Disordered and Partially Structured Models in Community Ecology: What are they? And how do we use them?

Mathew Leibold[*], Matthieu Barbier[†‡]

May 15, 2025

Abstract

Community ecology describes how species interact with each other and with their environment. In nature, processes can be very complex because they involve hundreds to thousands of species interacting with each other in complex environmental landscapes. Classical approaches that have provided key insights have largely focused on the study of tractable subsets of species and patches, but these do not always adequately address the wider scope of natural complexity. Alternate approaches that use specific parameters and/or that use simulations to study such highly diverse systems are problematic because they can become very detailed, system-specific, and easily divorced from general principles. Finally, 'minimal' models to explain data exist (e.g. null models, 'neutral theory' and 'entropy based' models), but they often do not provide adequate connections to experimental or mechanistic studies and results. Here we describe and discuss an alternate approach that seeks to link basic processes of community assembly (environmental heterogeneity, species interactions, dispersal, and stochasticity) with each other using 'disordered systems models' to make robust predictions about community structure, albeit without the detail of more system-specific approaches. We describe the logic of the approach, outline the methods involved, and identify important limitations. We also describe how this approach can be expanded to better incorporate additional nonrandom structure (such as intercorrelated parameters) in these basic processes and leading to 'partially structured models', and we introduce the idea that this could also be applied to metacommunities. Although implementing this approach in empirical studies will still be quite challenging, these approaches reduce the complexity of the overall problem by orders of magnitude, making it a promising approach to improve the study of biodiversity in realistic landscapes.

---

[*] University of Florida, Gainesville, Florida, USA. mleibold@ufl.edu

[†] CIRAD, UMR PHIM, 34090 Montpellier, France. contact@mrcbarbier.org¿

[‡] PHIM Plant Health Institute, University of Montpellier, CIRAD, INRAE, Institut Agro, IRD, 34090 Montpellier, France.

Introduction

Matthieu's proposal: I'd like to very clearly highlight, in a place that people won't miss, e.g. at the very start, what the purpose of this manuscript is and what is its relationship to other works that provide more technical details: Akjouj for mathematicians, my tutorial for modelers (REF impending), Cui and Mehta for a strong focus on random consumer-resource models.

Community ecology (and its relation to ecosystems ecology) is inherently a science about complexity. This important idea is captured by Darwin's 'entangled bank' that describes the complex nature of interspecific interactions and environmental factors (and, in Darwin's mind, their impacts on evolution). The nature of this complexity is exemplified by starting with simple approaches (e.g. exponential population growth) and gradually adding complexity (e.g. logistic growth, pairwise direct interactions, dispersal, etc.) that eventually reaches arbitrary levels of complexity. To a substantial degree, this approach describes the historical legacy of community ecology today.

An important step in illustrating the consequences of complexity involves the idea of indirect interaction effects among species, as described by Darwin's metaphor. In models that involve a small number of species (that is, 'modules' *sensu* Holt and Hochberg (2001), see also McPeek (2021)), the consequences of such indirect effects can be resolved, so that the overall 'net effects' can be derived from direct effects, resulting in concepts such as trophic cascades (Hairston et al., 1960) and apparent competition (Holt, 1977). However, building up from this to many more species becomes increasingly problematic, both theoretically and empirically, since the number of interactions increases as the square of the number of species. This approach effectively reaches its limit when the number of species starts to exceed five or so (McPeek, 2021).

Under some conditions, it is possible to get some insight about these net effects in highly diverse systems (Levine, 1976; Lawlor, 1979) but these insights can be strongly constrained because of the delicate context dependence that determines these net effects even when the direct effects are fixed (Schaffer, 1981). How can we realistically investigate the community ecology of ecosystems that involve many species (i.e., dozens, hundreds, or more)? Here we describe an approach based first on 'disordered systems models' (DSM) that can help us understand generic effects, and then on 'partially structured models' (PSMs) that further resolve possible deviations from DSMs in ecologically interesting ways. These approaches seek to provide robust insights about highly diverse communities and metacommunities without requiring completely detailed knowledge of the species and ecosystems involved, and thus bypassing the need to understand the detailed complexity of the 'entangled bank'. Although we describe some of the mathematics involved, we aim for a more general understanding that does not require substantial mathematical expertise and emphasize how this approach can be utilized by empirical ecologists. The mathematics involved can be found in the references we provide and, especially, in Bunin (2017); ?); ?.

To provide context for our claims, we imagine at least three plausible scenarios that can justify different possible approaches to the study of community ecology in the case of highly diverse communities and metacommunities:

1. One possibility is that much of this complexity is not important at all, and that there are only a few important components (e.g. species, traits or functional groups) to the dynamics. If so, we can justify focusing on those few particular details. We can call this approach 'classical simplicity', since the study of such key mechanisms is a traditional approach in community ecology, characterized by the 'modules' approach described above.
2. Another possibility is that many, or all, details affect community dynamics in a deeply inter-correlated and interconnected way, which cannot be easily dissected (or even dissected at all!). We posit that such a situation is fragile by definition: altering any single component of the community can impact all the others in major ways, somewhat like a computer program where removing one character in the code will cause the whole program to stop working. We can call this 'true complexity'. We acknowledge that it is possible that communities in nature are, in fact, characterized by such complexity. If so, the study of community ecology would be enormously challenging (Lawton, 1999). We question, however, if such ecosystems are likely to be sufficiently closed and co-evolved to allow the emergence and persistence of such fragile structures and, consequently, whether these structures are likely to be leading forces shaping the aggregate behavior of many entire ecosystems.
3. Finally, it is possible that many (or all) of the factors involved matter roughly equally and influence communities in heterogeneous but weakly correlated directions. If so, their combined impacts will often cancel each other out, so as to produce variation that resembles random variation, at least with respect to the ways in which they alter many of the resulting patterns, especially those that represent aggregate features of the ecosystem. Consequently, such aggregate properties can persist as important causes and consequences of the statistical moments of these details (e.g. mean and variance of interactions) rather than being detail-dependent as in 2 above. This perspective is the one we base our arguments about disordered (and partially ordered) models, and we call this "emergent simplicity".

If (meta-)communities actually show 'true complexity' , we may expect most of the variation within metacommunities to appear to be highly context-dependent and have little resolvable structure. Recent developments in parsing metacommunity variation indicate that this is not the case (Leibold and Peres-Neto in prep), but rather that the observable variation can be related to some mixture of environmental, spatial, and among-species covariation. While resolving the effects of these factors is difficult, this provides evidence that 'true complexity' is not likely, so that some aspects of the complexity could well be understood as 'emergent simplicty.'

The study of emergent simplicity is made possible by the careful invocation of randomness and probabilistic modeling. More specifically, we can identify two kinds of "randomness" that affect how we think about these models: "Stochasticity", i.e. random fluctuations over time (like environmental perturbations or stochastic demographic effects), and "disorder" that describes the apparent and effective randomness that results from the statistical behavior of diffuse causal networks of otherwise complex but intermingled deterministic fixed processes. Of course, there is information that is lost when we lump these "disordered" deterministic processes by using statistical descriptions of the parameters instead of the details. The problem is that identifying and quantifying the detailed information can be very difficult when there are more than a handful of species. Additionally, doing so can also lead to errors when we don't have precise knowledge of the parameters and functions involved to make the detailed behavior of the ecosystem. Instead, here, we look to identify and focus on patterns that are likely to be robust to these precise details, while acknowledging that our models will not make predictions about other patterns that are not as robust.

We describe the general approach with models that characterize 'classical simplicity' in the form of pairwise Lotka-Volterra models (we briefly discuss other options in the prospectus of this manuscript). Following May (1972) (see also Novak et al. (2016)), we then imagine that a large number of such pairwise Lotka-Volterra equations can be combined into a larger model that can be expressed in matrix form called the community matrix (we note that the community matrix is typically used to describe local communities, but that this matrix can be expanded to describe metacommunities as well (Gravel et al., 2016)). The key step in taking a disordered modeling approach is then to replace the actual community matrix with a representation that is only based on means and variances for the various components (figure 4). In practice this characterization is only valid if there are a sufficient number of species so that the means and variances of the matrix entries are sufficiently well characterized so as to be meaningful; e.g. one or several dozen species. We note that the modeling of intermediate-size systems (say, 5-12 species) remains a daunting challenge.

1 The strategy of modeling using DSMs and PSMs

Taking an approach that is based on DSMs (and subsequently on PSMs see below) involves a certain strategy to modeling that differs from the more conventional approaches in 'classical simplicity' or 'true complexity'. Rather than settling on a particular model structure and just 'filling in' parameters, we have to think carefully about which model components we want to aggregate and simplify using means and variances, and the types of patterns we hope to study. The strategy also differs from modeling approaches based on 'minimal' models, such as null models, neutral theory, or entropy-based approaches that start with data and ask what model

(a) General picture

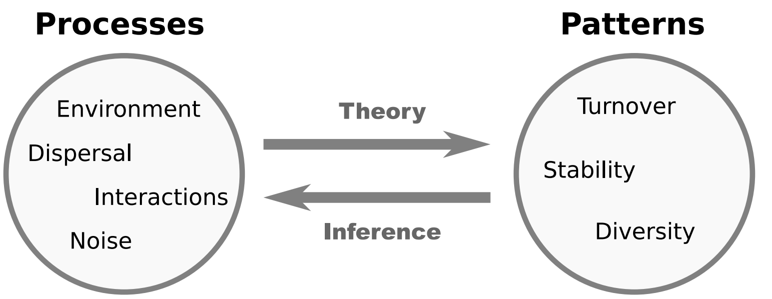

(b) Details of the modelling process

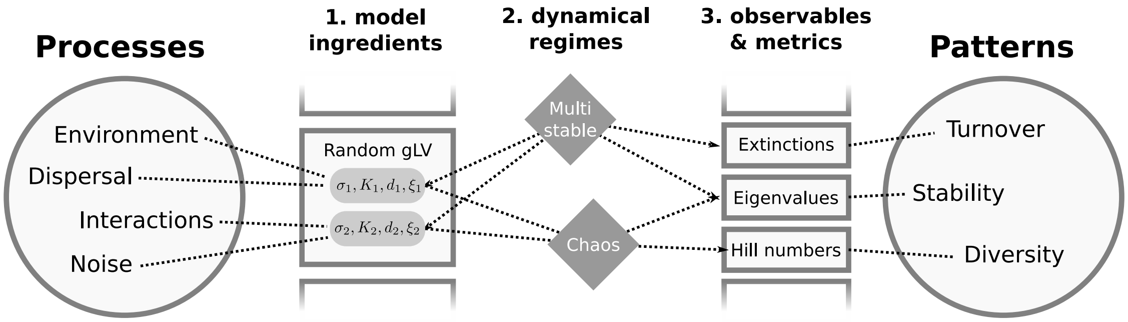

Figure 1: a) Our overall goals is to relate basic ecological processes to community level patterns. b) To do so, we break the process down into four parts: the model ingredients (Sec. 1), the different regimes, dynamics or behaviors it can display (Sec. 6) and the 'observables' chosen to characterize empirical patterns (Sec. 2) that we aim to predict from the model and use to validate it and infer its parameters. This tripartition is useful since different ingredients can produce the same behavior (e.g. many mechanisms can lead to chaotic fluctuations in abundances), and different behaviors can produce the same pattern (e.g. intrinsically chaotic dynamics or dynamics driven externally by environmental fluctuations can lead to similar-looking time series), see also McGill (2010).

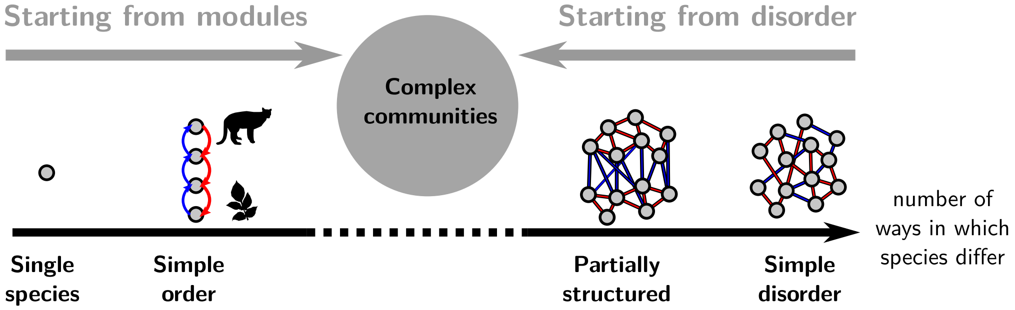

Figure 2: DSM and PSMs. When species differ importantly in just a few ways, e.g. trophic height, we can organize them in simple modules like food chains. When they differ in a vast number of ways, their interactions become largely unpredictable and, in effect, random-like. Real communities probably lie somewhere in-between, and can be approached from both sides, e.g. with Partially Structured Models.

structures can account for them, typically looking for minimally parameterized models. This is because we are interested in maintaining a strong connection to the fundamental dynamic processes and mechanisms of community dynamics.

Instead the intermediate perspective we describe asks how we can identify and focus on patterns that can be linked to processes and mechanisms in a robust way that minimally depends on detailed parameterization of complex models. We try to structure our models and identify patterns that are insensitive to factors other than those we chose to include in our approach based on 'emergent simplicity'. This means that many differences in model selection, even striking ones such as being discrete or continuous, deterministic or probabilistic, spatial or non-spatial, can be treated as irrelevant if/when they do not impact our selected target outcomes. One way of articulating this idea is: all details matter, but they do not necessarily matter for what we might care about.

Consequently, we focus on what might be called 'macroscopic' rather than detailed 'microscopic' model components (and applying this perspective to patterns as well as processes). Macroscopic predictions are thus likely to result from the aggregate behavior of model components, in the same way that the ideal gas law is a robust outcome of the complex behavior of many colliding molecules. It turns out, in fact, that trying to increase microscopic fidelity by adding more model ingredients does not always improve aggregate accuracy, and can even detract from it if poorly estimated or biased parameters are added. Hence, we want to think of models not in terms of how they fit the details of our microscopic intuition, but in terms of addressing the following issues:

- What kind of macroscopic (i.e. aggregate) behaviors and patterns can they robustly generate?
- Are these interesting to us?
- Can we adjust model parameters so that they can quantitatively match observed patterns?
- Can parameters that are fit using some patterns also correctly predict other patterns? − in which case the model is a workable description of macroscopic reality, even if its details turn out to be untrue to microscopic reality.

1.1 Why the Random Lotka-Volterra model?

To illustrate this approach, we focus on one basic model, the Random Lotka-Volterra (RLV) model (Bunin, 2017). This consists of using the well-known Lotka-Volterra equations for a diverse set of interacting species, and then assigning the parameters (i.e. intrinsic growth, carrying capacity, interaction coefficients, and dispersal) randomly among the species based on simple statistical features (e.g. means and variances). We do this because the RLV can be viewed as a highly simplified model that gives a first-order approximation to the most important and/or robust generic properties of more complex models. As such, it can serve as a 'baseline' model for understanding how arbitrarily complex systems might be regulated. We argue that they should be further elaborated, if need be, using Partially Structured Lotka-Volterra (PSLV) models as we describe later in section 3.

We also chose to build from the Lotka-Volterra model because it is a well-known generic model that has already been used in this way. However, there are alternate options that could be used for studying community and metacommunity dynamics (e.g. ecosystem-based models that describe stocks and fluxes of materials and energy, community assembly models based on colonization-extinction dynamics that ignore species abundances, etc.). At this point, we are unaware of published works that explicitly study such models with an explicit DSM formulation akin to the one we use here.

Although it seems relatively simple, the RLV model can span a very broad scope of behaviors, including single-point stable equilibrium, oscillations, alternate stable states, and more complex fluctuating dynamics such as high-dimensional and chaotic dynamics. This is because the entanglement of causal processes in the RLV can be extremely high and lead to behaviors that encompass the whole array of dynamics characterized by more complex non-linear models.

We formulate the RLV in a way that captures the idea that community assembly and community dynamics are driven by five basic processes (modified from Vellend (2010, 2016)). These consist of 'density-independent selection' that favors different species at a given site as a function of their local intrinsic population growth and carrying capacity, 'density-dependent selection' that accounts for feedback among component species, 'dispersal' that accounts for the movement of individual, 'stochasticity' that describes how chance events such as demographic drift and disturbances affect each species, and 'novelty' that describes how *de novo* species originate (here we however ignore this speciation process in our model since it is generally much slower than the others).

At the most general level, we use a spatialized Lotka-Volterra model that separately accounts for these processes by describing each species as having growth rates r_i, interactions A_{ij} (in which the carrying capacity is thus given by $K_i = 1/A_{ii}$ for intra-specific interactions), stochastic perturbations ξ_i, and dispersal among patches x and y with rates D_i

$$\frac{dN_i(x,t)}{dt} = r_i N_i(x,t)\left(1 - \frac{\sum_j A_{ij} N_j(x,t)}{K_i(x)}\right) + \xi_i(x,t) + \sum_y D_i\left(N_i(y,t) - N_i(x,t)\right) \tag{1}$$

The most basic version of the Lotka-Volterra model, without stochasticity or spatial fluxes ($\xi_i = D_i = 0$), allows at most one stable state per set of surviving species: for a given set Ω, assuming they can coexist at equilibrium, their abundances must verify:

$$0 = 1 - \sum_{j \in \text{survivors}} A_{ij} N_j^* \qquad \text{for all } i. \tag{2}$$

Such a linear system of equations can only have one solution (a different system and a different solution for each possible set of species Ω); however, that solution need not be a stable equilibrium (it could be an unstable fixed point, e.g. in case of a limit cycle or bistability between other fixed points), as we discuss in Sec. 6.

Populations tend to either grow or decay exponentially when rare, but this exponential growth/decay is impacted by various additional factors. Here, we further assume that we can separate the additional effects of these factors on instantaneous growth into contributions from a) density-dependent interactions (A_{ij}) among the species, and frequency-independent factors r_i and ξ_i which represent fixed and fluctuating impacts of an environment. We note however that these contributions of environment and interactions, while they are separated and simply additive here in short-term growth, become unavoidably entangled in their long-term effects on abundances (Barbier et al., 2022)

Spatial fluxes can be approximated by simple "diffusion-like" dispersal between adjacent patches, which is enough to qualitatively capture many of the consequences of accounting for space. We may also assume that growth rates r_i or interactions A_{ij} depend on spatial location to introduce spatial heterogeneity. We also note that, here, interactions only happen within one locality, thus the community (or site) scale can be defined as the spatial scale over which species interactions can be modeled as well-mixed – we do not consider here systems where, for instance, different species might interact over vastly different scales e.g. (Zelnik et al., 2021).

For the Random Lotka Volterra model, we proceed as follows:

1) We hypothesize a 'species pool' of size S (the number of species present that can potentially participate in community dynamics).

2) We assign parameters of the L-V equations at random to each of the species in the species pool; we describe useful ways to do this so as to optimize the study of the RLV model below.

3) We try to use analytical methods to derive long-term expected outcomes of the solution to the RLV model. We focus on 'aggregated' properties that reflect community-wide attributes rather than predictions about individual species because these are the most likely to be 'robust'. 4) If we cannot find analytical solutions, we use numerical or simulation approaches.

5) We compare the resulting outcomes to natural patterns (when possible). If we find that they match

reasonably well, we claim that the RLV is sufficient to explain such patterns.

6) If this comparison does not match adequately well, we work to develop modifications of the RLV model by using PSLV approaches that we describe later in section 3.

We use the RLV to study how it can predict a number of (meta) community properties that may be robust to the details that are omitted in this approach. For example, Barbier et al. (2018) show that under some conditions (mostly that interspecific interactions aren't too strong), purely local RLVs can predict properties such as the following:

a) local species richness: the size of the subset of species that co-occur (or co-exist) in local patches. It is often more convenient to express this as the ratio of local to global species numbers.

b) local evenness in the relative abundances of these species.

c) total abundance (or total biomass, depending on how the RLV is formulated) of organisms d) mean turnover (or production/loss) of organisms (or biomass)

e) mean local CV of species in time.

1.2 Applicability of disorder

A simple objection to the use of randomness is that species typically seem to be much more different from each other than simple random draws from a distribution. To clarify this point, we can think about *degrees* of heterogeneity:

- "variance-like" heterogeneity: quantitative variation within a group of species that can be adequately quantified by statistical moments (e.g. a mean and a variance) of their parameters, recognizing that we will expect to occasionally see some species that lay at the extremes of the distribution.
- "systematic" heterogeneity: irreducible qualitative differences among such groups (e.g. autotrophs versus heterotrophs).

Disorder is clearly applicable to "variance-like" heterogeneity[1]hereas "systematic" variance is more naturally suited for partially structured models (described below). Even so, an important lesson from experience is that differences that appear irreducible and fundamental can **sometimes** be ignored, and entities that seem like they belong to different classes can **sometimes** be bunched together for practical purposes. For example, even species that are very different on a *per capita* basis, such as bacteria & whales, may be more comparable in terms of biomass (as demonstrated e.g. by allometric relations holding over many classes of organisms (Hatton et al., 2021) at least for the purpose of understanding certain macroecological patterns.

---

[1] w

Consequently, whenever we think about using randomness in a model, we try to find the representation of the system that makes species as equivalent as possible so that their parameters can be represented by their 'variance like' heterogeneity.

To illustrate, we start with a simple theoretical example, using the non-spatial, unperturbed version of the equations (1) i.e.

$$\frac{dN_i}{dt} = r_i N_i \left(1 - \frac{N_i - \sum_{j\neq i}^{S} A_{ij} N_j}{K_i}\right)$$

The carrying capacity K_i is the equilibrium abundance in the absence of other species. Let us assume for now that we are looking at a competitive system where all species have a carrying capacity $K_i > 0$ (though everything that follows can be extended to e.g. trophic systems where predators cannot survive without prey).

We could then change variables to $\eta_i = N_i/K_i$, usually called the relative yield (see references and discussion in Barbier et al. (2021)), and still have a Lotka-Volterra model:

$$\frac{d\eta_i}{dt} = r_i \eta_i \left(1 - \eta_i - \sum_{j\neq i}^{S} \alpha_{ij} \eta_j\right) \tag{3}$$

with different interaction coefficients ($\alpha_{ij} = A_{ij} K_j / K_i$). If the A_{ij} are independent random parameters, then the α_{ij} are not independent (and conversely): they are correlated by row and by column due to the factors K_i. To apply disordered systems theory, the most appropriate choice is the one where parameters (α, or K and A, or others) are as uncorrelated as possible. We note that this is not necessarily the most intuitive representation, and its variables might not be as interpretable as abundances, biomasses or relative yields (see e.g. Spaak and De Laender (2020)).

Hereafter we will use terminology suggesting that our variables N_i represent abundances, but it is crucial to recognize that this will not always be the right choice; instead, they should be understood as a general proxy for a species' importance in community dynamics, perhaps expresses in biomass units or consumption rates etc. When we see apparent inequalities between species, we should first ask how much is really due to some species being more or less important, or whether we can restructure out model to capitalize on expressing it in other units.

2 What can be predicted, inferred and tested empirically

Above, we emphasized that disordered systems theory is focused on deriving robust predictions from dynamical models that are parameterized using a statistical approach to describe their parameters. What kinds

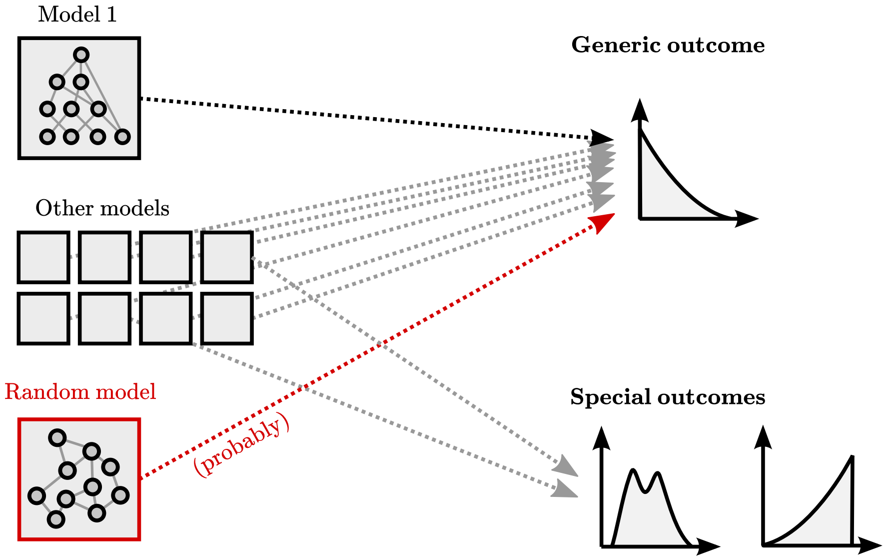

Figure 3: We may use random models because they are a simple way to access ‘robust’ or ’generic’ outcomes that we expect to arise in many nonrandom models (Barbier et al., 2018). Drawing parameters at random is thus different from ’exploring all the possibilities’ – on the contrary, when the number of parameters becomes large, the model is unlikely to stumble upon any non-generic possibility.

of data are likely to be robust in this way and what kinds of patterns can we predict about them from the disordered systems approach?

Communities can be described by an amazing variety of attributes (or ‘observables’), that can be analyzed for possibly regular arrangements; means, correlations, etc. that we can generally call ‘patterns’. As we pointed out earlier, many of these patterns can be very complicated in their details, for example, by often being very species-specific in nature and these are unavailable to study by our approach since we ignore species identities when we replace their attributes as drawn from statistical distributions. Many of these patterns, especially the more detailed ones, are also likely to be ‘fragile’ in that they depend on specific detailed assumptions that may or may not actually hold. For example, pairwise correlations in species abundances seem like they could be linked to pairwise interactions between them, but this is generally not the case due to the complex intertwined interactions involving multiple (and potentially all) species in the system. Schaffer (1981) showed that even if they could be measured, these correlations did not have any clear link to the pairwise interactions among the species, illustrating the fragility of relating patterns to the driving processes in the complex community.

By taking a disordered systems approach to community ecology, we accept that many possible patterns involving observable outcomes are likely to be fragile in this regard, and we try to focus on those that are not; see Figure 3. As discussed in earlier sections of this paper, the models we create sacrifice the need for detailed parameters (or predictors), which are often difficult to quantify anyway, and look for more ‘macroscopic’ predictors that can work equally well in making predictions about these more robust patterns.

Next, we discuss one of the most robust properties of RLV models, the existence of qualitatively distinct dynamical regimes, which informs which patterns we can study. We then discuss more extensively how and why we may select robust yet informative patterns to fit our model quantitatively.

2.1 Dynamical regimes

Contrary to simple statistical models (e.g. a linear relation or a normal distribution) where varying parameters will only lead to quantitative changes in predictions, dynamical models like RLV can transition between qualitatively different behaviors depending on their parameters, such as a stable equilibrium or limit cycles.

These regimes can be associated with entirely different patterns, rather than just quantitative differences: for instance, an oscillatory regime affords new patterns such as cycle period. Different regimes of the same model often entail different trends of pattern metrics with control parameters (e.g. diversity may increase then decrease as we keep increasing dispersal), which has important consequences for fitting. This is typically associated with a change in which terms dominate in the dynamical equation (for instance, dispersal rate may go from being small compared to extinction rate, to being large compared to it (Barbier et al., 2022)). But this is not always sufficient to distinguish meaningfully different regimes – for instance, strong competition may lead to either chaotic fluctuations or multiple stable states depending on slightly more specific features of that competition.

As shown in the Box ‘Diagram of dynamical regimes’, RLV models tend to have long-term behaviors that fall into a few broad classes; a single steady state point equilibrium, bounded fluctuations within a given domain or multiple outcomes that are “alternate attractors” (involving multiple possible combinations of either bounded fluctuations or steady-state points).

The existence of these distinct regimes, and how they relate to community-level parameters, is, in itself, a prediction amenable to empirical testing, as it is particularly robust to model details, and has recently received direct experimental support (Hu et al., 2022).

2.2 Selecting robust yet informative patterns

It is clear that we can only use our approach to try to predict ‘macroecological’ patterns that represent aggregated features and cannot predict more specific ones, see e.g. McGill (2010); Grilli (2020). Nevertheless, this still results in a wide array of possible features of the system that we can hope to study. The patterns of greatest interest should have two properties. First, they should be robust to the details we have omitted. We are thus not interested in patterns that are likely to be highly contingent on particular details, even if we can consider them as highly aggregated properties. This has been a criticism leveled against any specific measure of community stability, asymptotic resilience, where the rarest species tends to play a key role (Arnoldi et al., 2018). Secondly, we seek patterns that are likely to be sensitive to the driving variables we include (i.e. the means and variances of the parameters). We imagine that some patterns are so robust that they are not informative about anything. For example, McGill (2010) showed that properties such as the local Species Abundance Distribution (SAD) and the Species Area Relationship (SAR) are qualitatively (and perhaps to some degree quantitatively) robust to model formulations involving very different parameters. In this case, these patterns are consequently rather uninformative. Instead, we seek patterns that have some intermediate status so as to be robust to the details we ignore but sensitive to those we include.

Identifying these patterns can be an important step, although it can be challenging. In some cases, it is possible to use analytical models to identify them. In this case, there is an explicit formulation of both the robustness and characterization of the sensitivity (which patterns are sensitive to which parameter types). For example, Bunin (2017) used a technique called 'cavity modeling' to derive explicit equations relating the total biomass, total biomass turnover, species richness, species eveness, and population stability (figure 4), from a disordered model driven by variation in the carrying capacities of the species, the mean, variance, and mean asymmetry of interaction coefficients, and the total diversity of the species pool. This allowed them to specify just how the various patterns should relate to each of the driving variables.

Although optimal for insights into the causal effects in a system, such analytical solutions are often difficult or impossible to obtain. Instead, the search for relevant patterns will probably require simulations. Additionally, or alternatively, these patterns can be inspired by the interests of the researcher, perhaps based on previous knowledge such as might be obtained from previous work using e.g. 'classical simplicity' or previously documented macroecological patterns.

An important use of DSMs is to predict the generic effects of biodiversity. We may thus think about patterns observed when biodiversity is experimentally controlled. Among those, the relationship between diversity and total biomass provides an estimate of average competition (Fort, 2018; Wang et al., 2021), while a positive relationship between biodiversity and temporal variability would be indicative of a fluctuating dynamical regime (Roy et al., 2020; Hu et al., 2022). However, the informativeness of these patterns may break down when they are not experimental but correlative, i.e. when diversity is not directly manipulated.

So, why we should try to identify such 'robust-but-informative' patterns anyway? We can think of at least three reasons to seek to make such predictions:

1) We seek to make these predictions for their own sake. This might be useful in forecasting where we can ask questions about what could happen under different conditions and might be useful in predicting effects of e.g. climate change, eutrophication, repeated invasions, fragmentation etc. Under this scenario we wish

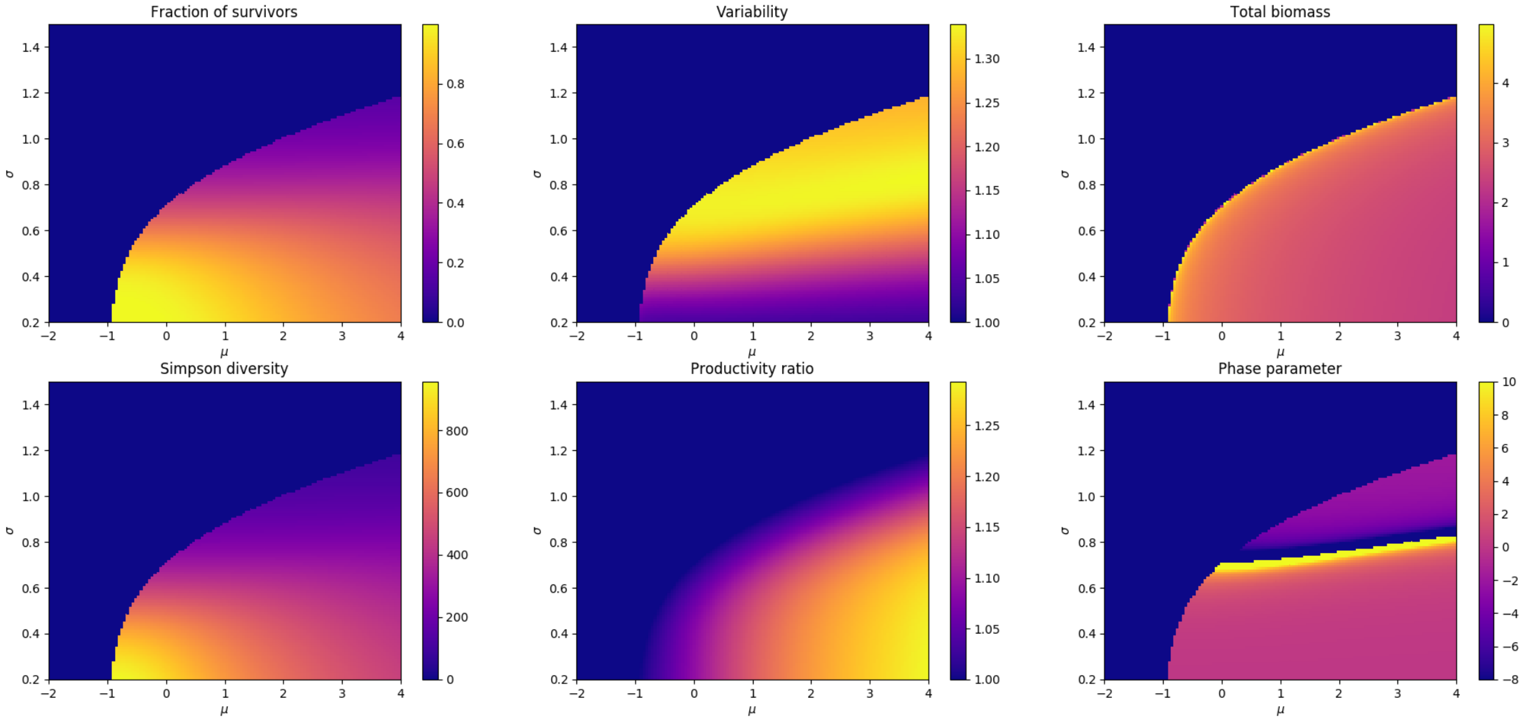

Figure 4: Theoretically predicted coexistence, abundance and stability properties for the assembled state of the random Lotka-Volterra model, as we vary $\mu = S\langle A_{ij}\rangle$, $\sigma = \sqrt{S\mathrm{var}(A_{ij})}$. The uniform area in the left of each graph signals the parameter region where abundances diverge (breakdown of the LV model due to strong mutualism). The bottom-right graph showcases the phase parameter: the sharp line where it diverges indicates the transition from the single-equilibrium regime (below the line) where our analytical results are exact for the reference model, to the chaotic regime (above the line) where they are approximate. Reproduced from Barbier et al. (2018).

to work from left to right in figure 2. We might be particularly interested in predictions that are robust to possible details since there are more likely to be reliable, see Figure 3. This approach may also, at least to some extent, allow us to extrapolate predictions for situations that lie outside the range of current variation since the robustness implies that they can do so unless there are fundamental changes to the structure of the models and when they might apply.

2) We seek to infer information about the predictor variables that we have modelled. In particular, measuring the carrying capacities of species, and even more so, the pairwise interactions coefficients is increasingly tedious as the number of species increases. Here we know we won't be able to obtain detailed information on these parameters, but there may be reasons to want to know how their means and variances could affect patterns. Under this scenario we would be working from right to left in Figure 2. Barbier et al. (2022) show that this can be done using a disordered systems model and further show that this is robust to the amount of dispersal in the system.

3) We seek use this approach in a hypothetico-deductive context to test or derive hypotheses about community dynamics (and metacommunity dynamics described below) by evaluating if predictor variables are consistent with observable patterns. Under this scenario we would be looking to see if the two approaches above are internally consistent (jointly working in both directions in figure 2). This is a substantially more ambitious question, especially in the absence of additional information (e.g. manipulative experiments, structured use of trait variation etc.), but at least it would allow us to apply the approach within a hypothetico-deductive research framework.

Barbier et al. (2018) found numerically that various straightforward aggregate predictions of DSMs were robust to a number of variations in fine-grained details of network structure, but not always to variations that could be interpreted as large-scale structures, e.g. when the network of interactions was either nested, directed, or bipartite. This suggests that it might be important to incorporate some general features of the networks as additional constraints on the purely disordered systems model, something that one might call 'partially structured' models; we describe and discuss these in a following section.

3 Partially structured models

Above we have focused on 'fully' disordered models, but of course this often seems like a strong constraint, perhaps applicable only to interactions within a guild of similar species (e.g. grasses, or soil bacteria (Barbier et al., 2021; Hu et al., 2022)). Often we can identify important aspects of community variation that we think might be important but find it difficult to formulate our model to have narrowly constrained random distributions for our parameters as described above. For example, if we wish to model food webs with distinct trophic levels, it seems highly problematic to imagine that interactions are adequately characterized by a simple mean and variance across all trophic levels. Instead we would be inclined to think that the predator-prey interactions between trophic levels would have very distinct parameter distributions (means, variances, asymmetries) than the interactions within any of the trophic levels. We can imagine the same sort of problem for any so-called set of 'functional groups' such as plant-pollinator systems. And one can also hypothesize other ways that interactions might be partitioned into different components (e.g. within and between habitats or in relation to the size of organisms). It seems extremely unlikely that we can entirely dismiss such empirically ubiquitous structures and often these seem particularly apparent to empirical ecologists who have documented the effects of such factors.

Is there any way to extend the insights we describe above to account for such 'functional groups' (for example, trophic levels)? This issue is closely related to the way food webs have classically been defined since such food webs often (perhaps always!) lump at least some species together based on functional considerations. For example, many food webs use the concept of 'trophospecies', defined as groups of species that qualitatively share the same resources and predators. Although this is done largely for practical reasons (the detailed differences among species in trophospecies is generally unknown, and formidable to study) this is also often justified by the idea that interactions within such groups are relatively uniform and that shared interactions with other groups are also relatively uniform even if they are distinct from each other and from the within group interactions. Extending disordered models in this way (within vs among groups), is what is meant by 'partially structured' (and thus also 'partially disordered') models.

Barbier et al. (2018) suggested that this could be done and illustrated how it might modify predictions from those based on the completely disordered approach. The idea is essentially to partition interactions into within and among functional groups and characterize each of these separately by their means and variances. If this approach can be justified and validated, it would provide a useful intermediate perspective between disordered systems models and fully specified ones. This seems like a conceptually straightforward, and to most community ecologists, meaningfully useful extension of disordered systems models. In fact, in the absence of such structure, many ecologists would feel that they may not even have a job!

Adding structure means adding parameters, typically ending up with as many parameters as classical few-species ecological models (e.g. tritrophic chains), or even more (since the presence of disorder means that we care not only about average interaction between e.g. trophic levels, but also about the (co-)variance of these interactions). Consequently, deciding just how interactions should be approximated by a combination of disorder and structure involves careful thinking. We discuss below how and when we expect a gain from this approach.

3.1 Types of structures

We can start by thinking of three possible ways to structure interactions with some degree of heterogeneity: keystone species or links, groups or modules, and continuous axes. All of these situations can be theoretically understood in simple settings with a slight increase in complexity from DSMs.

The basic intuition for the importance of the first possibility is that ecology is often a combination of a small numbers game and a large numbers one: for instance, it will often occur that a single species has a disproportionate and idiosyncratic role in the dynamics (e.g. Starfish in intertidal food webs (Paine, 1966), Daphnia in aquatic systems, (Mazumder, 1994) or wolves in temperate forests (Ripple et al., 2014), whereas many other species can modelled through their aggregate impact on the community. These scenarios can be theoretically understood with a 'focal species' perspective, often associated with studying the impacts of invasion and extinction (Arnoldi et al., 2022).

One of the most common types of structure that can be added to a random network is blocks or modules, such as functional groups, trophic levels, core and peripheral species, etc. This has been investigated theoretically, notably in the context of bitrophic systems (Yoshino et al., 2007; Advani et al., 2018; Feng et al., 2024). Stochastic Block Models (Vaca-Ramírez and Peixoto, 2022) are a favored inference tool to estimate these matrices of parameters. It might also be possible to detect the existence of groups more easily than we can infer their precise parameter values.

Another important type of structure is quantitative axes along which species are positioned, e.g. traits such as body size or degree of generalism, or trade-off surfaces such as the colonization-competition trade-off. They are in some ways more complex than groups since we can imagine arbitrary dependencies of parameters in the position of species along the axis, and in other ways simpler, since we expect continuity between adjacent values, allowing us to simply ask about, e.g., the slope of variation of parameters with species traits, as in models based on metabolic allometry (Brose et al., 2006; Barbier et al., 2018).

At this time, the use of partially structured models is rather scarce. However, it is a method that has tremendous potential. If we look at figure 4, we see that it is likely that reality is not very likely to resemble fully disordered systems models and that partially structured models represent the path leading to greater realism. It is also the most important element that allows us to interrogate the variety of possible constraints that might act on natural ecosystems. Just exactly how to proceed with the selection of different possible structuring elements is not clear, but it is an exciting possible way to add realism to the study of highly diverse communities without having to resort to the 'true complexity' approach described above.

3.2 Additional patterns

Introducing structure opens up new measurable properties, typically the same observables as before now resolved by positions in the structure. For instance, resolving a food web into a few trophic levels opens up many new properties to measure. Instead of total community biomass or abundance, we can have a more fine-grained description across groups, e.g. biomass pyramids. Instead of community stability, we can study response of one group to the perturbation of another, e.g. trophic cascades. And instead of species diversity versus some control factor like resources or area, we can look at the covariation between groups (e.g. predator diversity versus prey diversity) across the range of the control factor. And, when the structure is not made of groups but of continuous variation along a trait axis or trade-off surface, we can instead ask about the slope of change of these properties across the axis (e.g. the slope of a continuous size spectrum).

Unfortunately, the number of predictable patterns generally grows faster than the number of parameters which might be of increasing interest, but the patterns also become more and more fragile (dependent on precise assumptions) as we increase the level of detail of the structure. In the limit, we are back to 'irreducible complexity' where the number of parameters is at least the square of the number of species. This argument thus mirrors how complexity affects predictability and understanding when starting from the very simplest models.

Partially ordered models thus represent an important extension of disordered systems models because they address how potential deviations from fully disordered models may affect results. Seen in this light, fully disordered models can be thought of as 'null' hypotheses with which to compare any given partially ordered model. If no difference in predictions is observed between the two, then the structure may be irrelevant to the prediction and we may conclude that the result is especially robust to the structure we investigated (Barbier et al., 2018). Similarly, partially ordered models can be compared with simple models based on 'modules' where similar species are lumped together. If no differences in predictions are observed, then it is only the structure that is relevant and we may conclude that the complexity of species interactions is unimportant, at least for the predictions tested.

More generally, we might expect that partially structured models would often reveal important differences in predictions from fully disordered ones and thus allow us to reject the 'null hypothesis' and potentially replace it with the hypothesis implied by the partial structure we imposed (keeping in mind that there may be alternate structures that may do as well or even better). While the work involved in empirically testing (either via quantification of natural patterns or by experiments) is not trivial, one could imagine building a research program that uses partially structured models to identify which aspects of 'irreducible complexity' as we described in our introduction are most important by producing predictions that are most robust to the possible full complexity of the system.

4 Metacommunity

Our original equation of interest (Equation 1) specifies local dynamics using the RLV formulation, but also includes terms for dispersal with the term D_i as a species-specific 'dispersal' parameter. In equation 1 dispersal is modeled as a diffusion process where the net movement between two sites is proportional to the difference in abundance of species I in the two sites. This term converts our model into a 'metacommunity model' in which local communities are embedded in a set of analogous sites that comprise the 'metacommunity'. The biota that exists in the metacommunity can itself constrained by the cumulative dynamics of all the individual sites. This is in contrast with the more abstract approach we have taken so far consisting of an externally defined fixed 'regional biota' that simply and independently provides colonists to any given local community with no feedback from the dynamics of local communities.

Thinking about the dynamics of metacommunities can be traced back to early work at least in the late 1960s (Levins and Culver, 1971) but hasn't really flourished until the early part of the 2000s (Hubbell, 2001; Leibold et al., 2004). This work has greatly expanded our understanding of a wide array of ecological dynamics (reviewed by Leibold and Chase (2018)). Needless to say, these dynamics are more complex and complicated than those we can study with the simpler externally defined regional biota.

However, it is hard to argue that it isn't substantially more realistic. And work done to date in metacommunity ecology has certainly shown numerous ways by which dispersal can substantially modify our expectations about the distribution of species across sites (and often, through time as well). Developing our understanding about metacommunities based on basic principles of local selection (including both 'environmental' and species interactions), dispersal, and stochasticity (as well as speciation that we ignore here) also makes for a much more satisfyingly inclusive approach than the one we have studied so far.

Nevertheless, work to date indicates that there are many complex consequences of taking a metacommunity approach. Could an approach that builds on disordered systems modeling help identify robust and testable aspects of this complexity?

Doing so will take some thinking. If we stick to RLV models, we want to find some formulation of equations like the general equation we started with (Equation 1), either with our without the noise term $\xi_i(t)$. To start with, it immediately implies that we need some new aggregate parameters including at the very least: the variance of dispersal rates among species (D_i), and the spatial variance and auto correlations of growth rates or carrying capacities (r_i, K_i and/or A_{ii}), which represent the spatial features of the environment. It also seems likely we would want to include correlations of these traits with the other traits we have in the basic model. We also need to figure out how to structure the model to minimize these variances and correlations to optimize the utility of the disordered systems approach (as discussed in section 'Application of disorder' above) . Possible general consequences of taking this approach are, as of now, very poorly studied, and at this point our explanations outlined below should be seen as highly speculative even if we can point to some existing results.

However, it is likely that significant new insights can arise. For example, Gravel et al. (2016) studied how dispersal among local communities could alter how diversity and complexity are constrained by random interactions among the species. They start with the well-known disordered Jacobian model of May (1972) that predicts that $S < \left\langle J_{ij}^2 \right\rangle C$ in isolated communities. They then ask how making multiple copies of this model (one for each of the different sites) and connecting them with dispersal (d, which they hold fixed and equal among species) affects this inequality. They also ignore possible effects due to varying growth rates and carrying capacities (except in how they alter the values of J_{ij} via the transition to the relative yield formulation of the RLV. They find two analytical solutions: one that a) holds at very low dispersal levels, resulting in a modified inequality $S < \left\langle J_{ij}^2 \right\rangle Cd$ showing that dispersal can have a small stabilizing effect and thus allows for higher S than predicted by May; and b) one that applies at very high dispersal rates that results in a modified inequality $S < \left\langle J_{ij}^2 \right\rangle C/E$ in which E is the effective number of statistically independent matrices in the set of local matrices in the metacommunity (if all the local communities have the same matrix, $E = 1$ and we recover May's result; if each site is statistically independent of all the others, then we divide by the number of sites; if something in between, we have to calculate this effective value for E). They did not find any analytical solution for intermediate values of d, but numerical studies showed that dispersal was even more stabilizing than either of these limiting conditions. This study, while incomplete in addressing metacommunity dynamics in disordered systems models, nevertheless shows that they can dramatically change the predictions we would have made from the simpler non-metacommunity approach. They also studied additional aspects of metacommunity dynamics and studied how these might even vary with the type of dispersal (local vs global). These results show that multiple aspects of metacommunity dynamics respond in a structured way to variation in dispersal. They also show that these responses are relatively robust to the dispersal structure (at least in comparing local vs global dispersal).

Unfortunately, the fact that we focused our approach to he Jacobian matrix near a given equilibrium prevents other possible predictions (biomasses, turnover, etc) from being studied. We could nevertheless study this with the original spatial RLV model to quantify these effects, which is the focus of ongoing work (Roy et al., 2020; O'Sullivan et al., 2021; Garcia Lorenzana et al., 2024). Metacommunities are also more complex than single communities, and they can be characterized by numerous other metrics than single communities, especially those that involve cross-locality components such as spatial beta-diversity (mean and variance), gradient responses to productivity, etc. At this point, work is only beginning, and we anticipate exciting developments to come as the field better explores these consequences.

5 Conclusions

The assumption of disorder or emergent simplicity is that many model details matter only inasmuch as they contribute to aggregate processes at the community scale. Thus, using disordered systems models, and even more so partially structured models, can be a powerful and general way to extend insights that are based on basic mechanistic processes (e.g. dispersal, environment, interactions) to models that involve large numbers of species. Such models can therefore be critical tools that can allow the empirical study of natural communities (that virtually always contain dozens to thousands of species) while retaining insights that have come from typically small scale and few-species studies (both experimental and theoretical). To date, work using this approach have been largely restricted to idealized theoretical models. Our goal in this paper is to explain the approach in the hope that empirical scientists will be motivated to apply it to a wide array of basic and applied questions involving community and metacommunity ecology.

Here we have focused on applying the disordered (or 'partially structured') systems approach to Lotka-Volterra models because they represent the most obvious, and historically established, basis to approach (meta)community ecology. They can be used as a baseline to study a broad class of models that share the premise that population dynamics is the central driver of the patterns that we are interested in, and all other processes (e.g. species interactions, spatial fluxes) are modelled phenomenologically through their contributions to population growth or decline. However, we note that the approach could also be extended to other formulations of ecological systems (e.g. ecosystem models that are based on nutrient fluxes and stocks, which would more likely display linear than exponential dynamics or colonization-extinction models that ignore abundance patterns). Consequently, the idea of disordered systems is likely to be useful for a number of other ecologically relevant issues even if the Lotka-Volterra structure does not apply.

The idea of approximating species interactions as random is, of course, not new – it was famously employed by May (1972) to discuss the relationship between ecological complexity and stability. Yodzis (1981) challenged May's conclusions by pointing out that food web structure might be important in a way that closely reflects our thoughts about using partially structured models. Nevertheless, until recently, theoretical predictions based on (partially) random interactions had largely focused on a limited set of (often hard to measure) stability properties such as asymptotic resilience (Gravel et al., 2016; Allesina et al., 2015). The technical approach used by May and its successors makes less definite assumptions about ecological dynamics, but is restricted to these stability properties, whereas the Random Lotka-Volterra model can make predictions on many other observables such as abundance distributions across species and across space and time, complex dynamical regimes, etc. It potentially encompasses, qualitatively or even quantitatively, the predictions of a number of other randomness-based approaches that did not focus on species interactions: models of independent species such as the Stochastic Logistic Equation (Grilli, 2020) and stochastic geometry approaches (McGill, 2010), but also neutral theory (Hubbell, 2011) (except for speciation and phylogeny-related questions).

We argue that there are important advantages in using this method. Under full disorder, theory predicts (and simulations confirm) that the many parameters required to describe a community matter only in aggregate, through simple statistics such as mean and variance (Galla, 2006; Bunin, 2017). Partially structured models admit more parameters, but still far fewer than a full description of the community. Thus, it is possible to make predictions about (meta-)communities without having to estimate all their intricate details. Furthermore, using disordered models helps identify empirically robust patterns that are not likely to be fragile to the everchanging details of community processes.

As is true of any approach however, there also important caveats and remaining challenges:

- How to structure interactions in order to apply partially structured models?
- How to infer process from patterns (reverse modeling)?
- Generalizing beyond Lotka-Volterra (e.g. including material cycling or other ecological features that could significantly change the types of dynamics observed here)
- Long-term dynamics (all the regimes seen here, whether equilibria or chaos or noise-driven, are some sort of stationarity, i.e. long-term dynamics whose properties do not change anymore) vs transients

Nevertheless, it is worth keeping in mind that this overall approach has only just started and that improvements are likely in near term. One of the more important ongoing directions is the idea of partially structured models. An initial quantitative, albeit limited, exploration of this was proposed by Barbier et al. (2018) and here we point to other possible extensions. Work in this direction is particularly important in better addressing the various form of structure that empirical scientists have identified and because they can provide important testable hypotheses to determine if such structure actually matters for explaining natural systems. What is increasingly important is for empirical scientists to capitalize on these insight and develop methods to evaluate their utility in natural communities. It is our intent in this review to facilitate this effort by providing what we hope is a useful introduction to doing so.

References

Advani, M., Bunin, G., and Mehta, P. (2018). Statistical physics of community ecology: a cavity solution to macarthur's consumer resource model. *Journal of Statistical Mechanics: Theory and Experiment*, 2018(3):033406.

Allesina, S., Grilli, J., Barabás, G., Tang, S., Aljadeff, J., and Maritan, A. (2015). Predicting the stability of large structured food webs. *Nature communications*, 6(1):7842.

Arnoldi, J.-F., Barbier, M., Kelly, R., Barabás, G., and Jackson, A. L. (2022). Invasions of ecological communities: Hints of impacts in the invader's growth rate. *Methods in Ecology and Evolution*, 13(1):167–182.

Arnoldi, J.-F., Bideault, A., Loreau, M., and Haegeman, B. (2018). How ecosystems recover from pulse perturbations: A theory of short-to long-term responses. *Journal of theoretical biology*, 436:79–92.

Barbier, M., Arnoldi, J.-F., Bunin, G., and Loreau, M. (2018). Generic assembly patterns in complex ecological communities. *Proceedings of the National Academy of Sciences*, 115(9):2156–2161.

Barbier, M., Bunin, G., and Leibold, M. A. (2022). Getting more by asking for less: Linking species interactions to species co-distributions in metacommunities. *in prep*.

Barbier, M., De Mazancourt, C., Loreau, M., and Bunin, G. (2021). Fingerprints of high-dimensional coexistence in complex ecosystems. *Physical Review X*, 11(1):011009.

Bizeul, P. and Najim, J. (2021). Positive solutions for large random linear systems. *Proceedings of the American Mathematical Society*, 149(6):2333–2348.

Brose, U., Williams, R. J., and Martinez, N. D. (2006). Allometric scaling enhances stability in complex food webs. *Ecology letters*, 9(11):1228–1236.

Bunin, G. (2017). Ecological communities with lotka-volterra dynamics. *Physical Review E*, 95(4):042414.

Bunin, G. (2021). Directionality and community-level selection. *Oikos*, 130(4):489–500.

Feng, Z., Marsland III, R., Rocks, J. W., and Mehta, P. (2024). Emergent competition shapes top-down versus bottom-up control in multi-trophic ecosystems. *PLoS Computational Biology*, 20(2):e1011675.

Fort, H. (2018). Quantitative predictions from competition theory with an incomplete knowledge of model parameters tested against experiments across diverse taxa. *Ecological Modelling*, 368:104–110.

Fried, Y., Kessler, D. A., and Shnerb, N. M. (2016). Communities as cliques. *Scientific reports*, 6(1):1–8.

Galla, T. (2006). Random replicators with asymmetric couplings. *Journal of Physics A: Mathematical and General*, 39(15):3853.

Garcia Lorenzana, G., Altieri, A., and Biroli, G. (2024). Interactions and migration rescuing ecological diversity. *PRX Life*, 2(1):013014.

Gravel, D., Massol, F., and Leibold, M. A. (2016). Stability and complexity in model meta-ecosystems. *Nature communications*, 7(1):1–8.

Grilli, J. (2020). Macroecological laws describe variation and diversity in microbial communities. *Nature communications*, 11(1):1–11.

Hairston, N. G., Smith, F. E., and Slobodkin, L. B. (1960). Community structure, population control, and competition. *The American Naturalist*, 94(879):421–425.

Hatton, I. A., Heneghan, R. F., Bar-On, Y. M., and Galbraith, E. D. (2021). The global ocean size spectrum from bacteria to whales. *Science advances*, 7(46):eabh3732.

Holt, R. a. and Hochberg, M. (2001). Indirect interactions, community modules and biological control: a theoretical perspective. In *Evaluating indirect ecological effects of biological control. Key papers from the symposium'Indirect ecological effects in biological control', Montpellier, France, 17-20 October 1999*, pages 13–37. CABI Publishing Wallingford UK.

Holt, R. D. (1977). Predation, apparent competition, and the structure of prey communities. *Theoretical Population Biology*, 12(2):197–229.

Hu, J., Amor, D. R., Barbier, M., Bunin, G., and Gore, J. (2022). Emergent phases of ecological diversity and dynamics mapped in microcosms. *Science*, 378(6615):85–89.

Hubbell, S. P. (2001). *The unified neutral theory of biodiversity and biogeography.* Princeton University Press.

Hubbell, S. P. (2011). The unified neutral theory of biodiversity and biogeography. In *The Unified Neutral Theory of Biodiversity and Biogeography.* Princeton University Press.

Lawlor, L. R. (1979). Direct and indirect effects of n-species competition. *Oecologia*, 43:355–364.

Lawton, J. H. (1999). Are there general laws in ecology? *Oikos*, 84(2):177–192.

Leibold, M. A. and Chase, J. M. (2018). *Metacommunity ecology*, volume 59. Princeton University Press.

Leibold, M. A., Holyoak, M., Mouquet, N., Amarasekare, P., Chase, J. M., Hoopes, M. F., Holt, R. D., Shurin, J. B., Law, R., Tilman, D., et al. (2004). The metacommunity concept: a framework for multi-scale community ecology. *Ecology letters*, 7(7):601–613.

Levine, S. H. (1976). Competitive interactions in ecosystems. *The American Naturalist*, 110(976):903–910.

Levins, R. and Culver, D. (1971). Regional coexistence of species and competition between rare species. *Proceedings of the National Academy of Sciences*, 68(6):1246–1248.

Mallmin, E., Traulsen, A., and De Monte, S. (2024). Chaotic turnover of rare and abundant species in a strongly interacting model community. *Proceedings of the National Academy of Sciences*, 121(11):e2312822121.

Marcus, S., Turner, A. M., and Bunin, G. (2022). Local and collective transitions in sparsely-interacting ecological communities. *PLoS computational biology*, 18(7):e1010274.

May, R. M. (1972). Will a large complex system be stable? *Nature*, 238(5364):413–414.

Mazumder, A. (1994). Large-scale dynamics of zooplankton communities: multi-trophic interactions. *Canadian Journal of Fisheries and Aquatic Sciences*, 51(8):1893–1905.

McGill, B. J. (2010). Towards a unification of unified theories of biodiversity. *Ecology letters*, 13(5):627–642.

McPeek, M. A. (2021). Ecological pattern and process: the ecological theater and the evolutionary play. *The American Naturalist*, 198(5):592–605.

Novak, M., Yeakel, J. D., Noble, A. E., Doak, D. F., Emmerson, M., Estes, J. A., Jacob, U., Tinker, M. T., and Wootton, J. T. (2016). Characterizing species interactions to understand press perturbations: what is the community matrix? *Annual Review of Ecology, Evolution, and Systematics*, 47:409–432.

O'Sullivan, J. D., Terry, J. C. D., and Rossberg, A. G. (2021). Intrinsic ecological dynamics drive biodiversity turnover in model metacommunities. *Nature Communications*, 12(1):3627.

Paine, R. T. (1966). Food web complexity and species diversity. *The American Naturalist*, 100(910):65–75.

Ripple, W. J., Estes, J. A., Beschta, R. L., Wilmers, C. C., Ritchie, E. G., Hebblewhite, M., Berger, J., Elmhagen, B., Letnic, M., Nelson, M. P., et al. (2014). Status and ecological effects of the world's largest carnivores. *Science*, 343(6167):1241484.

Roy, F., Barbier, M., Biroli, G., and Bunin, G. (2020). Complex interactions can create persistent fluctuations in high-diversity ecosystems. *PLoS computational biology*, 16(5):e1007827.

Schaffer, W. M. (1981). Ecological abstraction: the consequences of reduced dimensionality in ecological models. *Ecological monographs*, 51(4):383–401.

Spaak, J. W. and De Laender, F. (2020). Intuitive and broadly applicable definitions of niche and fitness differences. *Ecology letters*, 23(7):1117–1128.

Vaca-Ramírez, F. and Peixoto, T. P. (2022). Systematic assessment of the quality of fit of the stochastic block model for empirical networks. *Physical Review E*, 105(5):054311.

Vellend, M. (2010). Conceptual synthesis in community ecology. *The Quarterly review of biology*, 85(2):183–206.

Vellend, M. (2016). The theory of ecological communities (mpb-57). In *The Theory of Ecological Communities (MPB-57)*. Princeton University Press.

Wang, S., Isbell, F., Deng, W., Hong, P., Dee, L. E., Thompson, P., and Loreau, M. (2021). How complementarity and selection affect the relationship between ecosystem functioning and stability. *Ecology*, 102(6):e03347.

Yodzis, P. (1981). The stability of real ecosystems. *Nature*, 289(5799):674–676.

Yoshino, Y., Galla, T., and Tokita, K. (2007). Statistical mechanics and stability of a model eco-system. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(09):P09003.

Zelnik, Y. R., Barbier, M., Shanafelt, D. W., Loreau, M., and Germain, R. M. (2021). Linking intrinsic scales of ecological processes to characteristic scales of biodiversity and functioning patterns. *bioRxiv*.

Box: The interpretation of disorder

Taking the goal of understanding aggregate properties, rather than detailed aspects at the species level, means that we might be able to parameterize our model with less detail than would be needed to make the more complex and precise predictions. Fully disordered systems models take an approach that is based on statistical moments of sets of parameters. We start by discussing 'fully disordered' models where the theory is simpler but we subsequently expand our approach to develop 'partially structured systems' models. Figure 2 contrasts the classical approach to complexity that starts with a single species and builds in increasingly complex dynamics via modules, networks, etc., with the disordered systems approach that starts with completely disordered RLV models, then modifies these with partially structured components as they may reveal themselves to be important (e.g. trophic levels or spatial structure). The two can even be seen as complementary methods that can cross-check each other.

Three non-exclusive justifications for the disordered systems approach are worth considering:

Objective motivation: a confluence of many factors

When there are many independent sources of variation, their combined effects often behave like random variables (see e.g. discussion in Barbier et al. (2021)). The canonical example is a dice throw where many deterministic microscopic factors (spin, speed, orientation, angle to substrate, etc.) actually determine the outcome but interfere so much with each other that none of them dominates, and the result is effectively random. Similarly, we may treat pair-wise species interactions as random if the impact of one species on another is not predictable from the interactions of these two species with other ones (i.e. pair-wise interactions are not inter-correlated or are affected by other factors that interfere with each other so as to make them so).

This implies that sufficiently complex interactions appear as 'random', in the sense that they cannot be resolved in detail, even though this aspect of the dynamics is actually deterministic (determined by the complexity of interactions that interfere with each other) and therefore may be hard to distinguish from other forms of stochasticity ('noise'). Such randomness can be seen as a conceptual "attractor" that can describe our expectations for the absence of any additional structuring features of the interactions and can consequently generate 'null hypotheses' by which to compare other possible effects (see the section below on Partially structured models). These null expectations thus sum both 'noise' on the one hand and 'unresolved determinism' on the other.

Pragmatic motivation: baseline models for robust phenomena

Even if we do not accept randomness as being “objectively” a good representation of the dynamics we are modelling, we can still adopt it on pragmatic grounds. We admit that we are dealing with intricate complex systems, possibly devoid of any true randomness, but we only focus on behaviors and patterns that are robust to these details because we believe these predictions would also hold even if we were to make changes in the details of the studied system (e.g. shuffle carying capacities, growth rates, or interaction coefficients). It is then plausible that the same phenomena or patterns will be very frequently be observed if these details are drawn at random and deviations would require very specific changes that would be very unlikely (Figure 3). Choosing a random model is useful simply because it is easier to manipulate than most non-random ones with the same behaviors, and it thus serves as a “baseline” model for this aspect of the dynamics.

Subjective motivation: aggregation and statistical equivalence

A third option is a subjective Bayesian viewpoint: 'Randomness' simply reflects the degree to which we are uncertain or uncaring about certain details. If so, aggregating implies that the variables we are aggregating over (e.g. the abundances of various species) are, in some sense, equivalent or exchangeable – that we do not mind adding apples and oranges if what we care about is total fruit biomass. This does not mean that these variables are identical in every way, but only that none is "special" in how it contributes to the aggregate. This clearly depends on which aggregate pattern is our focus, it may seem strange to add together the biomass of predators and prey to predict total biomass given that they have very different ecological consequences, but some predictable patterns assume an equivalence between them even if others many not.

Random models appear like a natural choice when we believe that we are indeed justified in treating species as statistically equivalent for a given pattern. In a random interaction network, species are not identical, but no species occupies a very unique role – even a well-connected "hub" species may not be a remarkable outlier because it occurs within the expected distribution of means and variances, but can be modeled as a representative sample of the overall community's distribution of connectedness. We come back to this notion of statistical equivalence below in Sec. 1.2.

6 Box: Diagram of dynamical regimes

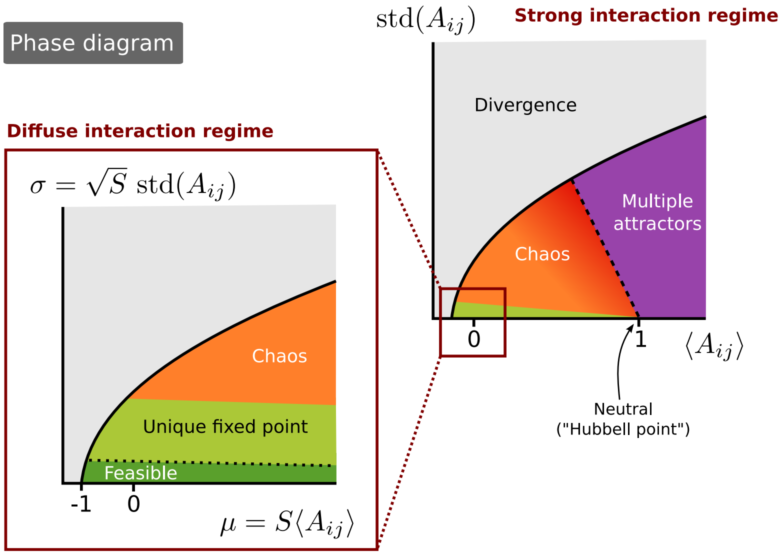

We show here dynamical regimes for the Random Lotka Volterra model without noise or space. Existing theory on these regimes crucially depends on the overall strength of interactions. We can define three main situations leading to distinct predictions: i) each species has a number of interactions α_{ij} that are each individually important (not negligible compared to $\alpha_{ii} = 1$ or stronger), ii) each pairwise interaction has limited impact on a species, but the total impact of all its interaction partners is important, and iii) important interactions are sparse and the others are either small or absent.

In the first case, each species can have a significant impact on many others' abundances, and negative interactions can easily lead to extinction. In speciose communities including such strong negative interactions, we expect that few species will coexist (Mallmin et al., 2024). In the second case, each species can be understood as interacting with the community as a whole, and a simple community-level description emerges readily. We can define aggregate parameters μ and σ^2 as the mean and variance of the sum of all interactions experienced by a species in the community

$$\mu \equiv \left\langle \sum_j^S \alpha_{ij} \right\rangle = S \langle \alpha_{ij} \rangle \,, \qquad \sigma^2 \equiv \mathrm{var}\left(\sum_j^S \alpha_{ij} \right) = S \ \mathrm{var}(\alpha_{ij}) \tag{4}$$

These parameters then turn out to be the main ones controlling all the outcomes (Bunin, 2017). Many other possible details such as which distributions interactions are drawn from (normal vs uniform vs other) do not affect the conclusions we make about their behavior. Finally, we may have very sparse interactions, where a substantial majority of pairwise coefficients A_{ij} are negligible, while some are strong. This regime is increasingly studied theoretically (Fried et al., 2016; Marcus et al., 2022). We note that interactions that can most plausibly be treated as fully disordered involve competitive or multifactorial interactions within a guild of similar species, e.g. competition between grasses in a field, or a mixture of positive and negative effects between soil bacteria. Contrary to e.g. food webs or plant pollinator networks, it is not obvious why interactions within a guild should be sparse, and thus we do not discuss this situation in depth here.

Most existing theoretical studies based on disordered systems are limited to the diffuse interaction regime shown in inset, where species interactions are important in aggregate but very weak individually (i.e. μ and σ remain moderate even for very large species number S), as this is the regime where theoretical tools are most applicable and the most robust predictions can be made. The strong interaction regime is less understood, mainly through simulations, except in some particular limits. The grey region indicates where abundances become infinite due to an excess of mutualistic interactions, i.e. a breakdown of the LV model, which can be solved in various ways.

One might expect that many dynamical behavior types might be possible in the Random Lotka-Volterra model. However, and perhaps surprisingly, the RLV only has a limited set of behaviors of interest, as these three occur predictably whereas all the other possibilities become increasingly unlikely when species numbers are 'large' (in practice greater than 15-20). Previous studies, e.g. Hu et al. (2022), have found suggestive empirical evidence for the first three phases.

Global coexistence: First, it is possible that all species in the biota can coexist and reach a stable point equilibrium together (Bizeul and Najim, 2021).

Unique fixed point with low diversity: Here, a fraction of species go extinct, but the remainder coexist stably in an equilibrium that cannot be invaded by any of the extinct species (if we reintroduce them, they go extinct again).

Chaotic turnover: Here, the dynamics tend to approach an equilibrium where a significant fraction of species thrive at significant abundances, but all such equilibria are unstable and can be invaded by other species that were previously going extinct. There is consequently a constant turnover of species through a kind of "pinball" dynamics (Roy et al., 2020; O'Sullivan et al., 2021). We note that when S is small, we can get limit cycles instead of chaos in that region.

Multiple attractors: When interactions are individually strong, the dynamics may reach multiple stable states, the simplest case being bistability between two competitors that can exclude each other. In the case of sparse interactions, this may lead to dynamics that are *directional*, i.e. progressing over long times toward more "mature" states (Bunin, 2021), contrary to chaotic turnover which is typically *adirectional*.

Box: Step-by-step guide to applying disordered systems theory

This box serves as a quick summary of the points made throughout the manuscript.

A step-by-step description of the process could be provided.

1. Find a parameterization of the system that tries to make the species as equivalent as possible (Sec. 1.2), e.g. expressing them in a common currency like total biomass or total resource consumption.

2. Introduce all elements in the model that are needed to resolve measurable behaviors and patterns (Sec. 2)

3. Try to fit the disordered model to data using the most discriminating patterns, bearing in mind that some parameters are hard to resolve depending on which patterns are available to measure (e.g. it is hard to distinguish temporal fluctuations due to environmental perturbations versus species interactions).

4. If disordered predictions are unsuccessful, introduce the minimal amount of structure needed to reach satisfactory predictions (Section 3),

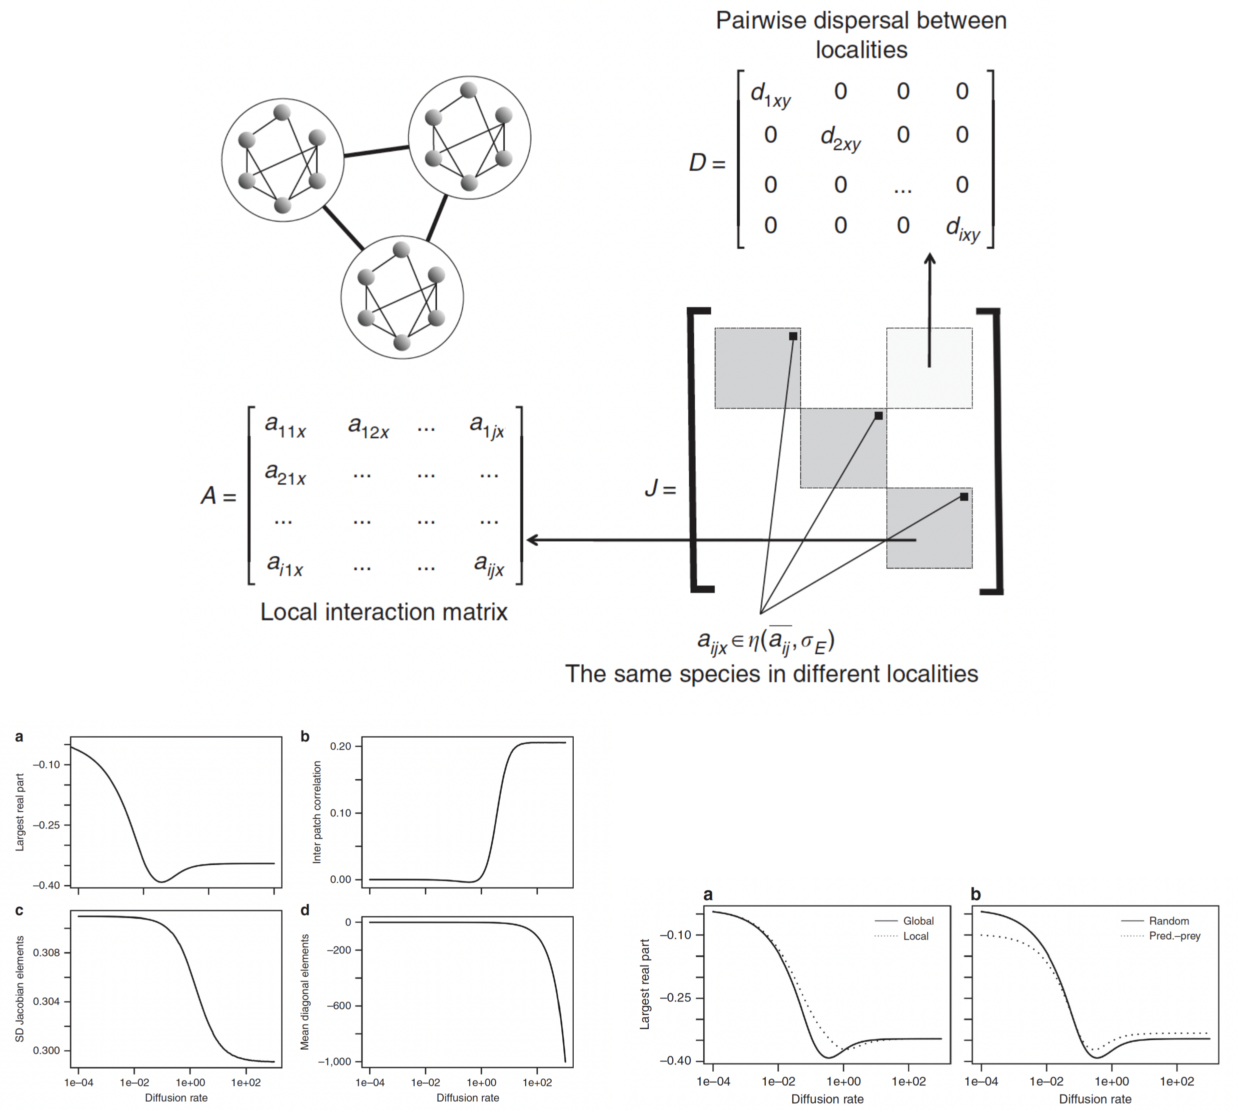

Figure 5: An archetypal representation of a metacommunity. Here we study the Jacobian Matrix foor a linearizeddynamic model, but the representation still captures the essence of the idea. Local interactions are represented by the A matrix that has intraspecific interactions (along the main diagonal that represent environmental selection), interspecific interactions (in the off-diagonal elements of the submatrices along the main diagonal), and dispersal in the off-diagonal matrices. A relevant formulation of DSM model for this representation would have each of these elements characterized by means and variances. And additional structure to generate relevant PSMs would involve modifying this general model to partition these matrices in relevant ways or by imposing correlations among the parameters. For example, the set of dispersal matrices could be constrained to allow only local dispersal (only matrices along the most proximal matrix to the main diagonal have non-zero values) vs global dispersal (there are no off diagonal matrices that are zero).