A semiparametric Bayesian method for detecting Allee effects

MASATOSHI SUGENO[1,3] AND STEPHAN B. MUNCH[2]

[1]*School of Marine and Atmospheric Sciences, Stony Brook University, Stony Brook, New York 11794-5000 USA*
[2]*Southwest Fisheries Science Center, 110 Shafer Rd., Santa Cruz, California 95060 USA*

Abstract. The importance of Allee effects has long been recognized both in theoretical studies of population dynamics and in conservation sciences. Although the necessary conditions for Allee effects to occur (e.g., difficulty in finding mates and mortality driven by generalist predators at low density) would seem to apply to many species, evidence for Allee effects in natural populations is equivocal at best. This apparent scarcity might be an artifact driven by poor power to detect them with traditional parametric models. To circumvent this potential problem, we developed a semiparametric Bayesian method based on a Gaussian process prior. We validated the method using simulated data sets and applied it to three herring data sets.

Key words: Allee effects; depensation; Gaussian process; semiparametric Bayesian modeling.

INTRODUCTION

Allee effects, which occur when population growth accelerates with density at low population sizes, have long been recognized as potential contributors to population extinction (Courchamp et al. 2008). Theoretical studies indicate that Allee effects may be important for conserving and managing wild populations: extinction probabilities for endangered species (Boukal and Berec 2002, Dennis 2002), vulnerability to invasion by introduced species, risk of disease outbreak (Taylor and Hastings 2005, Tobin et al. 2011), robust restoration programs for collapsed populations (Grevstad 1999, Deredec and Courchamp 2007, Armstrong and Wittmer 2011), and optimum harvesting rates for the sustainable use of wild populations (Lande et al. 1994), are all dependent on whether or not Allee effects are incorporated in model analysis. Allee effects may also offer an explanation for the ineffectiveness of moratoria to rebuild collapsed populations (Swain and Sinclair 2000, Walters and Kitchell 2001).

Several mechanisms contribute to generating Allee effects, including increased difficulty in finding mates (Gascoigne et al. 2009), increased mortality driven by generalist predators (Gascoigne and Lipcius 2004), and decreased fitness due to inbreeding depression (Willi et al. 2005). Each of these potential mechanisms has been observed at reduced population sizes (Berec et al. 2007, Courchamp et al. 2008, Gascoigne et al. 2009, Kramer et al. 2009).

Because the conditions for Allee effects to occur seem likely to apply in many species, we would expect Allee effects to manifest at the population level, referred to as "demographic" Allee effects (Stephens and Sutherland 1999). However, evidence for the presence of demographic Allee effects (hereafter Allee effects) is equivocal at best. Previous analyses of ecological time series for both terrestrial and aquatic species (Sæther et al. 1996, Sibly et al. 2005, Gregory et al. 2010) and fisheries data (Myers et al. 1995, Liermann and Hilborn 1997, Chen et al. 2002, Barrowman et al. 2003, Nash et al. 2009) all found limited evidence for Allee effects.

However, there are two potential pitfalls regarding previous methods for detecting Allee effects. First, the conclusions of the previous analyses are clearly conditional on the parametric models used. This is particularly important for noisy ecological data because many different models may fit equally well (or poorly) but produce qualitatively different predictions (see, e.g., Wood and Thomas 1999). Second, even when the correct model is *known*, Allee effects may be difficult to detect in noisy, short data sets (Shelton and Healey 1999). Using simulated data generated by an Allee effects model (the sigmoid Beverton-Holt model), Shelton and Healey (1999) demonstrated that likelihood ratio tests favored the true model over a model without Allee effects only when (1) simulated data visibly show the modeled shape, or (2) large samples ($n > 60$) are available.

Because we do not know the correct model and data clearly exhibiting Allee effects are rarely available, we hypothesize that the empirical rarity of Allee effects may be an artifact driven by poor power to detect them with traditional parametric models. To circumvent this limitation, we developed a semiparametric Bayesian (SB) method using a Gaussian process (GP) prior (Munch et al. 2005, Rasmussen and Williams 2006) to construct an index for assessing the presence of Allee effects. Although some non/semiparametric models are available for modeling simple population dynamics (e.g., Evans and Rice 1988, Cook 1998, Bravington et al. 2000, Munch et al. 2005), none of these directly

addressed the detection of Allee effects. We validated the method using simulated data sets and applied the method to data sets for three different Atlantic herring (*Clupea harengus*) populations.

METHODS

We begin this section by describing a general approach to modeling Allee effects and obtaining semiparametric Bayesian (SB) estimates of density dependence using Gaussian processes (GP). We then develop an approach to detecting Allee effects using the inferred density dependence. Finally, we describe the simulation studies used for method validation and three herring data sets used for case studies to illustrate the method. Mathematical details are given in the Appendices.

The Allee effects model

The relationship between present and future population size (N) may be written as

$$N_{t+1} = N_t \exp[f(N_t)]. \qquad (1)$$

This form ensures that population size is nonnegative ($N_t \geq 0$) and that once the population goes extinct, it stays there ($N_{t+1} = 0$ if $N_t = 0$). Strictly speaking this model applies to scalar populations, but may be thought of as a low-order approximation to the dynamics for structured populations. Allee effects are present if there is some population size (N_{Allee}) below which the population fails to replace itself, i.e., $N_{t+1} < N_t$ if $N_t < N_{Allee}$. Obviously, for the population to exist it must also be true that $N_{t+1} > N_t$ for $N_t > N_{Allee}$, from which we can deduce that N_{Allee} is an unstable equilibrium. Hence we must have $f(N_{Allee}) = 0$ and $f'(N_{Allee}) > 0$ (f' is the derivative of f with respect to N). Moreover, there must be some range of population sizes between 0 and N_{Allee} for which

$$f(N) \leq 0 \qquad f'(N) > 0. \qquad (2)$$

This is our general criterion for detecting Allee effects. For populations with age structure, this criterion provides a conservative bound on the presence of Allee effects (see Appendix A). Note that it is also feasible to have $f'(N) \leq 0$ for some $N < N_{Allee}$, corresponding to an interval where the per capita growth rate declines with density or remains the same. However, this is not a necessary condition for the presence of Allee effects and will not be considered further.

In situations where the total reproductive output can be approximated by the adult biomass, the relationship between adult (A) and juvenile (J) biomass can be written as follows:

$$J = A \exp[f(A)] \qquad (3)$$

where the arbitrary function, $f(A)$, determines the form of density dependence (Quinn and Deriso 1999, Courchamp et al. 2008). Allee effects occur at some small adult biomass where production of juvenile biomass falls below the level required for a population to persist
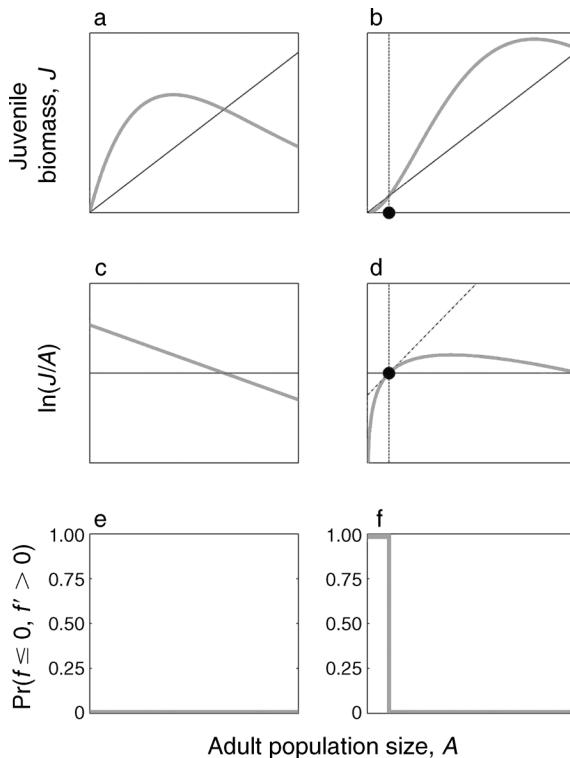

FIG. 1. (a) An example of a function for density dependence without Allee effects (gray solid line). (b) An example of a function with Allee effects (gray solid line) and the Allee effects threshold A_T (black dot). In panels (a) and (b), the horizontal axis indicates adult population size (A), and the vertical axis indicates juvenile biomass (J). The black solid lines correspond to $A = J$. (c) We plotted the example function shown in panel (a) in terms of $\ln(J/A)$ (the vertical axis) against A (the horizontal axis). (d) We plotted the example function shown in panel (b) in terms of $\ln(J/A)$ (the vertical axis) against A (the horizontal axis). We also plotted the slope of the example function at A_T (black broken line). (e) The probability of the presence of Allee effects, $\Pr[f(A) \leq 0, f'(A) > 0]$, obtained from the example function used in panels (a) and (c). (f) The probability of the presence of Allee effects obtained from the example function used in panels (b) and (d). In panels (e) and (f), the horizontal axis indicates adult population size, and the vertical axis indicates the probability of Allee effects.

(Quinn and Deriso 1999, Courchamp et al. 2008). More specifically, Allee effects occur when

$$f(A) \leq 0 \qquad f'(A) > 0 \qquad (4)$$

(see Fig. 1 for an illustration). This condition is analogous to the condition derived for a scalar population (Eq. 2). The largest value of A below which Eq. 4 is met is the Allee effect threshold (A_T) (Courchamp et al. 2008) and the population is considered to be at the edge of extinction when reduced below A_T. Throughout the remainder of the paper, we retain the juvenile–adult formulation of the model (Eqs. 3 and 4), but note that the methods can also be applied directly to other measures of populations size.

Note that our use of the term Allee effects refers exclusively to "strong" Allee effects. "Weak" Allee effects

differ in that f remains positive as population size declines. Therefore, populations with "weak" Allee effects can persist even at small population sizes. Although our methods readily can be extended to detect "weak" Allee effects, we focus our efforts on strong Allee effects because of their greater importance to conservation biology (Courchamp et al. 2008).

Modeling f with a GP prior

In the conventional approach to Allee effect detection, a parametric function is used to specify shape of f a priori. Common parametric models for Allee effects include the sigmoid Beverton-Holt (SBH) model (Myers et al. 1995) and the Saila-Lorda (SL) model (Saila et al. 1988), which are three-parameter generalizations of the Beverton-Holt and Ricker models, respectively. Other choices have also been used, although less frequently (Liermann and Hilborn 2001, Boukal and Berec 2002).

In contrast, our SB framework uses Bayesian inference for the shape of f starting from a GP prior. A GP prior is a generalization of the multivariate normal distribution to spaces of random functions, and analogously, the distribution is specified in terms of a mean function μ and a covariance function $\sum_{f,f}$ (Rasmussen and Williams 2006). Gaussian process regression was introduced to Bayesian analysis by O'Hagan and Kingman (1978), and applied to modeling density dependence by Munch et al. (2005). Rasmussen and Williams (2006) provide an excellent introduction to GP-based inference.

For mathematical convenience, we rewrite Eq. 3 using $y = \ln(J/A)$, and work with the log-transformed version:

$$y = f(A). \quad (5)$$

The analogous transformation for census data would be to use the log of population growth, $y = \ln(N_{t+1}/N_t)$ as the dependent variable (see, e.g., Sibly et al. 2005). In practice, Eqs. 1, 3, and 5 are only approximate. In keeping with a standard analysis of population time series (e.g., Myers et al. 1995, Gregory et al. 2010), we append an additive noise term ε to Eq. 5, which is normally distributed with the mean 0 and the variance σ^2.

Because the linear model is a common starting point for modeling density dependence (Myers et al. 1999, Sibly et al. 2005, Gregory et al. 2010), we assumed linearity in the prior mean function. Because we have no a priori reason to assume that the uncertainty of f or its curvature vary predictably with population size, we used an isotropic covariance function (i.e., a covariance function that only depends on the distance between population sizes). Specifically, we used the GP prior:

$$f \sim GP[\mu(A), \Sigma_{f,f}(A, A')]$$

$$\mu(A) = \ln(\alpha) + \beta A / \max(A)$$

$$\Sigma_{f,f}(A, A') = \tau^2 \exp\left[-\phi \left|\frac{A - A'}{\max(A)}\right|^2\right] \quad (6)$$

where $\ln(\alpha)$ and β are unbounded and τ^2 and ϕ are strictly positive. In this parameterization, τ^2 determines the vertical range of sampled values of f and ϕ determines its smoothness; a GP with a large (small) ϕ generates realizations of f that contain many (few) inflection points over the range of population sizes. Note that this choice for μ and $\Sigma_{f,f}$ does not severely limit the range of possible functional forms in the posterior within the range of the data (Munch et al. 2005, Rasmussen and Williams 2006); almost any functional shape that is clearly visible in the data may be recovered. We specified reference priors for $\ln(\alpha)$ and β and minimally informative priors for τ^2 and ϕ (Appendix B: *Prior specification*). After marginalizing over f (Rasmussen and Williams 2006), we used Metropolis-Hastings sampling (Munch et al. 2005; see Appendix B: *Parameter estimation*) to construct posterior distributions for the parameters, $Pr(\theta | data)$, where the parameters are collected in the vector $\theta = \{\ln\alpha, \beta, \phi, \tau^2, \sigma^2\}$.

Allee effects detection

The core idea of this paper is that we can use the SB inference for f described in the section *Modeling f with a GP prior* for calculating the probability of the presence of Allee effects. Specifically, we are interested in evaluating the posterior probability that the criterion (Eq. 4) is satisfied for some range of population sizes or estimates of adult biomass, i.e., $Pr[f(A) \leq 0, f'(A) > 0 | data]$.

Obtaining the posterior distribution of the presence of Allee effects involves a few steps. We first note that the derivative of a GP is also a GP (Rasmussen and Williams 2006), and the joint distribution for f and f' evaluated at a specific population size (and specific values of parameters) is bivariate normal (see Appendix C: *Construction of the joint distribution of a GP*). Given the observed population data, the joint conditional posterior distribution for f and f' evaluated at population size A, $Pr[f(A), f'(A) | \theta, data]$, is also bivariate normal (see Appendix D). Therefore, standard numerical integration can be applied to the bivariate normal distribution to obtain the posterior probability of the presence of Allee effects, i.e., $Pr[f(A) \leq 0, f'(A) > 0 | \theta, data]$ conditional on the parameters θ.

We are of course more interested in the posterior inference of Allee effects that are conditioned exclusively on data and not dependent on the specific choice of parameters, i.e., $Pr[f(A) \leq 0, f'(A) > 0 | data]$. However, this probability is analytically intractable and we obtained it via Monte Carlo integration using random samples generated from $Pr(\theta | data)$. For more details on these calculations, see Appendix D.

Our assessment framework for determining the presence of Allee effects is simply based on comparison of the prior and the posterior probabilities for the presence of Allee effects. Because the probability that Allee effects are present is most relevant at small population sizes, we used the posterior probability evaluated at the origin (i.e., $A = 0$ or $N = 0$) as an index, π, for the presence of Allee effects:

$$\pi = \Pr[f(0) \le 0, f'(0) > 0 \,|\, \text{data}]. \qquad (7)$$

Depending on the choice of parameters $\boldsymbol{\theta}$, π can take any value between 0 to 1 (i.e., $\Pr[f(0) \le 0, f'(0) > 0 \,|\, \boldsymbol{\theta}] \in (0, 1)$). However, marginalizing over the prior distribution for $\boldsymbol{\theta}$, the interval for $\Pr[f(0) \le 0, f'(0) > 0]$ collapses to 0.25 (for further details, see Appendix C: *Calculation of the benchmark probability*). To make this intuitive, recognize that the marginal prior over the two-dimensional space covered by $f(0)$ and $f'(0)$ is, more or less, the product of independent zero-mean normal distributions and is therefore symmetric such that each quadrant has equal prior probability. Thus, the marginal prior probability of being in the second quadrant (i.e., where $f(0) \le 0$ and $f'(0) > 0$) is 0.25. This prior on the presence of Allee effects is induced by our model structure and we use it as the benchmark probability. Because, in the posterior, π may be anywhere between 0 and 1, Allee effects are detected if $\pi > 0.25$. In cases where "weak" Allee effects are of interest, we could similarly calculate the probability of the presence of "weak" Allee effects, i.e., $\Pr[f(0) > 0, f'(0) > 0 \,|\, \text{data}]$, and compare it with the benchmark probability, 0.25.

Simulation study

We tested our method with simulated data sets. Because the strength of Allee effects, the noise level in data, and the data availability at low population sizes should influence our ability to correctly assess the presence of Allee effects (Hilborn and Walters 1992, Shelton and Healey 1999), our simulation study was designed to assess these factors comprehensively. For each treatment combination (i.e., Allee effect strength, noise level, and minimum population size), we generated 500 replicate data sets. For each simulated data set, we set the total number of observations to 30, which is representative of many ecological data sets.

We used the Saila-Lorda (SL) model (Saila et al. 1988), $J = \alpha A^{\gamma} \exp(-\beta A)$, to generate data with and without Allee effects (Table 1). The parameter γ controls the strength of Allee effects. We used three different strength levels: $\gamma = 1$ for data with negative density dependence (the Ricker model, hereafter "no effect"), $\gamma = 1.5$ for data with mild Allee effects ("mild"), and $\gamma = 2$ for data with strong Allee effects ("strong"). In the "no effect" case ($\gamma = 1$), we used a value of 4 for α and set $\beta = \alpha/e = 1.47$. For the "mild" and "strong" cases, we parameterized our simulation to isolate the influence of Allee effect strength on our ability to detect Allee effects. Specifically, we set parameters in the SL model so that the adult population size that can be replaced by recruitment (i.e., the largest A that satisfies $A = J$) and the slope, dy/dA at that population size were constant across all three models. Variability in recruitment was modeled with multiplicative log-normal noise without serial autocorrelation, $J = \alpha A^{\gamma} \exp(-\beta A + \omega)$, where $\omega \sim \mathcal{N}(0, \sigma_J^2)$ and σ_J^2 is the variance. We used σ_J^2 over the interval 0.1 to 1. A list of the parameters used is given in Table 1.

TABLE 1. Parameters for the data-generating models, $J = \alpha A^{\gamma} \exp(-\beta A)$.

Allee effect	α	β	γ	A_T
No effect	4	1.4712	1	N/A
Mild	6.7964	2.0023	1.5	0.0238
Strong	11.5416	2.533	2	0.1163

Notes: Here, A is adult population size, and J is juvenile biomass; α, β, and γ are parameters specifying the Saila-Lorda model. The Allee effects threshold, A_T, is not available in the "no effect" model because there are no Allee effects.

To evaluate method performance with respect to the availability of data at low population sizes independently of the range of observed values, we sampled A from a uniform distribution with constant width and variable minimum. Preliminary analysis of data for Atlantic cod (*Gadus morhua*) indicates that the range of observed adult population size is typically on the order of 1.5 times the density that generates the maximum recruitment (Sugeno and Munch, *in press*). In the "no effect" model with $\alpha = 4$ and $\beta = 1.47$, this is given by 0.68. Therefore, by multiplying this value with 1.5, we obtained 1.02 as the width and then set the minimum adult population size for this distribution as a proportion of 0.68, ranging from 0 to 0.34.

We compared the performance of our SB approach with the likelihood-based parametric approach in terms of Type I (false positive detection of Allee effects) and Type II (false negative detection of Allee effects) error rates. In keeping with previous parametric assessments of Allee effect frequency, we used the likelihood ratio to test for Allee effects, by comparing the fit of the SL model to the fit of the Ricker model (i.e., the SL model with $\gamma = 1$). To do so, we compared the likelihood ratio to a χ^2 distribution with 1 df and concluded that Allee effects are present if the nominal confidence level exceeded $P = 0.05$. To evaluate the sensitivity of our conclusions to parametric model choice, we also conducted likelihood ratio tests comparing the fit of the SBH model where $J = \alpha A^{\gamma}/(1 + \beta A^{\gamma})$ to the standard Beverton-Holt (BH) model (i.e., when $\gamma = 1$). For the SB method, we concluded that Allee effects were present whenever the posterior value for π exceeded the benchmark probability, 0.25. For each of these approaches (SL, SBH, and SB), we estimated error rates from the 500 replicates for each parameter combination. All calculations were carried out using code written in MATLAB 7 (MathWorks, Natick, Massachusetts, USA); see the Supplement.

Herring data sets

To illustrate application of the SB method to real data, we analyzed three different Atlantic herring (*Clupea harengus*) data sets from the Stock Recruitment Database (*available online*).[4] We used data on Iceland

[4] http://ramlegacy.marinebiodiversity.ca/ram-legacy-stock-assessment-database
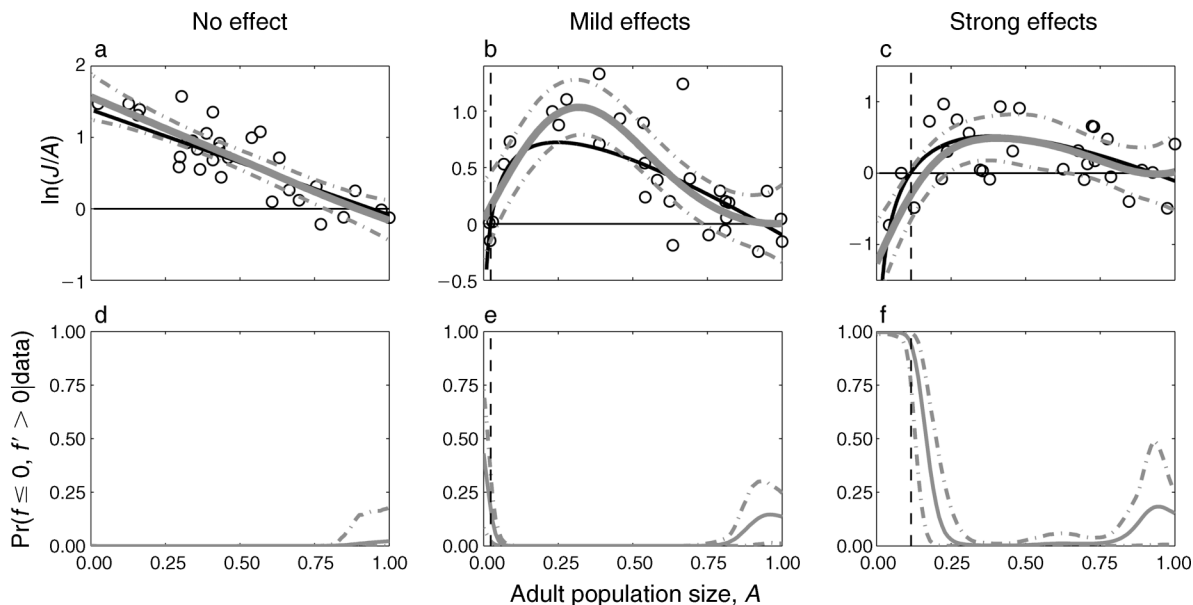
Fig. 2. An illustration of the SB (semiparametric Bayesian) method for detecting Allee effects using simulated data with 0 for minimum adult population size, A, with variance $\sigma_J^2 = 0.5$. The data models used to generate simulated data were (a, d) the "no effect" model ($\gamma = 1$ in the SL model); (b, e) the "mild" Allee effects model ($\gamma = 1.5$ in the SL model, $A_T = 0.024$); and (c, f) the "strong" Allee effects model ($\gamma = 2$, $A_T = 0.12$). (a–c) Fit of the GP prior to simulated data. We plotted the true functions (black solid line), simulated data (black-outlined circles), the mean (gray solid line), and the 95% confidence intervals (dash-dot lines) of the posterior predictive $f(A)$. The true A_T is indicated by the black vertical dashed line. The horizontal axis is adult population size (A), and the vertical axis is $\ln(J/A)$. (d–f) The posterior probability for the presence of Allee effects. We plotted the mean (gray solid line) and the 95% confidence intervals of the posterior probability (dash-dot lines). The horizontal axis is adult population size, and the vertical axis indicates the probability that Allee effects are present.

spring spawner (ICE, $n = 23$, 1947–1969), Downs stock (DOWN, $n = 65$, 1923–1987), and Georges Bank (GB, $n = 15$, 1961–1975). In keeping with previous work (Myers et al. 1995), the raw juvenile data (in the unit of the number of individuals) were multiplied by mass per individual so that y was dimensionless.

RESULTS

Simulation results: case studies

For clarity, we begin by describing a single set of simulations before describing the results of the broader simulation study. In this illustrative example (Fig. 2), we examined method performance under "no," "mild," and "strong" Allee effects, where the minimum population size was 0 and $\sigma_J^2 = 0.5$. In each case, the SB fit reasonably recaptures the shape of the data-generating models (Fig. 2a–c). More importantly, the probability of the presence of Allee effects clearly corresponds to the truth in the simulated data. In the "no effect" case, the probability is zero for all points close to the origin (Fig. 2d). In the "strong" effects case, there is a clear region in which the probability is close to 1 (Fig. 2f). In the "mild" case, there is a peak in the probability at the origin, but with wide confidence intervals (Fig. 2e). The probability of the presence of Allee effects evaluated at zero population size, π, is 0 for "no effects" data, 0.43 for "mild" data, and 1.00 for "strong" data, indicating that the presence of Allee effects was evaluated correctly.

However, a wide 95% confidence interval for π (0.084 for the lower bound and 0.73 for the upper bound) in "mild" data indicates that evidence for the presence of Allee effects is ambiguous.

Allee effects detection frequency

We turn now to summarizing the results of the simulation in general. Overall, averaging across data-generating models, noise levels, and the minimum population size, the SB method is the best method for correctly assessing for the presence of Allee effects (Table 2), although not by much. The error rate for the SB method is 0.49, whereas those for parametric models were 0.53 (SL) and 0.54 (SBH), respectively.

But the similarity of the error rates is an artifact of averaging, treating "no effect," "mild," and "strong" as equally likely. Parsing the error rates for each type of density dependence reveals that the SB method performed substantially better than parametric alternatives when data were "no effect" and "strong" (Table 2). In particular, for "strong" cases, the Type II error rate for the SB method is 0.58, followed by 0.70 (SL) and 0.73 (SBH).

To examine the performance of the models in greater detail, we generated Fig. 3 using subsets of the comprehensive analysis to show how the error rates were affected by the data-generating models, the noise level, and the minimum population size. For "no effect"

TABLE 2. Error rates of the SB (semiparametric Bayesian) method, the SL (Saila-Lorda) model, and the SBH (sigmoid Beverton-Holt) model in assessing the presence of Allee effects.

Approach	Error rate, by Allee effect			
	Average	No effect	Mild	Strong
SB method	**0.4890**	**0.0265**	0.8572	**0.5832**
SL model	0.5271	0.0284	**0.8508**	0.7021
SBH model	0.5432	0.0416	0.8580	0.7301

Notes: "Average" is the error rate averaged over the noise level, the minimum population size, and data-generating models. Error rates for the three models in "no effect," "mild," and "strong" data sets are also shown. The number indicated in boldface indicates the best model in each category.

cases, the Type I error rate for all three methods is quite low overall (Fig. 3a, d, g, j). With data close to the origin (0 for the minimum population size) and low noise ($\sigma_J^2 = 0.1$), the SB method is the only method that showed 0 error rate. The error rate for the SL model (0.02) is within acceptable levels, but that for the SBH model (0.18) is more than three times the nominal 0.05 level.

For "mild" and "strong" cases, the Type II error rates for all three models are greater and clearly increased with minimum population size. In keeping with previous results (Shelton and Healey 1999), the Type II error rate is larger in the "mild" cases than in the "strong" cases. In addition, as the noise level and minimum population size increase, detection probability goes down significantly and the performance of the parametric models is equivalent. The SB method performed better than the parametric models in general, particularly when the noise level was large, and was only outperformed by parametric alternatives at low noise ($\sigma^2 = 0.1$).

Empirical results: herring data sets

The probability of the presence of Allee effects evaluated at zero adult biomass (a proxy for zero population size), π, is 0.67 for Iceland spring spawner population (ICE), 0 for the Georges bank population (GB), and 0.70 for the Downs stock population (DOWN) (Fig. 4d–f). These results indicated that ICE and GB exhibit Allee effects whereas DOWN does not. We note that Myers et al. (1995) previously found evidence for Allee effects in ICE using the SBH model. However, wide 95% confidence intervals in π for ICE (0.13 for the lower bound and 0.91 for the upper bound) and GB (0.096 for the lower bound and 1.00 for the upper bound) suggest that results for ICE and GB are ambiguous, and that Allee effects are probably "mild." For the GB population, this result may be due to the limited sample size ($n = 15$).

DISCUSSION

In the context of Allee effects detection, Type I (false positive) and Type II (false negative) errors are not equivalent: the cost of a Type I error is reduced harvesting on a viable population, whereas the cost of
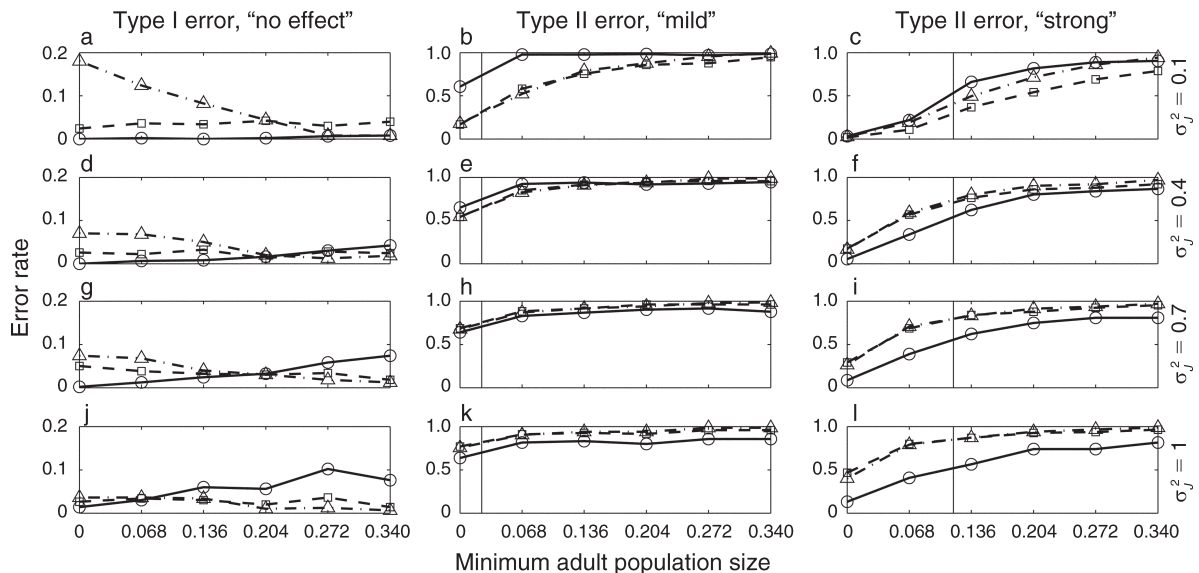

FIG. 3. Frequency of making incorrect assessment for the presence of Allee effects. The left column describes results for data without Allee effects ("no effect"). Hence we focus on a Type I error (false positive). The middle and right columns describe results for "mild" and "strong" Allee effects, and therefore we focused on a Type II error (false negative). In each panel, the horizontal axis is minimum adult population size in simulated data, and the vertical axis is the frequency of making an incorrect assessment for the presence of Allee effects. Rows of the plots correspond to noise levels $\sigma_J^2 = \{0.1, 0.4, 0.7, 1.0\}$. The solid line (with open circles) gives error rates for the SB method; the dashed line (with open squares) and dash-dot line (with open triangles) give error rates for the SL and the SBH likelihood ratio tests, respectively. In the middle and right columns, the vertical line indicates values of Allee effects thresholds ($A_T = 0.024$ for "mild" Allee effects and $A_T = 0.12$ for "strong" Allee effects).
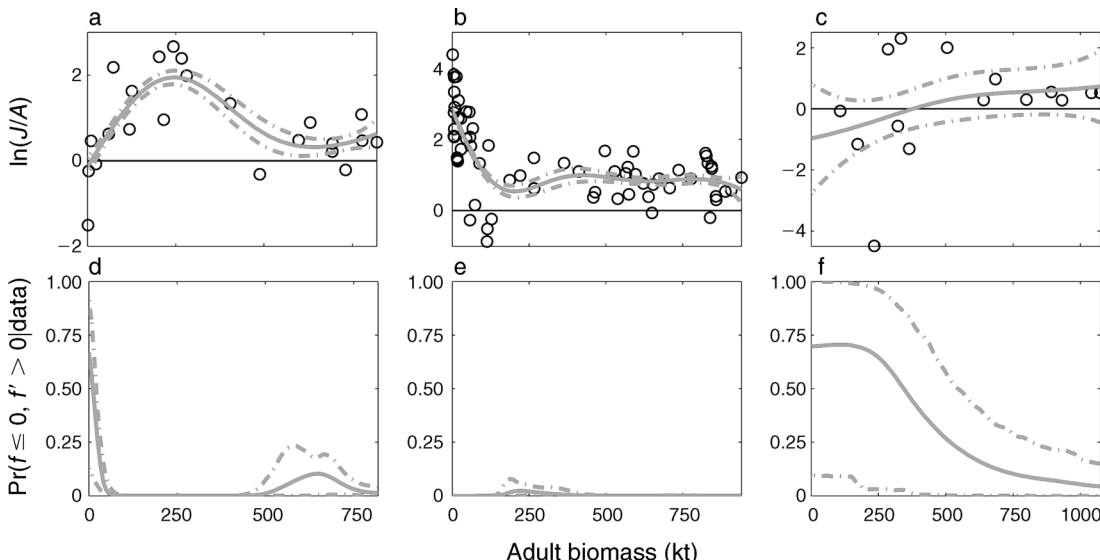
FIG. 4.   Analysis of three herring populations. They are (a, d) ICE (Iceland spring spawner, $n = 23$, 1947–1969); (b, e) DOWN (Downs stock, $n = 65$, 1923–1987); and (c, f) GB (Georges Bank, $n = 15$, 1961–1975). (a–c) The posterior predictive distributions for $f(A)$ were obtained in order to evaluate the fit of the SB model. We plotted data (black-outlined circles) and the mean (gray solid line) and the 95% confidence intervals of the posterior predictive $f(A)$ (dash-dot lines). The horizontal axis is adult biomass (A, a proxy for adult population size; kt = 1000 metric tons = 1 Gg), and the vertical axis is $\ln(J/A)$. (d–f) The posterior probability for the presence of Allee effects. We plotted the mean (gray solid line) and the 95% confidence intervals of the posterior probability (dash-dot lines). The horizontal axis is adult biomass, and the vertical axis indicates the probability that Allee effects are present.

a Type II error is extinction (Table 3; see Courchamp et al. 1999, Stephens and Sutherland 1999, Mieszkowska et al. 2009). Our results suggest that the SB method could reduce chance of overexploitation and population collapse due to making a Type II error: the SB method is clearly better at detecting Allee effects when present compared to parametric alternatives, including the true model, especially when "strong" Allee effects are present. Moreover, the SB method is effective regardless of the "true" model. This alone should make it broadly applicable in ecology where the underlying dynamics are rarely certain. That said, the success of the SB method is tempered somewhat by the fact that Type II error rates were still >50% for most cases, indicating the need for caution.

In this study, the SB method for detecting Allee effects was based on comparison of the probability for the presence of Allee effects, π, with the detection threshold of 0.25, which emerges from the prior specification of the SB method. Using this threshold, the SB method reduced the chance of making a Type II error compared with parametric models, but the error rate was still quite high. Because of asymmetric cost of incorrect assessment for the presence of Allee effects (Table 3) and need for precautionary approaches to managing natural populations (e.g., Hilborn et al. 2001), it may be desirable to set the detection threshold at some value lower than 0.25 in order to further reduce the chance of making a Type II error. This could be done in several ways, including using informative or nonindependent priors for $\ln(\alpha)$

and β, an anisotropic covariance function, or by choosing a threshold directly. In broad terms, a precautionary threshold might allow for a greater probability of a Type I error in order to avoid rejecting Allee effects when they are hard to detect. Of course, adopting an excessive safety margin would result in unnecessary losses to society as well (e.g., Hilborn et al. 2001), and we agree with many prior studies that a decision-theoretic approach (Berger 1985, Francis and Shotton 1997, Punt and Hilborn 1997, Wade 2000, Fenichel et al. 2008) should be taken to determine the most reasonable threshold level. This process usually requires considerable discussion between scientists, environmental practitioners, and local communities and is beyond the scope of the present paper.

Previous studies have concluded that Allee effects are quite rare in natural populations. Using parametric models and likelihood ratio tests, Gregory et al. (2010) analyzed 1198 data sets that include both terrestrial and aquatic species and Myers et al. (1995) analyzed 129 data sets on commercially harvested fish species. These

TABLE 3.   The consequence of incorrect assessment for the presence of Allee effects on populations.

	True form of density dependence	
Assessment	No effect	Allee effects
No effect	rebuild	Type II error (extinction)
Allee effects	Type I error (rebuild)	rebuild

studies concluded that Allee effects appear in <1.1% and <2.4% of populations, respectively. Given the results of our simulation study, however, we suspect that Allee effects are more prevalent than currently believed. Reanalysis of these data sets with a more sensitive tool such as the SB method developed here seems warranted.

LITERATURE CITED

Armstrong, D. P., and H. U. Wittmer. 2011. Incorporating Allee effects into reintroduction strategies. Ecological Research 26(4):687–695.

Barrowman, N. J., R. A. Myers, R. Hilborn, D. G. Kehler, and C. A. Field. 2003. The variability among populations of coho salmon in the maximum reproductive rate and depensation. Ecological Applications 13:784–793.

Berec, L., E. Angulo, and F. Courchamp. 2007. Multiple Allee effects and population management. Trends in Ecology and Evolution 22:185–191.

Berger, J. O. 1985. Statistical decision theory and Bayesian analysis. Springer Series in Statistics. Springer, New York, New York, USA.

Boukal, D. S., and L. Berec. 2002. Single-species models of the Allee effect: extinction boundaries, sex ratios and mate encounters. Journal of Theoretical Biology 219:375–394.

Bravington, M. V., T. K. Stokes, and C. M. O'Brien. 2000. Sustainable recruitment: the bottom line. Marine Freshwater Research 51:465–75.

Chen, D. G., J. R. Irvine, and A. J. Cass. 2002. Incorporating Allee effects in fish stock–recruitment models and applications for determining reference points. Canadian Journal of Fisheries and Aquatic Sciences 59:242–249.

Cook, R. M. 1998. A sustainability criterion for the exploitation of North Sea cod. ICES Journal of Marine Science 55:1061–1070.

Courchamp, F., L. Berec, and J. Gascoigne. 2008. Allee effects in ecology and conservation. Oxford University Press, New York, New York, USA.

Courchamp, F., T. Clutton-Brock, and B. Grenfell. 1999. Inverse density dependence and the Allee effect. Trends in Ecology and Evolution 14:405–410.

Dennis, B. 2002. Allee effects in stochastic populations. Oikos 96:389–401.

Deredec, A., and F. Courchamp. 2007. Importance of the Allee effect for reintroductions. Ecoscience 14:440–451.

Evans, G. T., and J. C. Rice. 1988. Predicting recruitment from stock size without the mediation of a functional relation. ICES Journal of Marine Science 44:111–122.

Fenichel, E. P., J. I. Tsao, M. L. Jones, and G. J. Hickling. 2008. Real options for precautionary fisheries management. Fish and Fisheries 9:121–137.

Francis, R. I. C. C., and R. Shotton. 1997. "Risk" in fisheries management: a review. Canadian Journal of Fisheries and Aquatic Sciences 54:1699–1715.

Gascoigne, J., L. Berec, S. Gregory, and F. Courchamp. 2009. Dangerously few liaisons: a review of mate-finding Allee effects. Population Ecology 51:355–372.

Gascoigne, J. C., and R. N. Lipcius. 2004. Allee effects driven by predation. Journal of Applied Ecology 41:801–810.

Gregory, S. D., C. J. A. Bradshaw, B. W. Brook, and F. Courchamp. 2010. Limited evidence for the demographic Allee effect from numerous species across taxa. Ecology 91:2151–2161.

Grevstad, F. S. 1999. Experimental invasions using biological control introductions: the influence of release size on the chance of population establishment. Biological Invasions 1(4):313–323.

Hilborn, R., J. J. Maguire, A. M. Parma, and A. Rosenberg. 2001. The precautionary approach and risk management: can they increase the probability of successes in fishery management? Canadian Journal of Fisheries and Aquatic Sciences 58:99–107.

Hilborn, R., and C. J. Walters. 1992. Quantitative fisheries stock assessment: choice, dynamics and uncertainty. Chapman and Hall, New York, New York, USA.

Kramer, A. M., B. Dennis, A. M. Liebhold, and J. M. Drake. 2009. The evidence for Allee effects. Population Ecology 51:341–354.

Lande, R., S. Engen, and B.-E. Sæther. 1994. Optimal harvesting, economic discounting and extinction risk in fluctuating populations. Nature 372:88–90.

Liermann, M., and R. Hilborn. 1997. Depensation in fish stocks: a hierarchical Bayesian meta-analysis. Canadian Journal of Fisheries and Aquatic Sciences 54:1976–1984.

Liermann, M., and R. Hilborn. 2001. Depensation: evidence, models and implications. Fish and Fisheries 2:33–58.

Mieszkowska, N., M. J. Genner, S. J. Hawkins, and D. W. Sims. 2009. Effects of climate change and commercial fishing on Atlantic cod Gadus morhua. Advances in Marine Biology 56:213–273.

Munch, S. B., A. Kottas, and M. Mangel. 2005. Bayesian nonparametric analysis of stock–recruitment relationships. Canadian Journal of Fisheries and Aquatic Sciences 62:1808–1821.

Myers, R. A., N. J. Barrowman, J. A. Hutchings, and A. A. Rosenberg. 1995. Population dynamics of exploited fish stocks at low population levels. Science 269:1106–1108.

Myers, R. A., K. G. Bowen, and N. J. Barrowman. 1999. Maximum reproductive rate of fish at low population sizes. Canadian Journal of Fisheries and Aquatic Sciences 56:2404–2419.

Nash, R. D. M., M. Dickey-Collas, and L. T. Kell. 2009. Stock and recruitment in North Sea herring (Clupea harengus); compensation and depensation in the population dynamics. Fisheries Research 95:88–97.

O'Hagan, A., and J. F. C. Kingman. 1978. Curve fitting and optimal design for prediction. Journal of the Royal Statistical Society B 40:1–42.

Punt, A. E., and R. Hilborn. 1997. Fisheries stock assessment and decision analysis: the Bayesian approach. Reviews in Fish Biology and Fisheries 7:35–63.

Quinn, T. J., and R. B. Deriso. 1999. Quantitative fish dynamics. Oxford University Press, New York, New York, USA.

Rasmussen, C. E., and C. K. I. Williams. 2006. Gaussian processes for machine learning. MIT Press, Cambridge, Massachusetts, USA.

Sæther, B.-E., T. H. Ringsby, and E. Roskaft. 1996. Life history variation, population processes and priorities in species conservation: towards a reunion of research paradigms. Oikos 77:217–226.

Saila, S. B., C. W. Recksiek, and M. H. Prager. 1988. Basic fishery science programs: a compendium of microcomputer programs and manual of operation. Elsevier, Amsterdam, The Netherlands.

Shelton, P. A., and B. P. Healey. 1999. Should depensation be dismissed as a possible explanation for the lack of recovery of the northern cod (Gadus morhua) stock? Canadian Journal of Fisheries and Aquatic Sciences 56:1521–1524.

Sibly, R. M., D. Barker, M. C. Denham, J. Hone, and M. Pagel. 2005. On the regulation of populations of mammals, birds, fish, and insects. Science 209:607–610.

Stephens, P. A., and W. J. Sutherland. 1999. Consequences of the Allee effect for behaviour, ecology and conservation. Trends in Ecology and Evolution 14:401–405.

Sugeno, M., and S. B. Munch. *In press*. A semiparametric Bayesian approach to estimating maximum reproductive rates at low population sizes. Ecological Applications 23:699–709.

Swain, D. P., and A. F. Sinclair. 2000. Pelagic fishes and the cod recruitment dilemma in the Northwest Atlantic. Canadian Journal of Fisheries and Aquatic Sciences 57:1321–1325.

Taylor, C. M., and A. Hastings. 2005. Allee effects in biological invasions. Ecology Letters 8:895–908.

Tobin, P. C., L. Berec, and A. M. Liebhold. 2011. Exploiting Allee effects for managing biological invasions. Ecology Letters 8:895–908.

Wade, P. R. 2000. Bayesian methods in conservation biology. Conservation Biology 14:1308–1316.

Walters, C., and J. F. Kitchell. 2001. Cultivation/depensation effects on juvenile survival and recruitment: implications for the theory of fishing. Canadian Journal of Fisheries and Aquatic Sciences 58:39–50.

Willi, Y., J. van Buskirk, and M. Fischer. 2005. A threefold genetic Allee effect: population size affects cross-compatibility, inbreeding depression and drift load in the self-incompatible *Ranunculus reptans*. Genetics 169:2255–2265.

Wood, S. N., and M. B. Thomas. 1999. Super-sensitivity to structure in biological models. Proceedings of the Royal Society B 266:565–570.

SUPPLEMENTAL MATERIAL

Appendix A

Allee threshold for populations with age structure (*Ecological Archives* E094-107-A1).

Appendix B

Prior specification and parameter estimation (*Ecological Archives* E094-107-A2).

Appendix C

The derivative of a GP, $f'(A)$ (*Ecological Archives* E094-107-A3).

Appendix D

Allee effects detection (*Ecological Archives* E094-107-A4).

Supplement

MATLAB software used in our paper (*Ecological Archives* E094-107-S1).