


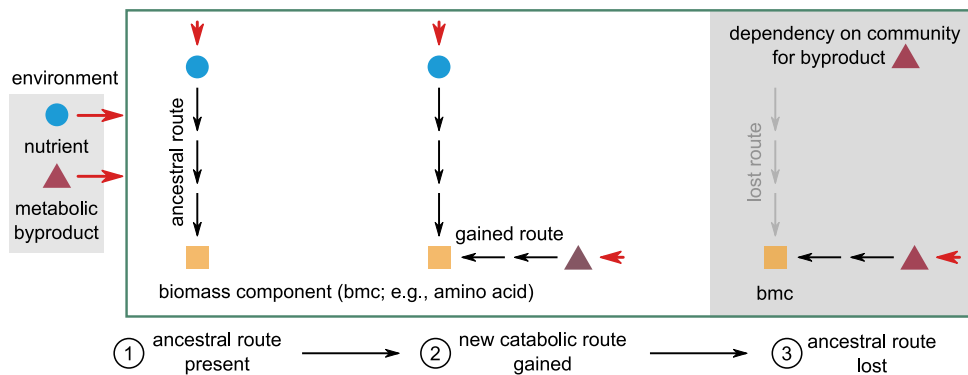
Article

# Horizontal gene transfer drives the evolution of dependencies in bacteria

We reconstructed the history of gain and loss of bacterial metabolic networks



Coupled gains and losses (CGLs) are a new mechanism for dependency evolution



Akshith Goyal

akshitg@mit.edu

Highlights

Metabolic dependencies are widespread across bacterial genomes

New genes expand bacterial catabolism via the process of horizontal gene transfer

During evolution, efficient pathways are gained, whereas redundant pathways are lost

Gained pathways often depend on the metabolic byproducts of the surrounding community



## Article

## Horizontal gene transfer drives the evolution of dependencies in bacteria

Akshit Goyal<sup>1,2,\*</sup>

## SUMMARY

Many naturally occurring bacteria lead a lifestyle of metabolic dependency for crucial resources. We do not understand what factors drive bacteria toward this lifestyle and how. Here, we systematically show the crucial role of horizontal gene transfer (HGT) in dependency evolution in bacteria. Across 835 bacterial species, we map gene gain-loss dynamics on a deep evolutionary tree and assess the impact of HGT and gene loss on metabolic networks. Our analyses suggest that HGT-enabled gene gains can affect which genes are later lost. HGT typically adds new catabolic routes to bacterial metabolic networks, leading to new metabolic interactions between bacteria. We also find that gaining new routes can promote the loss of ancestral routes (“coupled gains and losses”, CGLs). Phylogenetic patterns indicate that both dependencies—mediated by CGLs and those purely by gene loss—are equally likely. Our results highlight HGT as an important driver of metabolic dependency evolution in bacteria.

## INTRODUCTION

Naturally occurring bacteria lead one of two metabolic lifestyles: autonomy or dependency (Ochman and Moran, 2001; McCutcheon and Moran, 2007, 2012; Luo et al., 2013). Although autonomy reflects complete self-sufficiency in converting nutrients to biomass, bacteria with dependencies often require crucial metabolites from others (D’Onofrio et al., 2010; Davis, 1921; Morris et al., 2008; Suzuki et al., 1988; Watsuji et al., 2007). These metabolites are secreted by neighboring community members. Because such dependencies are common in bacterial communities, it is instructive to ask: what processes and factors affect their evolution; in other words, what drives the switch from metabolic autonomy to dependency?

To answer this question, recent studies have put forth the Black Queen hypothesis, which states that dependencies evolve through adaptive gene loss (Morris et al., 2012; Mas et al., 2016; Fullmer et al., 2015). Individuals lose costly, dispensable genes in leaky environments, trading autonomy for better growth (or fitness). As both experiments and models show, this is feasible—administering the loss of even a few specific biosynthetic genes in bacteria repeatedly leads to strong metabolic dependencies (Pande and Kost, 2017; Pande et al., 2014; D’Souza and Kost, 2016; Goyal, 2018). This also explains how endosymbionts undergo severe genome reduction (McCutcheon and Moran, 2012). These bacteria lack many biosynthetic pathways and instead depend on their hosts for the required biomass components.

However, many extant free-living bacterial species are also metabolically dependent, despite the “free-living” label (D’Souza et al., 2014; Monk et al., 2013; Goyal and Maslov, 2018; Wang et al., 2019; Enke et al., 2019). These species do not have merely reduced genomes, i.e., they do not differ from their ancestors only by gene losses, as expected under the Black Queen hypothesis (Zelezniak et al., 2015; Garcia et al., 2013; Giovannoni et al., 2014). Over time, they have also gained many genes, primarily by horizontal gene transfer (HGT) and at times by *de novo* gene birth (Pál et al., 2005; Vos et al., 2015; Szappanos et al., 2016; Press et al., 2016; Maslov et al., 2009). For these often-dependent bacteria, we ask: could gene gains also have contributed to which dependencies we observe today? Specifically, during dependency evolution, which genes will be gained, and do those gains affect and which genes will later be lost?

Here we explore the role of horizontal gene transfer in the evolution of metabolic dependencies in bacteria. Specifically, we measure the potential of HGT to drive dependency evolution by affecting the likelihood of subsequent gene loss events. To do this, we reconstructed the evolutionary history of 835 phylogenetically diverse, nonendosymbiont bacterial species. By inferring ancestral gene content, we mapped gene gains

<sup>1</sup>Physics of Living Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

<sup>2</sup>Lead contact

\*Correspondence: akshitg@mit.edu

<https://doi.org/10.1016/j.isci.2022.104312>



and losses along a large, deep-branching phylogeny and assessed their impact on bacterial metabolic capabilities. Our analyses suggest that horizontally transferred genes can indeed affect which genes are later lost and which dependencies emerge as a result. We have two lines of evidence to support this. First, we find that gene gains add new catabolic routes to bacterial metabolic networks. These gained catabolic routes increase the chance of new metabolic interactions between bacteria, a prerequisite for dependency evolution. Next, we show how these new routes can promote the loss of preexisting ancestral routes (a process we call “coupled gains and losses,” CGLs). We find that phylogenetic patterns indicate that both processes—CGLs and pure gene loss—are equally likely to lead to dependencies. Collectively, these results highlight horizontal gene transfer as an important driver of metabolic dependency evolution in bacteria.

## RESULTS

### Horizontal gene transfer (HGT) adds new catabolic routes to bacteria

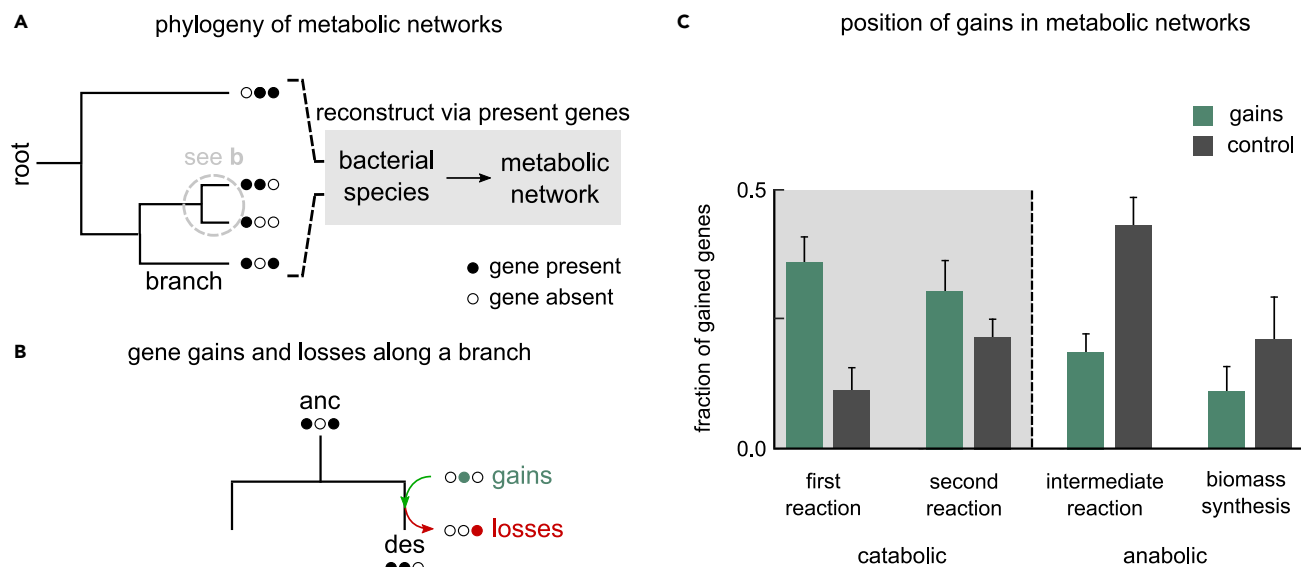
In a metabolic dependency, a donor organism secretes metabolites, which are in turn required by an acceptor organism. The secreted metabolites cannot be produced by the acceptor organism itself but are still necessary for survival and growth. We sought to understand how horizontal gene transfer, if at all, impacts the emergence of new dependencies. We hypothesized that newly acquired genes (primarily through HGT) lead to newer metabolic interactions. This could occur if gained genes allowed an acceptor organism to transport and break down previously unusable metabolites in its surroundings. We thus first asked: does horizontal gene transfer enhance the ability of a bacterial genome to utilize metabolites secreted by surrounding donors? In other words, does it add new catabolic routes? Analysis from *E. coli* suggests that this might indeed be the case, but the lack of systematic analyses across several bacterial species hinders us from generalizing these results (Pál et al., 2005).

To answer whether HGT typically adds new catabolic routes to bacteria, we used a two-pronged approach, combining bacterial metabolism and phylogeny. We first acquired a list of 1,031 bacterial species with complete genomes, whose metabolic data were available in the KEGG GENOME database. We explicitly removed from this list: (1) endosymbionts, due to their exotic metabolic lifestyles and genomes, and (2) closely related genomes, to avoid phylogenetic bias (see STAR Methods). This left us with the 835 species genomes we used for all our subsequent analyses (Table S1). For each genome, we extracted all metabolic genes present in at least one species, corresponding to a total of 3,022 unique genes.

Using these genomes, we inferred each species’ metabolic capabilities. For this, we reconstructed representative metabolic networks, one for each species, using gene presence-absence data (Figure 1A). Here, we mapped each gene to specific chemical reactions using the KEGG REACTION database (see STAR Methods). To identify gene gain and loss events during the evolution of these species, we inferred the most likely genetic makeup of their ancestors. For this, we first established evolutionary relationships using a well-known bacterial phylogenetic tree, and then applied ancestral reconstruction methods, to infer which of the 3,022 metabolic genes were likely present in each ancestor (see STAR Methods). With the presence-absence profiles of ancestors and descendants on each phylogenetic branch, we could infer which genes were gained and lost along them (Figure 1B).

We first tested whether HGT can expand the set of externally available metabolites that a metabolic network can catabolize. For this, we studied which position horizontally transferred genes typically occupied in each metabolic network they were gained in. Specifically, we were interested in whether the gained genes were in catabolic or anabolic parts of a metabolic network. We studied this across all phylogenetic branches. On each branch, we asked which position in the descendant’s metabolic network each gained gene occupied. Each position corresponded to metabolic reaction order, from catabolic to anabolic, as follows: first, second, intermediate, or biomass-synthesis reactions (see STAR Methods). If HGT was indeed likely to add new catabolic routes, we would expect gained genes to be concentrated in the catabolic parts of the network, i.e., first and second reactions. As controls, we measured the positions of randomly chosen genes in the same metabolic networks.

Along each branch, we measured the fraction of gained genes corresponding to each network position. We then calculated, across all branches, the average fraction of gains found to occupy each position. Consistent with our hypothesis, we found that horizontally acquired genes are more likely to be part of catabolic routes (mean; 69% catabolic versus 31% anabolic; Figure 1C, green bars); this is much more than



**Figure 1. Horizontal gene transfer adds new catabolic routes to bacteria**

(A) Schematic representation of our two-pronged approach: of combining phylogeny and bacterial metabolism. We used a well-known phylogenetic tree to infer the evolutionary relationships between the 835 bacterial species used in our analysis. For each extant species (shown on the tips of the tree), we used gene presence-absence data for 3,022 metabolic genes from the KEGG metabolic database. Filled circles indicate gene presence; empty indicate absence. (B) We inferred the gene presence-absence states of all internal nodes of the tree along each branch of the full tree (gray dashed circle in a); each branch connected an ancestor (anc) to a descendant (des). Along each branch, we inferred which genes were gained (green) and lost (red). (C) Bar plot showing the position of gained genes in bacterial metabolic networks. We split metabolic genes into catabolic (first or second reactions in a metabolic route) and anabolic (intermediate or biomass-synthesis reactions) based on the the chemical reactions they map to. Each green bar represents the average fraction of gained (horizontally transferred) genes at that position. Each black bar represents controls, i.e., the expected average fraction of gains at that position, given a random set of gene gains. Error bars show the SE, indicating the extent of variation across 1,669 phylogenetic branches.

expected by chance (in controls, we found mean 34% catabolic versus 66% anabolic; Figure 1B, black bars). This suggests that HGT can expand the number of external metabolites that bacteria can catabolize.

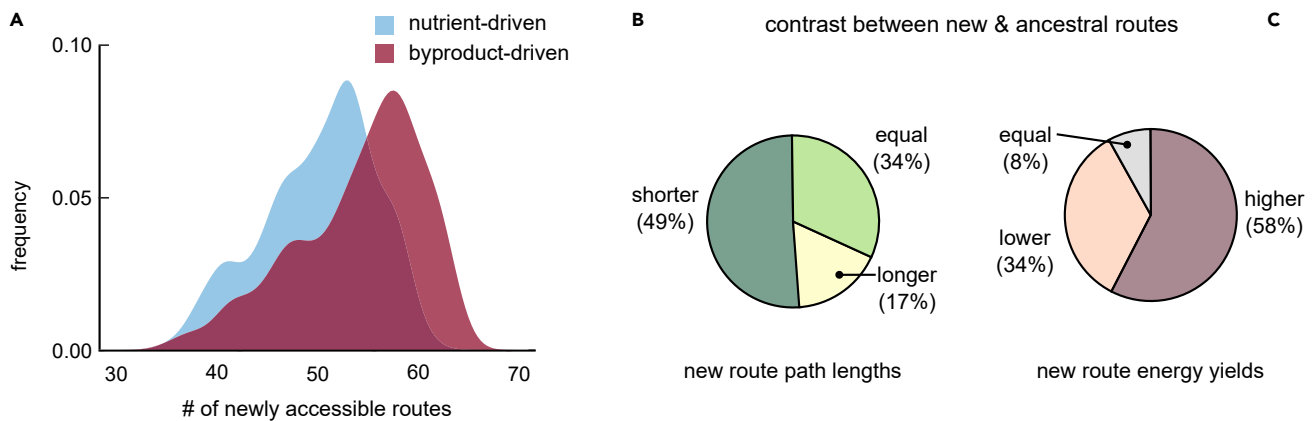
### HGT-enabled catabolic routes increase the likelihood of metabolic interactions

We next tested the possibility that newly acquired catabolic routes promote new metabolic interactions. For this, new routes should help break down the metabolic byproducts of other bacteria more often than nutrients available in the environment. To test this, we first curated a list of common external metabolites and classified them as byproducts or nutrients based on their presence on the exterior or interior of microbial metabolic networks (see Table S2 for the full list; see STAR Methods for a detailed procedure). In a metabolic interaction, the received metabolite should eventually help produce biomass components (such as pyruvate, ribose-5-phosphate, and alanine). For each ancestor-descendant pair, we analyzed how many such biosynthetic routes were added to the descendant's metabolic network, when compared with its ancestral network (see STAR Methods). We separately counted routes using byproducts as their starting point, from routes using nutrients.

We found that, on average, new biosynthetic routes, enabled by HGT, are more likely to be byproduct-driven than nutrient-driven (median number of routes, 56 and 51, respectively;  $P < 10^{-3}$ ; distributions compared via a Kolmogorov-Smirnov test; see Figure 2A). Moreover, these new routes could often be metabolically more efficient than their ancestral counterparts (see Figure 1D), i.e., they often had shorter path lengths (49%; Figure 2B, left) and yielded more energy (ATP; 58%; Figure 2C, right) than the corresponding routes in their ancestors (see STAR Methods). This suggests that newly acquired routes can indeed enable new metabolic interactions with other donor bacteria. In fact, some of these interactions can also have adaptive significance, which can encourage the evolution of metabolic dependencies.

### HGT can affect dependency evolution via coupled gains and losses of genes

Given that newly acquired routes have similar—and sometimes better—path lengths and energy yields, we wondered if their acquisition could promote the loss of corresponding ancestral routes. We thus



**Figure 2. HGT-enabled catabolic routes increase the likelihood of metabolic interactions**

(A) Distribution of the number of newly accessible catabolic routes (routes gained along each phylogenetic branch) across all 1,669 branches. The number of routes starting from nutrients are shown in blue, and those starting from byproducts are in red. New byproduct-driven routes would increase the chance of metabolic interactions with other bacteria (via their byproducts). We find that new catabolic routes are more likely to be byproduct-driven (median 56, versus 51 for nutrient-driven routes;  $P < 10^{-3}$ ; Kolmogorov-Smirnov test).

(B and C) Pie charts comparing new routes with their corresponding ancestral routes on each of the 1,669 branches. We compare new and ancestral routes based on their (B) path length (i.e., is the shortest new path shorter, longer, or the same length as the shortest ancestral path) and (C) energy yield (i.e., does the most energy-yielding new path have a higher, lower, or equal yield than the best ancestral path).

hypothesized the following mechanism through which, contingent on earlier HGT, metabolic dependencies could emerge by subsequent gene loss.

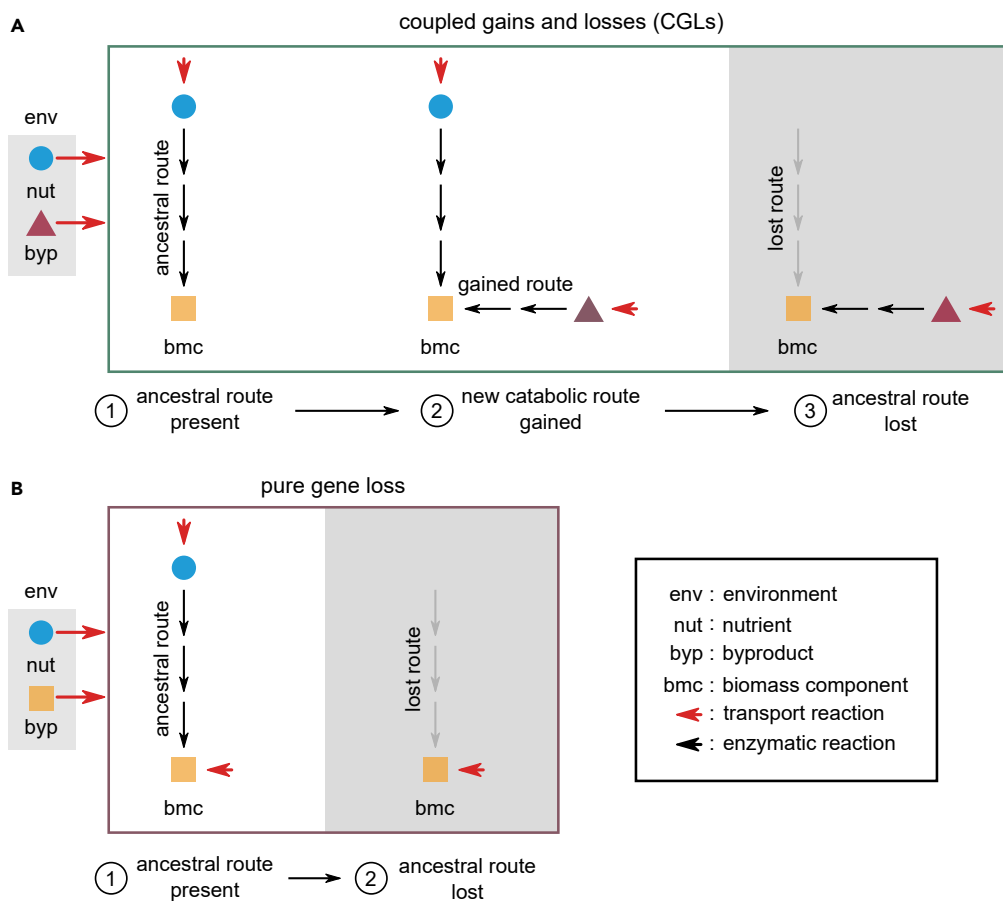
Consider the example illustrated in Figure 3A, with an environment that consists of a nutrient (*nut*, blue circle), and a byproduct (*byp*, purple triangle) secreted by a donor (not shown). Consider a specific acceptor organism, which requires the biomass component (*bmc*, yellow square) either directly or indirectly, to survive. We follow the modification of this acceptor's metabolic network, from ancestor to descendant, in three steps. First, the ancestor uses a specific metabolic pathway (labeled "ancestral route") to convert the available nutrient to the essential biomass component. Second, it gains a catabolic route (labeled "gained route") that uses the byproduct, *byp*, to produce *bmc*. Third, after such a gain, the acceptor loses the ancestral route to *bmc*. This is allowed because *bmc*—crucial for survival—can still be produced through a coupled (or alternate), byproduct-utilizing route. Once lost, however, the acceptor will become obligately dependent on its neighbors to receive the byproduct.

Note that here, until the gain of an alternate route, the ancestral route cannot be lost by the acceptor. Moreover, gaining such a route can not only allow the loss of the ancestral route but also promote it. This is most likely, for instance, when the gained route is more efficient than the preexisting route (which is often the case; Figure 2C). Collectively, we term this process (Figure 3A, steps 1 to 3) "coupled gains and losses" (CGLs). CGLs demonstrate how HGT can crucially affect metabolic dependency evolution.

Contrast CGLs with pure gene loss, which relies on the environmental availability of *bmc* for dependency evolution (Figure 3B). Unlike CGLs, these events are unlikely to be affected by, or depend on, HGT. Interestingly, the same genome (or microbial species) can evolve the same dependency via both mechanisms, but their likelihoods are crucially environment dependent. This is transparent in Figures 3A and 3B, where the primary difference between the two cases is which byproduct is available: in Figure 3A, it is *byp* (purple triangle); in Figure 3B, it is *bmc* (yellow square). The only other difference is the gain of a route to metabolize *byp*, which is likely in an environment where *byp* is available as a byproduct.

### Metabolic dependencies are equally likely to emerge via CGLs and pure gene loss

Given that both coupled gains and losses (CGLs) and pure gene loss are possible mechanisms for metabolic dependency evolution, we asked which of them was more likely to cause the metabolic dependencies observed in extant bacteria. To help answer this, we looked for two distinct but related phylogenetic signatures.




**Figure 3. HGT can affect dependency evolution via coupled gains and losses of genes**

(A) Schematic illustration of coupled gains and losses (CGLs), a new, alternate mechanism for dependency evolution driven by HGT. The gray box on the left indicates the environment, with a nutrient (*nut*, blue circle) and a byproduct (*byp*, purple triangle). Red arrows indicate the secretion and import of metabolites by bacteria. The three steps in the frame show how a bacterial species can evolve a dependency on another species that donates *byp* in the environment; by step 3, the acceptor species eventually depends on *byp* for survival. Each step follows the modification of a part of the acceptor's metabolic network. At step 1, the acceptor uses an ancestral route (black arrows) to convert *nut* in the environment to a key biomass component (*bmc*, yellow square). At step 2, it can alternatively use a newly gained route to convert *byp* to the same *bmc*. The ancestral and gained routes are coupled to each other, because they both produce *bmc*, which is crucial for survival and growth. At step 3, the acceptor loses the ancestral route (gray arrows) but can still produce *bmc* through the gained route. It thus becomes dependent on donors of *byp* for survival.

(B) Schematic illustration of dependency evolution via pure gene loss, not driven by HGT. The environment now has *bmc* available as a byproduct, instead of *byp* (gray box on the left). In this mechanism, step 1 is the same as (A), but unlike (A), at step 2, the acceptor can lose the ancestral route straight away, without requiring an alternate coupled route. However, this requires a particular environment where *bmc* is available.

First, we measured what fraction of ancestor-descendant transitions (each represented by a phylogenetic branch) displayed gain-loss patterns consistent with CGLs; we compared this with the fraction of transitions consistent with pure gene loss. Here, for CGL-consistent transitions, we asked whether a species gained a catabolic route for at least one biomass component that it also lost an alternate route for, during the ancestor-descendant transition (i.e., along a phylogenetic branch, how often did we detect any CGL events, like in Figure 3A; see STAR Methods). For transitions consistent with pure gene loss, we similarly asked how often we observed a species losing a preexisting route, without gaining an alternate catabolic route for the same biomass component (i.e., along a phylogenetic branch, how often did we detect any pure gene loss events, like in Figure 3B; see STAR Methods). To control the likelihood of both events occurring by chance, we also calculated the expected fractions of CGL and pure gene loss events, by repeating our measurements on simulated datasets (see STAR Methods). In each simulated dataset, we randomized



**Figure 4. Metabolic dependencies are equally likely to emerge via CGLs and pure gene loss**

(A) Bar plot showing the fraction of 1,669 phylogenetic branches in which we observed gene gain-loss patterns consistent with coupled gains and losses (CGLs; green) and with pure gene loss (red). Each gray bar represents the corresponding controls, i.e., the expected fraction of branches with patterns consistent with CGLs and pure gene loss, given a random set of gene gains and losses. Error bars show the SE indicating the level of variation across all branches.

(B) Line plot showing the likelihood of evolving dependency via CGLs (green) and pure gene loss (red) in simulated bacterial communities, as a function of the community diversity. The likelihood of dependency is the average fraction of communities in which the observed gains and losses along a branch led to a CGL-based on pure gene-loss-based dependency. Community diversity is the number of coexisting bacterial species in a simulated community. The gray region has  $\geq 7$  species, where CGLs are more likely than pure gene loss.

which genes were gained and lost along each ancestor-descendant transition or branch. Here, we ensured that the same number of genes were gained and lost along each branch as those in our observed dataset. We found at least one CGL event on 33% of all branches (Figure 4A; green); in comparison, we found pure gene loss events on 24% of all branches (Figure 4A; red). Both kinds of events were observed much more likely than expected by chance (Figure 3C; greys). This suggested that by merely observing gene gain-loss patterns, we would conclude that CGLs were more likely to lead to dependencies than pure gene loss.

Second, we added environmental context to the observed gain-loss patterns, by measuring how the likelihood of metabolic dependency evolution depended on a species' microbial community. Specifically, we asked how the chance of both events—CGLs and pure gene loss—depended on the number of species in a microbe's community (hereafter, community diversity). We hypothesized that more diverse communities would have a higher number of available nutrients in the environment, because more species would secrete metabolic byproducts. We expected that increasing community diversity would thus generally favor dependency evolution via both CGLs and pure gene loss; we did not know which of the two would be more favored. To measure the effect of community diversity on the chance of CGLs and pure gene loss across several environments, we asked how often the observed gains and losses would lead to dependencies because of CGLs and pure gene loss alone across hundreds of thousands of random, simulated microbial communities. We used simulated communities as proxies for environments, due to our lack of knowledge of the actual environments of different species across their evolutionary histories; in doing so, we were estimating the typical chance of dependencies emerging via CGLs and pure gene loss. For these simulations, we curated 1,035 environments, each with a different pair of nutrients present (Table S3; see STAR Methods for details). In each environment, we randomly chose unique sets of bacterial species from the 835 in our study as different communities; we chose 100 unique species sets at each level of diversity, from 2 to 10 (see STAR Methods). Because our results saturated beyond a diversity of 10, and because of computational feasibility, we did not continue simulations for more diverse and complex communities.

For each phylogenetic branch, and for each environment-community pair (roughly 100,000 per level of community diversity), we measured how often the observed gains and losses along the branch led to a

new dependency in the descendant; we measured this separately for CGLs and pure gene loss. Briefly, a dependency was CGL mediated when the following three conditions were satisfied: (1) the descendant gained an alternate route while also losing a coupled route (similar to [Figure 3A](#)), (2) the gained route used a metabolic byproduct from the community for biomass production, and (3) the biomass component produced via the gained route was not available as a community byproduct. In contrast, a dependency was pure gene loss mediated when (1) the descendant lost a preexisting route for the production of a biomass component (similar to [Figure 3B](#)), and (2) that biomass component was directly available as a metabolic byproduct from the community (see [STAR Methods](#)).

To illustrate our results, we plotted the fraction of simulations where we detected dependency evolution as a function of community diversity, i.e., we plotted the likelihood of dependencies via CGLs and pure gene loss with increasing community diversity ([Figure 4B](#)). Consistent with our hypothesis, we found that the likelihood of dependencies via both CGLs and pure gene loss increased with increasing diversity ([Figure 4B](#); CGLs in green; pure gene loss in red); the likelihood of both events saturated at high diversity. Strikingly, although pure gene loss was more likely at low diversity ( $<7$ ; white region in [Figure 4B](#)), we found that CGLs were more likely at high diversity ( $\geq 7$ ; gray region in [Figure 4B](#)). This was because although the number of byproducts increase with increasing diversity ([Figure S1](#)), byproducts are more likely to be pathway intermediates (favoring CGLs) than biomass components (favoring pure gene loss); see [Figure S2](#). Collectively, our analyses and simulations suggest that CGLs and pure gene loss are equally likely mechanisms for metabolic dependency evolution in bacteria.

## DISCUSSION

To summarize, here we showed that horizontal gene transfer (HGT) can play a significant role in metabolic dependency evolution in bacteria. Specifically, if an alternate metabolic pathway (or “route”) is gained by HGT, it can promote the loss a preexisting, otherwise indispensable route. Such alternate routes often catabolize metabolic byproducts from coexisting bacteria, thus making bacteria dependent on them. Overall, this is a new mechanism for dependency evolution: coupled gains and losses (CGLs). Phylogenetic evidence suggests that CGLs have occurred much more frequently across bacterial evolutionary history than expected by chance ([Figure 4A](#)). Further, phylogenetic evidence also suggests that CGLs can often be adaptive, as gained pathways are often shorter and more energy-efficient when compared with preexisting pathways ([Figures 2B and 2C](#)).

As a mechanism for metabolic dependency evolution, CGLs are contrasted with pure gene loss, also called the Black Queen hypothesis. We found that although in communities with low diversity, pure gene loss is the more likely cause of dependencies, in communities with high diversity, CGLs are more likely ([Figure 4B](#)). Our results thus enrich and supplement the Black Queen hypothesis, by explaining the role of prior gene gains on eventual gene loss. Note that both mechanisms assume that all metabolic intermediates in pathways can in principle be secreted by cells, and in turn all secreted byproducts may be imported by cells. Previous work has shown this to be a reasonable assumption, which can qualitatively predict metabolic dependencies (our focus), but has shown that it fails to quantitatively predict resulting growth rates, which is why we refrained from making such predictions ([Goyal and Maslov, 2018](#); [Zelezniak et al., 2015](#); [Dal Bello et al., 2021](#)). The reason that this assumption works qualitatively, but not quantitatively, is that there always exist generic transporters (such as several ABC transporters) that can enable the intake and outflow of most metabolites from cells (exceptions include phosphates), but the specific rates at which they enable flow for different metabolites remains unknown.

Contrasted with gene loss, which may only lead to dependencies, gene gain via HGT can result in a variety of outcomes, only one of which is evolving dependencies. Other outcomes include gaining a pathway (or part of a pathway) without losing an alternative pathway, resulting in metabolic redundancies (64% of HGT events), and even losing dependencies and becoming increasingly independent, from only being able to use byproducts to being able to use at least one nutrient for a biomass component (21% of HGT events). These two alternate outcomes are not mutually exclusive of each other but are exclusive of dependency evolution. Thus, metabolic network evolution by HGT takes genomes along richer and more varied evolutionary trajectories compared with gene loss.

We believe that our approach can also aid in a more accurate classification of bacterial lifestyles. Conventionally, bacteria are classified as either free-living or symbiotic in biological databases. Although this



classification suggests that free-living bacteria would often be independent (and symbiotic ones, dependent), these labels can be misleading. For instance, free-living bacteria are often metabolically dependent (D'Souza et al., 2014). In our analyses, we wanted to avoid relying on such a binary classification. We acknowledged that the degree of dependency of a bacterial genome lies along a spectrum and measured it by inferring which key biomass components each genome could synthesize in various nutrient environments. In this way, our approach is more precise and ecologically relevant.

The mechanism we proposed, CGLs, also makes the following prediction about experimental evolution: when co-evolved in a diverse community, bacteria are more likely to lose biosynthetic pathways that they have alternate pathways for; this is less likely when they are evolved alone. As a corollary, adding alternate pathways to bacteria will promote the loss of preexisting pathways. Crucially, pathways do not have to be completely lost. Our work suggests—but we did not quantitatively analyze—cases where only a part of a pathway may be lost when an externally available byproduct happens to be an intermediate in that pathway. Both predictions can be tested via laboratory evolution in a community context.

### Limitations of the study

The framework we used here, combining phylogenetic analyses with metabolic network analyses, can also help quantify the relative contributions of drift and selection to the reduction of bacterial genomes. Progressive genome reduction is often termed “genome streamlining,” and a key question in bacterial genome evolution asks how parallel, or repeatable, streamlining events are. The logic is that more parallel events reflect selection being dominant in genome reduction. We can systematically study these questions within our framework. For instance, we can measure how often we detect the same dependencies evolve along a phylogenetic branch and quantify how similar the corresponding gene loss events are. Similar, or repeatable, gene loss events would be consistent, with selection playing a major role in streamlining: perhaps “weeding out” genes no longer required in certain environments. Dissimilar gene loss events, on the other hand, would suggest that drift dominates. Such analyses are outside the scope of this study and the subject of future work.

Finally, our analyses focused on changes in metabolic network architecture, but dependency evolution can also occur via changes to gene regulatory networks. In experiments, we observe that both metabolic and regulatory changes are responsible for evolved dependencies (Lercher and Pál, 2007; D'Souza and Kost, 2016; Shitut et al., 2019). However, we do not understand how to incorporate the effect of regulatory changes on bacterial phenotypes as well as we do the effect of metabolic changes. Future work in this direction can help better understand the role of regulation on metabolic dependency evolution.

### STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- **KEY RESOURCES TABLE**
- **RESOURCE AVAILABILITY**
  - Lead contact
  - Materials availability
  - Data and code availability
- **METHOD DETAILS**
  - Mapping genomes to metabolic networks
  - Inferring gene gains and losses
  - Assessing gene positions in metabolic networks
  - Classifying metabolites as nutrients, byproducts, and biomass components
  - Calculating catabolic routes enabled by HGT
  - Detecting phylogenetic events consistent with CGLs and pure gene loss
  - Modeling the likelihood of dependency in simulated bacterial communities

### SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.isci.2022.104312>.

## ACKNOWLEDGMENTS

Part of this work was done at, and supported by, the National Centre for Biological Sciences (NCBS-TIFR), Bengaluru, India. I am grateful to Rohini Subrahmanyam, Deepa Agashe, and Saurabh Mahajan for valuable and encouraging discussions. A.G. is supported by the Gordon and Betty Moore Foundation as a Physics of Living Systems Fellow through grant number GBMF4513.

## AUTHOR CONTRIBUTIONS

A.G. conceptualized and designed the research, performed it, analyzed the data, and wrote the paper.

## DECLARATION OF INTERESTS

The author declares that there are no competing interests.

Received: October 21, 2021

Revised: February 17, 2022

Accepted: April 22, 2022

Published: May 20, 2022

## REFERENCES

- Benson, D.A., Karsch-Mizrachi, I., Lipman, D.J., Ostell, J., and Sayers, E.W. (2009). Genbank. *Nucleic Acids Res.* 37, D26–D31. <https://doi.org/10.1093/nar/gkn723>.
- Cohen, O., and Pupko, T. (2011). Inference of gain and loss events from phyletic patterns using stochastic mapping and maximum parsimony—a simulation study. *Genome Biol. Evol.* 3, 1265–1275. <https://doi.org/10.1093/gbe/evr101>.
- Dal Bello, M., Lee, H., Goyal, A., and Gore, J. (2021). Resource–diversity relationships in bacterial communities reflect the network structure of microbial metabolism. *Nat. Ecol. Evol.* 5, 1424–1434. <https://doi.org/10.1038/s41559-021-01535-8>.
- Davis, D.J. (1921). The accessory factors in bacterial growth: iv. the “satellite” or symbiosis phenomenon of pfeiffer’s bacillus (*b. influenzae*). *J. Infect. Dis.* 29, 178–186.
- D’Onofrio, A., Crawford, J.M., Stewart, E.J., Witt, K., Gavriš, E., Epstein, S., Clardy, J., and Lewis, K. (2010). Siderophores from neighboring organisms promote the growth of uncultured bacteria. *Chem. Biol.* 17, 254–264. <https://doi.org/10.1016/j.chembiol.2010.02.010>.
- D’Souza, G., Waschina, S., Pande, S., Bohl, K., Kaleta, C., and Kost, C. (2014). Less is more: selective advantages can explain the prevalent loss of biosynthetic genes in bacteria. *Evolution* 68, 2559–2570. <https://doi.org/10.1111/evo.12468>.
- D’Souza, G., and Kost, C. (2016). Experimental evolution of metabolic dependency in bacteria. *PLoS Genet.* 12, e1006364. <https://doi.org/10.1371/journal.pgen.1006364>.
- Enke, T.N., Datta, M.S., Schwartzman, J., Cermak, N., Schmitz, D., Barrere, J., Pascual-García, A., and Cordero, O.X. (2019). Modular assembly of polysaccharide-degrading marine microbial communities. *Curr. Biol.* 29, 1528–1535.e6. <https://doi.org/10.1016/j.cub.2019.03.047>.
- Fullmer, M.S., Soucy, S.M., and Gogarten, J.P. (2015). The pan-genome as a shared genomic resource: mutual cheating, cooperation and the black queen hypothesis. *Front. Microbiol.* 6, 728. <https://doi.org/10.3389/fmicb.2015.00728>.
- García, S.L., McMahon, K.D., Martínez-García, M., Srivastava, A., Szczyrba, A., Stepanauskas, R., Grossart, H.-P., Woyke, T., and Warnecke, F. (2013). Metabolic potential of a single cell belonging to one of the most abundant lineages in freshwater bacterioplankton. *ISME J.* 7, 137.
- Giovannoni, S.J., Cameron Thrash, J., and Temperton, B. (2014). Implications of streamlining theory for microbial ecology. *ISME J.* 8, 1553–1565. <https://doi.org/10.1038/ismej.2014.60>.
- Goyal, A. (2018). Metabolic adaptations underlying genome flexibility in prokaryotes. *PLoS Genet.* 14, e1007763. <https://doi.org/10.1371/journal.pgen.1007763>.
- Goyal, A., and Maslov, S. (2018). Diversity, stability, and reproducibility in stochastically assembled microbial ecosystems. *Phys. Rev. Lett.* 120, 158102. <https://doi.org/10.1103/physrevlett.120.158102>.
- Handorf, T., Ebenhö, O., and Heinrich, R. (2005). Expanding metabolic networks: scopes of compounds, robustness, and evolution. *J. Mol. Evol.* 61, 498–512. <https://doi.org/10.1007/s00239-005-0027-1>.
- Kanehisa, M., and Goto, S. (2000). Kegg: kyoto encyclopedia of genes and genomes. *Nucleic Acids Res.* 28, 27–30. <https://doi.org/10.1093/nar/28.1.27>.
- King, Z.A., Lu, J., Dräger, A., Miller, P., Federowicz, S., Lerman, J.A., Ebrahim, A., Palsson, B.O., and Lewis, N.E. (2015). Bigg models: a platform for integrating, standardizing and sharing genome-scale models. *Nucleic Acids Res.* 44, D515–D522. <https://doi.org/10.1093/nar/gkv1049>.
- Lercher, M.J., and Pál, C. (2008). Integration of horizontally transferred genes into regulatory interaction networks takes many million years. *Mol. Biol. Evol.* 25, 559–567. <https://doi.org/10.1093/molbev/msm283>.
- Luo, H., Csúros, M., Hughes, A.L., and Moran, M.A. (2013). Evolution of divergent life history strategies in marine alphaproteobacteria. *MBio* 4, e00373–13. <https://doi.org/10.1128/mbio.00373-13>.
- Mahajan, S., and Agashe, D. (2018). Translational selection for speed is not sufficient to explain variation in bacterial codon usage bias. *Genome Biol. Evol.* 10, 562–576. <https://doi.org/10.1093/gbe/evy018>.
- Mas, A., Jamshidi, S., Lagadeuc, Y., Eveillard, D., and Vandenkoornhuysen, P. (2016). Beyond the black queen hypothesis. *ISME J.* 10, 2085–2091. <https://doi.org/10.1038/ismej.2016.22>.
- Maslov, S., Krishna, S., Pang, T.Y., and Sneppen, K. (2009). Toolbox model of evolution of prokaryotic metabolic networks and their regulation. *Proc. Natl. Acad. Sci. U S A* 106, 9743–9748. <https://doi.org/10.1073/pnas.0903206106>.
- McCutcheon, J.P., and Moran, N.A. (2007). Parallel genomic evolution and metabolic interdependence in an ancient symbiosis. *Proc. Natl. Acad. Sci. U S A* 104, 19392–19397. <https://doi.org/10.1073/pnas.0708855104>.
- McCutcheon, J.P., and Moran, N.A. (2012). Extreme genome reduction in symbiotic bacteria. *Nat. Rev. Microbiol.* 10, 13–26. <https://doi.org/10.1038/nrmicro2670>.
- Monk, J.M., Charusanti, P., Aziz, R.K., Lerman, J.A., Premyodhin, N., Orth, J.D., Feist, A.M., and Palsson, B.O. (2013). Genome-scale metabolic reconstructions of multiple *escherichia coli* strains highlight strain-specific adaptations to nutritional environments. *Proc. Natl. Acad. Sci. U S A* 110, 20338–20343. <https://doi.org/10.1073/pnas.1307797110>.
- Morris, J.J., Kirkegaard, R., Szul, M.J., Johnson, Z.I., and Zinser, E.R. (2008). Facilitation of robust growth of *prochlorococcus* colonies and dilute liquid cultures by “helper” heterotrophic

- bacteria. *Appl. Environ. Microbiol.* 74, 4530–4534. <https://doi.org/10.1128/aem.02479-07>.
- Morris, J.J., Lenski, R.E., and Zinser, E.R. (2012). The black queen hypothesis: evolution of dependencies through adaptive gene loss. *MBio* 3, e00036–12. <https://doi.org/10.1128/mbio.00036-12>.
- Ochman, H., and Moran, N.A. (2001). Genes lost and genes found: evolution of bacterial pathogenesis and symbiosis. *Science* 292, 1096–1099. <https://doi.org/10.1126/science.1058543>.
- Pacheco, A.R., Moel, M., and Segrè, D. (2019). Costless metabolic secretions as drivers of interspecies interactions in microbial ecosystems. *Nat. Commun.* 10, 103. <https://doi.org/10.1038/s41467-018-07946-9>.
- Pál, C., Papp, B., and Lercher, M.J. (2005). Adaptive evolution of bacterial metabolic networks by horizontal gene transfer. *Nat. Genet.* 37, 1372–1375. <https://doi.org/10.1038/ng1686>.
- Pande, S., and Kost, C. (2017). Bacterial unculturability and the formation of intercellular metabolic networks. *Trends Microbiol.* 25, 349–361. <https://doi.org/10.1016/j.tim.2017.02.015>.
- Pande, S., Merker, H., Bohl, K., Reichelt, M., Schuster, S., De Figueiredo, L.F., Kaleta, C., and Kost, C. (2014). Fitness and stability of obligate cross-feeding interactions that emerge upon gene loss in bacteria. *ISME J.* 8, 953–962. <https://doi.org/10.1038/ismej.2013.211>.
- Plata, G., Henry, C.S., and Vitkup, D. (2015). Long-term phenotypic evolution of bacteria. *Nature* 517, 369–372. <https://doi.org/10.1038/nature13827>.
- Press, M.O., Queitsch, C., and Borenstein, E. (2016). Evolutionary assembly patterns of prokaryotic genomes. *Genome Res.* 26, 826–833. <https://doi.org/10.1101/gr.200097.115>.
- Segata, N., Börnigen, D., Morgan, X.C., and Huttenhower, C. (2013). PhyloPhlan is a new method for improved phylogenetic and taxonomic placement of microbes. *Nat. Commun.* 4, 2304. <https://doi.org/10.1038/ncomms3304>.
- Shitui, S., Ahsendorf, T., Pande, S., Egbert, M., and Kost, C. (2019). Nanotube-mediated cross-feeding couples the metabolism of interacting bacterial cells. *Environ. Microbiol.* 21, 1306–1320. <https://doi.org/10.1111/1462-2920.14539>.
- Sung, J., Kim, S., Cabatbat, J.J.T., Jang, S., Jin, Y.-S., Jung, G.Y., Chia, N., and Kim, P.-J. (2017). Global metabolic interaction network of the human gut microbiota for context-specific community-scale analysis. *Nat. Commun.* 8, 15393. <https://doi.org/10.1038/ncomms15393>.
- Suzuki, S., Horinouchi, S., and Beppu, T. (1988). Growth of a tryptophanase-producing thermophile, symbiobacterium thermophilum gen. nov., sp. nov., is dependent on co-culture with a bacillus sp. *Microbiology* 134, 2353–2362. <https://doi.org/10.1099/00221287-134-8-2353>.
- Szappanos, B., Fritzemeier, J., Csörgő, B., Lázár, V., Lu, X., Fekete, G., Bálint, B., Herczeg, R., Nagy, I., Notebaart, R.A., et al. (2016). Adaptive evolution of complex innovations through stepwise metabolic niche expansion. *Nat. Commun.* 7, 11607. <https://doi.org/10.1038/ncomms11607>.
- Vos, M., Hesselman, M.C., te Beek, T.A., van Passel, M.W., and Eyre-Walker, A. (2015). Rates of lateral gene transfer in prokaryotes: high but why? *Trends Microbiol.* 23, 598–605. <https://doi.org/10.1016/j.tim.2015.07.006>.
- Wang, T., Goyal, A., Dubinkina, V., and Maslov, S. (2019). Evidence for a multi-level trophic organization of the human gut microbiome. Preprint at bioRxiv, 603365.
- Watsuji, T.-O., Yamada, S., Yamabe, T., Watanabe, Y., Kato, T., Saito, T., Ueda, K., and Beppu, T. (2007). Identification of indole derivatives as self-growth inhibitors of symbiobacterium thermophilum, a unique bacterium whose growth depends on coculture with a bacillus sp. *Appl. Environ. Microbiol.* 73, 6159–6165. <https://doi.org/10.1128/aem.02835-06>.
- Zelezniak, A., Andrejev, S., Ponomarova, O., Mende, D.R., Bork, P., and Patil, K.R. (2015). Metabolic dependencies drive species co-occurrence in diverse microbial communities. *Proc. Natl. Acad. Sci. U S A* 112, 6449–6454. <https://doi.org/10.1073/pnas.1421834112>.

## STAR★METHODS

## KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
KEGG database	KEGG	<a href="https://www.genome.jp/kegg/pathway.html">https://www.genome.jp/kegg/pathway.html</a>
PhyloPhlAn	Segata et al., 2013	<a href="https://www.nature.com/articles/ncomms3304">https://www.nature.com/articles/ncomms3304</a>
Software and algorithms		
Relevant computer code and extracted data files	This study	<a href="https://github.com/eltanin4/black_queen_critique">https://github.com/eltanin4/black_queen_critique</a>
NetworkX	networkx.org	<a href="https://networkx.org/documentation/networkx-1.7">https://networkx.org/documentation/networkx-1.7</a>
GLOOME	Cohen and Pupko, 2011	<a href="http://gloome.tau.ac.il/">http://gloome.tau.ac.il/</a>

## RESOURCE AVAILABILITY

## Lead contact

Further information and requests should be directed to and will be fulfilled by the lead contact, Akshit Goyal ([akshitg@mit.edu](mailto:akshitg@mit.edu)).

## Materials availability

No experimental materials or data were generated during this study.

## Data and code availability

- This paper analyzes existing, publicly available data. All processed data files are available at: [https://github.com/eltanin4/black\\_queen\\_critique](https://github.com/eltanin4/black_queen_critique), and their sources listed in the [key resources table](#).
- All original code has been deposited at [https://github.com/eltanin4/black\\_queen\\_critique](https://github.com/eltanin4/black_queen_critique), and is publicly available as of the date of publication.
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

## METHOD DETAILS

## Mapping genomes to metabolic networks

We extracted a list of all 1,031 bacterial species whose complete genomes were available in the Kyoto Encyclopedia of Genes and Genomes (KEGG) GENOME database (Kanehisa and Goto, 2000). We then pruned this list to remove endosymbionts and closely related genomes. To remove endosymbionts, we used a curated list of endosymbiont genomes based on literature surveys (Mahajan and Agashe, 2018). To remove closely related genomes, when multiple genomes were available for a species, we chose one at random. This resulted in 835 genomes or species, which we used for all subsequent analyses (see [Table S1](#) for the full list). To infer which metabolic genes were present and absent in each species, we extracted a list of all the genes in that species which mapped to a corresponding metabolic reaction in the KEGG REACTION database. We found a total of 3,022 unique genes that were present in at least one of the 835 species in our dataset. We assumed that the set of all mapped metabolic reactions per species was its metabolic network ([Figure 1A](#)).

## Inferring gene gains and losses

To obtain species' phylogenetic relationships, we mapped our 835 genomes to PhyloPhlAn, a well-known phylogenetic tree by matching their GenBank accession numbers (Segata et al., 2013; Benson et al., 2008). To infer the most likely genetic make-up of each ancestor, i.e., the internal nodes of the phylogenetic tree, we used GLOOME, an ancestral state reconstruction method by Cohen and Pupko (2011). GLOOME has been commonly used to study the long-term evolutionary history of gene gains and losses on deep phylogenetic trees (Pál et al., 2005; Szappanos et al., 2016). The parameters we used while running the

method were consistent with previous large-scale studies of bacterial genome evolution, e.g., assuming a stationary root composition (Press et al., 2016); the full set of parameters is available in Table S4. We then calculated which genes were gained and lost along each phylogenetic branch, i.e., between an ancestor and its descendant(s), by comparing their gene presence-absence profiles. We assumed a gene was gained along a branch if it was absent in the ancestor, but present in the descendant; similarly, we assumed a gene was lost along a branch if it was present in the ancestor, but absent in the descendant (Figure 1B). For simplicity, we assumed that all gene gains were due to HGT, though a minority of gains could be due to *de novo* gene birth. Distinguishing between these two possibilities would require detailed sequence-level analysis, which was outside the scope of our manuscript. We verified that using an alternate, maximum-likelihood based method to infer gene gains and losses did not significantly affect our results (< 5% mismatch between gain-loss patterns across all 1,699 branches).

### Assessing gene positions in metabolic networks

To infer which position in a metabolic network a gained gene was likely to occupy, we first mapped each of the 3,022 unique metabolic genes in our study to known metabolic routes in the KEGG MODULE database. Each route in this database is a sequence of steps in the metabolism of a key biomass component; here, the first few steps are catabolic, and the next several steps are anabolic. To each gene, we assigned a position, as follows: (1) first reaction, if the gene corresponded to the first reaction in the route, (2) second reaction, if the gene corresponded to the second reaction, (3) intermediate reaction, for all other reactions (except the last) in the route, and (4) biomass synthesis, for the final reaction in the route. We assumed that genes in categories (1) and (2) were catabolic, and (3) and (4) were anabolic. This assumption is consistent with previous analyses of metabolic gene position (Pál et al., 2005). There are two main sources of ambiguity in this analysis: (1) many metabolic routes are not linear and contain cycles (e.g., the TCA cycle and pentose phosphate pathway), which makes the position of any reaction in them arbitrary; (2) many metabolic routes start not from an externally imported metabolite, but instead from a metabolite produced internally by another pathway (e.g., using ornithine and carbamoyl phosphate to produce arginine). To avoid any ambiguity stemming from these two issues, we only considered genes which were unique to one route. We verified that relaxing this constraint did not significantly affect our results on a statistical level (Figure S3, where genes present in multiple routes are assigned their most frequently observed position). Additionally, we also excluded short routes ( $\leq 3$  reactions, or steps) from our analysis, since it would be difficult to distinguish catabolism from anabolism in them. We calculated the distribution of gained genes in metabolic networks. For this, on each phylogenetic branch, we calculated the fraction of genes gained along that branch at each metabolic network position. We then averaged the fraction of gained genes at each position across all branches (Figure 1C; green bars). As a control, we plotted the expected fraction of gained genes at each position by calculating the average fractions if the genes gained along a branch were a random set of genes, picked from the 3,022 genes in our study; in choosing such random sets, we preserved the number of genes gained along each branch. The average fractions at each position across all branches are plotted as black bars on Figure 1C.

### Classifying metabolites as nutrients, byproducts, and biomass components

Of the 8,755 unique metabolites in our study, we classified certain metabolites as nutrients, byproducts and biomass components based on how likely functional roles in metabolic networks. First, to classify between nutrients and byproducts, we curated metabolites based on previously published large-scale metabolic network analyses (Pacheco et al., 2019; Plata et al., 2015; Sung et al., 2017). These analyses used both manual curation and metabolic modeling to distinguish between metabolites that were most likely to be environmentally available nutrients, from those likely to be the metabolic byproducts of other microbes. We found that metabolites on the exterior of metabolic networks were more likely to be nutrients, while those in the interior, byproducts (46 nutrients, 65 byproducts; Table S2). Second, to classify metabolites as biomass components, we used a database of experimentally-verified metabolic models, BiGG (King et al., 2015). We chose all metabolites listed in the biomass composition of different microbes as biomass components (total 137 metabolites, Table S2).

### Calculating catabolic routes enabled by HGT

We calculated the number of new catabolic routes enabled by HGT along each phylogenetic branch. For this, we first calculated the number of routes in each ancestral metabolic network, i.e., the joint network obtained by combining all metabolic genes and routes in the ancestor's genome. We distinguished between routes starting from nutrients (nutrient-driven routes) and byproducts (byproduct-driven routes).

We calculated the total number of unique paths in each network that started from nutrients and ended at one of the biomass components; similarly, we calculated the number of paths from byproducts. We used standard network analysis algorithms for these calculations, using the NetworkX package in Python. We then calculated the number of nutrient-driven and byproduct-driven routes in each descendant's metabolic network. Note that standard path-finding algorithms may overestimate the number of paths between two metabolites due to our network representation, which only captured pairwise relationships. Thus, if a reaction required two reactants, and one was unavailable, our algorithm would still count that reaction as leading to a feasible path. However, since we were interested in a comparison between nutrient- and byproduct-driven paths, not the absolute number of paths, we argued that this analysis would produce similar systematic differences in both calculations, and thus a comparison was still valid. Along each branch, we calculated the difference between the number of nutrient-driven and byproduct-driven routes between the descendant and ancestor. We plotted the distribution of this difference (the number of newly accessible routes) across all branches in [Figure 2A](#) (nutrient-driven in blue, byproduct-driven in red). To compare the path lengths (number of reaction steps) and energy yields (net number of ATP molecules produced) of the new routes, we did the following along each branch: (1) for path lengths, we compared the lengths of the shortest ancestral path with the shortest new path in the descendant, and asked if a new path was shorter, longer, or of equal length ([Figure 2B](#)); (2) for energy yields, we compared the net number of ATP molecules produced per nutrient or byproduct, along each route; here also we compared the most ATP-yielding ancestral path with the most ATP-yielding new path, and asked if the new path had a higher, lower or equal yield ([Figure 2C](#)). Including other energy currency molecules, such as GTP, did not affect our results, since their reactions can often be swapped with those involving ATP in the KEGG database.

### Detecting phylogenetic events consistent with CGLs and pure gene loss

Along each phylogenetic branch, we asked if there were at least one set of gene gains and losses consistent with coupled gains and losses (CGL-consistent transitions; described in [Figure 3A](#)) and at least one set consistent with pure gene loss (described in [Figure 3B](#)). We first calculated all routes that were lost and gained in the descendant (compared with the ancestor) as described the previous section. A new dependency arises when a biomass component can no longer be produced using only the environmentally-available nutrients. We considered the possibility of a dependency for a biomass component, one component at a time. We assumed there was a CGL-consistent transition on a branch if, for any biomass component: (1) the ancestor had only one nutrient-driven route to produce it and zero byproduct-driven routes, (2) the ancestor gained at least one byproduct-driven route to produce it, i.e., there was at least one such route in the descendant, and (3) the ancestor lost the nutrient-driven route during the transition to descendant. We assumed there was a pure gene loss-consistent transition on a branch if, for any biomass component: (1) the ancestor had only one nutrient-driven route to produce it and zero byproduct-driven routes, (2) the ancestor lost this route, and did not gain any byproduct-driven routes, i.e., the descendant had no routes to produce the biomass component. We calculated the fraction of branches where we detected CGL-consistent ([Figure 4A](#); green) and pure gene loss-consistent transitions ([Figure 4A](#); red). As controls, we calculated the expected fraction of branches with either transitions by using a random set of gains and losses instead ([Figure 4A](#); gray bars); in choosing such random sets, we preserved the number of gained and lost genes along each branch.

### Modeling the likelihood of dependency in simulated bacterial communities

Since environmental and community context is crucial to determining whether a given set of gene gains and losses will result in a metabolic dependency, we tested in how many environments and bacterial communities, the observed CGL and pure gene loss events on different branches (identified in the previous section) would result in an actual dependency. For this, we used metabolic models in  $\sim 1,000,000$  simulated environment-community combinations. We chose 1,035 environments, each with two of the 46 nutrients in [Table S2](#). We chose 900 communities, with 100 at each level of diversity (from two species to 10 species, in steps of 1); each community was a set of bacterial species chosen randomly from the 835 in our study. For each environment-community combination, we calculated the set of byproducts generated by the community by computing which metabolic pathway intermediates each species in the community could produce from the nutrients provided in the environment; we determined this using a popular "scope expansion" algorithm ([Handorf et al., 2005](#); [Goyal, 2018](#)).

For each phylogenetic branch, we then asked: in what fraction of environment-community combinations would the descendant evolve a dependency that the ancestor did not have, and through which

mechanism — CGLs or pure gene loss? For every level of community diversity, we plotted the fraction of examined cases where we detected a possible dependency through CGLs (Figure 4B; green); concurrently we plotted the fraction of cases where the dependency was through pure gene loss (Figure 4B; red). In each environment-community combination, we assumed we detected a CGL-mediated dependency if the following conditions were satisfied between the ancestor and descendant for any one biomass component: (1) the ancestor had only one nutrient-driven route to produce it and zero byproduct-driven routes, (2) the nutrient was available in that environment, (3) the ancestor gained at least one byproduct-driven route to produce it, i.e., there was at least one such route in the descendant, (4) the byproduct was available as a community byproduct, and (5) the ancestor lost the coupled nutrient-driven route during the transition to descendant. Similarly, in each environment-community combination, we assumed we detected a pure gene-loss mediated dependency if, for any biomass component: (1) the ancestor had only one nutrient-driven route to produce it and zero byproduct-driven routes, (2) the biomass component was available as a community byproduct, and (3) the ancestor lost the nutrient-driven route during the transition to descendant.