Ecological Applications, 27(7), 2017, pp. 2048–2060

Prediction in ecology: a first-principles framework

Michael C. Dietze[1]

Department of Earth and Environment, Boston University, 685 Commonwealth Avenue, Room 130, Boston, Massachusetts 02215 USA

Abstract. Quantitative predictions are ubiquitous in ecology, yet there is limited discussion on the nature of prediction in this field. Herein I derive a general quantitative framework for analyzing and partitioning the sources of uncertainty that control predictability. The implications of this framework are assessed conceptually and linked to classic questions in ecology, such as the relative importance of endogenous (density-dependent) vs. exogenous factors, stability vs. drift, and the spatial scaling of processes. The framework is used to make a number of novel predictions and reframe approaches to experimental design, model selection, and hypothesis testing. Next, the quantitative application of the framework to partitioning uncertainties is illustrated using a short-term forecast of net ecosystem exchange. Finally, I advocate for a new comparative approach to studying predictability across different ecological systems and processes and lay out a number of hypotheses about what limits predictability and how these limits should scale in space and time.

Key words: ecological forecasting; endogenous; exogenous; net ecosystem exchange; parameter; process error; random effects; scale; stability; uncertainty.

Introduction

How predictable is nature? This question is central to ecology, regardless of whether the goal is to comprehend nature from an empirical perspective, to seek a more general theoretical understanding, or to provide practical insight into the management of natural systems. Any ecologist making measurements in the field or the lab seeks to find patterns in the natural world. In doing so, they rely on observations and experiments to test hypotheses about what processes lead to consistent, predictable outcomes and to distinguish these from patterns that arise by chance. Generations of ecologists have filled journals, books, and dissertations with examples of predictable patterns in nature. Theorists and modelers have aimed to synthesize and abstract this understanding to better understand the general principles underlying processes. The best of these theories and models make explicit predictions, which can lead both to new testable hypotheses under novel conditions and allow applied ecologists to better manage and conserve the natural world.

However, while questions about predictability cross all disciplines of ecology, there has been comparatively little discussion on the nature of prediction in ecology, from either a theoretical or practical perspective (Evans et al. 2013, Mouquet et al. 2015, Petchey et al. 2015, Houlahan et al. 2017). The importance of ecological predictability for practical applications is clear: to make ecology more relevant to policy, management, and decision making, we need a better understanding of what we can forecast and how those forecasts can be improved (Clark et al. 2001, Dietze 2017, Dietze et al. 2017). Furthermore, understanding what factors affect predictability directly impacts what data we collect, how models are structured, and the statistical tools we use to link models to data. However, the need to understand predictability also has tremendous importance for advancing basic research. Prediction embodies the core principles of the scientific method, allowing us to make a specific, quantitative, and testable prediction about the future, to observe the outcome, and to refine our hypotheses. Furthermore, the lack of a general framework for prediction in ecology impedes both the analysis of specific problems and the search for generality across systems.

The goal of this paper is to present a general first-principles framework to understanding the predictability of dynamic ecological systems. Assessing predictability requires explicit metrics for measuring predictability, and herein I focus on prediction uncertainty, the rate at which prediction uncertainty increases into the future, and how prediction uncertainty scales in space and time. In the first section, a general framework is derived and linked to broad conceptual themes in ecology, as well as to practical challenges in data collection, statistical analysis, modeling, and forecasting. Qualitatively, this framework facilitates discussion of the factors that determine predictability, links these factors to classic debates in ecology, and makes a number of explicit predictions and recommendations. Quantitatively, it enables comparison of the absolute and relative contributions of different sources of uncertainty, both within a given problem and across different ecological processes. In the second

Manuscript received 7 February 2017; revised 17 May 2017; accepted 24 May 2017. Corresponding Editor: David S. Schimel.

[1] E-mail: dietze@bu.edu

section, as an example, this approach is applied to the prediction of short-term forest carbon fluxes. In the third section, I end with a discussion of new questions and research directions in predictability across ecology. Specifically, I advocate for a comparative approach to assessing the factors that determine predictability across systems, which allows us to look for patterns and generalities that span all subdisciplines. Given limited time and resources, this knowledge would likewise be eminently practical, as it would allow us to better target resources at understanding the dominant sources of variability for a particular class of problem, rather than having to tackle all sources of uncertainty equally. Furthermore, this framework has the potential to stimulate new basic research and accelerate progress across a wide range of fields, especially when applied through an iterative cycle of near-term prediction and hypothesis refinement.

A First-Principles Approach

Consider the general problem of trying to predict a specific ecological quantity of interest, Y, that describes the state of the system (e.g., population size, presence/absence at a specific location, species composition, number of infected individuals, biomass, nutrient pools) at some specific point in time, t. In this context, Y_t is known as a *state variable*. For most ecological processes the state of the system in the future, $Y_{t+\Delta t}$, will depend, in part, on the *current* state of the system, Y_t. This dependency of the future state on the current state defines what is known as a *dynamic system*. The relationship between these two states in time is driven by both the internal, or endogenous, dynamics and by external, or exogenous, drivers/covariates, which I'll denote here as X_t. Both Y_t and X_t may be either single values (scalars) or vectors describing multiple states and/or drivers. Let's next assume that there's some function, f, that describes our current understanding about how Y_t and X_t affect $Y_{t+\Delta t}$ and that this function requires a set of parameters θ. Because there will always be some difference between the predicted state of the system and its true state, we will also include a *process error*, ε_t. This process error is not the residual observation error, but the *dynamical* error in the predicted transition from one point in time to another. Finally, for most ecological predictions, the parameters θ are not physical constants, but empirically estimated coefficients. Because ecological processes can be quite variable across space, time, and taxa, let's split θ into two components, $\bar{\theta} + \alpha$, where $\bar{\theta}$ describes the *mean* of the parameters and α is the deviation from that mean experienced at a particular time, location, species, etc. In statistics, α is known as a *random effect* and is used to accommodate systematic variability in processes across multiple scales that is observed but not (yet) explained by the internal dynamics or external covariates.

Putting together the factors described above gives us the following general model of ecological dynamics

$$Y_{t+1} = f(Y_t, X_t|\bar{\theta} + \alpha) + \varepsilon_t \quad (1)$$

This modeling framework describes a large fraction of the statistical and process-based models used across many ecological subdisciplines, such as most models of population and community dynamics and ecosystem and biogeochemical pools and fluxes. The model also encompasses important special cases, such as $Y_{t+1} = f(X_t|\bar{\theta}) + \varepsilon_t$, which includes linear and nonlinear regression models, and $Y_{t+1} = f(X_t|\bar{\theta} + \alpha) + \varepsilon_t$, which includes mixed models. Furthermore, this formulation is agnostic to whether f takes on a specific parametric functional form or is represented by more general, flexible approaches (splines, GAMs, random forests, etc.).

Given this general dynamic model (Eq. 1), we can assess how different factors in the model affect the predictability of the system. While there are a number of different potential measures of predictability, one critical component is the uncertainty in the prediction, which we can quantify in terms of the variance in the prediction, $\mathrm{Var}[Y_{t+1}]$. Predictive variance generally increases with time, and the limit to predictability is determined by when the prediction is doing no better than chance (i.e., uncertainty exceeds natural variability or some other null model, Fig. 1A). The exact contributions of each factor in Eq. 1 to the predictive variance will depend on the specific functional form of f and the probability density functions for each of the factors considered, $\{Y_t, X_t\bar{\theta}, \alpha, \varepsilon\}$, but a general understanding can be found from a linear approximation of f. If we assume, for simplicity, that the covariance across the different factors is negligible, the predictive variance can be written as

$$\begin{aligned}\mathrm{Var}[Y_{t+1}] \approx & \underbrace{\left(\frac{\partial f}{\partial Y}\right)^2}_{\text{stability}} \underbrace{\mathrm{Var}[Y_t]}_{\text{IC uncert}} + \underbrace{\left(\frac{\partial f}{\partial X}\right)^2}_{\text{driver sens}} \underbrace{\mathrm{Var}[X]}_{\text{driver uncert}} \\ & + \underbrace{\left(\frac{\partial f}{\partial \theta}\right)^2}_{\text{param sens}} \left(\underbrace{\mathrm{Var}[\bar{\theta}]}_{\text{param uncert}} + \underbrace{\mathrm{Var}[\alpha]}_{\text{param variablility}} \right) \\ & + \underbrace{\mathrm{Var}[\varepsilon]}_{\text{process error}} \\ = & \,\mathrm{INTERNAL} + \mathrm{EXTERNAL} \\ & + (\mathrm{PARAMETERS} + \mathrm{RANDOM\ EFFECTS}) \\ & + \mathrm{PROCESS\ ERROR}\end{aligned} \quad (2)$$

This equation follows a simple and intuitive pattern; the contribution of each factor can be expressed in terms of the sensitivity of f to that factor and the factor's uncertainty. In general terms, this means that the most important factors affecting the predictability of any particular ecological process are those that are both highly sensitive and highly uncertain. More broadly, any of the factors can be important if they have high sensitivity or are highly uncertain, and knowing about sensitivity or

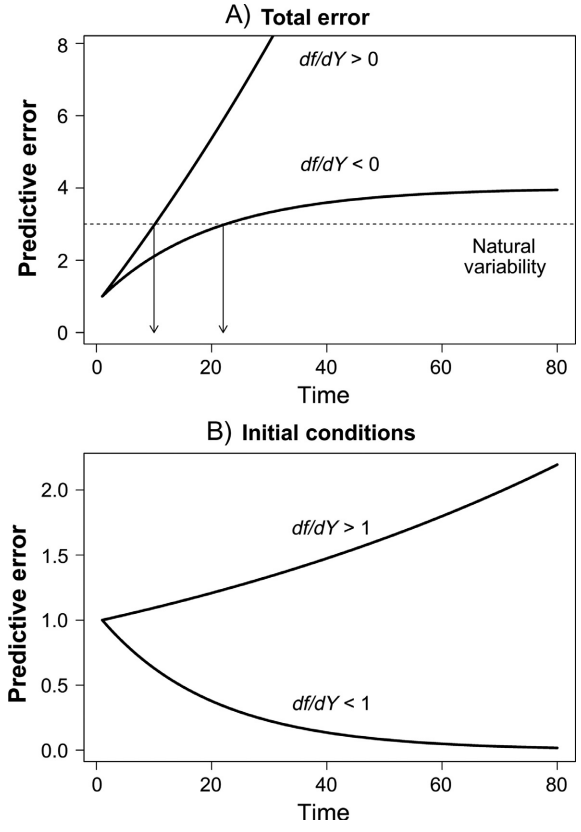

FIG. 1. Conceptual relationship between predictive error (square root of predictive variance) and predictability. (A) The limit to predictability (arrows) occur when forecast uncertainty exceeds our null expectation (typically determined by natural variability). (B) The internal stability of the system, *df/dy*, determines whether initial condition uncertainty grows or decays exponentially.

uncertainty alone is not sufficient to understand predictability (LeBauer et al. 2013).

In the sections that follow, I discuss some general insights that can be gleaned by considering each factor in turn. However, the power of Eq. 2 in any specific application comes from its ability to *quantitatively* break down an overall prediction into its components. There is much to be learned about the predictability of ecological systems by comparing the relative contributions of these factors to the predictability of different ecological processes.

Endogenous Stability and Initial Condition Uncertainty

The first term in Eq. 2 describes the contribution of the internal dynamics of an ecological system to its predictability. This term can be broken down into the uncertainty about the current state of the system, also known as *initial condition (IC) uncertainty*, and the internal sensitivity, also known as the *endogenous stability*. The stability in this context is the exact same concept of stability taught in general ecology textbooks when discussing classic ecological models (Lotka 1910, Volterra 1926), and frequently accompanied by pictures of balls rolling off hilltops (unstable equilibria) and into valleys (stable equilibria). That classic theoretical concepts show up prominently in Eq. 2 is heartening, as it means decades of ecological research and teaching are directly relevant to questions of predictability.

As with classic theory, the threshold between stable and unstable dynamics is important for understanding predictability. From the perspective of Y_t, Eq. 2 is a recursive model ($\mathrm{Var}[Y_{t+1}]$ is a function of $\mathrm{Var}[Y_t]$) directly analogous to the recursive model for exponential growth. Therefore, if $|\partial f/dY| > 1$ the system is unstable and the initial condition uncertainty, $\mathrm{Var}[Y_t]$, will grow exponentially with time and quickly dominate all other terms (Fig. 1B). This situation occurs in another familiar prediction problem, weather forecasting. Since the atmosphere is inherently unstable, any uncertainty about the current state of the atmosphere grows rapidly through time. Because of this, meteorological forecasting is generally considered an initial condition problem (Lorenz 1963). This is not just a theoretical curiosity. Massive amounts of atmospheric data are collected every day globally (ground measurements, radar and other upward-looking sensors, weather balloons, airborne measurements, and a wide range of weather satellites), most of which is collected primarily so that it can be fed into model-data assimilation systems to minimize initial condition uncertainty (Kalnay 2002, Lewis et al. 2006, Evensen 2009).

By contrast to unstable systems, where initial condition uncertainty dominates, when $|\partial f/dY| < 1$ a system is said to have stabilizing feedbacks and the initial condition uncertainty will decay exponentially over time (Fig. 1B). This exponential decay does not guarantee that the initial condition uncertainty will be small, indeed it may still drive near-term forecasts, but it does suggest that other sources of uncertainty will eventually become dominant. Some ecological systems, such as epidemic diseases, do appear to be chaotic or to exhibit critical periods of inherent instability (Bjørnstad et al. 2002, Ferrari et al. 2008). That said, research on chaos more generally in ecology has met with mixed results (Ellner and Turchin 1995, 2005, Pastor and Cohen 1997, Schimel et al. 2005, Bjørnstad 2015) and there are many examples of stabilizing feedbacks in ecology. Furthermore, truly chaotic populations would fluctuate wildly, resulting in frequent population crashes, genetic bottlenecks, and high risk of stochastic extinction. Therefore, as a working hypothesis, I posit that for most ecological forecasting problems we cannot ignore the other terms in Eq. 2.

Another important insight from Eq. 2 is that the first term is the only one that has a recursive feedback, and thus the only term that will grow or decay exponentially. By contrast, all other terms will respond linearly through time. Because variances and squared sensitivities are always positive, this implies that all other terms will

strictly increase with time. Therefore a general expectation is that forecast uncertainty will increase with time. More specifically, uncertainty will increase whenever

$$\left(\frac{\partial f}{\partial Y}\right)^2 > 1 - \frac{\left(\frac{\partial f}{\partial X}\right)\mathrm{Var}[X]^2 + \left(\frac{\partial f}{\partial \theta}\right)^2(\mathrm{Var}[\bar{\theta}] + \mathrm{Var}[\alpha]) + \mathrm{Var}[\varepsilon]}{\mathrm{Var}_{[Y_t]}} \quad (3)$$

Because the ratio on the right-hand side (RHS) is positive, uncertainty will always increase for unstable systems. For stable systems, if uncertainty is initially increasing then the denominator will increase through time, reducing the ratio, and vice versa. In either case, the uncertainty will converge to a steady state (Fig. 1 Top). A useful corollary to this analysis is that care should be taken when interpreting the predictability of models that only include internal dynamics (e.g., most classic population models), both because they lack all the terms in the numerator of Eq. 3 and because IC uncertainty should decline over time in most ecological systems.

Finally, many classic ecological debates have centered around the stability and predictability of system dynamics, such as arguments about climax communities (Clements 1936) vs. individualistic responses (Gleason 1926) vs. neutral processes (Hubbell 2001). In practice, the observed predictability of these dynamics can be very scale dependent (Levin 1992), in a way that is consistent with the expectations set in Eqs. 2 and 3. For example, when looking at vegetation dynamics at larger spatial and temporal scales, a lot of external and process variability is averaged over (Turner et al. 1993), reducing these terms in the numerator. This leads to a steady-state variance, as opposed to the increasing variance predicted for neutral systems $|\partial f/dY| = 1$ that lack intrinsic stability (Clark and McLachlan 2003). By contrast, at finer spatial and temporal scales, these same terms are larger, and thus the predictability of the dynamics would be lower (Norden et al. 2015), even if the intrinsic stability were to remain the same.

Exogenous Stability and Uncertainty

If the first term in Eq. 2 describes a system's *endogenous (internal) stability*, then the second term describes its *exogenous (external) sensitivity*, its sensitivity to external forcing. Consider first Var[X], the uncertainty in the model's covariates or drivers. Eq. 2 demonstrates that systems that are sensitive to unpredictable drivers will themselves be unpredictable. On the flip side, higher predictability (lower variance) occurs either when the system is insensitive to environmental variability (e.g., a resilient life-history stage, such as a seed bank, that buffers a population) or sensitive to drivers that are highly predictable (e.g., synchrony of an ecological process to diurnal, tidal, or annual cycles; Luo et al. 2015).

Predictability can likewise be context and scale dependent. For example, a process that is sensitive to spatial variability in some input (e.g., soils) may be highly predictable at a given location but hard to predict across space due to higher input uncertainty. Likewise, as noted earlier, spatial and temporal variability in drivers may average out at larger scales, especially if drivers at different locations or times are *decoupled*. This illustrates another important point: in deriving Eq. 2 we assumed that the covariances *among* the terms was negligible, but this is unlikely to be true when assessing the contributions *within* a term. The more general form, when considering multiple drivers, is

$$\sum\sum \frac{\partial f}{\partial X_i}\frac{\partial f}{\partial X_j}\mathrm{COV}[X_i, X_j] \quad (4)$$

from which we can see that the covariance between drivers can either increase the overall uncertainty (positive covariance and derivatives the same sign OR negative covariance and derivatives opposite sign) or decrease the uncertainty (vice versa) depending on the signs of the covariances and the derivatives. When considering the same driver at different times or locations, as in the above example, such drivers will typically be autocorrelated (Cohen and Pastor 1991). This implies that the uncertainty will be higher, and variability will be slower to average out, than if sites were independent. For any particular problem, the autocorrelation in time or space can be quantified (e.g., via a correlogram) to make a specific numerical (and potentially falsifiable) prediction about the rate at which uncertainties should average out when scaling up.

By contrast, if X_i and X_j are two different drivers, we will have little a priori expectation for whether their covariance will be positive or negative. At first glance our expectations for the signs of the sensitivities may likewise seem unknown as well, but all else being equal these responses to environmental variables often represent ecological and physiological trade-offs, in which case we might expect their slopes to be aligned in the way to reduce variability. That said, making strong assumptions here could easily lead us astray, and thus the more important recommendation is a reminder of the need to consider such interactions when making measurements and building models.

The other important observation about the exogenous term in Eq. 2 is that if a forecast of Y depends on X, then we must be able to forecast X as well. While the exogenous term is not itself recursive, the forecast of X generally will be. Therefore the same conclusions we reached in the previous section about the increase in forecast uncertainty for Y will likewise apply to X. Specifically, Var[X] is expected to increase with time, at a rate depending on the strength of its own stabilizing feedbacks and the uncertainties involved in its prediction. All else being equal, this implies that the relative importance of Var[X] will increase with time. That said, some drivers will increase in

uncertainty quickly (e.g., weather) while others may be so slow to change as to be essentially constant for many forecasts (e.g., topography, soils).

Because driver uncertainty tends to increase with time, the covariates useful for making predictions may be different from those used for explaining the same process. Consider two possible explanatory variables X_1 and X_2. Let's assume that, when using past data, X_1 is found to be a better predictor of Y than X_2. All else being equal, this implies that we should use X_1 to forecast Y into the future, but what if the future uncertainty in X_1, Var[X_1], is larger than that for X_2. How do we decide whether to use X_1 or X_2 when forecasting Y? Eq. 2 gives us a precise prediction on this matter. If Var[ε_1] and Var[ε_1] are the process errors associated with the models containing X_1 and X_2 respectively, then X_1 will produce a lower uncertainty prediction than X_2 if $(\partial f/\partial X_1)^2\text{Var}[X_1] + \text{Var}[\varepsilon_1] < (\partial f/\partial X_2)^2\text{Var}[X_2] + \text{Var}[\varepsilon_2]$. Conceptually, this implies that we favor X_1 if its explanatory power (lower process error) is sufficient to offset its higher uncertainty, but the exact threshold depends on the relative sensitivities of X_1 and X_2. Another way of looking at this is that the sensitivities are required to express Var[X] and Var[ε] on the same scale so that they can be compared.

Another important conclusion implied by Eq. 2 is that the experimental or observational design for prediction is often different from that used for hypothesis testing. In the classic hypothesis testing framework, the question being asked is "does X affect Y." By contrast, Eq. 2 implies that for prediction the central question is "*how much* does X affect Y," which is a question about the slope of the relationship between the two variables, $\partial Y/\partial X$. More generally, we are interested in the shape of this relationship. The answer to these questions cannot be reached with an ANOVA experimental design but rather requires a regression design, whereby X is varied along a continuum that spans its expected range of variability.

Finally, the first and second terms of Eq. 2 relate to a long-standing debate in ecology about the importance of endogenous (density-dependent) factors vs. exogenous factors in controlling the dynamics of ecological systems (Davidson and Andrewartha 1948). Eq. 2 does not provide any a priori expectations about the relative magnitudes of these terms, but it does provide a means of expressing different factors in the same terms (predictive variance), which allows us to make direct comparisons both across and within terms (e.g., relative importance of different drivers).

Parameter Uncertainty

The third term in Eq. 2, $(\partial f/\partial\theta)^2\text{Var}[\bar{\theta}]$, addresses the effects of parameter sensitivity and uncertainty. Unlike the previous two terms, which are closely tied to our understanding of the underlying ecological dynamics, this term captures statistical calibration and the weight of evidence in hand. In other words, do we have enough data to make a confident prediction? To address this question, we need to understand the difference between parameter sensitivity and parameter uncertainty and how each contributes to predictive uncertainty in a forecast. Parameter uncertainty centers around a basic statistical question: how well do we know the mean of θ? Basic statistics tells us that the uncertainty about the parameter mean, Var[$\bar{\theta}$], is largely a question of sample size.

When data are sparse, parameter uncertainty will be large and can often dominate ecological forecasts. This is particularly likely to occur for chronically data-limited problems, such as new invasive species and emerging diseases, where the system being forecast has not been observed before. A Bayesian perspective is particularly valuable in these cases, as it allows the formal incorporation of prior information, for example from meta-analyses (Koricheva et al. 2013, LeBauer et al. 2013) or expert elicitation (Morgan 2014).

The other situation where parameter uncertainty dominates is with models that are over-parameterized (i.e., have too many parameters relative to the amount of data used to calibrate them). Over-parameterization (also known as over-fitting) is typically addressed via some form of model selection, where alternative models are considered that vary in complexity and data are used to choose the most parsimonious fit (Hooten and Hobbs 2014). Most forms of model selection are trying, either explicitly or implicitly, to strike a balance between residual error and parameter uncertainty. Typically, as the number of parameters in a model is increased, the model is able to better fit the data, which causes the residual error to decrease. However, as the number of parameters increases, the parameter uncertainty likewise increases. Therefore, for a fixed amount of data, there is some intermediate model complexity that minimizes the total uncertainty (Gelfand and Ghosh 1998). When working with parametric statistical models (e.g., linear regression), most ecologists are familiar with the problem of over-fitting and are aware of how to avoid it. Where over-fitting is much more problematic is with both data mining and process-based models. The data mining community is likewise keenly aware of the over-fitting problem (Hawkins 2004, Radosavljevic and Anderson 2014), but over-fitting remains very easy to do when working with flexible semi-parametric models. The process-based modeling community, on the other hand, is prone to building models with hundreds of parameters without quantifying parameter uncertainty explicitly, and then failing to propagate parameter uncertainty into forecasts (Dietze et al. 2013). When parameter uncertainty is added to such models, forecasts that span or exceed the full range of biologically plausible outputs are not uncommon.

While forecasts may start out with high parameter uncertainty, classic sampling theory also tells us that Var[$\bar{\theta}$] will asymptotically decline to 0 as the amount of data increases, typically in proportion to $1/\sqrt{n}$. Therefore, provided any forecasting problem takes an iterative approach of incorporating new data as it becomes available, parameter uncertainty will tend to decline with time.

That said, in most models, parameter values tend to be constant in any given model run. Therefore, parameter error behaves like a bias: it does not average out and thus will tend to be more important when aggregating predictions to scale up in space or time. Thankfully, Eq. 2 provides important guidance for how to reduce this uncertainty most efficiently. Specifically, $(\partial f/\partial\theta)^2\mathrm{Var}[\bar{\theta}]$ implies that parameter uncertainty and parameter sensitivity contribute equally to predictive uncertainty. The value of this product, calculated for each parameter, can thus be used to determine the relative contributions of different parameters to the overall parameter uncertainty (LeBauer et al. 2013, Dietze et al. 2014). Once key parameters have been identified, additional data collection and synthesis can be directed specifically at the processes involved. Furthermore, $(\partial f/\partial\theta)^2\mathrm{Var}[\bar{\theta}]$ illustrates that a parameter can be important either because it is highly sensitive, or because it's poorly constrained, and that simply knowing one of these two components (uncertainty or sensitivity) is insufficient to determine which parameters are driving a model. However, once a parameter has been identified as important, the partitioning between uncertainty and sensitivity can be very helpful when prioritizing data collection. Specifically, parameters that are poorly constrained but less sensitive provide a faster return on investment because they are in the initial steep part of $1/\sqrt{n}$. By contrast, parameters that are well constrained (high n) but still important because of high sensitivity are in the region of diminishing returns where further reductions in uncertainty require ever greater investments of time and energy.

Process Error

The final two terms in Eq. 2, $(\partial f/\partial\theta)^2\mathrm{Var}[\alpha] + \mathrm{Var}[\varepsilon]$ represent different aspects of the same phenomenon: process error. Specifically, $\mathrm{Var}[\alpha]$ refers to the unexplained variability in model parameters (differences from site-to-site, year-to-year, etc.) and ε_t refer to errors in the model, not observation error. While the previous sections all considered different aspects of the variability that was explained by the model (internal dynamics, initial conditions, external drivers, and parameters), these two terms capture the unexplained variability. Unlike parameter uncertainty, which declines asymptotically with sampling, neither parameter variability or model error converge asymptotically because they represents variability in the underlying ecological process itself.

Process error is not a monolithic concept but rather subsumes within it a number of different sources of variability, such as model structural uncertainty, heterogeneity, and stochasticity. In this view, stochasticity refers to ecological processes, such as reproduction, mortality, dispersal, disturbance, that involve components that are modeled using random numbers. Like the flip of a coin, stochastic processes are those where the amount of knowledge about the physical system, and computation required to approach the problem deterministically, are so incredibly vast and detailed, and on a spatial and temporal scale so divorced from the process of interest, that there's no conceivable way to distinguishing the problem from true randomness, nor any practical reason to do so.

Model structural uncertainty captures that all models are approximations of reality and no model is perfect. While part of this error comes from computational approximations (e.g., discretizing time), in ecology, most of this error arises from the choice of equations. Most of the equations in ecological models are not deterministic physical laws, but empirical calibrations (e.g., regressions) and variations on a few simple equations for population growth and mass and energy fluxes. In practical terms, partitioning out model structural error often requires working with multiple models and making predictions based on model averaging (Hoeting et al. 1999). While beyond the scope of the current analysis, Eq. 2 could be extended explicitly to a multi-model context, for example using the probability of different functional forms to weight their different sensitivities, or in a nested model by including the probability of including any particular covariate in the exogenous term along with a matching probability for any coefficients in the parameter term.

Even when the equations in a model are reasonable, because the parameters in those models reflect biology rather than physical constants, there can be variability in the parameters themselves, $\mathrm{Var}[\alpha]$. Parameter variability thus reflects the considerable heterogeneity in ecological processes over a wide range of spatial, temporal, and phylogenetic scales. The unexplained differences among ecologically meaningful units, such as individuals, locations, years, and species, are often persistent, meaning that the differences are either permanent or slowly changing relative to the process of interest. For the purposes of prediction, it is important that we quantify this variability, even if we can't explain it, since we want to avoid falsely overconfident forecasts. That said, as we learn more about a system, we can often chip away at this process error and attribute more and more of the observed variability to the deterministic components of our models. Doing so requires that we know the source of variability and that the added complexity involved explicitly results in a net reduction in the total predictive uncertainty (Eq. 2).

Even when we can't explain parameter variability, we can improve our predictions by being able to partition this variability to different sources or scales (e.g., population, watershed, or individual; year; species). Statistically, we do this using random or hierarchical effects (Clark 2005). The amount of variability partitioned to different scales can help identify the processes responsible for such process error and prioritize additional data collection.

Consider two alternative scenarios for a study following 10 sites over the course of 10 years (Fig. 2). In both studies, the total amount of variability is the same, but in the first scenario, 75% of the variability is site-to-site, while in the second scenario, 75% of the variability is year-to-year. Each of these cases would point to very

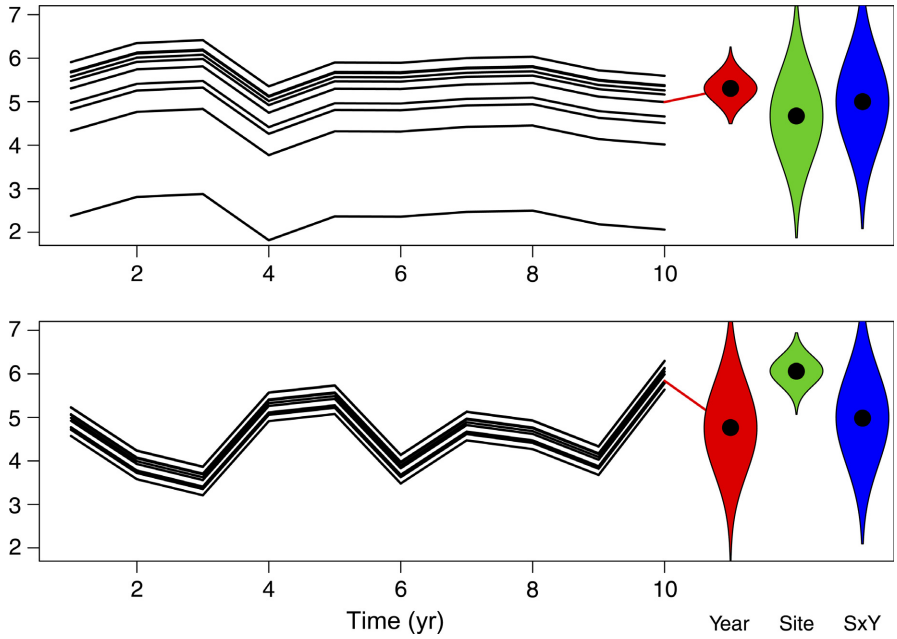

FIG. 2. Impact of random effects on forecasts. Both panels show the trajectories of 10 sites (individual lines) over 10 yr and have the same total variance. In the top panel, 75% of the variability is site-to-site, while in the bottom panel, 75% is year-to-year. On the right hand side, the violin plots under "Year" depict the prediction for a previously known site for the next new year. In this case, the top example (low year-to-year variability) is more predictable. "Site" depicts the prediction for a new site in year 10, and in this case, the bottom example (low site-to-site variability) is more predictable. Finally, S × Y depicts the prediction for a new site in a new year, which is identical for both panels since they have the same total variability. [Color figure can be viewed at wileyonlinelibrary.com]

different sets of potential explanatory variables and would imply very different future research directions. Furthermore, even if we don't know the causal agent behind this variability, how the variability is partitioned has direct impacts on forecasts. For example, a prediction for a known site in the following year would be more constrained under the first scenario (persistent site differences, low year-to-year variability) than under the second scenario (high synchrony among sites; Fig. 2, Year). By contrast, an out-of-sample forecast to a new site in the current year would be more constrained under the second scenario (Fig. 2, Site). Finally, the case of an out-of-sample prediction for a new year produces the same prediction under both scenarios. Overall, the estimation and partitioning of parameter variability represents an important mechanism for using measurements from one unit of measure (e.g., population, watershed, or individual; year; species) to make inferences about another in a way that accounts for both the potential differences and larger uncertainty when doing so.

Returning to our first-principles consideration of uncertainty propagation, it is challenging to determine whether process error and parameter variability will increase, decrease, or stay the same over time. If process errors are independent, they should stay the same through time; this is the case when process error is dominated by memory-less stochastic processes (e.g., background rates of recruitment and mortality). Other stochastic processes, such as disturbance, have memory (e.g., the probability of fire depends on fuel, and thus on time since last fire) so the variance would not be constant, and generally increases with time. At large spatial and long temporal scales, such processes typically converge to a steady-state distribution (e.g., a disturbance regime rather than an individual disturbance event). For most forms of model structural error and parameter variability our a priori expectation is that these process errors are autocorrelated. If the autocorrelation is a simple first-order autoregressive (AR[1]) process, as is commonly assumed in time-series models, then the process error will initially increase but asymptotically approach a steady state, $\mathrm{Var}[\varepsilon]/(1 - |\rho^2)$, where ρ is the correlation coefficient. It is worth noting that the same expectation applies in space (e.g., spatially correlated disturbances), as well as to integrating over driver variability in space and time, and in all cases cause variability at larger scales to average out more slowly. Taken as a whole, and integrating over multiple sources of process error, the arguments above imply that our general expectation is for process error to increase asymptotically to some steady-state, but the rate of increase and asymptote will be problem specific.

When forecasting is viewed as an iterative process, then our expectation is that process error will tend to decline to some non-zero asymptote. Specifically, as we make more measurements through time we expect not just for parameter estimates to converge asymptotically, but for models to be continually refined and improved. As the functional representations and computational

Table 1. Summary of key concepts and predictions.

Concept	Predictions
Endogenous (internal) stability, Y	Grows or declines exponentially, all other terms are linear
	Predictive uncertainty grows without bound or asymptotically
	Determined by classic stability thresholds
	Relative importance increases with larger scales
Exogenous (external) stability, X	Predictability increases when drivers are predictable or Dynamics are insensitive to variation
	Relative importance increases with time
	Covariates useful for prediction may be different from those used for explaining the same process
	Experimental design emphasizes *how much* X affects Y
	Autocorrelation slows how Var[X] averages out when scaling
	Model selection chooses models that are overly complex for prediction
Parameter uncertainty, $\bar{\theta}$	Dominates data-limited problems and over-parameterized models
	Within a forecast, does not necessarily increase or decrease
	With sampling, declines asymptotically to zero
	Does not average out when aggregating in space or time
	Sampling can be targeted though uncertainty analysis
Parameter variability, α, and process error, ε	Encompasses heterogeneity, model structural error, and stochasticity
	Does not decline asymptotically with sampling
	Like X, declines with scale but rate dependent on autocorrelation
	Declines to non-zero asymptote through model improvement
	Partitioning of parameter variability important for extrapolation from one unit of measure to another to account for potential differences and larger uncertainty

approximations in models are improved, we expect the model structural error to decline. Similarly, we also expect to chip away at the unexplained sources of parameter variability, shifting this into the explained endogenous and exogenous components. Nonetheless, models will always be approximations and there will always be something we didn't measure and didn't include in the model. Likewise, many models will always have processes that are best described as stochastic processes. Therefore, some portion of process error will always remain. What is hard to generalize is whether this asymptotic component will be large (e.g., dominated by irreducible stochasticity) or small at any particular scale.

In summary, the framework provided by Eq. 2 generates a number of predictions (Table 1). Some of the individual conclusions derived are already known, however the novelty of this first principles approach isn't the individual expectations, but the unifying framework that links these different expectations and shows how they arise from basic concepts. Collectively, the terms in Eq. 2 determine the rate at which the uncertainty increases, and thus the time horizon over which useful forecasts can be made, what has been called the ecological forecast horizon (Petchey et al. 2015). The relative magnitudes of these terms also tell us about the nature of the ecological forecasting problem. As we saw with meteorology, knowing the nature of the problem is important not just for understanding the theoretical nature of the system, but also for any practical attempt to make predictions. The data we collect, the models we develop, and the statistics we use are all tightly linked to which sources of uncertainty drive the predictability of the system.

While the above discussion has primarily been conceptual, the other strength of this framework is that it can be used quantitatively with real-world data to assess the absolute and relative magnitudes of the different terms. Furthermore, once such terms are computed for real world examples, it provides a framework for a comparative approach. The sections below will first address the application of this framework to real world data, and then lay out the future questions and directions for a comparative approach.

An Example: Net Ecosystem Exchange

To illustrate the application of the analytical concepts above, consider the example of predicting half-hourly Net Ecosystem Exchange (NEE). For purposes of illustration, I fit a simple dynamic linear model to 16 d of data from the Sylvania Wilderness Ameriflux tower and made predictions for the following 16 d based on the NOAA weather forecast. Sylvania is an old-growth northern hardwood–hemlock forest located in the upper peninsula of Michigan (46.242° N, 89.3476° W). Data were downloaded from Ankur Desai's real-time data server for 2016.[2]

The dynamic linear model was fit for using days 151–167 using top-of-tower air temperature (Ta) and photosynthetically active radiation (PAR) as covariates. PAR and Ta had the highest correlations with observed fluxes (−0.51 and −0.22, respectively) among the set of explanatory variables available in both the Ameriflux

[2] http://flux.aos.wisc.edu/twiki/bin/view/Main/LabData

data and weather forecasts. The model was fit in a state-space framework assuming normal error where NEE_o is the observed NEE values, NEE is the latent (unobserved) true value of NEE, τ_o is the observation error, τ_o is the process error, and β are the regression coefficients

$$NEE_{t+1} \sim N(\beta_0 NEE_t + \beta_1 + \beta_2 Ta + \beta_3 Ta + \beta_3 PAR, \tau_p)$$

$$NEE_{o,t} \sim N(NEE_t, \tau_o)$$

Model fits and forecasts were performed in R (version 3.2.2) using the ecoforecastR R package (version 0.1.0; *available online*)[3] and JAGS (Plummer 2010). Uninformative normal priors (mean 0, precision 0.001) were assumed for the β_s and uninformative Gamma(0.1,0.1) priors were assumed for the precisions.

Fig. 3 shows that this simple linear model with two covariates can capture the diurnal cycle of the flux data, though with considerable uncertainty during periods of missing data. Following the atmospheric sign convention for NEE data (negative values indicate uptake), we observe large uptakes of carbon during the day and moderate positive losses at night, with this cycle primarily driven by PAR (Table 2). The slope of the internal stability term is large (0.80), though stable and different from 0, indicating significant system memory in NEE. Observation error had a much higher precision (lower variance), than the process error, suggesting that much of the observed NEE variability represents real variation the model is not capturing. This conclusion is somewhat at odds with the flux literature, which acknowledges substantial observation errors in fluxes. A more in-depth analysis could be improved, for example by using a more precise observation error with informative priors, such as the asymmetric heteroscedastic Laplace (Hollinger and Richardson 2005, Richardson et al. 2006). Nonetheless, this simple proof-of-concept model represents a useful baseline for evaluating any nonlinear statistical or process-based model.

NEE was forecast for the next 16 d using the fit model and weather forecast data from NOAA's Global Ensemble Forecast System (GEFS). Driver uncertainty was captured by the spread of the 21 members of the ensemble forecast. Air temperature data was downscaled to 30 min by fitting a spline through the six-hourly forecast product while PAR was downscaled based on solar geometry. Confidence in the drivers was high for the first few days, but by the end driver forecasts show large uncertainties and little day-to-day variability (Fig. 4).

Forecasts were made using a 500-member ensemble, with each ensemble member sampling the set of drivers, parameters, and initial conditions with replacement. Forecasts had a clear diurnal cycle, but little day-to-day variability and high uncertainty. While the forecast uncertainty appears to be at a steady state, it contains a slowly increasing trend superimposed on a strong diurnal cycle (higher during the day).

[3] https://github.com/EcoForecast/ecoforecastR

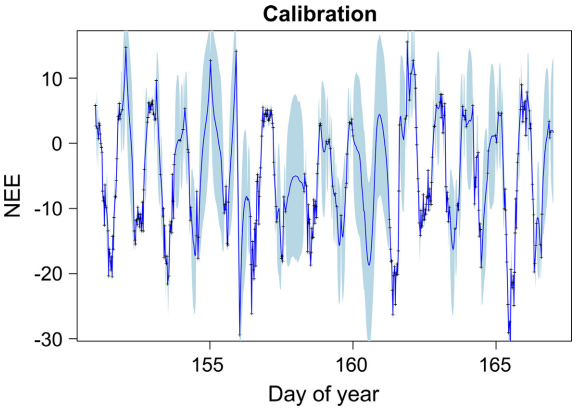

FIG. 3. Dynamical linear model median (solid line) and 95% CI (shaded area) fit to observed flux data (+). NEE, net ecosystem exchange. [Color figure can be viewed at wileyonlinelibrary.com]

TABLE 2. Posterior mean, standard deviation, and 95% CI for fit parameters (βs) and precisions (τs).

Parameter	Mean	SD	2.5%	97.5%
β_0 (NEE)	0.805	0.0288	0.745	0.858
β_1	0.962	0.575	−0.175	2.0725
β_2 (Ta)	−0.0552	0.0422	−0.1367	0.0284
β_3 (PAR)	−0.00240	0.000454	−0.00333	−0.00154
τ_p	0.0809	0.00692	0.0683	0.0954
τ_o	7.271	7.028	0.961	26.461

Note: See *An example: Net ecosystem exchange* for parameter definitions.

The uncertainties in model predictions were partitioned using two different approaches. The first used the analytical approximation discussed previously. The second was generated by running a series of forecasts that sequentially introduced different sources of uncertainty (parameter, driver, initial condition, process), and estimating the effect of each as the difference in variance between pairs of scenarios. The results of these two approaches were qualitatively similar but differ slightly because the second approach is sequential and accounts for nonlinear interactions. Results are summarized by the mean Table 3 and time series Fig. 5 of the relative proportion of variance attributable to each term. Overall, process error dominates the forecast from early on and remains relatively constant. Driver uncertainty shows an increasing trend, while parameter uncertainty shows a diurnal cycle but no trend and initial condition uncertainty decays quickly.

The uncertainty partitioning in cumulative NEE over the entire forecast is noticeably different than the average partitioning for each 30-min period (Fig. 5, Table 3). The contribution of process error is considerably smaller, as much of this variability is random and averages out over

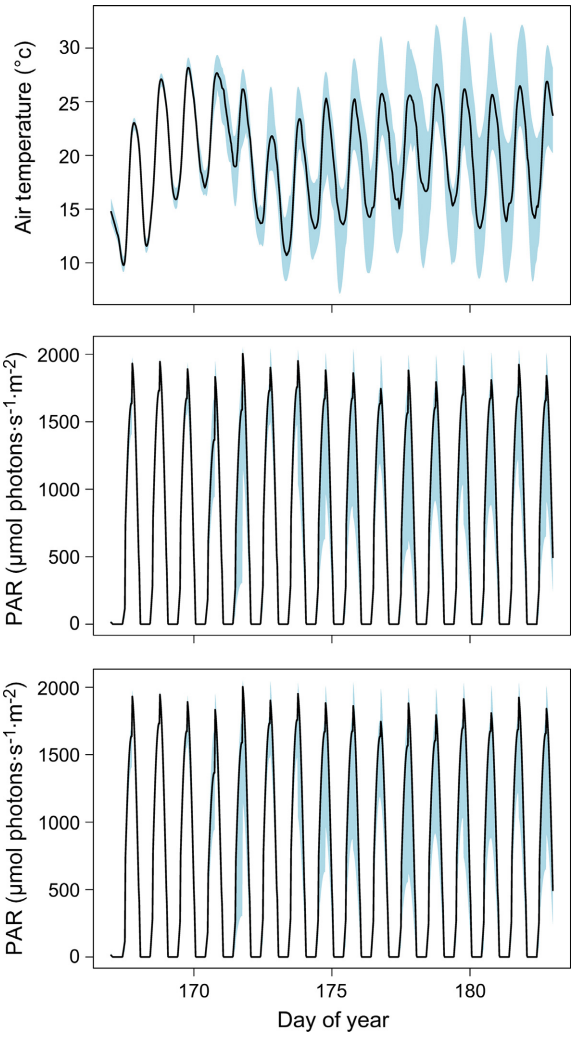

FIG. 4. Input weather forecast drivers and predicted NEE. Driver data show a noticeable pattern of increasing uncertainty with time. Forecast shows a clear diurnal pattern, but little variation in day-to-day predictions. PAR, photosynthetically active radiation. [Color figure can be viewed at wileyonlinelibrary.com]

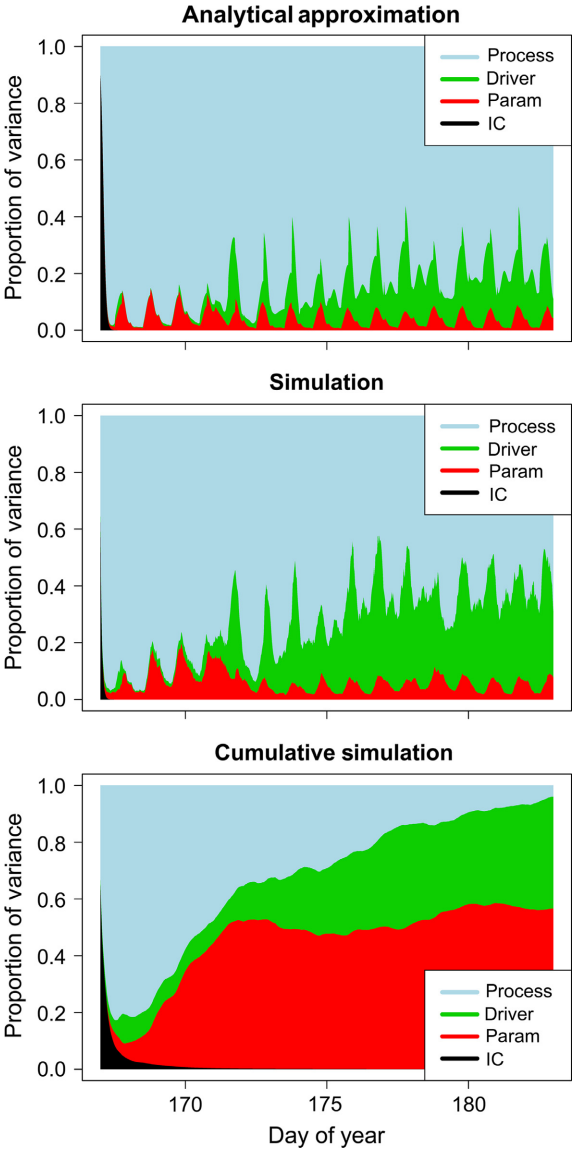

FIG. 5. Partitioning of forecast uncertainties by source using both analytical approximation and simulation-based approaches. Both approaches agree that uncertainty is dominated by process error, initial condition uncertainty decays to negligible rapidly, parameter uncertainty shows a consistent diurnal cycle but no trend, and driver uncertainty increases with time. By contrast, the uncertainty in the forecast cumulative NEE is dominated by parameter (Param) and driver (Driver) uncertainties, whose relative contributions start off small but increase in time, with parameter uncertainty leveling off after the first 5 d. Process error (Process) and initial condition uncertainty (IC) start large but decay over time, with initial condition error decaying rapidly. [Color figure can be viewed at wileyonlinelibrary.com]

TABLE 3. Mean partitioning of forecast uncertainty in instantaneous NEE according to both analytical and simulation approaches.

Approach	IC	Param	Driver	Process
Analytical	0.02883	0.042	0.115	0.814
Simulation	0.00269	0.061	0.211	0.726
Cumul. sim.	0.00033	0.566	0.394	0.040

Notes: Both approaches are in general agreement about the relative importance of the different sources of uncertainty. By contrast, the uncertainty in the cumulative NEE (Cumul. Sim.) shows that the contribution of process error (Process) is much lower, as many errors average out, and parameter uncertainty (Param) contributes much more to the overall uncertainty. Driver, driver uncertainty, IC, initial condition uncertainty.

time. As before contribution of driver uncertainty increases steadily over time, but for cumulative NEE this constitutes a larger fraction of overall uncertainty. Similarly, the contribution of parameter uncertainty is much larger and more persistent through time, contributing to

the forecast uncertainty *systematically*, and thus increases in relative importance for cumulative NEE. Overall, the contrast between instantaneous and cumulative NEE illustrates the impacts of aggregation across scale in understanding the contributions of different sources of uncertainty. A similar phenomena is expected when moving up spatial or taxonomic/functional scales (e.g., aggregating species into Families or functional types).

The conclusions of this analysis are conditional on the current model structure. A model with additional state variables (soil carbon, biomass, leaf area index [LAI], etc.) would include longer-term internal stability than NEE (days to centuries), which would increase the overall contribution of intrinsic stability to NEE dynamics over longer time scales. Next, increasing the number of covariates would reduce the process error but increase the parameter and driver uncertainties. Using a more complex, nonlinear model would similarly decrease process error but increase parameter uncertainty. Whether the overall predictive uncertainty increased or decreased in these cases would be a model selection question. Traditional model selection criteria (AIC, DIC, etc.) only capture the trade-off between process and parameter uncertainties, and would miss the impact of driver uncertainty. Therefore, from a predictive perspective, traditional model selection tends to select for models that are overly complex. Predictive uncertainty can be used explicitly as a model selection criteria, but results will depend on the time scale considered.

Questions and Directions

The strength of a first-principles approach (Eq. 2) is that it provides a useful general framework. The sections above illustrate how this framework can be used both qualitatively, to structure our understanding of predictability and long-standing conceptual debates in ecology, and quantitatively, to make explicit predictions and recommendations and to enable comparisons of the absolute and relative contributions of different sources of uncertainty. The example above illustrated a "within problem" application, using a simple dynamic linear model to assess the predictability of NEE at a single site. This analysis found initial condition uncertainty dissipated rapidly, parameter and driver uncertainties both contributed moderately, and the forecast was dominated by process error.

What is currently lacking is a systematic, quantitative understanding of how and why the relative contributions of these factors vary from problem to problem. A comparative approach, applying this framework to a wide variety of different ecological processes, would get at the heart of many key questions in ecology, help to determine what drives ecological dynamics, and allow us to look for patterns and generalities that span all subdisciplines. For example, the framework provides a common, cross-system approach to address long-standing debates about stability vs. drift, quantify the relative importance of endogenous (density dependent) factors vs. exogenous drivers, and scale predictions in space and time. A comparative synthesis across different systems and problems would likewise have immediate applications, allowing us to identify and target the dominant sources of variability for a particular class of problem, reducing the unproductive, paralyzing fear that we need to measure everything everywhere. Below I lay out some initial questions and hypotheses that emerge with a comparative approach.

The first question is "which sources of uncertainty drive ecological predictability for different systems?" I hypothesize (H1) that a partitioning of predictive uncertainties for different ecological processes will exhibit common patterns after accounting for sample size. The null hypothesis would be that the relative proportion of uncertainty attributable to these different sources varies randomly. Alternative hypotheses would be that certain sources of uncertainty are consistently more or less important, or that the relative importance of different uncertainties varies predictably with system attributes (terrestrial vs. aquatic; population, community, or ecosystem; etc.). Answering this question will improve our understanding of ecological systems, refine how we direct future research effort, and directly impact our ability to make predictions.

Next, when looking at any particular process, there is a general expectation that most ecological processes are autocorrelated in both space and time. Knowing the magnitude of that autocorrelation is necessary to make predictions; spatial and temporal autocorrelation thus affect the predictability of ecological processes. Furthermore, ecologists frequently attempt to use information from one of these scales to make inferences about the other (e.g., space for time substitution). However, when taking a comparative approach, we lack a general understanding of how uncertainty is partitioned across space and time for different ecological problems.

I hypothesize (H2) that across different ecological processes, the parameters describing autocorrelation in space and time (e.g., ρ, range) will themselves be correlated. While it is undoubtedly true that different factors drive ecological dynamics at different spatial and temporal scales (e.g., local competition vs. regional climate), it is also true that different processes vary considerably in their predictability. As noted in Table 1, highly predictable processes are likely insensitive to drivers that vary unpredictably over short and fast scales (e.g., buffering against variability) and are instead responding to drivers that are predictable, which can occur when drivers vary on both slow and large scales. Slow, predictable variation implies a high autocorrelation in drivers that translates to high autocorrelation in the ecological process. H2 posits that the underlying patterns in autocorrelation across processes explain more of the variability in predictability than the identity of the case-specific drivers. Furthermore, it predicts that we are not going to find ecological processes with high autocorrelation in time but high variability in space, and vice versa. This

hypothesis not only addresses a fundamental theoretical question, but also directly informs the practical forecasting problem of localization: how close do measurements need to be to a particular location to provide useful information? Finally, H1 and H2 together imply a further hypothesis about the limit of predictability, which we define here as the horizon in space or time at which a forecast decays to an appropriate null model: H3, the limit of predictability in space and time will scale with the autocorrelation parameters. Put more simply, H3 tests to what extent we can use pattern to make inferences about the predictability of a process.

Overall, these three hypotheses just scratch the surface at the range of questions raised by considering a comparative approach to assessing ecological predictability. The framework presented here is by no means a general unified theory of ecology, but it provides an important opportunity for assessing whether there are common patterns to ecological predictability. It is also a critical first step towards making our science more focused on prediction. A focus on prediction not only makes our science more relevant to society, at a time when the need for data-informed policy and management has never been greater, but it is also a means for accelerating our scientific understanding. Forecasting is ultimately a direct expression of the scientific method, whereby we make quantitative, testable predictions about the future and then use new observations to refine our understanding, update hypotheses, and iteratively improve our predictions. An iterative, near-term approach to prediction (Dietze 2017, Dietze et al. 2017) allows us to repeat this cycle frequently, and accelerate the pace at which we confront predictions with new data.

Acknowledgments

Thanks to Ankur Desai for sharing the Sylvania Wilderness Ameriflux tower data and for providing useful feedback on this paper. This paper and the concepts within also benefited from feedback from Peter Adler, Colin Averill, Shannon LaDeau, Joshua Mantooth, Jason McLachlan, Ann Raiho, Dave Schimel, Alexey Shiklomanov, Jennifer Talbot, Kathleen Weathers, and two anonymous reviewers. This project was supported by NSF MSB # 1638577.

Literature Cited

Bjørnstad, O. N. 2015. Nonlinearity and chaos in ecological dynamics revisited. Proceedings of the National Academy of Sciences of the USA 112:6252–6253. 201507708.

Bjørnstad, O. N., B. F. Finkenstadt, and B. T. Grenfell. 2002. Dynamics of measles epidemics: estimating scaling of transmission rates using a time series SIR model. Ecological Monographs 72:169–184.

Clark, J. S. 2005. Why environmental scientists are becoming Bayesians. Ecology Letters 8:2–14.

Clark, J. S., and J. S. McLachlan. 2003. Stability of forest biodiversity. Nature 423:635–638.

Clark, J. S., et al. 2001. Ecological forecasts: an emerging imperative. Science 293:657–660.

Clements, F. E. 1936. Nature and structure of the climax. Journal of Ecology 24:252–284.

Cohen, Y., and J. Pastor. 1991. The responses of a forest model to serial correlations of global warming. Ecology 72:1161–1165.

Davidson, J., and H. G. Andrewartha. 1948. The influence of rainfall, evaporation and atmospheric temperature on fluctuations in the size of a natural population of *Thrips imaginis* (Thysanoptera). Journal of Animal Ecology 17:200–222.

Dietze, M. C. 2017. Ecological forecasting. Princeton University Press, Princeton, New Jersey, USA.

Dietze, M. C., D. S. LeBauer, and R. Kooper. 2013. On improving the communication between models and data. Plant, Cell & Environment 36:1575–1585.

Dietze, M. C., et al. 2014. A quantitative assessment of a terrestrial biosphere model's data needs across North American biomes. Journal of Geophysical Research: Biogeosciences 119:1–15.

Dietze, M. C., et al. 2017. Iterative ecological forecasting: needs, opportunities, and challenges. National Ecological Observatory Network, Boulder, Colorado, USA.

Ellner, S., and P. Turchin. 1995. Chaos in a noisy world: new methods and evidence from time-series analysis. American Naturalist 145:343–375.

Ellner, S., and P. Turchin. 2005. When can noise induce chaos and why does it matter: a critique. Oikos 111:620–631.

Evans, M. R., et al. 2013. Predictive systems ecology. Proc. R. Sci. B, 280, 20131452

Evensen, G. 2009. Data assimilation: the ensemble Kalman filter. Second edition. Springer, Berlin.

Ferrari, M. J., R. F. Grais, N. Bharti, A. J. K. Conlan, O. N. Bjørnstad, L. J. Wolfson, P. J. Guerin, A. Djibo, and B. T. Grenfell. 2008. The dynamics of measles in sub-Saharan Africa. Nature 451:679–684.

Gelfand, A. E., and S. K. Ghosh. 1998. Model choice: a minimum posterior predictive loss approach. Biometrika 85:1–11.

Gleason, H. A. 1926. The individualistic concept of the plant association. Bulletin of the Torrey Botanical Club 53:7–26.

Hawkins, D. M. 2004. The problem of overfitting. Journal of Chemical Information and Computer Sciences 44:1–12.

Hoeting, J.A., D. Madigan, A.E. Raftery and C.T. Volinsky. 1999. Bayesian Model Averaging: A Tutorial. Stat. Sci. 14:382–417.

Hollinger, D. Y., and A. D. Richardson. 2005. Uncertainty in eddy covariance measurements and its application to physiological models. Tree Physiology 25:873–885.

Hooten, M. B., and N. T. Hobbs. 2014. A guide to Bayesian model selection for ecologists. Ecological Monographs 85(1): 3–28.

Houlahan, J. E., S. T. McKinney, T. M. Anderson, and B. J. McGill. 2017. The priority of prediction in ecological understanding. Oikos 126:1–7.

Hubbell, S. P. 2001. The unified neutral theory of biodiversity and biogeography. Princeton University Press, Princeton, New Jersey, USA.

Kalnay, E. 2002. Atmospheric modeling, data assimilation and predictability. Cambridge University Press, Cambridge, England.

Koricheva, J., J. Gurevitch, and K. Mengersen. 2013. Handbook of meta-analysis in ecology and evolution. Princeton University Press, Princeton, New Jersey, USA.

LeBauer, D., D. Wang, K. T. Richter, C. C. Davidson, and M. C. Dietze. 2013. Facilitating feedbacks between field measurements and ecosystem models. Ecological Monographs 83:133–154.

Levin, S. A. 1992. The problem of pattern and scale in ecology. Ecology 73:1943–1967.

Lewis, J. M., S. Lakshmivarahan, and S. K. Dhall. 2006. Dynamic data assimilation: a least squares approach. Cambridge University Press, New York, New York, USA.

Lorenz, E. N. 1963. Deterministic nonperiodic flow. J. Atmos. Sci. 20(2):130–141.

Lotka, A. J. 1910. Contribution to the theory of periodic reaction. Journal of Physical Chemistry 14:271–274.

Luo, Y., T. F. Keenan, and M. Smith. 2015. Predictability of the terrestrial carbon cycle. Global Change Biology 21:1737–1751.

Morgan, M. G. 2014. Use (and abuse) of expert elicitation in support of decision making for public policy. Proceedings of the National Academy of Sciences of the USA 111:7176–7184.

Mouquet, N., et al. 2015. REVIEW: predictive ecology in a changing world. Journal of Applied Ecology 52:1293–1310.

Norden, N., et al. 2015. Successional dynamics in Neotropical forests are as uncertain as they are predictable. Proceedings of the National Academy of Sciences of the USA 112:8013–8018. 201500403.

Pastor, J., and Y. Cohen. 1997. Nitrogen cycling and the control of chaos in a boreal forest model. Pages 304–319 *in* K. Judd, A. Mees, K. L. Teo, and T. L. Vincent, editors. Control and chaos. Birkhäuser, Boston, Massachusetts, USA.

Petchey, O. L., et al. 2015. The ecological forecast horizon, and examples of its uses and determinants. Ecology Letters 18: 597–611.

Plummer, M. 2010. JAGS version 2.2.0 user manual:0–39. https://sourceforge.net/projects/mcmc-jags/files/Manuals/4.x/jags_user_manual.pdf

Radosavljevic, A., and R. P. Anderson. 2014. Making better Maxent models of species distributions: complexity, overfitting and evaluation. Journal of Biogeography 41:629–643.

Richardson, A. D., et al. 2006. A multi-site analysis of random error in tower-based measurements of carbon and energy fluxes. Agricultural and Forest Meteorology 136:1–18.

Schimel, D. S., G. Churkina, B. H. Braswell, and J. Trenbath. 2005. Rememberance of weather past: ecosystem responses to climatic variability. Pages 350–368 *in* I. Baldwin, editor. A history of atmospheric CO_2 and its effects on plants, animals, and ecosystems. Springer, New York, New York, USA.

Turner, M., W. Romme, and R. Gardner. 1993. A revised concept of landscape equilibrium: disturbance and stability on scaled landscapes. Landscape Ecology 8:213–227.

Volterra, V. 1926. Variazioni e fluttuazioni del numero d'individui in specie animali conviventi. Memorie della R. Accademia dei Lincei 2:31–113.