance in samples that ONLY have the focal strain.

bai genes encode 7α-dehydroxylation pathway

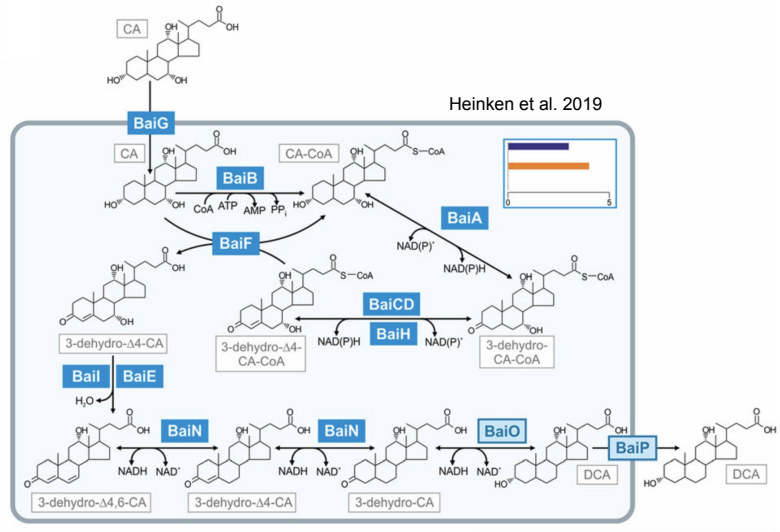

I then looked specifically for the presence of homologues for genes in the *bai* operon Which is responsible for the 7alpha-dehydroxylation of cholate to deoxycholate.

Preliminary genome reconstructions suggest deoxycholate associated genes

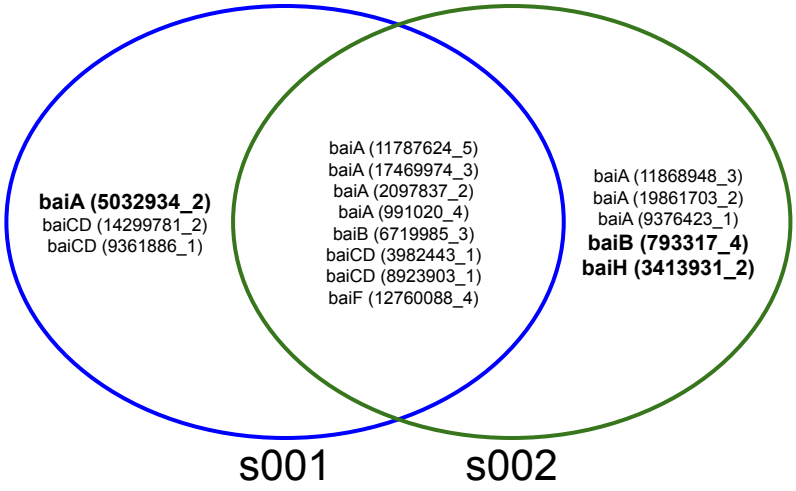

What we find is that there are indeed genes that differentiate Strain2 from Strain1, Including homologues of BaiB and BaiH

Integration of metagenomics and metabolomics

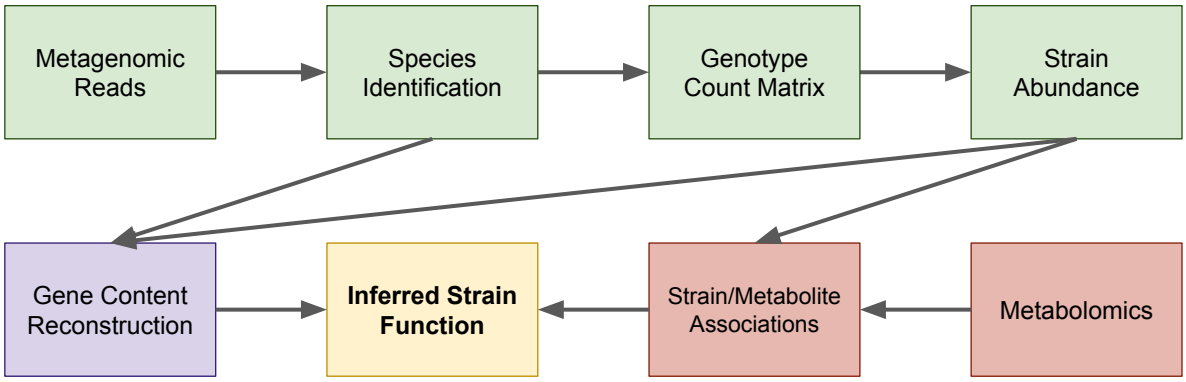

To quickly summarize the diverse analyses that I've covered today,
Here I'm showing diagramatically how we've integrated metagenomic data
Used for strain-tracking (in green) with metabolomic data to find
phenotypic associations (in red), and then cross-reference these with gene content reconstruction
To infer differences in strain function.

Next step

Culturing from feces with the goal of **isolating strains**:
validation of haplotypes, gene content, and phenotypic inferences

Extend MWAS-informed **comparative genomics** to more metabolites

Clearly, an important next step is to isolate strains from these individuals to test the hypotheses generated here.

What's more, I've only described one metabolite here. There are plenty more that I'd like to dig into.

Conclusion

Strain reconstruction may reveal clinically relevant differences in bacterial gene content among strains with implications for the effectiveness of FMT from different donors

In conclusion:

“Strain reconstruction may reveal clinically relevant differences in bacterial gene content among strains with implications for the effectiveness of FMT from different donors”

Effect of FMT protocol on microbiome transfer

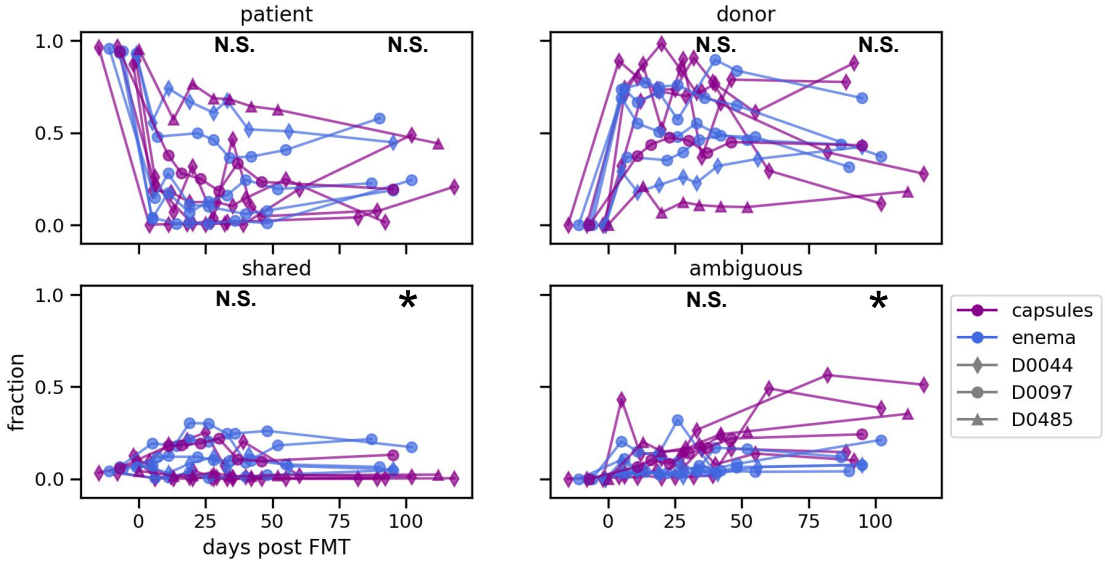

TODO: Consider dropping this slide entirely
TODO: Also consider responder status
TODO: Consider dropping capsule/enema distinction?

Changes (glaringly obvious and/or tautological) were sustained during followup (no significant difference).
P-value for the MWU test on differences between abx+ and abx- patients in relative abundance of **patient** taxa during maintenance was p=0.018
P-value for the MWU test on differences between abx+ and abx- patients in relative abundance of **donor** taxa during maintenance was p=0.085
Also, the BC distance from baseline was higher during maintenance in patients treated with abx. (and—borderline p-value—distance from donor was reduced during maintenance).

FMT protocol influences engraftment of microbiota

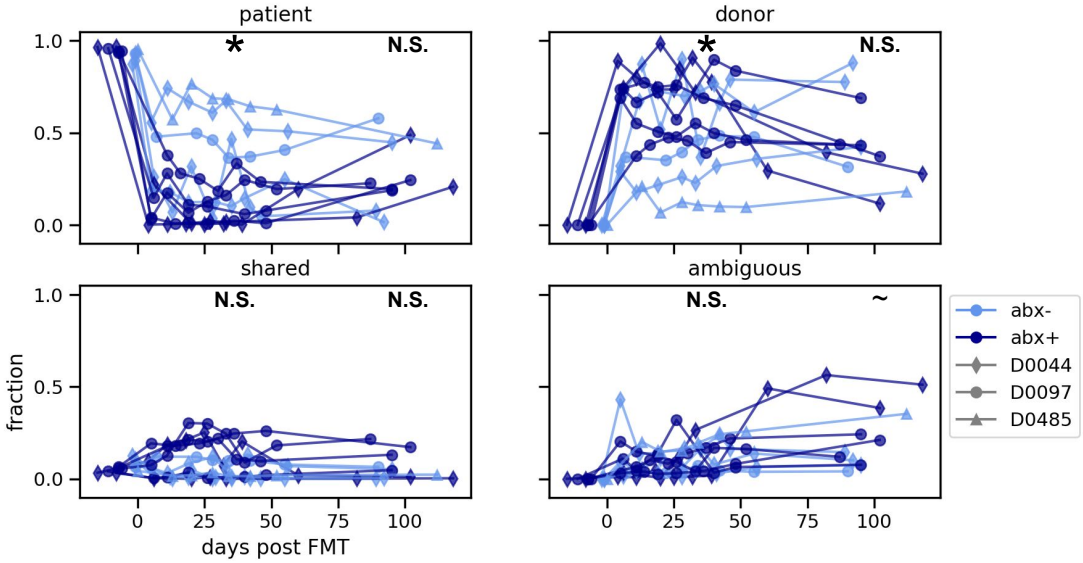

TODO: Add slides for responder status

Changes (glaringly obvious and/or tautological) were sustained during followup (no significant difference).
P-value for the MWU test on differences between abx+ and abx- patients in relative abundance of **patient** taxa during maintenance was p=0.018
P-value for the MWU test on differences between abx+ and abx- patients in relative abundance of **donor** taxa during maintenance was p=0.085
Also, the BC distance from baseline was higher during maintenance in patients treated with abx. (and—borderline p-value—distance from donor was reduced during maintenance).

TODO: Fraction of what
TODO: Add space between titles and x-axes.
TODO: COnsider dropping shared/ambig.

Overall engraftment not associated with response

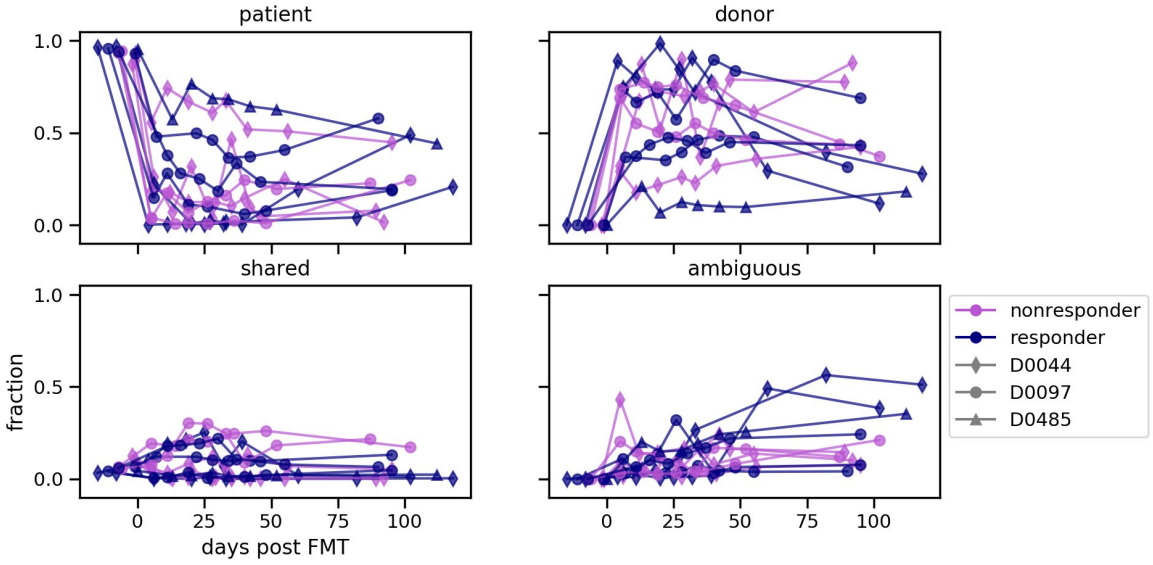

TODO: Add p-values

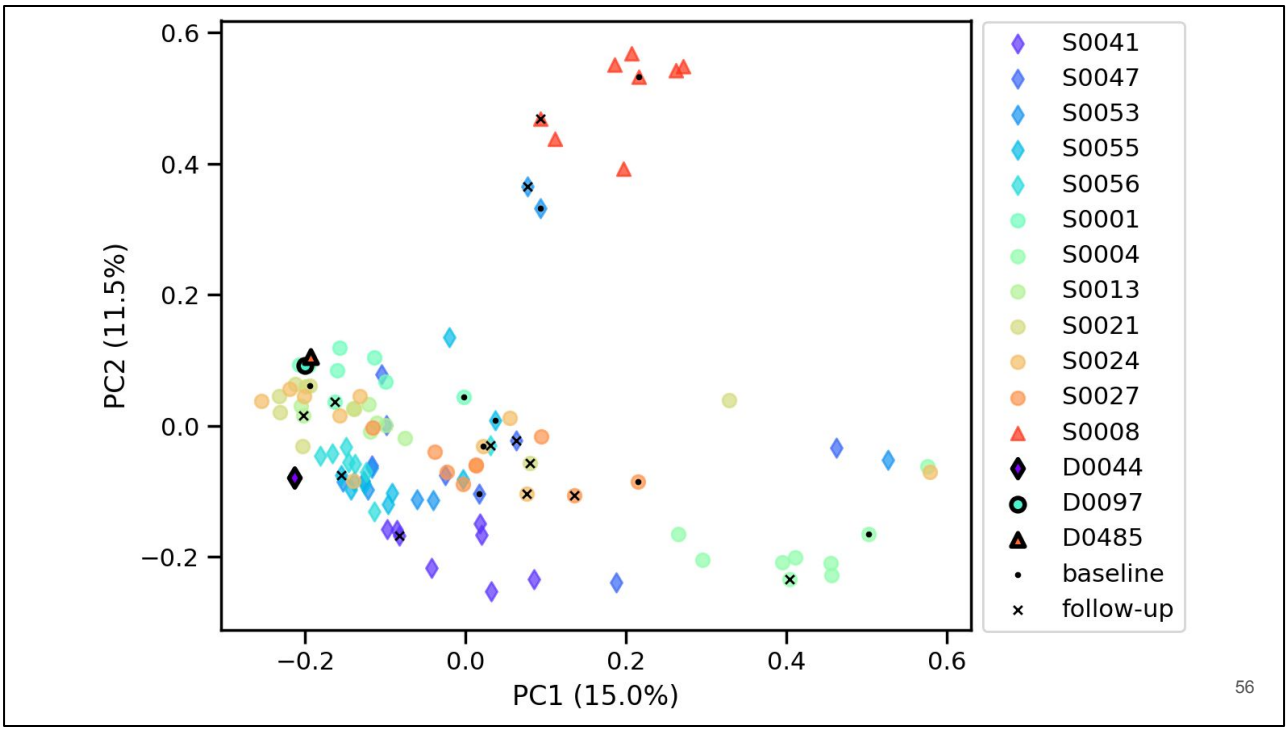

TODO: Write speaker notes for detailed description of this plot covering all dimensions that I'm trying to communicate about (especially shapes)
TODO: Does this slide and the following need a title/punchline?
TODO: Make pocket slides with responder status and other metadata annotating these ordinations
Ordination based on 16S does not greatly distinguish between three donors, and there's substantial overlap among subjects.

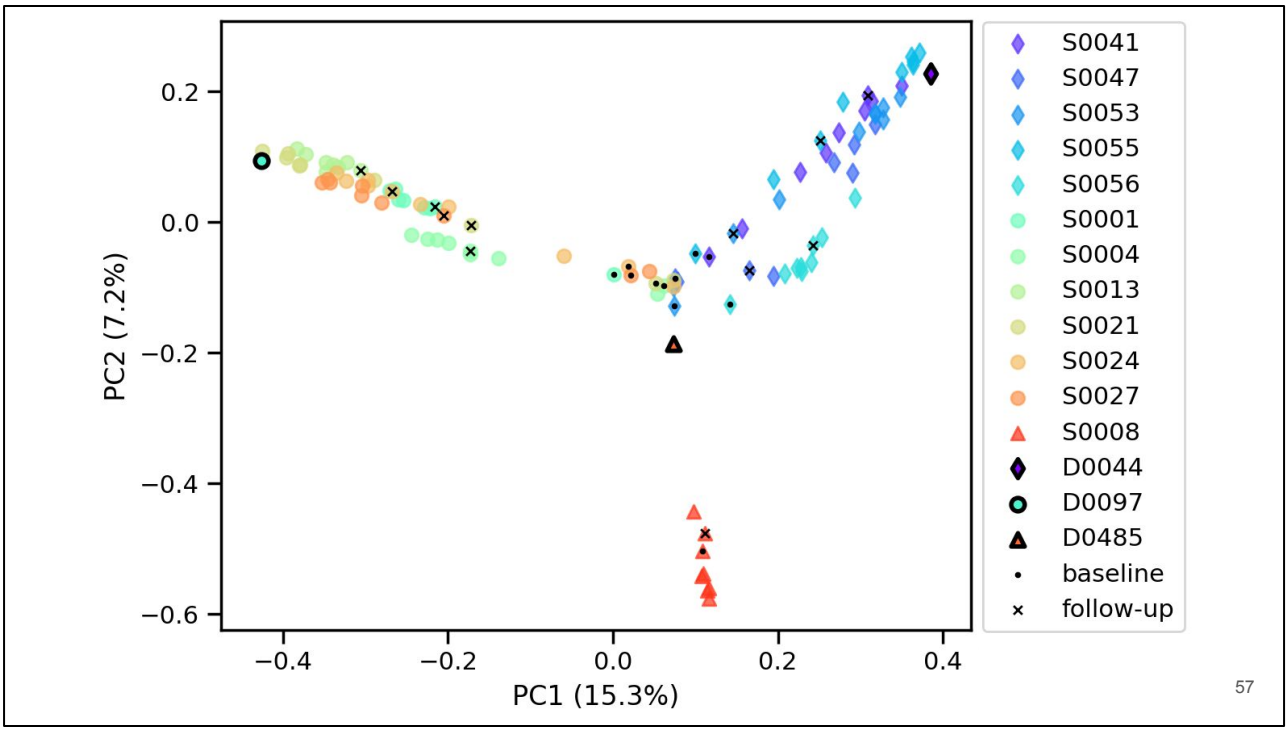

TODO: Consider dropping these two PCoA's entirely.
Strain-level has greater sensitivity for engraftment
Similarity among patients receiving the same donor entirely drives the first two principal coordinates

Can calculate the deoxycholate “potential” of donors

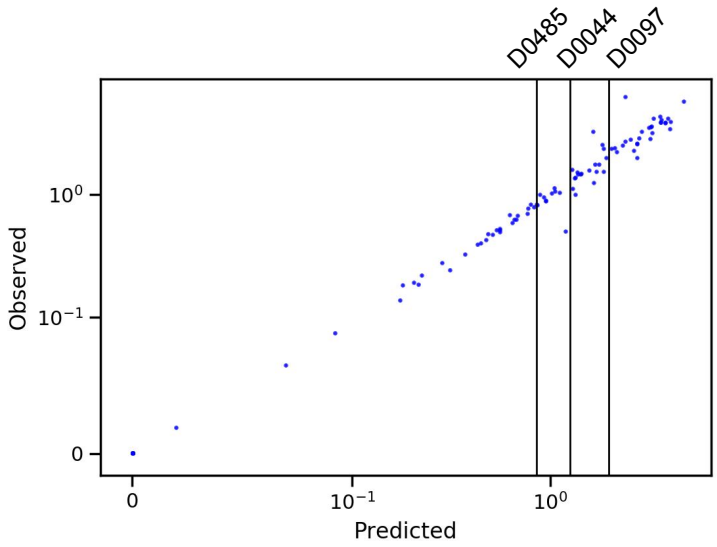

TODO: Log-log