Defining Coarse-Grainability in a Model of Structured Microbial Ecosystems

Jacob Moran
Department of Physics, Washington University in St. Louis, St. Louis, Missouri, USA

Mikhail Tikhonov*
Department of Physics and Center for Science and Engineering of Living Systems, Washington University in St. Louis, St. Louis, Missouri, USA

(Received 5 August 2021; revised 14 March 2022; accepted 13 April 2022; published 16 May 2022)

Despite their complexity, microbial ecosystems appear to be at least partially "coarse-grainable" in that some properties of interest can be adequately described by effective models of dimension much smaller than the number of interacting lineages. This is especially puzzling, since recent studies demonstrate that a surprising amount of functionally relevant diversity is present at all levels of resolution, down to strains differing by 100 nucleotides or fewer. Rigorously defining coarse-grainability and understanding the conditions for its emergence is of critical importance for understanding microbial ecosystems. To begin addressing these questions, we propose a minimal model for investigating hierarchically structured ecosystems within the framework of resource competition. We use our model to operationally define coarse-graining quality based on reproducibility of the outcomes of a specified experiment and show that a coarse-graining can be operationally valid despite grouping together functionally diverse strains. Furthermore, we demonstrate that a high diversity of strains (while nominally more complex) may, in fact, facilitate coarse-grainability and that, at least within our model, coarse-grainability is maximized when a community is assembled in its "native" environment. Our modeling framework offers a path toward building a theoretical understanding of which ecosystem properties, and in which environmental conditions, might be predictable by coarse-grained models.

DOI: 10.1103/PhysRevX.12.021038 Subject Areas: Biological Physics, Statistical Physics

I. INTRODUCTION

Microbial communities are complex dynamical systems composed of a highly diverse collection of interacting species, and yet they often appear to be at least partially "coarse-grainable," meaning that some properties of interest can be predicted by effective models of dimension much smaller than the number of interacting lineages. For example, industrial bioreactors consisting of hundreds of species are well described by models with $\lesssim 10$ functional classes [1,2]. What makes this possible? One potential explanation is that coarse-grainability is a direct consequence of the hierarchically structured trait distribution across organisms. If 100 interacting phenotypes are all close variants of only ten species, which can be further grouped into just two families, it is natural to expect that the diverse community might be approximately described by a two- or ten-dimensional model. Under this view, effective models are possible because ecosystems are less diverse than a naïve counting of microscopic strains might suggest.

However, recent data reveal this intuition to be too simplistic: A surprising extent of relevant diversity persists at all levels of resolution. Numerous studies highlight the role of strain-level variation in shaping the functional repertoire of a microbial population [3–8]. A recent work by Goyal *et al.* concludes that strains might indeed be "the relevant unit of interaction and dynamics in microbiomes, not merely a descriptive detail" [9]. Surprisingly, however, a greater strain diversity can sometimes enhance predictability instead of undermining it [10]. Equally puzzling, the notion of a bacterial species is undoubtedly useful, despite collapsing together strains that famously may collectively share only 20% of their genes [11]. Moreover, by some assessments, the species-level characterization of a community appears to be too *detailed* and can be coarse-grained further [12], e.g., to the level of a taxonomic family [13].

Rigorously defining coarse-grainability and understanding the conditions for its emergence is of critical importance: Harnessing coarse-grainability is our main instrument for understanding, predicting, or controlling the behavior of these complex systems. Can an ecosystem be coarse-grainable for some purposes but not others? Or in some

*tikhonov@wustl.edu

environments but not others? Can we ever expect the coarse-grained descriptions derived in the simplified environment of a laboratory to generalize to the complex natural conditions? Addressing this exciting set of general questions is an important challenge at the interface of theoretical microbial ecology and statistical physics.

Here, we introduce a theoretical framework to begin addressing these questions. The novelty of our approach is twofold. First, we propose a minimal model for investigating structured ecosystems. Much recent work studies the behavior of large microbial ecosystems in the unstructured regime, where the traits of interacting organisms are drawn randomly (see, e.g., [14–18]). However, real ecosystems assemble from pools of taxa whose trait distributions are highly nonrandom due to functional constraints, common selection pressures, or common descent. These factors create structure at all levels, from the distribution of genes across strains in microbial pangenomes [19–21] to the distribution of function across taxa [12,22,23], with important implications for dynamics, patterns of coexistence, or responses to perturbations [24–27]. In natural communities, taxa can often be grouped by identifiable functional roles, often represented by closely related species or strains. As we seek to define and characterize ecosystem coarse-grainability, it seems clear that this structure must play an important role. Our model implements such structure within a consumer-resource framework in a simple, principled way through trait interactions.

The second novelty of our approach is a framework for defining and evaluating a hierarchy of coarse-grained descriptions. The ultimate performance criterion for a coarse-graining scheme would be its ability to serve as a basis for a predictive model, capable of predicting ecosystem dynamics or properties. However, finding the "most predictive model" is a difficult problem. Here, as a simpler first step, we propose an operational approach which is inspired by the experiments in Ref. [13] and is based on the reproducibility of experimental outcomes. Specifically, we focus on a particular form of coarse-graining in which taxa are grouped together into putative functional groups. Grouping means omitting details, and we say that details are safe to ignore if they do not change the outcome of some specified experiment. Importantly, as we show, choosing different experiments changes which, or whether, details can be ignored.

Specifically, we define how ecosystems can be coarse-grainable in the weak sense, where a desired performance of a coarse-graining can be achieved in a given environment, and in the strong sense, where the performance of a given coarse-graining is *maintained* even as environment complexity is increased. We demonstrate that the same ecosystem can be coarse-grainable under one criterion—even in the strong sense—and not at all coarse-grainable under another. This reconciles the apparent paradox mentioned above, showing that a coarse-graining can be operationally valid for some purposes, despite grouping together functionally diverse strains. We explain how strong-sense coarse-grainability arises in the model considered here and show that this property is context specific: A coarse-graining that works in the organisms' natural ecoevolutionary context is easily broken if the community is assembled in the non-native environment or if the natural ecological diversity is removed. Finally, we discuss the extent to which our findings generalize beyond our model.

II. AN ECOEVOLUTIONARY FRAMEWORK FOR A HIERARCHICAL DESCRIPTION OF THE INTERACTING PHENOTYPES

In order to study the hierarchy of possible coarse-graining schemes for ecosystems, we need an ecoevolutionary framework that would describe players functionally, by a list of characteristics that can be made longer (more detailed) or shorter (more coarse-grained). In addition, for our purposes we also want an ability to tune the complexity of the environment, for example, to study the robustness of a coarse-graining between the simplified conditions of a laboratory and the more complex natural environment. In this section, we present our model implementing these two requirements.

A. The ecoevolutionary dynamics

A given environment presents various opportunities that organisms can exploit to gain a competitive advantage. Imagine a world where all such opportunities or "niches" are enumerated with index $i \in \{1 \ldots L_\infty\}$. The notation L_∞ highlights that, in general, one expects this to be a very large number, corresponding to a complete (and, in practice, unattainable) microscopic description. A strain μ is phenotypically described by enumerating which of these opportunities it exploits, i.e., by a string of numbers of length L_∞ which we denote $\sigma_{\mu i}$. For simplicity, we assume $\sigma_{\mu i}$ to be binary ($\sigma_{\mu i} \in \{0, 1\}$): Strain μ either can or cannot benefit from opportunity i. This allows us to think of evolution as acting via bit flips $0 \mapsto 1$ and $1 \mapsto 0$, corresponding to the acquisition or loss of the relevant machinery ("trait i") via horizontal gene transfer events or loss-of-function mutations.

We assume that the fitness benefit from carrying trait i is largest when the opportunity is unexploited and declines as the competition increases. For a given set of phenotypes present in the community, the ecological dynamics are determined by the feedback between strain abundance and opportunity exploitation [Fig. 1(a)]. Briefly, the strain abundances N_μ determine the total exploitation level $T_i \equiv \sum_\mu N_\mu \sigma_{\mu i}$ of opportunity i. The exploitation level determines the fitness benefit $h_i \equiv h_i(T_i)$ from carrying the respective trait; we choose $h_i(T_i)$ of the form $h_i(T_i) = [b_i/(1 + T_i/K_i)]$. These h_i, in turn, determine the growth or decline of the strains. Specifically, we postulate the following ecological dynamics:

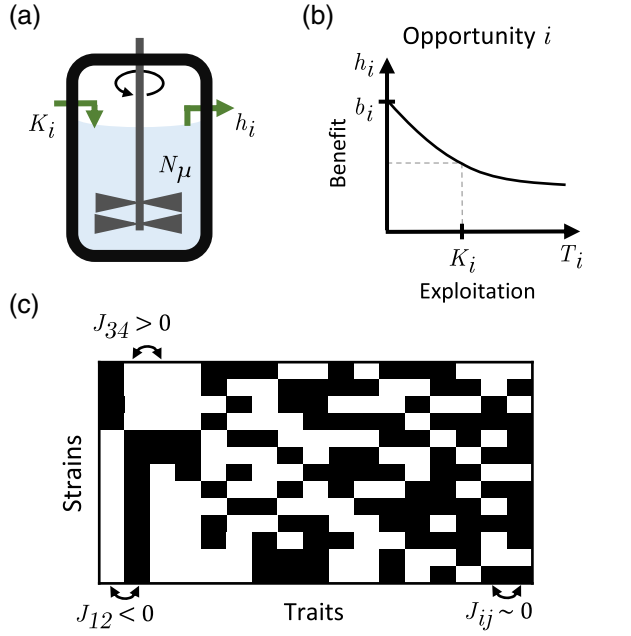

FIG. 1. Our ecoevolutionary framework modifies a standard model of resource competition. Organisms engage in ecological competition for limited resources and evolve by gaining or losing traits. Carrying a trait incurs a cost but enables the organism to benefit from the corresponding resource. Here, our novelty is to consider how traits interact with each other. Combinations that interact unfavorably are costly to maintain; as a result, not all phenotypes are competitive. (a) A metabolic interpretation of our model corresponds to an ecosystem in a chemostat. A set of strains with abundances $\{N_\mu\}$ compete for a set of substitutable resources indexed by i, e.g., alternative sources of carbon. In this interpretation, K_i correspond to resource supply rates, and h_i are the resource concentrations in the effluent. (b) For this work, we adopt a more general interpretation where the resources i need not be specifically metabolic. Instead, we think of i as enumerating any *depletable* environmental opportunities that the phenotypes can exploit, which confer a benefit h_i that declines with exploitation level T_i. We parametrize this dependence by the maximum benefit b_i and the carrying capacity K_i (the exploitation level where the benefit is halved); see the text. (c) In our model, phenotypes are binary vectors described by traits they carry. The most competitive phenotypes (rows in the cartoon) are not random but are shaped by pairwise trait interactions J_{ij}. Strongly synergistic traits ($J_{ij} > 0$) tend to cooccur, while strongly antagonistic traits ($J_{ij} < 0$) are likely not carried together. Such structured phenotypes lead to structured ecosystems, as we investigate.

$$\frac{\dot{N}_\mu}{N_\mu} = \sum_i \sigma_{\mu i} h_i - \chi_\mu \quad \text{strain abundance,} \tag{1a}$$

$$h_i = h(T_i) \equiv \frac{b_i}{1 + T_i/K_i} \quad \text{benefit from carrying } i, \tag{1b}$$

$$T_i \equiv \sum_\mu N_\mu \sigma_{\mu i} \quad \text{exploitation of } i. \tag{1c}$$

In these equations, the parameters b_i and K_i describe the environment, with b_i being the fitness benefit of being the first to discover the opportunity i (at zero exploitation $T_i = 0$) and the "carrying capacity" K_i describing how quickly the benefit declines as the exploitation level T_i increases [Fig. 1(b)]. The quantities χ_μ are interpreted as the "maintenance cost" of being an organism carrying a given set of traits; more on this below.

The dynamics (1) is basically the MacArthur model of competition for L_∞ substitutable "resources" [28–30]. To these dynamics, we add the stochastic arrival of new phenotypes arising through bit flips ("mutations"), as is standard in studies of adaptive dynamics. The combined ecoevolutionary process is simulated using a hybrid discrete-continuous method as described in Supplemental Material [31]. As presented so far, our ecoevolutionary model is similar to, e.g., Ref. [36]; our key novelty (trait interactions) is introduced in the next section. We note, however, that typically the interpretation of resources in models like (1) is metabolic [16,18,37–40]; for example, i might label the different forms of carbon available to a carbon-limited microbial community. Here, we adopt a more general perspective, where i labels any depletable environmental opportunity, which need not be specifically metabolic.

As an example, one way for a strain to survive in chemostat conditions is to develop an ability to adhere to the walls of the device [41]. The wall surface is finite and provides an example of a nonmetabolic limited resource. Similarly, being physically bigger or carrying a rare toxin could be a useful survival strategy, but in both cases the benefit decreases as the trait becomes widespread in the community. Unlike the forms of carbon, which may be numerous but are certainly countable and finite, the list of exploitable opportunities of this kind could be arbitrarily long ($L_\infty \to \infty$), especially when considering the complexity of natural microbial environments. Note that, by construction, our model allows coexistence of a very large number of phenotypes. In many studies, explaining such coexistence is the aim; here, it is our starting point. Rather than asking how a given environment enables coexistence of a diverse community, we start from the observation that natural communities are extremely diverse, interpret this as evidence for the existence of a very large number of (potentially unknown) limiting factors, and ask whether such diversity of types can be usefully coarse-grained.

Modeling fitness benefits as additive [Eq. (1a)] is certainly a simplification. It is also worth noting that the model (1) is special in that it possesses a Lyapunov function [42]; we return to this point below. Nevertheless, this is a good starting step for our program, namely, understanding the circumstances under which coarse-grained descriptions are adequate. Most crucially, a suitable choice of the cost model χ_μ allows us to naturally obtain communities with an

hierarchical structure of trait distributions across organisms mimicking that of natural biodiversity.

B. A simple cost model leads to hierarchically structured communities

Several studies investigated dynamics like (1) with costs assigned randomly (see, e.g., [15–18,40,43]). Here, we seek to build a model where the phenotypes in the community are not random but are hierarchically structured, reproducing phenomena such as divergent taxa belonging to identifiable functional groups, the fine-scale strain diversity found within a species, or the notion of "core" and "accessory" traits in a bacterial pangenome [44]. For this, consider the following cost structure:

$$\chi_\mu = c + \sum_i \chi_i \sigma_{\mu i} - \sum_{i<j} J_{ij}\sigma_{\mu i}\sigma_{\mu j}. \tag{2}$$

The parameter c encodes a baseline cost of essential housekeeping functions (e.g., DNA replication). χ_i is the cost of carrying trait i (e.g., synthesizing the relevant machinery); for most of our discussion, we set $c = 0.1$ and set all $\chi_i \equiv \chi_0 = 0.5$ for simplicity. The key object for us is the matrix J_{ij}, which encodes interactions between traits and shapes the pool of viable (low-cost) phenotypes [Fig. 1(c)]. As an example, the enzyme nitrogenase is inactivated by oxygen, so running nitrogen fixation and oxygen respiration in the same cell requires expensive infrastructure for compartmentalizing the two processes from each other; in our model, this corresponds to a strongly negative J_{ij} (carrying both traits is costly). An example for the opposite case of a beneficial interaction (positive J_{ij}) is a branched catabolic pathway, where sharing enzymes to produce common intermediates reduces the cost relative to running the two branches independently. Crucially, in our model, the parameters c, χ_i, and J_{ij} are the same for all organisms; we refer to them as encoding the "biochemistry" of our ecoevolutionary world.

We now make our key choice. To set J_{ij}, we generate a random matrix of progressively smaller elements, as illustrated in Fig. 2(a). Specifically, we draw the element J_{ij} out of a Gaussian distribution with zero mean and standard deviation $J_0 f[\max(i,j)]$, with a sigmoid-shaped $f(n) = 1/\{1 + \exp[(n - n^*)/\delta]\}$ [see Fig. 2(b)]. Throughout this work, we set $J_0 = 0.2$, $n^* = 10$, and $\delta = 3$. As we will see, this choice for the interaction matrix J implements a hierarchically structured distribution of traits. Intuitively, since high-cost phenotypes are poor competitors, we can think of the interactions J_{ij} as

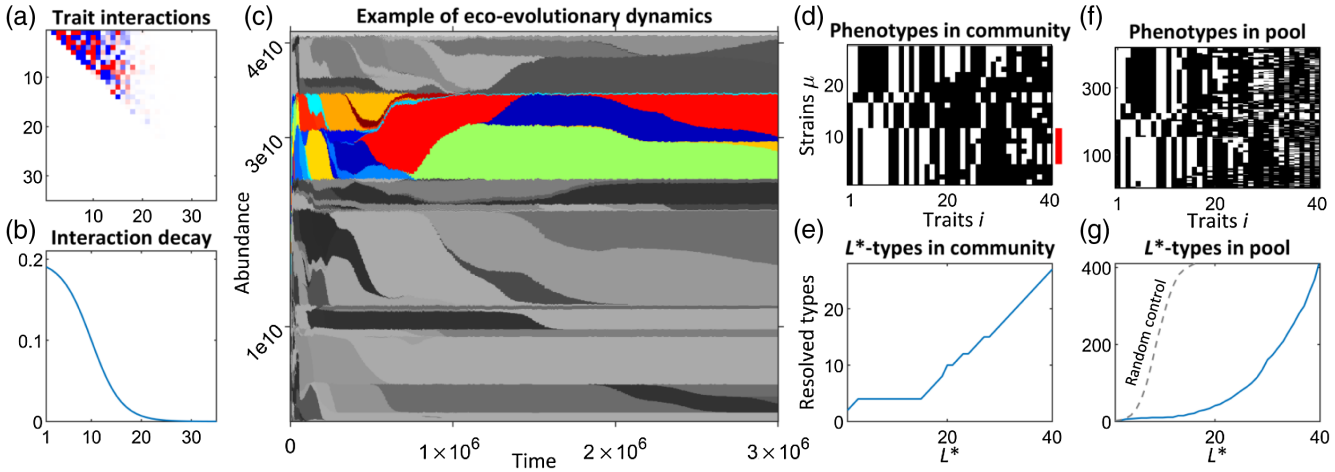

FIG. 2. A simple model of trait interactions leads to hierarchically structured ecosystems. (a),(b) In our model, the traits carried by a given phenotype interact with each other to determine its "maintenance cost" (see the text). The matrix of pairwise trait interactions J_{ij} is drawn randomly and is the same for all phenotypes, encoding the "biochemical constraints"; (a) shows an example (J_{ij} is triangular with one element per trait pair $i \neq j$). We assume an interaction structure such that a few traits interact strongly while others interact weaker and weaker (b). (c) An example of ecoevolutionary dynamics generated in our model. Shading corresponds to different phenotypes. Although new strains continue to emerge and die out throughout the period shown, they can be grouped into several coarse-grained types of approximately stable abundance (one is highlighted in color). (d) The phenotypes present at the end point of the trajectory shown in (c). Each of 27 phenotypes is a row of length $L_\infty = 40$ (white pixels are carried traits). The seven highlighted strains are identical in traits 1–24. We say that they belong to the same "L^*-type," for level of coarse-graining $L^* = 24$. (e) The number of L^*-types in the community in (d), shown as a function of L^*. At a coarse-grained level, the community appears to consist of only four types [one of these is highlighted in (c) using color]; resolving finer substructure requires $L^* > 15$. (f),(g) The same as (d) and (e) for a broader set of strains, pooled over $M_{\text{env}} = 50$ similar environments. The hierarchical structure is maintained (if the trait matrix were randomized, the number of L^*-types would grow exponentially; see the dashed line). Here, we ask: In what sense, if any, could the phenotypic details beyond $L^* \approx 20$–25 be coarse-grained away in this model?

determining the "sensible" trait associations. For strongly interacting traits, only some combinations are competitive, resulting in traits that are mutually exclusive ($J_{ij} < 0$) or that frequently cooccur ($J_{ij} > 0$) in low-cost (viable) phenotypes [Fig. 1(c)]. In contrast, a weakly interacting trait can be gained, be lost, or remain polymorphic, as dictated by the environment. An example might be a gene encoding a costly pump that enables the organism to live in otherwise inaccessible (toxin-laden) regions of the habitat. Such a trait is "weakly interacting" if the cost of running the pump does not significantly depend on the genetic background. As we will see, our model naturally gives rise to hierarchically structured sets of phenotypes that share some "core" functions but differ in others to form finer-scale diversity, resembling the notions of core and accessory traits of a bacterial pangenome [44].

C. Environment defines a strain pool

To build some intuition about the model defined above, consider Fig. 2(c) that shows an example of these ecoevolutionary dynamics for one random biochemistry, and an environment where we set $b_i \equiv b_0 = 1$ for simplicity, and $K_i = K_0 = 10^{10}$ to set the scale of population size as appropriate for bacteria. Grayscale shading corresponds to distinct phenotypes; the community is initialized with a single (randomly drawn) phenotype. The dynamics of Fig. 2(c) illustrate that our framework allows us to define a form of ecosystem stability where all the original phenotypes may have gone extinct and were replaced by others, and yet at a coarse-grained level the ecosystem structure remains recognizably "the same." Here, starting from about $t \simeq 10^5$, the dynamics resemble a stable coexistence of several coarse-grained "species" (one is highlighted in color), whose overall abundance remains roughly stable even as individual strains continue to emerge and die out. To formalize this observation, we need the notion of coarse-grained "L^*-types", which we now introduce.

As we continue the simulation, the dynamics converge to an ecoevolutionary equilibrium (a state where the coexisting types are in ecological equilibrium and no single-bit-flip mutant can invade). In this example, it consists of 27 coexisting phenotypes and is shown in Fig. 2(d). Note that, confirming our expectations, it appears to possess a hierarchical structure. The seven highlighted strains are identical over the first 24 components and differ only in the "tail" (components 25–40). A coarse-grained description that characterizes organisms only by the first $L^* = 24$ traits would be unable to distinguish these strains; we say that these strains belong to the same L^*-type with $L^* = 24$. Figure 2(e) plots the number of L^*-types resolved at different levels of coarse-graining L^* [within the community shown in Fig. 2(d)]. For $L^* = 3$–15, the number of types remains stable at just 4; the color in Fig. 2(c) highlights one of them. Beyond $L^* = 15$, adding more details begins to resolve additional types, up until $L^* = L_\infty$ when the number of L^*-types coincides with the total number of microscopic strains.

Of course, when discussing the diversity of strains one expects to find in a given environment, it is important to remember that no real environment is exactly static, and no real community is in evolutionary equilibrium. To take this into account while keeping the model simple, we consider not a single equilibrium but a collection of communities assembled in $M_{\text{env}} = 50$ similar environments where we randomly perturb the carrying capacity of all opportunities [$K_i = K_0(1 + \epsilon\eta_i)$, with $\epsilon = 0.1$ and η_i are independent identically distributed from a standard Gaussian]; see Supplemental Material [31]. Figure 2(f) shows the set of strains pooled over the 50 ecosystems assembled in this way. This *strain pool* is the central object we seek to coarse-grain. We stress that its construction explicitly depends on the environment. (Or, more specifically, the particular random set of M_{env} similar environments, but $M_{\text{env}} = 50$ is large enough that the results we present are robust to their exact choice.)

As we see in Fig. 2(f), adding more strains to the pool makes its hierarchical structure even more apparent. Quantitatively, the number of L^*-types [Fig. 2(g)] grows much slower than if the traits of each phenotype were randomly permuted (the dashed control curve): Microscopically, perturbing the environment favors new strains, but at a coarse-grained level, these new strains are variations of the same few types. This is precisely the behavior that we aim to capture in our model. Beyond $L^* \approx 20$–25, the number of resolved types begins to grow rapidly. Can this diversity be coarse-grained away? Is there a precise sense in which these tail-end traits are "just details"? To answer this question, we must begin by making it quantitative.

III. COARSE-GRAINING

A. Methodology for defining coarse-grainability

The L_∞-dimensional description we define represents the complete list of niches and opportunities present in a natural habitat. Any recreation in the laboratory is simplified, retaining only some of the relevant factors. We model simplified environments as including resources or opportunities 1 through L [Fig. 3(a)]. The parameter L represents environment complexity. The other key parameter is the level of coarse-graining detail, L^* [Fig. 3(b)]. For each L^*, the identity and combined abundance of L^*-types provides a candidate coarse-grained description of the ecosystem. We seek a quantitative metric for assessing its quality.

Ideally, this assessment would be a comparison of performance of two models—one highly detailed, the other coarse-grained—and our test would evaluate the prediction error for a given property of interest. However, what we build is not a coarse-grained *model* but a hierarchy of coarse-grained variables. These variables could be used to

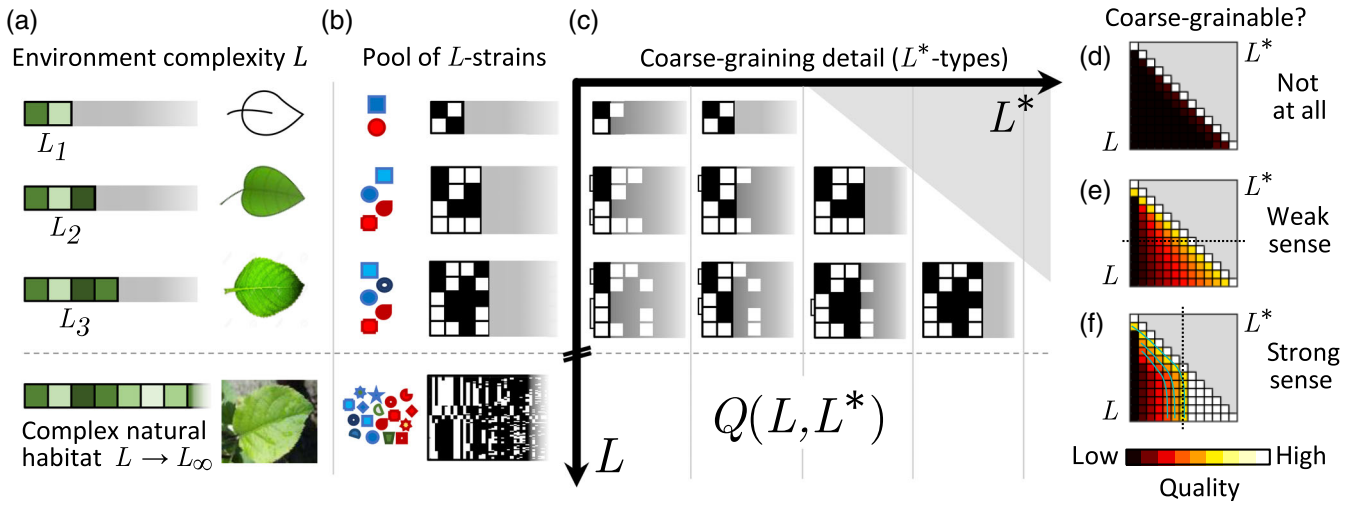

FIG. 3. Defining weak and strong coarse-grainability. (a) The complex natural habitat is modeled as including a large number L_∞ of exploitable resources or opportunities. In a laboratory, we can consider a sequence of ever-more-detailed approximations including resources $1, \ldots, L$ (with the remaining ones set to zero). (b) For each environment, the model describes the pool of strains we expect to encounter [the pool of "L-strains"; see Fig. 2(f)]. For a given L, the strains are unlikely to carry traits i for resources not provided ($i > L$). As environment complexity L increases, the pool becomes increasingly diverse. (c) The set of L-strains can be coarse-grained to a varying level of detail $L^* \leq L$. Let $Q(L, L^*)$ be any quantitative metric (to be defined later) scoring the quality of the L^*-coarse-graining in the environment of complexity L. At $L^* = L$, the strain diversity is fully resolved (no coarse-graining). The "coarse-grainability" of the ecosystem is encoded in the behavior of $Q(L, L^*)$ at $L^* < L$. Different metrics Q encode different operational definitions of coarse-grainability. (d) A non-coarse-grainable ecosystem (*sensu* quality metric Q). The coarse-graining quality remains poor unless the microscopic strain diversity is fully resolved (at $L^* = L$). (e) Weak-sense coarse-grainability: In any given environment (a fixed L, highlighted), a desired quality can be achieved with a coarser-than-microscopic description ($L^* < L$). (f) Strong-sense coarse-grainability: The same coarse-graining (a fixed L^*, highlighted) provides the desired quality even as the environment complexity is increased.

build any number of models, and identifying the most predictive of these is a highly nontrivial task. Here, we sidestep this problem by proposing an operational approach that evaluates a coarse-graining based on the reproducibility of outcomes of a specified experimental protocol.

We will describe and contrast two protocols, each of which could be seen as verifying the validity of the coarse-graining and each yielding its own metric of coarse-graining quality $Q(L, L^*)$; see Fig. 3(c). The "diagonal" entries of Q (with $L^* = L$) correspond to an absence of coarse-graining: The description of strains resolves *all* the traits relevant in a given environment. Coarse-grainability is encoded in the behavior of $Q(L, L^*)$ with $L^* < L$ [Figs. 3(d)–3(f)]. Consider first the behavior of $Q(L, L^*)$ as a function of L^*, with L fixed. If we observe that, in a given environment, sufficient quality can be achieved already with $L^* < L$, we say that the ecosystem is coarse-grainable in the weak sense. For strong-sense coarse-grainability, we ask if the same coarse-grained description continues to perform well even as the environment is made more complex (i.e., instead of fixing L and varying L^*, we fix L^* and vary L). Strong-sense coarse-grainability would be a highly desirable property, but *a priori* it is unclear if it is even theoretically possible.

Crucially, these definitions depend on the choice of the operational criterion for assessing coarse-graining validity (the experiment whose results we require to be reproducible). Below, we show that the same ecosystem can be coarse-grainable in the strong sense under one criterion and yet not coarse-grainable at all under another.

B. Operational definitions of coarse-graining quality $Q(L, L^*)$

In this section, we describe two "experimental" protocols, each of which could be seen as a sensible test of the quality of a coarse-graining. They establish two alternative criteria for a coarse-graining to be operationally valid, which we then contrast.

1. The reconstitution test

One possible criterion is the *reconstitution test*. Drawing a random representative for each of the L^*-types in the strain pool, we seed an identical environment with the representatives we choose, allowing them to reach an ecological equilibrium [Fig. 4(b)]. If the details ignored by the coarse-graining are indeed irrelevant, we expect such "reconstituted" replicates to all be alike. If the reconstituted communities are found to be highly variable depending on exactly which representative we happen to pick, this signals that the distinctions we attempt to ignore are, in fact, significant.

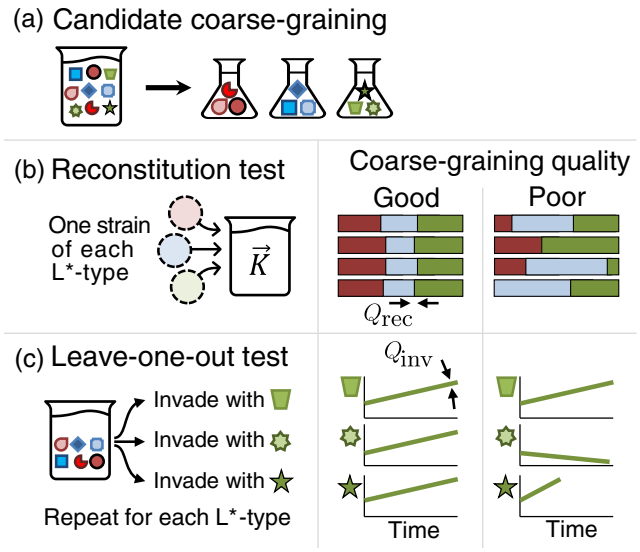

FIG. 4. Specific criteria for assessing coarse-graining quality $Q(L, L^*)$. (a) In this cartoon, the community is coarse-grained into three operational taxonomic units (OTUs), implemented in our model as L^*-types. (b) The reconstitution test. Under this criterion, grouping strains into coarse-grained OTUs is justified if reconstituting a community from a single representative of each OTU yields similar communities regardless of which representatives we pick. As a quantitative measure, we compare the OTU abundances across replicates. (c) The "leave-one-out test." Under this criterion, grouping strains into coarse-grained OTUs is justified if the strains constituting OTU X (green in this cartoon) all behave similarly when introduced into a community missing X. As a quantitative measure, we compare the invasion rates of the left-out strains.

Quantitatively, for each L^*-type μ_*, let us denote $n_{\mu_*}^{(\alpha)}$ its final relative abundance (i.e., the fraction of total population size) in the reconstituted replicate α. The coefficient of variation of $n_{\mu_*}^{(\alpha)}$ over α (denoted $\mathrm{CV}_\alpha[n_{\mu_*}^{(\alpha)}]$) provides a natural measure of variability across replicates. To combine these into a single number, we compute the average such variability over all L^*-types μ_*, weighted by their mean relative abundance across replicates (denoted $\langle n_{\mu_*}^{(\alpha)}\rangle_\alpha$):

$$Q_{\rm rec} = \sum_{\mu_*} \langle n_{\mu_*}^{(\alpha)}\rangle_\alpha \mathrm{CV}_\alpha[n_{\mu_*}^{(\alpha)}].$$

Since the coefficient of variation is, by definition, $\mathrm{CV}_\alpha[n_{\mu_*}^{(\alpha)}] = \mathrm{std}_\alpha[n_{\mu_*}^{(\alpha)}]/\langle n_{\mu_*}^{(\alpha)}\rangle_\alpha$, our metric simplifies to $Q_{\rm rec} = \sum_{\mu_*} \mathrm{std}_\alpha[n_{\mu_*}^{(\alpha)}]$. With this definition, a perfect reconstitution has $Q_{\rm rec} = 0$. Conveniently, this is automatically the case if $L^* = L$ (no coarse-graining).

2. The leave-one-out test

As we will see, the criterion defined above is extremely stringent and is rarely satisfied. In this section, we introduce a weaker version. Instead of the composition of the entire community, we explicitly focus on one particular property of interest (below, the invasion rate of a strain). Furthermore, instead of requiring the grouped-together strains to be interchangeable in absolute terms, we ask that they behave similarly *in the context of the assembled community.*

Specifically, for a given scheme grouping strains into coarse-grained types, consider assembling a community missing a particular coarse-grained type μ_* [the ecological equilibrium reached when combining all the strains in the pool, except those belonging to type μ_*; see Fig. 4(c)]. We judge the coarse-graining as valid if the different strains constituting the missing type μ_* all behave similarly when introduced into this community. As one example, we can compare their initial growth rates if introduced into the community at low abundance, called henceforth "invasion rate" (other possible choices include the abundance the strain reaches if established or the level of niche exploitation h_i in the resulting community; these are shown in Supplemental Material [31]). If the invasion rates are similar, describing the community as missing the coarse-grained type μ_* indeed is consistent. If, however, the invasion rates vary strongly, we conclude that the features our coarse-graining is neglecting are, in fact, important.

Quantitatively, denote the invasion rate of strain μ into a community missing type μ_* as r_{μ,μ_*}. We define

$$Q_{\rm inv} = \sum_{\mu_*} \bar{n}_{\mu_*} \mathrm{std}_{\mu\in\mu_*} r_{\mu,\mu_*},$$

where $\bar{n}_{\mu_*}$ is the relative mean abundance of strains belonging to type μ_* in the pool and $\mathrm{std}_{\mu\in\mu_*}$ denotes the standard deviation over all strains belonging to μ_* weighted by strain abundance in the pool (i.e., a strain's combined abundance observed across the set of $M_{\rm env}$ environments used to define the pool). Once again, at $L^* = L$ we automatically have $Q_{\rm inv} = 0$, as this corresponds to the fully microscopic description (each type μ_* is represented by exactly one strain). Note that this averaging convention (weighted by abundance in the pool) is slightly different from that used in the previous section (using average abundance across the assembled replicates). Using the same convention for both $Q_{\rm inv}$ and $Q_{\rm rec}$ does not change our results but artificially inflates the latter with noise from low-abundance (rare) strains. (For details, see Supplemental Material [31].)

To illustrate the difference between the two criteria, consider the statement that a community consisting of *Tetrahymena thermophila* and *Chlamydomonas reinhardtii* cannot be invaded by *Escherichia coli* [45]. What meaning should we ascribe to this statement when phrased in terms of coarse-grained units rather than specific strains? Under the first criterion, we require that if we combine any single strain of *T. thermophila*, any strain of *C. reinhardtii*, and

any strain of *E. coli*, only the first two survive. Under the second criterion, we combine a vial labeled *T. thermophila*, containing the entire diverse ensemble of its strains, with a similarly diverse vial of *C. reinhardtii* and verify that the resulting community cannot be invaded by any individual strain of *E. coli* [46].

Note that, in our model, the existence of a Lyapunov function [42] means the ecological equilibrium is uniquely determined by the environment and the identity of the competing strains; their initial abundance or the order of their introduction does not matter. While this is a simplification, this property is very useful for our purposes, since any lack of reproducibility between reconstituted communities is then clearly attributable to faulty coarse-graining. In a model where even identical phenotypes could assemble into multiple steady states, distinguishing this variability from the variability due to strain differences adds a layer of complexity to our analysis.

IV. RESULTS

A. A coarse-graining may be operationally valid despite grouping functionally diverse strains

Throughout this section, we continue to use an environment with $K_i \equiv K_0$ and $b_i \equiv b_0$ (all L_∞ opportunities are equally lucrative). In practice, when approximating a complex environment in the laboratory, we try to capture the most salient features first. Thus, it would have been perfectly natural to instead let K_i and/or b_i decline with i; one would expect this to improve coarse-grainability, and this is indeed the case (see Supplemental Material [31]). The motivation for our choice is twofold: First, keeping all K_i and b_i the same requires fewer parameters than choosing a particular functional form of decline with i. Second, the regime where no niches are obviously negligible only makes it more striking to find that an ecosystem can be not only coarse-grainable, but coarse-grainable in the strong sense.

Figure 5(a) plots $Q_{\mathrm{inv}}(L, L^*)$ for the leave-one-out test comparing the invasion rates of different strains falling into the same coarse-grained types. We find that any desired coarse-graining quality can be achieved by a sufficient L^* and is almost unaffected by L. As environment complexity increases and becomes capable of sustaining an ever-growing number of microscopic strains, each L^*-type becomes increasingly diverse. Nevertheless, all the strains in the same L^*-type continue to behave similarly by our invasion-rate-based metric; in other words, under this criterion, the ecosystem is coarse-grainable in the strong sense.

And yet, it would be wrong to conclude that the traits beyond a given L^* are "negligible" in any absolute sense. This is clearly demonstrated by the reconstitution test [Fig. 5(b)]. If we attempt to reconstruct the community from its members, *every* detail matters: No amount of coarse-graining is acceptable. We now explain this apparent paradox within our model.

Consider a community at an ecological equilibrium, and let us focus on a particular phenotype σ carrying one of the

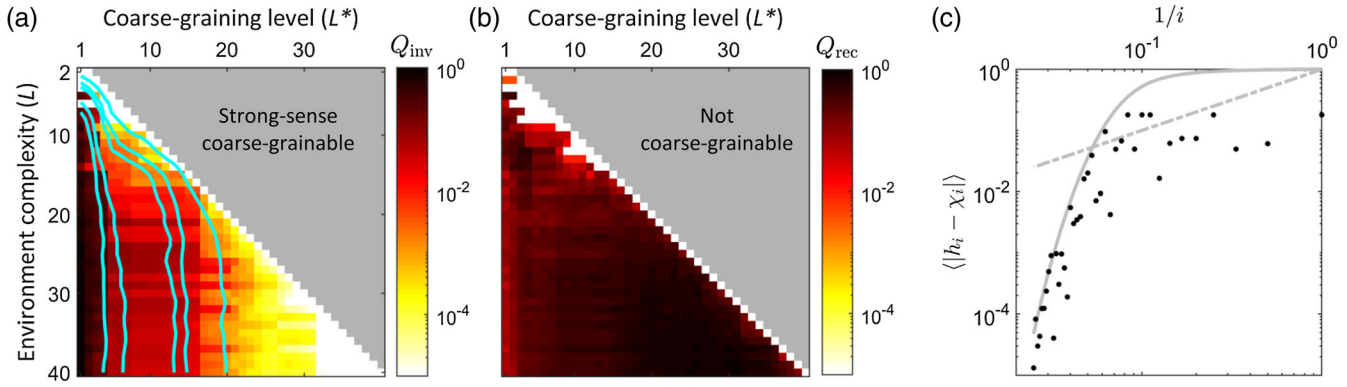

FIG. 5. The same ecosystem can be coarse-grainable under one criterion, but not under another. (a) If coarse-graining quality is evaluated using the leave-one-out test (assessing reproducibility of strain invasion rates), our ecosystem model is coarse-grainable in the strong sense: The acceptable level of coarse-graining, determined by the desired quality score (isolines of Q), is robust to environment complexity [compare to Fig. 3(f)]. (b) In contrast, under the reconstitution test criterion, no amount of coarse-graining is acceptable [compare to Fig. 3(d)]. This comparison shows that a coarse-graining can be operationally valid for a given purpose (a) even when the strains it groups together are functionally diverse (b). Both heat maps represent a single random biochemistry, the same in both. Isolines in (a) are averaged over 20 biochemistries to demonstrate robustness (see Supplemental Material [31]). (c) Explaining the origin of strong-sense coarse-grainability in our model. The plot shows the scaling with i of $|h_i - \chi_i|$ [computed for $L^* = 30$, $L = 40$, and averaged across communities assembled for the leave-one-out test in (a)]. The strong-sense coarse-grainability in (a) is ensured whenever the decay is faster than $1/i$ (dashed gray line). Intuitively, this makes the tail-end traits approximately neutral in the assembled community; see the text. We expect this scaling to be controlled by the sigmoidal decay of trait interaction magnitude $|J_{ij}|$, as confirmed here [solid gray line; the same as Fig. 2(b) but normalized to a maximum of 1 to show the decay of interaction strength rather than their absolute magnitude].

weakly interacting (tail-end) traits i_0: $\sigma_{i_0} = 1$. What would be the fitness effect of losing this trait? Losing the benefit h_{i_0} from opportunity i_0 is offset by the reduction in maintenance cost; for a weakly interacting trait, the contribution from the term $\sum_j J_{ji_0}\sigma_j\sigma_{i_0}$ is negligible, and the change in cost is simply χ_{i_0}. We conclude that the fitness effect of losing the trait is $\delta f = \chi_{i_0} - h_{i_0}$. At an evolutionary equilibrium, we therefore have $h_{i_0} = \chi_{i_0}$ (the "functional attractor" state [40]). When this condition is satisfied, we say that the opportunity or niche i_0 is "equilibrated." If a weakly interacting niche is equilibrated, carrying the respective trait becomes approximately neutral.

Here, our community is not at the evolutionary equilibrium; nevertheless, a sufficiently diverse strain pool similarly ensures that the opportunities corresponding to the weakly interacting (tail-end) traits become approximately equilibrated: $h_i \approx \chi_i$. For a simpler model where the phenotype costs χ_μ are drawn randomly, the mechanism for this can be understood analytically (the "shielded phase" in Ref. [18]; see also Ref. [47]). Here, the costs are not random, but, as long as trait interactions are weak, one expects the behavior to be similar (see Supplemental section S6.2 in Ref. [18]). This expectation is confirmed in simulations. Figure 5(c) shows the observed niche disequilibrium $h_i - \chi_i$ as a function of $1/i$. The plot confirms that the tail-end niches ($1/i \to 0$) are increasingly well equilibrated ($|h_i - \chi_i|$ decays with i). The strong-sense coarse-grainability in Fig. 5(a) is ensured whenever the decay is faster than $1/i$ (dashed gray line). This is because, with this scaling, the sum of contributions from the omitted tail-end traits is bounded (see Supplemental Material [31]). The analytical argument in Ref. [18] leads us to expect the disequilibrium to be controlled by the decaying typical magnitude of interactions $|J_{ij}|$ (solid gray line). If the tail-end niches are equilibrated, carrying the respective traits becomes approximately neutral, and the ability of a strain to invade is entirely determined by its phenotypic profile over nonequilibrated niches, explaining the observations in Fig. 5(a). We conclude that, in our model, the strong-sense coarse-grainability is a consequence of the faster-than-$1/i$ decay of interaction strength in Fig. 3(b).

Crucially, however, this approximate neutrality applies only in the environment created by the assembled community and does not mean that the distinctions are functionally negligible. For instance, consider the (Lotka-Volterra-style) interaction term for a given pair of strains $\mu \neq \nu$:

$$A_{\mu\nu} \equiv \frac{1}{N_\mu}\frac{\partial \dot{N}_\mu}{\partial N_\nu} = \sum_i \sigma_{\mu i}\sigma_{\nu i}\frac{h_i^2}{b_i K_i} = \frac{\sum_i \sigma_{\mu i}\sigma_{\nu i}h_i^2}{b_0 K_0},$$

where we substitute $b_i \equiv b_0$ and $K_i \equiv K_0$ for our environment. Even when tail-end niches are equilibrated with $h_i \approx \chi_i = \chi_0$, we find that each of them contributes equally to the interaction term: No detail is negligible.

This argument directly relates the observed effect to the distinction between a trait that is truly neutral and one that is effectively neutral in the assembled community only. A truly neutral trait, one incurring almost no cost and bringing almost no benefit, would have $h_i \to 0$ and its contribution to the interaction term $A_{\mu\nu}$ would indeed be small. And, indeed, if we repeat our analysis for a scenario where both b_i and χ_i decline with i, we find that neglecting the tail-end traits becomes an adequate coarse-graining also for the reconstitution test (see Supplemental Material [31]).

The conclusion from contrasting Figs. 5(a) and 5(b) is worth emphasizing. In the example we construct, the coarse-grained description is valid *sensu* Fig. 5(a). This means that, for instance, we can meaningfully say that "a community assembled of OTU#1 and OTU#2 can be invaded by OTU#4." We can even measure, e.g., the invasion rate and be assured that it is quantitatively reproducible, with a bounded error bar, across the many strains that constitute OTU#4 at the microscopic level. Despite all this, the *interaction* between the OTUs as coarse-grained units is not actually definable: Any specific pair of strains of OTU#1 and OTU#4 may interact differently with each other, as is indeed observed experimentally [9].

Our focus on reproducibility of L^*-type abundances across replicates is inspired by the experiments in Ref. [13]. To complete this parallel, we should mention that, besides inoculating the same environment with a set of similar inocula, as we do for our reconstitution test [cf. Fig. 4(b)], one could also use the same inoculum to seed a set of similar environments. To implement this in our model, we use the strain pool constructed as described in Sec. II C to inoculate a set of environments with slight variations in the carrying capacities $K_i \approx K_0$ drawn from a Gaussian distribution of width $\epsilon = 0.1$. This is meant to represent the unavoidable variability present in any experimental replicates of the "same" environment $K_i \approx K_0$, which can affect fitness even when subtle [48]. After assembling the replicate communities, we find that community composition is more reproducible at coarser levels of description [Figs. 6(b) and 6(c)], consistent with the experimental observations of Goldford *et al.* [13] and with the interpretation of this pattern as resulting from functional redundancy within coarse-grained types [12,49].

B. Using non-native strain pool reduces coarse-grainability

The previous section describes a mechanism by which strain diversity can aid coarse-grainability. As we explain, in our model ecosystem the diverse set of strains contained within the coarse-grained units is able to successfully equilibrate the weakly interacting niches, rendering them effectively neutral and leading to the behavior shown in

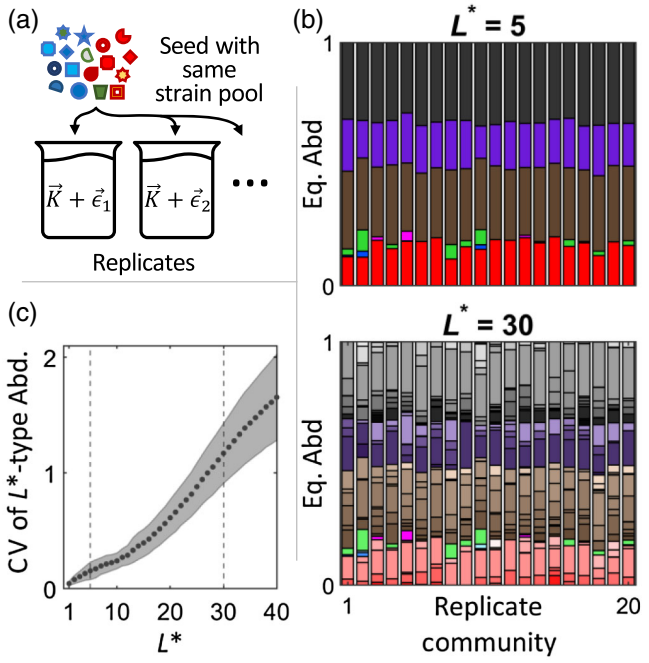

FIG. 6. Replicate communities assembled in similar environments are more reproducible at coarser level of description. (a) A set of similar environments $\vec{K}+\vec{\epsilon}$ (each carrying capacity modified by 10% Gaussian noise) is inoculated with the same strain pool and brought to ecological equilibrium. (b) Equilibrium relative abundances of coarse-grained L^*-types across 20 replicates, shown for two levels of coarse-graining. A coarser description ($L^*=5;7$ resolved types) is more reproducible, consistent with experimental observations [13]. (c) The variability of coarse-grained descriptions increases with the level of detail. Variability is measured as the average coefficient of variation (CV) in relative abundance of an L^*-type over 100 replicates, weighted by L^*-type mean relative abundance across replicates. Dashed lines mark $L^*=5$, 30 shown in (b). Data points and shading show mean ±SD over 20 random choices of biochemistry $\{J_{ij}\}$. All simulations performed with $L=40$.

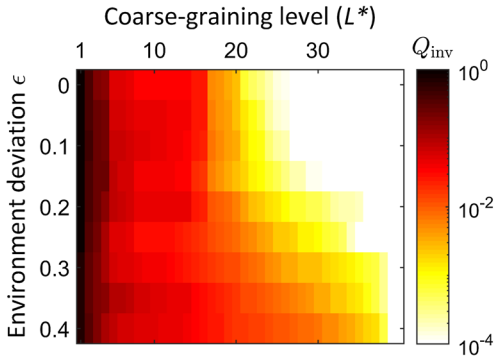

FIG. 7. A coarse-graining scheme works best when the environment is populated by the native strain pool. The same test environment as in Fig. 5(a) is inoculated with strain pools that evolve in environments increasingly further away (see the text). The coarse-graining quality is assessed by leave-one-out experiments and shown as a function of L^* and environment deviation ϵ from the test condition. L is fixed at $L=40$ for comparison with the last row in Fig. 5(a). As the environments for generating strain pools are modified, the traits that were previously negligible can no longer be coarse-grained. The same random biochemistry as in Fig. 5(a) is used, and each pixel is averaged over 20 random environments.

Fig. 5(a). However, for this to occur, the strain pool diversity needs to be derived from a sufficiently similar set of environments, as we now show.

To see this, we repeat the leave-one-out analysis in Fig. 5(a), except now we inoculate the same test environment of complexity $L=40$ (using $K_i=K_0$, $b_i=b_0$ as before) with strain pools derived from *other* environments that are increasingly dissimilar to it. Specifically, following the procedure described in Sec. II C, we generate strain pools in environments with $K_i=K_0(1+\epsilon\eta_i)$, where η_i are drawn from the standard normal distribution and ϵ is the parameter we vary. (The b_i are left at $b_i=b_0$ for simplicity.) The results are presented in Fig. 7, which shows the performance of different L^*-coarse-grainings under the leave-one-out test.

At $\epsilon=0$, this is identical to the protocol in Fig. 5(a). We see that describing phenotypes by 20 traits is sufficient for the invasion rates of grouped-together strains to be consistent within an error bar of $Q<10^{-2}$. However, as ϵ is increased and the strain pools we use are derived from increasingly distant environments, the same coarse-graining becomes insufficient. Instead, a substantially higher level of coarse-graining detail L^* is required to maintain the desired quality. In summary, we find that, in our model, a coarse-graining scheme works best when the environment is populated by the native strain pool.

V. DISCUSSION

The interface of statistical physics and theoretical ecology has a long and highly influential tradition of studying large, random ecosystems, starting from the work of May [50]. The key insight of this approach is that patterns that are typical to some *ensemble* of ecosystems are more likely to be generalizable and reproducible than the details specific to any one realization. However, the choice of the ensemble (and, in particular, adding constraints relevant for natural ecosystems) can affect predictions significantly [51–57]. Which predictions of random-interaction models are robust to introducing more realistic structures and, conversely, which phenomena cannot be explained without invoking structural constraints is an active area of research [58].

Resource competition models—one of the simplest frameworks explicitly linking composition to function—offer a highly promising context to begin addressing these questions, with much recent progress. For example, it was recently shown that cross-feeding interactions structured by shared "rules of metabolism" (but otherwise random) can

reproduce a surprising range of experimental observations [13,43]. This work made it possible to begin disentangling which experimental observations can be seen as evidence for nontrivial underlying mechanisms and which can be reproduced already in the simplest models.

In this work, we presented a simple framework that allows generating random ecosystems with community structure as a tunable control parameter. Instead of postulating a fixed architecture, such as a number of discrete "families" of phenotypes [43], we use a biologically motivated approach to derive it from functional trade-offs, parametrized by a matrix of trait-trait interactions J. Simple (few-parameter) choices for J generate communities with complex structures, including hierarchical architectures which, at least superficially, appear to mimic those of natural biodiversity. Perhaps the most immediate benefit from such a framework would be to help develop new ways to quantify the highly multidimensional concept of community structure across scales, such as, for example, the structure of microbial pangenomes.

In this spirit, here we used this framework to quantify the notion of coarse-grainability. We proposed a way to operationally define the quality of a coarse-grained description based on the reproducibility of outcomes of a specified experiment. We demonstrated that an ecosystem can be coarse-grainable under one criterion while also not at all coarse-grainable under another.

Specifically, one way to approach the coarse-graining problem is to group together only the individuals that are to a sufficient extent interchangeable. This is the criterion we introduced as a "reconstitution test" and is the criterion implicitly assumed by virtually all compositional models of ecosystem dynamics [59]. However, experimental evidence [3–6,9] suggests that, unless we are willing to resolve types differing by as few as 100 bases, this criterion is likely violated in most practical circumstances. It is certainly violated when grouping strains into functional groups, or taxonomic species or families [11,12,60–63]. One expects, therefore, that explaining the practical successes of such descriptions would require a different definition of what makes a coarse-graining scheme adequate.

We proposed that this can be achieved with only a subtle change to the criterion: namely, by requiring that the grouped strains be approximately interchangeable not in all conditions but in the conditions created by the assembled community itself. As long as the strains we study remain in a diverse ecological context, and as long as this diversity is derived from a sufficiently similar environment, we find that the coarse-grained description can be consistent in the sense that the strains grouped together possess similar properties of interest (e.g., invasion rate and postinvasion abundance).

In this paper, we focused on a case where the traits were differentiated only by the strength of their interactions, which established a unique hierarchy among them (a clear order in which to include them in the hierarchy of coarse-grained descriptions). In the more general case, the trait cost χ_i or the trait usefulness in a given environment (b_i and K_i) will set up alternative, potentially conflicting hierarchies. We expect the model to have a rich phenomenology in this regime, which we have not considered here. Another obvious limitation of our analysis is that our model includes only competitive interactions. A simple way to extend our framework would be to include cross-feeding interactions; we leave this extension for future work.

Our analysis introduced a distinction between weak-sense and strong-sense coarse-grainability based on whether the performance of a coarse-graining scheme is robust to increasing the environment complexity. We explained how strong-sense coarse-grainability arises in our model, linking it to a previously described phenomenon, namely, that a sufficiently diverse community may "pin" resource concentrations (here, the exploitation of environmental opportunities) at values that are robust to compositional details [16,18,37,47]. Tracing its origin makes it clear that strong-sense coarse-grainability in our model is only as good as the assumption that the cost of carrying weakly interacting traits is independent of the phenotypic background. Whether this assumption is ever a good approximation in natural ecosystems remains to be seen. Still, our argument provides an explicit mechanism for how coarse-grainability can not only coexist, but may, in fact, be facilitated by diversity.

The fact that strong-sense coarse-grainability is at least theoretically possible is intriguing also for the following reason. Throughout this work, we interpreted L as indexing a sequence of ever-more-complex environments (e.g., a minimal medium with one carbon source; a mixture of several carbon sources; resuspended homogenized leaf matter; or an actual leaf). An alternative perspective, however, is to think of a single environment of interest and take L to be the level of detail at which it is modeled. Any model we could ever consider, however detailed, is necessarily incomplete. Consider the example of the human gut: How important is the exact geometry of the gut epithelium, the effect of peristalsis and flow on small-scale bacterial aggregates, or the exact role of the vast diversity of uncharacterized secondary metabolites [64,65]? It seems plausible that the complete list of factors shaping this ecosystem includes many we will never even know about, let alone include in our models. Our analysis raises an intriguing—though, at this point, purely speculative—question of whether the tremendous diversity of natural ecosystems might afford our models some unexpected degree of robustness to such unknown details.

In conclusion, there are many reasons to believe that analyzing a species in artificial laboratory environments might be of limited utility for understanding its function or interactions in the natural environment [66]. Usually, however, the concern is that the laboratory conditions are

too simple, and, in reality, many more details may matter. Here, we use our model to propose that, at least in some conditions, the opposite can be true: Understanding the interaction of two strains in the foreign conditions of the Petri dish may require a much more detailed knowledge of microscopic idiosyncracies. Removing individual strains of a species from their natural ecoevolutionary context may eliminate the very reasons that make a species-level characterization an adequate coarse-graining of the natural diversity.

All simulations were performed in MATLAB (Mathworks, Inc.). The associated code, data, and scripts to reproduce all figures in this work are available at Mendeley Data [67]. A PYTHON implementation of the model is also available [68].

ACKNOWLEDGMENTS

We thank J. Grilli, C. Holmes, R. S. McGee, and C. Strandkvist for helpful discussions. This research was supported in part by National Science Foundation Grant No. PHY-1748958, the Gordon and Betty Moore Foundation Grant No. 2919.02, and the Kavli Foundation.

[1] A. Jeglot, J. Audet, S. R. Sorensen, K. Schnorr, F. Plauborg, and L. Elsgaard, *Microbiome Structure and Function in Woodchip Bioreactors for Nitrate Removal in Agricultural Drainage Water*, Front. Microbiol. **12**, 678448 (2021).

[2] S. Bertacchi, M. Ruusunen, A. Sorsa, A. Sirviö, and P. Branduardi, *Mathematical Analysis and Update of ADM1 Model for Biomethane Production by Anaerobic Digestion*, Fermentation **7**, 237 (2021).

[3] T. D. Lieberman, K. B. Flett, I. Yelin, T. R. Martin, A. J. McAdam, G. P. Priebe, and R. Kishony, *Genetic Variation of a Bacterial Pathogen within Individuals with Cystic Fibrosis Provides a Record of Selective Pressures*, Nat. Genet. **46**, 82 (2014).

[4] A. Tett, K. D. Huang, F. Asnicar, H. Fehlner-Peach, E. Pasolli, N. Karcher, F. Armanini, P. Manghi, K. Bonham, M. Zolfo *et al.*, *The Prevotella Copri Complex Comprises Four Distinct Clades Underrepresented in Westernized Populations*, Cell Host Microbe **26**, 666 (2019).

[5] S. Zhao, T. D. Lieberman, M. Poyet, K. M. Kauffman, S. M. Gibbons, M. Groussin, R. J. Xavier, and E. J. Alm, *Adaptive Evolution within Gut Microbiomes of Healthy People*, Cell Host Microbe **25**, 656 (2019).

[6] M. A. Lawson, I. J. O'Neill, M. Kujawska, S. G. Javvadi, A. Wijeyesekera, Z. Flegg, L. Chalklen, and L. J. Hall, *Breast Milk-Derived Human Milk Oligosaccharides Promote Bifidobacterium Interactions within a Single Ecosystem*, ISME J. **14**, 635 (2020).

[7] M. Roodgar, B. H. Good, N. R. Garud, S. Martis, M. Avula, W. Zhou, S. M. Lancaster, H. Lee, A. Babveyh, S. Nesamoney *et al.*, *Longitudinal Linked-Read Sequencing Reveals Ecological and Evolutionary Responses of a Human Gut Microbiome during Antibiotic Treatment*, Genome Res. **31**, 1433 (2021).

[8] B. A. Niccum, E. K. Kastman, N. Kfoury, A. Robbat, Jr., and B. E. Wolfe, *Strain-Level Diversity Impacts Cheese Rind Microbiome Assembly and Function*, Msystems **5**, e00149 (2020).

[9] A. Goyal, L. S. Bittleston, G. E. Leventhal, L. Lu, and O. X. Cordero, *Interactions between Strains Govern the Eco-evolutionary Dynamics of Microbial Communities*, Elife **11**, e74987 (2022).

[10] M. Imhof and C. Schlotterer, *E. coli Microcosms Indicate a Tight Link between Predictability of Ecosystem Dynamics and Diversity*, PLoS Genet. **2**, e103 (2006).

[11] O. Lukjancenko, T. M. Wassenaar, and D. W. Ussery, *Comparison of 61 Sequenced Escherichia coli Genomes*, Microb. Ecol. **60**, 708 (2010).

[12] S. Louca, L. W. Parfrey, and M. Doebeli, *Decoupling Function and Taxonomy in the Global Ocean Microbiome*, Science **353**, 1272 (2016).

[13] J. E. Goldford, N. Lu, D. Bajić, S. Estrela, M. Tikhonov, A. Sanchez-Gorostiaga, D. Segrè, P. Mehta, and A. Sanchez, *Emergent Simplicity in Microbial Community Assembly*, Science **361**, 469 (2018).

[14] M. Advani, G. Bunin, and P. Mehta, *Statistical Physics of Community Ecology: A Cavity Solution to MacArthur's Consumer Resource Model*, J. Stat. Mech. (2018) 033406.

[15] W. Cui, R. Marsland III, and P. Mehta, *Effect of Resource Dynamics on Species Packing in Diverse Ecosystems*, Phys. Rev. Lett. **125**, 048101 (2020).

[16] R. Marsland III, W. Cui, J. Goldford, A. Sanchez, K. Korolev, and P. Mehta, *Available Energy Fluxes Drive a Transition in the Diversity, Stability, and Functional Structure of Microbial Communities*, PLoS Comput. Biol. **15**, e1006793 (2019).

[17] M. Tikhonov, *Community-Level Cohesion without Cooperation*, eLife **5**, e15747 (2016).

[18] M. Tikhonov and R. Monasson, *Collective Phase in Resource Competition in a Highly Diverse Ecosystem*, Phys. Rev. Lett. **118**, 048103 (2017).

[19] M. R. Domingo-Sananes and J. O. McInerney, *Mechanisms That Shape Microbial Pangenomes*, Trends Microbiol. **29**, 493 (2021).

[20] O. X. Cordero and M. F. Polz, *Explaining Microbial Genomic Diversity in Light of Evolutionary Ecology*, Nat. Rev. Microbiol. **12**, 263 (2014).

[21] O. M. Maistrenko, D. R. Mende, M. Luetge, F. Hildebrand, T. S. B. Schmidt, S. S. Li, J. F. M. Rodrigues, C. von Mering, L. Pedro Coelho, J. Huerta-Cepas, S. Sunagawa, and P. Bork, *Disentangling the Impact of Environmental and Phylogenetic Constraints on Prokaryotic Within-Species Diversity*, ISME J. **14**, 1247 (2020).

[22] H. Wu and E. Moore, *Association Analysis of the General Environmental Conditions and Prokaryotes' Gene Distributions in Various Functional Groups*, Genomics **96**, 27 (2010).

[23] R. Starke, P. Capek, D. Morais, N. Jehmlich, and P. Baldrian, *Explorative Meta-analysis of 377 Extant Fungal Genomes Predicted a Total Mycobiome Functionality of 42.4 Million Kegg Functions*, Front. Microbiol. **11**, 143 (2020).

[24] C. O. Webb, D. D. Ackerly, M. A. McPeek, and M. J. Donoghue, *Phylogenies and Community Ecology*, Annu. Rev. Ecol. Systematics **33**, 475 (2002).
[25] A. Eng and E. Borenstein, *Taxa-Function Robustness in Microbial Communities*, Microbiome **6**, 1 (2018).
[26] K. Isobe, N. J. Bouskill, E. L. Brodie, E. A. Sudderth, and J. B. H. Mertiny, *Phylogenetic Conservation of Soil Bacterial Responses to Simulated Global Changes*, Phil. Trans. R. Soc. B **375**, 20190242 (2020).
[27] C. A. Serván, J. A. Capitán, Z. R. Miller, and S. Allesina, *Effects of Phylogeny on Coexistence in Model Communities*, bioRxiv, 10.1101/2020.09.04.283507.
[28] R. MacArthur and R. Levins, *The Limiting Similarity, Convergence, and Divergence of Coexisting Species*, Am. Nat. **101**, 377 (1967).
[29] R. MacArthur, *Species Packing and Competitive Equilibrium for Many Species*, Theor. Popul. Biol. **1**, 1 (1970).
[30] P. Chesson, *MacArthur's Consumer-Resource Model*, Theor. Popul. Biol. **37**, 26 (1990).
[31] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevX.12.021038 for a detailed description of the ecoevolutionary process in our model, a comparison of ecosystem coarse-grainability in terms of other example ecological properties and other technical details, including Refs. [32–35].
[32] D. T. Gillespie, *Exact Stochastic Simulation of Coupled Chemical Reactions*, J. Phys. Chem. **81**, 2340 (1977).
[33] E. L. Haseltine and J. B. Rawlings, *Approximate Simulation of Coupled Fast and Slow Reactions for Stochastic Chemical Kinetics*, J. Chem. Phys. **117**, 6959 (2002).
[34] J. G. Lopez and N. S. Wingreen, *Noisy Metabolism Can Promote Microbial Cross-Feeding*, Elife **11**, e70694 (2022).
[35] M. M. Desai and D. S. Fisher, *Beneficial Mutation-Selection Balance and the Effect of Linkage on Positive Selection*, Genetics **176**, 1759 (2007).
[36] B. H. Good, S. Martis, and O. Hallatschek, *Adaptation Limits Ecological Diversification and Promotes Ecological Tinkering during the Competition for Substitutable Resources*, Proc. Natl. Acad. Sci. U.S.A. **115**, E10407 (2018).
[37] A. Posfai, T. Taillefumier, and N. S. Wingreen, *Metabolic Trade-Offs Promote Diversity in a Model Ecosystem*, Phys. Rev. Lett. **118**, 028103 (2017).
[38] A. Altieri and S. Franz, *Constraint Satisfaction Mechanisms for Marginal Stability and Criticality in Large Ecosystems*, Phys. Rev. E **99**, 010401(R) (2019).
[39] P.-Y. Ho, B. H. Good, and K. Huang, *Competition for Fluctuating Resources Reproduces Statistics of Species Abundance over Time across Wide-Ranging Microbiotas*, Elife **11**, e75168 (2022).
[40] L. Fant, I. Macocco, and J. Grilli, *Eco-evolutionary Dynamics Lead to Functionally Robust and Redundant Communities*, bioRxiv, 10.1101/2021.04.02.438173.
[41] D. E. Dykhuizen and D. L. Hartl, *Selection in Chemostats*, Microbiol. Rev. **47**, 150 (1983).
[42] R. MacArthur, *Species Packing, and What Competition Minimizes*, Proc. Natl. Acad. Sci. U.S.A. **64**, 1369 (1969).
[43] R. Marsland, W. Cui, and P. Mehta, *A Minimal Model for Microbial Biodiversity Can Reproduce Experimentally Observed Ecological Patterns*, Sci. Rep. **10**, 1 (2020).
[44] H. Tettelin, V. Masignani, M. J. Cieslewicz, C. Donati, D. Medini, N. L. Ward, S. V. Angiuoli, J. Crabtree, A. L. Jones, A. S. Durkin *et al.*, *Genome Analysis of Multiple Pathogenic Isolates of Streptococcus Agalactiae: Implications for the Microbial "Pan-genome,"* Proc. Natl. Acad. Sci. U.S.A. **102**, 13950 (2005).
[45] H. Mickalide and S. Kuehn, *Higher-Order Interaction between Species Inhibits Bacterial Invasion of a Phototroph-Predator Microbial Community*, Cell Syst. **9**, 521 (2019).
[46] Although we use this as an example here, we should note that in the original reference [45] this statement is not actually meant as a species-level claim but indeed refers to the three specific strains used in the experiment.
[47] T. Taillefumier, A. Posfai, Y. Meir, and N. S. Wingreen, *Microbial Consortia at Steady Supply*, eLife **6**, e22644 (2017).
[48] G. Kinsler, K. Geiler-Samerotte, and D. A. Petrov, *Fitness Variation across Subtle Environmental Perturbations Reveals Local Modularity and Global Pleiotropy of Adaptation*, eLife **9**, e61271 (2020).
[49] S. Estrela, J. C. Vila, N. Lu, D. Bajic, M. Rebolleda-Gomex, C.-Y. Chang, J. E. Goldford, A. Sanchez-Gorostiaga, and A. Sanchez, *Functional Attractors in Microbial Community Assembly*, Cell Syst. **13**, 29 (2022).
[50] R. M. May, *Will a Large Complex System Be Stable?*, Nature (London) **238**, 413 (1972).
[51] R. May, *Stability and Complexity in Model Ecosystems* (Princeton University, Princeton, NJ, 1973).
[52] D. Šiljak, *When Is a Complex Ecosystem Stable?*, Math. Biosci. **25**, 25 (1975).
[53] A.-M. Neutel, J. A. Heesterbeek, and P. C. De Ruiter, *Stability in Real Food Webs: Weak Links in Long Loops*, Science **296**, 1120 (2002).
[54] N. Rooney, K. McCann, G. Gellner, and J. C. Moore, *Structural Asymmetry and the Stability of Diverse Food Webs*, Nature (London) **442**, 265 (2006).
[55] U. Brose, R. J. Williams, and N. D. Martinez, *Allometric Scaling Enhances Stability in Complex Food Webs*, Ecol. Lett. **9**, 1228 (2006).
[56] S. Allesina, J. Grilli, G. Barabás, S. Tang, J. Aljadeff, and A. Maritan, *Predicting the Stability of Large Structured Food Webs*, Nat. Commun. **6**, 1 (2015).
[57] G. Bunin, *Interaction Patterns and Diversity in Assembled Ecological Communities*, arXiv:1607.04734.
[58] M. Barbier, J.-F. Arnoldi, G. Bunin, and M. Loreau, *Generic Assembly Patterns in Complex Ecological Communities*, Proc. Natl. Acad. Sci. U.S.A. **115**, 2156 (2018).
[59] D. I. Bolnick, P. Amarasekare, M. S. Araújo, R. Bürger, J. M. Levine, M. Novak, V. H. Rudolf, S. J. Schreiber, M. C. Urban, and D. A. Vasseur, *Why Intraspecific Trait Variation Matters in Community Ecology*, Trends Ecol. Evol. **26**, 183 (2011).
[60] S. Louca, S. M. Jacques, A. P. Pires, J. S. Leal, D. S. Srivastava, L. W. Parfrey, V. F. Farjalla, and M. Doebeli, *High Taxonomic Variability Despite Stable Functional*

Structure across Microbial Communities, Nat. Ecol. Evol. **1**, 1 (2016).

[61] C. Burke, P. Steinberg, D. Rusch, S. Kjelleberg, and T. Thomas, *Bacterial Community Assembly Based on Functional Genes Rather than Species*, Proc. Natl. Acad. Sci. U.S.A. **108**, 14288 (2011).

[62] B. Segerman, *The Genetic Integrity of Bacterial Species: The Core Genome and the Accessory Genome, Two Different Stories*, Front. Cell. Infect. Microbiol. **2**, 116 (2012).

[63] R. G. Beiko, *Microbial Malaise: How Can We Classify the Microbiome?*, Trends Microbiol. **23**, 671 (2015).

[64] P. Vernocchi, F. Del Chierico, and L. Putignani, *Gut Microbiota Profiling: Metabolomics Based Approach to Unravel Compounds Affecting Human Health*, Front. Microbiol. **7**, 1144 (2016).

[65] M. R. Wilson, L. Zha, and E. P. Balskus, *Natural Product Discovery from the Human Microbiome*, J. Biol. Chem. **292**, 8546 (2017).

[66] J. Bergelson, M. Kreitman, D. A. Petrov, A. Sanchez, and M. Tikhonov, *Functional Biology in Its Natural Context: A Search for Emergent Simplicity*, eLife **10**, e67646 (2021).

[67] J. Moran and M. Tikhonov, *Defining Coarse-Grainability in a Model of Structured Microbial Ecosystems*, Mendeley Data 10.17632/gpfw4y46c8 (deposited 10 March 2022).

[68] https://github.com/tikhonov-group/ecoevocrm.