

IDEA AND PERSPECTIVE

An expanded modern coexistence theory for empirical applications

Stephen P. Ellner,¹ Robin E. Snyder,² Peter B. Adler³ and Giles Hooker⁴

Abstract

Understanding long-term coexistence of numerous competing species is a longstanding challenge in ecology. Progress requires determining which processes and species differences are most important for coexistence when multiple processes operate and species differ in many ways. Modern coexistence theory (MCT), formalised by Chesson, holds out the promise of doing that, but empirical applications remain scarce. We argue that MCT's mathematical complexity and subtlety have obscured the simplicity and power of its underlying ideas and hindered applications. We present a general computational approach that extends our previous solution for the storage effect to all of standard MCT's spatial and temporal coexistence mechanisms, and also process-defined mechanisms amenable to direct study such as resource partitioning, indirect competition, and life history trade-offs. The main components are a method to partition population growth rates into contributions from different mechanisms and their interactions, and numerical calculations in which some mechanisms are removed and others retained. We illustrate how our approach handles features that have not been analysed in the standard framework through several case studies: competing diatom species under fluctuating temperature, plant–soil feedbacks in grasslands, facilitation in a beach grass community, and niche differences with independent effects on recruitment, survival and growth in sagebrush steppe.

Keywords

Coexistence, competition, environmental variability, model, theory.

Ecology Letters (2018)

INTRODUCTION

Understanding the indefinite coexistence of numerous competing species is a longstanding, central question in community ecology (e.g. Hutchinson 1959, 1961; Grubb 1977; Shmida & Ellner 1984; Hubbell 2001). This issue is particularly acute for plants, which as Grubb (1977 p.107) noted “all need light, carbon dioxide, water and the same mineral nutrients.” But it also arises when co-occurring animal species compete for a resource or habitat that is essential to all, such as coral reef fish competing for territories (e.g. Sale 1979; Caley 1995; Munday 2004; Volkov *et al.* 2007) and desert rodents competing for seeds (e.g. Kotler & Brown 1988; Brown 1989; Ziv *et al.* 1993; Abu Baker & Brown 2014).

The solution Grubb proposed was the “regeneration niche”: even if the trees are very similar, seeds and seedlings may be very different. Moreover, species may regenerate at different times (e.g. Usinowicz *et al.* 2017). But there are many other hypotheses, all with theoretical and many with empirical support, including: predator limitation (Holt & Bonsall 2017), specialist pathogens (Bagchi *et al.* 2014; Comita *et al.* 2014), hydrological niches (Silvertown *et al.* 1999), resource ratio

differences (Dybzinski & Tilman 2007), spatial environmental variation (Sears & Chesson 2007), and life history trade-offs (Lönnerberg & Eriksson 2013).

Rarely will only one of these processes be operating in a real species-rich communities. Progress therefore requires more than just identifying which processes that might contribute to coexistence operate in a community. If we observe that two warbler species forage in different parts of the tree, is this crucial for coexistence, or irrelevant because neither species is resource-limited? We need ways to determine which differences and processes are most important when species differ in many ways and multiple processes operate.

Modern coexistence theory (MCT), formalised by Chesson (1994) and Chesson (2000a), holds out the promise of doing exactly that, by quantifying the contributions of different mechanisms to species persistence. The over 2300 citations of Chesson (2000b) – nearly half in the last 5 years – attest to the conceptual importance of MCT (Web of Science, accessed April 25 2018). However, empirical applications of MCT remain scarce. One difficulty for empirical applications is that applying MCT to a new study system often requires a case-specific model, and deriving the necessary formulas for

¹Department of Ecology and Evolutionary Biology, Cornell University, Ithaca, NY, USA

²Department of Biology, Case Western Reserve University, Cleveland, OH, USA

³Department of Wildland Resources and the Ecology Center, Utah State University, Logan, UT, USA

⁴Department of Biological Statistics and Computational Biology, Cornell University, Ithaca, NY, USA

*Correspondence: E-mail: spe2@cornell.edu

coexistence mechanisms in a new model entails a complex mathematical analysis requiring a deep mathematical understanding of MCT. Annual plants with variable germination are a classic example in MCT (Chesson 1994), but the first empirical application in Angert *et al.* (2009) required a new model, entailing a 17-page mathematical appendix to derive the necessary formulas.

Another challenge is that restrictive assumptions are often used to make the mathematics tractable, which can bias the results. For example, Angert *et al.* (2009) assumed that all species are equally affected by competition and that inter- and intra-specific competition were equal in strength for all species. MCT's assumption of small environmental variance is often problematic for empirical applications, because fluctuation-dependent mechanisms only become important when environmental variation is large. Mathematical results are also largely limited to unstructured models in which populations are described by total abundance, while many empirical population studies use matrix or integral projection models because there are substantial demographic differences between individuals of different ages or sizes. Similarly, formulas for stabilising and equalising components of niche differences have been cited and applied (e.g. Godoy & Levine 2014, eqns A.4 and A.5) without realising that they are specific to two-species communities with Lotka–Volterra competition. So while analytic results have yielded important insights, empirical applications often require less restrictive assumptions.

Finally, MCT analyses coexistence in terms of a few conceptual “mechanisms” (Box 1). Instead, many ecologists might prefer an analysis focusing on observed processes amenable to direct experimental study, such as resource competition, indirect competition via predators or pathogens, life history tradeoffs or other system-specific processes. Theory with the flexibility to analyse coexistence in terms of multiple species differences and multiple mechanisms, in ways that readily adapt to various study systems, would substantially broaden applicability.

We propose a broad extension of MCT that removes these obstacles. Our previous paper (Ellner *et al.* 2016) provided a partial solution for one mechanism, the storage effect. The general approach here applies to all the standard MCT coexistence mechanisms, to process-defined mechanisms such as those listed above and any others thought to operate in a community. Our approach is an *extension* of current MCT, not a replacement for it, designed for detailed analysis of particular communities rather than general principles and insights.

In many situations, the dominant coexistence mechanisms do not involve temporal fluctuations, so long-term data on population responses to temporal variability is not necessary for applying our approach. And with appropriate data, our approach can also be applied to components of population growth rate, such as survival or fecundity.

We begin by summarising the core ideas which form the conceptual basis for MCT. For clarity we gloss over some subtleties that become important in applications (a thorough, exact exposition is provided by Barabás *et al.* (2018)). We then introduce our approach through a simple case study, coexistence maintained by variable temperature in an experimental community of two diatom species. Next, we present our general approach, and illustrate it through a series of empirical case studies that include spatial coexistence mechanisms, facilitation, and structured

populations with stage-specific niche differences. In all of these cases, we quantify coexistence mechanisms not covered by MCT, though for the diatom case study, we also quantify the traditional MCT mechanisms of storage effect and relative non-linearity. R scripts to replicate all figures and tables are available online (see DATA ACCESSIBILITY STATEMENT).

CORE IDEAS OF MCT

Modern coexistence theory posits that coexistence is stabilised by processes that give any species, when rare, a population growth rate advantage over other “resident” species that remain at typical steady-state abundances. A species' average instantaneous population growth rate when rare is called its *invasion growth rate*. A positive invasion growth rate buffers a species against extinction, maintaining its persistence in the community. If a species relies on different resources than its competitors, or is limited by different enemies or different environmental conditions, then its invasion growth rate may be positive. All species persist and coexistence is stable if *all species have a positive invasion growth rate*.

Modern coexistence theory quantifies coexistence mechanisms by asking how they contribute to each invader's population growth rate advantage over resident species. It does this by comparing observed population growth rates with those that would occur if one or more mechanisms were absent. How much would invader and resident growth rates change if all enemies were perfect generalists, or if the environment were constant?

The core approach in MCT is *decompose and compare*. Decompose population growth rates into a sum of terms for the effects of different factors, and then compare invader and residents term-by-term. Considering invader–resident differences is essential because we want to say that a mechanism stabilises coexistence of species A and species B if it gives each, when it is rare, an advantage over the other. This can happen two ways: the mechanism can help whichever species is rare, or it can hurt whichever species is common. The invader's growth rate includes only the former; to capture both, we need to make an invader–resident comparison. A residents' average population growth rate is necessarily zero, because they are neither increasing nor decreasing in the long run. However the contribution of any particular mechanism to that growth rate (e.g. the effect on its growth rate of predator switching or a variable environment) could be positive or negative, depending on whether the mechanism helps or hurts.

Standard MCT uses Taylor-series expansions to decompose and to evaluate invader–resident differences; in section SI.1, we give a simple example to illustrate the procedures. The resulting term-by-term differences in the growth rate decomposition are then grouped into the canonical coexistence mechanisms of standard MCT (Box 1).

In the invader–resident comparisons, it is essential that residents are not allowed to re-equilibrate when we ask, for example, how does variance in temperature contribute to coexistence? It seems natural to answer that question by doing a simulation or experiment with temperature held constant. But constant temperature helps or hurts each species, changing their abundance, age structure, etc., thus altering the competition experienced by each species. When we compare the outcomes with and

Box 1 Summary of the canonical coexistence mechanisms in MCT, following Barabás et al. (2018).

In MCT for purely temporal variation (Chesson 1994), the population growth rate r_j of species j is assumed to depend on the environment, E_j , and competition, C_j . E and C are not direct measures of the physical environment and competition, but parameters that represent population responses (e.g. C may be the proportional reduction in cell division rate due to resource scarcity). Note that C includes all frequency- and density-dependent feedbacks. MCT then asks how differences among species in the distributions and impacts of the E_j s and C_j s affect the long-run growth rate of each species in both resident and invader states. Second-order Taylor expansion is used to partition the effects on growth rates of different moments of \mathcal{E}_j and \mathcal{C}_j , the direct effects of E_j and C_j on growth rates. The invader-resident differences in each of the resulting terms are then grouped into the following “mechanisms” for species i as invader:

- r'_i , the sum of all terms that do not include density-dependent feedbacks.
- $\Delta\rho_i$, the sum of all terms involving means of the \mathcal{C}_j .
- ΔN_i , the sum of all terms involving variances of the \mathcal{C}_j .
- ΔI_i , the sum of all terms involving $\mathcal{E}_j, \mathcal{C}_j$ covariances.

r'_i gives the invasion growth rate in the absence of direct density-dependent feedbacks. However, it includes $Var(E)$ terms that measure effects of environmental fluctuations on population growth rates, which we generally make a separate term in our analyses. Whenever environment and competition both vary, there are two possible sources of nonlinear averaging, but only one of them (competition) gets a stand-alone term in standard MCT.

$\Delta\rho_i$ represents all fluctuation-independent mechanisms, such as resource partitioning or species-specific enemies. In standard MCT, the scaling factors q_{ir} used in invader–resident comparisons (see eqn 20) are chosen so that $\Delta\rho_i = 0$.

ΔN_i is called *relative nonlinearity in competition*. This fluctuation-dependent term reflects differences in the degree of nonlinearity of r in each species' response to limiting competitive factors. If these differences are present, and the limiting factors fluctuate, nonlinear averaging can benefit some species and hurt others.

ΔI_i is called the *storage effect* because this mechanism's importance was first recognised in models where gains during good years were “stored” in a long-lasting life stage with low sensitivity to environment or competition. However, ΔI can make a positive contribution to coexistence whenever invaders have low environment-competition covariance (letting them increase rapidly in good conditions) and they are buffered against equally rapid decrease in bad conditions.

The canonical growth-rate decomposition for purely temporal variability is then

$$\bar{r}_i = r'_i + \Delta N_i + \Delta I_i.$$

With spatial rather than temporal variability, there is an additional term (fitness-density spatial covariance) and the spatial storage effect involves spatial rather than temporal covariance of environment and competition.

Our growth-rate decompositions differ in several ways. Because we do not use Taylor approximations, our decompositions usually include higher-order interaction terms. Where standard MCT reduces the number of terms by grouping and weighting, we decompose finely. We also prefer to work directly with the measurable state variables and environmental covariates that characterise the community (e.g. resource concentration and temperature) rather than the effects of resource limitation and temperature on population growth rates. However, users who prefer can express population growth rates as functions of E and C variables as they would be defined in standard MCT, and do decompositions using our methods that align more closely with the canonical mechanisms. We discuss these options for our diatoms case study in section SI.3.

without temperature variability, we would not know how much of the different is due to temperature variability *per se*, and how much is due to its cascading effects on competitive interactions, age structure, etc. So instead, each term in the invader–resident comparison is evaluated in the situation where all processes are operating, and terms involving $Var(\text{Temp})$ quantify the direct effect of variance in temperature. This point can be hard to understand in the abstract, so in our case studies below we highlight where it comes up.

HOW DOES FLUCTUATING TEMPERATURE MAINTAIN DIATOM COEXISTENCE?

The most important question about our new approach is, can it tell us something useful? Here, we use an empirical case

study to argue that it does, by quantifying the mechanisms contributing to coexistence of two diatom species in experiments by Descamps-Julien & Gonzalez (2005). We do this in two ways, the first analogous to standard MCT for temporally fluctuating environments, the second based on the trait differences between the species.

Descamps-Julien & Gonzalez (2005) demonstrated that two diatom species, *Cyclotella pseudostelligera* and *Fragilaria crotonensis*, competing for a single limiting resource (silicate), could coexist in a chemostat with periodic variation in temperature, but not at any constant temperature. Ellner *et al.* (2016) showed that the storage effect was not sufficient to explain the diatoms' coexistence and therefore suggested that relative nonlinearity of competition, the only other fluctuation-dependent mechanism in standard MCT, was an essential

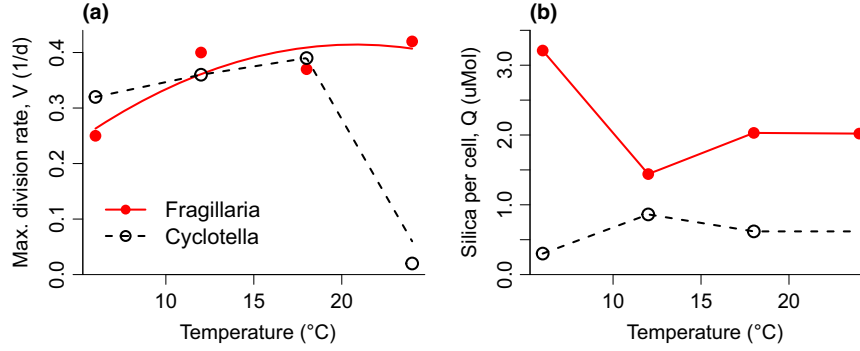


Figure 1 Species-specific temperature responses of the parameters (a) V , and (b) Q , governing nutrient uptake and conversion efficiency. Points (closed circles: *Fragillaria*, open circles: *Cyclotella*) are parameter estimates derived from 9-day single-species batch experiments on each species at a constant temperature (Table 1 of Descamps-Julien & Gonzalez (2005)). The plotted lines and curves were used to simulate the model with continuously varying temperature. Q for *Cyclotella* could not be estimated at 24 °C because of its very low growth rate in the batch experiments. Because *Cyclotella*'s growth at 24 °C was much better in chemostats than in batch experiments, our V function for *Cyclotella* (dashed line in panel a) uses a higher value of V at 24 °C, chosen so that the model matches better the average abundance of *Cyclotella* in chemostat experiments; but even without this modification the model predicts coexistence in the variable temperature experiment. Figure generated by `PlotForcedChemo.R`, `ForcedChemoSubs.R`

coexistence mechanism. We were wrong. We forgot about nonlinear averaging of environmental fluctuations, which does not get a stand-alone term in standard MCT.

Our model, closely following Descamps-Julien & Gonzalez (2005), is a two-species chemostat with some parameters depending on temperature θ ,

$$\begin{aligned} \frac{dS}{dt} &= D(S_0 - S) - Q_1(\theta)x_1 \frac{V_1(\theta)S}{K_1 + S} - Q_2(\theta)x_2 \frac{V_2(\theta)S}{K_2 + S} \\ \frac{dx_j}{dt} &= x_j \frac{V_j(\theta)S}{K_j + S} - Dx_j, \quad j = 1, 2. \end{aligned} \quad (1)$$

Here S is extracellular silicate concentration in the chemostat, x_1 and x_2 are population densities of *Fragillaria* and *Cyclotella*, respectively, S_0 is silicate concentration in the inflow, and D is dilution (outflow) rate. Temperature varies periodically,

$$\theta(t) = \theta_0 + a \sin(2\pi t/P). \quad (2)$$

with mean θ_0 , amplitude a , period P . Coexistence was observed experimentally with $\theta_0 = 18^\circ\text{C}$, $a = 6^\circ\text{C}$, $P = 60\text{d}$.

Functions specifying how Q_j (silicate per cell) and V_j (maximum cell division rate) depend on temperature were estimated from batch experiments (Fig. 1). As half-saturation constants K_j are nearly constant over the range of temperatures in the experiments (18–24°C), we model them as constant: $K_1 = 0.25 \mu\text{M}$ (*Fragillaria*), $K_2 = 0.14 \mu\text{M}$ (*Cyclotella*). At temperatures of 18°C or lower where both species have similar V values, *Cyclotella* has a significant advantage (smaller K , hence faster nutrient uptake) but at high temperatures *Fragillaria* wins because its V remains high while *Cyclotella*'s falls precipitously.

In the model, species j has silicate- and temperature-dependent instantaneous population growth rate

$$r_j(\theta, S) = \frac{V_j(\theta)S}{K_j + S} - D. \quad (3)$$

Our first analysis begins by partitioning the long-term average population growth rate \bar{r} of each species, in both its invader and resident states, into: the growth rate that would occur

without variance in silicate or temperature; the main effects of variance in silicate and in temperature; the interaction between the two variances; and effects of covariance between silicate and temperature.

To do the analysis for species 1, values of $S(t)$ and $\theta(t)$ are recorded from a long simulation of the model with species 1 invading (absent, or kept very rare at all times), and species 2 resident at steady state, using empirically estimated parameters under the experimental conditions. The same analysis could be done with data from a long experiment. Values need to be recorded at times t_k ($k = 1, 2, \dots, m$) spaced closely enough to capture all relevant population dynamics, and for long enough to accurately estimate average growth rate; in practice this means that doubling the simulation/experiment duration or doubling the observation frequency has no meaningful effects.

Denote the average values of silicate and temperature across the times t_k by \bar{S} and $\bar{\theta}$, respectively. The average population growth rates of each species are then estimated by time-averaging,

$$\bar{r}_j = \frac{1}{m} \sum_{k=1}^m r_j(\theta(t_k), S(t_k)), \quad j = 1, 2. \quad (4)$$

A population grows if $\bar{r} > 0$.

The growth rate partitioning is depicted in Fig. 2, and formulas for each term are in Table 1. To compute the main effect of silicate variability, for each species define

$$\begin{aligned} \varepsilon_j^0 &= r_j(\bar{\theta}, \bar{S}) \\ \varepsilon_j^S(S) &= r_j(\bar{\theta}, S) - \varepsilon_j^0. \end{aligned} \quad (5)$$

ε_j^0 is the population growth rate with temperature and silicate constant at their means, while ε_j^S is the main effect of silicate concentration varying around its mean, relative to the “null” conditions $(\bar{\theta}, \bar{S})$ (Means are a natural variance-free null, but other choices are possible, as we explain below). Similarly the main effect of temperature variability is

$$\varepsilon_j^\theta(\theta) = r_j(\theta, \bar{S}) - \varepsilon_j^0 \quad (6)$$

The effect of having variability in both S and θ generally will not equal the sum of the main effects. The difference between

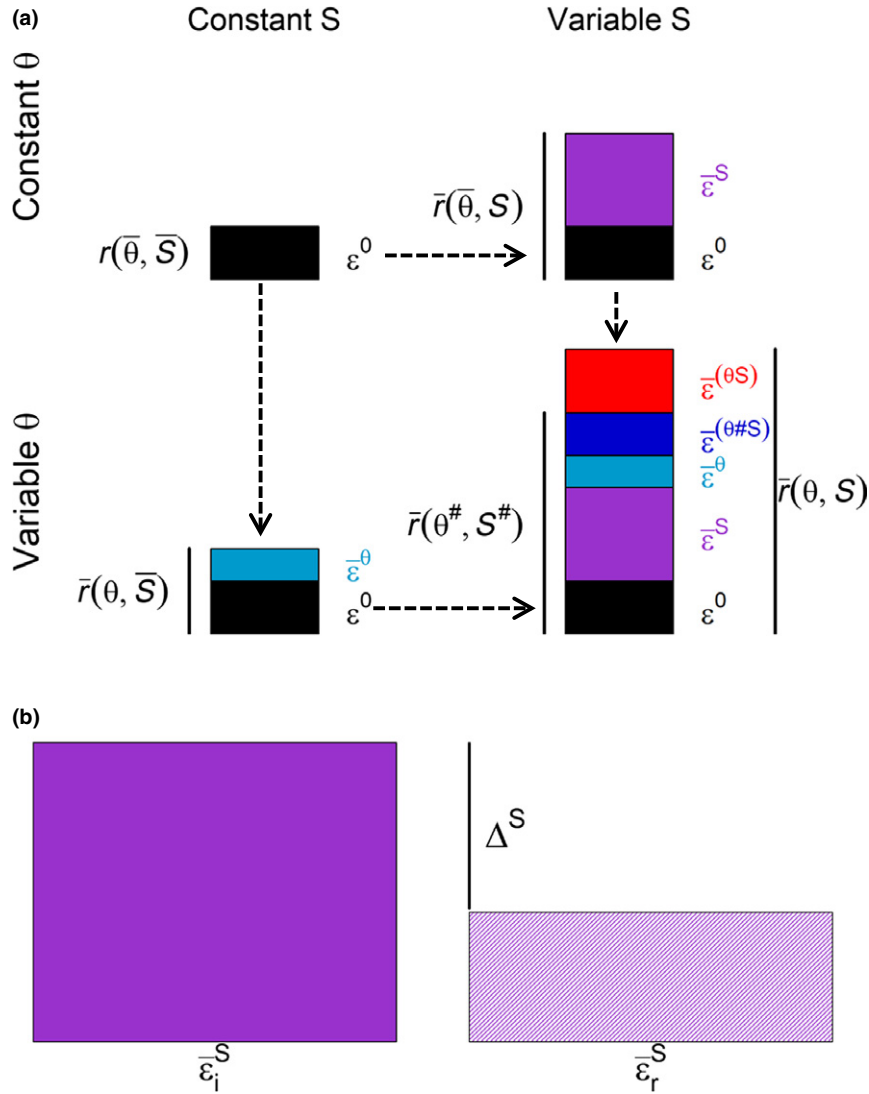


Figure 2 (a) Conceptual representation of how population growth rate is partitioned for any one species, as invader or resident, in the diatoms case study. ϵ^0 (black, top left) is growth rate when temperature θ and silicate S are held constant at their average values. ϵ^S (purple, top right) is the change in growth rate \bar{r} when silicate varies but temperature remains constant at its average value, and ϵ^θ (cyan, bottom left) is the change in growth rate when temperature varies but silicate remains constant at its average value. The further change in growth rate when both vary (bottom right), beyond the combined effect of each varying on its own, is partitioned into the effect of uncorrelated joint variation $\bar{\epsilon}^{(\theta \# S)}$ (dark blue) and the additional effect of correlations between silicate and temperature $\bar{\epsilon}^{(\theta S)}$ (red); $\bar{r}(\theta^\#, S^\#)$ is the long-run growth rate when θ and S vary in an uncorrelated way, given by the first term in the formula for $\bar{\epsilon}_j^{(\theta \# S)}$ in (10). For clarity, this figure is drawn for a hypothetical case where all terms in the partitioning are positive. (b) Conceptual representation of how a term-by-term comparison of invader species i and resident species r defines the contribution Δ of a coexistence mechanism to invader growth rate. This illustrates the case where there is only one resident species. Figure generated by `AwesomeSchematic.R` and `AwesomeSchematicPart2.R`

the actual effect and the sum of main effects is the interaction term,

$$\epsilon_j^{\theta S}(\theta, S) = r_j(\theta, S) - [\epsilon_j^0 + \epsilon_j^S + \epsilon_j^\theta]. \quad (7)$$

Re-arranging (7),

$$r_j(\theta, S) = \epsilon_j^0 + \epsilon_j^S(S) + \epsilon_j^\theta(\theta) + \epsilon_j^{\theta S}(\theta, S). \quad (8)$$

Averaging both sides of (8) as in eqn (4) gives a partition of average population growth rate into the variance-free growth rate, the main effects of variability in S and in θ , and the interaction between variability in S and in θ :

$$\bar{r}_j = \epsilon_j^0 + \bar{\epsilon}_j^S + \bar{\epsilon}_j^\theta + \bar{\epsilon}_j^{\theta S}. \quad (9)$$

Following analytic MCT, we further decompose $\bar{\epsilon}_j^{\theta S}$ into the effect of variance *per se* in θ and S , and the effect of covariance between them. To accomplish this, let $\bar{\epsilon}_j^{(\theta \# S)}$ denote the expectation of $\epsilon_j^{\theta S}$ (eqn 7) when both have their true univariate distributions but the covariance between them is removed. And, let $\bar{\epsilon}_j^{(\theta S)}$ be the effect of restoring the covariance

Table 1 Calculation formulas for the diatoms case study

Term	Formula	Meaning
$\bar{\theta}$	$\frac{1}{m} \sum_{k=1}^m \theta(t_k)$	Mean temperature
\bar{S}	$\frac{1}{m} \sum_{k=1}^m S(t_k)$	Mean silicate, varying temperature experiment or simulation
\bar{S}^*	$\frac{1}{m} \sum_{k=1}^m S^*(t_k)$	Mean silicate, constant temperature experiment or simulation
$\varepsilon_j^0, \varepsilon_j^*$	$r_j(\bar{S}, \bar{\theta}), r_j(\bar{S}^*, \bar{\theta})$	Population growth rates at mean temperature and silicate
ε_j'	$\varepsilon_j^0 - \varepsilon_j^*$	Effect of fluctuation-driven change in mean S
$\bar{\varepsilon}_j^S$	$\frac{1}{m} \sum_{k=1}^m r_j(\bar{\theta}, S(t_k)) - \varepsilon_j^0$	Main effect of variation in silicate
$\bar{\varepsilon}_j^\theta$	$\frac{1}{m} \sum_{k=1}^m r_j(\theta(t_k), \bar{S}) - \varepsilon_j^0$	Main effect of variation in temperature
$\bar{\varepsilon}_j^{\theta S}$	$\frac{1}{m} \sum_{k=1}^m r_j(\theta(t_k), S(t_k)) - [\varepsilon_j^0 + \bar{\varepsilon}_j^S + \bar{\varepsilon}_j^\theta]$	Interaction of silicate and temperature variation
$\bar{\varepsilon}_j^{(\theta \# S)}$	$\frac{1}{m^2} \sum_{k=1}^m \sum_{i=1}^m r_j(\theta(t_k), S(t_i)) - [\varepsilon_j^0 + \bar{\varepsilon}_j^S + \bar{\varepsilon}_j^\theta]$	Independent variation component of $\bar{\varepsilon}_j^{\theta S}$
$\bar{\varepsilon}_j^{(\theta S)}$	$\bar{\varepsilon}_j^{\theta S} - \bar{\varepsilon}_j^{(\theta \# S)}$	Covariance component of $\bar{\varepsilon}_j^{\theta S}$
$r(V, K, S)$	$\frac{VS}{K+S} - D$	Instantaneous population growth rate as a function of traits and silicate concentration
ε_j^0	$\frac{1}{m} \sum_{k=1}^m r(\bar{V}(\theta(t_k)), \bar{K}, S(t_k))$	Baseline (equal-traits) mean population growth rate
$\bar{\varepsilon}_j^V$	$\frac{1}{m} \sum_{k=1}^m r(V_j(\theta(t_k)), \bar{K}, S(t_k)) - \varepsilon_j^0$	Main effect of differences in $V(\theta)$
$\bar{\varepsilon}_j^K$	$\frac{1}{m} \sum_{k=1}^m r(\bar{V}(\theta(t_k)), K_j, S(t_k)) - \varepsilon_j^0$	Main effect of differences in $K(\theta)$
$\bar{\varepsilon}_j^{VK}$	$\frac{1}{m} \sum_{k=1}^m r(V_j(\theta(t_k)), K_j, S(t_k)) - [\varepsilon_j^0 + \bar{\varepsilon}_j^V + \bar{\varepsilon}_j^K]$	Interaction of V and K differences

$S(t_k)$ are values from a model run or experiment with time-varying temperature, $S^*(t_k)$ from a run or experiment with constant temperature $\bar{\theta}$. In these experiments $S^*(t_k)$ is constant so there is no need to time-average, but in other experiments that might not be true so we give the general formulas. j is the species index in all formulas. Formulas above the dashed line are the environment-centric E decomposition; formulas below the dashed line are the trait-centered T decomposition. To compute the invasion growth rate components Δ_i for species i , the formulas are applied to both species ($j = 1, 2$) using data or model simulations with species j as the invader and the other species resident, and each Δ_j is the difference between corresponding $\bar{\varepsilon}_j$ for the invader and the resident.

$$\bar{\varepsilon}_j^{(\theta S)} = \bar{\varepsilon}_j^{\theta S} - \bar{\varepsilon}_j^{(\theta \# S)}. \quad (10)$$

Ellner *et al.* (2016) evaluated $\bar{\varepsilon}_j^{(\theta \# S)}$ by temporal randomisation to remove correlations; the formulas in Table 1 are the expected value of that process (i.e. the average across infinitely many randomisations). With few main effects the formulas are computationally more efficient, but with many, randomisation may be preferable.

Combining eqns (9) and (10) gives the full decomposition

$$\bar{r}_j = \varepsilon_j^0 + \bar{\varepsilon}_j^S + \bar{\varepsilon}_j^\theta + \bar{\varepsilon}_j^{(\theta \# S)} + \bar{\varepsilon}_j^{(\theta S)}. \quad (11)$$

We call (11) an *E-decomposition* because it decomposes population growth rate into contributions from different aspects of the species' environment.

Next, we compute invader–resident differences by applying the formulas in Table 1 to both species, using $S(t)$ and $\theta(t)$ values from simulations (or experiments) with species 1 invading and species 2 as the resident. Let $\bar{\varepsilon}_{k \setminus j}$ denote a term

computed for species k when species j is the invader, and Δ_j the invader–resident difference between corresponding terms when species j is the invader. For example (with $j = 1$ invading, $k = 2$ resident)

$$\Delta_1^S = \bar{\varepsilon}_{1 \setminus 1}^S - \bar{\varepsilon}_{2 \setminus 1}^S \text{ and } \Delta_1^{(\theta \# S)} = \bar{\varepsilon}_{1 \setminus 1}^{(\theta \# S)} - \bar{\varepsilon}_{2 \setminus 1}^{(\theta \# S)}. \quad (12)$$

Being a resident at steady state, species 2 has $\bar{r}_2 = 0$. We therefore have

$$\bar{r}_1 = \bar{r}_1 - \bar{r}_2 = \Delta_1^0 + \Delta_1^S + \Delta_1^\theta + \Delta_1^{(\theta \# S)} + \Delta_1^{(\theta S)}. \quad (13)$$

The growth rate $\varepsilon_1^0 = r_1(\bar{\theta}, \bar{S})$ still reflects the temperature fluctuations during the experiment, because the resident's response to temperature affects \bar{S} . An alternative, completely fluctuation-independent growth rate is $\varepsilon_1^* = r_1(\bar{\theta}, \bar{S}^*)$, where \bar{S}^* is the mean S in an experiment or simulation at constant temperature $\bar{\theta}$ with species 2 absent. Then $\varepsilon_1^0 = \varepsilon_1^* + \varepsilon_1'$ in (11), where $\varepsilon_1' = \varepsilon_1^0 - \varepsilon_1^*$ is the effect of fluctuation-driven changes in mean S . The term Δ_1^0 in (13) is then replaced by $\Delta_1^* + \Delta_1'$.

Table 2 E-decomposition of coexistence mechanisms for experiments with two diatom species (Descamps-Julien & Gonzalez 2005)

Growth rate contributions	<i>Fragilaria</i> $r_{inv} = 0.061 \text{ d}^{-1}$	<i>Cyclotella</i> $r_{inv} = 0.007 \text{ d}^{-1}$	Stabilising component
Fluctuation-free growth rate, Δ^*	-0.031	0.041	0.005
Fluctuation-driven change in mean S , Δ'	0.020	0.001	0.011
Relative nonlinearity in temperature θ , Δ^θ	0.092	-0.037	0.028
Relative nonlinearity in silicate S , Δ^S	-0.014	-0.001	-0.007
θ, S variance interaction, $\Delta^{(\theta \# S)}$	-0.045	0.000	-0.022
θ, S covariance (storage effect), $\Delta^{(\theta S)}$	0.038	0.003	0.021

Values were calculated from the last 1200 days of a 3600 day simulation, recording 10 values each day. Growth rate contributions (Δ) are invader-resident pairwise differences in the decompositions of invader and resident growth rates; for each species, the sum of all Δ s equals the invasion growth rate r_{inv} . Recall that $\Delta^* + \Delta' = \Delta^0$, as explained following eqn (13). The stabilising component of each is simply the average of the first two numerical columns; equalising components (not tabulated) are the deviations between each species' Δ and that average. The storage effect estimates here differ slightly from Ellner *et al.* (2016) because here we are not using the scaling factors in invader-resident comparisons. Values calculated by `FORCEDCHEMO_FUNC_COVVAR.R`, and `FORCEDCHEMOsubs.R`.

Like the analytic formulas in standard MCT, eqn (13) expresses species 1's invasion growth rate as a sum of contributions from different aspects of the biotic and abiotic environments. (Readers familiar with MCT may ask, where are the scaling factors? We answer that question below, but for this case study it doesn't matter because they are all very near 1 (Ellner *et al.* 2016)). Δ_1^0 is the difference in population growth rates at mean temperature and silicate. Δ_1^S is the difference in the main effects of silicate variability, between species 1 as invader and species 2 as resident. This difference results from the nonlinear response of cell division rate to silicate concentration so we call it *relative nonlinearity in silicate*. Similarly, Δ_1^θ measures relative nonlinearity in temperature. We call $\Delta_1^{(\theta S)}$ the storage effect because, as in standard MCT, storage effect is the contribution to population growth rate of covariance between the environment and competitive factors determining r (Ellner *et al.* 2016).

However, inexact correspondence with MCT is unavoidable. MCT terms such as “storage effect” or “relative nonlinearity” refer to terms in small variance approximations that we do not use. Similarly, our variance-interaction terms such as $\bar{e}^{(\theta \# S)}$ are absent in standard MCT, because they come from the third- and higher-order terms that are dropped from MCT's Taylor series approximations.

There are also avoidable differences reflecting choices in applying our method. We prefer to decompose in terms of environmental drivers, here temperature and silicate, whereas standard MCT uses an environment-dependent parameter E such as $V(\theta)$, and a competition parameter C , such as the effect of silicate limitation on cell division rate. In section SI.3, we present an alternate analysis of this case study using E and C parameters. Also, standard MCT lumps terms that we keep separate. Specifically, the baseline growth rate (\bar{r} or $\bar{\lambda}$) in MCT that is described as representing “variation-independent coexistence mechanisms” (Chesson 2000a, p.224) and “mechanisms operating on a shorter time-scale than the unit of time considered explicitly in the model” (Chesson 1994, p. 249) includes the direct effects of fluctuations in E that are not mediated through variance in C or E, C covariance. Failure to appreciate that led to our incorrect conclusion (Ellner *et al.* 2016) that relative nonlinearity of competition must be important in this community.

Results from our E-decomposition are given in Table 2. For each species, the growth rate contributions (Δ s) add up (by definition) to equal the invasion growth rates. A negative invasion growth rate would imply that the species could not invade the other. We can therefore assay the importance of each mechanism for coexistence by asking what the invasion growth rate would become if the corresponding contribution is set to zero. The stabilising component of each mechanism is defined to be its average contribution across species, and the equalising components are each species' deviation from the average (as in Chesson (2003), but we set the scaling coefficients to 1). Thus, a component is stabilising (in this sense) simply if its average contribution to the invasion growth rates of all species is positive. (Contrary to intuition, by this definition a mechanism can be stabilising even if it does not benefit each species when rare. Despite this concern, we follow current usage rather than inventing new terms).

The results in Table 2 show that storage effect, relative nonlinearity in temperature, fluctuation-driven changes in mean silicate, and the fluctuation-free null growth rate are all stabilising, i.e. they all increase average invasion growth rate. *Cyclotella* has positive invasion growth rate because its fluctuation-free growth rate is large enough to offset the negative contribution from relative nonlinearity in temperature. *Fragilaria*'s positive invasion growth rate is crucially dependent on the positive contribution from relative nonlinearity in temperature, without which its invasion growth rate would become negative. The two largest fluctuation-dependent terms for *Fragilaria* – relative nonlinearity in temperature, and the interaction between temperature and silicate variability – are absent from standard MCT for the reasons explained above.

The E-decomposition identifies how different features of the species' biotic and abiotic environments promote or impede coexistence, given how species respond to their environment. Using the same approach we can additionally do a species-centric decomposition to provide complementary information about which attributes of species promote or impede coexistence under the abiotic and biotic conditions that they experience. We call this a T-decomposition, T standing for “traits.”

Given the environment ($S(t), \theta(t)$), population growth rates of *Fragilaria* and *Cyclotella* differ because they have different half-saturation constants K_j (which are constant over the experiment's

temperature range) and different temperature-dependent maximum division rates $V_j(\theta)$, which determine their response to silicate concentration. We therefore decompose the invader and resident population growth rates into the main effects of differences in K and in V , and their interaction, relative to the “null” growth rate that results if both species are given the average responses $\bar{K} = 0.5(K_1 + K_2)$, $\bar{V}(\theta) = 0.5(V_1(\theta) + V_2(\theta))$. Formulas for the terms are in Table 1; we again use Δ to denote a difference between invader and resident terms.

The main effects (Table 3) reiterate the biology: *Cyclotella* benefits from its lower K and is harmed by its lower V at high temperatures, and the reverse is true for *Fragilaria*. Both of these are stabilising. More unexpected is the size of the interaction terms. Although they are not necessary for coexistence, the invasion growth rates would be quite different without them (nearly an order of magnitude larger for *Cyclotella*).

This T-decomposition could also be applied separately to each component in the E-decomposition, to show which traits generate each growth rate component, but we do not pursue that here.

GENERAL THEORY

General functional decomposition

The E- and T-decompositions in our diatoms case study (Table 1) are examples of a general *functional decomposition* applicable to any collection of two or more processes, mechanisms, or species differences affecting population growth rate.

The first step is to select the features of interest. A “feature” is some aspect of reality that affects population growth rates. The features in our diatom E-decomposition were variance in S and θ and covariance between them. The features in our diatom T-decomposition were interspecific differences in V and K . In Angert *et al.* (2009) the features were temporal variation in seed germination fraction, temporal variation in seedling growth and survival, and temporal variation in competition. Features in the case studies below include the presence of plant–soil feedbacks, and facilitation of other species by each member of a community.

The decomposition consists of breaking up the long-run growth rate of each species (as invader and then as resident) into (0) a null growth rate in the absence of all selected features; (1) a set of “main effect” terms representing the effect of adding one and only one feature; (2) a set of two-way interaction terms representing the effect of adding each possible pair

of features, above and beyond the sum of their main effects; (3) and so on, until all features are represented. The null term can contribute to coexistence when it includes the stabilising effects of features that were not selected for the decomposition. Term-by-term invader–resident comparisons then measure the contribution of each growth rate component to invasion growth rates.

In section SI.4, we give a mathematical definition of the decomposition; here we explain it by describing general E- and T-decompositions of population growth rate. We then discuss invader–resident comparisons, and define stabilising and equalising components.

General E-decomposition

An E-decomposition follows standard MCT in focusing on coexistence maintained by environmental variability. The features are variances and covariances of biotic or abiotic variables affecting population growth rates. For example, if population growth rate r is a function of environmental variables X , Y , and Z , we write

$$r(X, Y, Z) = \varepsilon^0 + \varepsilon^X + \varepsilon^Y + \varepsilon^Z + \varepsilon^{XY} + \varepsilon^{YZ} + \varepsilon^{XZ} + \varepsilon^{XYZ}. \quad (14)$$

The null growth rate $\varepsilon^0 = r(\bar{X}, \bar{Y}, \bar{Z})$ is the growth rate when all variables or traits are set to their averages. Terms with superscripts represent the marginal effects of letting all superscripted variables vary while fixing all other variables at their average values — they are the difference between long-run growth with all superscripted variables free, and the sum of all lower-order terms (i.e. terms where fewer variables are free to vary). For example,

$$\begin{aligned} \varepsilon^X &= r(X, \bar{Y}, \bar{Z}) - \varepsilon^0 \\ \varepsilon^{XY} &= r(X, Y, \bar{Z}) - [\varepsilon^X + \varepsilon^Y + \varepsilon^0] \\ \varepsilon^{XYZ} &= r(X, Y, Z) - [\varepsilon^{XY} + \varepsilon^{YZ} + \varepsilon^{XZ} + \varepsilon^X + \varepsilon^Y + \varepsilon^Z + \varepsilon^0]. \end{aligned} \quad (15)$$

Taking expectations of all terms in (14) and (15), each r becomes \bar{r} and each ε becomes $\bar{\varepsilon}$. Mean population growth rate is thus decomposed into the main effects of variation in each argument, and interactions among variation in different arguments.

It is important to recall that each ε only includes the direct effects of variation in superscripted variables; for example $\bar{\varepsilon}(X, Y, \bar{Z})$ is computed using the distribution of X and Y , and

Table 3 T-decomposition of coexistence mechanisms for experiments with two diatom species (*Descamps-Julien & Gonzalez, 2005*), using the formulas in Table 1

Growth rate contributions	<i>Fragilaria</i> $r_{inv} = 0.061 \text{ d}^{-1}$	<i>Cyclotella</i> $r_{inv} = 0.007 \text{ d}^{-1}$	Stabilising component
Difference in K , Δ^K	−0.057	0.079	0.011
Difference in V , Δ^V	0.079	−0.017	0.031
Interaction, Δ^{VK}	0.039	−0.054	−0.007

All species necessarily have the same null growth rate so the corresponding invader–resident comparison Δ^0 is zero. The stabilising component of each growth rate contribution is the average of the first two numerical columns; equalising components (not tabulated) are the deviations between each species' Δ and that average. Tabulated values are calculated in `ForcedChemo_TraitDecomp.R` and `ForcedChemoSubs.R`.

the value of \bar{Z} that occurs when Z is fluctuating. Similarly, in a structured population model, presence versus absence of fluctuations in any variable will affect population structures, but terms for the effect of those fluctuations would be evaluated using population growth rates computed with the population structures observed under natural conditions, as in our *Structured populations* example below. If the effect of changes in population structure is a feature of interest, it can be included in the decomposition, producing terms for its main effect, interactions, and so on. *Different decisions about which features are of interest lead to different decompositions.*

Next, we can do a *covariance decomposition* of any $\bar{\varepsilon}$ term with multiple free variables, breaking it up into the effects of their variation *per se* and effects of covariances among them. Our notation convention is that a superscript with parentheses is a term in this secondary decomposition, and a # symbol separates subsets of variables that have within-subset covariation preserved while variables separated by one or more # symbols vary independently. For example, if r is a function of X, Y, Z, W , and U , then $\bar{\varepsilon}^{(XY\#ZW)}$ is the expectation of r with U set at its average value, X and Y covarying, Z and W covarying, and X and Y independent of Z and W , *minus* all terms with the same non-constant variables but fewer covariances.

In the diatoms example, we had $\bar{\varepsilon}^{XY} = \bar{\varepsilon}^{(X\#Y)} + \bar{\varepsilon}^{(XY)}$. Higher-order terms are broken down similarly with higher-order covariance terms representing the additional affect of the correlations present, beyond the combined effect of all possible lower-order covariance terms. For example,

$$\bar{\varepsilon}^{(XYZ)} = \bar{\varepsilon}^{XYZ} - [\bar{\varepsilon}^{(XY\#Z)} + \bar{\varepsilon}^{(X\#YZ)} + \bar{\varepsilon}^{(XZ\#Y)} + \bar{\varepsilon}^{(X\#Z\#Y)}]. \quad (16)$$

and $\bar{\varepsilon}^{(XY\#Z)}$ is the effect on population growth rate of restoring the observed covariance between X and Y , with Z independent of both, relative to the growth rate when X, Y and Z are mutually independent.

Researchers may wish to decompose only some terms into effects of independent variation and covariation. Alternatively, covariance between two variables could be separated into several components (e.g. rainfall-competition covariance for small individuals and for large individuals), leading to a finer decomposition of the contribution of covariance to population growth.

General T-decomposition

The features in a T-decomposition are attributes that differ among species. The null growth rate is the population growth rate that results from giving all species the across-species average value for each trait; this depends on the scale of measurement, e.g. length versus log length. The main effects are restoring one trait to its true value in all species. Let Θ denote the vector of parameters or variables characterising the biotic and abiotic environment, possibly time-varying. In the diatoms example, Θ is $(S(t), \theta(t))$ during the experiment. Then if the traits are X_1, X_2, \dots, X_n the null growth rate is

$$\varepsilon^0 = \mathbb{E}_{\Theta} r(\bar{X}_1, \bar{X}_2, \dots, \bar{X}_n, \Theta) \quad (17)$$

the hypothetical long-term growth rate when each species has the average value for all traits, but the environment (abiotic and biotic) varying as it actually did with observed traits. The main effect for trait J is

$$\varepsilon^J = \mathbb{E}_{\Theta} r(\bar{X}_1, \dots, X_J, \dots, \bar{X}_n, \Theta) - \varepsilon^0. \quad (18)$$

The interaction between traits J and K is

$$\varepsilon^{JK} = \mathbb{E}_{\Theta} r(\bar{X}_1, \dots, X_J, \dots, X_K, \dots, \bar{X}_n, \Theta) - [\varepsilon^J + \varepsilon^K + \varepsilon^0] \quad (19)$$

and so on, exactly as in an E-decomposition. Invader-resident term-by-term differences quantify the contribution of each term to the invading species' advantage when rare. Note, we again measure only the direct effects of each trait by evaluating terms using Θ taken from data or simulations where all traits have their observed values.

Alternatively, the decomposition can include (as additional features) indirect effects mediated by species' effect on their environment. In SI.9.4 we illustrate this alternative, with indirect effects mediated by population structure as the additional feature.

Analogous to covariance decomposition, a step from average to true trait values could be broken into two steps, by considering the intermediate situation where traits vary among species without trait-trait correlations. We do not pursue this refinement because many terms will often be zero (see SI.6.)

Resident weightings in invader-resident comparisons

When comparing invader and resident population growth rates, we generally weight all residents equally. Instead, or additionally, one can ask how an invader gains an advantage over each resident individually. Or, to recover the canonical MCT mechanisms, residents can be weighted by the *scaling factors* q_{ir} (Chesson 1994, 2000a). These alternative resident weightings are as follows:

$$\begin{aligned} \Delta_{i,=} &= \bar{\varepsilon}_{i\setminus i} - \frac{1}{S-1} \sum_{r \neq i} \bar{\varepsilon}_{r\setminus i} && \text{Equal weight} \\ \Delta_{i,k} &= \bar{\varepsilon}_{i\setminus i} - \bar{\varepsilon}_{k\setminus i}, k \neq i && \text{Pairwise with resident species } k \\ \Delta_{i,q} &= \bar{\varepsilon}_{i\setminus i} - \sum_{r \neq i} q_{ir} \bar{\varepsilon}_{r\setminus i} && \text{Scaling factors} \end{aligned} \quad (20)$$

where S is the total number of species. The three weightings are three different but equally valid ways of breaking one number (invasion growth rate of species i) into a sum of interpretable components. The same $\bar{\varepsilon}_{i\setminus i}$ and $\bar{\varepsilon}_{r\setminus i}$ values are used in all weightings, calculated from data or simulations where species i is invading a community with all other species resident. Thus, a pairwise comparison between invader A and resident B will include effects of the other residents C, D, E, etc.

In many models, the scaling factors q_{ir} do not exist (for example, when there are more limiting factors than species, Barabás *et al.* (2018)). In other cases, they are not unique (Ellner *et al.* 2016), or they can become negative (Snyder *et al.* 2005, p. E92), turning an invader-resident "difference" into a

Table 4 Calculation formulas for analysis of spatial coexistence mechanisms

Term	Formula	Meaning
$\bar{N}_j(t)$	$Q^{-1} \sum_{x=1}^Q n_{j,x}(t)$	Average within-patch abundance
$v_{j,x}(t)$	$n_{j,x}(t)/\bar{N}_j(t)$	Relative abundance in patch x
$\lambda_{j,x}(t)$	$\lambda_j(E_{j,x}(t), C_{j,x}(t))$	Per-capita fitness of patch x individuals
$\bar{v}_j(t), \bar{\lambda}_j(t)$	$Q^{-1} \sum_{x=1}^Q v_{j,x}(t), Q^{-1} \sum_{x=1}^Q \lambda_{j,x}(t)$	Spatial averages of v, λ
$\bar{E}_j(t), \bar{C}_j(t)$	$Q^{-1} \sum_{x=1}^Q E_{j,x}(t), Q^{-1} \sum_{x=1}^Q C_{j,x}(t)$	Average environment and competition factors
$Cov_x(v_j(t), \lambda_j(t))$	$[Q^{-1} \sum_{x=1}^Q v_{j,x}(t)\lambda_{j,x}(t)] - \bar{v}_j(t)\bar{\lambda}_j(t)$	Fitness-density covariance
$\varepsilon_j^0(t)$	$\lambda_j(\bar{E}_j(t), \bar{C}_j(t)) - 1$	Baseline (zero variance) population growth rate
$\bar{\varepsilon}_j^C(t)$	$[Q^{-1} \sum_{x=1}^Q \lambda_j(\bar{E}_j(t), C_{j,x}(t)) - 1] - \varepsilon_j^0(t)$	Main effect of spatial variance in C
$\bar{\varepsilon}_j^E(t)$	$[Q^{-1} \sum_{x=1}^Q \lambda_j(E_{j,x}(t), \bar{C}_j(t)) - 1] - \varepsilon_j^0(t)$	Main effect of spatial variance in E
$\bar{\varepsilon}_j^{EC}$	$[\bar{\lambda}_j(t) - 1] - [\varepsilon_j^0 + \bar{\varepsilon}_j^E + \bar{\varepsilon}_j^C]$	Interaction of E and C variation
$\bar{\varepsilon}_j^{(E\#C)}$	$[Q^{-2} \sum_{x=1}^Q \sum_{y=1}^Q \lambda_j(E_{j,x}(t), C_{j,y}(t)) - 1] - [\varepsilon_j^0 + \bar{\varepsilon}_j^E + \bar{\varepsilon}_j^C]$	Independent variation component of $\bar{\varepsilon}_j^{EC}$
$\bar{\varepsilon}_j^{(EC)}$	$\bar{\varepsilon}_j^{EC} - \bar{\varepsilon}_j^{(E\#C)}$	Covariance variance component of $\bar{\varepsilon}_j^{EC}$
<hr/>		
$\tilde{\lambda}_j(t)$	(total pop. at $t + 1$)/(total pop. at t)	Population growth rate
$\bar{\varepsilon}_j^0(t)$	$\tilde{\lambda}_j(F = 0, c_{ii} = 1)(t) - 1$	Baseline growth rate
$\bar{\varepsilon}_j^c(t)$	$[\tilde{\lambda}_j(F = 0, c_{ii} = 0.5)(t) - 1] - \varepsilon_j^0(t)$	Main effect of plant–soil feedbacks c
$\bar{\varepsilon}_j^F(t)$	$[\tilde{\lambda}_j(F \neq 0, c_{ii} = 1)(t) - 1] - \varepsilon_j^0(t)$	Main effect of local retention F
$\bar{\varepsilon}_j^{Fc}(t)$	$[\tilde{\lambda}_j(F \neq 0, c_{ii} = 0.5)(t) - 1] - [\varepsilon_j^c(t) + \bar{\varepsilon}_j^F(t) + \bar{\varepsilon}_j^0(t)]$	Interaction of c and F

j is the species index in all formulas. Formulas above the dashed line are a spatial E-decomposition for the Chesson (2000a) model with purely spatial variation. Formulas below the dashed line are a T-decomposition for the Petermann *et al.* (2008) model. The argument $c_{ii} = 1$ means that all entries of c equal 1 (no plant–soil feedbacks), while $c_{ii} = 0.5$ means that all diagonal entries of c equal 0.5 (negative plant–soil feedbacks). Similarly, F is either zero for all species (no local retention) or has the same positive value for all species. Formulas for the various $\lambda_j(t)$ are given in section SI.8. All formulas here are for spatial “snapshot” data at one time t , or one observation of population growth. If data at multiple times are available, each ε or $\bar{\varepsilon}$ term is calculated for each time point, then averaged over time.

weighted sum of invader and resident growth rates. Ellner *et al.* (2016) discuss these conceptual difficulties in detail. And in experimental practice, numerical values of the scaling factors may be very sensitive to choices that should be immaterial in the analytic theory (Letten *et al.* 2018). Because of these issues, we generally avoid use of the scaling factors.

Stabilising and equalising components

We define the *stabilising* component of any mechanism (i.e. of any Δ term) to be the average of that term across all species as invader. The *equalising* component of a mechanism is, for each species, its deviation from the average.

Chesson (2000b, 2003) uses similar definitions but first scales population growth rates such that the scaled invasion growth rates sum to zero in the absence of the coexistence mechanisms considered, which eliminates any stabilising component apart from those mechanisms. While this is aesthetically pleasing, scaling coefficients with this property are only certain to exist when species compete for a single limiting

factor (Barabás *et al.* 2018, section 2.7), so we use unscaled population growth rates. The stabilising component of the null term, if it is non-zero, reflects the stabilising component of all features not included in the chosen decomposition.

APPLICATIONS

Spatial coexistence mechanisms in the Chesson (2000) model

To highlight our approach’s generality, we show how it can be used to analyze coexistence in the Chesson (2000a) model for environments with purely spatial variation. Let $n_{j,x}$ denote abundance of species j in patch $x = 1, 2, \dots, Q$. The expected contribution of patch x individuals to the global population at time $t + 1$ is given by

$$\lambda_j(E_{j,x}(t), C_{j,x}(t))n_{j,x}(t), \quad (21)$$

where $E_{j,x}$ and $C_{j,x}$ are the environment and competition factors affecting species j in patch x . Apart from space the

population is unstructured, so λ_j includes new recruits and survivors. The function λ_j is the same for all patches; intrinsic patch differences are incorporated into $E_{j,x}$ and $C_{j,x}$. For example, $C_{j,x}$ could be population $n_{j,x}$ divided by local carrying capacity $K_{j,x}$.

Chesson (2000a) showed that (with notation as defined in Table 4) the total population of species j has annual growth rate

$$\tilde{\lambda}_j(t) = \bar{\lambda}_j(t) + Cov_x(v_j(t), \lambda_j(t)), \quad (22)$$

the sum of average patch-specific fitness and fitness-density covariance.

A growing population is represented by $\tilde{\lambda} > 1$. To parallel the continuous-time case we therefore make the response variable $\tilde{\lambda} - 1$ as in Chesson (2000a). We can then decompose $\tilde{\lambda}_j(t) - 1$ using the two-factor case of our general E-decomposition. The result is

$$\begin{aligned} \tilde{\lambda}_j(t) - 1 = & \bar{\epsilon}_j^0(t) + \bar{\epsilon}_j^E(t) + \bar{\epsilon}_j^C(t) + \bar{\epsilon}_j^{(E\#C)}(t) + \bar{\epsilon}_j^{(EC)}(t) \\ & + Cov_x(v_j(t), \lambda_j(t)). \end{aligned} \quad (23)$$

With $\bar{\epsilon}$ here denoting a spatial average at a particular time. Formulas for all terms are in Table 4.

As in the non-spatial case, mechanism contributions (Δ s) are the difference between corresponding $\bar{\epsilon}$ terms for invader and resident species: e.g. $\Delta_j^C = \bar{\epsilon}_{jV}^C - \frac{1}{S-1} \sum_{k \neq j} \bar{\epsilon}_{kV}^C$. The invader-resident difference in Cov_x terms is *growth-density covariance* and the difference in $\epsilon^{(EC)}$ terms corresponds to the *spatial storage effect* as defined by Chesson (2000a). The term Δ^C aligns closely with spatial relative nonlinearity of competition, and Δ^E is spatial relative nonlinearity in environment (absent from the standard decomposition).

Case study: Janzen-Connell effect in grasslands

Petermann *et al.* (2008) developed and parameterised experimentally a spatial model for species coexistence in grasslands through local plant-soil feedbacks mediated by soil microbes. The landscape consists of sites, each containing a single legume, grass, or forb individual, which have a species-specific death rate. Individuals produce seeds each year, which are retained locally (in the parent's site) with probability F (*local retention fraction*), and otherwise disperse at random across all sites. At sites becoming open through death of the occupant, there is lottery competition among seeds, with a twist: a seed's probability of capturing the site depends not only on the identity of the seed and competing seeds, but on the identity of the adult previously occupying the site. Because of persisting species-specific pathogens, a seed is less likely to win a site previously held by a conspecific adult. Specifically, the probability that a seed of type i will capture a site formerly occupied by an adult of type j is $c_{ij} s_i / \sum_k c_{kj} s_k$, where s_k is the number of seeds of type k in a site. Based on experiments, Petermann *et al.* modelled the soil pathogen effect by assuming $c_{ii} = 0.5$, $c_{i \neq j} = 1$.

In individual-based simulations of the model, Petermann *et al.* (2008) observed long-term coexistence of all species, so long as the local retention fraction F is not too high. To better understand this result in terms of underlying mechanisms,

we used a decomposition in which the selected features are the presence of local retention ($F > 0$ vs. $F = 0$) and plant-soil feedback ($c_{ii} = 0.5$ vs. $c_{ii} = 1$). We present here simulation-based results; the decomposition can also be done analytically (section SI.8).

Table 4 gives the formulas for decomposing population growth rates; invader-resident comparisons Δ_j used equal weighting of residents. When estimating the $\tilde{\lambda}_j(t)$, we kept the resident population totals at the steady-state values they assume in the presence of both local seed retention and plant-soil feedbacks (estimated by running a long simulation for each pair of species as residents, with no invader). Each of the $\tilde{\lambda}_j(t)$ in Table 4 was estimated by repeatedly simulating one time-step forward from those resident densities with the invader occupying one additional site, and averaging over replicates. This ensures that, e.g. Δ^c measures only the direct effects of having or not having plant-soil feedbacks, not the indirect effects mediated by changes in resident species abundance due to presence or absence of plant-soil feedbacks.

Figure 3 shows the estimated contributions to invasion growth rate of local retention (Δ^F), plant-soil feedbacks (Δ^c), and their interaction, as a function of local retention (F). As expected, local retention in the absence of plant-soil feedbacks (Δ^F) has little effect, but local retention combined with plant-soil feedbacks (Δ^{Fc}) reduces invader growth rates (anti-stabilising). Plant-soil feedbacks alone (Δ^c) increase the invader growth rate for all species (stabilising), because in the absence of local retention almost no invader seeds fall into an invader-occupied site, while many resident seeds fall into a site occupied by a conspecific. These results let us understand Petermann *et al.*'s findings as follows: coexistence occurs when the stabilising effect of plant-soil feedbacks is not dominated by the anti-stabilising effect of the interaction between plant-soil feedbacks and local seed retention.

Facilitation and coexistence among beach grasses

Many ecologists have called for coexistence theory to better integrate positive interactions alongside the traditional focus on competition (Bruno *et al.* 2003; McIntire & Fajardo 2013; Bulleri *et al.* 2015). While mechanistic consumer-resource models can readily incorporate positive interactions (Gross 2008), facilitation in phenomenological, Lotka-Volterra type competition models poses problems for coexistence theory. Positive interspecific interactions (negative competition coefficients, in the Lotka-Volterra convention) can lead to infinite population growth (Gause & Witt 1935), and additionally Chesson's (2013) expression for niche overlap becomes invalid as it can include the square root of a negative number. In this section we show how our approach makes it possible to quantify the impacts of positive interactions on coexistence in Lotka-Volterra type models.

Zarnetske *et al.* (2013) used field and experimental data to parameterise a Lotka-Volterra model for beach grass communities in the US Pacific Northwest comprised of *Ammophila arenaria* (AA), *Ammophila breviligulata* (AB), and *Elymus mollis* (EM). Interactions among these species are a mix of competition and facilitation. The Zarnetske *et al.* (2013) model is

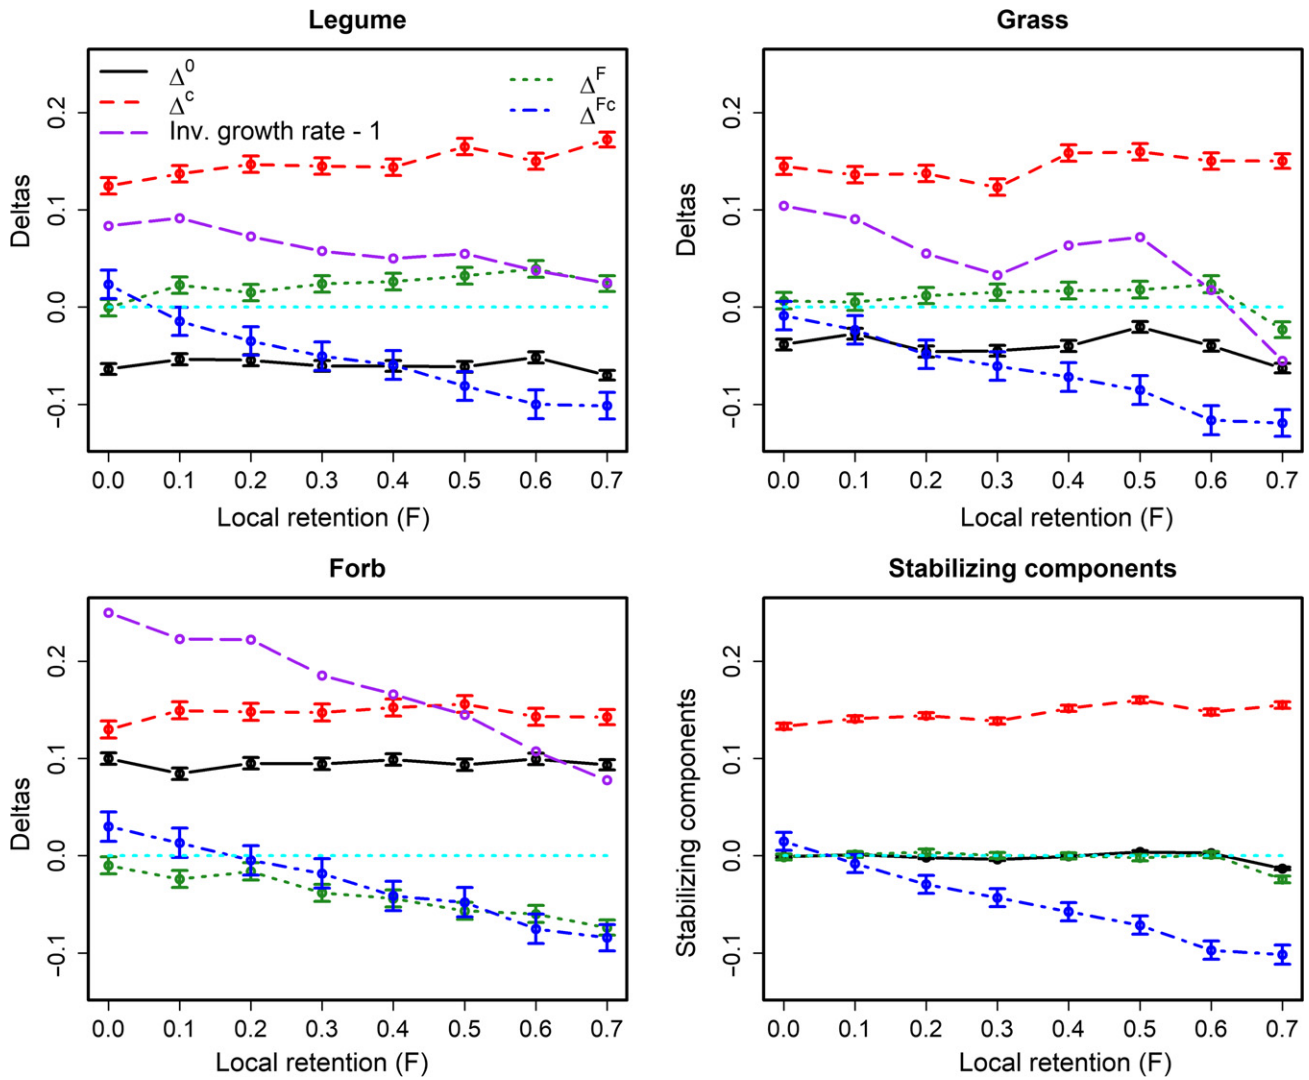


Figure 3 Strength of coexistence mechanisms as a function of the probability of local retention F in the plant–soil feedback model of Petermann *et al.* (2008). The dotted cyan curve at 0 is drawn to provide a visual baseline. We only consider $F \leq 0.7$ because stable coexistence is lost when F is slightly larger. Comparison with the analytic decomposition (section SI.8) shows that values for grass at $F = 0.7$ are slightly inaccurate, probably because one resident (legume) has very low steady-state abundance. Figure generated by `PetermannPartitioning.R` and `plotPetermannDeltas.R`

$$\frac{dx_i}{dt} = r_i x_i \left(1 - K_i^{-1} \sum_{j=1}^3 \alpha_{ij} x_j \right) \quad (24)$$

with $\alpha_{ii} = 1$ for $i = 1, 2, 3$. We used parameter estimates from their Table S3-A for low sand input (because those were most robust to the assumed time to reach equilibrium), and the maximum assumed time to equilibrium (because those estimates use the weakest assumptions about unmeasured densities). The estimated competition coefficients are

$$\alpha = \begin{pmatrix} \text{by AA} & \text{by AB} & \text{by EM} \\ 1.000 & -0.214 & 0.131 \\ -0.358 & 1.000 & 0.370 \\ 0.089 & -0.139 & 1.000 \end{pmatrix} \begin{matrix} \text{on AA} \\ \text{on AB} \\ \text{on EM} \end{matrix} \quad (25)$$

Negative coefficients indicate facilitation: AA facilitates AB, AB facilitates AA and EM.

We performed a T-decomposition with two main effects, facilitation by AA and facilitation by AB, relative to the no-facilitation situation in which all negative α_{ij} are set to 0. As in the T-decomposition for competing diatoms, we change traits (α values) but leave the environment (equilibrium species abundances) “as is”. So for example, paralleling the diatom T-decomposition formulas in Table 1, with species i invading we have

$$\varepsilon_j^0 = r_j \left(1 - K_j^{-1} \sum_{m \neq i} \tilde{\alpha}_{jm}^0 \bar{x}_m \right), \quad j = 1, 2, 3. \quad (26)$$

where \bar{x}_m is the equilibrium value of x_m when $x_i = 0$ in (24) with the estimated competition coefficients, and $\tilde{\alpha}^0$ are modified coefficients with facilitation removed by setting all negative values to zero. Because per-capita population growth rates are linear in the competition coefficients, all interaction terms are zero.

Table 5 Results for the beach grass model of Zarnetske *et al.* (2013) with different resident weightings

	<i>Ammophila arenaria</i> (AA) $\bar{r}_{inv} = 0.17/\text{month}$				<i>Ammophila breviligulata</i> (AB) $\bar{r}_{inv} = 0.25/\text{month}$				<i>Elymus mollis</i> (EM) $\bar{r}_{inv} = 0.36/\text{month}$			
	Equal	q_{ir}	–AB	–EM	Equal	q_{ir}	–AA	–EM	Equal	q_{ir}	–AA	–AB
Δ^0	0.19	0.14	0.14	0.25	0.19	0.19	0.19	0.19	0.24	0.62	0.23	0.26
Δ^{AA}	0.00	0.00	0.00	0.00	0.06	0.06	0.06	0.06	–0.04	–0.38	0.00	–0.08
Δ^{AB}	–0.02	0.03	0.03	–0.08	0.00	0.00	0.00	0.00	0.15	0.11	0.12	0.18

Growth rate contribution Δ^0 is invader–resident comparison of “null” growth rates when all facilitation is eliminated, and Δ^{AA} and Δ^{AB} are the main effects of facilitation by AA and AB, respectively. The different invader–resident weightings (column headings) are defined in eqn (12); “Equal” denotes equal weighting of residents, q_{ir} denotes weighting by Chesson’s scaling factors, and “–XY” denotes a pairwise comparison between the species in the table heading as the invader, with competing species XY as resident. The scaling factors q_{ir} are derived in sect. SI.7. Tabulated values are calculated by `Beach-grass_TraitDecomp.R` and `BeachgrassFuns.R`.

Table 6 Main effects and interactions of niche differences impacting Recruitment, Growth, and Survival in the four dominant species in Idaho sagebrush steppe

Species	\bar{r}_{inv} (1/yr)	Δ^0	Δ^R	Δ^G	Δ^S	Δ^{RG}	Δ^{RS}	Δ^{GS}	Δ^{RGS}
<i>Artemisia tripartita</i>	0.015	–1.16	0.00	0.55	0.68	–0.01	–0.02	–0.03	0.01
<i>Hesperostipa comata</i>	0.25	0.12	0.13	–0.06	0.09	–0.01	–0.03	0.01	0.00
<i>Poa secunda</i>	0.43	0.15	0.29	0.03	0.04	–0.04	–0.06	0.00	0.01
<i>Pseudoroegneria spicata</i>	0.20	0.35	0.10	–0.20	–0.04	–0.00	–0.01	0.01	0.00
Stabilising			0.13	0.08	0.19	–0.02	–0.03	–0.00	0.01

\bar{r}_{inv} is invasion growth rate when niche differences at all stages are present, and Δ_0 is the no-niche-differences growth rate. Tabulated values were calculated by `partition_simulations.R` and scripts that it sources.

Results (Table 5) were calculated for all three resident weightings (eqn 20). For each species, each column is a partition of the species’ invasion growth rate r_{inv} (i.e. each column sums to r_{inv} for the species, apart from rounding errors). For AA and AB the three weightings give equivalent results: facilitation is relatively unimportant, and all species would persist without facilitation because intraspecific competition outweighs inter-specific competition (all resident weightings give identical results for AB, because with AB as a rare invader, neither of the residents AA and EM experiences any direct impacts of facilitation). For EM, the equal weighting and pairwise comparisons are similar: AB facilitation contributes 1/3 to 1/2 of invasion growth rate but is not essential for persistence, while AA facilitation may have a small negative effect. But in the comparison using scaling factors, Δ_{q_i} , facilitation by AA makes a large negative contribution. This occurs because $q_{EM,AB} = 4.7$, so facilitation of AB by AA (which is a detriment to EM) is weighted very heavily. The result of this large q_{ir} value is that a feature of low importance in both pairwise comparisons between EM and one resident becomes important when EM is compared to both residents weighted by their scaling factors. Such outcomes are one of our reasons for emphasising equal weight and pairwise invader–resident comparisons.

Structured populations: process-specific niche differences in sagebrush steppe

Chu & Adler (2015) used multispecies size-structured integral projection models (IPM) to compare the importance of

stabilising features that act independently on the recruitment (R), survival (S), and growth (G) of the four dominant species in a sagebrush steppe community. Here we revisit their analysis using our approach. The main difference is that the Chu & Adler (2015) analysis was based on invader growth rates, not invader–resident comparisons. See Chu & Adler (2015) for full details of the models, data, and parameter estimation; the main features of the model are summarised in SI.9.

The key model feature for our analysis is that for each demographic rate $V = R, S$ or G there is a matrix α^V of interaction coefficients that determine the impact on that rate of competition with neighbouring plants of each species,

$$\bar{w}_j^V(u, t) = \sum_{k=1}^4 \alpha_{jk}^V w_{jk}(u, t). \quad (27)$$

Here $w_{jk}(u, t)$ is a measure of average species- k cover within the competition neighbourhood of a size- u individual in species j , and $\bar{w}^V(u, t)$ is the overall impact of all neighbours on demographic rate V . Differences between intra- and inter-specific interaction coefficients in α^V generate measures of pairwise “niche difference,” defined in SI.9.

We performed a T-decomposition with three main effects, the niche differences affecting recruitment, growth, and survival, represented by the competition coefficients α_{ij}^V . Following Chu & Adler (2015), the “no-niche” state is defined by modifying all between-species α_{ij} values (in α^R, α^S and α^G) so that each pairwise niche overlap ρ_{ij} equals 1, without changing the fitness differences κ_j/κ_i . Formulas for ρ, κ and how they were adjusted are in SI.9. In a two-species Lotka-Volterra

community, the values of ρ_{ij} and κ_{ij} determine whether or not all species coexist stably. But with more species, indirect effects can be important (e.g. A facilitates B by harming C). Nonetheless, we regard $\rho_{ij} = 1$ as a reasonable definition of what it means for species i and j to have no niche differences in this model, and Chesson (2013) argued that κ_j/κ_i is a valid general measure of pairwise fitness differences in multispecies communities.

Details of the calculations are in SI.9. The results (Table 6) are qualitatively congruent with the conclusions in Chu & Adler (2015): niche differences affecting recruitment (Δ^R) contribute most to persistence of the grasses, while those affecting survival (Δ^S) are most important for the shrub *Artemisia*. However, one notable difference is that effects of niche differences on growth are detrimental to *Hesperostipa* and *Pseudoroegneria*, and effects on survival are detrimental to *Pseudoroegneria*, whereas Chu & Adler (2015, fig. 5C) concluded that all effects of niche differences are at least mildly helpful. The detrimental effects arise in our analysis because benefits to an invading grass species are outweighed by larger benefits to their competitor *Artemisia*. When our analysis is redone with *Artemisia* removed from the community (results not shown), all effects of niche differences are helpful to all three grasses. This difference between our analysis and Chu & Adler (2015) illustrates the importance of invader–resident comparison: a seemingly beneficial feature may hurt you by helping your competitor more than it helps you.

DISCUSSION

We believe that the mathematical complexity and subtleties of standard MCT, and the approximate analytic formulas it requires for empirical applications, have obscured the simplicity, generality, and power of the underlying ideas, and obstructed the path to empirical applications. We have tried to break this logjam by separating the underlying concepts from the mathematical implementation, and providing a more general computational implementation. Our approach is not a replacement for analytic MCT – it is an additional tool, for detailed analysis of specific systems.

We abandon analytic MCT's requirement to use scaling factors q_{ir} in invader–resident growth rate comparisons. The scaling factors remove a term from the analytic growth rate decompositions for fluctuation-dependent mechanisms, allowing some necessary calculations even when the dynamics of the limiting factors for which species compete are unknown (Barabás *et al.* 2018). That term is not problematic in our approach, because all terms are evaluated using experimental or simulated data. Similarly, we abandon growth rate scalings in stabilising components, which have a firm theoretical foundation only for the case of a single limiting factor (Barabás *et al.* 2018). We also do not use the small-variance approximations that eliminate many interaction terms.

The main benefit of our approach is that analytic formulas are replaced by simulations or data, so new case studies or coexistence mechanisms do not require new math. This benefit is illustrated by our case studies, which include features (higher-order interactions between temperature and resource variability, interacting fluctuation-independent spatial

mechanisms, facilitation, and stage-specific niche differences in structured populations) that have not been analyzed within the standard framework. Many of these probably could be analyzed in the standard framework, but our approach gives more accurate and more complete answers, with much less effort.

The main weakness of our approach is that analytic formulas are replaced by simulations or data, so we do not get the generality and qualitative insights that can come from analytic formulas for coexistence mechanisms. Levins (1966) contrasted three incompatible goals of population modelling: generality, realism, and precision. Our approach provides realism and precision, giving a complete decomposition when the processes in a community have been modelled quantitatively, but the results are specific to the modelled community. Generality, a strength of analytic MCT, can only emerge with our approach by identifying general patterns in numerical results.

Like analytic MCT, our approach requires a way to calculate population growth rates as a function of the factors that determine their values at any time, and data (empirical, or simulation-derived) on the patterns of variation in those factors. There is no methodological cure for this requirement. To fully understand how species coexist we need to know what factors limit population growth rates, and quantify the impacts of each factor. Measuring and modelling population responses to multiple limiting factors in the field is especially challenging for long-lived species with complex life cycles, but even for short-lived species such as annual plants, describing population responses to variation in community composition and spatial and temporal environmental variation can be difficult. However, like analytic MCT our approach can also be applied to a single component of population growth rate, in the manner of Sears & Chesson (2007) who compared the contributions of spatial storage effect and local resource competition to seed yield in two desert annuals. This is accomplished by using the focal component (e.g. per capita fecundity) as the response variable, rather than population growth rate, and empirical data on the driving variables.

Discussion about how to quantify coexistence mechanisms will benefit from applying our approach to more case studies. Several times a new case study revealed to us that a seemingly wonderful idea was fatally flawed. Abandoning the scaling factors q_{ir} may be contentious, and therefore is likely to evolve. We have noted conceptual and practical difficulties with the q_{ir} , and suggested equal weighting of residents as a simple alternative, but equal weighting may not be appropriate when some resident species are far more common than others. Our approach also needs to be expanded further to encompass many important situations that we cannot yet address, including spatiotemporal variation, many individual- or grid-based spatial models (e.g. Adler *et al.* 2006), and processes occurring continuously in time as organisms move through space.

ACKNOWLEDGEMENTS

We thank our pre-submission and anonymous reviewers for comments and insightful questions that led to new analyses, substantial restructuring and a more user-friendly paper. We

thank the authors of Barabás *et al.* (2018) for sharing their paper prior to publication and for discussions that deepened our understanding of MCT. This research was supported by US NSF grants DEB 1353039 (SPE and GJH), DEB 1354041 (RES), and DEB 1353078 (PBA).

AUTHORSHIP

SPE, RES and GJH led the theory development; SPE and PBA led the data analysis and modelling; SPE, PBA and RES wrote scripts to simulate and analyze models; all authors discussed all aspects of the research and contributed to writing and revising the paper.

REFERENCES

- Abu Baker, M.A. & Brown, J.S. (2014). Foraging in space and time structure an African small mammal community. *Oecologia*, 175, 521–535.
- Adler, P.B., HilleRisLambers, J., Kyriakidis, P., Guan, Q. & Levine, J.M. (2006). Climate variability has a stabilizing effect on coexistence of prairie grasses. *Proc. Natl Acad. Sci. USA*, 103, 12793–12798.
- Angert, A.L., Huxman, T.E., Chesson, P. & Venable, D.L. (2009). Functional tradeoffs determine species coexistence via the storage effect. *Proc. Natl Acad. Sci. USA*, 106, 11641–11645.
- Bagchi, R., Gallery, R.E., Gripenberg, S., Gurr, S.J., Narayan, L., Addis, C.E. *et al.* (2014). Pathogens and insect herbivores drive rainforest plant diversity and composition. *Nature*, 506, 85–88.
- Barabás, G., D'Andrea, R. & Stump, S.M. (2018). Chesson's coexistence theory. *Ecol. Monographs*, 88, 277–303.
- Brown, J. (1989). Desert rodent community structure - a test of 4 mechanisms of coexistence. *Ecol. Monographs*, 59, 1–20.
- Bruno, J.F., Stachowicz, J.J. & Bertness, M.D. (2003). Inclusion of facilitation into ecological theory. *Trends Ecol. Evol.*, 18, 119–125.
- Bulleri, F., Bruno, J.F., Silliman, B.R., Stachowicz, J.J. & Michalet, R. (2015). Facilitation and the niche: implications for coexistence, range shifts and ecosystem functioning. *Funct. Ecol.*, 30, 70–78.
- Caley, M. (1995). Reef-fish community structure and dynamics: an interaction between local and larger-scale processes? *Mar. Ecol. Prog. Ser.*, 129, 19–29.
- Chesson, P. (1994). Multispecies competition in variable environments. *Theor. Popul. Biol.*, 45, 227–276.
- Chesson, P. (2000a). General theory of competitive coexistence in spatially-varying environments. *Theor. Popul. Biol.*, 58, 211–237.
- Chesson, P. (2000b). Mechanisms of maintenance of species diversity. *Am. Rev. Ecol. Syst.*, 31, 343–366.
- Chesson, P. (2003). Quantifying and testing coexistence mechanisms arising from recruitment fluctuations. *Theor. Popul. Biol.*, 64, 345–357.
- Chesson, P. (2013). Species competition and predation. In: *Ecological Systems: Selected Entries from the Encyclopedia of Sustainability Science and Technology* (ed. Leemans, R.). Springer, New York, chap. 13, pp. 223–256.
- Chu, C. & Adler, P.B. (2015). Large niche differences emerge at the recruitment stage to stabilize grassland coexistence. *Ecol. Monographs*, 85, 373–392.
- Comita, L.S., Queenborough, S.A., Murphy, S.J., Eck, J.L., Xu, K., Krishnadas, M. *et al.* (2014). Testing predictions of the JanzenConnell hypothesis: a meta-analysis of experimental evidence for distance- and density-dependent seed and seedling survival. *J. Ecol.*, 102, 845–856.
- Descamps-Julien, B. & Gonzalez, A. (2005). Stable coexistence in a fluctuating environment: an experimental demonstration. *Ecology*, 86, 2815–2824.
- Dybzinski, R. & Tilman, D. (2007). Resource use patterns predict long-term outcomes of plant competition for nutrients and light. *Am. Nat.*, 170, 305–318.
- Ellner, S.P., Snyder, R.E. & Adler, P.B. (2016). How to quantify the temporal storage effect using simulations instead of math. *Ecol. Lett.*, 19, 1333–1342.
- Gause, G.F. & Witt, A.A. (1935). Behavior of mixed populations and the problem of natural selection. *Am. Nat.*, 69, 596–609.
- Godoy, O. & Levine, J.M. (2014). Phenology effects on invasion success: insights from coupling field experiments to coexistence theory. *Ecology*, 95, 726–736.
- Gross, K. (2008). Positive interactions among competitors can produce species-rich communities. *Ecol. Lett.*, 11, 929–936.
- Grubb, P.J. (1977). The maintenance of species-richness in plant communities: the importance of the regeneration niche. *Biol. Rev.*, 52, 107–145.
- Holt, R.D. & Bonsall, M.B. (2017). Apparent competition. *Ann. Rev. Ecol. Evol. Syst.*, 48, 447–471.
- Hubbell, S.P. (2001). *The Unified Neutral Theory of Biodiversity and Biogeography*. Princeton University Press, Princeton, NJ.
- Hutchinson, G.E. (1959). Homage to santa rosalia or why are there so many kinds of animals? *Am. Nat.*, 93, 145–159.
- Hutchinson, G. (1961). The paradox of the plankton. *Am. Nat.*, 95, 137–145.
- Kotler, B. & Brown, J. (1988). Environmental heterogeneity and the coexistence of desert rodents. *Ann. Rev. Ecol. Syst.*, 19, 281–307.
- Letten, A.D., Dhimi, M.K., Ke, P.J. & Fukami, T. (2018). Species coexistence through simultaneous fluctuation-dependent mechanisms. *Proc. Natl Acad. Sci.*, 201801846.
- Levins, R. (1966). The strategy of model building in population ecology. *Am. Sci.*, 54, 421–431.
- Lönnerberg, K. & Eriksson, O. (2013). Rules of the seed size game: contests between large-seeded and small-seeded species. *Oikos*, 122, 1080–1084.
- McIntire, E.J.B. & Fajardo, A. (2013). Facilitation as a ubiquitous driver of biodiversity. *New Phytol.*, 201, 403–416.
- Metcalfe, C.J.E., Ellner, S.P., Childs, D.Z., Salguero-Gómez, R., Merow, C., McMahon, S.M. *et al.* (2015). Statistical modelling of annual variation for inference on stochastic population dynamics using Integral Projection Models. *Methods Ecol. Evol.*, 6, 1007–1017.
- Munday, P. (2004). Competitive coexistence of coral-dwelling fishes: the lottery hypothesis revisited. *Ecology*, 85, 623–628.
- Petermann, J.S., Fergus, A.J.F., Turnbull, L.A. & Schmid, B. (2008). Janzen-Connell effects are widespread and strong enough to maintain diversity in grasslands. *Ecology*, 89, 2399–2406.
- Sale, P. (1979). Recruitment, loss and coexistence in a guild of territorial coral-reef fishes. *Oecologia*, 42, 159–177.
- Sears, A.L.W. & Chesson, P. (2007). New methods for quantifying the spatial storage effect: an illustration with desert annuals. *Ecology*, 88, 2240–2247.
- Shmida, A. & Ellner, S. (1984). Coexistence of plant species with similar niches. *Vegetatio*, 58, 29–55.
- Silvertown, J., Dodd, M., Gowing, D. & Mountford, J. (1999). Hydrologically defined niches reveal a basis for species richness in plant communities. *Nature*, 400, 61–63.
- Snyder, R.E., Borer, E.T. & Chesson, P. (2005). Examining the relative importance of spatial and nonspatial coexistence mechanisms. *Am. Nat.*, 166, E75–E94.
- Usinowicz, J., Chang-Yang, C.H., Chen, Y.Y., Clark, J.S., Fletcher, C., Garwood, N.C. *et al.* (2017). Temporal coexistence mechanisms contribute to the latitudinal gradient in forest diversity. *Nature*, 550, 105–108.
- Volkov, I., Banavar, J.R., Hubbell, S.P. & Maritan, A. (2007). Patterns of relative species abundance in rainforests and coral reefs. *Nature*, 450, 45–49.
- Zarnetske, P.L., Gouhier, T.C., Hacker, S.D., Seabloom, E.W. & Bokil, V.A. (2013). Indirect effects and facilitation among native and non-native species promote invasion success along an environmental stress gradient. *J. Ecol.*, 101, 905–915.

Ziv, Y., Abramsky, Z., Kotler, B. & Subach, A. (1993). Interference competition and temporal and habitat partitioning in 2 gerbil species. *Oikos*, 66, 237–246.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Editor, Jessica Metcalf

Manuscript received 4 May 2018

First decision made 11 June 2018

Second decision made 30 July 2018

Manuscript accepted 14 August 2018