

The mathematics of extinction across scales: from populations to the biosphere

The sixth mass extinction poses an unparalleled quantitative challenge to conservation biologists. Mathematicians and ecologists alike face the problem of developing models that can scale predictions of extinction rates from populations to the level of a species, or even to an entire ecosystem. We review some of the most basic stochastic and analytical methods of calculating extinction risk at different scales, including population viability analysis, stochastic metapopulation occupancy models, and the species area relationship. We also consider two major extensions of theory: the possibility of evolutionary rescue from extinction in a changing environment, and the posthumous assignment of an extinction date from sighting records. In the case of the latter, we provide an example using data on Spix's macaw (*Cyanopsitta spixii*), the "rarest bird in the world," to demonstrate the challenges associated with extinction date research.

The Mathematics of Extinction Across Scales: From Populations to the Biosphere

Colin J. Carlson, Kevin R. Burgio, Tad A. Dallas, & Wayne M. Getz

Abstract

The sixth mass extinction poses an unparalleled quantitative challenge to conservation biologists. Mathematicians and ecologists alike face the problem of developing models that can scale predictions of extinction rates from populations to the level of a species, or even to an entire ecosystem. We review some of the most basic stochastic and analytical methods of calculating extinction risk at different scales, including population viability analysis, stochastic metapopulation occupancy models, and the species area relationship. We also consider two major extensions of theory: the possibility of evolutionary rescue from extinction in a changing environment, and the posthumous assignment of an extinction date from sighting records. In the case of the latter, we provide an example using data on Spix's macaw (*Cyanopsitta spixii*), the "rarest bird in the world," to demonstrate the challenges associated with extinction date research.

Keywords: Sixth mass extinction, species area relationship, sighting records, population viability analysis, mean time to extinction.

MSC 2010 Codes: 92B05, 92D40

Colin J. Carlson

The National Socio-environmental Synthesis Center (SESYNC)
1 Park Place, Suite 300, Annapolis MD 21401, United States
e-mail: ccarlson@sesync.org

Kevin R. Burgio

Department of Ecology & Evolutionary Biology, University of Connecticut,
75 N. Eagleville Road, Unit 3043, Storrs CT 06269 United States
e-mail: kevin.burgio@uconn.edu

Tad A. Dallas

Department of Environmental Science and Policy, University of California Davis,
One Shields Ave. Davis, CA 95616, United States
e-mail: tdallas@ucdavis.edu

Wayne M. Getz

Department of Environmental Science, Policy and Management, University of California Berkeley
130 Mulford Hall #3114 Berkeley, CA 94720-3114, United States
e-mail: wgetz@berkeley.edu

It's easy to think that as a result of the extinction of the dodo, we are now sadder and wiser, but there's a lot of evidence to suggest that we are merely sadder and better informed.

– Douglas Adams, *Last Chance to See*

1 Introduction

Most species, like most living organisms on Earth, have a finite lifespan. From the origin of a species onward, every species changes and adapts to its environment. Some species exist longer than others, but all eventually face extinction (or, are replaced by their descendants through evolution). Currently, there are approximately 8.7 million eukaryote species alone. But in the history of Earth, it is estimated that there have been a daunting 4 billion species altogether, and at least 99 percent of them are now gone.¹

How long can a species exist? Of the species currently on Earth, some are deeply embedded in the geological record and have changed very little over the span of million years, such as coelacanths or ginkgo trees. Most species persist for a few millions of years or more, and in periods of environmental stability, extinctions typically occur at a low and steady baseline rate. But at various points in the history of the Earth, extinction rates have suddenly accelerated for brief and eventful periods that biologists term *mass extinction events*. In 1982, based on the marine fossil record, David Raup and Jack Sepkoski suggested that five of these mass extinctions have occurred over the past half billion years.² In all five, more than half of all contemporary species disappeared,³ each sufficiently drastic to be considered the end of a geological era: the Ordovician 444 million years ago (*mya*), Devonian 375 *mya*, Permian 251 *mya*, Triassic 200 *mya* and Cretaceous 66 *mya*.

But in recent years, ecologists have reached the consensus that the biosphere is currently experiencing, or at the very least entering, the sixth mass extinction.⁴ Unlike the previous five, which were caused by planetary catastrophes and other changes in the abiotic environment, the sixth mass extinction is the undeniable product of human activities. While anthropogenic climate change is one of the most significant contributors, a number of other factors have recently exacerbated extinction rates, including habitat loss and fragmentation, biological invasions, urbanization, over-harvesting, pollution, pests, and emerging diseases.

¹ Camilo Mora et al., “How many species are there on Earth and in the ocean?,” *PLoS Biol* 9, no. 8 (2011): e1001127.

² David M Raup and J John Sepkoski Jr, “Mass extinctions in the marine fossil record,” *Science* 215, no. 4539 (1982): 1501–1503.

³ Michael L McKinney and Julie L Lockwood, “Biotic homogenization: a few winners replacing many losers in the next mass extinction,” *Trends in ecology & evolution* 14, no. 11 (1999): 450–453.

⁴ Elizabeth Kolbert, *The sixth extinction: An unnatural history* (A&C Black, 2014).

How does the sixth mass extinction scale up against the last five? The number of extinctions alone is an unhelpful metric, as species richness changes over time. A more convenient unit of measurement commonly used by scientists is the number of *extinctions per millions of species-years* (E/MSY). From a landmark study by Gerardo Ceballos and colleagues, we know that in the geological record, vertebrates normally go extinct at a rate of 2 E/MSY in the periods in-between mass extinctions. But since 1900, that rate is an astounding 53 times higher.⁵ One study has suggested that the sixth mass extinction is comparable to other mass extinctions in E/MSY rates, meaning that with enough time, the geological definition of a mass extinction (three quarters extinction) could be achieved in hundreds to thousands of years.⁶ Or, to consider another framing : a 1970 study estimated that at a baseline, one species goes extinct per year,⁷ while just a decade later that estimate was revised to one species per hour.⁸ Plants, insects, and even micro-organisms all face similarly catastrophic threats; and these across-the-board losses of biodiversity pose a threat to human survival that some argue could even threaten our own species with extinction.

The crisis of extinction is, for scientists, a crisis of prediction. While extinction is a natural part of ecosystem processes and of the history of the planet, the job of conservation biologists is to protect species that would otherwise be brought to an untimely and avoidable end. To do that, conservationists must sort and prioritize the 8.7 million eukaryotes (and, even, some prokaryotes) to assess which species face the greatest threat—and which can, and cannot, be saved by human intervention. Assessment is easiest at the finest scales: by marking and tracking all the individuals in a region, a population ecologist can make a statistically-informed estimate of the probability of imminent extinction. But above the population level, assessment is much more challenging, requiring sophisticated (and complicated) meta-population models that are typically data-intensive. If a species is rare enough and the data are “noisy,” its extinction may seem uncertain even after the fact; but mathematical models can help assign a probability to the rediscovery of a species once thought extinct, and resolve when (and even why) a species has disappeared long after it is gone. Above the level of a single species, measuring extinction is an altogether different problem, requiring a different type of model to explain how biodiversity arises and is maintained over time.

Each of these modeling approaches represents a different aspect of a connected problem, and we deal with each in turn in this chapter. We begin by discussing the basic mechanics of extinction as a demographic process at the population scale, including population viability analysis, with a case study on evolutionary rescue processes. We progress up to the metapopulation scale, including patch occupancy models and island biogeography. At the species scale, we dive deeper into the issue

⁵ Gerardo Ceballos et al., “Accelerated modern human-induced species losses: Entering the sixth mass extinction,” *Science advances* 1, no. 5 (2015): e1400253.

⁶ Anthony D Barnosky et al., “Has the Earth’s sixth mass extinction already arrived?,” *Nature* 471, no. 7336 (2011): 51–57.

⁷ R Levins, “Extinction,” *Lectures on Mathematics in the Life Sciences* 2 (1970): 77–107.

⁸ Norman Myers, *The sinking ark* (Pergamon Press, Oxford, 1979).

of evolutionary rescue, including the potential for plasticity to buffer species from extinction in a changing environment. Expanding at the species level, we discuss the recently-growing literature on using the sighting record to determine the odds that species are extinct, with a handful of case studies including Spix's macaw and the ivory-billed woodpecker. Finally, we discuss how extinction scales up to the community level, and how extinction rates are inferred from habitat loss using macroecological theory. The models we present are seminal and well-known, but extinction risk modeling is a dynamic and rapidly-growing field. Consequently, these models only present a handful of the many different approaches that link different temporal and spatial scales of extinction together.

2 The Population Scale

Even though many make a terminological distinction between *extinction* (the loss of a species) and *extirpation* (the eradication of a population), extinction is still fundamentally a process that begins at the population scale. With the exception of sudden, unexpected catastrophes, extinction at the population scale is almost always the product of either a declining population or of stochastic variations in an already-small population, both of which follow mathematical rules that can be used to quantify extinction risk. Perhaps the most significant body of theory about population extinction deals with the estimation of a population's *mean time to extinction* (MTE, typically T_E in mathematical notation), an important quantity to both theoretical ecologists and conservation practitioners. For both theoretical and applied approaches to extinction, understanding the uncertainty around T_E requires an understanding of the shape of the extinction time distribution, including developing and testing demographic theory that accurately captures both the central tendencies⁹ and the long tail¹⁰ of empirical extinction times. We begin by reviewing some of the basic population-scale approaches that scale up to ecosystem-level theory of extinction.

2.1 Stochasticity and the Timing of Extinction

In the most basic terms, a population declining at a steady rate will eventually become extinct; the simplest deterministic equation governing the size of a population N , as it grows over time t (generally measure in units of either years or generations) is given by

⁹ John M Drake, "Extinction times in experimental populations," *Ecology* 87, no. 9 (2006): 2215–2220.

¹⁰ John M Drake, "Tail probabilities of extinction time in a large number of experimental populations," *Ecology* 95, no. 5 (2014): 1119–1126.

$$\frac{dN}{dt} = rN \quad (1)$$

where if r is positive the population is growing, while if r is negative, the population heads rapidly towards extinction. A slightly more complex model that captures the phenomenological capping of the growth of a population at a population ceiling termed K is:

$$\frac{dN}{dt} = \begin{cases} rN & \text{if } 1 < N < K \\ 0 & \text{if } N = K \end{cases} \quad (2)$$

While K is often called a carrying capacity, this is perhaps misleading, as in this context it only introduces density dependence when $N = K$, and not before. Eqns. 1 and 2 both imply that if $r < 0$, $\ln(N)$ declines linearly with slope r . For shrinking populations (i.e., $r < 0$) these equations imply that the mean time to extinction (T_E) can be derived analytically as the amount of time before the population reaches one individual (i.e. $N = 1$ at $t = T_E$):

$$T_E(N_0) = -\ln(N_0)/r \quad (3)$$

Consequently, for a given population with a fixed r the maximum achievable extinction time given a starting stable population size would be

$$\max(T_E) = -\ln(K)/r \quad (4)$$

However, in this model if $r > 1$, the population never reaches extinction and simply grows forever.

Deterministic models only tell a part of the story. In the history of conservation biology, two paradigms emerged that separately explain the process of population extinctions. The *declining population paradigm* explains that populations shrink and vanish due to a combination of internal and external failures, and suggests that the key to conserving populations is to identify and prevent those failures. In contrast, the *small population paradigm* is rooted in ideas of stochasticity, suggesting that even without factors like environmental degradation or disease, smaller, more fragmented populations simply face higher extinction risk due to stochastic population processes.¹¹ For one thing, stochasticity produces populations with log-normally distributed sizes (i.e. most populations are comparatively small relative to a few larger ones). The underlying reason for this can be traced back to Jensen's inequality, which suggests the expected value of a convex function applied to a random variable x is greater than, or equal to, that function applied to the expected value of the random variable (below, $E[\cdot]$ is the expectation operator):

$$E[f(x)] \geq f(E[x])$$

¹¹ Graeme Caughley, "Directions in conservation biology," *Journal of animal ecology*, 1994, 215–244.

Applied to stochastic population growth, if r is stochastic, the expectation of r will always be greater than the expected real growth rate of the population:¹²

$$E[r] > E[(N_t/N_0)^{1/t}]$$

Iterating these lower growth rates over an infinite amount of time, populations that are growing randomly with $\bar{r} \leq 1$ (i.e. less than exponential growth) all tend eventually to extinction.

In general, r can be decomposed into two component processes; births and deaths. In their foundational work on the ecology of invasion and extinction—*The Theory of Island Biogeography*—Robert MacArthur and E.O. Wilson proposed a simple model with discrete per-capita birth and death rates, λ and μ respectively. With $\lambda + \mu$ changes expected per time step, the estimated time until a single change (birth or death) is given $1/(\lambda + \mu)$. Thus the time to extinction for a population of x individuals, $T_E(x)$ can be intuitively understood (with a more detailed derivation in MacArthur and Wilson) as: i) the expected time for one change to occur (birth or death); plus ii) the probability the change is a birth (i.e., $\frac{\lambda}{\lambda + \mu}$) multiplied by the time to extinction if the population is of size $x - 1$; plus iii) the probability the change is a death (i.e., $\frac{\mu}{\lambda + \mu}$) multiplied by the time to extinction if the population is of size $x + 1$.

This reasoning produces the relationship:

$$T_E(x) = \frac{1}{\lambda + \mu} + \frac{\lambda}{\lambda + \mu} T_E(x + 1) + \frac{\mu}{\lambda + \mu} T_E(x - 1)$$

This simple but elegant relationship can be used to produce an expression for $T_E(K)$ using the method of induction; in particular, $T_E(x)$ can be expressed as a function of $T_E(1)$, noting that $T_E(0) = 0$. To do this, MacArthur and Wilson add a population ceiling K as before, and consider two cases of density dependence. If births are density dependent, then (using the notation $\tilde{\lambda}(x)$ to distinguish between the function $\tilde{\lambda}$ and the constant λ and similarly for μ)

$$\tilde{\lambda}(x) = \begin{cases} \lambda x & \text{if } x < K \\ 0 & \text{if } x \geq K \end{cases}$$

$$\tilde{\mu}(x) = \mu x$$

and (through and inductive procedure not shown here) the time to extinction is

$$T_E(K) = \frac{\lambda}{\lambda - \mu} T_E(1) + \frac{\lambda}{\mu(K + 1)(\lambda - \mu)} - \frac{1}{\lambda - \mu} \sum_{i=1}^K \frac{1}{i}$$

where

¹² Mark S Boyce, "Population growth with stochastic fluctuations in the life table," *Theoretical Population Biology* 12, no. 3 (1977): 366–373.

$$T_E(K) = \frac{1}{\mu}(K+1)$$

and

$$T_E(1) = \sum_{i=1}^K \left(\frac{\lambda}{\mu}\right)^i \frac{1}{i\lambda} + \left(\frac{\lambda}{\mu}\right)^K \frac{1}{\mu(K+1)}$$

In contrast, if and when deaths are density dependent,

$$\tilde{\lambda}(x) = \lambda x$$

$$\tilde{\mu}(x) = \begin{cases} \mu x & \text{if } x < K \\ 0 & \text{if } x \geq K \end{cases}$$

and the time to extinction is

$$T_E(K) = \frac{\lambda}{\lambda - \mu} T_E(1) - \frac{1}{\lambda - \mu} \sum_{i=1}^K \frac{1}{i}$$

In this scenario, $T_E(K) = T_E(K+1)$; with some induction (not shown here), $T_E(1)$ can also be expressed as

$$T_E(1) = \sum_{i=1}^K \left(\frac{\lambda}{\mu}\right)^i \frac{1}{i\lambda}$$

This provides an explicit method for calculating $T_E(K)$, the maximum achievable time to extinction with these rates. MacArthur and Wilson made a handful of key observations about the behavior of these functions as they relate both to island biogeography and to the population process of extinction. First, $T_E(1)$ can be surprisingly large if $\lambda > \mu$, meaning that a net tendency for growth has incredibly long times before extinction, even with stochasticity. Second, if populations start with a single propagule (as their work is framed in the context of island colonists), roughly μ/λ go extinct almost immediately while roughly $(\lambda - \mu)/\lambda$ grow to K and take $T_E(K)$ years to go extinct. (This means that even though density dependence is not introduced until $N = K$, the effects of the population ceiling are still emergent on the dynamics of the whole system.) Third, “established populations” ($N = K$) have a readily calculated extinction time:

$$T_E(K) \approx \frac{\lambda}{\lambda + \mu} T_E(1) = \frac{\lambda}{r} T_E(1)$$

When $\lambda > \mu$, the time to extinction scales exponentially with the population ceiling, and does so at a hyperbolically accelerating rate with r . In short, bounded random birth-death processes still approach extinction, but do so incredibly slowly if populations tend towards growth.

To more explicitly determine time to extinction in an exponentially growing population, consider a population subject to simple *Weiner process* type stochas-

tic fluctuations $W(t)$.¹³ Specifically, if dW represents the derivative of $W(t)$ such that $W(0) = 0$, then $W(t)$ is normally distributed around 0 such that

$$W(t) \sim \mathcal{N}(0, t) \quad (5)$$

and the model is written as

$$dN = rNdt + \int N dW$$

This stochastic differential equation implies that for moderate population sizes, where environmental stochasticity prevails over demographic stochasticity (discussed more fully in the next section), then for levels of infinitesimal environmental variance σ^2 , the expected change in log population size $X = \log(N)$ over a small interval $[t, t+h]$ is¹⁴

$$E[X(t+h) - X(t)] \sim \mathcal{N}(\mu h, \sigma^2 h) \text{ where } \mu = r - \sigma^2/2 \quad (6)$$

Solving the stochastic differential equation provides a distribution for X at time t :

$$g(X) = \frac{1}{\sigma\sqrt{2\pi t}} \left(1 - \exp\left(\frac{-2XX_0}{\sigma^2 t}\right) \right) \exp\left(-\frac{(X - X_0 - \mu T)^2}{2\sigma^2 T}\right)$$

Consequently, the distribution of the time to extinction (a population size of $X = 0$, i.e. $N = 1$) is

$$f(T) = \frac{X_0}{\sigma\sqrt{2\pi t}} \exp\left(-\frac{(X_0 + \mu T)^2}{2\sigma^2 T}\right)$$

If $\mu \leq 1$, this integrates to zero; otherwise, it integrates to $1 - \exp(-2\mu X_0/\sigma^2)$. Combining these expressions,

$$P(T < T_E) = \int_0^\infty g(X) dX = \int_1^\infty f(T) dT$$

gives the probability that the population persists to time T without going extinct.

In reality, populations show a combination of deterministic and stochastic behavior over time, and their extinction is a product of both. In the late 1980s, the field of *population viability analysis* (PVA) emerged from the need to find appropriate analytical and simulation methods for predicting population persistence over time. According to one history of PVA, Mark Shaffer's work on grizzly bears in Yellow-

¹³ Steinar Engen and Bernt-Erik Sæther, "Predicting the time to quasi-extinction for populations far below their carrying capacity," *Journal of theoretical Biology* 205, no. 4 (2000): 649–658; Bernt-Erik Sæther and Steinar Engen, "Including uncertainties in population viability analysis using population prediction intervals," *Population viability analysis*, 2002, 191–212.

¹⁴ Engen and Sæther, "Predicting the time to quasi-extinction for populations far below their carrying capacity"; Sæther and Engen, "Including uncertainties in population viability analysis using population prediction intervals."

stone helped birth the field through two important developments, which we break down in turn below.¹⁵

2.1.1 Demographic and Environmental Stochasticity

Shaffer's first major contribution was the use of extinction risk simulations that account for—and differentiate between—two major kinds of stochasticity. *Demographic stochasticity* is defined at the scale of the individual and occurs through random variation in demography and reproduction, while *environmental stochasticity* occurs at a synchronized scale for an entire population (e.g., a bad year may change vital rates uniformly for all individuals in a population). While the impact of environmental stochasticity is ultimately scale-independent, larger populations become less sensitive to demographic stochasticity as they grow. This is due to the integer-based nature of birth-death processes, where populations made up of fewer individuals will suffer a disproportionate effect from a birth or death event.

Demographic and environmental stochasticity have measurably different effects on T_E in basic population models. A simple modeling framework distinguishing between them was laid out in a 1993 paper by Russell Lande.¹⁶ That framework begins again with Eq. 2, except that we now regard r as an explicit function of time $r(t)$ with a mean \bar{r} . In the case of demographic stochasticity, individual variations have no temporal autocorrelation and at the population scale,

$$r(t) \sim \mathcal{N}(\bar{r}, \sigma_d^2/N)$$

where σ_d^2 is the variance of a single individual's fitness per time. As above, the population can be expressed as a diffusion process from the initial population size N_0 :

$$\frac{1}{2}\sigma^2(N_0)\frac{d^2T_E}{dN_0^2} + \mu(N_0)\frac{dT_E}{dN_0} = -1$$

The solution of that differential equation for T_E (where extinction happens at $N = 1$) is given as a function of the initial population size:

$$T_E(N_0) = 2 \int_1^{N_0} e^{-G(z)} \int_z^K \frac{e^{G(z)}}{\sigma^2(y)} dy dz$$

where

$$G(y) = 2 \int_1^y \frac{\mu(N)}{\sigma^2(N)} dN$$

¹⁵ Steven R Beissinger, "Population viability analysis: past, present, future," *Population viability analysis*, 2002, 5–17.

¹⁶ Russell Lande, "Risks of population extinction from demographic and environmental stochasticity and random catastrophes," *American Naturalist*, 1993, 911–927.

For populations experiencing demographic stochasticity and starting at their carrying capacity, this gives us an expression for extinction time that is perhaps slightly clearer:

$$T_E = \left(\frac{1}{\bar{r}} \int_1^K \frac{e^{2\bar{r}(N-1)N/\sigma_d^2}}{N} dN \right) - \frac{\ln K}{\bar{r}}$$

Thus Lande argues in¹⁷ that when \bar{r} is positive, MTE scales exponentially with carrying capacity, while when \bar{r} is negative it scales logarithmically with carrying capacity (i.e., $T_E \propto \ln(K)$), much like in the deterministic decline given by Eqs. 3 & 4). In contrast, in the case of environmental stochasticity, the variance acts on the entire population at once (cf. Eq. 6):

$$E[\ln N(t)] = \ln N_0 + (\bar{r} - \sigma_e^2/2)t$$

and the mean time to extinction is now given by¹⁸

$$T_E = \frac{2}{V_e c} \left(\frac{K^c - 1}{c} - \ln K \right)$$

where

$$c = \frac{2\bar{r}}{\sigma_e^2} - 1$$

In the case of environmental stochasticity, if the “long-run growth rate” ($\bar{r} = \bar{r} - \sigma_e^2/2$) is zero or negative, MTE again scales logarithmically with K . When long-run growth is positive, the dynamic is a bit more complicated:

$$T_E \approx 2K^c / (\sigma_e^2 c^2) \quad \text{if} \quad c \ln K \gg 1$$

In this case, the scaling of MTE with K bends up if and only if $\bar{r}/\sigma_e^2 > 1$ (i.e., if and only if the intrinsic growth rate exceeds environmental variation).

2.1.2 Minimum Viable Populations and Effective Population Size

The second major contribution of Shaffer’s work was the introduction of the concept of a *minimum viable population* (MVP). In Shaffer’s original work, MVP is defined as the smallest possible population for which there is a 95% chance of persistence (a 5% or lower chance of extinction) after 100 years. In their foundational treatment of the minimum viable population concept, Gilpin and Soulé¹⁹ identify four special cases—*extinction vortices*—in which a population is likely to tend below the MVP and towards ultimate extinction.

¹⁷ Lande, “Risks of population extinction from demographic and environmental stochasticity and random catastrophes.”

¹⁸ Ibid.

¹⁹ Michael E Gilpin, “Minimum viable populations: processes of species extinction,” *Conservation biology: the science of scarcity and diversity*, 1986, 19–34.

The first, the *R Vortex*, is perhaps the most obvious: demographic stochasticity (variation in r) reduces populations and increases variation in r , a positive feedback loop of demographic stochasticity directly driving populations to extinction. The *D Vortex* occurs when the same processes—potentially in concert with external forces—produce increased landscape fragmentation (see §3.1.1 for an explanation of D), which not only reduces local population sizes (increasing local extinction rate) but also has subtle effects on population genetic diversity. The final two vortices—the *F Vortex* and *A Vortex*—both concern the genetic and evolutionary trajectories of small stochastic populations. In the first, inbreeding and demographic stochasticity form a feedback cycle, while in the latter, maladaptation is the underlying mechanism of extinction. Both are especially relevant in research surrounding phenomena like climate change, but fully understanding them requires a mathematical language for the genetic behavior of near-extinction populations.

In heavily subdivided populations with low dispersal, increased inbreeding can lead to decreased genetic diversity and the accumulation of deleterious or maladapted alleles that make the total population less viable than its size might indicate. As a consequence, intermediate-sized populations with low genetic diversity can behave, demographically, like small populations. *Effective population size*, or N_e , quantifies that phenomenon, expressing the genetically or reproductively “effective” number of individuals in a population. In some cases, measuring population size with N_e may more readily allow the computation of a meaningful and predictive MVP, by removing some of the variability between different populations of the same size, and by more accurately capturing the long-term reproductive potential of the available genetic material. (Relatedly, it is worth noting that in one unusual study, it was found that there is no statistical link between species MVP and global conservation status.²⁰)

A number of different approaches exist for the estimation of N_e . Sewall Wright, who created the concept of effective population size, offered one interpretation based on neighborhoods. In his model, offspring move a distance away from their parent based on a two-dimensional spatial normal distribution with the standard deviation σ .²¹ If individuals have a density D , then

$$N_e = 4\pi\sigma^2 D$$

Wright²² also provides a more commonly invoked method of calculating N_e based on sex structure, using N_m and N_f to respectively denote the number of breeding females and males in the population:

$$N_e = \frac{4N_m N_f}{N_m + N_f}$$

²⁰ Barry W Brook, Lochran W Traill, and Corey JA Bradshaw, “Minimum viable population sizes and global extinction risk are unrelated,” *Ecology letters* 9, no. 4 (2006): 375–382.

²¹ Sewall Wright, “Isolation by distance under diverse systems of mating,” *Genetics* 31, no. 1 (1946): 39.

²² Sewall Wright, “The interpretation of population structure by F-statistics with special regard to systems of mating,” *Evolution*, 1965, 395–420.

In such an approach, a population of all males or all females would have an N_e of 0 (because no new offspring could be produced in the next generation, rendering the population functionally extinct). That method of deriving N_e is still frequently cited in population conservation work to the present day, as small populations tend to stochastically deviate from a 50:50 sex ratio, sometimes severely impacting long-term survival.

A more genetics-based method of calculating N_e comes from the Wright-Fisher model of a two-allele one-locus system, referred to as the *variance effective population size*.²³ In that model, variance between generations $\sigma^2(a)$, for allele A with frequency a , is given by $a(1-a)/2N$, yielding an effective population size of

$$N_e = \frac{a(1-a)}{2\sigma^2}$$

Alternatively, for a locus with a greater degree of polymorphism, or multi-locus microsatellite data, genetic diversity θ and mutation rate μ are related by

$$N_e = \frac{\theta}{4\mu}$$

A more commonly used metric in current literature is *inbreeding effective population size*. To construct that metric, we start by defining population-level measures of heterozygosity. In the simplest Hardy-Weinberg formulation for a two allele system with allele frequencies a and $1-a$, the expected fraction of heterozygote offspring $E(H) = 2a(1-a)$. By counting the real fraction of heterozygotes and comparing, we can measure the assortiveness of mating:

$$f = \frac{E(H) - H}{H}$$

That value f is called the inbreeding coefficient, ranging from 0 to 1; again according to Wright,²⁴ N_e should be calculated such that it satisfies

$$N_e = \frac{1}{2\Delta f}$$

where Δf is the change per generation (in a declining or small population, genetic diversity decreases at a rate determined by the population size and inbreeding).

Returning to the extinction vortex concept with N_e in mind clarifies the genetic component of those extinction processes. While the *D Vortex* reduces N_e as a byproduct of fragmentation (in fact, decreasing neighborhood size), the last two extinction vortices bring N_e below the MVP through specifically genetic modes of extinction. In the *F Vortex*, a positive feedback loop between increased inbreeding (hence f , the inbreeding coefficient) and decreases in effective population size drive

²³ Brian Charlesworth and Deborah Charlesworth, "The Evolutionary Effects of Finite Population Size: Basic Theory," chap. 5 in *Elements of Evolutionary Genetics* (2012), 195–244.

²⁴ Fred W Allendorf and Nils Ryman, "The role of genetics in population viability analysis," *Population viability analysis*. University of Chicago Press, Chicago, 2002, 50–85.

a population to extinction over a few generations. A notorious real-world example of such a process might be the near-extinction (or extinction, depending on your species concept) of the Florida panther, a subspecies of *Puma concolor* ultimately rescued through outbreeding with Texas panthers. All things considered, their rescue was both fortuitous and improbable, as the species was assigned a 5% or less chance of avoiding imminent extinction in 1995.²⁵ Finally, in the *A vortex* (i.e., adaptation), decreased N_e acts as a buffer to the strength of selection acting on phenotypes that are closely paired with environmental variation or change, leading to mismatch between them that reduces both r and N (and N_e) until extinction (a process we cover in much greater detail in §4.1). Obviously, the four vortices are non-independent processes, and probably often exist in combination in real-world cases.

2.1.3 Population Viability Analysis: Theory and Practice

Population viability analysis is conventionally implemented by modeling the dynamics of different compartmental classes within a population, such as age and sex structure. The foundations of that method date as far back as P. H. Leslie's population analyses in the late 1940s in the framework of discrete matrix models and linear systems theory. Formulations of the Leslie model and the theory behind such models can be found in several expository texts,²⁶ with a brief outline provided here. In the Leslie model, the population is divided into n age classes, where $N_i(t)$ is used to denote the number of individuals in age class i at time t . In each age class, the parameter s_i ($0 < s_i \leq 1$) is used to represent the proportion of individuals aged i that survive to age $i + 1$, in which case the variables $N_i(t)$ and $N_{i+1}(t + 1)$ are linked by the equation

$$N_{i+1}(t + 1) = s_i N_i(t) \quad (7)$$

At some point we either terminate this series of equations at age n by assuming that $s_n = 0$ (i.e. no individuals survive beyond age n) or we interpret N_n as the group of individuals in the population aged n and older and use the equation

$$N_n(t + 1) = s_{n-1} N_{n-1}(t) + s_n N_n(t) \quad (8)$$

to imply that all individuals aged n and older are subject to the survival parameter s_n (i.e., individuals older than age n are indistinguishable from individuals aged n). If we now interpret $N_0(t)$ as all newborn individuals born just after individuals have progressed one age class, then $N_0(t)$ can be calculated using the formula

²⁵ Warren E Johnson et al., "Genetic restoration of the Florida panther," *Science* 329, no. 5999 (2010): 1641–1645.

²⁶ Hal Caswell, *Matrix population models* (Wiley Online Library, 2001); Wayne M Getz and Robert G Haight, *Population harvesting: demographic models of fish, forest, and animal resources*, vol. 27 (Princeton University Press, 1989).

$$N_0(t) = \sum_{i=1}^n b_i N_i(t) \quad (9)$$

where b_i is the average (expected) number of progeny produced by each individual aged i . In this model we have not differentiated between the sexes; so, for example, if each female aged i is expected to produce 3 young and the population has a 1:1 sex ratio (same number of males to females) then $b_i = 1.5$ for this age class. If we now apply Equation 7 for the case $i = 0$, we obtain the equation

$$N_1(t+1) = s_0 N_0(t) = s_0 \sum_{i=1}^n b_i N_i(t) \quad (10)$$

Equations 7 to 10 can be written compactly in matrix notation (a *Leslie matrix*) as

$$\mathbf{N}(t+1) = L\mathbf{N}(t) \quad (11)$$

$$\text{where } \mathbf{N} = \begin{pmatrix} N_1 \\ \vdots \\ N_n \end{pmatrix} \text{ and } L = \begin{pmatrix} s_0 b_1 & \cdots & s_0 b_{n-1} & s_0 b_n \\ s_1 & \cdots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & s_{n-1} & s_n \end{pmatrix}$$

The matrix L is a nonnegative matrix since all its elements are non-negative, with at least one positive element. Further, if there exists some integer $p > 0$ such that L^p is positive (i.e. all its elements are positive), then it is known from the Perron-Frobenius Theorem that the matrix L has a dominant positive eigenvalue λ_p (known as the Perron root) and a corresponding eigenvector \mathbf{v}_p whose elements are all positive. These values λ_p and \mathbf{v}_p characterize the long term behavior of \mathbf{N} such that

$$\mathbf{N}(t) \sim (\lambda_p)^t \mathbf{v}_p$$

This equation implies that as t gets very large $\mathbf{N}(t)$ grows like $(\lambda_p)^t$ and the ratio of different age classes matches the ratio of elements of \mathbf{v}_p . Thus, if $\lambda_p > (<)1$, $\mathbf{N}(t)$ will grow (decline) geometrically at the rate λ_p and approach the so-called *stable age-distribution*, as characterized by the ratio of consecutive elements of \mathbf{v}_p . In other words, this model predicts that the population will go extinct whenever the largest eigenvalue of L is less than one (i.e., $0 < \lambda_p < 1$). On the other hand, if $\lambda_p > 1$, then we expect density-dependent effects at some point to rein in the unfettered growth by causing survival rates to decline. In particular, if the survival rate s_0 of the youngest age class is the most sensitive of the survival rates to increases in the total biomass density

$$B = \sum_{i=1}^n w_i N_i \quad (12)$$

of the population, where $w_i > 0$ is the average weight of an individual in age class i , then we should replace s_0 in Eqn. 10 with an expression such as




Fig. 1 An example PVA without (A) and with (B) the influence of demographic stochasticity, and with no (blue), medium (red) or high (purple) environmental stochasticity. Based on many numerical simulations, an “extinction curve” can be plotted from the probability of population survival over time (C). This analysis can be used to make decisions about management and conservation: here, illustrating that three populations with migration between them survive for much longer in a poached population of rhinos than a single population. An interactive tutorial of PVA, which can be adjusted to produce anything from the simplest population dynamics to a stochastic, structured metapopulation experiencing harvesting, can be found at <http://www.numerusinc.com/webapps/pva>.

$$s_0 = \frac{\hat{s}_0}{1 + (B/K_0)^\gamma} \quad (13)$$

where \hat{s}_0 is the density-independent survival rate, K_0 is the density at which \hat{s}_0 is halved, and $\gamma > 1$ is termed the “abruptness” (as it controls the abruptness in the onset of density, approaching a step down function as γ gets large²⁷). Similar modifications can be made to the other survival parameters s_i , depending on their sensitivity to changes in population density.

Stochastic equivalents of these deterministic models typically treat the survival rates s_i as probabilities that each individual survives each time period, rather than as the proportion of individuals surviving each time period; and b_i itself is a random variable drawn from an appropriately defined distribution (usually the binomial distribution). Stochastic models of this sort can be made even more complex by adding more population structure (e.g. genetic variability) or increased levels of complexity (e.g. modeling at the metapopulation scale, discussed in §3, or adding underlying environmental variation or other landscape structure). Though MVP or extinction rates might be difficult to calculate analytically for models of this level of complexity, repeated simulation can easily allow empirical derivation of these properties of a system,²⁸ and is perhaps the most widespread practice in existence for estimating population extinction risk in conservation research. An example using an interactive web app²⁹ is shown in Figure 1.

²⁷ Wayne M Getz, “A hypothesis regarding the abruptness of density dependence and the growth rate of populations,” *Ecology* 77, no. 7 (1996): 2014–2026.

²⁸ Brett A Melbourne and Alan Hastings, “Extinction risk depends strongly on factors contributing to stochasticity,” *Nature* 454, no. 7200 (2008): 100–103.

²⁹ Wayne M Getz et al., “A web app for population viability and harvesting analyses,” *Natural Resource Modeling*, 2016,

Is population viability analysis the perfect tool for studying extinction? PVA is currently the gold standard for most applied conservation research, both by virtue of being an all-encompassing term for quantitative extinction risk modeling at the population scale, and the absence of any suitable alternative. But PVA, like any quantitative tool, is tremendously sensitive to assumptions, parameterization, and data availability. Imprecise parameterization, from noisy data or tenuous assumptions, proportionally reduces the precision of PVA, to a degree that may be hard to characterize; it is consequently important to report uncertainty from PVA estimates.³⁰ Similarly, given the challenges of developing an accurate and precise model, it has been widely agreed that PVA should be treated as more of a relative or comparative tool (for instance, between different management or conservation scenarios), and authors should refrain from treating minimum viable population or extinction time estimates as absolute, precise estimates.³¹ Despite this, many managers still use PVA as an absolute estimate of extinction risk, a pervasive problem with no clear solution.

2.2 Case Study: PVA, Disease, and Evolutionary Rescue

In 2015, an epidemic of unknown identity eliminated more than half of the population of the critically endangered saiga antelope (*Saiga tatarica*), in the short span of three weeks. While the causative agent was ultimately identified as a species of *Pasteurella*, the mechanism by which a normally asymptomatic non-pathogenic bacterium killed at least 130,000 antelopes is still in question.³² Literature explaining the die-off, or predicting the consequences for the species, remains comparatively limited; the fate of the species remains uncertain, and it may yet face extinction in the coming years.

Disease is rarely responsible for the extinction of a cosmopolitan species; but for already-threatened species like the saiga, it can be one of the most rapid, unpredictable and unpreventable mechanisms of extinction. Disease has been implicated in a handful of notable wildlife extinctions, like that of the thylacine (*Thylacinus cynocephalus*) and Carolina parakeet (*Conuropsis carolinensis*), and has been the definitive mechanism of extinction for species like the eelgrass limpet (*Lottia alveus*).³³ While most diseases co-evolve with their hosts to an optimal virulence

³⁰ Stephen P Ellner et al., "Precision of population viability analysis," *Conservation Biology* 16, no. 1 (2002): 258–261.

³¹ J Michael Reed et al., "Emerging issues in population viability analysis," *Conservation biology* 16, no. 1 (2002): 7–19; Tim Coulson et al., "The use and abuse of population viability analysis," *Trends in Ecology & Evolution* 16, no. 5 (2001): 219–221; Steven R Beissinger and M Ian Westphal, "On the use of demographic models of population viability in endangered species management," *The Journal of wildlife management*, 1998, 821–841.

³² EJ Milner-Gulland, "Catastrophe and hope for the saiga," *Oryx* 49, no. 04 (2015): 577–577.

³³ Francisco De Castro and Benjamin Bolker, "Mechanisms of disease-induced extinction," *Ecology Letters* 8, no. 1 (2005): 117–126.

that prevents the species from reaching extinction, diseases that can persist in the environment may be released from such constraints and be more likely to evolve “obligate killer” strategies (like that of anthrax³⁴). Fungal pathogens in particular tend to grow rapidly in hosts and spread rapidly between them, which can result in population collapses before optimal virulence levels can be attained.³⁵

Two notable fungal diseases have recently demonstrated the destructive potential of environmentally transmitted pathogens. Perhaps the most significant example of disease-driven extinctions is the trail of destruction caused by the chytrid fungus *Batrachochytrium dendrobatidis* (Bd). Bd has been found in at least 516 species of amphibian³⁶ and has driven decline or extinction in at least 200,³⁷ including at least two thirds of the genus *Atelopus* alone.³⁸ According to some estimates, current extinction rates that amphibians face (largely but not entirely due to chytrid) are roughly 200 times the background rate; including declining species, that estimate is closer to an even more staggering 25-45,000.³⁹ White nose syndrome (WNS; *Geomyces destructans*), a similar fungal epizootic, has similarly spread through bat populations in the eastern United States, causing widespread population-level die-offs since the mid-2000s. While white-nose syndrome has yet to drive any entire species to extinction, significant concern remains regarding its ongoing spread; one study in 2010 using population viability analysis suggested a 99% extinction risk for the little brown bat (*Myotis lucifugus*) in under two decades.⁴⁰ Even in a best-case scenario where white-nose mortality was reduced to one twentieth of its rate, substantially reducing extinction risk, bats would still be reduced to one percent of their original population size.

White-nose syndrome has also become a potential case study for evolutionary rescue, one of the most controversial phenomena in extinction research. The premise that rare genes for resistance or tolerance can bring a disease-ridden population back from the brink of extinction has theoretical support, and potentially indicated from the rapid evolutionary response of certain hosts documented throughout the literature.⁴¹ But WNS constitutes one of the most interesting and controversial examples because, while populations show some sign of recovery from the disease, at the

³⁴ SA Frank and P Schmid-Hempel, “Mechanisms of pathogenesis and the evolution of parasite virulence,” *Journal of evolutionary biology* 21, no. 2 (2008): 396–404.

³⁵ Matthew C Fisher et al., “Emerging fungal threats to animal, plant and ecosystem health,” *Nature* 484, no. 7393 (2012): 186–194.

³⁶ Deanna H Olson et al., “Mapping the global emergence of *Batrachochytrium dendrobatidis*, the amphibian chytrid fungus,” *PloS one* 8, no. 2 (2013): e56802.

³⁷ Lee Francis Skerratt et al., “Spread of chytridiomycosis has caused the rapid global decline and extinction of frogs,” *EcoHealth* 4, no. 2 (2007): 125–134.

³⁸ J Alan Pounds et al., “Widespread amphibian extinctions from epidemic disease driven by global warming,” *Nature* 439, no. 7073 (2006): 161–167.

³⁹ Malcolm L McCallum, “Amphibian decline or extinction? Current declines dwarf background extinction rate,” *Journal of Herpetology* 41, no. 3 (2007): 483–491.

⁴⁰ Winifred F Frick et al., “An emerging disease causes regional population collapse of a common North American bat species,” *Science* 329, no. 5992 (2010): 679–682.

⁴¹ Sonia Altizer, Drew Harvell, and Elizabeth Friedle, “Rapid evolutionary dynamics and disease threats to biodiversity,” *Trends in Ecology & Evolution* 18, no. 11 (2003): 589–596.

time of writing, no definitive genetic mechanism for resistance has been isolated, a necessary component of demonstrating evolutionary rescue from disease-induced extinction.⁴² Consequently, speculation about evolutionary rescue is controversial and so far has been conducted in primarily theoretical settings. In an age-structured matrix population model proposed by Maslo and Fefferman, two scenarios for recovery from WNS are considered.⁴³ In one, bats' adaptive immunity leads to re-stabilization at much lower levels overall, but a much faster recovery to a stable balance of juveniles (J) and adults (A), with subscript t denoting the number of individuals in these two age classes at time t . In that model, in the absence of white-nose,

$$\begin{pmatrix} J_{t+1} \\ A_{t+1} \end{pmatrix} = \begin{pmatrix} 0.95 & 0.35 \\ 0.95 & 0.87 \end{pmatrix} \begin{pmatrix} J_t \\ A_t \end{pmatrix}$$

In a second model they propose, recovery comes not from adaptive immunity but from innate immunity through a genetic mechanism for resistance. In that scenario a robust type (R) is present in the gene pool with frequency p_t and protects from white nose infection; the remainder of individuals are wild type (WT). In the evolutionary rescue model, all individuals have lower survivorship, but wild type bats fare much worse and reproduce at slightly slower rates (imposing strong selection against WT):

$$\begin{pmatrix} J_{t+1} \\ A_{t+1} \end{pmatrix} = p_t \begin{pmatrix} 0.86 & 0.32 \\ 0.86 & 0.78 \end{pmatrix} \begin{pmatrix} J_t^R \\ A_t^R \end{pmatrix} + (1 - p_t) \begin{pmatrix} 0.52 & 0.27 \\ 0.52 & 0.46 \end{pmatrix} \begin{pmatrix} J_t^{WT} \\ A_t^{WT} \end{pmatrix}$$

In this model, an 11-year stabilization period ultimately leads to population recovery with a positive net growth rate (calculated as the dominant eigenvalue $\lambda = 1.05$), potentially saving populations from extinction. Despite the lack of genetic evidence for evolutionary rescue, Maslo and Fefferman propose that observed similarities between the dynamics they observe and real data on white-nose outbreaks suggests that evolutionary rescue may be happening in real time. Other work since has similarly supported the idea that bat populations may be recovering. Validating these results requires that researchers identify genetic variation between populations associated with differential outcomes, and develop models more directly informed by those mechanisms.

3 The Metapopulation Scale

Populations rarely exist in isolation, but are often connected to other populations through dispersal processes, creating a *metapopulation*. Metapopulations are considered to be in a relatively constant state of flux, as local extinctions of species

⁴² Altizer, Harvell, and Friedle, "Rapid evolutionary dynamics and disease threats to biodiversity."

⁴³ Brooke Maslo and Nina H Fefferman, "A case study of bats and white-nose syndrome demonstrating how to model population viability with evolutionary effects," *Conservation Biology* 29, no. 4 (2015): 1176–1185.

in habitat patches are buffered by re-colonization from local dispersal. In this way, dispersal can be beneficial or detrimental to metapopulation persistence. Under high dispersal, patches become homogeneous and population dynamics tend to become synchronous. This synchrony is destabilizing, in that periods of low population sizes will be experienced by all patches, increasing the likelihood of stochastic extinction of the entire metapopulation. On the other hand, too little dispersal will result in spatial clustering of a species, as the species will be confined to the set of patches that can be successfully reached and colonized and similarly potentially increasing extinction risk.⁴⁴

The importance of dispersal to patch-level colonization and metapopulation persistence highlights that extinction processes occur at two scales: the local patch-level (i.e., a single population in the network of habitat patches) or at the entire metapopulation level (i.e., either through catastrophic events or cascading local extinctions). Extinctions of single patches can occur as a result of demographic, environmental, or genetic stochasticity (addressed in more detail in §2.1.1), or through extrinsic events related to habitat loss or natural enemies.⁴⁵ Metapopulation level extinction can also result from environmental stochasticity at the regional scale,⁴⁶ provided this stochasticity is spatially autocorrelated, such that it is expected to promote synchronous dynamics among habitat patches.⁴⁷

3.1 Basic Metapopulation Models and Extinction

In the classic metapopulation model described by Richard Levins, the balance between patch colonization (c) and local extinction (e) determines patch occupancy dynamics. In this case, local habitat patches are either occupied or unoccupied, and both patch number and the spatial orientation of patches are undescribed. Dispersal among habitat patches can rescue patches from extinction, or allow for the recolonization of extinct patches. All patches are treated as equal, so that any patch is suitable for a species, and (as a simplifying assumption) all habitat patches can be reached from all other patches. This simplified representation treats space as implicit, and patch quality and size as constant; rather than an explicit population size, patch occupancy is just a 0 or 1 state. The dynamics of the proportion occupied patches, P , are given by a differential equation:

⁴⁴ Karen C Abbott, “Does the pattern of population synchrony through space reveal if the Moran effect is acting?,” *Oikos* 116, no. 6 (2007): 903–912; Karen C Abbott, “A dispersal-induced paradox: synchrony and stability in stochastic metapopulations,” *Ecology letters* 14, no. 11 (2011): 1158–1169.

⁴⁵ Ilkka Hanski, “Metapopulation dynamics,” *Nature* 396, no. 6706 (1998): 41–49.

⁴⁶ James C Bull et al., “Metapopulation extinction risk is increased by environmental stochasticity and assemblage complexity,” *Proceedings of the Royal Society of London B: Biological Sciences* 274, no. 1606 (2007): 87–96.

⁴⁷ Ana R Gouveia, Ottar N Bjørnstad, and Emil Tkadlec, “Dissecting geographic variation in population synchrony using the common vole in central Europe as a test bed,” *Ecology and evolution* 6, no. 1 (2016): 212–218.

$$\frac{dP}{dt} = cP(1-P) - eP \quad (14)$$

In that equation, extinction is a random process for every occupied patch that is entirely independent of the state of the system. In contrast, colonization rates depend both on the fraction of occupied and unoccupied patches, as emigrants move from occupied patches to re-colonize unoccupied ones. The balance between the two processes of extinction and colonization determines long-term persistence of the metapopulation;⁴⁸ that is, a necessary condition for metapopulation persistence in this model is

$$\frac{e}{c} < 1$$

At a non-trivial equilibrium, the patch occupancy is given as

$$\hat{P} = 1 - \frac{e}{c} \quad (15)$$

This suggests that the equilibrium fraction of occupied patches is a simple function of colonization (c) and extinction (e). If extinction rates are greater than zero, this implies that the equilibrium occupancy is less than one even if colonization exceeds extinction; that is to say, not every patch will ever be stably filled if extinction is nontrivial. This shows that even a metapopulation in equilibrium is still in a constant state of patch-level flux. In real applications, this implies that just because a patch of habitat is empty, that may not imply it is uninhabitable; and similarly, just because a population goes extinct, it may not be indicative of broader declines or instability.

While admittedly a simple representation of a metapopulation, the Levins model can yield important insights into spatial population dynamics.⁴⁹ For instance, the mean time to extinction of any given population/patch is the inverse of the rate (i.e., $T_E = 1/e$), providing a link to the models at the population scale discussed above. We can take the Levins model a step further to explicate the relationship between patch occupancy and overall mean time to extinction T_M at the metapopulation scale. Starting with the assumption that the total H patches each have their own average extinction time T_L (which should be $1/e$),

$$T_M = T_L \exp \left(\frac{(\hat{P}H)^2}{2H(1-\hat{P})} \right)$$

Consequently, using Eq. 15, we can also express T_M as

$$T_M = T_L \exp \left(\frac{H}{2} \left(cT_L + \frac{1}{cT_L} - 2 \right) \right)$$

⁴⁸ Richard Levins, "Some demographic and genetic consequences of environmental heterogeneity for biological control," *Bulletin of the Entomological society of America* 15, no. 3 (1969): 237–240.

⁴⁹ F Elías-Wolff et al., "How Levins' dynamics emerges from a Ricker metapopulation model," *Theoretical Ecology* 2, no. 9 (2016): 173–183.

showing that metapopulation extinction time increases exponentially, not linearly, with the MTE of individual habitat patches.⁵⁰

The simplicity of the Levins model has resulted in a sizable body of literature surrounding and extending the model. For instance, in the original Levins' model all patches are equidistant from one another, identical in quality, and can only be in one of two potential states (occupied or unoccupied), but each of these conditions is frequently adjusted in derivative stochastic patch occupancy models (SPOMs). Researchers have shown that despite the simplicity, Levins-type dynamics can emerge from more complicated stochastic metapopulation models,⁵¹ and extensions of the Levins model continue to provide insight into the influence of habitat patch size and topography (i.e., spatial orientation of habitat patches) on metapopulation persistence.⁵²

3.1.1 Island Biogeography and Metapopulation Capacity

A simple extension of the Levins model considers a set of spatially explicit patches of variable size, where a distance matrix D describes the distance between all patches in the metapopulation. The model borrows elements of MacArthur and Wilson's *Theory of Island Biogeography*,⁵³ such that distance between patches (D_{ij}) and patch area (A_i) influence extinction and colonization processes, where the patch extinction rate scales with patch area ($e_i = e/A_i$), and colonization (c_i) becomes a property of distance (D_{ij}), patch area (A_i), and dispersal rate (α) where

$$c_i = \sum_{j \neq i} e^{-\alpha D_{ij}} A_j p_j(t)$$

This suggests that the mean time to extinction of a habitat patch ($1/e_i$) is determined by the area of the patch. This makes the occupancy probability of each patch in the metapopulation, described in terms of matrix M

$$M_{ij} = e^{-\alpha D_{ij}} A_i A_j$$

and the leading eigenvalue of this matrix M describes the persistence of the metapopulation (also known as *metapopulation capacity*⁵⁴ or λ_M). The condition for metapopulation persistence is that the dominant eigenvalue of M must be greater than the ratio between extinction and colonization rates:

⁵⁰ Ilkka Hanski, "Single-species metapopulation dynamics: concepts, models and observations," *Biological Journal of the Linnean Society* 42, nos. 1-2 (1991): 17–38.

⁵¹ Elías-Wolff et al., "How Levins' dynamics emerges from a Ricker metapopulation model."

⁵² Luis J Gilarranz and Jordi Bascompte, "Spatial network structure and metapopulation persistence," *Journal of Theoretical Biology* 297 (2012): 11–16.

⁵³ Robert H MacArthur and Edward O Wilson, *Theory of Island Biogeography* (MPB-1), vol. 1 (Princeton University Press, 2015).

⁵⁴ Ilkka Hanski and Otso Ovaskainen, "The metapopulation capacity of a fragmented landscape," *Nature* 404, no. 6779 (2000): 755–758.

$$\lambda_M > e/c$$

While spatially explicit, this approach assumes that dispersal among habitat patches is determined by patch area and distance to other patches, ignoring population dynamics in each patch. However, since habitat patches vary in their size and connectedness to other patches, it is possible to determine the relative importance of each habitat patch to metapopulation persistence in this framework,⁵⁵ potentially informing conservation and management decisions.⁵⁶

3.1.2 Incorporating Patch Dynamics

The above extension of the Levins' model allows for patches to vary in size and connectedness. Another extension is to consider the abundances of habitat patches within the metapopulation, thus considering the dynamics of each patch, and the effects of dispersal among local populations.⁵⁷

$$N_i(t+1) = R_i(t)N_i(t)e^{-N_i/K}$$

This expression assumes that the growth rate of each habitat patch is R_i , and that the carrying capacity is a constant K . If we assume that the population growth rates (r_i) are independent and identically-distributed Gaussian random variables, this causes R_i values to be log-normally distributed, and allows us to define persistence thresholds for the metapopulation based on the variance in the population growth rates r_i . The threshold for metapopulation persistence relies on exceeding a threshold value ($\sigma_{\text{threshold}}$) in terms of the variance among local patch population growth rates (r_i). If μ is the mean local population growth rate over time, this threshold is

$$\sigma_{\text{threshold}} > \sqrt{2|\mu|}$$

This model can be extended to yield many interesting conclusions. For instance, if populations have influence on where their offspring go, population growth rates may be maximized by seeding offspring in less than suitable "sink" habitat if habitat quality fluctuates with time, and when the "source" habitat occasionally experiences catastrophes.⁵⁸ The complexity of metapopulation dynamics in the face of environ-

⁵⁵ Hanski and Ovaskainen, "The metapopulation capacity of a fragmented landscape"; Jacopo Grilli, György Barabás, and Stefano Allesina, "Metapopulation persistence in random fragmented landscapes," *PLoS Comput Biol* 11, no. 5 (2015): e1004251.

⁵⁶ Subhashni Taylor et al., "Applications of Rapid Evaluation of Metapopulation Persistence (REMP) in Conservation Planning for Vulnerable Fauna Species," *Environmental management* 57, no. 6 (2016): 1281–1291.

⁵⁷ Manojit Roy, Robert D Holt, and Michael Barfield, "Temporal autocorrelation can enhance the persistence and abundance of metapopulations comprised of coupled sinks," *The American Naturalist* 166, no. 2 (2005): 246–261.

⁵⁸ V AA Jansen and Jin Yoshimura, "Populations can persist in an environment consisting of sink habitats only," *Proc. Natl. Acad. Sci. USA* 95 (1998): 3696–3698.

mental stochasticity, variable patch quality, dispersal, and competition has fueled expansive theoretical work.⁵⁹ An obvious next step is to scale from single species metapopulations to multi-species communities (i.e., metacommunities), which allows for the modeling of how species interactions, predator-prey dynamics, and community assembly relate to persistence.⁶⁰

4 The Species Scale

Extinction is defined at the scale of the species, but it is also at this level of taxonomic resolution that it is perhaps hardest to quantify—and, to summarize—due to considerable diversity of approaches and applications. We explore in this chapter two applied extensions of that body of theory, corresponding to two common quantitative frameworks for species-level extinctions. In the first, the complete loss of suitable habitat leads to an inevitable—if not immediate—extinction. Species can escape extinction through three primary channels: acclimation, adaptation, and migration. Species distribution models are often used to calculate extinction risk at the community scale in that framework (described in greater detail below), but they can only at best include the last of those three rescue processes. Evolutionary models, on the other hand, can link demography and genetics to the overall risk of extinction in a changing environment; we explore that application here in the context of both adaptation and phenotypic plasticity.

The second framework is based in the notion that population extinctions become species extinctions; and so the framework for population (and metapopulation) viability analysis described above acts as a sufficient method for estimating species extinction risk. In many cases, that may be a safe assumption, as near-extinction species are reduced down to a single persistent population or a handful in isolated refugia. But in real applications, persistence in small isolated refugia may be difficult to study, or even observe with any regularity; consequently, an entire body of literature has been developed to relate extinction risk to the sightings of rare species. That body of theory allows two applications: the posthumous assignment of extinction dates to extinct species, and sighting-based hypothesis testing for a species of unknown extinction status. We explore both applications below.

⁵⁹ David P Matthews and Andrew Gonzalez, “The inflationary effects of environmental fluctuations ensure the persistence of sink metapopulations,” *Ecology* 88, no. 11 (2007): 2848–2856; Michael B Bonsall and Alan Hastings, “Demographic and environmental stochasticity in predator–prey metapopulation dynamics,” *Journal of Animal Ecology* 73, no. 6 (2004): 1043–1055.

⁶⁰ Mathew A Leibold et al., “The metacommunity concept: a framework for multi-scale community ecology,” *Ecology letters* 7, no. 7 (2004): 601–613.

4.1 Adaptation and Plasticity in a Changing Environment

Bounding uncertainty is the seminal challenge in extinction research, and in the real world, species' potential to acclimate and adapt to changing environments confers an unknown degree of robustness that may give species a chance at evading extinction. As discussed above, evolutionary rescue has been a particularly tantalizing—and controversial—idea in the context of disease research. But more broadly, evidence suggests that extinction risk is heavily complicated by species' variable ability to track changing environments.

Most models of evolutionary rescue approach the problem by explicitly modeling fitness curves and the speed of natural selection. In a foundational paper by Gomulkiewicz & Holt,⁶¹ population size N_t changes over time in response to its mean fitness W_t such that

$$N_t = W_{t-1}N_{t-1} = \prod_{i=1}^{t-1} W_i N_0$$

If fitness is below one (i.e., populations are reproducing at a rate below replacement), then the population will tend towards extinction. The model they present uses a pseudoextinction threshold N_c such that if the initial fitness W_0 is held constant over the entire interval,

$$T_E = \frac{\ln N_c - \ln N_0}{\ln W_0}$$

Without adaptation (i.e. W_t increase above $W_0 < 1$), the population declines to extinction. To model adaptation, Gomulkiewicz & Holt assume that environmental change begins at time 0, adapting a system of equations for describing natural selection on a single phenotypic trait originally proposed by Russell Lande.⁶² In that notation, the trait z has an optimum phenotype z_{opt} . The population mean phenotype is expressed as d_t , the distance of the average z from z_{opt} at each timestep, with an initial value d_0 . As for any quantitative trait, individual phenotypic values z are normally distributed around the population mean with some variance σ_z^2 :

$$z_t \sim \mathcal{N}(d_t, \sigma_z^2)$$

The corresponding fitness function with width ω_z is expressed as a bell curve around the optimum:

$$W(z) = W_{\max} e^{-z^2/2\omega_z}$$

where W_{\max} is the fitness of z_{opt} . The width of the fitness function (which can be interpreted as the strength of selection), the existing variance in the trait, and the dis-

⁶¹ Richard Gomulkiewicz and Robert D Holt, "When does evolution by natural selection prevent extinction?," *Evolution* 49, no. 1 (1995): 201–207.

⁶² Russell Lande, "Natural selection and random genetic drift in phenotypic evolution," *Evolution*, 1976, 314–334.

tance from the optimum, determine how quickly the population evolves; the changing fitness of the population can be expressed as:

$$W_t = W_{\max} \sqrt{\omega_z / (\sigma_z^2 + \omega_z)} e^{-d_t^2 / (2\sigma_z^2 + 2\omega_z)} \quad (16)$$

Even a population with a mean at z_{opt} does not have perfectly maximized fitness, because of the variance around the mean; the actual growth rate of the population when $d_t = 0$ can be expressed as

$$\hat{W} = W_{\max} \sqrt{\omega_z / (\sigma_z^2 + \omega_z)}$$

This provides a clear way to simplify Eq. 16:

$$W_t = \hat{W} e^{-d_t^2 / (2\sigma_z^2 + 2\omega_z)}$$

In this expression, the changing fitness of the population is expressed only as a function of the optimum and the strength of selection on the trait z .

How does the actual distribution of phenotypes change over time? In real systems, evolution is seldom a direct progression towards the optimum, even under hard selection with ample genetic variation. If the trait z has a heritability h^2 (where a heritability of 1 means the trait is perfectly heritable, and 0 would indicate perfect plasticity or no genetic basis), Gomulkiewicz & Holt define a scaleless “evolutionary inertia”

$$k = \frac{\omega_z + (1 - h^2)\sigma_z^2}{\omega_z + \sigma_z^2}; 0 \leq k \leq 1$$

which in turn simplifies how fast the population shifts towards its optimum phenotype:

$$d_t = k^t d_0$$

Together, this set of equations produces the governing expression for the system:

$$t \ln \hat{W} - \frac{d_0^2}{2(\omega_z + \sigma_z^2)} \frac{1 - k^{2t}}{1 - k^2} = \ln \frac{N_c}{N_0}$$

If this equation has no roots when solving for t , then this indicates the population will fall and rise without any real extinction risk; but when it does, the roots are estimates of the time until the population falls below the critical threshold (T_E) and the time until recovery could be evolutionarily possible (T_P in their notation, where N_t passes back above N_c). The interval between these two values is characterized by a small population that, due to demographic stochasticity, would require much more intensive conservation efforts (e.g., managed *ex situ* breeding) than normal to possibly survive that interval. The time to recovery (growth switches from negative to positive even though $N_t < N_c$) is

$$T_R = \frac{1}{\ln k^2} \left(\ln \ln \hat{W} - \ln \frac{d_0^2}{2(\omega_z + \sigma_z^2)} \right)$$

From this expression, Gomulkiewicz and Holt derive a useful finding: “ t_R increases logarithmically with the degree of initial maladaptation ... but is independent of the initial population density.” Or, to rephrase: the possibility and speed of evolutionary rescue depends on the initial phenotypic distribution, the evolutionary inertia, and the speed of selection, but is scale invariant across population sizes; even small populations with high enough genetic diversity and low inertia can be rescued by evolutionary rescue.

The model developed by Gomulkiewicz and Holt sets useful theoretical bounds on the genetically-coded evolution of a trait; but in the real world, phenotypic plasticity complicates this pattern, and presents one of the hardest challenges for predicting how species might escape extinction. In a similar model developed by Chevin *et al.*,⁶³ the trait in question z has a developmental trajectory with both a genetic component and the potential for phenotypic plasticity in response to an environmental variable ε . Their model uses a “reaction norm” approach to plasticity (popularized by Schlichting, Pigliucci and others⁶⁴), breaking down that phenotypic trait into an adaptive genetic component a and a plastic component b that responds to the environmental gradient. They express the distribution of the phenotype $p(z)$ at generation n in an environment changing at rate $\varepsilon(t) = \eta t$ as:

$$p(z) \sim \mathcal{N}(\bar{z}, \sigma_z^2)$$

Here the population mean \bar{z} , expressed in terms of the generation time T under the assumptions that i) developmental plasticity takes effect at time τ during ontogeny and ii) the strength of plasticity b (the slope of a phenotypic reaction norm), takes the form

$$\bar{z} = \bar{a} + b\eta T(n - \tau)$$

Assumption (ii) is of course a limiting one, given that plastic reaction norms are in fact evolvable; but extensions of quantitative theory that incorporate this idea are underdeveloped. We also assume that the variance associated with z has both environmental and genetic components: i.e.,

$$\sigma_z^2 = \sigma_a^2 + \sigma_e^2$$

Assuming there is an optimum phenotype $\theta = B\varepsilon$, where B is the optimal rate of change to plastically track the changing environment, Gomulkiewicz and Holt define a changing population size with a maximum growth rate W_{\max} , such that

$$W(z) = W_{\max} \exp\left(-\frac{(z - \theta)^2}{2\omega_z} - \frac{b^2}{2\omega_b}\right)$$

⁶³ Luis-Miguel Chevin, Russell Lande, and Georgina M Mace, “Adaptation, plasticity, and extinction in a changing environment: towards a predictive theory,” *PLoS Biol* 8, no. 4 (2010): e1000357.

⁶⁴ Carl D Schlichting, Massimo Pigliucci, et al., *Phenotypic evolution: a reaction norm perspective*. (Sinauer Associates Incorporated, 1998).

where both ω 's represent the strength of stabilizing selection (the width of fitness curves, comparable to above). From there they make the link to overall population dynamics, where the intrinsic rate of growth of the population r can be scaled with generation time and related to selection on z as

$$r = \frac{\ln(W)}{T} = \frac{\ln(W_{\max})}{T} - \frac{\ln(1 + \sigma_z^2/\omega_z) + b^2/\omega_b}{2T} - \frac{(\bar{z} - \theta)^2}{2T} \gamma$$

where γ , the strength of stabilizing selection, is given by

$$\gamma = \frac{1}{\omega_z + \sigma_z^2}$$

The first two terms become the maximum possible growth rate r_{\max} if z reaches the optimum θ :

$$r_{\max} = \frac{\ln(W_{\max})}{T} - \frac{\ln(1 + \sigma_z^2/\omega_z) + b^2/\omega_b}{2T}$$

From this expression for population dynamics, Chevin *et al.* derive a formula for the critical rate of environmental change, above which ($\eta > \eta_c$), plasticity and adaptation combined still fail to prevent extinction (recalling that B is the optimal rate of change to plastically track the changing environment and b is the slope of the phenotypic reaction norm):

$$\eta_c = \sqrt{\frac{2r_{\max}\gamma}{T}} \frac{h^2\sigma_z^2}{|B - b|}$$

From this expression, it is very easy to determine the long term tendency of the population to extinction or survival as a function only of the degree of plasticity and the associated strength of costs (ω_b). The greater the extent of plasticity, the more the costs of plasticity separate out population trajectories; but when plasticity has a weak slope, the extinction isoclines converge towards the same threshold. This conceptualization of adaptation to environmental change as a single-trait system with readily measured costs of adaptive plasticity is obviously an idealization, but also clearly illustrates a number of important points. While adaptive genetic variation has a clear direct relationship to evolutionary rescue, plasticity also plays an important role; and quantifying plasticity without quantifying its costs can provide a misleading perspective on the feasibility of adaptation and acclimation.

4.1.1 Is Evolutionary Rescue Real?

Evolutionary rescue is not a “silver bullet,” and the application of evolutionary theory to real populations and metapopulations is far from straightforward. For one thing, evolutionary rescue requires a sufficiently large population that a species is buffered against demographic and environmental stochasticity long enough for

higher-fitness phenotypes to become predominant.⁶⁵ Additional complications include, but are not limited to:

- **Initial environmental conditions.** Bell and Gonzalez showed that populations that begin at intermediate stress levels may react the slowest to environmental “deterioration,” producing a U-shaped curve in adaptive rescue.⁶⁶ They explain this as a product of two competing processes driving evolutionary rescue: as baseline stress increases, overall mutation rates decline, but the proportion of beneficial mutations (or, perhaps more accurately, the associated fitness differential) increases. Populations beginning in “mildly stressful conditions” may simply be at the low point of both processes. Bell and Gonzalez similarly show that populations with a history of minor environmental deterioration have a much greater probability of evolutionary rescue in a fast-changing environment.
- **The velocity of environmental change.** As Chevin *et al.*’s model highlights, environmental changes that are too rapid almost invariably drive species to extinction, when selection simply cannot operate fast enough to keep pace; this finding is readily confirmed in environmental settings. Rapid environmental changes can also functionally reduce mutation rates at a population scale. A study of *E. coli* by Lindsey *et al.* showed that “The evolutionary trajectory of a population evolving under conditions of strong selection and weak mutation can be envisioned as a series of steps between genotypes differing by a single mutation,” and some “priming mutations” may be necessary to arrive at further genotypic combinations with substantially higher fitness.⁶⁷ Consequently, if environmental changes are too rapid, higher fitness genotypes may be “evolutionary inaccessible.”
- **Dispersal rates and metapopulation connectivity.** Simulated metapopulation models by Schiffers *et al.* showed that higher dispersal rates can severely limit the propensity of populations to experience local adaptation, especially in a heterogeneous environment (a phenomenon they refer to as “genetic swamping”), and thereby potentially limit evolutionary rescue.⁶⁸ However, for an entire species to persist, intermediate (local) dispersal may be necessary to allow adaptive mutations to spread, a finding shown experimentally by Bell and Gonzalez.
- **Linkage disequilibrium.** Schiffers *et al.*’s study, which simulated genomes in an “allelic simulation model,” produced an unusual result suggesting that linkage between adaptive loci may not actually increase the rate of adaptation. The interaction this could have with the “priming mutation” process is complex and poorly explored in a theoretical context.

⁶⁵ Yi-Qi Hao *et al.*, “Evolutionary rescue can be impeded by temporary environmental amelioration,” *Ecology letters* 18, no. 9 (2015): 892–898.

⁶⁶ Graham Bell and Andrew Gonzalez, “Adaptation and evolutionary rescue in metapopulations experiencing environmental deterioration,” *Science* 332, no. 6035 (2011): 1327–1330.

⁶⁷ Haley A Lindsey *et al.*, “Evolutionary rescue from extinction is contingent on a lower rate of environmental change,” *Nature* 494, no. 7438 (2013): 463–467.

⁶⁸ Katja Schiffers *et al.*, “Limited evolutionary rescue of locally adapted populations facing climate change,” *Philosophical transactions of the Royal Society of London B: Biological sciences* 368, no. 1610 (2013): 20120083.

A final important consideration should be made with regard to what Schiffers *et al.* distinguish as *complete* vs. *partial evolutionary rescue*. In their models, they find that when adaptive traits originated but spread poorly (as a combination of linkage disequilibrium, habitat heterogeneity, and dispersal limitations), it substantially reduced population sizes and ultimately produced an “effective reduction in the suitable habitat niche.” This type of partial evolutionary rescue could be most common in real-world scenarios, where adaptation in larger populations experiencing the slowest rates of environmental change may allow persistence, but not maintain a species throughout its entire range, and may still be followed by a substantial reduction in overall habitat occupancy.

If current research on global climate change is any indication, this type of partial evolutionary rescue may ultimately be a poor buffer against extinction. Climate change may set the events of an extinction in motion, but research suggests that habitat loss from climate change is rarely the direct and solitary causal mechanism of an extinction.⁶⁹ Instead, climate change may reduce a population to small enough levels at which other mechanisms drive extinction. Small populations are especially susceptible to stochastic crashes in population size, and may also be especially susceptible to stochastic collapse due to other factors within-species (Allee effects in breeding, inbreeding) or from interactions with other species (competition, invasion, disease). Ultimately, the synergy between these drivers may produce a greater overall extinction risk that many modeling approaches might not directly quantify, but that could be most likely to drive species to extinction, and drive ecosystems into novel assemblages.⁷⁰

4.2 After Extinction: Lazarus Species, Romeo Errors, and the Rarest Birds in the World

The job of conservation biologists and extinction researchers is far from over after the extinction of a species. The *autoecology* of an extinct species (its basic biology, ecology, natural history, distribution and other species-level characteristics) often becomes a permanent unknown, assumed to be lost to the annals of history. But as statistical tools for ecological reconstruction become more sophisticated, researchers have the chance to explore basic questions about extinction in retrospect. In particular, the same body of theory that governs the timing of extinction in a declining population can be applied in a retrospective sense as well, to estimate the likely extinction date of a species. (Or, more formally, the estimation of the MTE from a given point can be used to pinpoint T_E , even with the same data, after extinction has already occurred.) These methods have been used both for ancient species

⁶⁹ Barry W Brook, Navjot S Sodhi, and Corey JA Bradshaw, “Synergies among extinction drivers under global change,” *Trends in ecology & evolution* 23, no. 8 (2008): 453–460.

⁷⁰ Lewis J Bartlett et al., “Synergistic impacts of habitat loss and fragmentation on model ecosystems,” in *Proc. R. Soc. B*, vol. 283, 1839 (The Royal Society, 2016), 20161027.

like the megalodon,⁷¹ and for more recent extinctions like that of the dodo⁷² or the thylacine.⁷³ But perhaps most interestingly that theory can be applied when the uncertainty bounds on T_E contain the present date, meaning that the extinction of a species is not taken as a certain part of history. Even ancient “Lazarus species” can be rediscovered, like the coelacanth, believed to have gone extinct 66 million years ago but rediscovered in the last century. How can we confidently say the coelacanth continues to exist, but the megalodon is likely to never be rediscovered?

4.2.1 Statistical Methods for the Sighting Record

Once a species is suspected to be extinct, at what point do we stop looking for them? With limited resources for conservation, trying to find and conserve a species that is no longer around wastes resources better used elsewhere; but making a Type I error and assuming a species is falsely extinct (and abandoning conservation efforts) can lead to a “Romeo Error,” whereby giving up on the species can lead to actual extinction.⁷⁴ Since 1889, 351 species thought to be extinct have been “rediscovered,”⁷⁵ highlighting just how big of a problem this may be. In order to answer these questions, determining the probability that a species is still extant, despite a lack of recent sightings, is an important step in making evidence-based decisions conservation managers must make about allocating resources.

But how do we determine the likelihood that a species is extinct? How long does it have to be since the last time an individual was seen before we can say, with some certainty, that the species is gone? The most obvious step is to assemble all available evidence from when the species was around. The first place to look is in the specimen record, which conventionally acts as the “gold-standard” of evidence. However, other data can be brought to bear, including observations, photos, and audio recordings. All these forms of evidence are collectively referred to as *sightings*. For a dataset of sightings $\mathbf{t} = (t_1, \dots, t_n)$, perhaps the simplest approach is to wait at least as long as the last interval during which the species was apparently absent before declaring a species extinct. One could formalize the estimate of the extinction date, T_E , as:

$$\widehat{T}_E = t_n + (t_n - t_{n-1})$$

⁷¹ Catalina Pimiento and Christopher F Clements, “When did Carcharocles megalodon become extinct? A new analysis of the fossil record,” *PloS one* 9, no. 10 (2014): e111086.

⁷² David L Roberts and Andrew R Solow, “Flightless birds: when did the dodo become extinct?,” *Nature* 426, no. 6964 (2003): 245–245.

⁷³ Colin J Carlson, Alexander L Bond, and Kevin R Burgio, “Estimating the extinction date of the thylacine with mixed certainty data,” *Conservation Biology* 32, no. 2 (2018): 477–483.

⁷⁴ NJ Collar, “Extinction by assumption; or, the Romeo Error on Cebu,” *Oryx* 32, no. 4 (1998): 239–244.

⁷⁵ Brett R Scheffers et al., “The world’s rediscovered species: back from the brink?,” *PloS one* 6, no. 7 (2011): e22531.

This approach, formalized by Robson and Whitlock,⁷⁶ is accompanied by a $(1 - \alpha)\%$ confidence interval with a lower bound at the last sighting t_n and the upper bound

$$T_E^u = t_n + \frac{1 - \alpha}{\alpha}(t_n - t_{n-1})$$

and accompanying p -value for testing the hypothesis that the species is extinct at the current time T :

$$p = \frac{t_n - t_{n-1}}{T - t_{n-1}}$$

The reasoning behind this method is fairly sound: if a large gap exists between the last two sightings, conservation biologists should wait at least that long before pronouncing a species certain to be extinct. But this estimator is also severely conservative, and has very limited theoretical grounding.

In 1993, Andrew Solow developed a more explicitly probabilistic approach,⁷⁷ which assumes sightings are generated by a random process with by a fixed sighting rate m that becomes 0 at T_E , the true date of extinction. The probability of the data conditional on a current time T and an extinction date T_E , is

$$P(T_n \leq t_n | T_E \geq T) = (t_n/T)^n$$

In that light, Solow says, hypothesis testing is easy: against the null hypothesis that extinction has yet to happen (i.e., $T_E > T$), we can test the alternate hypothesis that the species is extinct ($T_E < T$). For a given last sighting at T_N , we can provide a p -value for the test with desired significance level α equivalent to

$$P(T_N \leq \alpha^{1/n}T | T_E < T) = \alpha(T/T_E)^n$$

for values of $\alpha^{1/n}T < T_E < T$; for values of T_E lower than or equal to that critical value $\alpha^{1/n}T$, the value of that P is equal to 1 and the null hypothesis is rejected with full certainty. Solow explains, by way of example, that with 10 sightings and 95% confidence, the critical value of T_E/T is 0.74, and so the null hypothesis is sure to be rejected (extinction is confidently confirmed) if the true extinction date occurs within the first 74% of the $(0, T)$ window. Based on this approach, the maximum likelihood estimate \widehat{T}_E would be t_n , but this is clearly biased, and performs poorly as an estimation method. Instead, he suggests an alternate non-parametric estimator:⁷⁸

$$\widehat{T}_E = \frac{n+1}{n}t_n$$

And, in addition, a $1 - \alpha$ upper confidence interval bound:

⁷⁶ DS Robson and JH Whitlock, "Estimation of a truncation point," *Biometrika* 51, nos. 1/2 (1964): 33–39.

⁷⁷ Andrew R Solow, "Inferring extinction from sighting data," *Ecology* 74, no. 3 (1993): 962–964.

⁷⁸ Andrew R Solow, "Inferring extinction from a sighting record," *Mathematical biosciences* 195, no. 1 (2005): 47–55.

$$T_E^u = t_n / \alpha^{1/n}$$

Solow also proposed a foundational Bayesian approach based on the likelihood of a given sighting rate m would generate an observed density of data. Hypothesis testing in the Bayesian format where the likelihood of the sighting data given H_0 (the species is extant at time T) is

$$\int_0^\infty m^n e^{-mT} dP(m)$$

and given H_A (the species went extinct at T_E) is

$$\int_0^\infty m^n e^{-mT_E} dP(m)$$

From those and other assumptions, he derived the Bayes factor for the hypothesis test:

$$B(t) = \frac{n-1}{(T/t_n)^{n-1} - 1}$$

In Bayesian statistics, the Bayes factor is used to express the relative support between these two hypotheses. It bypasses the problem of setting a prior on the data or, in fact, of the two hypotheses; and instead, just expresses the posterior:prior odds of H_0 . If H_0 is small ($B \ll 1$), that suggests there is strong evidence against the null hypothesis.

Do these approaches make sense? If an extinction happens abruptly on the scale of sightings data (say, an epidemic wipes a species out within a year), then sighting rates might remain relatively constant throughout the sighting record. Similarly, applying this method to paleontological records may make sense, as prior information about variation in specimen preservation might be limited (and so a constant rate parameter is the best possible prior). But there are also a number of situations where the constant sighting rate m simply does not suffice. Lessons from population ecology remind us that extinction is, at its most fundamental scale, a process of declining abundance. If sightings are proportional to abundance (which they generally are), replacing m with a non-constant function has the potential to sharply refine the process of extinction date estimation.

Two additional methods have been suggested by Solow to account for the changing rate of sightings. Both assume that sightings are a declining process, which will make at least some of the above estimators prone to Type I errors. The first method assumes sighting rates decline exponentially at a rate β , so that the sighting density for $0 \leq t \leq T_E$ can be expressed as:

$$f(t) = \frac{\beta e^{-\beta t}}{1 - e^{-\beta T_E}},$$

In this model, if we express

$$s = \sum_{i=1}^n t_i$$

and k is the integer part of s/t_n (which we can write as $\lfloor s/t_n \rfloor$), then the estimated extinction date can be given as:

$$\widehat{T}_E = t_n + \frac{\sum_{i=0}^k (-1)^i \binom{n}{i} (s - it_n)^{n-1}}{n(n-1) \sum_{i=0}^{k-1} (-1)^i \binom{n-1}{i} (s - (i+1)t_n)^{n-2}}$$

The p -value is given as $p = F(t_n)/F(T)$ where

$$F(x) = 1 - \sum_{i=1}^{\lfloor s/x \rfloor} (-1)^{i-1} \binom{n}{i} \left(1 - \frac{ix}{s}\right)^{n-1}$$

The confidence interval can be determined using this p -value. However, computationally, that upper bound does not always exist—a major problem with this method.

In contrast, the second and far more complex method, implemented by Roberts and Solow in their 2003 study of the dodo,⁷⁹ accounts for the fact that the last few sightings of the species should, in most circumstances, follow a Weibull distribution. The method, *optimal linear estimation* (OLE), estimates T_E through linear algebra, with

$$\begin{aligned} T_E &= \sum_{i=1}^k w_i t_{n-i+1} \\ w &= (e' \Lambda^{-1} e)^{-1} \Lambda^{-1} e \\ e &= \begin{pmatrix} 1 \\ \vdots \\ \vdots \\ \vdots \\ 1 \end{pmatrix} \quad (\text{dimension } k) \end{aligned}$$

and Λ is a k by k matrix, for which

$$\begin{aligned} \Lambda_{ij} &= \frac{\Gamma(2\hat{v} + i) \Gamma(\hat{v} + j)}{\Gamma(\hat{v} + i) \Gamma(j)} \\ \hat{v} &= \frac{1}{k-1} \sum_{i=1}^{k-2} \ln \frac{t_n - t_{n-k+1}}{t_n - t_{i+1}} \end{aligned}$$

The OLE's upper 95% confidence interval is given by:

$$\begin{aligned} \widehat{T}_{ci}^u &= T_n + \frac{T_n - T_{n-k+1}}{c(\alpha) - 1} \\ c(\alpha) &= \left(\frac{-\log(\alpha/2)}{k} \right)^{-\hat{v}} \end{aligned}$$

⁷⁹ Roberts and Solow, "Flightless birds: when did the dodo become extinct?"

The OLE method has been recorded as one of the most successful methods available for predicting extinction,⁸⁰ and has the added bonus of being adjustable through sensitivity analysis to examine how different extent of sighting data changes the overall estimate.

4.2.2 Case Study: Spix's Macaw

Perhaps the most fruitful body of research concerning extinction date estimation has been within ornithology, where data on the last sightings of rare species is often more available than for other groups, due to tremendous global interest in bird sightings and observation by non-scientists. The most popular methods for sighting date research have often been developed in association with data on notable extinct birds, including the dodo⁸¹ and the ivory-billed woodpecker.⁸² In fact, one of the most expansive reviews of sighting date estimators, conducted by Elphick and colleagues, estimated the extinction date of 38 extinct or near-extinct birds from North America (including Hawaii, a hotspot of bird extinction).⁸³ But for rarer birds around the world, basic data on their extinction may be somewhat more lacking.

One such bird, the Spix's macaw (*Cyanopsitta spixii*) has been called "the world's rarest bird"⁸⁴ and has been the subject of two popular animated movies (Rio and Rio 2). Currently, Spix's macaw is considered critically endangered (possibly extinct in the wild) by the IUCN, with a small number of captive individuals (~130) found around the world. Not seen in the wild since 2000, a video of a Spix's macaw in Brazil made headlines in 2016. The video was subsequently examined by ornithologists and the consensus that the bird was, in fact, a Spix's macaw, though many believe it was likely an escaped captive bird.

Sightings of the Spix's macaw are sporadic, and after the first known specimen being shot in 1819 by Johann Baptist Ritter von Spix (though he believed the bird to be a Hyacinth Macaw), it was not recorded again until a wild-caught individual was procured by the Zoological Society of London in 1878. Collecting sighting records of the Spix's macaw relies mostly on data from trappers/poachers and inferring data from captive individuals. Given the illicit nature of wildlife poaching, better data may exist in the husbandry records of the wild-caught individuals currently in captivity, but those data are not freely available. Verifiable observations are few and far between, as this species was not subject to any intensive study or searches until the mid-1980s, when only a handful of individuals were found and, of those remaining, most were caught by poachers.

⁸⁰ Christopher F Clements et al., "Experimentally testing the accuracy of an extinction estimator: Solow's optimal linear estimation model," *Journal of Animal Ecology* 82, no. 2 (2013): 345–354.

⁸¹ Roberts and Solow, "Flightless birds: when did the dodo become extinct?"

⁸² Andrew R Solow and Andrew R Beet, "On uncertain sightings and inference about extinction," *Conservation Biology* 28, no. 4 (2014): 1119–1123.

⁸³ Chris S Elphick, David L Roberts, and J Michael Reed, "Estimated dates of recent extinctions for North American and Hawaiian birds," *Biological Conservation* 143, no. 3 (2010): 617–624.

⁸⁴ Tony Juniper, *Spix's Macaw: the race to save the world's rarest bird* (Simon / Schuster, 2004).




Fig. 2 Estimates of likely extinction date of the Spix's macaw based on extinction estimating equations in Solow (1993). The lines represent the estimated probability the species is extant each year; the blue line is the results using physical evidence only (specimens / wild-caught individuals), the orange line for uncontroversial sightings *and* physical evidence, and the green line is the results for all sightings, including controversial. The dotted line is a significance level of 0.05. Once the probability drops below this level, the species is considered likely extinct.

For this case study, we collected sighting and specimen data from GBIF (Global Biodiversity Information Facility; www.gbif.org) and Juniper's authoritative book on Spix's macaw. We found physical evidence (specimens / wild-caught captive birds) for sightings in the years: 1819, 1878, 1884, 1901, 1928, 1954, 1977, 1984, 1985, 1986, and 1987. Due to their rarity and the demand for them, we assumed individuals were caught in the wild the same year they were procured by the receiving institution / zoo. We considered all observations of the Spix's macaw reported in Juniper's book to be as uncontroversial as physical specimens, as there aren't many and these few have been rigorously scrutinized: 1903, 1927, 1974, 1989, 1990, and 2000. Our only controversial sighting is the recent video taken in 2016, which we omit from the model. By eliminating the controversial sighting (in analyses 1 and 2), we inherently test a methodological question: would extinction date estimators have pronounced the apparently-extant species dead?

Our analysis was conducted using the beta version of R package *sExtinct*, which allows a handful of different extinction analyses to be implemented (and we encourage prospective users to test the demos available with the package).⁸⁵ Our analysis uses two of the most common methods. First, we used the original Solow maximum likelihood approach, plotting the probability of persistence in Figure 2. The maximum likelihood estimates are given in that method as:

- Specimens only (1819-1987): $T_E = 2040$
- Uncontroversial sightings (1819-2000): $T_E = 2035$
- All sightings (1819-2016): $T_E = 2052$

⁸⁵ C Clements, *sExtinct: Calculates the historic date of extinction given a series of sighting events*. R package version 1.1, 2013.

That method suggests, even with the most limited dataset, that the species still appears to exist. In contrast, the OLE method tells a different story:

- Specimens only: $T_E = 1988$ (95% CI: 1987 - 2006)
- Uncontroversial sightings: $T_E = 2002$ (95% CI: 2000 - 2018)
- All sightings: $T_E = 2021$ (95% CI: 2016 - 2045)

All things considered, both analyses would have suggested at least a marginal chance the 2016 sighting may have been legitimate, and there is a possibility that a wild population of Spix's macaws may be out there, yet undiscovered in the Amazon rainforest. Ultimately, the 2016 sighting was determined by ornithologists and conservation experts to most likely be an escaped bird. With no new sightings in recent memory, and additional modeling work further refining these estimates, most experts now consider Spix's macaw extinct in the wild.⁸⁶ Hope for the species exists, with significant interest in captive breeding programs.

4.2.3 Hope Springs Eternal: Addressing False Sightings

Consider the plight of the ivory-billed woodpecker (*Campephilus principalis*), a charismatic and iconic part of the North American fauna. The ivory-billed woodpecker's decline was gradual, and unlike its gregarious and easily-spotted compatriots (such as the passenger pigeon, *Ectopistes migratorius*, or the Carolina parakeet, *Conuropsis carolinensis*, both extinct in a similar time period), sightings of the woodpecker were already rare previous to its decline. So while the bird's last "credible" sighting was in 1944, the precise date of its extinction remains controversial, and some believe the bird still exists based on unverified observations as recent as 2004 (with audiovisual evidence reviewed in a highly controversial 2005 paper in *Science*⁸⁷). These controversial observations led to one of the most costly surveys in history, yet yielded no new evidence. In some circles, the search continues; in 2016, two ornithologists—Martjan Lammertink and Tim Gallagher—traveled through Cuba searching for remaining populations of the elusive woodpecker. Was Lammertink and Gallagher's search justified from a statistical standpoint? And perhaps, more importantly, how can we address the problem of inaccurate sightings?

Not all sightings are created equal. Holding a dead body of an individual of the species constitutes good evidence the species was present the year the specimen was collected; but if some person claims they saw an extremely rare species with no corroborating evidence, they may have misidentified the individual, or in some cases, even lied, meaning that this sighting could be invalid. Roberts *et al.* found that extinction date estimators are sensitive to the data used, and can, unsurprisingly, lead

⁸⁶ Stuart HM Butchart *et al.*, "Which bird species have gone extinct? A novel quantitative classification approach," *Biological Conservation* 227 (2018): 9–18.

⁸⁷ John W Fitzpatrick *et al.*, "Ivory-billed Woodpecker (*Campephilus principalis*) persists in continental North America," *Science* 308, no. 5727 (2005): 1460–1462.

to very different estimates of extinction dates.⁸⁸ They partitioned sighting data into three categories: 1) physical evidence, 2) independent expert opinion, and 3) controversial sightings in order of certainty. They found that adding independently-verified observations to the analysis can sometimes lead to earlier predicted extinction times, since the “gaps” within the sighting record are closed up, whereas, by nature, later controversial sightings, if treated as legitimate (i.e., on par with physical evidence), can greatly push the estimates of extinction to later years.

To account for this uncertainty, a few approaches have been proposed recently. These approaches largely expand on Solow’s 1993 Bayesian equation above, modified to consider multiple levels of uncertainty in the data.⁸⁹ In the most advanced of these models, valid and invalid sightings are generated by separate Poisson processes. If valid sightings occur at rate Λ and invalid sightings at Θ , the proportion of valid sightings is

$$\Omega = \frac{\Lambda}{\Lambda + \Theta}$$

The sightings data can be split into certain data (t_c , with n_c certain records) and uncertain data (t_u , with n_u uncertain sightings, and $n_u(\hat{T}_E)$ sightings before \hat{T}_E). The conditional likelihood of the data is that of the two datasets multiplied:

$$\begin{aligned} p(t|\hat{T}_E) &= p(t_c|\hat{T}_E)p(t_u|\hat{T}_E) \\ p(t_c|\hat{T}_E) &= \frac{(n_c - 1)!}{(\hat{T}_E)^{n_c}} \\ p(t_u|\hat{T}_E) &= \int_0^1 \omega^{-n_u} (1 - \omega)^{n_u - n_u(\hat{T}_E)} \left(\hat{T}_E + \frac{1 - \omega}{\omega} T \right)^{-n_u} d\omega \end{aligned}$$

where ω is a dummy variable representing the certainty rate Ω . Bayes’ theorem gives the probability the species went extinct in the observation interval $(0, T]$, an event E with probability

$$p(E|t) = \frac{p(t|E)p(E)}{p(t)} = \frac{p(t|E)p(E)}{p(t|E)p(E) + p(t|\bar{E})(1 - p(E))}$$

The conditional likelihood of the data is

$$p(t|E) = p(t|\hat{T}_E)p(\hat{T}_E)$$

and conversely $p(t|\bar{E})$ is evaluated using the same function but replacing \hat{T}_E with T . The prior probability of the extinction date $p(\hat{T}_E)$ is a key part of successfully

⁸⁸ David L Roberts, Chris S Elphick, and J Michael Reed, “Identifying anomalous reports of putatively extinct species and why it matters,” *Conservation Biology* 24, no. 1 (2010): 189–196.

⁸⁹ Solow and Beet, “On uncertain sightings and inference about extinction”; Tamsin E Lee et al., “Inferring extinctions from sighting records of variable reliability,” *Journal of applied ecology* 51, no. 1 (2014): 251–258; CJ Thompson et al., “Inferring extinction risks from sighting records,” *Journal of theoretical biology* 338 (2013): 16–22.

implementing Bayesian analyses, and has a significant effect on the estimated extinction date.⁹⁰ Solow & Beet (2014) suggest three possible priors: uniform, linear, or exponential decline after the last certain sighting.

As previously noted, setting a prior probability for $p(E)$ is even more challenging. Instead of explicitly calculating the probability the species is still in existence, one can simply calculate the Bayes factor

$$B(t) = \frac{p(t|E)}{p(t|\bar{E})}$$

This bypasses the problem of assigning a prior probability of extinction $p(E)$, which is often unsatisfyingly solved by assigning an equal probability of extant or extinct ($p(E) = 0.5$). Ecological hypotheses can still be easily tested via the Bayes factor, though, as in a recent study that re-evaluated avian extinctions on Tristan de Cunha, and found that two of three species considered were probably extinct before the arrival of rats (the conventionally-assumed cause of their extinction).⁹¹

Solow & Beet's model is one of a handful of models that all use Bayesian approaches to estimate the extinction date, and test whether a species is extant. For an overview of the assumptions and relative strengths of these approaches, see Boakes *et al.*⁹² We note that while some more complex methods exist that, for instance, assign different intermediate levels of certainty to different kinds of evidence,⁹³ this may ultimately be superfluous. In many cases, the expenditure and effort required to obtain expert opinions may only have a marginal benefit, contributing little extra certainty to the models.⁹⁴ Consequently, the choice of model should depend on the available data, the operational power of any given study, and the degree of certainty needed for decision making.

4.2.4 The Ivory-Billed Woodpecker, and the Hunt for More Lazarus Species

To briefly reconsider Lammertink and Gallagher's continuing search for the ivory-billed woodpecker: regardless of how the sighting record for the ivory-billed woodpecker is analyzed, all indications point to an extremely low likelihood that the species is extant.⁹⁵ In the work of Elphick *et al.*, estimates based on physical evidence suggested a T_E of 1941 (upper 95% CI: 1945) and including expert opinion

⁹⁰ Andrew R Solow, "On the prior distribution of extinction time," *Biology Letters* 12, no. 6 (2016): 20160089.

⁹¹ Alexander L Bond, Colin J Carlson, and Kevin R Burgio, "Local extinctions of insular avifauna on the most remote inhabited island in the world," *Journal of Ornithology*, 2018, 1–12.

⁹² Elizabeth H Boakes, Tracy M Rout, and Ben Collen, "Inferring species extinction: the use of sighting records," *Methods in Ecology and Evolution* 6, no. 6 (2015): 678–687.

⁹³ Lee *et al.*, "Inferring extinctions from sighting records of variable reliability."

⁹⁴ Tamsin E Lee, Clive Bowman, and David L Roberts, "Are extinction opinions extinct?," *PeerJ* 5 (2017): e3663.

⁹⁵ Elphick, Roberts, and Reed, "Estimated dates of recent extinctions for North American and Hawaiian birds"; Nicholas J Gotelli *et al.*, "Specimen-Based Modeling, Stopping Rules, and the

sightings only moves T_E towards 1945 (upper 95% CI: 1948). Solow & Beet reanalyzed this problem with their Bayesian models that differentiate between certain and uncertain sightings. With a uniform prior on T_E over 1897-2010, if valid and invalid sightings are treated as generated from the same Poisson process, the calculated Bayes factor of 0.13 strongly suggests the species persists; but in the model that treats the processes as separate, which they suggest is the more accurate and appropriate one, the Bayes factor of 4×10^6 indicates almost no chance the species might be extant. In summation, the hard evidence available to modelers casts serious doubts on the validity of the species' "rediscovery" in 2004,⁹⁶ and finds little justification for the subsequent, costly search to find more conclusive evidence of the ivory-billed woodpecker's existence.

Some argue the search continues as long as hope does, and even to the most critically minded statistical ecologists, sighting models may not be an adequate source of information. In a recent controversy, a paper using the Solow & Beet model suggested astronomically low odds that the thylacine (*Thylacinus cynocephalus*) might be rediscovered.⁹⁷ Other scientists raised a number of issues with the methodology: spatial and temporal heterogeneity in sighting rates are ignored by Solow & Beet's model, there may have been active disincentives to report sightings after the species became protected, and perhaps the species' behavior had changed to become more cryptic.⁹⁸ Whether or not these factors change the results of the model, these challenges speak to a broader tension in conservation: with other species potentially within the bounds of rescue, the resources of conservation organizations might be better devoted to saving those species than to chasing the ghosts of woodpeckers or thylacines past.⁹⁹ But this is far from universally taken as true, and is a value judgment, not a statistical one.

In cases where rediscovery seems legitimately possible, statistical ecology still has an important role to play in those efforts. Once it is determined that there is an acceptable level of probability that a species is extant, one possible way to further leverage the data collected would be use the data to build species distribution models (SDMs) to aid in the search and rescue effort. In basic terms, SDMs use information about the conditions where a species has occurred (and where it has not occurred) to determine the realized ecological niche of the species. This niche can be projected onto geographic space to help identify areas that appear highly

Extinction of the Ivory-Billed Woodpecker," *Conservation Biology* 26, no. 1 (2012): 47–56; Solow and Beet, "On uncertain sightings and inference about extinction."

⁹⁶ David A Sibley et al., "Ivory-billed or pileated woodpecker?," *Science* 315, no. 5818 (2007): 1495–1496.

⁹⁷ Carlson, Bond, and Burgio, "Estimating the extinction date of the thylacine with mixed certainty data."

⁹⁸ Barry W. Brook et al., "Deficiencies in estimating the extinction date of the thylacine with mixed certainty data," *Conservation Biology*, doi:10.1111/cobi.13186.

⁹⁹ Colin J Carlson, Alexander L Bond, and Kevin R Burgio, "Reevaluating sighting models and moving beyond them to test and contextualize the extinction of the thylacine," *Conservation Biology*, 2018,

suitable for the species but perhaps have not been searched yet.¹⁰⁰ This approach has been successful in identifying new populations of threatened species (e.g.¹⁰¹), with the author identifying new populations of four of the eight rare plant species in the study. While SDMs are commonly used in a variety of different ecological and conservation applications, there is a deep literature on comparisons of SDM methods (see Qiao *et al.*¹⁰² for an overview), so caution must be exercised in selecting which methods are best for the available occurrence and environmental data. This approach—of determining the probability a species is still extant and using SDMs to identify the areas they are most likely to be—may provide a way forward for conservation agencies for making cost-effective decisions of which species to pursue and where to look for them. And with recent work advancing the idea of spatial extinction date estimators, to be used in combination with SDMs, significant advances in theory and practice are hopefully on the horizon.¹⁰³

5 The Community Scale and Beyond

Suppose that, in a twisted experiment motivated by an ecology-related childhood trauma, a mad scientist was developing a scheme to reduce global biodiversity to one half of the Earth's total species. Hunting, fishing and poaching could achieve that goal slowly, but would be particularly inefficient for eradicating insects; and while a generalist disease might help eradicate a handful of mammals or a sizeable fraction of amphibians, the majority of species would still remain. But perhaps realizing that habitat loss might be the most efficient tool to destruction, that scientist might cut the Gordian knot, by simply bisecting the Earth and destroying half. Would his plan come to fruition?

Our mad scientist's plan is riddled with flaws. If half of the species were endemic to each half of the Earth with no overlap, his plan would succeed. But a handful of species in any clade of life are globally-cosmopolitan; and no matter how his plan was executed, the handful of species occurring on both halves of the Earth would leave him with far, far more than half the species he started with.

With renewed vigor, the mad scientist sets out on a newly ambitious project: what percentage scorched earth would be required to achieve his goal? He begins by counting every species on his sidewalk block, then in his neighborhood, and up

¹⁰⁰ Jonathon C Dunn et al., "Mapping the potential distribution of the Critically Endangered Himalayan Quail *Ophrysia superciliosa* using proxy species and species distribution modelling," *Bird Conservation International* 25, no. 4 (2015): 466–478.

¹⁰¹ JL McCune, "Species distribution models predict rare species occurrences despite significant effects of landscape context," *Journal of Applied Ecology* 53, no. 6 (2016): 1871–1879.

¹⁰² Huijie Qiao, Jorge Soberón, and Andrew Townsend Peterson, "No silver bullets in correlative ecological niche modelling: insights from testing among many potential algorithms for niche estimation," *Methods in Ecology and Evolution* 6, no. 10 (2015): 1126–1136.

¹⁰³ Colin J Carlson et al., "Spatial extinction date estimation: a novel method for reconstructing spatiotemporal patterns of extinction and identifying potential zones of rediscovery," *bioRxiv*, 2018, 279679.

to bigger scales. With enough grant funding and undergraduate assistants, he has eventually covered a measly 6.25% of the Earth when he realizes he has counted half of Earth's species. To enact his master plan, he's tasked with destroying the remaining 93.75%. Going by land area alone (his grudges, we suppose, do not extend to the ocean), he only needs preserve 3.6 million square miles of land - roughly (conveniently?) the land area of the United States.

The process our nationalist, isolationist villain has enacted is the empirical construction of the species-area relationship (SAR), one of the oldest and most powerful scaling laws in macroecology. Because the synthesis of different factors at global scales is challenging, and habitat loss is one of the easiest extinction drivers to measure, the SAR gives us a powerful tool for approximating extinction rates - at the price of not knowing specifically which species will go extinct.

5.1 The Species Area Relationship

The biogeographer Olof Arrhenius began the process of formalizing the SAR in a classic 1921 paper in the *Journal of Ecology* titled "Species and Area."¹⁰⁴ In it he observed that, by expanding the area of focus, the number of species continues to increase at a diminishing rate (but, never reaching an asymptote¹⁰⁵). The canonical formula for the SAR has come to be called the Arrhenius SAR, and is formulated as

$$S = cA^z$$

where c is a constant fit to the data, and z is a slope conventionally treated as 0.25. The application of that formula to extinction rate estimation is relatively obvious: by changing the amount of area, we can change the number of species:

$$S' = c(A')^z$$

and calculate the number of extinctions

$$E(A') = S - S'$$

In our mad scientist's failed scheme, reducing the area of the Earth by half would leave us with far more than half the species:

$$\frac{S'}{S} = \left(\frac{0.5A}{A}\right)^{0.25} = (0.5)^{0.25} = 0.84$$

In a 2004 *Nature* paper that has become perhaps the most cited study on extinction since the millennium, a group of researchers led by Chris Thomas refined the global extinction rate estimate by analyzing species' habitat losses from climate change

¹⁰⁴ Olof Arrhenius, "Species and area," *Journal of Ecology* 9, no. 1 (1921): 95–99.

¹⁰⁵ Mark Williamson, Kevin J Gaston, and WM Lonsdale, "The species–area relationship does not have an asymptote!," *Journal of Biogeography* 28, no. 7 (2001): 827–830.

and applying the SAR. Their extinction-area relationship took three forms applied to n species, with a given A_i area per species before change, and A'_i subsequent to habitat loss:

$$E_1 = 1 - \left(\frac{\sum_{i \in (1,n)} A'_i}{\sum_{i \in (1,n)} A_i} \right)^{0.25}$$

$$E_2 = 1 - \left(\frac{1}{n} \sum_{i \in (1,n)} \frac{A'_i}{A_i} \right)^{0.25}$$

$$E_3 = \frac{1}{n} \sum_{i \in (1,n)} \left(1 - \left(\frac{A'_i}{A_i} \right)^{0.25} \right)$$

Using those three methods in combination with species distribution models, the authors estimated that 15-37% of species on Earth might face climate-driven extinction by 2050. That result is by far one of the most important ones produced in any study of extinction, and has supported a number of the most expansive conservation programs worldwide.

5.2 Refining and Reformulating the SAR

Like many “laws” of ecology, the conventional SAR has problems and pitfalls, and with the tremendous array of approaches developed to study it, it has even been called ecology’s “most protean pattern.”¹⁰⁶ Subsequent to the publication of Thomas *et al.*’s study, one of the most seminal debates in extinction research has centered around its conclusion that climate change is likely to act as the most consequential driver of the sixth mass extinction. Different approaches to the species area relationship, and comparable or derivative macroecological methods, have sprung up in the wake of Thomas’s work. Here, we review a few of the different approaches that can be used to predict extinction rates at the community level.

5.2.1 z : a Dynamic Scaling Property

The most immediate problem with applying the species area relationship is that the slope z , normally set to 0.25, is neither universal nor scale-independent. In part, this is because of two different constructions of the SAR. The slope of 0.25 derives from the experimental work of Macarthur and Wilson on island ecosystems, which applied the SAR to the richness of species on islands of different sizes. For islands

¹⁰⁶ Mark V Lomolino, “Ecology’s most general, yet protean pattern: The species-area relationship,” *Journal of Biogeography* 27, no. 1 (2000): 17–26.

(and for application of the island SAR to extinction), a slope of 0.25 is justified under a set of three (relatively common) circumstances delineated by Harte and Kitzes: “(i) total abundance in the new area A is proportional to area, (ii) individuals found in A are chosen by a random draw from all individuals in A_0 [the initial area], and (iii) the number of individuals of each species in A_0 follows a canonical lognormal abundance distribution.”¹⁰⁷

However, the continental “nested” SAR (constructed from nested areas on a continental scale) does not always follow the same property. This is in part because the conventionally-used SAR assumes self-similarity (or in more tangible terms, picking two patches of different area always yields a roughly-the-same-slope difference in species). As it turns out, self-similarity works within some sites but not others, and within the Western Ghats mountains of India alone, scaling up from vegetation sampling plots to broader scales brings z down from values closer to 0.5, down to values approaching 0.¹⁰⁸ Selecting an appropriate slope based on scale is an important part of appropriate use of the SAR to predict extinction rates, and as analyses approach the continental scale, the appropriateness of the SAR method decreases as z approaches zero.

5.2.2 An Alternate Approach Based on the Endemics Area Relationship

In the Thomas *et al.* study, the application of the species area relationship followed three methods, and while some explicitly predicted extinction risk at the scale of a single species, all rely on the prediction of reduced species richness based on habitat loss. In place of this indirect calculation of decreased richness, a more direct approach uses what is called the *endemics area relationship*, which calculates the number of endemic species restricted to a given area (all of which should be committed to extinction when the area is destroyed). As pointed out by He and Hubbell, the SAR and the EAR are not mirror curves except in a single special case when species are completely randomly distributed in space; else, the “forwards” and “backwards” methods of extinction calculation are not, they argue, comparable.¹⁰⁹

Prediction of extinction based on the EAR may be more appropriate for measuring the immediate effects of habitat loss, and is likely to better account for the “geometry of habitat clearing.”¹¹⁰ Storch *et al.*¹¹¹ developed an approach to the SAR and the EAR that scales the area by the mean geographic range size in the

¹⁰⁷ John Harte and Justin Kitzes, “The use and misuse of species-area relationships in predicting climate-driven extinction,” in *Saving a Million Species* (Springer, 2012), 73–86.

¹⁰⁸ John Harte, Adam B Smith, and David Storch, “Biodiversity scales from plots to biomes with a universal species–area curve,” *Ecology letters* 12, no. 8 (2009): 789–797.

¹⁰⁹ Fangliang He and Stephen P Hubbell, “Species-area relationships always overestimate extinction rates from habitat loss,” *Nature* 473, no. 7347 (2011): 368–371.

¹¹⁰ Ann P Kinzig and John Harte, “Implications of endemics–area relationships for estimates of species extinctions,” *Ecology* 81, no. 12 (2000): 3305–3311.

¹¹¹ David Storch, Petr Keil, and Walter Jetz, “Universal species-area and endemics-area relationships at continental scales,” *Nature* 488, no. 7409 (2012): 78–81.

focal clade/area and scales richness by the average number of species in that mean geographic range. When plotted, the SAR curves upwards while the EAR is roughly linear with a slope of 1 across most scales. Starting from basic knowledge about the average geographic range size of a given species, this result indicates that extinction from habitat loss can be predicted based on the EAR across scales fairly accurately.

5.2.3 An Alternate Approach Based on Maximum Entropy

Two “unifying” theories have dominated discussions about macroecology. The first is the unified neutral theory (UNT) of biogeography and ecology (proposed by Stephen Hubbell), which is beyond the scope of this chapter; the second is the maximum entropy theory of ecology (METE) proposed by John Harte. The METE deserves special mention here, due to a particular focus in the METE literature on improving the applicability of the SAR to extinction rate prediction. What differentiates both the UNT and the METE from more general conceptions of the SAR is the explicit treatment of species abundance as a component of community assembly. The theory of the METE is far too complex to encapsulate in this chapter (and an entire book by Harte exists for that purpose), but a few useful derivations are worth mentioning. One is the derivation by Kitze and Harte of an extinction probability that is applicable at the species scale¹¹² based on proportional area loss (A_0/A , shortened to β) and corresponding reduction in abundance (n from n_0) with a general probability distribution

$$P(n|n_0, A_0, A) = ce^{-\lambda n}$$

for which they provide rough approximations

$$c \approx \frac{1}{(An_0/A_0) + 1}$$

and

$$\lambda \approx \ln \left(1 + \frac{A_0}{An_0} \right)$$

Drawing on similar concepts from the pseudoextinction thresholds we discuss above in §4.1, they suggest the probability a remainder $r_c = n/n_0$ will be left after habitat loss is equivalent to

$$\text{Prob} \left[\frac{n}{n_0} > r_c \right] = \int_{r_c n_0}^{n_0} ce^{-\lambda n} dn = \frac{[n_0\beta/(1+n_0\beta)]^{r_c n_0} - [n_0\beta/(1+n_0\beta)]^{n_0}}{(1+n_0\beta) \ln(1+1/n_0\beta)}$$

¹¹² Harte and Kitze, “The use and misuse of species-area relationships in predicting climate-driven extinction.”

Given a starting population and a critical population size, analogous results can be derived for the Thomas *et al.* calculations; and higher level predictions can be made based on the distribution of abundances and critical abundances within the community.

In a subsequent publication,¹¹³ this *extinction area relationship* is extended even further to extrapolate a MaxEnt-based probability that a given number of species will remain after habitat loss. It assumes a logseries distribution ϕ of abundance for species with a mean μ_ϕ with a single shape parameter p

$$\phi(n_0) = \frac{-p^{n_0}}{\ln(1-p)n_0}$$

$$\mu_\phi = \frac{-p}{(1-p)\ln(1-p)}$$

They similarly propose an upper-truncated geometric species specific abundance distribution, which provides the probability n individuals remain in a fractionally reduced area a (β in their other notation) based on a shape parameter q :

$$\Pi(n|a, n_0) = \frac{(1-q)q^n}{1-q^{n_0+1}}$$

where q is solved implicitly based on a and n_0 from

$$an_0 = \frac{q}{1-q} - \frac{(n_0+1)q^{n_0+1}}{1-q^{n_0+1}}$$

The probability a species is found in area A after habitat loss follows a distribution g that takes the form

$$g(a, n_c) = \sum_{n_0=1}^{\infty} (1 - \Pi(n \leq n_c | a, n_0)) \phi(n_0)$$

which scales up to a community-level richness after area loss

$$p(S|S_0, g) = \binom{S_0}{S} g^S (1-g)^{S_0-S}$$

where

$$g(a, n_c, \mu_\phi) = \sum_{n_0=1}^{\infty} \left(1 - \frac{q^{n_c+1} - 1}{q^{n_0+1} - 1} \right) \left(\frac{-p^{n_0}}{n_0 \ln(1-p)} \right)$$

or if the pseudoextinction threshold is set to zero (i.e. no species has 0% survival odds until all individuals are dead) and area loss is severe, that expression can be reduced to eliminate the q term:

¹¹³ Justin Kitzes and John Harte, "Beyond the species-area relationship: improving macroecological extinction estimates," *Methods in Ecology and Evolution* 5, no. 1 (2014): 1-8.

$$g(a, n_c, \mu_\phi) = -\frac{a}{\ln(1-p)} \sum_{n_0=1}^{\infty} \frac{p^{n_0}}{an_0 + 1}$$

This METE approach thus provides a *probabilistic species area relationship* (PSAR) that can be used to provide not only an expected extinction rate under habitat loss but also a range of confidence. This becomes an especially important tool in a small community of only a few dozen species or fewer (or in communities with pervasive low abundance across species), where deviations from SAR-based predictions may be greater due to stochastic processes.

How does the PSAR scale up against the Thomas-SAR? It has a clear advantage in the prediction of individual species extinction risk (but correspondingly requires more data on abundance/demography that may be absent for many poorly-known taxa). Kitze and Harte provide two illustrations; first, assuming the normal slope of 0.25, the PSAR predicts a 44% chance of extinction for a species that loses 90% of its habitat. Second, if we assume a pseudoextinction threshold of 50 individuals, by comparison to the predictions of the PSAR, the Thomas-SAR approach underpredicts extinction risk if $n_0 \leq 1000$ but overpredicts if n_0 is greater—supporting the notion that the 15-37% extinction rates that Thomas *et al.* study predicted could be an overestimate.

5.2.4 Tying Up Loose Threads, Thinking Across Scales

The various different approaches to predicting extinction at the broadest scales have driven substantial controversy among different interpretations of macroecological theory. But one of the most important problems is that estimates of extinction from these methods are still poorly connected, by and large, to the rest of the extinction literature, and to the other types of models we discuss above. One of the most innovative and unusual approaches in the literature was presented by Rybicki and Hanski,¹¹⁴ who simulated a stochastic patch occupancy model (similar to those presented in §3.1) with spatially heterogeneous environmental conditions across patches. While their model incorporates the standard mainstays of an SPOM (colonization, extinction, a dispersal kernel), it also incorporates a phenotype and niche breadth that produce a Gaussian fitness function (much like the models in §4.1)

Tying together a number of the important ideas discussed above, the work of Rybicki and Hanski made several advances into new territory. For one, they make a semantic distinction between the endemics-area relationship (EAR, which they define as the $S = cA^z$ relationship applied to the area lost $A = a$) and the “remaining species-area relationship” (RAR),

$$S - S_{\text{loss}} = c(A_{\text{new}}/A)^z$$

¹¹⁴ Joel Rybicki and Ilkka Hanski, “Species–area relationships and extinctions caused by habitat loss and fragmentation,” *Ecology letters* 16, no. s1 (2013): 27–38.

The EAR and RAR, as two methods of calculating extinction risk, are not interchangeable or symmetric counterparts. Rybicki and Hanskii highlight a discrepancy between Storch *et al.*'s suggested EAR slope of roughly 1, and He and Hubbell's values which were a tenth smaller,¹¹⁵ which they suggest can be resolved by the fact that Storch fit the EAR while He and Hubbell were calculating the RAR; and their simulations agree with the results of He and Hubbell that the slope of the RAR may be half or less that of the SAR.

Their empirical approach to simulation leads to a valuable conclusion that stands in opposition to previous work. While Kinzig and Harte¹¹⁶ and He and Hubbell¹¹⁷ both strongly suggest that the SAR overestimates extinction risk; the results of Rybicki and Hanskii's simulations suggest that in the short term, the RAR underestimates extinction while the continental SAR ($z \approx 0.1$) is adequate. Their result ties the population scale to the community scale, as they attribute it to species' populations *outside* destroyed or fragmented habitat falling below critical thresholds and facing extinction despite the lack of total endemic extirpation. In the long term, they suggest, the island SAR ($z = 0.25$) may be the best predictor of total losses. Finally, they explore the difference between leaving a single patch of habitat and fragmenting habitat, and conclude all models underestimate extinction risk in scenarios of extreme fragmentation. To address that problem, they propose a modified species area relationship

$$S = cA^z e^{-b/\lambda_M}$$

where λ_M is the metapopulation capacity (see §3.1) and b is another scaling parameter like c and z . If n is the number of habitat fragments, they suggest, the metapopulation capacity scales linearly with A^3/n^2 , meaning that the *fragmented landscape species area relationship* (FL-SAR) can be expressed as:

$$S_{\text{new}}/S = (A_{\text{new}}/A)^2 e^{-bn^2/A^3}$$

While the data to fit such an expression might be challenging to collect (and so the FL-SAR may not be an immediately useful conservation planning tool), the FL-SAR provides an important and much needed link between the population and metapopulation processes we discuss above, and our broader understanding of the rate of extinction at landscape and community scales.

¹¹⁵ He and Hubbell, "Species-area relationships always overestimate extinction rates from habitat loss."

¹¹⁶ Kinzig and Harte, "Implications of endemics-area relationships for estimates of species extinctions."

¹¹⁷ He and Hubbell, "Species-area relationships always overestimate extinction rates from habitat loss."

6 Last Chance to See

What don't we know about extinction yet?

As predictive tools gain precision, our estimates of the extinction rates of well-known groups like mammals and birds also become more precise. But the majority of the world's species are not yet known; most animal diversity is harbored by insects or parasites (especially nematodes), and the vast majority of species in those groups are undiscovered or undescribed. Their extinction rates are just as poorly quantified as their diversity or the hotspots of their biodiversity. But some basic estimates suggest that 7% of the planet's invertebrates may have already gone extinct—at which rate evidence would suggest that 98% of extinctions on Earth are currently going undetected.¹¹⁸ It's also especially difficult to compare these extinction rates to historical baselines, because the fossil record for most invertebrates and other taxa are incomplete or nearly absent.

An especially poignant problem is the detection and estimation of coextinction rates—the secondary extinction of species dependent on others for their ecological niche—which Jared Diamond suggested in 1989 was one of the four horsemen of mass extinction (in his words, “overhunting, effects of introduced species, habitat destruction, and secondary ripple effects”).¹¹⁹ Among the most obvious candidates for coextinction are two main groups: pollinators (which can have a strict dependency on host plants) and endosymbionts (parasites and mutualists, which may exhibit strict specificity in their association with plant or animal hosts). While both groups are believed to be severely at risk of secondary extinction, quantifying their extinction rate can be challenging, as there is rarely a 1:1 correspondence between hosts and dependent species. An approach popularized by Koh simulates host extinctions in a random order and predicts the number of corresponding coextinctions from the *affiliation matrix*; by fitting a function to real affiliation matrices, Koh *et al.* found that if host specificity is 1:1 then the slope is linear, but when affiliates use a greater number of hosts, the coextinction function is concave upward:

$$\bar{A} = (0.35\bar{E} - 0.43)\bar{E} \ln \bar{s} + \bar{E}$$

where E gives primary extinction risk, A gives secondary extinction risk, and s is host specificity.¹²⁰ Subsequent work has shown that even though parasites and mutualists may experience a reduced rate of extinction from host switching, the majority of threatened species on Earth might still be mutualists and parasites (due to the tremendous diversity of such species, e.g. the estimated 300,000 species of helminth

¹¹⁸ Claire Régner et al., “Mass extinction in poorly known taxa,” *Proceedings of the National Academy of Sciences* 112, no. 25 (2015): 7761–7766.

¹¹⁹ Jared M Diamond, NP Ashmole, and PE Purves, “The present, past and future of human-caused extinctions [and discussion],” *Philosophical Transactions of the Royal Society B: Biological Sciences* 325, no. 1228 (1989): 469–477.

¹²⁰ Lian Pin Koh et al., “Species coextinctions and the biodiversity crisis,” *science* 305, no. 5690 (2004): 1632–1634.

alone¹²¹). One recent study using the Thomas species-area relationship approach estimated that, from the synergistic pressures of climate change and coextinction, up to one third of all helminth parasites might be threatened with extinction by 2070.¹²²

Most affiliate extinctions are poorly cataloged, if recorded at all,¹²³ and only limited conservation frameworks exist for their study. More data is needed on host-symbiont association networks to better inform the role that non-random structure in those networks might play in increasing or decreasing extinction rates; some work has suggested that species preferentially favor more stable host species, the underlying cause of a “paradox of missing coextinctions.”¹²⁴ Similarly, the potential for species to switch hosts and thereby avoid extinction is unknown, but likely mitigates global extinction risk. In parasitology, the Stockholm Paradigm suggests that host-parasite associations diversify in changing climates and environments as a function of (1) phenotypic plasticity, (2) trait integration and (3) phylogenetic conservatism of “latent potential,” which together produce a pattern of *ecological fitting* that might benefit parasites (and thereby other symbionts) in the face of the sixth mass extinction.¹²⁵ A more in-depth treatment of the theoretical ecology of ecological fitting can be found in the recent work of Araujo *et al.*¹²⁶

Is saving microbes and parasites from extinction a reasonable goal? Some argue that it is,¹²⁷ but others have recently suggested it’s “time to get real about conservation” and focus on our failure to adequately prevent catastrophic population crashes in megafauna like elephants.¹²⁸ Regardless of animal type or conservation status, the development of demographic theory and predictive modeling are our best options to understand and mitigate extinction risk in natural populations. One such advance deserving of special mention is the development of *early warning signals* of population collapse. This is a developing body of literature that is built around the fact that populations on the verge of collapse often produce detectable statistical signals¹²⁹ that, detected far enough in advance, might allow mitigation efforts and prevention of population collapse.

¹²¹ Andy Dobson *et al.*, “Homage to Linnaeus: How many parasites? How many hosts?,” *Proceedings of the National Academy of Sciences* 105, no. Supplement 1 (2008): 11482–11489.

¹²² Colin J Carlson *et al.*, “Parasite biodiversity faces extinction and redistribution in a changing climate,” *Science Advances* 3, no. 9 (2017): e1602422.

¹²³ Robert R Dunn *et al.*, “The sixth mass coextinction: are most endangered species parasites and mutualists?,” *Proceedings of the Royal Society of London B: Biological Sciences* 276, no. 1670 (2009): 3037–3045.

¹²⁴ Giovanni Strona, Paolo Galli, and Simone Fattorini, “Fish parasites resolve the paradox of missing coextinctions,” *Nature communications* 4 (2013): 1718.

¹²⁵ Daniel R Brooks and Eric P Hoberg, “How will global climate change affect parasite–host assemblages?,” *Trends in parasitology* 23, no. 12 (2007): 571–574.

¹²⁶ Sabrina BL Araujo *et al.*, “Understanding host-switching by ecological fitting,” *PloS one* 10, no. 10 (2015): e0139225.

¹²⁷ Eric R Dougherty *et al.*, “Paradigms for parasite conservation,” *Conservation Biology*, 2015,

¹²⁸ Aaron M Ellison, “It’s time to get real about conservation.,” *Nature* 538, no. 7624 (2016): 141.

¹²⁹ Marten Scheffer *et al.*, “Early-warning signals for critical transitions,” *Nature* 461, no. 7260 (2009): 53–59.

The majority of early warning signals for extinction currently rely on identifying *critical slowing down*, a process by which as the dominant eigenvalue of the system tends towards zero, populations return to equilibriums lower after perturbations, with increasing variance and temporal autocorrelation. Critical slowing down is often a sign of a dynamical system approaching a bifurcation point, which may sometimes indicate a non-catastrophic shift to an alternative stable state,¹³⁰ but more usefully, may indicate an impending extinction emerging from a shift into a sharp decline or into chaos. A foundational experiment by Drake & Griffen showed that critical slowing down can readily be detected from time series data for populations facing decreasing food availability up to eight generations before they reach the bifurcation point.¹³¹ Conventionally, this is done by taking a set of metrics such as the autocorrelation, coefficient of variation, or skewness (termed leading indicators), scaling and adding them into a single metric (composite indicators), and tracking them over time. Given any indicator w , the standardized statistic is calculated as a function of the mean and standard deviation of the time series to that point:

$$\hat{w}_t = \frac{w_t - \bar{w}_{1:t}}{\sigma(w_{1:t})}$$

When the test statistic \hat{w} passes a threshold level of deviation from the running average $\bar{w}_{1:t}$, such as 2σ , this can be taken as an early warning signal.¹³²

Early warning signals are far from a perfect tool. Most research has focused on detecting critical slowing down, but not all types of dynamical systems exhibit critical slowing down.¹³³ Even if critical slowing down is expected, these types of early warning signals are far from perfectly accurate. Ecological data, especially from population abundance estimates, often has a high signal:noise ratio,¹³⁴ to which early warning signals are still sensitive, leading to an expected mix of both Type I and II errors, depending on the quality and quantity of data. In some cases, this problem can be accommodated for by evaluating early warning signals as an iterative process over the time series, rather than taking the first warning as the only required evidence.¹³⁵

Simulation work has shown early warning signals to be fairly robust to incomplete sampling, but proportionally data intensive to a degree that may be impossible

¹³⁰ Sonia Kéfi et al., “Early warning signals of ecological transitions: methods for spatial patterns,” *PloS one* 9, no. 3 (2014): e92097.

¹³¹ John M Drake and Blaine D Griffen, “Early warning signals of extinction in deteriorating environments,” *Nature* 467, no. 7314 (2010): 456.

¹³² Christopher F Clements and Arpat Ozgul, “Including trait-based early warning signals helps predict population collapse,” *Nature communications* 7 (2016).

¹³³ Carl Boettiger, Noam Ross, and Alan Hastings, “Early warning signals: the charted and uncharted territories,” *Theoretical ecology* 6, no. 3 (2013): 255–264; Alan Hastings and Derin B Wysham, “Regime shifts in ecological systems can occur with no warning,” *Ecology letters* 13, no. 4 (2010): 464–472.

¹³⁴ Boettiger, Ross, and Hastings, “Early warning signals: the charted and uncharted territories.”

¹³⁵ Christopher F Clements et al., “Body size shifts and early warning signals precede the historic collapse of whale stocks,” *Nature ecology & evolution* 1, no. 7 (2017): 188.

to reach with available ecological time series data.¹³⁶ When population data is lacking or incomplete, trait-based data can be used as a proxy or a supplement, if traits such as body size are expected to correlate with fitness, and are responsive to ecological shifts.¹³⁷ One particularly interesting demonstration by Clements *et al.* showed that body size data could be used to predict the collapse of whale populations 10–40 years before the whaling industry depleted stocks.¹³⁸ Other cutting-edge work is attempting to scale the detection of early warning signals to the metapopulation level by developing spatial early warning signals,¹³⁹ which could be used to optimize reserve design and address the influence of dispersal, stochasticity, and local population dynamics on metapopulation persistence. Some work has even suggested critical slowing down could be used to identify tipping points of network collapse due to serial coextinctions in symbiont networks.¹⁴⁰

The pressure for more accurate, predictive tools will only grow in the next few decades of research. A recent review by Mark Urban surveyed studies of climate change-driven extinction risk and found that, despite the variation between different modeling methods and scopes, projected extinction rates are not only rising but one in six species might be imminently threatened with extinction.¹⁴¹ Similarly, in a study of roughly 1000 species of plants and animals, about half had experienced population extinctions driven by climate change.¹⁴² As extinction rates accelerate due to global change and we fully enter the sixth mass extinction, the need for better analytical and simulation tools—that produce precise estimates from limited data—will only grow. In light of the constant need to test, revise and re-test models of extinction, to a mathematically-trained ecologist or an ecologically-minded mathematician, this field of research is a critical opportunity to apply the principles of ecosystem science towards a high-impact and worthy goal.

Acknowledgements

Thanks to Alex Bond and Andrew Beet for helpful comments and conversations throughout the writing process.

¹³⁶ Christopher F Clements *et al.*, “Factors influencing the detectability of early warning signals of population collapse,” *The American Naturalist* 186, no. 1 (2015): 50–58.

¹³⁷ Clements and Ozgul, “Including trait-based early warning signals helps predict population collapse.”

¹³⁸ Clements *et al.*, “Body size shifts and early warning signals precede the historic collapse of whale stocks.”

¹³⁹ Kéfi *et al.*, “Early warning signals of ecological transitions: methods for spatial patterns.”

¹⁴⁰ Vasilis Dakos and Jordi Bascompte, “Critical slowing down as early warning for the onset of collapse in mutualistic communities,” *Proceedings of the National Academy of Sciences* 111, no. 49 (2014): 17546–17551.

¹⁴¹ Mark C Urban, “Accelerating extinction risk from climate change,” *Science* 348, no. 6234 (2015): 571–573.

¹⁴² John J Wiens, “Climate-related local extinctions are already widespread among plant and animal species,” *PLOS Biology* 14, no. 12 (2016): e2001104.