UC Santa Barbara

UC Santa Barbara Previously Published Works

Title
Estimating the magnitude of environmental stochasticity in survivorship data

Permalink
https://escholarship.org/uc/item/9kt8g5vj

Journal
Ecological Applications, 8(1)

Author
Kendall, Bruce E.

Publication Date
1998

DOI
10.1890/1051-0761(1998)008[0184:ETMOES]2.0.CO;2

Peer reviewed

ESTIMATING THE MAGNITUDE OF ENVIRONMENTAL STOCHASTICITY IN SURVIVORSHIP DATA

Bruce E. Kendall

*Department of Ecology and Evolutionary Biology, University of Arizona, Tucson, Arizona 85721 USA, and
National Center for Ecological Analysis and Synthesis,[1] University of California,
Santa Barbara, California 93106 USA*

Abstract. Small populations are often at risk of extinction through processes that are effectively stochastic. Prediction of this extinction risk requires that the observed temporal variation in demographic rates be accurately partitioned between demographic stochasticity (variation among individuals) and environmental stochasticity (variation among years, correlated across individuals). However, studies of population viability analysis that include both forms of stochasticity parameterize the magnitude of environmental stochasticity incorrectly (they overestimate it). I describe and evaluate tests (1) to determine whether all the year-to-year variation in observed survivorship can be explained by demographic stochasticity alone, and (2) if not, to estimate the magnitude of environmental stochasticity in survival. The first issue can be resolved with a G test. I used simulated data to show that this test has an appropriate type I error rate, unless the individual survival probability is either very low or very high. Small amounts of environmental stochasticity often are not detected by the G test (type II error), but the hypothesis of demographic stochasticity alone is consistently rejected when environmental stochasticity is large. In contrast, estimating the magnitude of environmental stochasticity requires explicit hypotheses about the nature of the underlying variation, but I provide a flexible framework in which many such hypotheses can be examined. In particular, I show, using simulated data, that if the temporal variability in individual survival probabilities is distributed according to a beta distribution, then the maximum likelihood estimate of the variance of the survival probability is biased, but in a consistent and correctable way. The estimate obtained with my method is usually superior to an estimate that assumes that all of the variation in the observed survivorship is due to environmental stochasticity. I show how to include deterministic sources of variability, such as density dependence, and how to apply different assumptions about the underlying environmental stochasticity. I illustrate these tests with data from a population of Acorn Woodpeckers (*Melanerpes formicivorus*). With these data, I can determine that there is strong environmental stochasticity in juvenile survival, whereas variation in adult survival can be explained either by density dependence or by weak environmental stochasticity.

Key words: Acorn Woodpecker; demographic stochasticity; density dependence; environmental stochasticity; Melanerpes formicivorus; *stochastic population modeling; survivorship.*

INTRODUCTION

A central problem in conservation biology is the understanding and prediction of the dynamics and persistence of small populations (Shaffer 1987). Conservation biologists frequently must assess the probability that a population will decline to a critically small size, and estimate the probability of subsequent population extinction if it does so (Gilpin and Soulé 1986, Shaffer 1990). When the number of individuals is small, variation in the demographic rates reduces the long-term growth rate of the population and can have drastic short-term effects (such as extinction). These effects reduce the relevance of the large body of deterministic population dynamic theory, and require that we employ stochastic theories.

Ecologists use the term ''stochasticity'' to refer to a variety of factors. In the context of population dynamics, several classes of factors affect the demographic processes of survival, reproduction, and dispersal: intrinsic variables whose values are either (1) unknown or (2) deliberately ignored (such as age, nutritional state, experience, or genotype); (3) intrinsic variables that are inherently unpredictable (such as ''accidents''); and (4) extrinsic variables that are inherently unpredictable (such as weather). The first three categories represent unknown variation among individuals at a given time, which is commonly called ''demographic stochasticity'' (Shaffer 1987, Durant and Harwood 1992); the fourth represents unknown variation over time (or space) that affects all individuals in the population in a similar way, and is called ''environmental stochasticity'' (Lacy 1993, Akçakaya and Baur 1996).

There usually will not be enough information to predict the fate of a given individual at a given time. The

best one can do in this situation is to estimate the "probability" of, for example, survival over a stated interval. This use of probabilities refers not to the process being inherently random, but to the fact that the investigator lacks relevant information. Furthermore, it often will not be possible to assess the variation among individuals, so the most parsimonious (although not necessarily the most accurate) assumption is that all individuals are identical, at a given time, in their demographic processes. Nevertheless, because individual demography is framed in probabilities rather than rates, this identity among individuals (as long as the individual fates are themselves uncorrelated) can lead to demographic variability at the population level. This phenomenon is exactly analogous to statistical sampling issues, such as tossing coins: if there are few coins, the proportion that come up heads can be far from 0.5. This process is also called "demographic stochasticity," and is what almost always is meant when demographic stochasticity is included in models (Lacy 1993, Akçakaya and Baur 1996).

Having made this assumption that all individuals have the same demographic parameters, then environmental stochasticity clearly represents unpredictable temporal variation in these parameters. This variation in demography is caused by variation in the environment, but the two are not the same: individuals can buffer the environmental variation through behavioral and physiological responses. Thus, environmental stochasticity in demography represents the filtering (possibly highly nonlinear) of environmental variation by the organisms in question.

Demographic and environmental stochasticity can be quite different. Early analyses suggested that they had qualitatively different effects on the relationship between population size and population extinction time (Shaffer 1987). Although these analyses may be flawed (Lande 1993), the two processes differ in the way in which variance in the demographic parameters scales with population size. Furthermore, the variance due to both types of stochasticity reduces the long-term growth rate of the population, whereas only environmental stochasticity affects long-term individual fitness. Thus, it is important to include both types of stochasticity in models of small populations.

Increasing numbers of models being used for population viability analysis are structured in just this way (Armbruster and Lande 1993, Lacy 1993, Lindenmayer et al. 1993, Akçakaya and Baur 1996, Bustamante 1996, McCarthy 1996, Song 1996). However, in every published example I have found, the models are parameterized incorrectly. To model survival, for example, each individual is given a survival probability over the time step of the model (typically one year); all individuals are given the same probability at a given time (the "sampling" definition of demographic stochasticity), but the value fluctuates over time (environmental stochasticity). The magnitude of the demographic sto-

chasticity at any given time is determined by the population size and the survival probability at that time. The parameters that must be estimated empirically are the mean and variance of the survival probability over time: that is, the properties of the environmental stochasticity. The best data for parameterizing this part of the model are a sequence of estimates of survivorship over multiple times. Survivorship is the the proportion of the population that survives over each time interval, and it fluctuates through time. The universal approach has been to calculate the sample variance in survivorship and to use that as the variance of the survival probability in the model. However, demographic stochasticity will cause the survivorship to vary even if the survival probability is constant; in general, the variation in survivorship is a nonadditive combination of the effects of demographic and environmental stochasticity. Thus, the magnitude of environmental stochasticity has consistently been overestimated. This is likely to have produced overly pessimistic estimates of the times to population extinction.

In this paper, I show how to extract a more accurate estimate of the magnitude of environmental stochasticity from data on population sizes and survivorships. First, I describe simulated data, based on the model of survival in the previous paragraph, which I will use to test the accuracy and error rates of the tests. I then show how to answer the question, "Is there temporal variation in survival?" This is a simple hypothesis-testing exercise. Third, I describe and analyze a technique to estimate the magnitude of environmental stochasticity, if it exists, using likelihood models. Finally, I apply these techniques to data, published by Stacey and Taper (1992), from a population of Acorn Woodpeckers (*Melanerpes formicivorus*). With these data, I illustrate how to use alternative models of environmental stochasticity and how to incorporate density dependence in survival.

SIMULATED DATA

I generated two types of simulated data, reflecting models with demographic stochasticity alone and in combination with environmental stochasticity. These models embody the verbal models developed in the *Introduction,* and were designed to be qualitatively similar to the Acorn Woodpecker data to be described. In all cases, I generated 1000 Y-year data sets, where Y was one of 10, 20, 30, or 40. In all data sets, the number of individuals for each year, N_t, was drawn at random from a uniform distribution on [10, 50].

To simulate the effects of demographic stochasticity alone, I assumed that there was a fixed individual survival probability, π. For each year, the number of survivors, m_t, was drawn at random from a binomial distribution with parameters N_t and π; I call this the binomial model. I created nine sets of simulations, with $\pi = \{0.1, 0.2, \ldots, 0.9\}$.
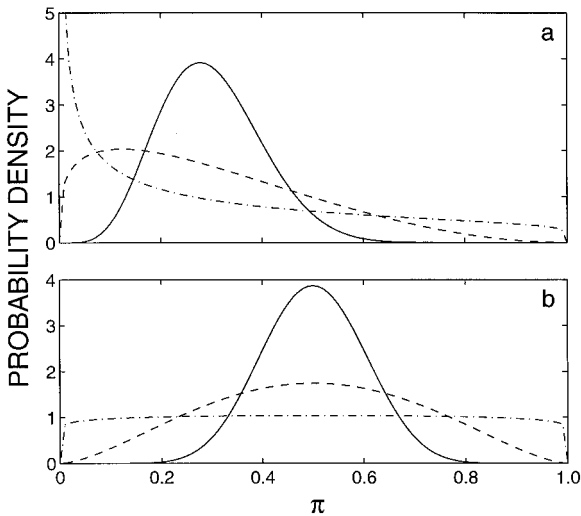
I simulated the combined effects of demographic and

FIG. 1. Beta-distributed probability densities of survival probability (π), showing variances of 0.01 (solid line), 0.04 (dashed line), and 0.08 (dash-dotted line). Means are (a) 0.3 and (b) 0.5.

environmental stochasticity by allowing π_t, the individual survivorship probability in year t, to vary from year to year. I drew π_t from a beta distribution with parameters chosen to give the desired mean and variance. I then drew m_t from a binomial distribution with parameters N_t and π_t. This is known as a beta-binomial model (Cox and Snell 1989). I used $\bar{\pi} = 0.5$ and created eight sets of simulations, with $\sigma^2(\pi) = \{0.01, 0.02, \ldots, 0.08\}$.

A variance in survival of 0.01 may seem small, but it is not (Fig. 1): the maximum possible variance for a variable restricted to [0, 1] is 0.25, and a uniform distribution has a variance of $\sim$0.08.

I generated random numbers with the package RANLIB, which generates uniform deviates (for population sizes) with a pair of multiplicative linear congruential generators (L'Ecuyer and Côté 1991) and binomial random numbers (for numbers of survivors) with the BTPE algorithm (Kachitvichyanukul and Schmeiser 1988).

IS THERE TEMPORAL VARIATION IN SURVIVAL?

As described in the *Introduction,* the most parsimonious model of individual variation in survival is that each individual has the same probability of surviving. If there is no environmental stochasticity, then this probability is the same each year. This model of demographic stochasticity is mathematically equivalent to sampling error: the observed variation in survivorship is simply a consequence of drawing an integer number of survivors from a finite population. A straightforward way to examine this hypothesis is to use the G test as a test of independence among years. Sample sizes are often small, so I use the Williams correction (Williams 1976).

In this context, the G test is an approximate test of whether the number of survivors each year is drawn

from a binomial distribution, with an individual survival probability given by $\Sigma m_t / \Sigma N_t$. The G test is unreliable if the expected number of deaths or survivors is ever less than five (Sokal and Rohlf 1995). In small populations, this might often be the case, so either the G test would have to be abandoned or those years of data would have to be discarded.

I used the simulated data to test both the type I error (the probability of incorrectly rejecting the hypothesis of demographic stochasticity alone) and the type II error (the probability of incorrectly failing to reject the hypothesis). To test the former, I subjected data from the binomial model with $Y = 10$ to the G test with $\alpha = 0.05$. Except when π is close to zero or one, the type I error rate clusters around 0.05, the desired result (Fig. 2). To test the type II error, I used data from the beta-binomial model, for which the hypothesis of demographic stochasticity alone is false. The type II error is large when the environmental stochasticity is small, but declines exponentially with the true variance in survival probability; the error also declines rapidly with the length of the data set (Fig. 3).

The G test, as used here, is an approximation to the underlying binomial distributions of the null model. It is possible to construct an exact test that models those distributions directly. For the data used here, the exact test produces results broadly similar to those from the G test; thus, I do not discuss it further.

ESTIMATING THE MAGNITUDE OF ENVIRONMENTAL STOCHASTICITY

If the previous analysis shows that demographic stochasticity alone cannot explain the observed variability in survivorship, then the natural question to ask is, "How much environmental stochasticity is present?" As reviewed in the *Introduction,* the most common approach is simply to calculate the variance of the observed survivorships. There are many ways of calcu-
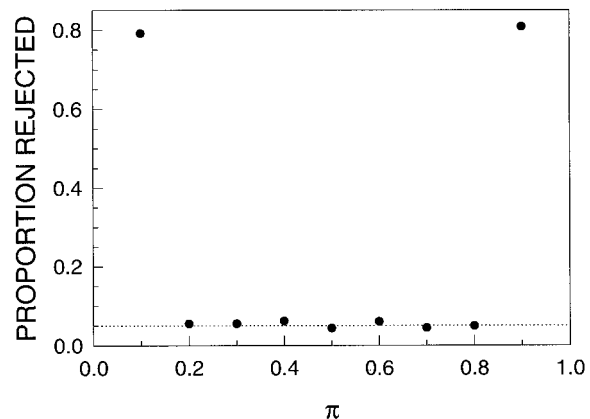

FIG. 2. Type I error in the G test, showing the proportion of data sets for which the null hypothesis of demographic stochasticity alone was rejected at $\alpha = 0.05$, as a function of survival probability, π, for data generated by the binomial model.
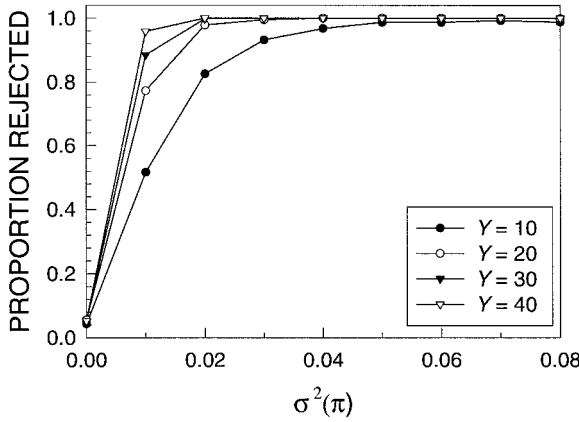
FIG. 3. Type II error in the *G* test, showing the proportion of data sets for which the null hypothesis of demographic stochasticity alone was rejected at $\alpha = 0.05$, as a function of the variance in survivorship probability, for data generated by the beta-binomial model. *Y* is the number of years of simulated data.

TABLE 1. Results of fitting the beta-binomial model to simulated data with environmental stochasticity. The estimated variance in the survival probability for each simulation is $\widehat{\sigma^2(\pi)}$; V_p is the variance in survivorship. Means are taken over the 1000 simulations. The correction is described in the text.

$\sigma^2(\pi)$	Mean $\widehat{\sigma^2(\pi)}$	Corrected mean $\widehat{\sigma^2(\pi)}$	Mean V_p
0.00	0.0010	0.0011	0.0076
0.01	0.0078	0.0087	0.0157
0.02	0.0169	0.0188	0.0248
0.03	0.0256	0.0285	0.0334
0.04	0.0346	0.0384	0.0423
0.05	0.0431	0.0479	0.0502
0.06	0.0533	0.0592	0.0603
0.07	0.0610	0.0678	0.0675
0.08	0.0703	0.0781	0.0768

lating this quantity, reflecting different assumptions about the reliability of individual estimates of survival probability. However, it seems reasonable to weight, in some way, by population size, because very few values of *p* are possible when *N* is small. The formulation that seems to have the least bias is

$$\text{var}(p) = \frac{\sum_{t=1}^{Y} N_t (p_t - \bar{p})^2}{\sum_{t=1}^{Y} N_t} \qquad (1)$$

where $\bar{p} = \sum_{t=1}^{Y} m_t / \sum_{t=1}^{Y} N_t$. When applied to the simulated data, however, this estimate of the variance in survival probability is positively biased at low levels of environmental stochasticity and negatively biased at high levels (Table 1).

A better answer to this question requires assumptions about how the value of π_t varies through time. Assume, for example, that π_t is a random number drawn from a probability distribution $P(s)$. What is the estimate of π_t, given the data in year *t*? Bayes' Theorem states that

$$\Pr(\theta \mid D) = \frac{\Pr(D \mid \theta)\Pr(\theta)}{\Pr(D)} \qquad (2)$$

where θ is the value of the parameter of interest and *D* is the data. In this case, θ is the value of π_t, and the data are N_t and m_t. $\Pr(D \mid \theta)$ is then just the binomial function, $\binom{N_t}{m_t} s^{m_t} (1 - s)^{N_t - m_t}$, and $\Pr(\theta)$ is the underlying probability distribution $P(s)$. $\Pr(D)$, the "probability of the data," is just a normalization constant; by considering likelihoods instead of probabilities, I can ignore it. Thus, $L(s)$, the relative likelihood that $\pi_t = s$, given the data and $P(s)$, is

$$L(s) = \binom{N_t}{m_t} s^{m_t} (1 - s)^{N_t - m_t} P(s). \qquad (3)$$

The real quantity of interest is $P(s)$. Let λ represent the set of parameters that control the shape of *P*, and denote the dependency by $P_\lambda(s)$. Substituting $P_\lambda(s)$ into Eq. 3 gives a joint likelihood of λ and *s* in year *t*. To get the likelihood of λ, given the data in year *t*, I integrate Eq. 3 over all possible values of the true survival probability *s*:

$$L_t(\lambda) = \int_0^1 \binom{N_t}{m_t} s^{m_t} (1 - s)^{N_t - m_t} P_\lambda(s) \, ds. \qquad (4)$$

Finally, to take into account all of the years of data, I calculate the joint likelihood of the L_t's by taking their product:

$$L(\lambda) = \prod_{t=1}^{Y} L_t(\lambda). \qquad (5)$$

Each value of λ has associated with it a mean and variance of the distribution $P(s)$. Thus, I can transform the likelihood function that is estimated for λ into a likelihood function for the variance of *P*, which is what I call the magnitude of environmental stochasticity.

The choice of a function for *P* is somewhat arbitrary, as there will rarely be enough information to determine what it truly is. It depends not only on the nature of the environmental variability, but also on the response of the organism to that variability. It is valuable to parameterize *P* in such a way that the binomial model is a special case of it, and that *P* itself is a special case of any further extensions, such as density dependence. Choosing *P* in this way allows comparisons to be made between models of differing complexity, using the likelihood ratio test. Lebreton et al. (1992) describe this test in detail; in brief, suppose λ_0 is a special case of λ: certain parameters are fixed, usually to zero. If $2 \log(L(\lambda)/L(\lambda_0))$ is not significantly large with respect to the χ^2 distribution with degrees of freedom equal to the number of constrained parameters, then the full model $P(\lambda)$ is not significantly more likely than the constrained model $P(\lambda_0)$.
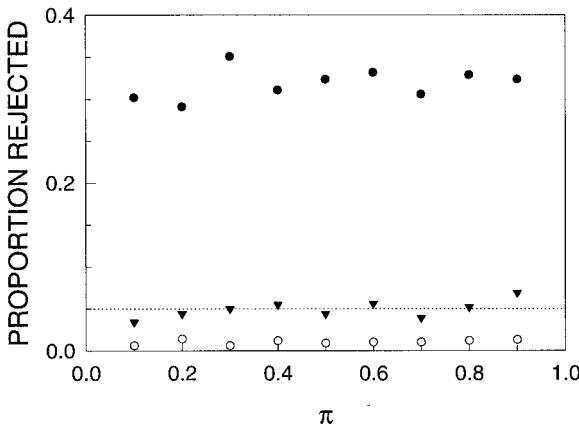
FIG. 4. Type I error fitting the beta-binomial model to the binomial data. Shown is the proportion of data sets for which the null hypothesis of demographic stochasticity alone was rejected by the criteria $\sigma^2(\pi) \neq 0$ (●), $P < 0.05$ (○), and $P < 0.19$ (▼), where the P values refer to the likelihood ratio test between the binomial and beta-binomial models. The third criterion has a type I error rate of ~0.05.

Beta-distributed environmental stochasticity

In the following analysis, I use the beta distribution, which is a two-parameter distribution with a central tendency that is restricted to the interval [0, 1] (for more details see any introductory probability text, such as Ross 1994; I show some examples in Fig. 1). The beta distribution is convenient in this context, because Eq. 4 reduces to

$$L_t(\bar{\pi}, \sigma^2(\pi)) = \binom{N_t}{m_t} \frac{B(m_t + a, N_t - m_t + b)}{B(a, b)} \quad (6)$$

where B is the beta function and a and b are the parameters of the beta distribution, related to the mean and variance by

$$a = \bar{\pi}\left[\frac{\bar{\pi}(1 - \bar{\pi})}{\sigma^2(\pi)} - 1\right] \quad (7)$$

$$b = (1 - \bar{\pi})\left[\frac{\bar{\pi}(1 - \bar{\pi})}{\sigma^2(\pi)} - 1\right]. \quad (8)$$

This is the beta-binomial model and, hence, it is a true model of the simulated data. I can therefore use these data to test the reliability and bias of the likelihood estimates.

First, consider the case where there is no environmental stochasticity. Ideally, the maximum likelihood estimate (MLE) of $\sigma^2(\pi)$ would be zero; this in fact occurs roughly two-thirds of the time (Fig. 4). In most other cases, the MLE is very small and I need to determine whether it is significantly different from zero. I do this by using the likelihood ratio test to compare the beta-binomial model with the binomial model. Typically, an appropriate P value for this test is larger than the desired α (Lebreton et al. 1992). When I use a rejection criterion of $P = 0.05$, then the type I error

rate is ~1%, and a criterion of $P = 0.19$ is required to bring the error rate up to 5% (Fig. 4). I will use this value of P for comparisons of the binomial and beta-binomial models throughout the rest of the paper.

When this criterion is applied to the simulated data from the beta-binomial model, then the type II error rate is similar to that from the G test (Fig. 5). This result is not surprising, because the G test is a form of likelihood ratio test; the main difference here is that a specific alternative hypothesis is being compared to the null model.

Applied to the data from the beta-binomial model, both the MLE of the mean survival probability and the mean survivorship give unbiased estimates of the true mean survival probability. The picture is rather different for the true variance, however (Table 1). The variance estimate from the survivorships is too high at low levels of environmental stochasticity, and too low at high levels. In contrast, the MLE from fitting the beta-binomial model is consistently too low. This bias can be substantially reduced by multiplying the MLE by $Y/(Y - 1)$. The remaining small negative bias seems independent of the true variance. Why does this correction work? Recall that an ordinary sample variance is negatively biased, which is why we divide by $(n - 1)$ rather than n (Sokal and Rohlf 1995). In this case, there are Y estimates of π, and the same bias phenomenon seems to hold, despite the fact that the individual estimates are likelihood distributions rather than single values.

AN EXAMPLE: ACORN WOODPECKERS

Stacey and Taper (1992) studied a population of Acorn Woodpeckers (*Melanerpes formicivorus*) in Water Canyon, central New Mexico, United States, from 1975 through 1984. They marked all individuals in the population (which fluctuated in size from 32 to 52
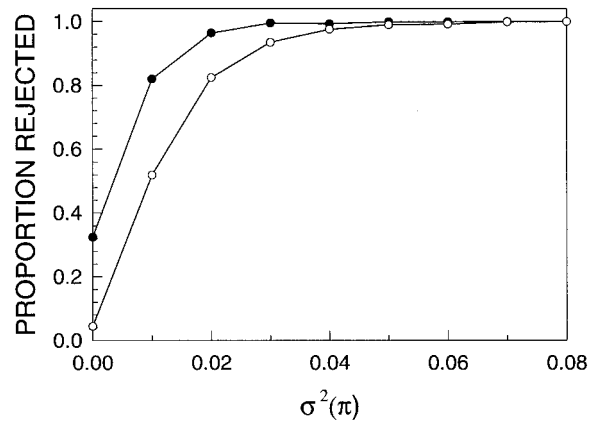

FIG. 5. Type II error fitting the beta-binomial model to the beta-binomial data. Shown is the proportion of data sets for which the null hypothesis of demographic stochasticity alone was rejected by the criteria $\sigma^2(\pi) \neq 0$ (●) and $P < 0.19$ (○), where the P value refers to the likelihood ratio test between the binomial and beta-binomial models.

TABLE 2. Number of individuals (N), number surviving to the following spring (m), and survivorships (p) of juvenile and adult Acorn Woodpeckers in the Water Canyon, New Mexico population, reconstructed from Stacey and Taper (1992). The survivorships reported here are recalculated from the estimated N and m.

Year	Juveniles			Adults		
	N	m	p	N	m	p
1975	59	33	0.569	46	24	0.522
1976	22	14	0.636	46	31	0.674
1977	43	13	0.302	40	28	0.700
1978	42	17	0.405	51	19	0.372
1979	1	0	0.000	52	28	0.538
1980	48	18	0.375	32	22	0.688
1981	39	7	0.179	46	30	0.652
1982	8	2	0.250	49	24	0.490
1983	25	11	0.440	35	21	0.600

adults), and recorded reproduction and survival. Both the presence (Stacey and Ligon 1987) and absence (Stacey and Taper 1992) of age-dependent adult survivorship have been inferred from the data; I assume the latter for this analysis, and use the annual survivorships reported in Stacey and Taper (1992). I treat juveniles separately, because their survivorship is often substantially lower than that of adults.

Stacey and Taper (1992) reported the number of adults, survivorships for juveniles and adults, and fecundity (fledged young per breeding pair) for each year. I estimated the number of juveniles by dividing the number of adults by 2.64 (the mean number of adults per breeding pair; Koenig and Stacey 1990), and multiplying the result by the fecundity. In some cases, I then adjusted the number by a small amount until I could find an (N, m) pair that matched the reported survivorships. For both juveniles and adults, I estimated the number surviving by multiplying the number of individuals by the survivorship. For the adults, there was not always an integer m that generated the reported survivorship; I chose the value that gave the closest value of survivorship to that reported. Juvenile survivorship ranged from 0.00 to 0.64, and adult survivorship ranged from 0.37 to 0.70 (Table 2).

Is there environmental stochasticity?

In the juvenile Acorn Woodpecker data, the year 1979 must clearly be discarded, because the G test cannot handle zero values; the year 1982 is borderline (the expected number of survivors is 4.8). The latter turns out not to matter: if 1982 data are included, $P = 3.72 \times 10^{-5}$, df = 7 ($G_{adj} = 32.205$); if excluded, $P = 2.41 \times 10^{-5}$, df = 6 ($G_{adj} = 31.112$). This is a clear rejection of the hypothesis that juvenile survival probability is homogeneous across years. For adults, $P = 0.015$, df = 8 ($G_{adj} = 18.937$). Thus, the adult survival probability is also heterogeneous among years, although the unexplained variation is not as large as it is for the juveniles.
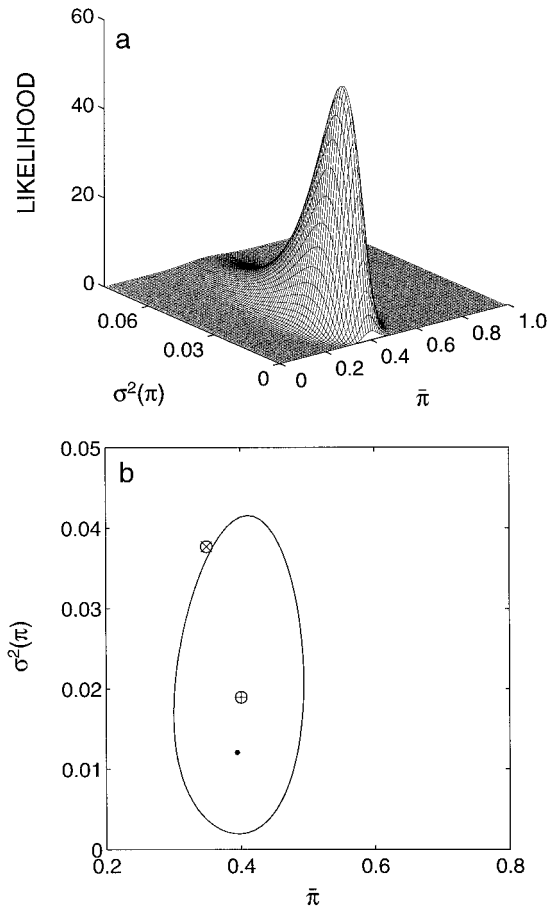

FIG. 6. Joint likelihoods of the mean ($\bar{\pi}$) and variance ($\sigma^2(\pi)$) of the individual survival probability for juvenile Acorn Woodpeckers, as fit by the beta-binomial model. (a) Likelihood surface; (b) the maximum likelihood value (●) and the 95% confidence region. Symbols: ⊕ marks the weighted mean and variance of the observed survivorships; ⊗ marks the unweighted mean and variance of the observed survivorships.

Fitting with the beta-binomial model

The likelihood calculations using the beta distribution for environmental stochasticity confirm that the magnitude of environmental stochasticity is higher for juvenile than for adult Acorn Woodpeckers (Figs. 6 and 7). The likelihood ratio test indicates that the beta-binomial model is significantly more likely than the binomial model for both juveniles (Table 3; $P = 0.010$) and adults (Table 4; $P = 0.081$). The maximum likelihood estimates of the mean survival probability were similar to the mean survivorship (juveniles: 0.401 vs. 0.395; adults: 0.576 vs. 0.571). However, the maximum

TABLE 3. Maximum likelihood parameter estimates for juvenile Acorn Woodpeckers.

Model	$\hat{\bar{\pi}}$	$\widehat{\sigma^2(\pi)}$	log L
Beta-binomial	0.395	0.012	-23.627
Binomial	0.401		-26.945

likelihood estimates of the variance in survival probability due to environmental stochasticity were substantially less than the total variance in the survivorships (juveniles: 0.0121 vs. 0.0190; adults: 0.0056 vs. 0.0111). Figure 8 shows the probability distributions of π embodied by these estimates.

Density dependence

Environmental stochasticity is not the only potential source of between-year variation in survival probability. Survival often decreases with increasing density, for reasons such as fewer resources provisioned, smaller metabolic reserve for migration, smaller size, or fewer safe sites. The relationship between density and survival is often straightforward enough that it can be estimated by a regression of survival on density. The analysis then focuses on the residual variation about the regression: can it be explained by demographic stochasticity, or is there some additional year-to-year variation in survival?

A logistic regression indicates that, in adults, the survival probability declines with density (Fig. 9). The question then becomes: does the apparent environmental stochasticity in adult survival merely reflect the variation in density? To address this question, I created a density-dependent model for which the previous simpler models are special cases. That is, the number surviving in any given year is drawn from a beta-binomial distribution, but the mean of that distribution is a logistic function of density, $\overline{\pi}_t = \exp(\alpha + \beta N_t)/(1 + \exp(\alpha + \beta N_t))$. The third parameter of the model is $\sigma^2(\pi)$, the variance of the beta distribution, which I assume to be independent of population size. I can then use the likelihood ratio test to compare this full model to a density-dependent model without environmental stochasticity and to the density-independent model with environmental stochasticity.

The parameter estimates and likelihoods are in Table 4. The full model is not significantly more likely than the simple density-dependent model ($P = 0.706$, df = 1), indicating that there is no evidence for environmental stochasticity on top of density dependence. The full model is significantly more likely than the density-independent beta-binomial model ($P = 0.030$, df = 1), and the density-dependent binomial model is more likely than the simple binomial model ($P = 0.006$, df = 1). Thus, density dependence improves the model, independently of whether environmental stochasticity is included. Unfortunately, there is no way to directly compare the density-dependent binomial model with the density-independent beta-binomial model.

There is no evidence of density dependence in juvenile survival (likelihood ratio test between density-dependent beta-binomial and beta-binomial models: $P = 0.614$, df = 1).
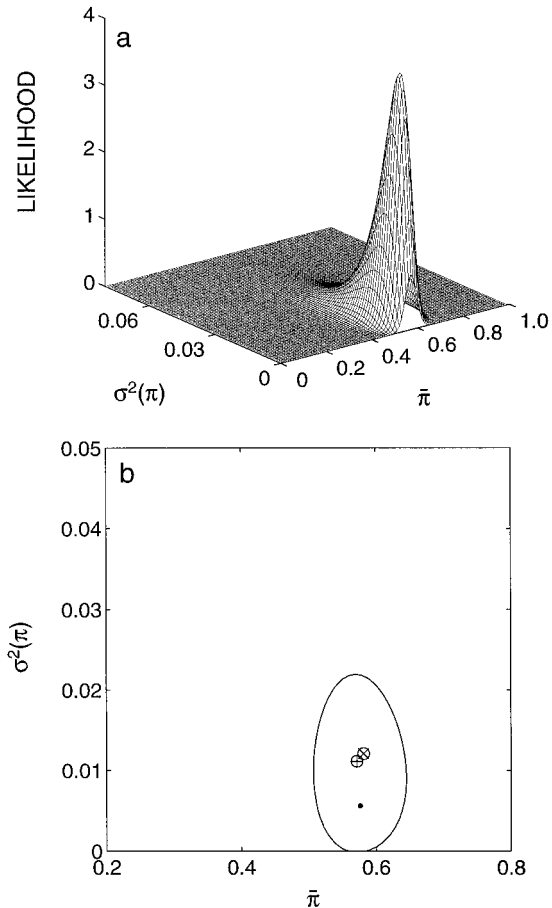

FIG. 7. Joint likelihoods of the mean ($\bar{\pi}$) and variance ($\sigma^2(\pi)$) of the individual survival probability for adult Acorn Woodpeckers, as fit by the beta-binomial model. (a) Likelihood surface; (b) the maximum likelihood value (●) and the 95% confidence region. Symbols: ⊕ marks the weighted mean and variance of the observed survivorships; ⊗ marks the unweighted mean and variance of the observed survivorships.

TABLE 4. Maximum likelihood parameter estimates for adult Acorn Woodpeckers. The column marked $\hat{\pi}$ is the estimate of π for the binomial model, $\bar{\pi}$ for the beta-binomial model, and π_0 for the catastrophe model.

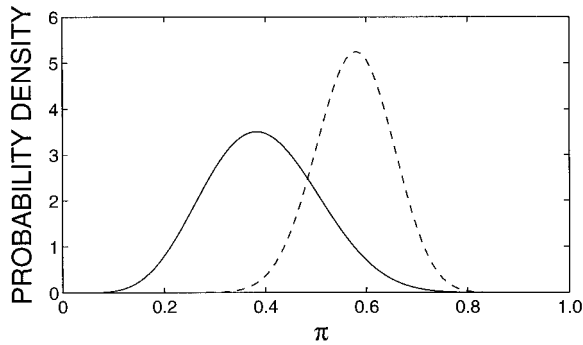| Model	$\hat{\pi}$	$\hat{\alpha}$	$\hat{\beta}$	$\widehat{\sigma^2(\pi)}$	$\hat{\pi}_f$	$\hat{c}$	$\log L$
Density-dependent beta-binomial		2.359	−0.046	0.001			−23.946
Density-dependent binomial		2.388	−0.046				−24.017
Beta-binomial	0.576			0.006			−26.303
Binomial	0.572						−27.827
Density-dependent catastrophe		2.388	−0.046		…	0	−24.017
Catastrophe	0.622				0.457	0.274	−25.968

FIG. 8. The maximum likelihood distributions of the survival probability in Acorn Woodpeckers, as estimated from the beta-binomial model: solid line, juveniles; dashed line, adults.

A different model of environmental stochasticity

A simultaneous test for homogeneity of replicates, using the G test, reveals that either the set of years excluding 1978 or the set of years excluding 1977 and 1980 appears homogeneous for adult survival. The studies of this population do not record anything unusual about 1977 and 1980, but the acorn crop, which forms the main part of the birds' winter diet, failed in the fall of 1978 (Koenig and Stacey 1990). This is also the year of lowest adult survivorship. This suggests that the only substantial source of environmental stochasticity for adults is crop failure; all other variation in survivorship during "good years" can be explained by demographic stochasticity alone. However, 1978 was also a high-density year. Reanalyzing the density-dependent models with the 1978 data removed suggests that there might still be a density effect (likelihood ratio test between density-dependent binomial and binomial models: $P = 0.080$, df $= 1$), but the evidence is weakened.

The phenomenon of crop failure suggests an alternate model of environmental stochasticity in adult woodpeckers: there is one survival probability, π_0, in normal years, and another, π_f, in crop failure years; a crop failure occurs with probability c. I allow π_0 to be density dependent; with only one crop failure observed, there is
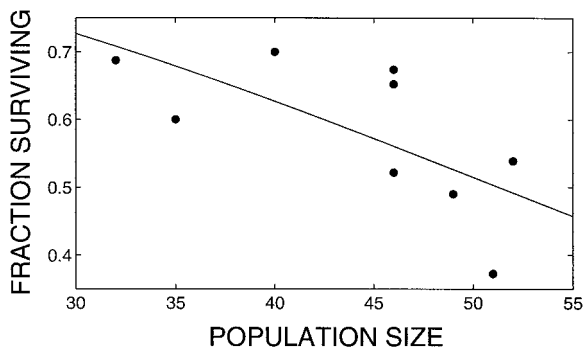

FIG. 9. Survival declines with density in adult Acorn Woodpeckers. The curve shows the maximum likelihood estimate of the logistic regression for the density-dependent beta-binomial model.

no way to assess density dependence in π_f. The MLE of the density-independent catastrophe model (Table 4) fits marginally better than that of the binomial model (likelihood ratio test: $P = 0.156$, df $= 2$; although the likelihood is higher than the beta-binomial model, it has one more parameter). Adding density dependence to the catastrophe model increases the likelihood significantly ($P = 0.048$, df $= 1$), but the maximum likelihood estimate of c is zero: there are no bad years! Thus, the catastrophe model does not seem to be a particularly good description of adult woodpecker survival.

DISCUSSION

The first step in building a mixed model of environmental and demographic stochasticity is to ensure that the observed variation is not all due to demographic stochasticity. The tests described in this paper perform this task for survivorship. The G test is quite effective at distinguishing demographic stochasticity from all but weak environmental stochasticity, as long as the mean survivorship is not too extreme. The exact test (results not reported) is more reliable in these extreme cases; however, it is more biased in favor of the null hypothesis of demographic stochasticity than is the G test. Thus, it would be helpful to reduce the type II error in this range; a more detailed examination of the distribution of P values may shed light on this issue. Whenever the variance falls in this range and the null hypothesis is close to being rejected, then simulations can help. For example, one could generate 1000 (or more) demographic stochasticity data sets with the observed population sizes to generate the expected distribution of G under the hypothesis of demographic stochasticity alone. The observed G would then be compared to this distribution.

This approach to looking for temporal heterogeneity in demographic parameters is sometimes applied in long-term demographic studies (such as Berg 1994 and Prince et al. 1994). However, it seems not to have been used in the context of modeling small populations.

Any estimate of the magnitude of environmental stochasticity requires that the alternative model be more specific than just the negation of the null model of demographic stochasticity alone. The cost of this requirement is high, inasmuch as a poor choice for the underlying distribution of the demographic parameter may give spurious results; but the payoff of getting an improved estimate of the magnitude of environmental stochasticity is also high. Furthermore, incorporating environmental stochasticity into population models requires a specific functional form for the probability distribution. Thus, in a sense, the analyses described here represent an improved method for fitting those models to the data. The technique is also extremely flexible: many hypotheses about the sources of variation in the demographic parameter can be implemented, as I illustrated with the density-dependent and catastrophic models of adult woodpecker survival.

The analysis of simulated data is essential to this approach. Without the simulated data, I could not have found the appropriate P values for the likelihood ratio test (indeed, to be rigorous, I should have done the same sort of analysis for each model comparison), nor would I have known about the bias in the maximum likelihood estimate of the variance in the survival probability. The simulated data that I used are fairly representative of small populations. For particular applications, it may be beneficial to tie the simulations even more closely to the real data, using the actual set of observed population sizes, for example.

There are two major unanswered questions from the analysis of the simulated beta-binomial data. The first concerns the source of the remaining bias in the MLE of the variance. Is it, in fact, a constant, or does it depend on the number of years of data or the minimum population size? Preliminary analysis suggests that the bias might be reduced by integrating over all values of the mean survival probability, to get an unconditional MLE of the variance. The second question is whether the type II error at low levels of environmental stochasticity can be reduced, given a fixed length of data (the value of more years of data is clear from Fig. 3). I suspect that the possibility for improvement depends, at least in part, on the minimum population size: for example, when $\pi_t = 0.5$ and $N_t = 10$, the variance in survivorship due to demographic stochasticity is 0.025; when $N_t = 25$, it is 0.01. Thus, in the simulated data with $\sigma^2(\pi) = 0.01$, the contributions of demographic and environmental stochasticity are often approximately the same and, hence, are inherently difficult to distinguish. This point reinforces the value of tying the simulated data closely to the observed data.

The analyses of the Acorn Woodpecker data show clear evidence for environmental stochasticity in juvenile survival, but the magnitude of variation in the survival probability is likely to be less than the variation in survivorship. Adult survival is subject to either environmental stochasticity or density dependence; the data do not support including both simultaneously. In either case, the magnitude of environmental stochasticity is smaller for adults than it is for the juveniles. Because both groups experience the same general environment, this reinforces the point I made in the *Introduction* that environmental stochasticity is a function of the organism as well as the underlying environmental variability.

Adult survival was unusually low in a year of acorn crop failure, and there is no evidence of environmental stochasticity in the remaining years. Curiously, however, there is not strong support for a catastrophic model of adult survival; perhaps this is because there is only one such ''bad year'' in the data.

There are a number of limitations and critical simplifying assumptions in this form of analysis. First, it will often be necessary to lump adult age classes, as I have done for the adult Acorn Woodpeckers, to obtain a sufficient sample size. This assumed age independence of survival implies a type II survivorship function. Thus, the method will be most useful for organisms such as birds, shrubs, and trees, many of which show type II survivorship of mature individuals.

Second, a crucial assumption of these analyses is that the years are independent of one another. In reality, the survival probability can be correlated among years, both because of autocorrelations in the relevant environmental variables and because some individuals may be present for several years in a row, carrying along any unobservable traits that might affect their fates (e.g., if there are good survivors and poor survivors, the former will tend to persist longer than the latter). If these phenotypic differences are genetically determined, then the problem is even worse, as certain genotypes may persist for multiple generations. I do not know how such autocorrelations might bias the results; this is an important open question.

Third, any hidden variability among individuals violates the assumption I have made of uniform survival probability within years. There may even be variation among individuals in the ability to buffer the effects of environmental variability.

Fourth, the analyses presented here have assumed that the population is completely censused and survivorship is known exactly for each year. When this assumption is not true, then there is an additional level of uncertainty: each datum is only an estimate of the true value of survivorship that is to be entered into the model. For most sampling procedures, it is possible to calculate a likelihood function for the value of the true survivorship. Then, for example, Eq. 4 would have to be integrated over all likely (N_t, m_t) pairs, weighted by the likelihood of (N_t, m_t) given the data. Clearly, this will make the overall likelihood function flatter, and it will probably increase the variance of the maximum likelihood estimate; whether it introduces new biases into the MLE is an open question.

Finally, I have focused in this paper on the maximum likelihood estimates of the model parameters, but it is clear that, in many cases, the uncertainty in these estimates is high (consider the area encompassed by the 95% confidence region for the parameters of the beta-binomial model of juvenile woodpeckers). When using these models to project population fates, it will be important to take this uncertainty into account. For example, a likelihood distribution of times to extinction is usually calculated using a point estimate of population parameters, but it might be better calculated by integrating over all likely parameter combinations, weighted by their relative likelihoods.

I have focused on mortality in this analysis; reproduction is the other important demographic process. However, unlike mortality, there is no widely applicable null model for demographic stochasticity in fecundity. Not only do species differ in their potential for variation in fecundity, and in the structure of that variation, but there may be circumstances in which

variable fecundity is itself an optimal strategy (Schaffer 1974, Grey 1980). Nevertheless, with sufficient biological information, it should be possible to construct specific models. One approach might be to look at the distribution of fecundities among individuals within a year, and then to test whether that distribution differs among years. A complete population model would also require some estimates of immigration and emigration. With such a full model, one could make sophisticated estimates of extinction risk, using techniques such as those in Ludwig (1996).

This type of analysis can add important insights to the biology of the organism under study and point toward further research. For example, in the Acorn Woodpeckers of Water Canyon, why is adult survival, but not juvenile survival, density dependent? Why is the variance in juvenile survival larger than that in adult survival? Investigations into these sorts of questions might lead to a deeper understanding of the effects of behavior and physiology on population dynamics.

There has long been a substantial gap between the theory of stochastic population dynamics and those processes that are measurable in the field. Expanding our conception of stochasticity to include the limits of an observer's information as well as the "true" variability among individuals and environments will help to guide this theory in more directly applicable directions.

Literature Cited

Akçakaya, H. R., and B. Baur. 1996. Effects of population subdivision and catastrophes on the persistence of a land snail population. Oecologia **105**:475–483.

Armbruster, P., and R. Lande. 1993. A population viability analysis for African elephant (*Loxodonta africana*): how big should reserves be? Conservation Biology **7**:602–610.

Berg, Å. 1994. Maintenance of populations and causes of population changes of curlews *Numenius arquata* breeding on farmland. Biological Conservation **67**:233–238.

Bustamante, J. 1996. Population viability analysis of captive and released Bearded Vulture populations. Conservation Biology **10**:822–831.

Cox, D. R., and E. J. Snell. 1989. Analysis of binary data. Second edition. Chapman and Hall, London, UK.

Durant, S. M., and J. Harwood. 1992. Assessment of monitoring and management strategies for local populations of the Mediterranean monk seal *Monachus monachus*. Biological Conservation **61**:81–92.

Gilpin, M. E., and M. E. Soulé. 1986. Minimum viable populations: processes of species extinction. Pages 19–34 *in* M. E. Soulé, editor. Conservation biology: the science of scarcity and diversity. Sinauer Associates, Sunderland, Massachusetts, USA.

Grey, D. R. 1980. Minimisation of extinction probabilities in reproducing populations. Theoretical Population Biology **18**:430–443.

Kachitvichyanukul, V., and B. W. Schmeiser. 1988. Binomial random variate generation. Communications of the ACM **31**:216–222.

Koenig, W. D., and P. B. Stacey. 1990. Acorn Woodpeckers: group-living and food storage under contrasting ecological conditions. Pages 415–453 *in* P. B. Stacey and W. D. Koenig, editors. Cooperative breeding in birds. Cambridge University Press, Cambridge, UK.

Lacy, R. C. 1993. Vortex: a computer simulation model for population viability analysis. Wildlife Research **20**:45–65.

Lande, R. 1993. Risks of population extinction from demographic and environmental stochasticity and random catastrophes. American Naturalist **142**:911–927.

Lebreton, J.-D., K. P. Burnham, J. Clobert, and D. R. Anderson. 1992. Modeling survival and testing biological hypotheses using marked animals: a unified approach with case studies. Ecological Monographs **62**:67–118.

L'Ecuyer, P., and S. Côté. 1991. Implementing a random number with splitting facilities. ACM Transactions on Mathematical Software **17**:98–111.

Lindenmayer, D. B., R. C. Lacy, V. C. Thomas, and T. W. Clark. 1993. Predictions of the impacts of changes in population size and environmental variability on Leadbeater's possum, *Gymnobelideus leadbeateri* McCoy (Marsupialia: Petauridae) using population viability analysis: an application of the computer program Vortex. Wildlife Research **20**:67–86.

Ludwig, D. 1996. The distribution of population survival times. American Naturalist **147**:506–526.

McCarthy, M. A. 1996. Extinction dynamics of the helmeted honeyeater: effects of demography, stochasticity, inbreeding, and spatial structure. Ecological Modelling **85**:151–163.

Prince, P. A., P. Rothery, J. P. Croxall, and A. G. Wood. 1994. Population dynamics of Black-browed and Grey-headed Albatrosses *Diomedea melanophris* and *D. chrysostoma* at Bird Island, South Georgia. Ibis **136**:50–71.

Ross, S. 1994. A first course in probability. Fourth edition. Prentice Hall, Englewood Cliffs, New Jersey, USA.

Schaffer, W. M. 1974. Optimal reproductive effort in fluctuating environments. American Naturalist **108**:783–790.

Shaffer, M. 1987. Minimum viable populations: coping with uncertainty. Pages 69–86 *in* M. E. Soulé, editor. Viable populations for conservation. Cambridge University Press, Cambridge, UK.

———. 1990. Population viability analysis. Conservation Biology **4**:39–40.

Sokal, R. R., and F. J. Rohlf. 1995. Biometry. Third edition. W. H. Freeman, New York, New York, USA.

Song, Y.-L. 1996. Population viability analysis for two isolated populations of Hainan eld's deer. Conservation Biology **10**:1467–1472.

Stacey, P. B., and J. D. Ligon. 1987. Territory quality and dispersal options in the Acorn Woodpecker, and a challenge to the habitat-saturation model of cooperative breeding. American Naturalist **130**:654–676.

Stacey, P. B., and M. Taper. 1992. Environmental variation and the persistence of small populations. Ecological Applications **2**:18–29.

Williams, D. A. 1976. Improved likelihood ratio tests for complete contingency tables. Biometrika **63**:33–37.