A better way to estimate population trends

Jean-Yves Humbert, L. Scott Mills, Jon S. Horne and Brian Dennis

J.-Y. Humbert (jean-yves.humbert@art.admin.ch) and L. S. Mills, Wildlife Biology Program, College of Forestry and Conservation, Univ. of Montana, Missoula, MT 59812, USA. Present address for JYH: Agricultural Landscapes and Biodiversity, Agroscope Reckenholz-Tänikon Research Station ART, Reckenholzstrasse 191, CH-8046 Zurich, Schwitzerland. – J. S. Horne and B. Dennis, Dept of Fish and Wildlife Resources, Univ. of Idaho, Moscow, ID 83844-1136, USA.

Estimation of a population trend from a time series of abundance data is an important task in ecology, yet such estimation remains logistically and conceptually challenging in practice. First, the extent to which unequal intervals in the time series, due to missing observations or irregular sampling, compromise trend estimation is not well-known. Furthermore, the predominant trend estimation method (loglinear regression of abundance data against time) ignores the possibility of process noise, while an alternative method (the 'diffusion approximation') ignores observation error in the abundance data. State-space models that account for both process noise and observation error exist but have been little used. We study an adaptation of the exponential growth state-space (EGSS) model for use with missing data in the time series, and we compare its trend estimation to the status quo methods. The EGSS model provides superior estimates of trend across wide ranges of time series length and sources of variation. The performance of the EGSS model even with half of the counts in the time series missing implies that trend estimates may be improved by diverting effort away from annual monitoring and towards increasing time series length or improving precision of the abundance estimates for years that data are collected.

Estimating a population's rate of change, or trend, is a fundamental challenge in basic and applied ecology. Population trend can be defined as the average change in log-abundance per unit of time (Dennis et al. 1991, Link and Sauer 1998). Although trend can be estimated from age or stage-specific vital rates, it is commonly estimated from abundance data collected over time via enumeration, abundance estimates or abundance indices (Morris et al. 2002, Marsh and Trenham 2008). Ecological time series models, derived from abundance data over time for single populations, can incorporate additional factors such as environmental covariates (Dennis and Otten 2000), observer-related covariates (Link and Sauer 1997), and density dependence (James et al. 1996, Brook and Bradshaw 2006, Dennis et al. 2006). However, data to support such analyses are often not readily available (e.g. covariates are mostly absent in the nearly 5000 population time series maintained by the Global Population Dynamics Database; see Fagan et al. (2001), Inchausti and Halley (2001) and Brook et al. (2006)). Therefore, exponential growth models lacking covariates are most commonly used for trend estimation in applications (Sabo et al. 2004).

Ecologists interested in estimating simple exponential trend from abundance data typically face four primary challenges. First, the lengths of time series are short, often consisting of only 5–10 time steps (e.g. years) and seldom more than 30. Second, time series often have unequal time intervals, arising from observations missing due to funding, logistical or personnel constraints. Third, variance in the abundance data arises from an often unknown combination of environmental noise and observation or estimation error in the abundance data themselves. Finally, different statistical methods exist to estimate trend for time series abundance data, with little guidance as to relative performance.

At present two methods that accommodate unequal intervals in the time series are most commonly used to estimate trend (Table 1). The oldest, and predominant, method is a log-linear regression of counts against time, where the slope of the regression gives the population trend (Caughley 1977, Gerrodette 1987, Eberhardt and Simmons 1992). We refer to this as the exponential growth observation error (EGOE) model, because it tacitly assumes that variability in the data arises purely as sampling or 'observation error', with the population itself governed by deterministic exponential growth (Supplementary material Appendix 1).

A second method for estimating exponential trend parameters assumes that the population is censused (i.e. no observation error), and that variability in abundances is entirely due to growth rate fluctuations caused by environmental variability or 'process noise' (Dennis et al. 1991, Lande et al. 2003). The method takes abundance on the logarithmic scale to be described by a Brownian motion diffusion process with a constant drift rate. We refer to this as the exponential growth process noise (EGPN) model.

Table 1. Examples of typical applications of EGOE (exponential growth observation error), EGPN (exponential growth process noise), and EGSS (exponential growth state space) methods to estimate trend parameters for animal populations.

Application	Taxa	Citation
EGOE (log-linear approach)		
Determine contemporary declines	koala	Phillips 2000
Assess declines	common frog	Meyer et al. 1998
Determine population growth rate	wildebeest, buffalo, zebra	Grange et al. 2004
Estimate growth rate for an endangered species	wallaby	Fisher et al. 2000
Estimate population growth rate	ibex	Largo et al. 2008
EGPN (diffusion approximation approach)		
Predict extinction probabilities	35 rare breeding bird species	Gaston and Nicholls 1995
Predict extinction probabilities	mount Graham red squirrel	Buenau and Gerber 2004
Predict extinction probabilities	cape mountain zebra	Watson et al. 2005
EGSS (state-space approach)		
Estimate population growth and extinction parameters	can. sea otter and Yellowstone grizzly bear	Lindley 2003
Determine how spatial correlations affect PVA	chinook salmon	Hinrichsen 2009

Interestingly, EGPN estimates have been used widely to estimate the trend parameter for population viability analysis (PVA) but rarely for simply describing the trend of populations (Table 1).

A recently developed third method for estimating trend uses a stochastic 'state-space' exponential growth model that assumes both observation error and environmental process noise (Holmes 2001, Lindley 2003, Staples et al. 2004, Dennis et al. 2006). We term this the exponential growth state space (EGSS) model. Estimation is based on the insight that the EGSS model can be written as a linear mixed model (Staples et al. 2004), thereby making calculations possible through software for analysis of variance with mixed random and fixed effects. The EGOE and EGPN models occur as limiting special cases when the corresponding variance parameters approach zero. More recently, Staudenmayer and Buonaccorsi (2006) scale the observations in the linear mixed model framework to allow for unequally spaced time intervals. That the EGSS model can accommodate missing data is not well known in ecological practice.

In this paper we explain how to obtain maximum likelihood (ML) and restricted maximum likelihood (REML) parameter estimates for the EGSS model based on the scaled observations of Staudenmayer and Buonaccorsi (2006). We then use computer simulations to compare the performance of EGOE, EGPN and EGSS approaches for estimating population trend. We document the statistical properties of point and interval estimates of the trend parameter when each of the three models are applied to data generated by stochastic exponential growth with varying time series lengths, numbers of missing values in the time series, and ratios of observation error versus process noise. In light of our findings of substantial robustness of trend estimates under the EGSS model, we provide recommendations for future biomonitoring study design and analysis.

Methods

We use lower case to denote data as well as constants, and upper case to denote the stochastic process (random variable) that generates data. So, we write $n(0)$, $n(t_1)$, ..., $n(t_q)$ for a recorded time series of population abundances at times 0 $(=t_0)$, t_1, ..., t_q, and $N(t)$ for a random population abundance at time t with some associated probability distribution. 'Abundance' refers to numbers, biomass, or density, and may be determined from a complete census, an estimate, or an index that reliably and proportionally tracks population fluctuations.

The three models we consider are based on familiar deterministic exponential growth:

$$n(t) = n_0 \lambda^t \qquad (1)$$

Here $n(t)$ is population abundance at time t, with $n_0 = n(0)$, and λ is a positive constant indicating an increasing ($\lambda > 1$) or decreasing ($\lambda < 1$) population. On the logarithmic scale:

$$\ln n_t = \ln n_0 + (\ln \lambda)t \qquad (2)$$

A general stochastic version of the exponential growth model with both process noise and observation error takes the form of a state space model with an unobserved population component and a component representing the observed or estimated abundance values. Let $X(t)$ be the unobserved log-abundance of the population (now assumed to be a stochastic process) at time t and $Y(t)$ be the estimated or observed value of $X(t)$. We write the EGSS model as:

$$dX(t) = (\ln \lambda)dt + dB(t) \qquad (3)$$

$$Y(t_i) = X(t_i) + F_i \qquad (4)$$

Here $dB(t) \sim normal(0, \sigma^2 dt)$, and $F_i \sim normal(0, \tau^2)$. The term $dB(t)$ is a random perturbation representing the process noise (environmental variability), and F_i represents the observation error, assumed to have no auto- or cross-correlations. The quantity $\mu = \ln\lambda$ is the expected change of $X(t)$ in one time unit; it is our trend parameter (see Supplementary material Appendix 1 for more details). The model defines $X(t)$ as a Brownian motion process with drift rate μ and represents a generalization of the earlier EGSS version to continuous time. Equation 3 provides a simple recipe for simulating an increment of a population trajectory over a small time interval dt as an increment of deterministic exponential growth on the log scale perturbed by a normal random quantity; an entire population trajectory would be constructed by accumulating such increments. Sampling times

t_0 (=0), t_1, ..., t_q are not necessarily equally spaced. Special cases of the EGSS model are EGOE ($\sigma^2 = 0$) and EGPN ($\tau^2 = 0$). Additional properties of the EGOE, EGPN, and EGSS models, along with details of ML and REML parameter estimation, are given in Supplementary material Appendix 1.

For each set of conditions evaluated, we simulated 5000 replicate time series from the EGSS model with specified μ (from -0.2 to 0.2) and variance parameters (from 0 to 0.25 for both σ^2 = process noise and τ^2 = observation error). The duration of the simulated time series ranged from 5 to 50, bracketing from well below the minimum duration deemed acceptable for estimating population trend to approximately the longest time series currently available (Bence 1995, Holmes 2004).

The log-population abundances X(t), (t = 1, 2, ...) were generated with the discrete time autoregressive model:

$$X(t) = X(t-1) + \mu + E_t \qquad (5)$$

where $E_t \sim \text{normal}(0, \sigma^2)$. The autoregressive model has statistical properties identical to the continuous time model (Eq. 3) evaluated at discrete times. Observation errors were then added to each X(t):

$$Y(t) = X(t) + F_t \qquad (6)$$

where $F_t \sim \text{normal}(0, \tau^2)$. The simulated population values were calculated as $\exp(Y(t))$. Also, if a simulated population decreased to ≤ 2, it was deleted and simulated again in order to limit estimation to populations that do not become extinct before the requisite number of observations were obtained. The size of X(0) was always $\ln(1000)$, and the initial value of Y(t) was randomly chosen by $Y(0) = \ln(1000) + F_0$. (Results were insensitive to changes in X(0); Supplementary material Appendix 2). To explore the effects of missing observations in the time series, we removed randomly a fixed number of the abundance data from each simulated time series (the first and last observations were not allowed to be removed to keep the length of the survey constant). Comparing the estimates of trend from the same time series under complete and missing data isolates the effect of missing observations from other effects such as process stochasticity.

To each simulated time series, we fitted the EGOE and EGPN models using ML and the EGSS model using REML. We created box plots of point estimates of μ across the 5000 replicates for each of the three methods (EGOE, EGPN, and EGSS). We also explored the coverage of the interval estimators of μ by plotting the percent of 5000 replicates whose confidence intervals (CI) actually included μ. (See Supplementary material Appendix 1 for details on confidence interval construction). Because it may be equally problematic, in an applied context, for a CI of $\hat{\mu}$ to either over- or under- include the true μ, we use a 50% CI to provide symmetry for assessing over and underestimation. For example, a model reaching 40% coverage can be interpreted as equally biased to another model with 60% coverage. By contrast, a more traditional 90% or 95% CI would obscure comparisons because of the asymmetry of over- and under-inclusion. All simulations and computations were performed with R 2.8.1 for Windows (code available in Supplementary material Appendix 1).

Results

The same general patterns were found for all combinations of input trend parameters (positive, zero, or negative μ, with a range of process noise and observation error). Therefore, we only present detailed results for one scenario, a small true population decline with comparatively large process and sampling variance ($\mu = -0.02$, $\sigma^2 = 0.01$, $\tau^2 = 0.01$). (Results of other combinations of input parameters can be found in Supplementary material Appendix 2).

All three models provided relatively unbiased estimates of μ regardless of the ratio of process noise (σ^2) to observation error (τ^2) (Fig. 1A). Additionally, estimates of μ were robust to time series lengths from 5 to 50 years (Fig. 1B) and the number of missing observations in the time series (Fig. 1C). Although the additional parameter in the EGSS model might lead to an expectation of higher variability in the estimate of μ compared to the EGOE and EGPN models, we found no evidence for this under any conditions. For all 3 methods, variability increased with the ratio of process variation to observation error (Fig. 1A) and as the time series length decreased (Fig. 1B), but was mostly unaffected by missing observations (Fig. 1C).

While performance of the three models was good relative to bias in the estimates of μ, there were substantial differences in confidence interval coverage (Fig. 2). Confidence intervals constructed using the EGOE model were good (i.e. μ was contained in about 50% of the intervals) only when process noise was absent or small in relation to observation error (σ^2/τ^2 ratio near 0; Fig. 2A). However, as the σ^2/τ^2 ratio increased, confidence interval coverage became $< 50\%$ indicating that the intervals were too narrow. The coverage was improved by decreasing the time series length (Fig. 2B). Similarly, coverage also improved as the time series became less complete (i.e. more missing values; Fig. 2C). Improved coverage for smaller samples often happens when estimates are statistically inconsistent (i.e. converge to the wrong values asymptotically).

Confidence intervals constructed using the EGPN model had excellent coverage when process noise overwhelmed observation error (e.g. $\sigma^2/\tau^2 > \sim 5$; Fig. 2A). However, when process noise was minimal relative to observation error ($\sigma^2/\tau^2 < \sim 1$), confidence intervals constructed using the EGPN model became too wide ($>> 50\%$). As with the EGOE model, coverage improved (slightly) with decreased time series length (Fig. 2B) and more missing data (Fig. 2C).

Empirical coverage of confidence intervals constructed using the EGSS model tended to be less than 50% (CIs too narrow) but only slightly so as coverage rarely strayed more than 10% from the nominal value (Fig. 2). Coverage was consistently better than EGOE and EGPN models across all combinations of process noise and observation error, time series lengths, and number of missing observations. Unlike EGOE and EGPN, coverage improved
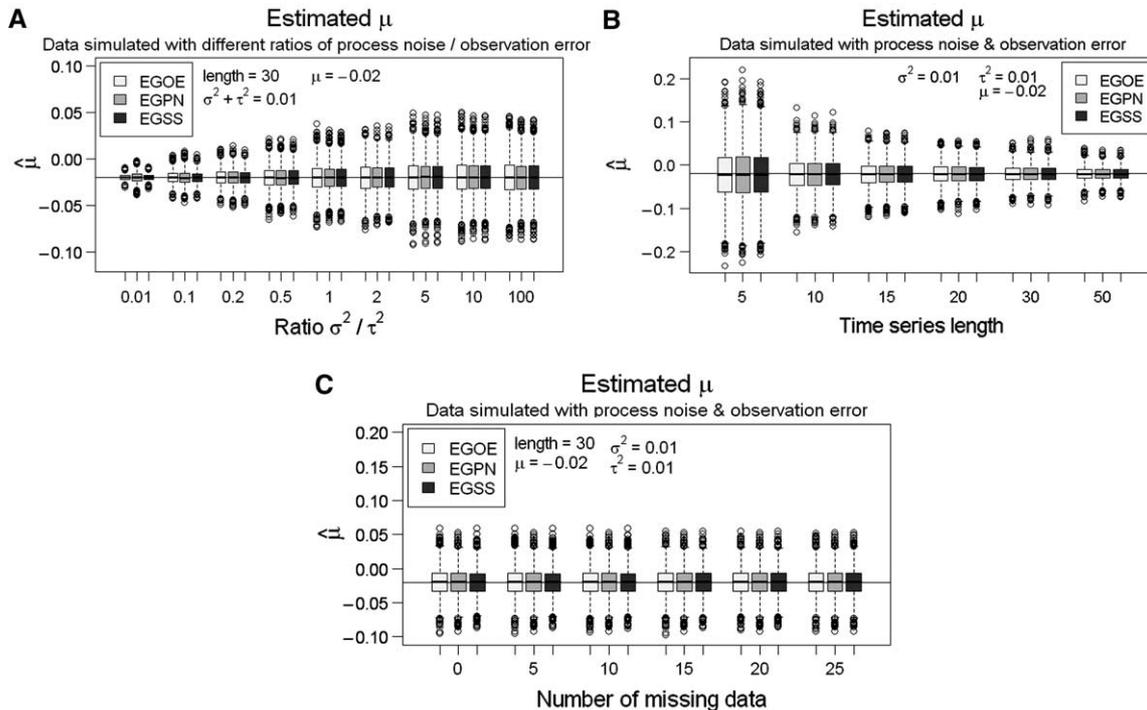
Figure 1. Estimation of trend (μ) using three estimators of exponential growth from time series of abundance data: EGOE (exponential growth with observation error), EGPN (exponential growth with process noise), EGSS (exponential growth state space). The horizontal line shows the true $\mu = -0.02$ (error bars extend to the most extreme data point in 5000 simulations which is no more than 1.5 times the interquartile range from the box). Performance is evaluated under varying conditions: (A) ratio of process (σ^2) to sampling (τ^2) variance ranging from 0.01 (trivial process variance) to 100 (trivial sampling variance); (B) time series length from 5 to 50 time steps; (C) for a time series of 30 time steps, missing abundance data ranging from none to 25 missing out of 30. Results here are typical of a range of values of μ, σ^2, τ^2 and time series lengths (Supplementary material Appendix 2).

with increasing time series length (Fig. 2B) and decreased number of missing values (Fig. 2C).

Discussion

The status quo for monitoring simple exponential trend in wild populations is to analyze the time series data with methods that assume either no process noise (EGOE model) or no observation error (EGPN), and to strive to collect abundance data each year. We recommend two fundamental shifts to this paradigm. First, we encourage the use of the EGSS model, extended to accommodate missing values (Supplementary material Appendix 1). Secondly, we suggest that researchers need not be reluctant to skip sampling in some years if the saved sampling effort could be focused on improving the remaining abundance estimates or on increasing time series length.

For a wide range of realistic conditions we found EGSS to be a superior estimator, performing well even with short time series, missing data, and a range of process and sampling variance. The most popular method for estimating exponential trend (i.e. EGOE) performed the worst of the three methods evaluated, with strongly biased small confidence intervals when process noise became even 1/10 as large as observation error. From an applied perspective this means that using EGOE will impart false confidence in what could be a qualitatively wrong trend, inferring for

example that a population is declining when it is actually increasing or stationary. This underscores the fact that the exponential growth model underlying EGOE is deterministic, assuming that none of the variance arises from environmental noise, so CIs will only attain the claimed coverage when process variance is lacking. Also disturbing was our finding that increasing the time series length or ensuring that no abundances are missing from the time series did not improve coverage of the EGOE estimate of trend; in fact the coverage only got worse.

The EGPN, commonly called the diffusion approximation, has received critical evaluation as an estimator of exponential trend parameters in PVA predictions based on abundance data (Holmes and Fagan 2002, Elderd et al. 2003, Holmes 2004). We find that the EGPN performs reasonably well in estimating trend as long as the ratio of process to observation variance is greater than 5. However, because the EGPN model assumes that all variation in the time series comes from process noise (Supplementary material Appendix 1), the presence of substantial observation error leads to confidence intervals of trend that are too wide. The resulting management inferences would too often include 0 and thereby miss actual increases or decreases. Like EGOE, coverage of trend estimates with EGPN becomes worse (although only slightly) with longer or more complete time series.

In most cases a time series will be influenced by both process and observation error. The EGSS both estimates
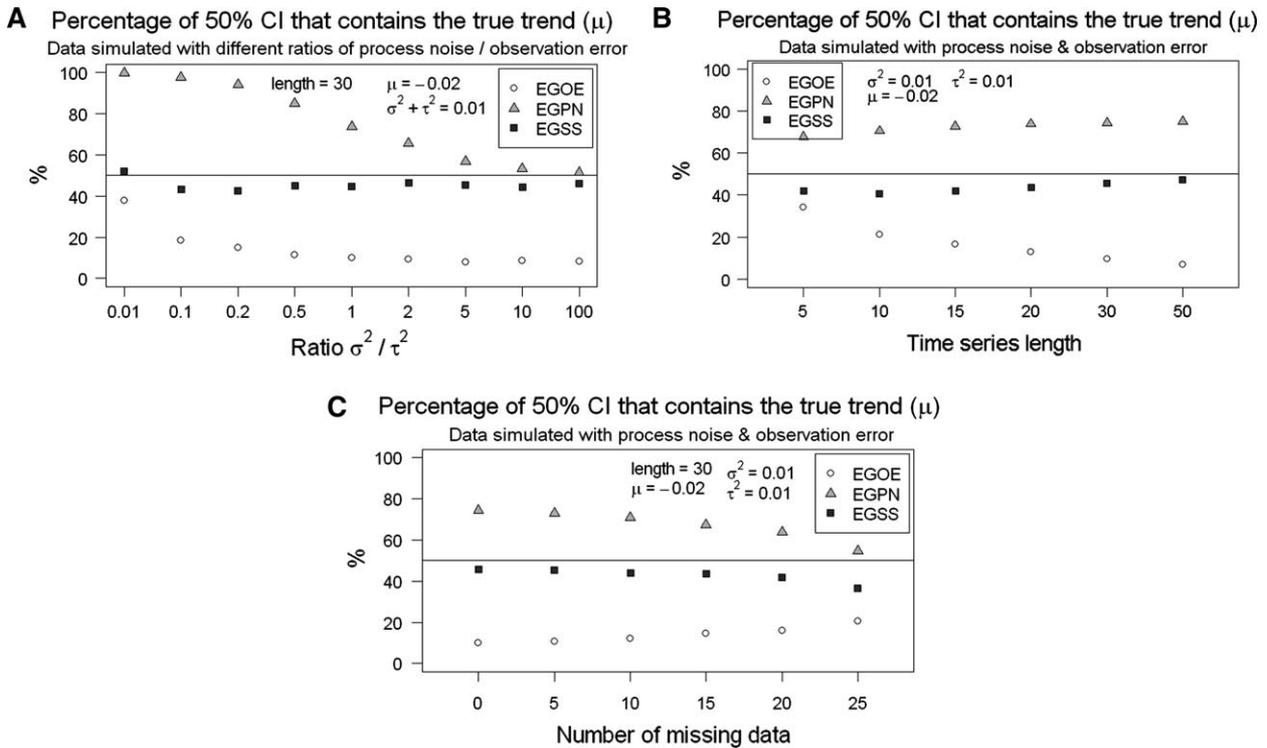
Figure 2. Estimation of coverage of confidence intervals of trend (μ). Scenarios are similar to Fig. 1, except that dependent variable is the percentage (of 5000 simulations) of trend 50% CI's that contains the true trend. The horizontal line shows the expected 50% coverage.

and incorporates both forms of variation. If observation error predominates, EGSS performs only slightly worse than EGOE, and if process noise predominates, EGSS is only slightly worse than EGPN. Except for those rare cases where an investigator knew for certain that a time series was entirely dominated by either process noise or observation error, the EGSS should be used to estimate trend. Furthermore, coverage of the CI for trend under the EGSS model improves with time series length. It also improves with fewer missing observations in the time series.

The EGSS model, as adapted for missing data, performed well for trend estimation even with approximately half of the abundance estimates missing from the time series. This implies that in some cases the collection of data every time step (e.g. year) may be less important than the quality of abundance estimates and length of time series. In other words, exponential trend estimation may well be improved if the money spent on data collection each consecutive year could instead be used to collect fewer, better estimates of abundance over the same total time duration, or to extend the time duration of the monitoring farther. Such a strategy is in stark contrast to many monitoring programs in management agencies, which strive to collect data every year to keep the time series complete, even if the data are poor (Hauser et al. 2006).

As for the chronic question of 'How much data is enough?', the minimum data requirements for trend parameter estimation are four data points for the EGSS and three for the EGPN and EGOE. Our simulations showed that the EGSS performed well (i.e. unbiased estimate of μ and proper confidence intervals) with as short a time series as 10 and approximately half of the

abundances missing; this implies that a 10-year time series with five data points could be considered a minimum for trend estimation using EGSS under the conditions we considered, and assuming that the density independent model is adequate. However, a better approach to determining sample design would calculate confidence intervals for μ for data simulated with trial parameter values under different time series lengths.

We emphasize that we were by no means comprehensive in our assessment of various ways that exponential population growth rate or trend might be estimated from count data (Pradel 1996, Clark and Bjornstad 2004, Ryding et al. 2007). For example, much more detailed trend analyses derived from log-linear models exist (Thomas and Martin 1996, Link and Sauer 1998, Bart et al. 2003, Sæther and Engen 2004), and for linear state-space models in the presence of autocorrelated environmental variables in density dependent populations (Lindén and Knape 2009). Our goal was to focus on the simple case where a series of abundance data without covariates for a single location are assessed for estimating trend and calculating an associated CI. This scenario is widely used for introducing students to the practice of estimating growth rate from field data, and is also applicable to many surveillance monitoring programs and population assessments (Table 1).

Trend analyses of monitoring data will likely proliferate with continued human-caused stresses on plants and animals (Balmford et al. 2003, Marsh and Trenham 2008), and yet both our review of first principles of model construction and our simulations indicate that the most widely used method of log-linear regression of abundance data against time (i.e. EGOE) performs poorly if variance in

the time series includes meaningful amounts of process variation. We find that a state space model (EGSS) was superior to either the EGOE or EGPN for the same length of sampling. The robustness of the EGSS model to the absence of more than half of the counts in the time series implies that in some cases improved estimates of trend and its variance may be obtained by skipping some consecutive years in a monitoring program and putting the money saved into extending the time series or improving estimates for each year that data are collected.

References

Balmford, A. et al. 2003. Measuring the changing state of nature. – Trends Ecol. Evol. 18: 326–330.

Bart, J. et al. 2003. Estimating population trends with a linear model. – Condor 105: 367–372.

Bence, J. R. 1995. Analysis of short time series: correcting for autocorrelation. – Ecology 76: 628–639.

Brook, B. W. and Bradshaw, C. J. A. 2006. Strength of evidence for density dependence in abundance time series of 1198 species. – Ecology 87: 1445–1451.

Brook, B. W. et al. 2006. Minimum viable population sizes and global extinction risk are unrelated. – Ecol. Lett. 9: 375–382.

Buenau, K. E. and Gerber, L. R. 2004. Developing recovery and monitoring strategies for the endemic mount Graham red squirrels (*Tamiasciurus hudsonicus grahamensis*) in Arizona. – Anim. Conserv. 7: 17–22.

Caughley, G. 1977. Analysis of vertebrate populations. – Wiley.

Clark, J. S. and Bjornstad, O. N. 2004. Population time series: process variability, observation errors, missing values, lags and hidden states. – Ecology 85: 3140–3150.

Dennis, B. and Otten, M. R. M. 2000. Joint effects of density dependence and rainfall on abundance of San Joaquin kit fox. – J. Wildlife Manage. 64: 388–400.

Dennis, B. et al. 1991. Estimation of growth and extinction parameters for endangered species. – Ecol. Monogr. 61: 115–143.

Dennis, B. et al. 2006. Estimating density dependence, process noise, and observation error. – Ecol. Monogr. 76: 323–341.

Eberhardt, L. L. and Simmons, M. A. 1992. Assessing rates of increase from trend data. – J. Wildlife Manage. 56: 603–610.

Elderd, B. D. et al. 2003. The problems and potential of count-based population viability analyses. – In: Brigham, C. A. and Schwartz, M. W. (eds), Population viability in plants: conservation, management, and modeling of rare plants. Springer, pp. 173–202.

Fagan, W. F. et al. 2001. Characterizing population vulnerability for 758 species. – Ecol. Lett. 4: 132–138.

Fisher, D. O. et al. 2000. Population dynamics and survival of an endangered wallaby: a comparison of four methods. – Ecol. Appl. 10: 901–910.

Gaston, K. J. and Nicholls, A. O. 1995. Probable times to extinction of some rare breeding bird species in the United Kingdom. – Proc. R. Soc. Lond. B 259: 119–123.

Gerrodette, T. 1987. A power analysis for detecting trends. – Ecology 68: 1364–1372.

Grange, S. et al. 2004. What limits the Serengeti zebra population? – Oecologia 140: 523–532.

Hauser, C. E. et al. 2006. Should managed populations be monitored every year? – Ecol. Appl. 16: 807–819.

Hinrichsen, R. A. 2009. Population viability analysis for several populations using multivariate state-space models. – Eur. Conf. Ecol. Modell., pp. 1197–1202.

Holmes, E. E. 2001. Estimating risks in declining populations with poor data. – Proc. Natl Acad. Sci. USA 98: 5072–5077.

Holmes, E. E. 2004. Beyond theory to application and evaluation: diffusion approximations for population viability analysis. – Ecol. Appl. 14: 1272–1293.

Holmes, E. E. and Fagan, W. E. 2002. Validating population viability analysis for corrupted data sets. – Ecology 83: 2379–2386.

Inchausti, P. and Halley, J. 2001. Investigating long-term ecological variability using the global population dynamics database. – Science 293: 655–657.

James, F. C. et al. 1996. New approaches to the analysis of population trends in land birds. – Ecology 77: 13–27.

Lande, R. et al. 2003. Stochastic population dynamics in ecology and conservation. – Oxford Univ. Press.

Largo, E. et al. 2008. Can ground counts reliably monitor ibex *Capra ibex* populations? – Wildlife Biol. 14: 489–499.

Lindén, A. and Knape, J. 2009. Estimating environmental effects on population dynamics: consequences of observation error. – Oikos 118: 675–680.

Lindley, S. T. 2003. Estimation of population growth and extinction parameters from noisy data. – Ecol. Appl. 13: 806–813.

Link, W. A. and Sauer, J. R. 1997. New approaches to the analysis of population trends in land birds: comment. – Ecology 78: 2632–2634.

Link, W. A. and Sauer, J. R. 1998. Estimating population change from count data: application to the North American Breeding Bird Survey. – Ecol. Appl. 8: 258–268.

Marsh, D. M. and Trenham, P. C. 2008. Current trends in plant and animal population monitoring. – Conserv. Biol. 22: 647–655.

Meyer, A. H. et al. 1998. Analysis of three amphibian populations with quarter-century long time-series. – Proc. R. Soc. Lond. B 265: 523–528.

Morris, W. F. et al. 2002. Population viability analysis in endangered species recovery plans: past use and future improvements. – Ecol. Appl. 12: 708–712.

Phillips, S. S. 2000. Population trends and the koala conservation debate. – Conserv. Biol. 14: 650–659.

Pradel, R. 1996. Utilization of capture–mark–recapture for the study of recruitment and population growth rate. – Biometrics 52: 703–709.

Ryding, K. E. et al. 2007. Using time series to estimate rates of population change from abundance data. – J. Wildlife Manage. 71: 202–207.

Sabo, J. L. et al. 2004. Efficacy of simple viability models in ecological risk assessment: does density dependence matter? – Ecology 85: 328–341.

Sæther, B. E. and Engen, S. 2004. Stochastic population theory faces reality in the laboratory. – Trends Ecol. Evol. 19: 351–353.

Staples, D. F. et al. 2004. Estimating population trend and process variation for PVA in the presence of sampling error. – Ecology 85: 923–929.

Staudenmayer, J. and Buonaccorsi, J. P. 2006. Measurement error in a random walk model with applications to population dynamics. – Biometrics 62: 1178–1189.

Thomas, L. and Martin, K. 1996. The importance of analysis method for breeding bird survey population trend estimates. – Conserv. Biol. 10: 479–490.

Watson, L. H. et al. 2005. Population viability of Cape mountain zebra in Gamka Mountain Nature Reserve, South Africa: the influence of habitat and fire. – Biol. Conserv. 122: 173–180.

Supplementary material (available online as Appendix O17839 at www.oikos.ekol.lu.se/appendix). Appendix 1. Description of three exponential growth models allowing unequal intervals in the time series, with computer program in R. Appendix 2. Results from some other parameter combinations.

Appendix 1

Description of three exponential growth models allowing unequal intervals in the time series, with computer program in R

Here we first describe the assumptions and statistical methods entailed in fitting the EGOE and EGPN models to data, and then describe the EGSS model, including a novel extension of its statistical inferences to accommodate unequal time intervals in the trajectory of abundance observations. As in the main text, we use the statistical notation convention of lower case to denote data (a particular outcome of a stochastic model) as well as constants, and upper case to denote the stochastic process (random variable) that generates data. So, we write $n(0)$, $n(t_1)$, …, $n(t_q)$ for a recorded time series of population abundances (the data) at times 0 (= t_0), t_1, …, t_q, and $N(t)$ for a random population abundance at time t with some associated probability distribution.

EGOE model

We first consider the case where the population abundances in the time series $(n(0), n(t_1), …, n(t_q))$ contain observation or sampling error but negligible process variance in the form of environmental or individual variation. The lognormal distribution is an often reasonable model of observation error because it reflects the heterogeneity of ecological sampling conditions (Dennis et al. 2006). Let $x(t) = \ln n(t)$, and let $Y(t_i)$ be a value of $x(t)$ estimated or observed with error at time t_i. Under these assumptions, we can write the EGOE model as

$$X(t) = x(0) + (\ln \lambda)t \qquad (A1)$$

$$Y(t_i) = x(t_i) + F_i \qquad (A2)$$

where F_i has a normal distribution with mean 0 and variance τ^2 ($F_i \sim \text{normal}(0, \tau^2)$). Generally, pure observation or sampling error is independent across sampling occasions, so $F_0, F_1, …, F_q$ can be assumed to be uncorrelated. Presence of autocorrelated errors would tend to suggest that the observations are influenced by stochastic process noise in the population in addition to observation error and would warrant use of a different model. An equivalent way of writing the model for $Y(t_i)$ is

$$Y(t_i) \sim \text{normal}(\beta \mu t_i +, \tau^2) , \qquad (A3)$$

where $\beta = x(0)$ and $\mu = \ln \lambda$, with $Y(0), Y(t_1), …, Y(t_q)$ independent. The parameter μ is the trend parameter and can be interpreted as the expected difference in observations separated by one time unit:

$$E[Y(t) - Y(t - 1)] = \mu \qquad (A4)$$

The variance of such a difference would be $2\tau^2$.

This statistical model for log-scale linear trend and observation-error-only is that of ordinary linear regression. For convenience we denote the time series of estimated log-abundances data by $y(0) = y_0$, $y(t_1) = y_1$, …, $y(t_q) = y_q$. The likelihood function for the unknown parameters β, μ, and τ^2 is a product of normal probability density functions:

$$L\left(\beta, \mu, \tau^2\right) = \prod_{i=0}^{q} \left(\tau^2 2\pi\right)^{-1/2} \exp\left\{ -\frac{\left[y_i - \left(\beta + \mu t_i\right)\right]^2}{2\tau^2} \right\} \qquad (A5)$$

This likelihood function is identical to the likelihood for a linear regression using the y_i's as the response variable values and the t_i's as the predictor variable values. The maximum likelihood (ML) estimates for β and μ are the values that jointly maximize the likelihood and are the familiar least squares estimates; that is, trend (μ) is estimated by the slope of the linear regression. The bias-corrected ML estimate, $\hat{\tau}^2$, of the variance parameter is the sum of squared model residuals divided by $q - 1$ (mean squared error in most regression packages), the total number of observations being $q + 1$.

The confidence interval is calculated as $\hat{\mu} \pm t_{\alpha/2, q-1} \text{SE}(\hat{\mu})$, where $t_{\alpha/2, q-1}$ is the $100\left[1 - \left(\alpha/2\right)\right]$th percentile of a student's t-distribution with $q - 1$ degrees of freedom and $\text{SE}(\hat{\mu})$ is:

$$\text{SE}(\hat{\mu}) = \sqrt{\frac{\hat{\tau}^2}{\sum_{i=0}^{q}\left(t_i - \bar{t}\right)^2}} \qquad (A6)$$

If the data production process justifies use of the EGOE model, that is, if stochastic variability in the underlying population can be assumed negligible, then the statistical analysis has some conveniences. All the inferences from standard linear regression, such as confidence intervals and hypothesis tests, remain valid. The point estimate and confidence interval boundaries for λ are found by exponentiating those for μ. In addition, if population abundance is appropriately indexed, that is, if the observation process has a mean proportional to population abundance, or log-abundance is estimated with an additive sampling bias (if for instance a consistent proportion of animals elude sampling), then the bias or index constant will be absorbed into the intercept parameter β. The trend estimate will remain the same.

EGPN model

The EGPN model is the "diffusion process" model of Dennis et al. (1991). The model assumes that during any small time interval dt, the population's log-scale growth rate experiences a random perturbation due to environmental variability:

$$d \ln N(t) = \left(\ln \lambda\right)dt + dB(t) \qquad (A7)$$

or

$$dX(t) = (\ln \lambda)dt + dB(t) \tag{A8}$$

Here $dB(T) \sim normal(0, \sigma^2 dt)$. The random quantity $dB(t)$ is an increment of a Brownian motion stochastic process over a small time interval, with the correlation between $dB(t_i)$ and $dB(t_j)$ equal to 0 if $t_i \neq t_j$, and σ^2 is a positive constant. Written this way as a 'stochastic differential equation,' Eq. A8 can be regarded as just a recipe for simulating a trajectory of population abundance for specified values of λ and σ^2 (for instance, Allen 2003, Higham 2001): (1) Over a tiny time interval dt, generate an increment $dX(t)$ of log-abundance from a normal distribution with a mean of $(\ln \lambda)dt$ and a variance of $\sigma^2 dt$. (2) Update log-population abundance as $X8t) + dX(t)$ and update time as $t + dt$. (3) Return to step (1) and repeat the process for a new time interval, generating a new normal random number uncorrelated with the previous one.

The process $X(t)$ is the same as Brownian motion with drift (Dennis et al. 1991); the original scale process $N(t) = \exp(X(t))$ is often termed geometric Brownian motion, especially in investment analysis (Ross 2002). The model is a time series model that induces dependence among the population abundances N)0), $N(t_1)$, ... $N(t_q)$. However, the process $N(t)$ (along with $X(t)$) is a Markov process, that is, the statistical properties of the distribution of $N(t + s)$ given the value of $N(t)$ do not depend on observations previous to $N(t)$. Furthermore, it can be shown that the conditional distribution of $X(t + s)$ given that $X(t)$ is fixed at the value $x(t)$ is a normal distribution:

$$X(t + s) \mid \{X(t) = x(t)\} \sim normal(x(t) + \mu, \sigma^2 s)) \tag{A9}$$

where $\mu = \ln \lambda$ and $s \geq 0$. In discrete time, with equal time intervals between observations, the process $X(t)$ is a form of nonstationary autoregressive model that has served as the null hypothesis in statistical tests of density dependence (Dennis and Taper 1994). Our trend parameter defined by $\mu = \ln \lambda$ in the EGPN model is interpreted as the expected change of a population's log-abundance in one time unit. For the EGSS model, the geometric mean of $N(t)$ given by

$$\exp\{E[\log N(t)]\} = \exp\{E[X(t)]\} = \exp(\mu t) \tag{A10}$$

characterizes 'typical' sample paths of $N(t)$ better than does the mean population size

$$E[N(t)] = \exp[(\mu + \sigma^2/2)t] \tag{A11}$$

The geometric mean happens to be the same as the median for the highly skewed lognormal distribution of population size. The quantity e^μ gives a better portrait of the growth rate of the bulk of the sample paths than does $e^{\mu + \sigma^2/2}$ (Tuljapurkar 1989, Dennis et al. 1991). The environmental 'process noise' is additive on the logarithmic scale and produces proportional variability at all population abundances (unlike 'demographic' process noise which produces essentially deterministic behavior at large population abundances). Unlike observation error, process noise causes stochastic fluctuations regardless of whether or not the population is observed.

The observations $x_0, x_1, \ldots, x_q$ recorded at times $0, t_1, \ldots, t_q$, are now assumed to be log-transformed population census values (i.e. no observation error). The likelihood function used for this model is the joint probability density of $x_1, \ldots, x^q$, given the starting population of x_0. Due to the above-mentioned Markov property, the likelihood function can be conveniently written as a product of conditional normal densities:

$$L(\mu, \sigma^2) = \prod_{j=1}^{q} (\sigma^2 s_j 2\pi)^{-1/2} \exp\left\{ -\frac{[x_j - (x_{j-1} + \mu s_j)]^2}{2\sigma^2 s_j} \right\} \tag{A12}$$

Here $s_j = t_j - t_{j-1}$. This is the likelihood function for the log-transformed abundances; the likelihood displayed in Dennis et al. (1991) is a product of lognormal densities for the untransformed abundances. Either likelihood yields the same parameter estimates. Here the initial population abundance is not an unknown parameter, because there is no observation error in the model. As described by Dennis et al. (1991), the ML estimates for μ and σ^2 can be calculated with formulas (Eq. 24–26 in Dennis et al. 1991) or by simple linear regression with the intercept fixed at zero using the transformed values $(x_j - x_{j-1})/\sqrt{s_j}$ as the response variable, and the transformed time intervals $\sqrt{s_j}$ as the predictor variable. The resulting slope estimate is the ML estimate of μ, and the mean squared error is the bias-corrected ML estimate of σ^2. The time intervals $(s_1, s_2, \ldots)$ do not need to be equal to obtain the ML estimates of μ and σ^2.

Note that this maximum likelihood estimator of trend reduces (see equation 24 in Dennis et al. (1991) to a form using only the first $(n(t_0))$ and last $(n(t_q))$ abundance estimates in the time series, and the total duration of the survey (i.e. the time interval between the first and last observations):

$$\hat{\mu} = \frac{1}{t_q - t_0} \ln\left(\frac{n(t_q)}{n(t_0)} \right) \tag{A13}$$

All counts are used, however, to estimate the variance.

The confidence interval is calculated as $\hat{\mu} \pm t_{\alpha/2, q-1} SE(\hat{\mu})$, where $t_{\alpha/2, q-1}$ is the $100[1 - (\alpha/2)]$th percentile of a student's t-distribution with $q - 1$ degrees of freedom and $SE(\hat{\mu})$ is:

$$SE(\hat{\mu}) = \sqrt{\frac{\hat{\sigma}^2}{t_q - t_0}} \tag{A14}$$

EGSS model, extended to unequal time intervals between observations

An exponential growth model with both process noise and observation error can be constructed by combining the EGOE and the EGPN models, creating a state space model with an unobserved stochastic population component, $X(t)$, and an observed or estimated component, $Y(t)$:

$$DX(t) = (\ln \lambda)dt + dB(t) \tag{A15}$$

$$Y(t_i) = X(t_i) + F_i \tag{A16}$$

Here $dB(t) \sim normal(0, \sigma^2 dt)$ and $F_i \sim normal(0, \tau^2)$, with no auto- or cross-correlations. In other words, the unobserved population is governed by the EGPN model, but the data are created with observation error as in the EGOE model. The EGSS model was proposed by Holmes (2001), and parameter estimation was studied by Lindley (2003), Staples et al. (2004) and Dennis et al. (2006).

These authors defined the model in discrete time, with equal intervals between samples. Equation A15 and A16 incorporate the continuous time version of $X(t)$ and constitute an extension of the EGSS model to unevenly spaced sampling times.

In simulations, the ML parameter estimates for the EGSS model tend to have a persistent small-sample bias (Staples et al. 2004). An improvement to ML estimation can be made, akin to restricted maximum likelihood (REML) estimation in mixed effects models. Staples et al. (2004) showed that REML estimates based on second differences of the observations were superior to ML estimates for the EGSS model; their ML and REML estimates, however, were based on equally spaced time intervals between observations. The ML and REML estimates are extended to unequally spaced sampling times as follows. Staudenmayer and Buonaccorsi (2006) provide a more theoretical treatment.

The model has four unknown parameters: β (= x_0), μ (=ln λ), σ^2 and τ^2. The observations $Y(0)$, $Y(t_1)$, …, $Y(t_q)$ are dependent but not Markovian, and they have a joint multivariate normal distribution in which the mean of $Y(t)$ is $\beta + \mu t$, the variance of $Y(t)$ is $\sigma_2 t + \tau_2$, and the covariance of $Y(t)$ and $Y(t + s)$ is $\sigma^2 t$, for any times $t, s \geq 0$. For ML estimation, the multivariate normal log-likelihood function given by

$$\ln L\left(\beta, \mu, \sigma^2, \tau^2\right) = -\frac{(q+1)}{2}\ln(2\pi) - \frac{1}{2}\ln\left(|V|\right) - \frac{1}{2}\left(y - m\right)'V^{-1}\left(y - m\right) \quad \text{(A17)}$$

is used. Here the data values $y_0, y_1, …, y_q$ recorded at times 0, t_1, …, t_q are the elements of the column vector $\mathbf{y}$, the means β, $\beta + \mu t_1$, $\beta + \mu t_2$, … are the elements in the column vector $\mathbf{m}$, the variances $\tau^2, \sigma^2 t_1 + \tau^2, \sigma^2 t_2 + \tau^2, …$ are the elements on main diagonal in the matrix $\mathbf{V}$, and the (i, j)th and (j, i)th covariance elements in $\mathbf{V}$ are both $\sigma^2 t_i$, where $t_i < t_j$:

$$V = \begin{bmatrix} \tau^2 & 0 & 0 & 0 & \cdots \\ 0 & \sigma^2 t_1 + \tau^2 & \sigma^2 t_1 & \sigma^2 t_1 & \sigma^2 t_1 \\ 0 & \sigma^2 t_1 & \sigma^2 t_2 + \tau^2 & \sigma^2 t_2 & \sigma^2 t_2 \\ 0 & \sigma^2 t_1 & \sigma^2 t_2 & \sigma^2 t_3 + \tau^2 & \sigma^2 t_3 \\ \vdots & \sigma^2 t_1 & \sigma^2 t_2 & \sigma^2 t_3 & \ddots \end{bmatrix} \quad \text{(A18)}$$

The ML parameter estimates are calculated from the log-likelihood function with numerical maximization (for instance, with the 'optim' function in R; R Core Development Team 2006).

The REML estimates are found by rescaling the differences of the observations. The scaled first differences of the observations, defined by

$$W_i = \frac{1}{s_i}\left[Y(t_i) - Y(t_{i-1})\right], i = 1, 2, …, q \quad \text{(A19)}$$

with $s_i = t_i - t_{i-1}$, have a multivariate normal distribution with a constant mean vector equal to $\mathbf{j}\mu$ ($\mathbf{j}$ is a column vector of ones), that is, $E(W_i) = \mu$. The variance-covariance matrix $\mathbf{V}_1$ of the W_i's is found as the matrix product

$$V_1 = D_1 V D_1',$$

where

$$D_1 = \begin{bmatrix} -\dfrac{1}{s_1} & \dfrac{1}{s_1} & 0 & 0 \\ 0 & -\dfrac{1}{s_2} & \dfrac{1}{s_2} & 0 \\ 0 & 0 & \ddots & \ddots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} \quad \text{(A20)}$$

is a $q \times (q + 1)$ matrix with all elements zero except for those on the two long diagonals, and with $\mathbf{V}$ being the variance-covariance matrix of the $Y(t_i)$'s. By carrying out the matrix multiplications we find the resulting variances and covariances of the W_i's to be: $V(W_i) = \left(\sigma^2/s_i\right) + \left(2\tau^2/s_i^2\right)$, $\text{Cov}(W_i, W_{i+1}) = -\tau^2/\left(s_i s_{i+1}\right)$, $\text{Cov}(W_i, W_{i+k}) = 0$ ($k \geq 2$). Note that one parameter, the initial condition β, has been eliminated from the distribution of the W_i's. When the intervals s_i are equal, the model for the W_i's is equivalent to a normal linear mixed model in which the variance-covariance matrix has constant main diagonal, equivalent and constant subdiagonals, and zeros elsewhere. As Staples et al. (2004) noted, existing software packages for linear mixed models often allow such a 'banded Toeplitz' structure to be specified , and so REML estimates of EGSS parameters can be readily obtained (programs provided by Staples et al. 2004). When the intervals s_i are unequal, however, the variances and covariances of the W_i's are unequal as well. It is not clear how to adapt current software packages for linear mixed models to accommodate the varied structure. Instead, REML estimates for the EGSS model with varying time intervals are obtained by numerical maximization.

The REML estimates are constructed from the second differences of the observations (first differences of the W_i's):

$$U_i = W_{i+1} - W_i, i = 1, 2, …, q - 1 \quad \text{(A21)}$$

It can be shown that $U_1, U_2, …, U_{q-1}$ have a multivariate normal distribution in which the mean of each U_i is 0, and the variance-covariance matrix is found as

$$V_2 = D_2 V_1 D_2' = D_2 D_1 V D_1' D_2' \quad \text{(A22)}$$

where

$$D_2 = \begin{bmatrix} -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & \ddots & \ddots \\ \vdots & \vdots & \vdots & \vdots \end{bmatrix} \quad \text{(A23)}$$

is an ordinary $(q - 1) \times q$ differencing matrix. Let

$$u_i = w_{i+1} - w_i = \frac{1}{s_{i+1}}\left(y_{i+1} - y_i\right) - \frac{1}{s_i}\left(y_i - y_{i-1}\right), i = 1, 2, …, q - 1 \quad \text{(A24)}$$

be the twice-differenced data. The log-likelihood for the REML estimates is then given by

$$\ln L\left(\sigma^2, \tau^2\right) = -\frac{(q-1)}{2}\ln(2\pi) - \frac{1}{2}\ln\left(|V_2|\right) - \frac{1}{2}u'V_2^{-1}u \quad \text{(A25)}$$

in which $\mathbf{u}$ is a column vector containing the values $u_1, u_2, …, u_{q-1}$, and $\mathbf{V}_2$ is the variance-covariance matrix given above. Now

3

two parameters, β and μ, have been eliminated by the differencing; however, the REML log-likelihood must still be maximized numerically for the REML estimates $\hat{\sigma}^2$ and $\hat{\tau}^2$. The estimate of the trend parameter μ then is calculated as

$$\hat{\mu} = \frac{j'V_1^{-1}w}{j'V_1^{-1}j} \tag{A26}$$

where **w** is the column vector containing w_1, w_2, …, w_q, with the elements in V_1 (the variance-covariance matrix for the W_i's) evaluated using the REML estimates $\hat{\sigma}^2$ and $\hat{\tau}^2$. The variance in the trend estimate is:

$$\mathrm{Var}(\hat{\mu}) = \frac{1}{\left(j'V_1^{-1}j\right)} \tag{A27}$$

and this variance can be estimated using the REML estimates $\hat{\sigma}^2$ and $\hat{\tau}^2$. Finally, the estimate of the initial condition β becomes

$$\hat{\beta} = \frac{j'V^{-1}\left(y - t\hat{\mu}\right)}{j'V^{-1}j}. \tag{A28}$$

Here $t = [t_0, t_1, ..., t_q]'$, with the elements in **V** (the variance-covariance matrix for the $Y(t_i)$'s) evaluated using the REML estimates $\hat{\sigma}^2$ and $\hat{\tau}^2$.

For the numerical maximizations required for ML or REML estimates, initial parameter values are required for starting the iterations. Ideally, for a given data set the ML or REML estimates should be verified by trying many initial parameter values, because the EGSS model can produce likelihoods with multiple local maxima. However some strategy for automating the calculation of initial values is important for processing many data sets or for repetitive fitting techniques such as bootstrapping. One possibility is to fit both the EGOE and EGPN models and then cut in half the resulting estimates of τ^2 and σ^2. Another possibility is to use the one-step covariance of the vector of first differences W_i to estimate an initial value for τ^2 (see Eq. 33 in Dennis et al. 2006) and then use the variance of W_i to estimate an initial value for σ^2 (Eq. 32 in Dennis et al. 2006). Also, a standard numerical trick to prevent negative parameter values in iterations is to use transformed parameters, i.e. $\sigma^2 = \exp(\theta)$ and $\tau^2 = \exp(\delta)$, and maximize log-likelihoods for θ and δ as real-valued parameters.

An asymptotic $100\left(1 - \alpha\right)\%$ confidence interval is calculated as $\hat{\mu} \pm z_{\alpha/2} \mathrm{SE}(\hat{\mu})$ where $z_{\alpha/2}$ is the $100\left[1 - \left(\alpha/2\right)\right]$th percentile of the standard normal distribution, and $\mathrm{SE}(\hat{\mu})$ is:

$$\mathrm{SE}(\hat{\mu}) = \sqrt{\mathrm{Var}(\hat{\mu})} \tag{A29}$$

If a statistical test for zero trend is desired, we suggest using $\mathrm{SE}(\hat{\mu})$ and a standard normal percentile in an equivalence testing framework (Dixon and Pechmann 2005). The null hypothesis is that a substantial trend is present (i.e. μ is outside of a fixed, specified interval containing zero), and the alternative hypothesis is that trend is negligible for practical purposes (μ is inside the specified interval containing zero). Dixon and Pechmann (2005) give further details about equivalence testing and illustrate the concept with tests for the trend parameter in the EGOE model.

Although finite samples in many models can lead to ML (or REML) estimates with skewed sampling distributions accompanied by finite-sample bias, leading to poor CI coverage, we found that the sampling distribution of the REML estimates of μ were symmetrical and normal. This suggests that μ is being estimated well and that the asymptotic theory of ML/REML estimation is providing useful approximations for CI construction.

References

Allen, L. J. S. 2003. An introduction to stochastic processes with applications to biology. – Pearson Education, Upper Saddle River.

Dennis, B. and Taper, M. L. 1994. Density dependence in time series observations of natural populations: estimation and testing. – Ecol. Monogr. 64: 205–224.

Dennis, B. et al. 1991. Estimation of growth and extinction parameters for endangered species. – Ecol. Monogr. 61: 115–143.

Dennis, B. et al. 2006. Estimating density dependence, process noise, and observation error. – Ecol. Monogr. 76: 323–341.

Dixon, P. M. and Pechmann, J. H. K. 2005. A statistical test to show negligible trend. – Ecology 86: 1751–1756.

Higham, D. J. 2001. An algorithmic introduction to numerical simulation of stochastic differential equations. – Siam Rev. 43: 525–546.

Holmes, E. E. 2001. Estimating risks in declining populations with poor data. – Proc. Natl Acad. Sci. USA 98: 5072–5077.

Lindley, S. T. 2003. Estimation of population growth and extinction parameters from noisy data. – Ecol. Appl. 13: 806–813.

Ross, S. M. 2002. An elementary introduction to mathematical finance: options and other topics (2nd ed.). – Cambridge Univ. Press.

Staples, D. F. et al. 2004. Estimating population trend and process variation for PVA in the presence of sampling error. – Ecology 85: 923–929.

Staudenmayer, J. and Buonaccorsi, J. P. 2006. Measurement error in a random walk model with applications to population dynamics. – Biometrics 62: 1178–1189.

Tuljapurkar, S. 1989. An uncertain life – demography in random environments. – Theor. Popul. Biol. 35: 227–294.

Computer program, in the R language, for calculating ML and REML estimates of parameters in the EGSS model

```
#  Exponential Growth State Space model:
#  Version Nov.28.2009
#
#  R program for calculating maximum likelihood (ML) and restricted maximum
#  likelihood (REML) estimates of unknown parameters for the Exponential
#  Growth State Space (EGSS) model of stochastic population growth.
#  The model is
#
#  dX(t) = mu*dt + dB(t)
#                with dB(t) ~ normal(0,ssq*dt),
#  Y(t) = X(t) + F(t)
#                with F(t) ~ normal(0,tsq).
#
#  Here X(t) is log-population abundance, Y(t) is observed or estimated value
#  of X(t), x0, mu, ssq, tsq are parameters.  The parameter ssq is the
#  variance of the process noise, and tsq is the variance of the observation
#  error.
#
#  The model takes population abundance N(t) = exp(X(t)) to be governed by
#  a stochastic, density independent model, with the observed abundances
#  O(t) = N(t)*exp(F(t)) arising from lognormal sampling error.
#
#  User provides time series of observed population abundances o(0), o(1),
#  ..., o(q), which are log-transformed by the program into y(0), y(1), ...,
#  y(q), assumed to be a time series realization of Y(t).  Likelihood
#  function of y(0), y(1), ..., y(q) is that of a multivariate normal
#  distribution.  The observation times t_0, t_1, t_2, ..., t_q can have
#  unequal intervals.
#
#  Program computes initial parameter values for iterations.  The program
#  should be re-run for several sets of initial values, as the likelihood
#  function for the model frequently has multiple local maxima, see program
#  section 4.
#
#  Alternative programs, for observation times with equal intervals,
#  are available as an online appendix to Staples et al. (2004).
#  See also Staudenmayer and Buonaccorsi (2006) for a more theoretical
#  development.
#
#  Program citations:
#    Dennis et al. 2006.   Ecological Monographs.
#    Humbert et al. 2009.  Oikos.
#    Staples et al. 2004.  Ecology.
#    Staudenmayer and Buonaccorsi. 2006.  Biometrics.
#
```

```
#------------------------------------------------------------------
#        1. USER INPUT SECTION
#------------------------------------------------------------------
#
#  The best way to conduct these analyses is by preparing data in Excel, then
#  saving as a text file.
#
#  In Excel, in cell A1, type (Exactly) the word «Observed.t», without the
#  quotes. Put the observed abundance in column A, starting with row 2.
#  In cell B1, type (Exactly) the word «Time.t» without the quotes. Put the
#  time step identifiers that correspond to the observed population size in
#  column B, starting  with row 2.
#  Your first time step can be 0, or 1, or anything else (e.g. a year).
#
#  ** NOTE: If you have years with no data, that is fine.
#  Just omit the year and the associated abundance. However, if you have a
#  year where you sampled but got zero abundance, you cannot use these
#  approaches. **
#
#  Once the data sheet is prepared in Excel, save it as on the C:\ drive
#  as "c:\my_data.txt". It must be a tab delimited text file.
#  Excel may try to name it my_data.txt.txt or my_data.txt.xls but don't
#  let it! You can confirm that it is, indeed, a text file by double clicking
#  on it and confirming that it opens with Notepad, not Excel.
#
#  ** NOTE: You can change the drive or the name of the input file by
#  changing the second next line. **

rm(list=ls(all=TRUE))            #  Clears all objects from memory
my_data <- data.frame(read.table(«C:/my_data.txt», header=TRUE, sep=»\t»))
Observed.t <- my_data$Observed.t
Time.t <- my_data$Time.t
print.table(cbind(Observed.t,Time.t))

#  OUTPUT: 2 Files will go the same place as the input file:
#     GRAPHICS                  my_graph.png
#      INTERPRET OUTPUT         my_results.txt

#  REST OF THIS SECTION CAN BE IGNORED IF USING EXCEL INPUT
#     Example data below are American Redstart counts from North American
#     Breeding Bird Survey, record # 02014 3328 08636, 1966-95 (Table 1 in
#     Dennis et al. 2006).
#     To run this example, comment out the 5 lines immediately
#     above this paragraph, using the pound sign, and uncomment the 4 lines
#     immediately below this paragraph, but do not change anything outside
#     of the user input section.  When you uncomment the lines below, remove
#     only the pound signs at the left side of this document.  Do not remove
#     the pound signs after the semicolons.
#  Observed.t=c(18,10,9,14,17,14,5,10,9,5,11,11,4,5,4,8,2,3,9,2,4,7,4,1,2,
#     4,11,11,9,6);     #  No zeros!  (With zeros, you must use another model)
#  Time.t=c(0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23,24,
```

```
#    25,26,27,28,29);  #  Initial time can be nonzero.


#----------------------------------------------------------------------
#         2. PROGRAM INITIALIZATION SECTION
#----------------------------------------------------------------------
library(MASS);           #  Loads miscellaneous functions (ginv, etc.)
T.t=Time.t-Time.t[1];    #  Time starts at zero.
Y.t=log(Observed.t);     #  Log-transform the observations.
q=length(Y.t)-1;         #  Number of time series transitions, q.
qp1=q+1;                 #  q+1 gets used a lot, too.
S.t=T.t[2:qp1]-T.t[1:q]; #  Time intervals.
m=rep(1,qp1);            #  Will contain Kalman means for Kalman calculations.
v=rep(1,qp1);            #  Will contain variances for Kalman calculations.
sink(file = «C:/my_results.txt», append = FALSE, type = «output», split = T)
                         #  Tells the program where to write the output
#  ** CHANGE LINE ABOVE IF YOU WANT RESULTS TO GO SOMEWHERE OTHER THAN C:/ DRIVE


#----------------------------------------------------------------------
#         3. SECTION FOR DEFINING ML & REML LOG-LIKELIHOODS
#----------------------------------------------------------------------


#  ML objective function «negloglike.ml» is negative of log-likelihood;
#  the Nelder-Mead optimization routine in R, «optim», is a minimization
#  routine.  The ML objective function uses equations 24-26 from Dennis et
#  al. (2006).  The three function arguments are:  theta, vector of
#  parameters (transformed to the real line), yt, vector of time series
#  observations, and tt, vector of observation times.
negloglike.ml=function(theta,yt,tt)
{
   muu=theta[1];
   sigmasq=exp(theta[2]);    #  Constrains ssq > 0.
   tausq=exp(theta[3]);      #  Constrains tsq > 0.
   xzero=theta[4];
   q=length(yt)-1;
   qp1=q+1;
   yt=matrix(yt,nrow=qp1,ncol=1);
   vx=matrix(0,qp1,qp1);
   for (ti in 1:q)
   {
      vx[(ti+1):qp1,(ti+1):qp1]=matrix(1,1,(qp1-ti))*tt[ti+1];
   }
   Sigma.mat=sigmasq*vx;
   Itausq=matrix(rep(0,(qp1*qp1)),nrow=q+1,ncol=q+1);
   diag(Itausq)=rep(tausq,q+1);
   V=Sigma.mat+Itausq;
   mu=matrix((xzero+muu*tt),nrow=qp1,ncol=1);
   ofn=((qp1)/2)*log(2*pi)+(0.5*log(det(V)))+
      (0.5*(t(yt-mu)%*%ginv(V)%*%(yt-mu)));
   return(ofn);
}
```

```r
#  REML objective function „negloglike.reml" is negative of log-likelihood
#  for second differences of the log-scale observations.  The REML objective
#  function uses equations A18-A25 from Humbert et al. (2009).  The three
#  function arguments are:  theta, vector of parameters (transformed to the
#  real line), yt, vector of time series observations (log scale), and
#  tt, vector of observation times.  Function performs the differencing.
negloglike.reml=function(theta,yt,tt)
{
   sigsq=exp(theta[1]);           #  Constrains ssq > 0.
   tausq=exp(theta[2]);           #  Constrains tsq > 0.
   q=length(yt)-1;
   qp1=q+1;
   vx=matrix(0,qp1,qp1);
   for (ti in 1:q)
   {
      vx[(ti+1):qp1,(ti+1):qp1]=matrix(1,1,(qp1-ti))*tt[ti+1];
   }
   Sigma.mat=sigsq*vx;
   Itausq=matrix(rep(0,(qp1*qp1)),nrow=q+1,ncol=q+1);
   diag(Itausq)=rep(tausq,q+1);
   V=Sigma.mat+Itausq;
   ss=tt[2:qp1]-tt[1:q];
   D1mat=cbind(-diag(1/ss),matrix(0,q,1))+cbind(matrix(0,q,1),diag(1/ss));
   D2mat=cbind(-diag(1,q-1),matrix(0,q-1,1))+
      cbind(matrix(0,q-1,1),diag(1,q-1));
   Phi.mat=D2mat%*%D1mat%*%V%*%t(D1mat)%*%t(D2mat);
   wt=(yt[2:qp1]-yt[1:q])/ss;
   ut=wt[2:q]-wt[1:q-1];
   ofn=(q/2)*log(2*pi)+(0.5*log(det(Phi.mat)))+
      (0.5*(ut%*%ginv(Phi.mat)%*%ut));
   return(ofn);
}


#-------------------------------------------------------------------
#        4. SECTION FOR CALCULATING EGOE AND EGPN ESTIMATES
#        (FOR USE AS INITIAL VALUES)
#-------------------------------------------------------------------
# The EGOE estimates
Ybar=mean(Y.t);
Tbar=mean(T.t);
mu.egoe=sum((T.t-Tbar)*(Y.t-Ybar))/sum((T.t-Tbar)*(T.t-Tbar));
x0.egoe=Ybar-mu.egoe*Tbar;
ssq.egoe=0;
Yhat.egoe=x0.egoe+mu.egoe*T.t;
tsq.egoe=sum((Y.t-Yhat.egoe)*(Y.t-Yhat.egoe))/(q-1);


# The EGPN estimates
Ttr=sqrt(S.t);
Ytr=(Y.t[2:qp1]-Y.t[1:q])/Ttr;
mu.egpn=sum(Ttr*Ytr)/sum(Ttr*Ttr);
Ytrhat=mu.egpn*Ttr;
```

8

```
ssq.egpn=sum((Ytr-Ytrhat)*(Ytr-Ytrhat))/(q-1);
tsq.egpn=0;
x0.egpn=Y.t[1];


# Initial values for EGSS are averages of EGOE and EGPN values
mu0=(mu.egoe+mu.egpn)/2;      #  For ML only
ssq0=ssq.egpn/2;             #  For ML and REML
tsq0=tsq.egoe/2;             #  For ML and REML
x00=x0.egoe;                 #  For ML only


# To set different initial values for iterations, enter manually a value
#   after the equal sign of the concern parameter instead of the
#   automatically generated value. Then run again the line and the program
#   section 5 below.
#   Initial values near the EGOE and EGPN models are good for exploring
#   possible alternative local maxima. The values which produce the largest
#   log-likelihood should be used. To see the log-likelihood for the REML
#   estimates type:
#   EGSSreml$value[1];
#   See Dennis et al. 2006 for more details.


#---------------------------------------------------------------------
#          5. SECTION FOR CALCULATING ML & REML PARAMETER ESTIMATES
#---------------------------------------------------------------------

# The ML estimates.
EGSSml=optim(par=c(mu0,log(ssq0),log(tsq0),x00),
   negloglike.ml,NULL,method="Nelder-Mead",yt=Y.t,tt=T.t);
params.ml=c(EGSSml$par[1],exp(EGSSml$par[2]),exp(EGSSml$par[3]),
   EGSSml$par[4]);
lnlike.ml=-EGSSml$value[1];
AIC.egss=-2*lnlike.ml+2*length(params.ml);

mu.ml=params.ml[1];             # These are the ML estimates.
ssq.ml=params.ml[2];            #           --
tsq.ml=params.ml[3];            #           --
x0.ml=params.ml[4];             #           --


# The REML estimates.
EGSSreml=optim(par=c(log(ssq0),log(tsq0)),
   negloglike.reml,NULL,method="Nelder-Mead",yt=Y.t,tt=T.t);
params.reml=c(exp(EGSSreml$par[1]),exp(EGSSreml$par[2]))

ssq.reml=params.reml[1];        #  These are the REML estimates.
tsq.reml=params.reml[2];        #           --

vx=matrix(0,qp1,qp1);
for (ti in 1:q)
{
   vx[(ti+1):qp1,(ti+1):qp1]=matrix(1,1,(qp1-ti))*T.t[ti+1];
}
```

```
Sigma.mat=ssq.reml*vx;
Itausq=matrix(rep(0,(qp1*qp1)),nrow=q+1,ncol=q+1);
diag(Itausq)=rep(tsq.reml,q+1);
V=Sigma.mat+Itausq;
D1mat=cbind(-diag(1/S.t),matrix(0,q,1))+cbind(matrix(0,q,1),diag(1/S.t));
V1mat=D1mat%*%V%*%t(D1mat);
W.t=(Y.t[2:qp1]-Y.t[1:q])/S.t;
j1=matrix(1,q,1);
V1inv=ginv(V1mat);
mu.reml=(t(j1)%*%V1inv%*%W.t)/(t(j1)%*%V1inv%*%j1);
j=matrix(1,qp1,1);
Vinv=ginv(V);
x0.reml=(t(j)%*%Vinv%*%(Y.t-mu.reml*T.t))/(t(j)%*%Vinv%*%j);
Var_mu.reml=1/(t(j1)%*%V1inv%*%j1);         #  Variance of mu
mu_hi.reml=mu.reml+1.96*sqrt(Var_mu.reml);  #  95% CI for mu
mu_lo.reml=mu.reml-1.96*sqrt(Var_mu.reml);  #       --


#  Calculate estimated population sizes for EGSS model
#    with Kalman filter, for plotting.
#
#  Choose ML or REML estimates here for calculating model values
#  for plotting (by commenting out the unwanted, default is REML).
#  mu=mu.ml;  ssq=ssq.ml;  tsq=tsq.ml;  x0=x0.ml;
mu=mu.reml;  ssq=ssq.reml;  tsq=tsq.reml;  x0=x0.reml;


m[1]=x0;             #  Initial mean of Y(t).
v[1]=tsq;            #  Initial variance of Y(t).


for (ti in 1:q)   #  Loop to generate estimated population abundances
{                 #     using Kalman filter (see equations 6 & 7,
                  #     Dennis et al. (2006)).
   m[ti+1]=mu+(m[ti]+((v[ti]-tsq)/v[ti])*(Y.t[ti]-m[ti]));
   v[ti+1]=tsq*((v[ti]-tsq)/v[ti])+ssq+tsq;
}


#  The following statement calculates exp{E[X(t) | Y(t), Y(t-1),...,Y(0)]};
#    see equation 54 in Dennis et al. (2006).
Predict.t=exp(m+((v-tsq)/v)*(Y.t-m));


#  Plot the data & model-fitted values
plot(Observed.t ~ Time.t,xlab=»time»,ylab=»population abundance»,
   type=»o»,lty=»solid»,pch=1,cex=1);
                  #  Population data are circles.
points(Predict.t ~ Time.t,type=»l»,lty=»dashed»,lwd=1);
                  #  Estimated abundances are dashed line.
legend(«top», c(«Observed.t»,»Predict.t»),lty=c(1,2),pch=c(«o»,»»),bty=»n»)
                  #  Graph legend


#  Print the parameter estimates
parms.egoe=c(mu.egoe,ssq.egoe,tsq.egoe,x0.egoe); #  Collect for printing
parms.egpn=c(mu.egpn,ssq.egpn,tsq.egpn,x0.egpn); #         --
```

10

```
parms.reml=c(mu.reml,ssq.reml,tsq.reml,x0.reml);  #          --
parms.ml=c(mu.ml,ssq.ml,tsq.ml,x0.ml);            #          --
names=c("mu","ssq","tsq","x0");                   #          --
types=c("EGOE","EGPN","EGSS-ML","EGSS-REML");     #          --


#  Print stuff
matrix(cbind(parms.egoe,parms.egpn,parms.ml,parms.reml),
   nrow=4,ncol=4,byrow=TRUE,dimnames=list(types,names));


#  Print CI, default is for EGSS-REML
matrix(cbind(mu_lo.reml,mu_hi.reml),nrow=1,ncol=2,byrow=TRUE,
   dimnames=list("95% CI for MU",c("LO","HI")));


#  Print log-likelihood and AIC for EGSS ML
matrix(cbind(lnlike.ml,AIC.egss),nrow=1,ncol=2,byrow=TRUE,
   dimnames=list("EGSS ML RESULTS",c("LN-LIKELIHOOD","AIC")));


#  Plot the data & model-fitted values to a png file
png(file = "C:/my_graph.png")                     #  Open a png file for plotting
#  ** CHANGE LINE ABOVE IF YOU WANT RESULTS TO GO SOMEWHERE OTHER THAN C:/ DRIVE


plot(Observed.t ~ Time.t,xlab="time",ylab="population abundance",
   type="o",lty="solid",pch=1,cex=1);
                  #  Population data are circles.
points(Predict.t ~ Time.t,type="l",lty="dashed",lwd=1);
                  #  Estimated abundances are dashed line.
legend("top", c("Observed.t","Predict.t"),lty=c(1,2),pch=c("o",""),bty="n")
                  #  Graph legend
graphics.off()    #  Close graphics file
sink()            #  Remove output diversion to results file so output will
                  #   be sent back to the screen
```

Appendix 2

Results from some other parameter combinations

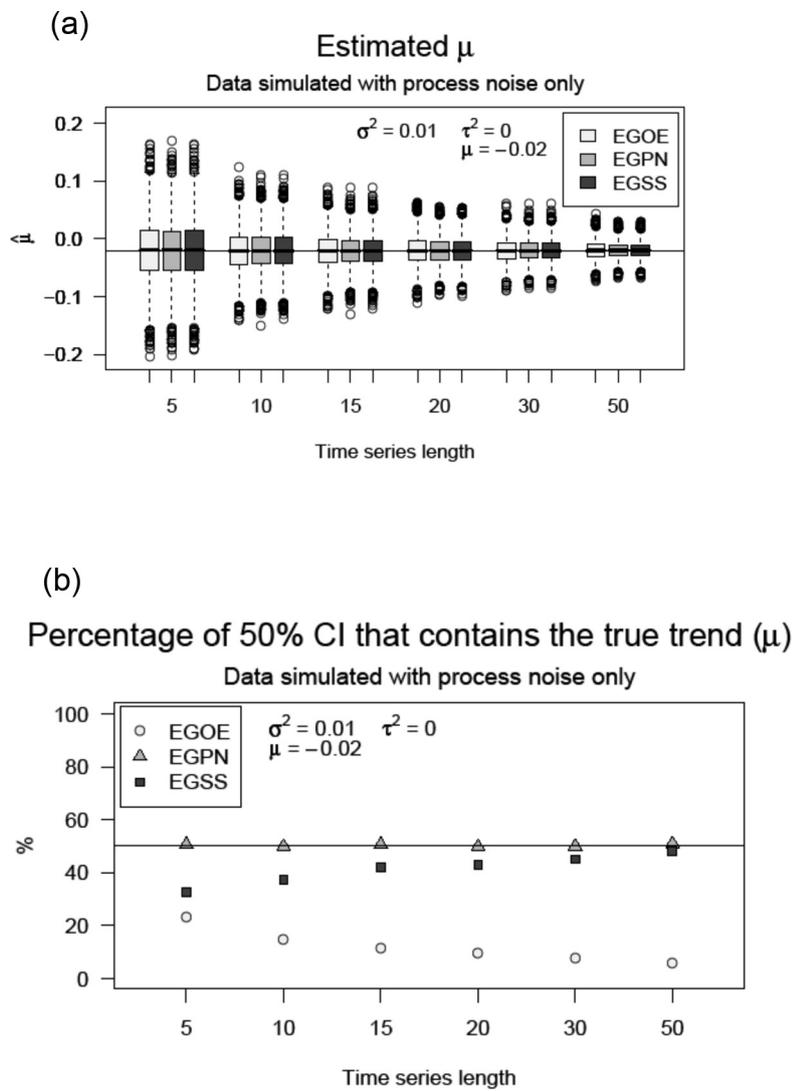Showing that inferences from the main text Fig. 1 and 2 are robust and generalizable.

(a)

(b)

Figure A1a–b. Data with process noise only: $\mu = -0.02$, $\sigma^2 = 0.01$ and $\tau^2 = 0$.
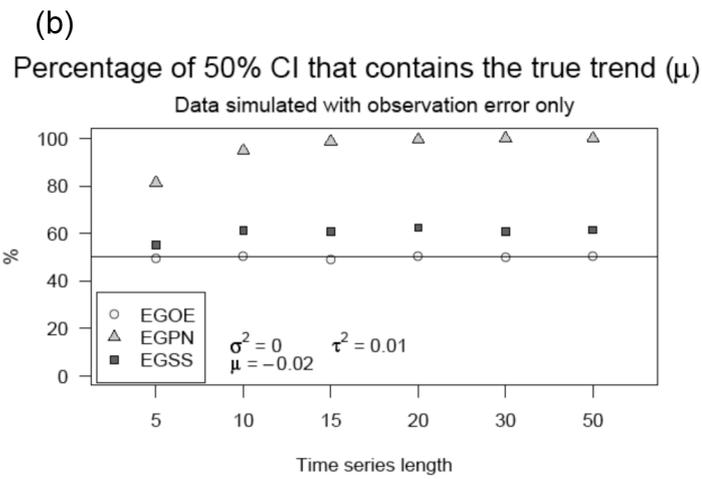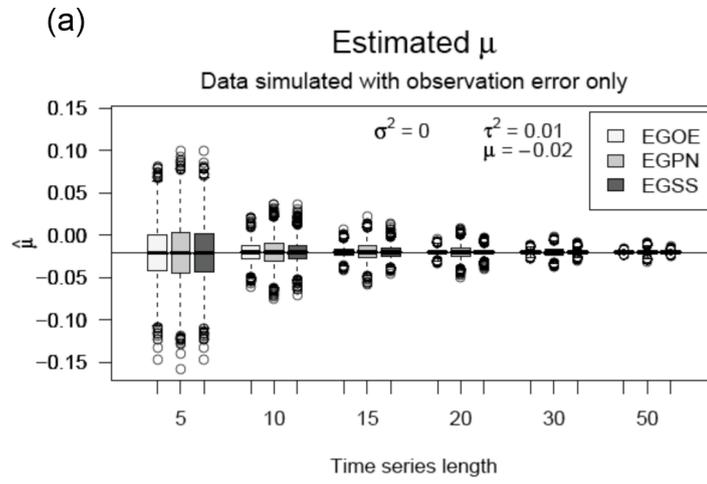
(a)
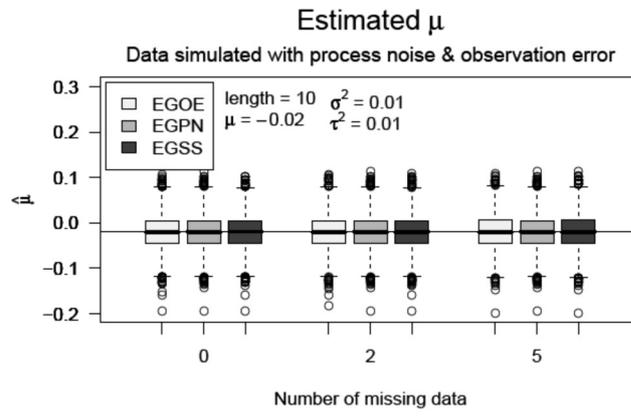
Estimated μ

Data simulated with observation error only

(b)

Percentage of 50% CI that contains the true trend (μ)

Data simulated with observation error only

Figure A2a–b. Data with observation error only: $\mu = -0.02$, $\sigma^2 = 0$ and $\tau^2 = 0.01$.

(a)

Estimated μ

Data simulated with process noise & observation error

(b)

Percentage of 50% CI that contains the true trend (μ)
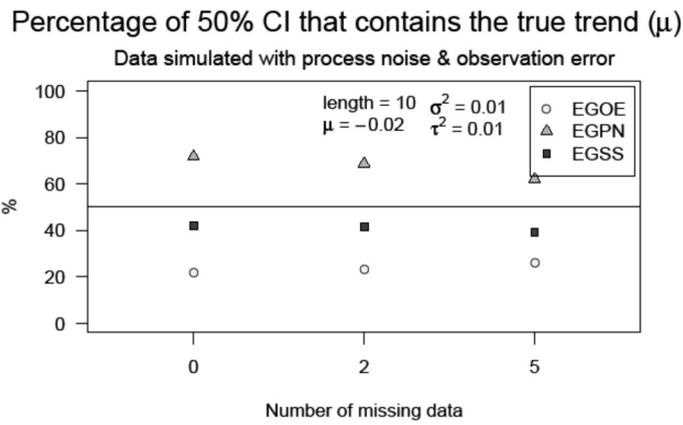
Data simulated with process noise & observation error

Figure 3a–b. Shorter time series: time series length = 10, μ= –0.02, σ^2= 0.01 and τ^2= 0.01.

(a)

Estimated μ
Data simulated with different ratios of process noise / observation error

(b)

Percentage of 50% CI that contains the true trend (μ)
Data simulated with different ratios of process noise / observation error
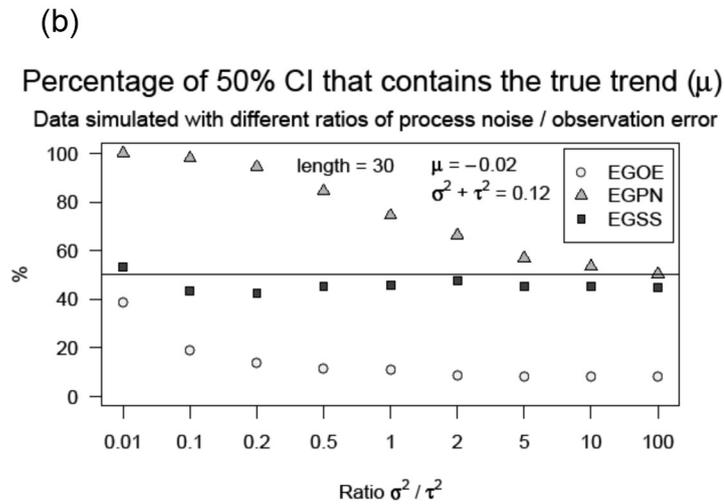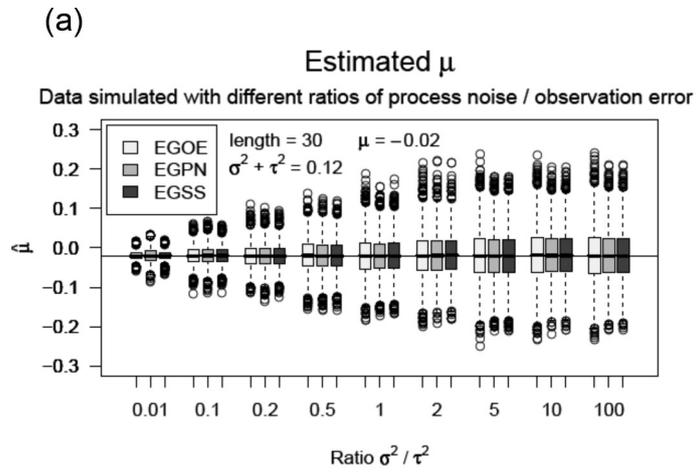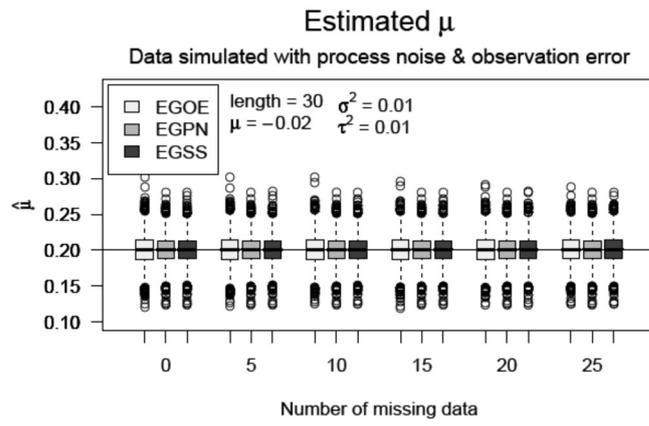
Figure 4a–b. Higher variance: time series length = 30, μ= 0, $\sigma^2 + \tau^2 = 0.12$. Ratio of process (σ^2) to sampling (τ^2) variance ranging from 0.01 (trivial process variance) to 100 (trivial sampling variance).

(a)

Estimated μ

Data simulated with process noise & observation error

(b)

Percentage of 50% CI that contains the true trend (μ)

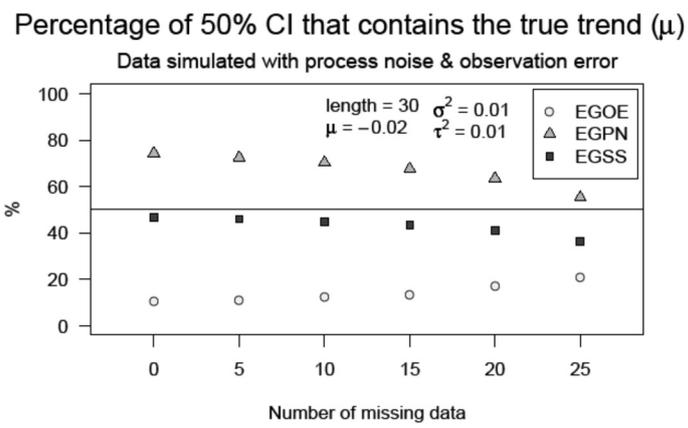Data simulated with process noise & observation error

Figure 5a–b. Large positive trend: time series length = 30, μ= 0.20, σ^2= 0.01 and τ^2= 0.01.