ARTICLE

Methods, Tools, and Technologies

Measuring stability in ecological systems without static equilibria

Adam Thomas Clark[1] | Lina K. Mühlbauer[1] | Helmut Hillebrand[2,3,4] | Canan Karakoç[5]

[1]Institute of Biology, University of Graz, Graz, Austria

[2]Institute for Chemistry and Biology of Marine Environments, Carl-von-Ossietzky University Oldenburg, Wilhelmshaven, Germany

[3]Helmholtz-Institute for Functional Marine Biodiversity at the University of Oldenburg, Oldenburg, Germany

[4]Alfred Wegener Institute, Helmholtz-Centre for Polar and Marine Research, Bremerhaven, Germany

[5]Department of Biology, Indiana University, Bloomington, Indiana, USA

Correspondence
Adam Thomas Clark
Email: adam.clark@uni-graz.at

Abstract

Ecological stability refers to a range of concepts used to quantify how species and environments change over time and in response to disturbances. Most empirically tractable ecological stability metrics assume that systems have simple dynamics and static equilibria. However, ecological systems are typically complex and often lack static equilibria (e.g., predator–prey oscillations, transient dynamics, chaos). Failing to account for these factors can lead to biased estimates of stability, in particular, by conflating effects of observation error, process noise, and underlying deterministic dynamics. To distinguish among these processes, we combine three existing approaches: state space models; delay embedding methods; and particle filtering. Jointly, these provide something akin to a deterministically "detrended" version of the coefficient of variation, separately tracking variability due to deterministic dynamics versus stochastic perturbations. Moreover, these variability estimates can be used to forecast dynamics, classify underlying sources of stochastic dynamics, and estimate the "exit time" before a state change takes place (e.g., local extinction events). Importantly, the time-delay embedding methods that we employ make very few assumptions about the functions governing deterministic dynamics, which facilitates applications in systems with limited data and a priori biological knowledge. To demonstrate how complex dynamics without static equilibria can bias ecological stability estimates, we analyze simulated time series of abundance dynamics in a system with time-varying carrying capacity and empirically observed abundance dynamics of the green algae *Chlamydomonas terricola* grown in a diverse microcosm mixture under variable temperature conditions. We show that stability estimates based on raw observations greatly overestimate temporal variability and fail to accurately forecast time to extinction. In contrast, joint application of state space modeling, delay embedding, and particle filters were able to: (1) correctly quantify the contributions of deterministic versus stochastic variability; (2) successfully estimate "true" abundance dynamics; and

(3) correctly forecast time to extinction. Our results therefore demonstrate the importance of accounting for effects of complex, nonstatic dynamics in studies of ecological stability and provide an empirically tractable and flexible toolkit for conducting these measurements.

INTRODUCTION

Stability has long been a central focus in ecology (Holling, 1973; Ives et al., 1999; Lewontin, 1969; May, 1973; Pimm, 1984). Although many different definitions and aspects of stability exist, a unifying theme is that they aim to quantify variability and sensitivity to disturbance in dynamical systems (Donohue et al., 2016; Grimm & Wissel, 1997). These metrics are therefore vital for, among other things, estimating the likelihood of invasions or extinctions, forecasting population dynamics, or predicting effects of anthropogenic disturbances.

Despite decades of work and major advances in understanding of stability in ecological systems (Donohue et al., 2013; Tilman, 1995), at least two related methodological hurdles currently hinder progress. First, although many flexible theoretical stability metrics have been developed, most of these tend to be difficult to apply in practice, leading theoretical versus empirical studies to apply disparate, and generally incompatible, toolsets (DeAngelis & Yurek, 2015; Donohue et al., 2016; Evans, 2012). Second, most existing empirically tractable metrics assume that systems have a single, static equilibrium (e.g., a carrying capacity), which is used as a "baseline" against which effects of disturbances and recovery rates are gauged (Donohue et al., 2016; Grimm & Wissel, 1997; Ives et al., 1999; but see Chesson, 2017). However, in real-world ecological systems, baseline conditions are not always clear, and dynamics often vary greatly across space and time (Clark et al., 2021; Coulson, 2021; DeAngelis & Waterhouse, 1987; Hastings et al., 1993; Pimm et al., 2019; Shoemaker et al., 2022).

A major challenge is that variability in complex systems can be attributed to at least three different sources: observation error, process noise, and deterministic variation (de Valpine & Hastings, 2002). These describe, respectively: deviations between imperfect measurements and the true state of a system; stochastic fluctuations that actually influence dynamics; and variability caused by repeatable system processes (i.e., such that two identical systems starting in the same state will follow exactly the same dynamics). Because observation error results only

in changes in our perception of a system but not its actual state, it can lead to overestimates of variability, and thus underestimates of stability (de Mazancourt et al., 2013). In contrast, both process noise and deterministic variation can influence stability, although their relative contributions can be difficult to disentangle. For example, fluctuations in species abundances might be driven by repeated environmental perturbations, which could indicate instability and that local extinction is imminent (Scheffer et al., 2009), or they might result from stable oscillatory cycles (e.g., predator–prey dynamics), which can be consistent with long-term persistence and coexistence (Blasius et al., 2020).

In this study, we seek to overcome these methodological hurdles by combining a suite of established methods that are both empirically tractable and suitable for analyzing dynamic systems even in the absence of static equilibria. By "systems in the absence of static equilibria," we mean those that do not fluctuate around a single fixed plot, including, for example, systems undergoing limit cycles, oscillations, or chaotic dynamics, those with time-varying equilibria, and even those with no equilibrium state at all. Jointly, the methods presented here return estimates of: (1) the true dynamic trajectory of the system after accounting for effects of observation error; (2) the relative contributions of observation error, process noise, and deterministic dynamics to overall temporal variability; and (3) average "exit time" until a state change occurs (e.g., time to extinction). To test this approach, we first apply it to simulated abundance dynamics for a system with time-varying carrying capacity, for which these three attributes are already known a priori. We then apply the methods to analyze empirically observed abundance dynamics of the green algae *Chlamydomonas terricola* grown under temporally oscillating temperature conditions. Taken together, our results show how classic stability metrics based on raw observations can substantially overestimate temporal variability in complex systems without static equilibria. However, we find that in many cases these biases can be mitigated by properly separating the influences of different sources of variability.

METHODS

The methods we present here build on three well-established toolsets. First, observations of a single long time series, or of multiple spatially replicated time series, are collected (Figure 1a). Then, nonparametric delay embedding methods are employed to reconstruct the underlying deterministic dynamics of the system (Figure 1b,c). Next, estimates from the delay embedding methods are combined with hypothesized parametric distributions describing the effects of process noise and observation error as part of a state space model (Figure 1d). Then, the parameters for the stochastic distributions are fitted using particle filtering and Markov chain Monte Carlo (MCMC) optimization (Figure 1e). Finally, the outputs of this optimizer can be used to estimate, for example, the relative contributions of deterministic variation, observation error, and process noise, or to forecast system dynamics and expected time to local extinction (Figure 1f,g).

We call this joint approach "Particle-Takens filtering," following from the related, existing method of "Kalman-Takens" filtering (Hamilton, Berry, & Sauer, 2017; Hamilton, Lloyd, & Flores, 2017). Functions for implementing these methods and detailed step-by-step guides of the underlying functions are available in the accompanying pttstability (Particle-Takens stability) package for the R programming language (R Development Core Team, 2019). More information on these methods, and a detailed workflow for the package, are described in the *Method definitions* and *Detailed workflow*.

Method definitions

State space models

State space models are commonly applied to account for stochasticity in time series data and are the foundation of many methods in population ecology (de Valpine & Hastings, 2002; Knape & de Valpine, 2012; Plard et al., 2019). The approach decomposes dynamics into two "states": one representing imperfect observations of the system at time t, $N_{obs}(t)$, and a second, unobservable state, representing the "true" value $N_{true}(t)$. These states can represent any dynamical variable, though in population ecology, $N_{obs}(t)$ typically represents observed abundance dynamics, and $N_{true}(t)$ represents hypothetical "true" abundances in the absence of observation error. Jointly, these yield a model of the form:

$$N_{obs}(t) = N_{true}(t) + O, \tag{1a}$$

$$N_{true}(t+1) = g(N_{true}(t)) + P. \tag{1b}$$

Here, g is a deterministic function that governs changes in N_{true} over time, and O and P are functions representing effects of observation error and process noise, respectively. Because Equation (1b) separates the effects of variability driven by deterministic processes in function g from stochastic processes driven by P, the process noise function effectively represents a "detrended" analogue of the classic coefficient of variation stability metric (CV, i.e., $\sqrt{var(N_{obs})}/<N_{obs}>$, where $<x>$ is the arithmetic mean of variable x) after separating the effects of deterministic variability and observation error. Although all three sources of variability are important for understanding system dynamics, as we will discuss below, separating the contributions of these underlying processes can yield substantially more useful and tractable forecasts of system dynamics.

In theory, the model in Equations (1a) and (1b) is both simple and general: unbiased estimates of $N_{true}(t)$ can be generated by maximizing the joint likelihood of observations given in Equation (1a), versus the likelihood of subsequent estimates given in Equation (1b). In practice, however, fitting these equations to observations is challenging and usually requires long time series (Massoud et al., 2018). Additionally, model predictions are strongly influenced by the functional form and underlying processes chosen for g, O, and P (Ovaskainen & Meerson, 2010). In particular, if deterministic dynamics modeled by g do not accurately describe those of the real system, then this modeling error will be incorporated into P, leading to an overestimate of process noise. Thus, especially with limited data and little a priori knowledge of underlying biology, it can be difficult to separate robust model predictions from spurious effects of model choice (Auger-Méthé et al., 2016).

Delay embedding methods

A partial solution to this problem is to integrate delay embedding methods into state space models, as demonstrated by the recently developed Kalman-Takens filter (Hamilton, Berry, & Sauer, 2017; Hamilton, Lloyd, & Flores, 2017). This method builds on Kalman filtering, a classic approach for parameterizing state space models, by replacing the deterministic function g with a nonparametric estimate rooted in Takens' delay embedding theorem (Sugihara et al., 1990; Takens, 1981). Takens' theorem effectively shows that future states of dynamical systems can be forecast by averaging across historically observed dynamics, weighted by their similarity to the current state. These historical dynamics can represent repeated measurements of a single variable or of multiple variables and covariates.
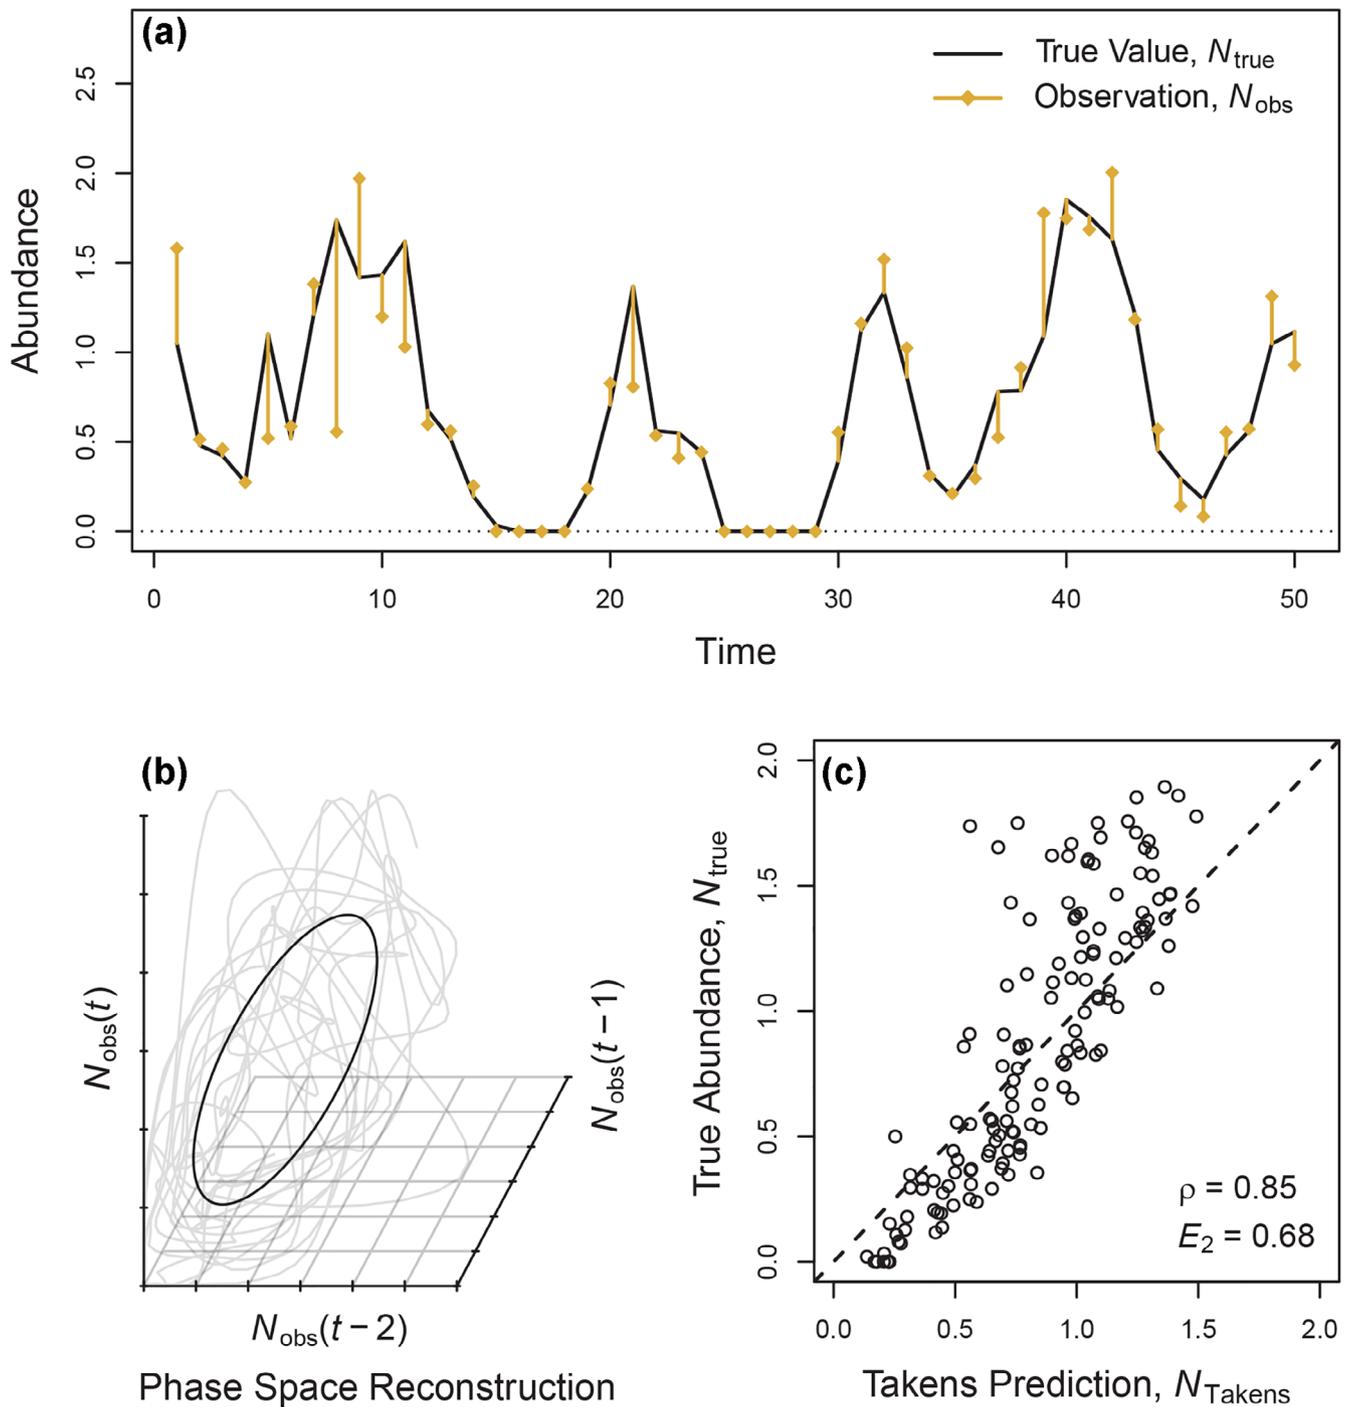
FIGURE 1 Workflow for the pttstability package. (a) A noisy time series of abundance dynamics simulated from Equations (6a) and (6c) with $r = 1.2$, $K = 1$, $\sigma_O = 0.3$, and $\sigma_P = 0.15$. (b) Delay embedding of abundance dynamics. $N_{obs}(t)$ shows observed abundance at time t. Gray lines show the trajectory of observations in phase space, and black circle shows the deterministic trajectory. (c) True abundances versus predictions from empirical dynamic model trained on observations. Pearson correlation coefficient (ρ) and coefficient of efficiency (E_2) show goodness of fit. (d, e) Particle filter results for a subset of dynamics. Thick blue line in (d) shows average trajectory predicted by the filter, whereas thin blue lines and points show individual particles representing prediction uncertainty. (e) True abundance at $t = 30$ (black dashed line), and distributions showing effects of process noise and observation error (blue and yellow, respectively), and posterior estimate from the filter (gray). (f) Effects of process noise. Green points show what abundances would have been in the absence of process nose (N_{true}'), while blue-shaded region shows estimates of these states from the filter. Note stochastically driven extinction event at $t = 34$. (g) Relationship between abundance, process noise, and probability of mortality.
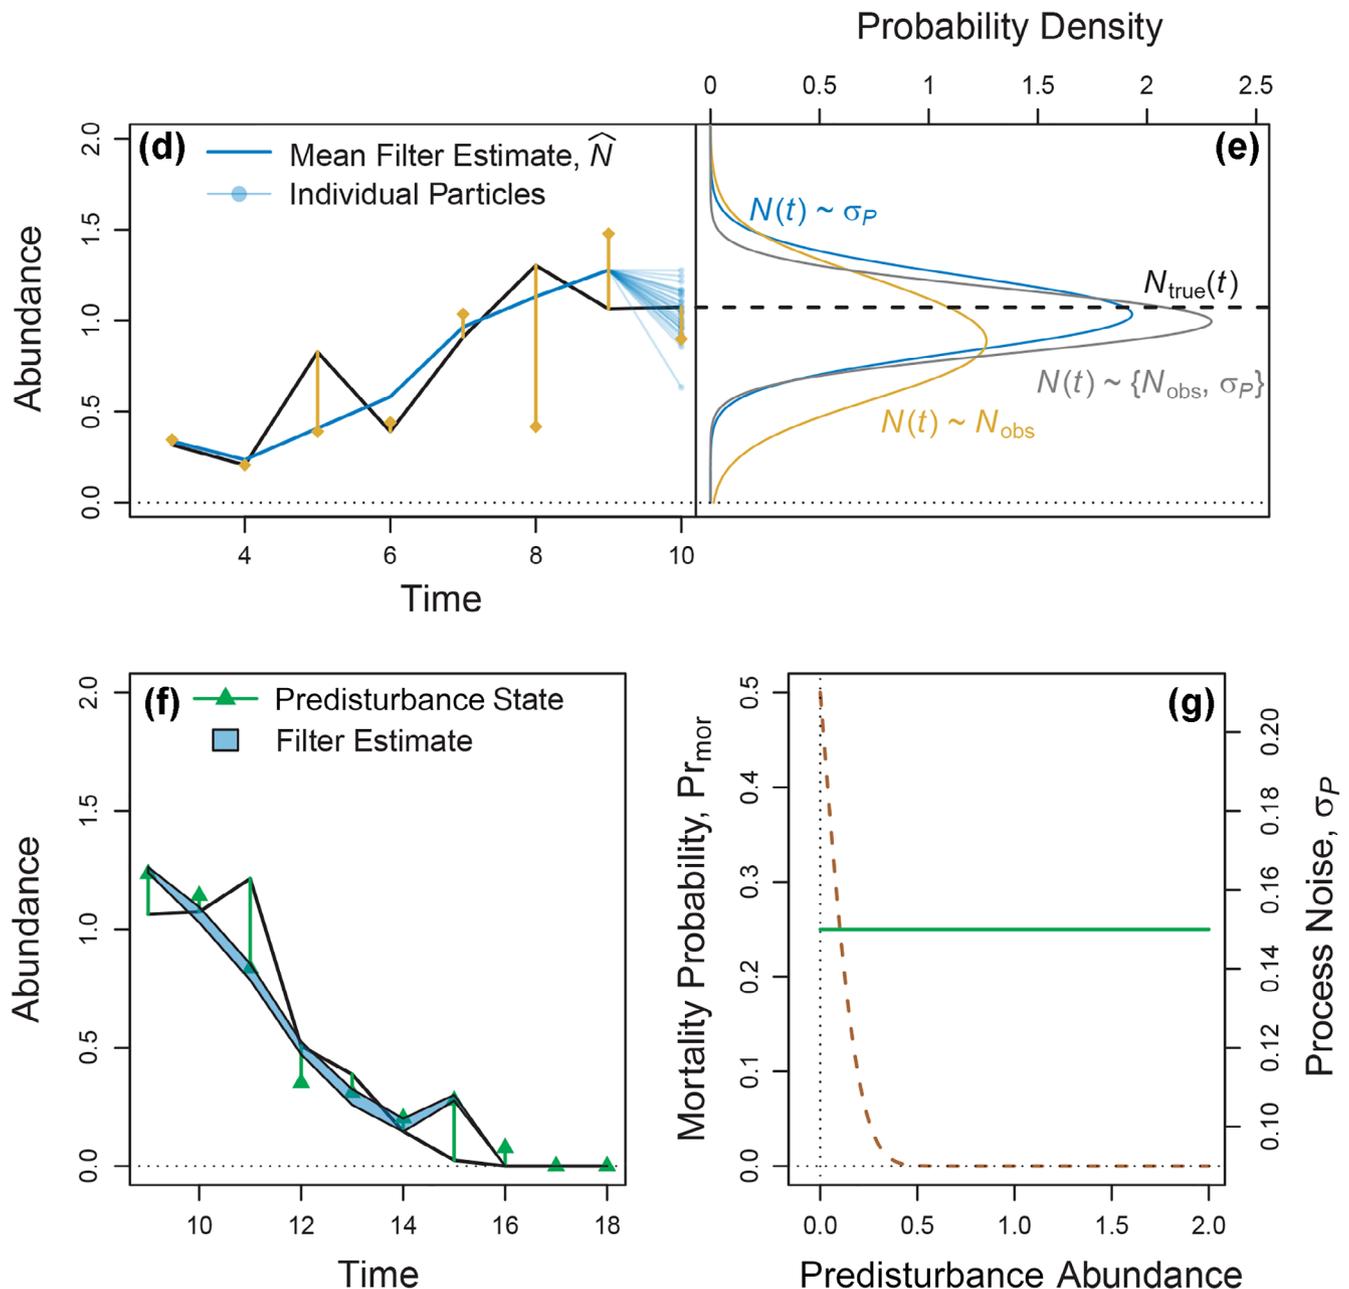
FIGURE 1 (Continued)

In practice, delay embedding methods have yielded spectacularly good predictive performance across a wide range of systems (Clark et al., 2015; Hsieh et al., 2008; Perretti et al., 2012; Sugihara et al., 1990). Additionally, this approach has two major advantages over other nonparametric methods. First, given sufficiently long time series, delay embedding methods can, in theory, accurately predict the deterministic component of almost any system, including those with nonstatic, nonlinear, or chaotic dynamics (Deyle & Sugihara, 2011; Sugihara et al., 1990). Second, even in complex, multivariate systems, forecasts can be made based on limited data,

because time-lagged observations can be substituted for missing variables (Schaffer, 1984; Sugihara et al., 2012).

Particle filtering

Although delay embedding methods replace the deterministic function g in Equation (1b), stochastic effects of observation error and process noise must still be represented parametrically. A limitation of Kalman filters is that they apply only when O and P follow relatively simple functional forms, for example, normally distributed

fluctuations with a fixed mean and standard deviation (Hamilton, Berry, & Sauer, 2017). This presents a problem, as realistic observation error and process noise functions in ecological systems are often complex, for example, for abundance dynamics, distributions must be censored at zero, since abundances cannot become negative. Moreover, although Kalman filters excel at generating unbiased estimates of true dynamics, they are typically not designed to accurately estimate the parameters of O and P.

To overcome these problems, we instead apply particle filtering to fit the state space models (Knape & de Valpine, 2012). Particle filtering simulates a swarm of synthetic time series based on Equations (1a) and (1b), each called a "particle," where O and P can be chosen to represent any hypothesized functions regardless of functional form. The likelihood of observations given the hypothesized model can then be calculated directly from the resulting estimates of the true system state. The state space model can then be fit to data using established methods such as MCMC (Hartig et al., 2019; Wilkinson, 2018). The advantage of particle filtering over other methods, such as Kalman filtering, is that it can be used to fit much more complex statistical distributions, providing more flexibility for users in defining stochastic functions; however, this flexibility does come at the cost of much longer computational times, and in some cases lower statistical power when applied to simple systems.

Detailed workflow

Our workflow proceeds in five steps, discussed below. A step-by-step walk-through of a worked example is available in help documentation for the pttstability package.

Specifying functions

First, we specify parametric functions for O and P representing the effects of observation error and process noise as follows:

$$
O(\sigma_O, <N_{obs}>, N_{est}(t))
$$
$$
= \mathrm{pmax}(-N_{est}(t), \mathrm{rnorm}(\mu = 0,
$$
$$
\sigma = \mathrm{pmax}(0.01 \times <N_{obs}>, \ \sigma_O \times N_{est}(t)))), \quad (2a)
$$

$$
P(\sigma_P, N_{est}{}'(t)) = \mathrm{pmax}(-N_{est}{}'(t), \mathrm{rnorm}(\mu = 0, \ \sigma = \sigma_P)), \quad (2b)
$$

where rnorm and mean are functions in the R stats and base packages, $<N_{obs}>$ is the average state observed across

all timesteps, $N_{est}(t)$ is estimated state at time t, and $N_{est}{}'(t)$ is an estimate of what the state would have been in the absence of process noise, that is, $g(N_{est}(t-1))$. To make the functions applicable for abundance dynamics, we censor the distributions using the pmax command from the base R package, which ensures that Equations (1a) and (1b) return only nonnegative results. For Equation (2a), we assume that the standard deviation of the observation error function is a fixed fraction of the current abundance, σ_O, with a minimum value of 1% of mean observed abundance, which is necessary to prevent infinitely small likelihoods when abundance falls to zero (see Equations 6a and 6b). For Equation (2a), we assume that process noise is drawn from a normal distribution with fixed standard deviation σ_P and minimum resulting abundance 0.

Delay embedding

Next, a delay embedding model must be fitted to replace function g in Equation (1b). This model jointly incorporates the effects of deterministic dynamics (e.g., growth, predator–prey oscillations, etc.) and persistent effects of press perturbations. To fit these functions, we apply the S-mapping empirical dynamic modeling (EDM) routine from the rEDM package (Park et al., 2020; Sugihara et al., 1990). The algorithm requires user-defined parameters describing optimal embedding dimension (E) and nonlinearity (θ) as inputs, which are selected via leave-one-out cross validation using the s_map function (Ye et al., 2018).

S-mapping has two advantages over classic delay embedding methods. First, it often results in better predictions and extrapolations, because it fits a series of local regressions that allow the relative weighting of embedding dimensions to vary across state space (Sugihara et al., 1990). Second, it produces a static prediction matrix rather than recalculating neighbor distances for every particle, which substantially increases computational efficiency. Finally, it allows for inclusion of covariates (Ye & Sugihara, 2016). Following Hamilton, Lloyd and Flores (2017), we generate delay embedding predictions directly from noisy observations (i.e., N_{obs}). This approach also increases computational efficiency and is theoretically justified for most observation error functions (Deyle & Sugihara, 2011). Recall that the resulting predictions are solely used to replace the deterministic function g and are not direct estimates of N_{true}.

Particle filtering

The observation error, process noise, and EDM functions are then passed to the particle filter, in order to estimate the

likelihood of observations given in the model. We apply a modified version of the filters proposed by Knape and de Valpine (2012)—see the function particleFilterLL in the pttstability package for details. As is standard in state space models, this likelihood is calculated using the inverse of the observation error function. To account for the censored distributions in Equation (2a), we apply a Tobit model centered on N_{est} with lower bound zero. The Tobit model has the advantage of jointly accounting for zero and nonzero abundances as part of a single distribution with only three parameters (i.e., mean, standard deviation, and lower limit), which simplifies the optimization process. For any given parameter set, likelihood for each particle i is calculated as follows:

$$u_i = \exp(\text{LL}(N_{obs}|N_{est,i}) - \max[\text{LL}(N_{obs}|N_{est,i})]), \quad (3)$$

where LL is the log-likelihood from the inverse observation error function and the max function prevents rounding errors for small likelihoods (Wilkinson, 2011). These likelihood estimates are then used to generate weights for each particle as follows:

$$w_i = \exp(u_i)/\Sigma_j \exp(u_j), \quad (4)$$

which also cancels out the maximum function in Equation (3). Particles are then resampled with replacement over every timestep based on these weights, generating a joint estimate of N_{true} and corresponding likelihoods resulting from the combined effects of the deterministic EDM function, the stochastic process noise and observation error functions, and the observations. Resampling is conducted in order to prevent "particle degeneracy," that is, cases were all particles in the set being tracked have such low likelihood that they can no longer be used to model system dynamics (Knape & de Valpine, 2012; Wilkinson, 2011).

Optimization

To fit parameter estimates for the observation error and process noise functions in Equations (2a) and (2b), log-likelihoods from the filter must be passed to an optimizer. We used the differential-evolution MCMC zs (DEzs) sampler from the BayesianTools package (Hartig et al., 2019) with bounded flat priors (details below) for σ_O and σ_P. We chose this solver because particle filtering produces noisy likelihood estimates (i.e., estimates for a fixed parameter set fall within a distribution, rather than on a fixed point), which can cause issues for some other more commonly used MCMC algorithms that include a gradient assent step. We then substituted these estimates into the

particle filter to produce estimates of true states for each timestep N_{est} and of $N_{est}'(t)$, that is, what the system state would have been in the absence of process noise.

State changes

For models that undergo state changes (e.g., extinction events or "critical transitions"), results from the filter can also be used to estimate the average waiting time before the next state change takes place, even if none have been observed yet. This estimate is equivalent to "exit time" sensu Arani et al. (2021). However, while existing methods for calculating exit times are limited to long, low-dimensional time series (Arani et al., 2021), exit times can be computed for our methods directly from the fitted EDM and process noise functions. In particular, for abundance dynamics, average per-timestep mortality probability can be calculated as follows:

$$\text{Pr}_{mor} = \text{mean}(\text{pnorm}(0, \ \mu = N_{est,i}'(t), \ \sigma = \sigma_P))$$
$$\text{for all } N_{est,i}'(t) \neq 0, \quad (5)$$

with expected time to extinction $1/\text{Pr}_{mor}$. In other words, Equation (5) is the average probability per timestep of process noise driving abundance to zero, that is, a measure of persistence sensu Law and Morton (1996). This metric is of particular interest for stability analysis, as it summarizes species abilities to persist in communities over time, based on the joint effects of both deterministic variation (i.e., in N_{est}) and stochastic variation (i.e., σ_P) in the system. Thus, in Equation (5), high Pr_{mor} can arise from either high rates of process noise or from low average abundances (or both). Note that Equation (5) can be applied even for short time series where no extinctions have been observed by simulating the particle filter forward in time and analyzing the resulting dynamics.

Tests on simulated data

To test the ability of our methods to accurately estimate true dynamics, parameter values for σ_P and σ_O, and extinction probabilities, we applied it to simulated time series of species abundance dynamics from the model:

$$N_{true}(t+1) = N_{true}(t) \exp[r(1 - N_{true}(t)/K(t))] + P(\sigma_P, N_{true}'(t)), \quad (6a)$$

$$N_{obs}(t) = N_{true}(t) + O(\sigma_O, <N_{obs}>, N_{true}(t)), \quad (6b)$$

$$K(t) = (\sin(t/2) + 1.5)(2/3). \quad (6c)$$

The oscillatory K function in Equation (6c) ensures that the model equilibrium varies over time (Figure 2a) and roughly matches the period observed in the empirical example described below.

For each of 1255 simulations, we drew parameter values from random uniform distributions, ranging from 0.01 to 0.5 for σ_O and from 0.0065 to 0.65 for σ_P (ranges were chosen because these produced a wide range of dynamics; iteration count was determined by the number of simulations that completed within the time frame allocated to us on our high-performance computing cluster). We then used Equations (6a) and (6c) to generate dynamics, with growth rate $r = 1.2$ for all simulations (i.e., the nonchaotic regime of the Ricker model). If extinction events occurred, the simulation was "recolonized" with a probability of 20% per timestep, with initial abundance of

0.1. All simulations were run for 150 timesteps (roughly matching the total composite time series length per treatment in the empirical examples below). Note that because $<K(t)> = 1$ over long time spans, σ_O and σ_P roughly describe the average size of fluctuations due to observation error or process noise, relative to the average population size.

To test the performance of the delay embedding methods, we compared two different approaches for fitting state space models to N_{obs}. First, we applied the "correct" underlying deterministic function, that is, by substituting g in Equation (1b) with Equations (6a) and (6c) and the true r and colonization parameter values. This approach might represent, for example, a well-tuned integrated population model. Second, we replaced g with the nonparametric estimate derived from EDM, representing a
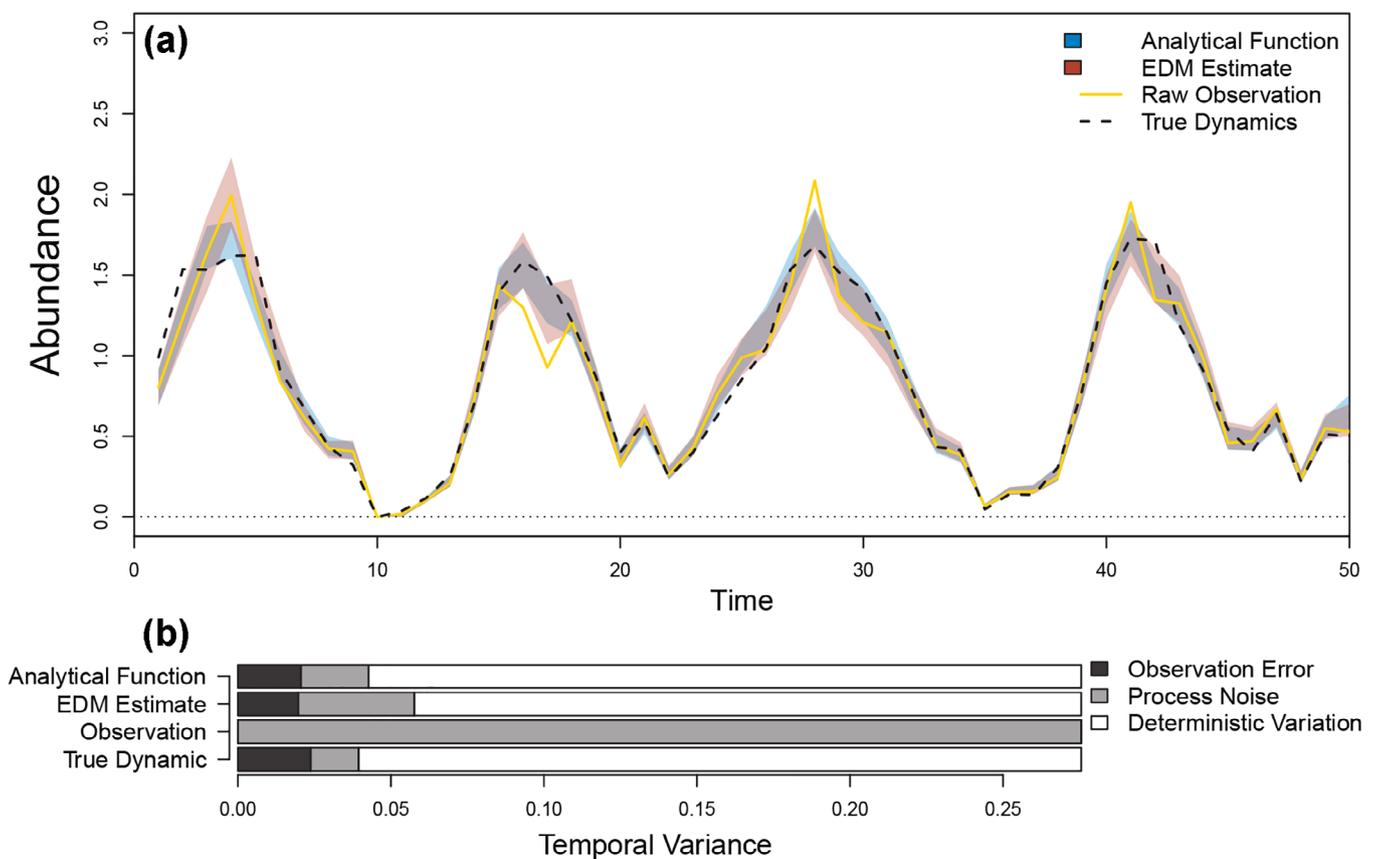

FIGURE 2 Example of different sources of temporal variation in a system without static equilibria. (a) A noisy time series representing hypothetical abundance dynamics with oscillating carrying capacity, simulated following Equations (6a) and (6c) in the main text ($r = 1.2$, $K_0 = 1$, $\sigma_O = 0.2$, and $\sigma_P = 0.1$). Black dashed line shows true dynamics, yellow line shows noisy observations, and shaded regions show mean estimate $\pm$ 1 standard deviation from the particle filters based on the correct analytical model (red), or empirical dynamic model (EDM) (blue), as described in the main text. (b) Partitioning of total temporal variance into observation error, process noise, and deterministic variation. "True dynamic" shows actual underlying contribution of each type of variation: Process noise magnitude is determined by function P in Equation (6a), observation error by function O in Equation (6b), and deterministic variation as the total remaining variation, driven primarily by the time-varying carrying capacity in Equation (6c). "Observation" assigns all observed variation to process noise, as is typically implied in standard CV calculations. "Analytical function" and "EDM estimate" show results from two particle filters.

case where the underlying deterministic function is not known a priori.

For both approaches, we specified uniform bounded priors over 0.01 to 0.5 for σ_O, and 0.0065 to 0.65 for σ_P, simulated trajectories for 1000 particles per time series, and 6000 MCMC steps over three chains including a 1000-step burn-in. We retained only estimates from models for which good MCMC convergence was indicated, that is, R-hat $\leq$ 1.1. This criterion was met for approximately 74% of simulations fitted using the true deterministic function and approximately 79% using EDM ($\sim$65% met the criteria for both approaches). Note that higher convergence rates could have been achieved by hand-tuning MCMC options for each simulation. To compare observed versus predicted values, we used two metrics: Pearson correlation coefficient (ρ) and the coefficient of efficiency (E_2), which is similar to R^2, but tracks correspondence along the 1–1 line rather than a fitted regression line, that is, as $1 - \Sigma(x_{\text{true}} - x_{\text{est}})^2/(x_{\text{true}} - <x_{\text{true}}>)^2$.

Empirical example

We applied our methods to analyze empirically observed abundance dynamics of the green algae *C. terricola* in a microcosm experiment (Burgmer & Hillebrand, 2011). In this experiment, 19 phytoplankton species were grown together in a semicontinuous culture over 70 weeks. Abundance for each species was measured as cell counts via microscopy 17 times over the experiment, at approximately four-week intervals. We focus on *C. terricola* because it was consistently one of the most abundant species in mixture and present results for two experimental treatments: cultures raised at high- versus low-oscillating temperature conditions, each replicated eight times. Both treatments were subject to slow, smooth, oscillatory changes in temperature ranging from 22 to 8°C, or from 20 to 4°C, with an oscillatory period of 52 weeks, meant to mimic seasonal effects, and with predators present ("HSP" and "LSP" in the original paper). We choose these treatments because they represent similar conditions, but display divergent dynamics (long-term oscillations vs. decline toward extinction).

For both treatments, we combined observed dynamics across replicates to generate long composite time series of $17 \times 8 = 136$ observations per treatment. Methods for analyzing composite time series are well developed for EDM and are already implemented in the rEDM package (Clark et al., 2015; Hsieh et al., 2008). For each subset, we simulated between 1000 and 512,000 particles, choosing the minimum number needed to avoid particle degeneracy (i.e., infinitely low likelihoods). See the help documentation in the pttstability package for the particleFilterLL_piecewise function for more details.

Prior to analysis, composite time series for each treatment were standardized by dividing abundances by their mean (i.e., such that mean(N_{obs}) = 1). We again used uniform priors, with lower limit 0.001 and upper limit 2 for both σ_O and σ_P, and ran 10,000 MCMC steps, including 2000 burn-in steps. Because systems were "closed," we set colonization probability arbitrarily near zero (1e–6 per timestep). We then applied the workflow described above to fit parameter estimates for σ_O and σ_P, estimate true abundance dynamics for each replicate, and forecast expected time to extinction.

RESULTS

Analysis of the simulated time series demonstrated that both when fitted using the "correct" underlying deterministic function and when fitted using nonparametric EDM methods, the state space models performed well across a wide range of process noise and observation error regimes. For particle filters fitted to simulated dynamics, predictions of abundance dynamics from both approaches significantly outperformed raw observations, although predictive ability declined with higher observation error, σ_O (Figures 2 and 3a,b). In general, predictions from the correct deterministic function were somewhat better than those from EDM, and differences among methods were more pronounced for E_2 (i.e., correspondence of predictions vs. observations along the 1–1 line) than for ρ (i.e., Pearson correlation).

Predictions for observation error strength, σ_O, closely matched true values for the simulated system, both for particle filters fitted using the correct deterministic function and fitted using EDM (Figure 4a). For filters fitted with the correct deterministic function, estimates of process noise strength, σ_P, closely matched true values across all simulations (Figure 4b,c). For those fitted with EDM, predictions of σ_P became worse under stronger observation error, although estimates remained relatively accurate for observation error $\sigma_O < 0.3$ (see Appendix S1: Figure S1). In contrast, regardless of observation error strength, total variability observed in the raw time series (e.g., as might be used to calculate CV) greatly overestimated σ_P (yellow points and intervals in Figure 4b,c).

For estimates of mean per-timestep mortality probability Pr_{mor} as derived from Equation (5), patterns in prediction accuracy were similar to those for σ_P. Both particle filters yielded good predictions when observation error was low (Figure 5a,b), but when observation error was high, filters fitted with EDM tended to overpredict mortality probabilities, especially when the true mortality probability was low (Figure 5d,e). This bias likely resulted from errors in the EDM predictions of
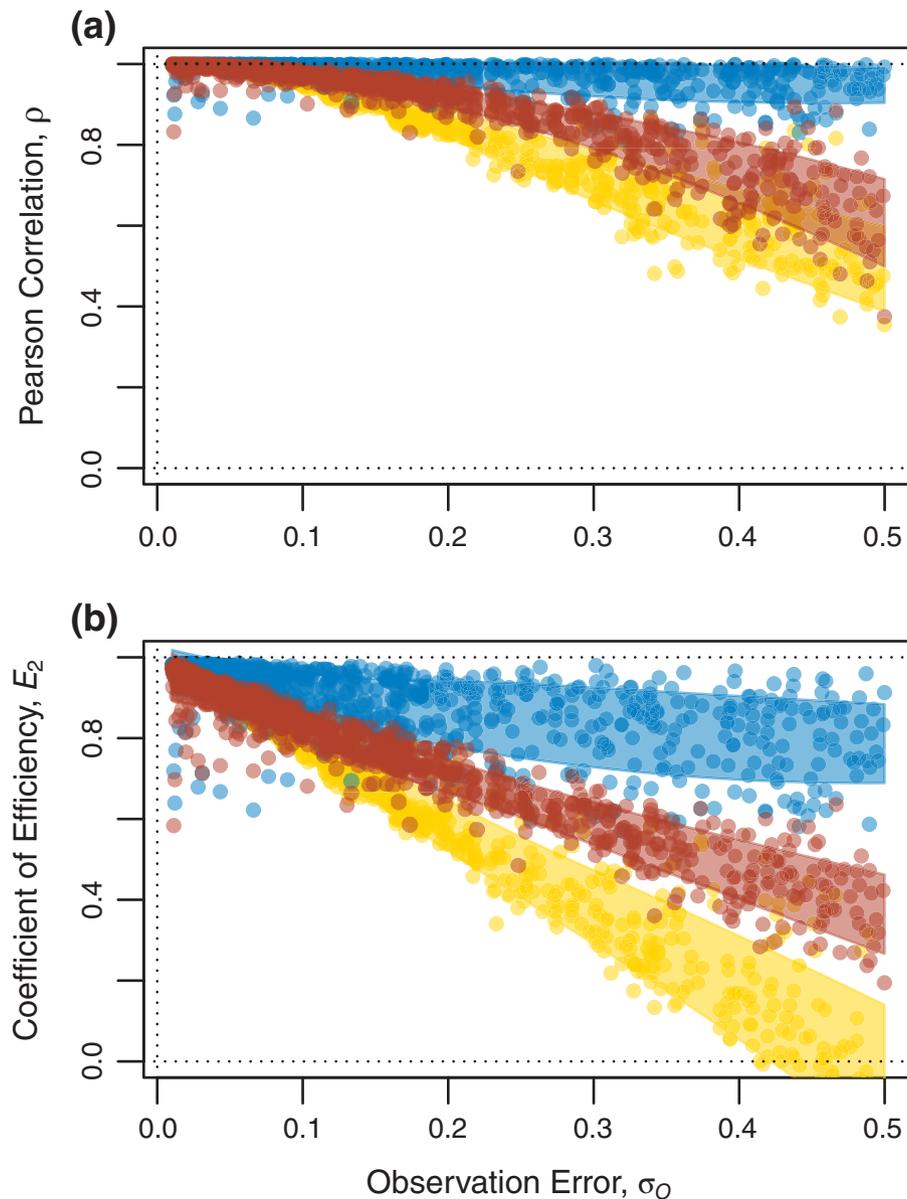
(a)

(b)

FIGURE 3   Goodness of fit between true abundances for the simulated dynamics described in the main text versus observations (yellow), and predictions from the particle filters based on the correct analytical model (blue), or on empirical dynamic model (red). Points show results from individual simulations, and shaded regions show mean trend ± 1 standard deviation. Variable σ_O describes the magnitude of observation error following Equation (6b) and roughly relates to the fraction of overall observed variability that is attributable to observation error. Coefficient of efficiency (E_2) is conceptually similar to R^2, but measures scatter around the 1–1 line rather than around a fitted regression, as $1 - \Sigma(x_{\text{true}} - x_{\text{est}})^2/(x_{\text{true}} - <x_{\text{true}}>)^2$.

deterministic function g, which began to overestimate the frequency of "true" zero abundances in very noisy time series. In contrast, estimates of mortality probability derived by counting the number of times that species fell to zero abundance in raw observations only remained accurate so long as the average expected time to extinction was shorter than the observed time series length (i.e., such that on average at least one mortality event occurred over the observed time span). For $1/\text{Pr}_{\text{mor}} <$ time series length, no estimates were possible

(Figure 5c,f). Note that we do not account for recolonization in these raw estimates, as colonization rates cannot be calculated from raw data where no extinctions have been observed.

Analysis of the empirically observed *C. terricola* abundance dynamics revealed relatively similar effects of observation error σ_O and process noise σ_P on dynamics across both the high- and low-temperature treatments (Figure 6). However, compared with the high-temperature treatment, in the lower temperature treatment, abundance was
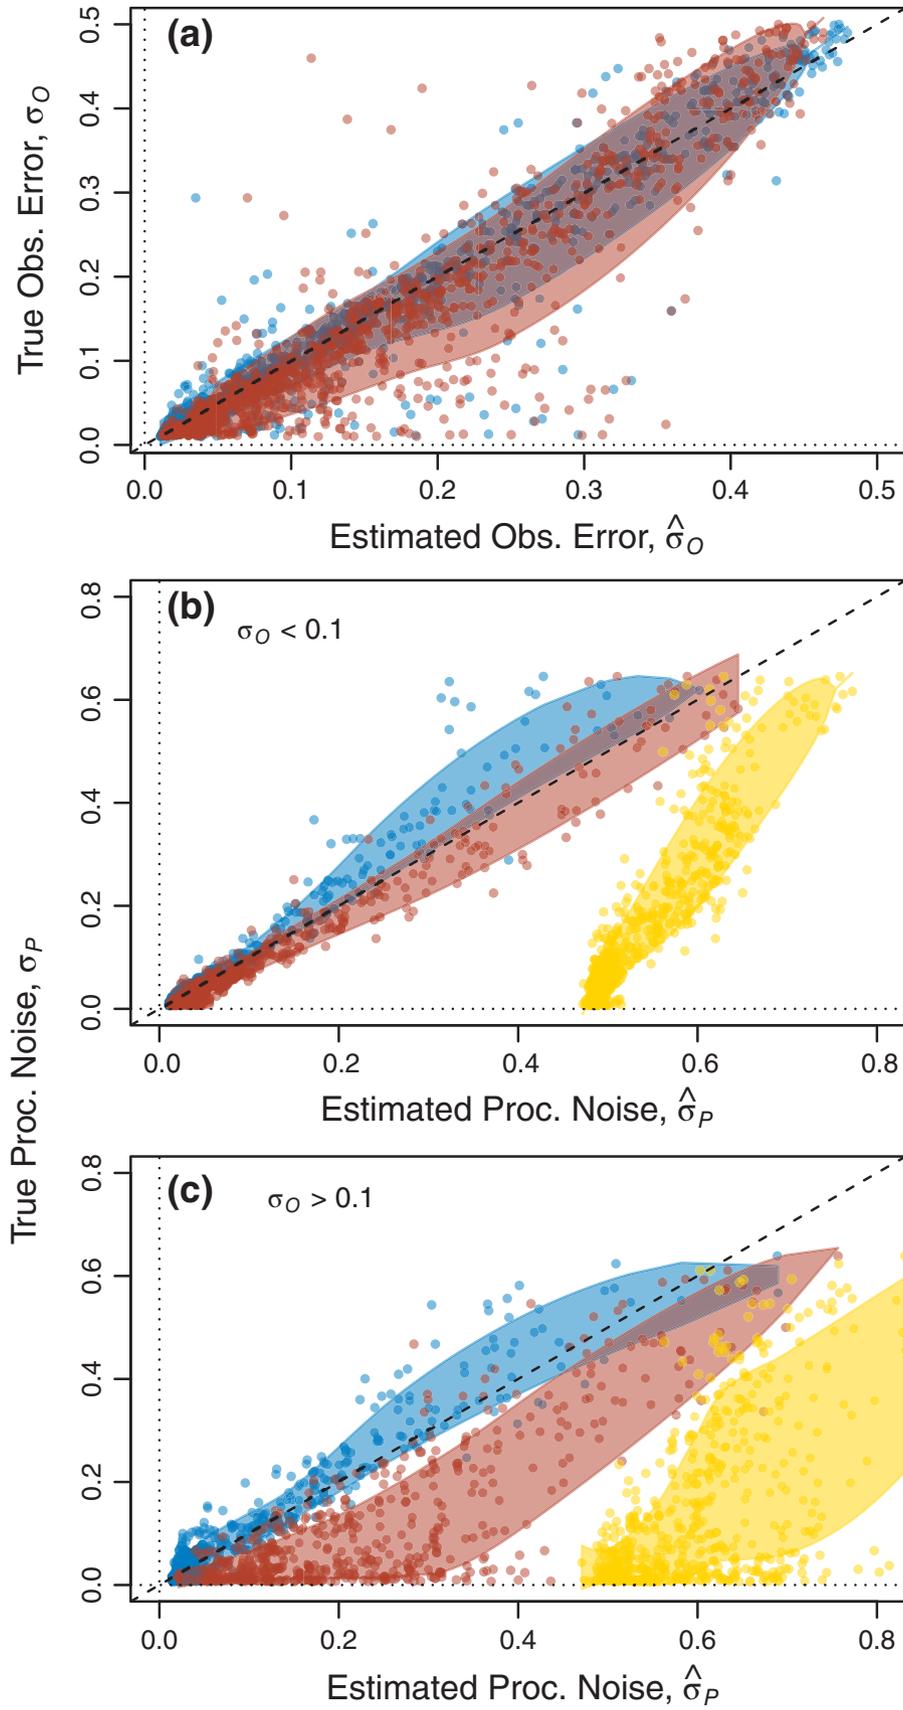
FIGURE 4 Legend on next page.

generally lower (Figure 6a vs. 6b), observation error and process noise were stronger (Figure 6c,e), and projected time to extinction was shorter (Figure 6f vs. 6g). Overall, these estimates also accorded with observed dynamics, which showed generally stable oscillations in the high-temperature replicates up to about Day 400 (Figure 6a) versus apparent extinction of *C. terricola* in all low-temperature replicates after Day 200 (Figure 6b).

DISCUSSION

Our results demonstrate the problems that complex dynamics without static equilibria pose for classic ecological stability metrics. When multiple underlying sources of variation are not properly accounted for, stability estimates based on raw observations of temporal variance overestimated variability, and therefore underestimated stability. Similarly, raw observations were generally not very good indicators of true system state, nor could they forecast the timing of rare state transitions such as extinctions. In contrast, the Particle-Takens filtering methods that we introduce here were able to correctly separate the effects of observation error, process noise, and deterministic dynamics, and produced substantially more accurate estimates of the strength of stochastic disturbances and of time to extinction. Jointly, these results highlight the importance of properly accounting for different sources of variation in ecological stability analysis (Loreau & de Mazancourt, 2013).

Importantly, by separating deterministic versus stochastic dynamics using the EDM particle filter, process noise parameter σ_P provides an estimate of variability that is akin to the classic CV metric, but can be estimated even in systems with nonstatic equilibria. In effect, these methods "detrend" deterministic variation and observation error from observed time series data before calculating variability. Ours is by no means the first method to develop a deterministically detrended analogue to CV (e.g., Arani et al., 2021; Guiz et al., 2016; Hamilton, Berry, & Sauer, 2017), but our approach has the great advantage that it does not require specifying the functional forms underlying deterministic dynamics and applies even in complex and high-dimensional systems. Critically, this link allows us to leverage existing ecological theory across a wide range of systems, even in cases with limited data and where average states vary greatly over time (Coulson, 2021; Pimm et al., 2019).

Note that in order to separate deterministic versus stochastic variation, our methods effectively define deterministic dynamics as those that are *repeatable* and *predictable* based on past system behavior, that is, given multiple replicates that start in exactly the same state, deterministic dynamics describe changes in system state that would play out identically across all replicates. In contrast, we define process noise as variability that is not repeatable or predictable, that is, causing replicates that start at the same state to diverge over short time spans (Shoemaker et al., 2019). Consequently, σ_P describes overall variability resulting from the joint effects of disturbances and of system recovery from those disturbances (Arnoldi et al., 2016). Theoretical analyses of simple linear systems have shown that, in some cases, this process noise term can be further decomposed into separate parameters representing resilience (i.e., the rate at which systems recover from disturbances) and resistance (i.e., the direct effect of a disturbance on system state) (Arnoldi et al., 2019; Clark et al., 2021). While it might be possible to apply similar decompositions to σ_P in the Particle-Takens framework (especially over short time spans), doing so would require making strong assumptions about the functional form of the deterministic function g.

Perhaps of greater practical importance, our results show that the process noise term σ_P can be combined with estimates of deterministic dynamics to predict exit times even if no regime shifts have yet been observed, for example, as demonstrated for per-timestep probability of extinction (Pr_{mor}) following Equation (5). Note that Pr_{mor} is closely related to the concept of persistence (Lande, 1987; Law & Morton, 1996) or estimates of time to extinction from integrated population models (Plard et al., 2019). For the process noise function in Equation (2a), extinction probability rises as a function of the strength of process noise relative to species abundance, that is, as $N(t)/\sigma_P$. Thus, high extinction risk can result either from deterministic dynamics that hold populations close to zero or from strong stochastic forcing caused by process noise. Pr_{mor} therefore captures at least one definition of press perturbations, that is, effects

FIGURE 4 Estimates of (a) observation error (σ_O) and (b, c) process noise (σ_P) for the simulated dynamics described in the main text. Black dashed line shows 1–1 relationship (i.e., "perfect" estimates). Blue shows particle filter estimates based on the correct analytical model, red shows estimates from the empirical dynamic model-based filter, and yellow shows estimated variability based on raw observations (i.e., based on the incorrect assumption that all observed variability is attributable to process noise). Shaded regions show mean ± 1 standard deviation. Panels (b) versus (c) separate process noise estimates for simulated time series with low (b) versus high (c) observation error. See Appendix S1: Figure S1 for a more detailed comparison. Obs., observation; Proc., process.
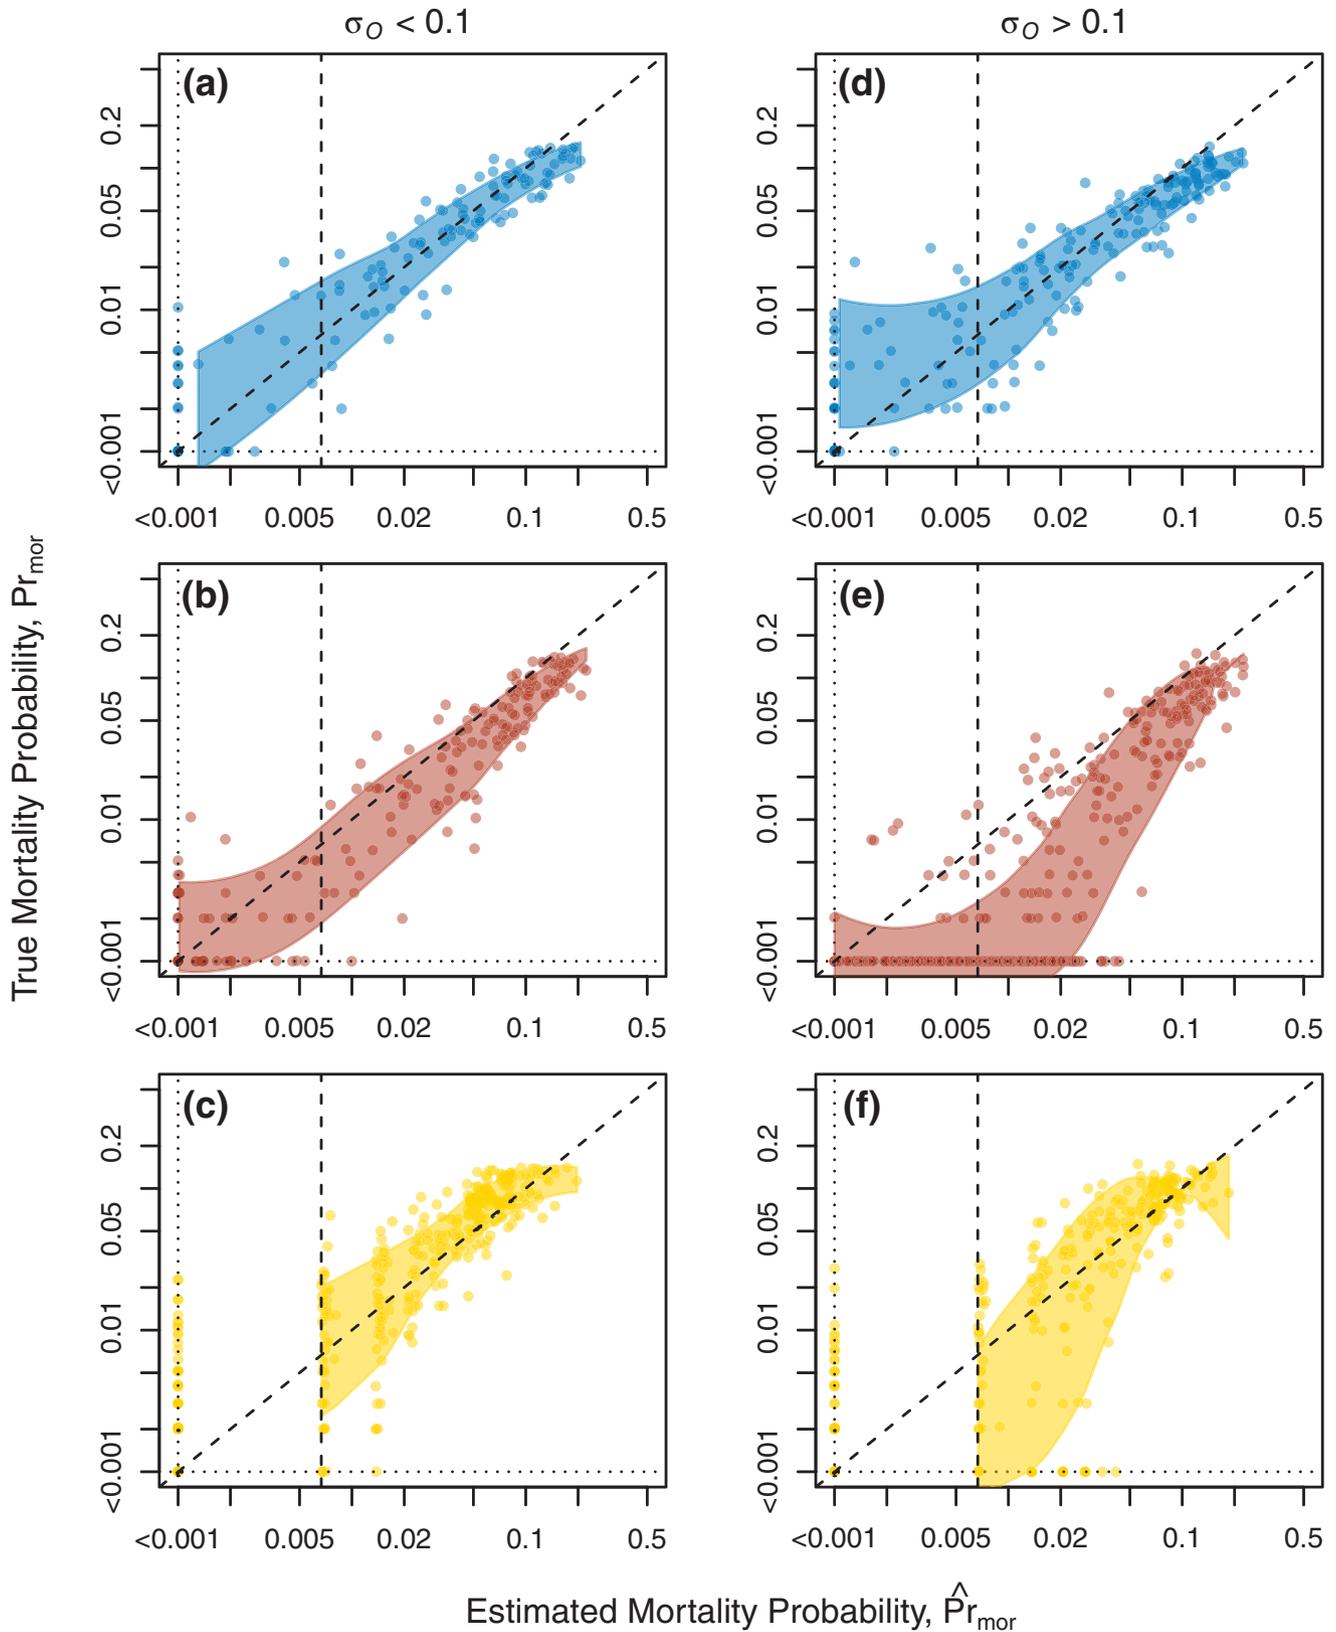
FIGURE 5 Performance for per-timestep mortality probability (Pr_{mor}) estimates for the simulated dynamics described in the main text. Blue shows particle filter estimates based on the correct analytical model, red shows estimates from the empirical dynamic model-based filter, and yellow shows estimates based on raw observations. Vertical dashed line shows "observational limit," that is, $1/t_{max}$, the smallest possible mortality probability that can still be detected based on raw observations. Otherwise, points, lines, and intervals are as described in the legend in Figure 4. Panels (a–c) versus (d–f) separate mortality probability estimates for simulated time series with low versus high observation error.
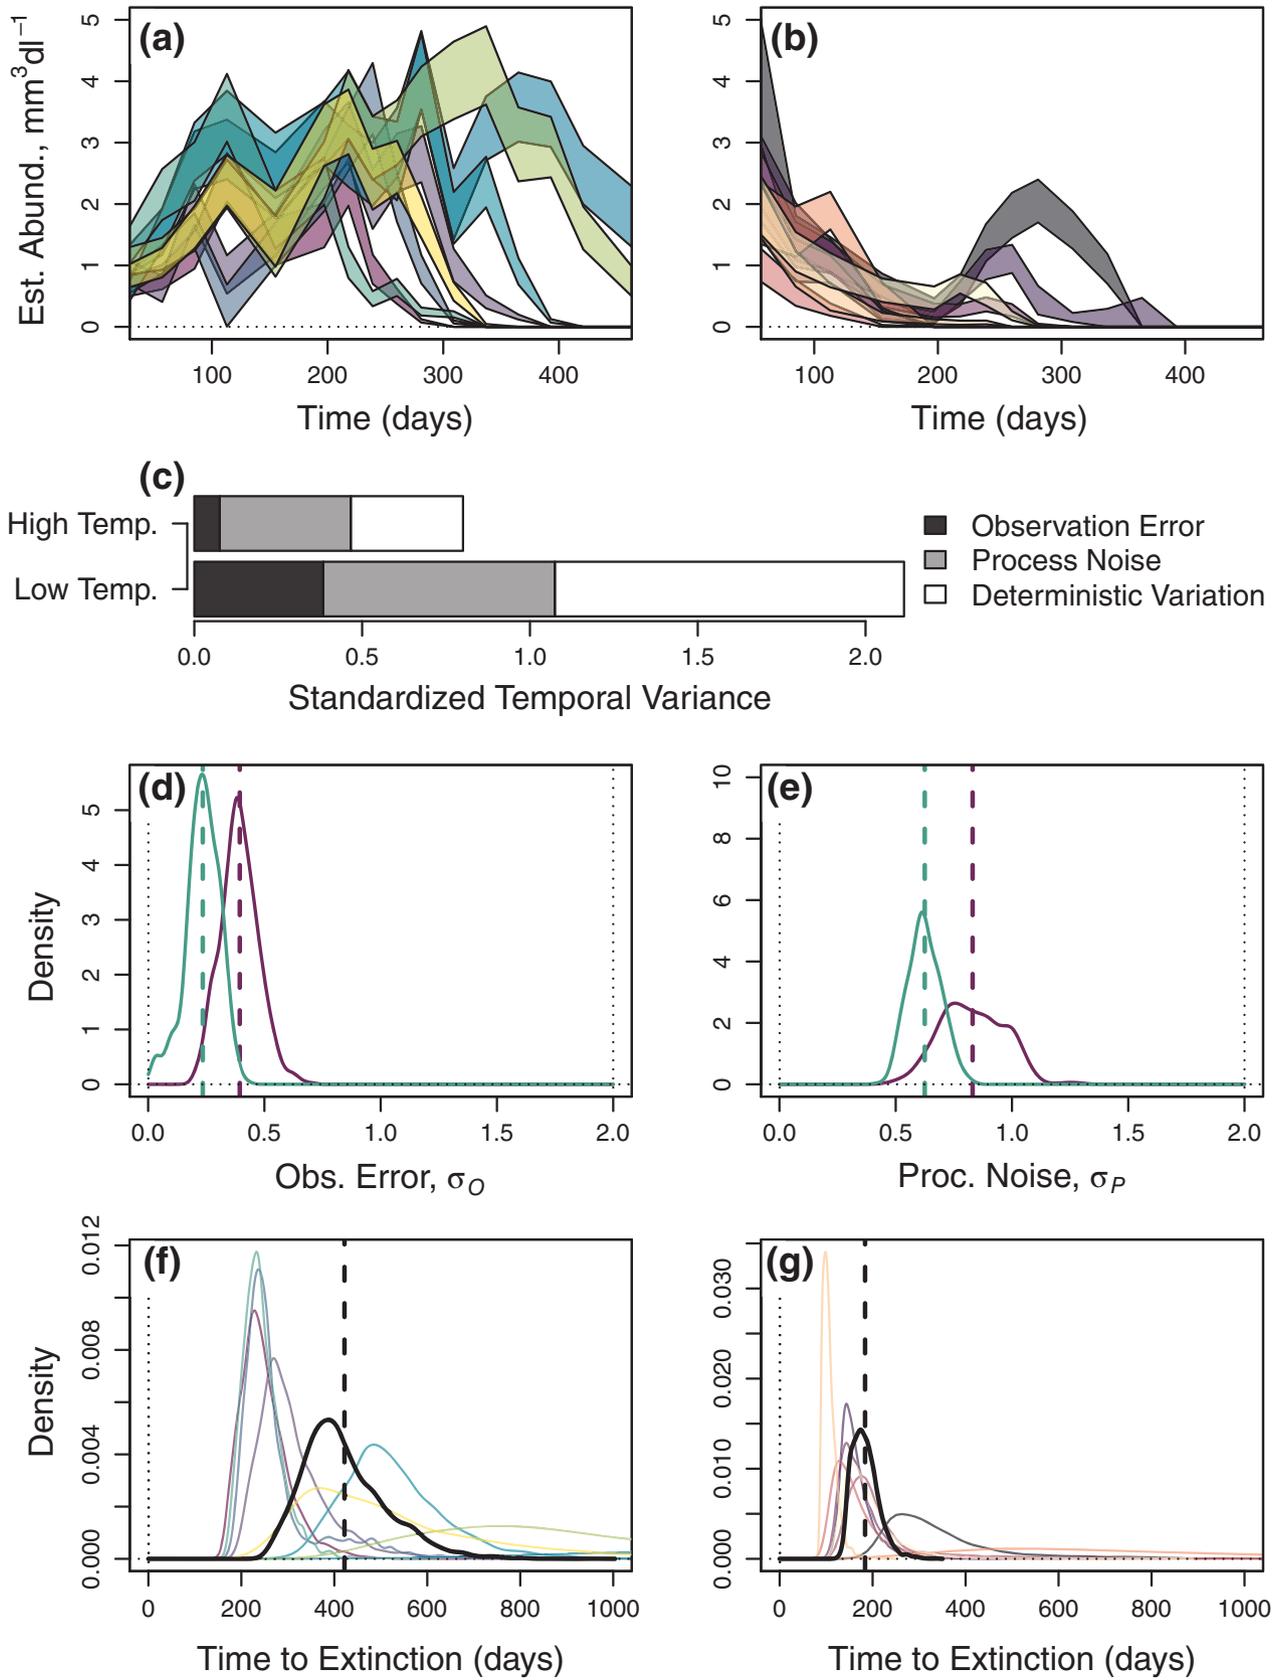
FIGURE 6 Legend on next page.

of slow, persistent changes in abundances, such as those caused by the time-varying K in our simulated examples, or the seasonal temperature fluctuations in our empirical examples. This insight accords with our results for *C. terricola* grown under cool conditions, for which the average displacement in abundances caused by press perturbations (σ_P) was roughly the same size as the average abundance (i.e., $<N(t)> \approx \sigma_P$), thus leading to relatively rapid extinction across all replicates.

Again, we emphasize that the primary contribution of the method presented here is that it applies in cases with very complex dynamics, and arbitrarily complex process noise and observation error functions, which is beyond the scope of most other existing methods. An important caveat is that this flexibility necessarily comes at the cost of decreased computational speed, lower precision, and poorer extrapolative ability. For example, each time series presented here took about 1–2 h to analyze with 10,000 MCMC iterations, whereas applications of the Kalman-Takens filter often run in just a few seconds (Hamilton, Berry, & Sauer, 2017). Similarly, accuracy and precision of estimates from our approach will usually be lower than for a well-tuned integrated population model—note, for example, that the particle filter fitted with the correct underlying deterministic function always outperformed the EDM (but see Perretti et al., 2012). In general, estimates from Particle-Takens filtering will only ever be as good as the EDM model and parametric stochastic functions used to fit them. For example, as EDM predictions grew poorer due to high levels of observation error in our simulations, more of the total observed variation was incorrectly subsumed into σ_P (see Appendix S1: Figure S2). Although such poorly tuned models may still provide reasonably good predictions of true dynamics, and even of extinction probabilities, they will no longer yield accurate estimates of the relative strength of deterministic versus stochastic forcing (Auger-Méthé et al., 2016). In some cases, it may be possible to improve performance by conducting repeated observations per timestep to help separate the effects of observation error (Knape & de Valpine, 2012), collecting longer time series or more spatial replicates (Clark et al., 2015), or incorporating covariates (Chang et al., 2021; Ye & Sugihara, 2016). Alternatively, in cases with very weak deterministic forcing, the S-mapping routine can be reduced to a simple estimate of the overall mean state, that is, by setting $E = 1$ and $\theta = 0$.

Despite these limitations, the Particle-Takens methods we apply here appear to work very well in practice, even in the presence of high observation error or strong process noise. In particular, for true abundance (Figures 3 and 6a,b) stochastic function parameters (Figures 4 and 6d,e), and even extinction probability (Figures 5 and 6f,g), the approach returned unbiased estimates for systems where observation error accounted for up to about 30% of total variability. Above this threshold, process noise strength and extinction probabilities tended to be overestimated, although estimates were still positively and significantly correlated with their true values (Appendix S1: Figure S1). In total, these results suggest that the approach should perform well across a wide range of empirical systems—and that results will almost always be less biased than those generated from raw observations of temporal variance.

Method extensions

Beyond the functionalities discussed in detail here, there are several additional features built into the pttsstability package that can be used to analyze more complex dynamics. First, the default function for describing process noise effects, procfun0, can be fit with two parameters rather than just one, yielding process noise that grows in intensity as a function of abundance following a Taylor power law (i.e., of the form $P(a_P, b_P, N_{est}'(t)) = \text{pmax}(-N_{est}'(t), \text{rnorm}(\mu = 0, \sigma = \sqrt(a_P N_{est}'(t)^\wedge b_P)))$) (Taylor, 1961). This functionality can be used, among other things, to incorporate the effects of demographic stochasticity versus environmental variability into fitted models, with $b_P < 2$ consistent with variability in demographic rates that decline as a function of abundance (Arnoldi et al., 2019). To use this feature, users simply need to input three sets of priors into the optimization function rather than two (see example analyses in Appendix S1: Figure S3, and in the analyze_burgmer_casestudy_powerscaling.R script in the ptstability_analyze_burgmer code repository).

FIGURE 6 Results from the empirical dynamic model-based particle filter, fitted to dynamics of the green algae species *Chlamydomonas terricola*. (a, b) Estimated true abundances under conditions with high mean temperature (a) or low mean temperature (b); shaded regions show mean $\pm$ 1 standard deviation. (c) Partitioning of temporal variance: Process noise magnitude is determined by function P in Equation (6a), observation error by function O in Equation (6b), and deterministic variation as the total remaining variation not attributable to P or O. (d, e) Distributions showing posterior estimates for the observation and process noise parameters under low (teal) and high (purple) temperatures. (f, g) Estimated time to extinction for each replicate (colored lines) and averaged across replicates (black line) under conditions with high mean temperature (f), or low mean temperature (g); dashed lines show mean estimates. Est. Abund., estimated abundance; Obs., observation; Proc., process; Temp., temperature.

Similarly, the default observation error function, obsfun0, can also be fit with two parameters instead of one, yielding observation error that grows linearly with abundance, but with a fixed lower bound (i.e., of the form $O(a_O, b_O, <N_{\text{obs}}>, N_{\text{est}}(t)) = \text{pmax}(-N_{\text{est}}(t), \text{rnorm}(\mu = 0, \sigma = a_O + b_O \times N_{\text{est}}(t))))$. This feature can be used to capture the effects of zero-inflated distributions (e.g., zero-inflated Poisson) and is helpful for modeling populations with low average abundances or frequent extinction events. To use this feature, users simply need to supply four sets of priors to the optimizer (which by default fits both the two-parameter process noise and the two-parameter observation error functions).

Lastly, for users who wish to go beyond the package defaults, the observation error, process noise, and deterministic functions are all editable, and can be replaced with arbitrarily complex user-defined functions. For details of necessary function inputs and outputs, see the help files for procfun0, obsfun0, detfun0 (for analytical deterministic models), EDMfun0 (for EDM deterministic models), and parseparam0 (the function used to pass parameters between the particle filter and simulation functions). For users who wish to apply multivariate EDM embeddings (e.g., to account for species interactions or environmental covariates) (Deyle et al., 2016; Ye & Sugihara, 2016; Ye et al., 2018), multivariate S-mapping can also be handled by the default EDM function, whereas multivariate simplex methods need to be custom-coded by the user. See Appendix S1 and the EDMfun0 help file for more details.

CONCLUSIONS

Our study shows that failing to account for multiple sources of variability can bias ecological stability analyses and presents a potential toolkit to overcome these challenges. We expect that these methods could be especially useful for analyzing systems with nonstatic equilibria (e.g., predator–prey oscillations), or undergoing rapid changes due to anthropogenic disturbances (Coulson, 2021; Pimm et al., 2019), and are therefore hopeful that our work will facilitate studies of stability across a much wider range of systems than has previously been possible.

CONFLICT OF INTEREST
The authors declare no conflict of interest.

DATA AVAILABILITY STATEMENT
All data and code (Clark, 2022a, 2022b) are available from Zenodo: https://doi.org/10.5281/zenodo.7215828 and https://doi.org/10.5281/zenodo.7215824.

ORCID
Adam Thomas Clark https://orcid.org/0000-0002-8843-3278
Lina K. Mühlbauer https://orcid.org/0000-0002-6493-2742
Helmut Hillebrand https://orcid.org/0000-0001-7449-1613
Canan Karakoç https://orcid.org/0000-0002-3921-3535

REFERENCES
Arani, B. M. S., S. R. Carpenter, L. Lahti, E. H. van Nes, and M. Scheffer. 2021. "Exit Time as a Measure of Ecological Resilience." *Science* 372(6547): eaay4895. https://doi.org/10.1126/science.aay4895.

Arnoldi, J.-F., M. Loreau, and B. Haegeman. 2016. "Resilience, Reactivity and Variability: A Mathematical Comparison of Ecological Stability Measures." *Journal of Theoretical Biology* 389: 47–59. https://doi.org/10.1016/j.jtbi.2015.10.012.

Arnoldi, J.-F., M. Loreau, and B. Haegeman. 2019. "The Inherent Multidimensionality of Temporal Variability: How Common and Rare Species Shape Stability Patterns." *Ecology Letters* 22(10): 1557–67. https://doi.org/10.1111/ele.13345.

Auger-Méthé, M., C. Field, C. M. Albertsen, A. E. Derocher, M. A. Lewis, I. D. Jonsen, and J. M. Flemming. 2016. "State-Space Models' Dirty Little Secrets: Even Simple Linear Gaussian Models Can Have Estimation Problems." *Scientific Reports* 6(1): 26677. https://doi.org/10.1038/srep26677.

Blasius, B., L. Rudolf, G. Weithoff, U. Gaedke, and G. F. Fussmann. 2020. "Long-Term Cyclic Persistence in an Experimental Predator–Prey System." *Nature* 577(7789): 226–30. https://doi.org/10.1038/s41586-019-1857-0.

Burgmer, T., and H. Hillebrand. 2011. "Temperature Mean and Variance Alter Phytoplankton Biomass and Biodiversity in a Long-Term Microcosm Experiment." *Oikos* 120(6): 922–33. https://doi.org/10.1111/j.1600-0706.2010.19301.x.

Chang, C.-W., T. Miki, M. Ushio, P.-J. Ke, H.-P. Lu, F.-K. Shiah, and C.-h. Hsieh. 2021. "Reconstructing Large Interaction Networks from Empirical Time Series Data." *Ecology Letters* 24: 2763–74. https://doi.org/10.1111/ele.13897.

Chesson, P. 2017. "AEDT: A New Concept for Ecological Dynamics in the Ever-Changing World." *PLOS Biology* 15(5): e2002634. https://doi.org/10.1371/journal.pbio.2002634.

Clark, A. 2022a. "adamtclark/pts_r_package: Code for pttstability R Package (1.0)." Zenodo. https://doi.org/10.5281/zenodo.7215828.

Clark, A. 2022b. "adamtclark/pttstability_analyze_burgmer: Empirical Datasets and Code for pttstability Package (1.0)." Zenodo. https://doi.org/10.5281/zenodo.7215824.

Clark, A. T., J.-F. Arnoldi, Y. R. Zelnik, G. Barabas, D. Hodapp, C. Karakoç, S. König, et al. 2021. "General Statistical Scaling Laws for Stability in Ecological Systems." *Ecology Letters* 24(7): 1474–86. https://doi.org/10.1111/ele.13760.

Clark, A. T., H. Ye, E. R. Forest Isbell, J. C. Deyle, D. Tilman, and G. Sugihara. 2015. "Spatial 'Convergent Cross Mapping' to Detect Causal Relationships from Short Time-Series." *Ecology* 96(6): 1174–81. https://doi.org/10.1890/14-1479.1.

Coulson, T. 2021. "We Live in a Changing World, but That Shouldn't Mean We Abandon the Concept of Equilibrium." *Ecology Letters* 24(1): 3–5. https://doi.org/10.1111/ele.13629.

de Mazancourt, C., F. Isbell, A. Larocque, F. Berendse, E. De Luca, J. B. Grace, B. Haegeman, et al. 2013. "Predicting Ecosystem Stability from Community Composition and Biodiversity." *Ecology Letters* 16(5): 617–25. https://doi.org/10.1111/ele.12088.

de Valpine, P., and A. Hastings. 2002. "Fitting Population Models Incorporating Process Noise and Observation Error." *Ecological Monographs* 72(1): 57–76. https://doi.org/10.1890/0012-9615(2002)072[0057:FPMIPN]2.0.CO;2.

DeAngelis, D. L., and J. C. Waterhouse. 1987. "Equilibrium and Nonequilibrium Concepts in Ecological Models." *Ecological Monographs* 57: 1–21. https://doi.org/10.2307/1942636.

DeAngelis, D. L., and S. Yurek. 2015. "Equation-Free Modeling Unravels the Behavior of Complex Ecological Systems." *Proceedings of the National Academy of Sciences* 112(13): 3856–7. https://doi.org/10.1073/pnas.1503154112.

Deyle, E. R., R. M. May, S. B. Munch, and G. Sugihara. 2016. "Tracking and Forecasting Ecosystem Interactions in Real Time." *Proceedings of the Royal Society B: Biological Sciences* 283(1822): 20152258. https://doi.org/10.1098/rspb.2015.2258.

Deyle, E. R., and G. Sugihara. 2011. "Generalized Theorems for Nonlinear State Space Reconstruction." *PLoS One* 6(3): e18295.

Donohue, I., H. Hillebrand, J. M. Montoya, O. L. Petchey, S. L. Pimm, M. S. Fowler, K. Healy, et al. 2016. "Navigating the Complexity of Ecological Stability." *Ecology Letters* 19(9): 1172–85. https://doi.org/10.1111/ele.12648.

Donohue, I., O. L. Petchey, J. M. Montoya, A. L. Jackson, L. McNally, M. Viana, K. Healy, M. Lurgi, N. E. O'Connor, and M. C. Emmerson. 2013. "On the Dimensionality of Ecological Stability." *Ecology Letters* 16(4): 421–9. https://doi.org/10.1111/ele.12086.

Evans, M. R. 2012. "Modelling Ecological Systems in a Changing World." *Philosophical Transactions of the Royal Society B: Biological Sciences* 367(1586): 181–90. https://doi.org/10.1098/rstb.2011.0172.

Grimm, V., and C. Wissel. 1997. "Babel, or the Ecological Stability Discussions: An Inventory and Analysis of Terminology and a Guide for Avoiding Confusion." *Oecologia* 109(3): 323–34. https://doi.org/10.1007/s004420050090.

Guiz, J., H. Hillebrand, E. T. Borer, M. Abbas, A. Ebeling, A. Weigelt, Y. Oelmann, et al. 2016. "Long-Term Effects of Plant Diversity and Composition on Plant Stoichiometry." *Oikos* 125(5): 613–21. https://doi.org/10.1111/oik.02504.

Hamilton, F., T. Berry, and T. Sauer. 2017. "Kalman-Takens Filtering in the Presence of Dynamical Noise." *The European Physical Journal Special Topics* 226(15): 3239–50. https://doi.org/10.1140/epjst/e2016-60363-2.

Hamilton, F., A. L. Lloyd, and K. B. Flores. 2017. "Hybrid Modeling and Prediction of Dynamical Systems." *PLoS Computational Biology* 13(7): e1005655. https://doi.org/10.1371/journal.pcbi.1005655.

Hartig, F., F. Minunno, and S. Paul. 2019. "BayesianTools: General-Purpose MCMC and SMC Samplers and Tools for Bayesian Statistics." https://CRAN.R-project.org/package=BayesianTools.

Hastings, A., C. L. Hom, S. Ellner, P. Turchin, H. Charles, and J. Godfray. 1993. "Chaos in Ecology: Is Mother Nature a Strange Attractor?" *Annual Review of Ecology and Systematics* 24(1): 1–33. https://doi.org/10.1146/annurev.es.24.110193.000245.

Holling, C. S. 1973. "Resilience and Stability of Ecological Systems." *Annual Review of Ecology and Systematics* 4(1): 1–23. https://doi.org/10.1146/annurev.es.04.110173.000245.

Hsieh, C.-h., C. Anderson, and G. Sugihara. 2008. "Extending Nonlinear Analysis to Short Ecological Time Series." *American Naturalist* 171(1): 71–80.

Ives, A. R., K. Gross, and J. L. Klug. 1999. "Stability and Variability in Competitive Communities." *Science* 286(5439): 542–4. https://doi.org/10.1126/science.286.5439.542.

Knape, J., and P. de Valpine. 2012. "Fitting Complex Population Models by Combining Particle Filters with Markov Chain Monte Carlo." *Ecology* 93(2): 256–63. https://doi.org/10.1890/11-0797.1.

Lande, R. 1987. "Extinction Thresholds in Demographic Models of Territorial Populations." *The American Naturalist* 130(4): 624–35. https://doi.org/10.1086/284734.

Law, R., and R. D. Morton. 1996. "Permanence and the Assembly of Ecological Communities." *Ecology* 77(3): 762–75. https://doi.org/10.2307/2265500.

Lewontin, R. C. 1969. "The Meaning of Stability." *Brookhaven Symposia in Biology* 22: 13–24.

Loreau, M., and C. de Mazancourt. 2013. "Biodiversity and Ecosystem Stability: A Synthesis of Underlying Mechanisms." *Ecology Letters* 16: 106–15. https://doi.org/10.1111/ele.12073.

Massoud, E. C., J. Huisman, E. Benincà, M. C. Dietze, W. Bouten, and J. A. Vrugt. 2018. "Probing the Limits of Predictability: Data Assimilation of Chaotic Dynamics in Complex Food

Webs." *Ecology Letters* 21(1): 93–103. https://doi.org/10.1111/ele.12876.

May, R. M. 1973. "Stability and Complexity in Model Ecosystems." In *Monographs in Population Biology*, Vol 6. Princeton, NJ: Princeton University Press.

Ovaskainen, O., and B. Meerson. 2010. "Stochastic Models of Population Extinction." *Trends in Ecology & Evolution* 25(11): 643–52. https://doi.org/10.1016/j.tree.2010.07.009.

Park, J., C. Smith, G. Sugihara, and E. R. Deyle. 2020. "REDM: Empirical Dynamic Modeling ('EDM')." https://CRAN.R-project.org/package=rEDM.

Perretti, C. T., G. Sugihara, and S. B. Munch. 2012. "Nonparametric Forecasting Outperforms Parametric Methods for a Simulated Multi-Species System." *Ecology* 94(4): 794–800.

Pimm, S. L. 1984. "The Complexity and Stability of Ecosystems." *Nature* 307(5949): 321–6. https://doi.org/10.1038/307321a0.

Pimm, S. L., I. Donohue, J. M. Montoya, and M. Loreau. 2019. "Measuring Resilience Is Essential to Understand It." *Nature Sustainability* 2(10): 895–7. https://doi.org/10.1038/s41893-019-0399-7.

Plard, F., R. Fay, M. Kéry, A. Cohas, and M. Schaub. 2019. "Integrated Population Models: Powerful Methods to Embed Individual Processes in Population Dynamics Models." *Ecology* 100: e02715. https://doi.org/10.1002/ecy.2715.

R Development Core Team. 2019. *R: A Language and Environment for Statistical Computing*. Vienna: R Foundation for Statistical Computing.

Schaffer, W. M. 1984. "Stretching and Folding in Lynx Fur Returns: Evidence for a Strange Attractor in Nature?" *The American Naturalist* 124(6): 798–820. https://doi.org/10.1086/284318.

Scheffer, M., J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. van Nes, M. Rietkerk, and G. Sugihara. 2009. "Early-Warning Signals for Critical Transitions." *Nature* 461(7260): 53–9. https://doi.org/10.1038/nature08227.

Shoemaker, L. G., L. M. Hallett, L. Zhao, D. C. Reuman, S. Wang, K. L. Cottingham, R. J. Hobbs, et al. 2022. "The Long and the Short of It: Mechanisms of Synchronous and Compensatory Dynamics across Temporal Scales." *Ecology* 103: e3650. https://doi.org/10.1002/ecy.3650.

Shoemaker, L. G., L. L. Sullivan, I. Donohue, J. S. Cabral, R. J. Williams, M. M. Mayfield, J. M. Chase, et al. 2019. "Integrating the Underlying Structure of Stochasticity into

Community Ecology." *Ecology* 101: e02922. https://doi.org/10.1002/ecy.2922.

Sugihara, G., B. Grenfell, and R. M. May. 1990. "Distinguishing Error from Chaos in Ecological Time-Series." *Philosophical Transactions of the Royal Society B: Biological Sciences* 330(1257): 235–51.

Sugihara, G., R. May, H. Ye, E. R. Chih-hao Hsieh, M. F. Deyle, and S. Munch. 2012. "Detecting Causality in Complex Ecosystems." *Science* 338(6106): 496–500.

Takens, F. 1981. "Detecting Strange Attractors in Turbulence." In *Dynamical Systems and Turbulence, Warwick 1980*, Vol 898, edited by D. Rand and L.-S. Young, 366–81. Lecture Notes in Mathematics. Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/BFb0091924.

Taylor, L. R. 1961. "Aggregation, Variance and the Mean." *Nature* 189(4766): 732–5. https://doi.org/10.1038/189732a0.

Tilman, D. 1995. "Biodiversity: Population Versus Ecosystem Stability." *Ecology* 77(2): 350–63. https://doi.org/10.2307/2265614.

Wilkinson, D. J. 2018. "Smfsb: Stochastic Modelling for Systems Biology. R Package Version 1.3." https://CRAN.R-project.org/package=smfsb.

Wilkinson, D. J. 2011. *Stochastic Modelling for Systems Biology*. Mathematical and Computational Biology, 2nd ed., Vol 44. Boca Raton, FL: Chapman & Hall.

Ye, H., A. T. Clark, E. R. Deyle, S. Munch, O. Keyes, J. Cai, E. White, et al. 2018. "REDM: Applications of Empirical Dynamic Modeling from Time Series." Zenodo. https://doi.org/10.5281/ZENODO.1294063.

Ye, H., and G. Sugihara. 2016. "Information Leverage in Interconnected Ecosystems: Overcoming the Curse of Dimensionality." *Science* 353(6302): 922–5. https://doi.org/10.1126/science.aag0863.

SUPPORTING INFORMATION

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